

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Provable and Efficient Algorithms for Federated, Batch and Reinforcement Learning

Permalink

<https://escholarship.org/uc/item/8km4g7sq>

Author

Ghosh, Avishek

Publication Date

2021

Peer reviewed|Thesis/dissertation

Provable and Efficient Algorithms for Federated, Batch and Reinforcement Learning

by

Avishek Ghosh

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering – Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Kannan Ramchandran, Co-chair

Professor Aditya Guntuboyina, Co-chair

Professor Thomas Courtade

Professor Anil Aswani

Spring 2021

Provable and Efficient Algorithms for Federated, Batch and Reinforcement Learning

Copyright 2021
by
Avishek Ghosh

Abstract

Provable and Efficient Algorithms for Federated, Batch and Reinforcement Learning

by

Avishek Ghosh

Doctor of Philosophy in Engineering – Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Kannan Ramchandran, Co-chair

Professor Aditya Guntuboyina, Co-chair

We propose and analyze iterative algorithms that are computationally efficient, statistically sound and adaptive (in some settings). We consider three different frameworks in which data is presented to the learner. First, we consider the Federated (Distributed) Learning (FL) setup, where data is only available at the edge, and a center machine learns various models via iteratively interacting with the edge nodes. Second, we study the canonical setting of supervised batch learning, where all the data and label pairs are available to the learner at the beginning. Third, we examine the framework of online learning, where data is presented in a streaming fashion. In particular, we focus on specific settings like Bandit and Reinforcement Learning (RL).

In the Federated Learning (FL) framework, we address the canonical problems of device heterogeneity, communication bottleneck and adversarial robustness for large scale high dimensional problems. We propose efficient and provable first and second order algorithms, and use ideas like quantization of information and apply several robust aggregation schemes to address the above-mentioned problems, while retaining the optimal statistical rates simultaneously. For the (supervised) batch learning framework, we use an efficient and statistically sound algorithm, namely Alternating Minimization (AM) and address the problem of max-affine regression; a non convex problem that generalizes the classical phase retrieval and closely resembles convex regression. We give convergence guarantees of AM, with near optimal statistical rate. Finally, in the online learning setup, we address the problem of adaptation (model selection) for contextual bandits (linear and beyond) and later extend these techniques to Reinforcement Learning (RL). Our algorithms here are efficient, provable and more importantly adaptive to the problem complexity.

Dedicated to Maa and Baba

Contents

Contents	ii
1 Introduction	1
1.1 Algorithms for Federated (Distributed) Learning (FL)	1
1.2 Learning in Supervised Batch Setup	7
1.3 Efficient and Adaptive Algorithms for Bandit and Reinforcement Learning (RL) . .	10
I Learning in Federated Framework	16
2 Clustered Federated Learning	17
2.1 Introduction	17
2.2 Related work	18
2.3 Problem formulation	19
2.4 Algorithm	20
2.5 Theoretical guarantees	22
2.6 Experiments	26
2.7 Conclusion and Open Problems	29
2.8 Proof of Theorem 1	30
2.9 Proof of Theorem 2	35
3 Communication-Efficient and Byzantine-Robust Distributed First Order Learning	41
3.1 Introduction	41
3.2 Problem Formulation	45
3.3 Compression At Worker Machines	46
3.4 Robust Compressed Gradient Descent	47
3.5 Distributed Learning with Restricted Adversaries	48
3.6 Distributed Optimization with Arbitrary Adversaries	51
3.7 Byzantine Robust Distributed Learning with Error Feedback	52
3.8 Experiments	57
3.9 Conclusion and Open Problems	60
3.10 Analysis of Algorithm 4	61

3.11	Proof of Theorem 5	71
4	Efficient and Robust FL with Newton Algorithm	80
4.1	Introduction	80
4.2	Problem Formulation	83
4.3	COMRADE Can Communicate Less	83
4.4	COMRADE Can Resist Byzantine Workers	86
4.5	COMRADE Can Communicate Even Less and Resist Byzantine Workers	87
4.6	Experimental Results	88
4.7	Conclusion	89
4.8	Analysis of Section 4.3	90
4.9	Analysis of Section 4.4	97
4.10	Auxiliary Lemmas	100
4.11	Analysis of Section 4.5	102
5	Escaping Saddle Points in FL: Distributed Cubic Regularized Newton Method	108
5.1	Introduction	108
5.2	Related Work	111
5.3	Problem Formulation	112
5.4	Distributed Cubic Regularized Newton	113
5.5	Byzantine Resilience	116
5.6	Experimental Results	118
5.7	Proofs	121
II	Efficient and Tractable Algorithms for Non-Convex Batch Learning	136
6	Max Affine Regression with Gaussian Design	137
6.1	Introduction	137
6.2	Background and problem formulation	143
6.3	Main results	148
6.4	Discussion	157
6.5	Proofs	160
6.6	Proof of Theorem 12	163
6.7	Proof of Theorem 13	179
6.8	Proof of Theorem 14	185
6.9	Technical Lemmas and Background	188
7	Max-affine Regression Beyond Gaussian Design—Small Ball covariates	207
7.1	Introduction	207
7.2	Main results	210
7.3	Discussion and Open Problems	216

7.4	Proof of Theorem 18	217
7.5	Proof of Corollary 3	227
7.6	Technical Lemmas and Background	227

III Learning and Adaptation in Bandit and RL Framework 234

8 Problem-Complexity Adaptive Model Selection for Stochastic Linear Bandits 235

8.1	Introduction	235
8.2	Related Work	238
8.3	Norm as a measure of Complexity	239
8.4	Dimension as a Measure of Complexity - Continuum Armed Setting	242
8.5	Dimension as a Measure of Complexity - Finite Armed Setting	245
8.6	Simulations	246
8.7	Conclusion	248
8.8	Proofs	249

9 Model Selection for Contextual Bandits Beyond Linear Structure 264

9.1	Problem formulation	264
9.2	Algorithm—Adaptive Contextual Bandits (ACB)	266
9.3	Explore-then-Commit algorithm for model selection	268
9.4	Model Selection with infinite function classes	271
9.5	Proofs	274

10 Model Selection for Reinforcement Learning with Function Approximation 283

10.1	Formulation and Linear Model	283
10.2	Norm as Complexity measure	285
10.3	Dimension as complexity measure	287
10.4	Conclusion	290
10.5	Proofs	290

Bibliography 300

Acknowledgments

I would like to acknowledge the people who have played a crucial part toward this thesis. This thesis would not have been possible without their help and support over the years. This list is nowhere comprehensive, and there are many whose presence I appreciated greatly but did not get a chance to thank here.

I would like to express my heartfelt gratitude to my PhD co-advisors, Prof. Kannan Ramchandran and Prof. Aditya Guntuboyina. Kannan has been a great support and inspiration for me throughout my doctoral journey. First, I was amazed by his dedication as a teacher. I had the good opportunity to work as a teaching assistant with him in an undergraduate probability and a graduate information theory course, and the amount of work he puts towards improving the courses is simply mind blowing. Furthermore, he was always available for stimulating discussion sessions on research, brainstorming about new directions to pursue. He is always open to explore new areas, and motivated me to do the same. Kannan gave me the freedom to explore areas I like, even when it fell outside his research interest. Many of the problems I have in this thesis actually started from our brainstorming sessions. One thing I learnt from Kannan is that, it is of fundamental interest to ask the correct set of questions, even though the answers to them are partially known. He also helped me to improve my presentation skills and taught me how to keep the audience engaged. His enthusiasm for research and devotion to his work are forever going to be an inspiration for me in my career and personal life.

I would like to extend my gratitude to Prof. Aditya Guntuboyina. Most of my research activities can be attributed to a course Aditya taught, Theoretical Statistics (Stat. 210 B), in Spring, 2017. I have started working with Aditya immediately afterwards. Having insightful discussions with Aditya on research problems helped me a lot. While Kannan is mostly about big pictures, Aditya is interested in the technical side of things, and I have the good fortune to work with both of them. I got to experience the *best of both worlds*. Whenever I had a doubt, Aditya was always available for meeting, and even when my doubts are trivial or even sometimes stupid, he always encouraged me.

I would also like to acknowledge the contributions of the amazing faculties who inspired me over the years. I benefited a lot from taking fundamental courses on probability, optimization and statistics from Prof. Venkat Anantharam, Prof. Martin Wainwright and Prof. Aditya Guntuboyina. I would also like to take this opportunity to thank the members of my quals as well as thesis committee—Prof. Thomas Courtade, Prof. Anil Aswani, Prof. Kannan Ramchandran and Prof. Aditya Guntuboyina for giving me valuable feedback and support. I would also like to acknowledge my masters advisors, Prof. Anurag Kumar and Prof. Aditya Gopalan of Indian Institute of Science, Bangalore who believed in me and gave me the opportunity to work in his group, and fortified my interest in academic research by being an amazing advisor and person alike.

Over my graduate studies, I had the privilege to work with extraordinarily talented people on research projects, and this thesis would not have been possible without their integral contributions. In my initial years, I started collaborating with Dong Yin and Ashwin Pananjady. I learnt a lot from them, and had enlightening discussions. In the later years, I had the good fortune of working with Prof. Arya Mazumdar, Raj Kumar Maity and Abishek Sankararaman. I really enjoyed our virtual brainstorming sessions, and their deep knowledge on areas like Distributed Learning and Bandits really helped me a lot. Furthermore, I have had the good fortune to work with students and postdocs

in and outside Berkeley—Swanand Kadhe, Sinho Chewi, Jichan Chung, Sayak Ray Chowdhury, Forest Yang, Abhay Parekh, Vipul Gupta and Justin Hong to name a few. I would also like to thank my mentor Dean Foster at Amazon SCOT, with whom I interned in the summer of 2020.

PhD is a long journey, and one definitely needs friends to make this journey smooth. I would like to thank my awesome friends at BLISS Lab—Vipul Gupta, Soham Phade, Dong Yin, Nived Rajaraman, Abishek Sankararaman, Banghua Zhu, Orhan Ocal, Vidya Muthukumar, Ashwin Pananjady, Payam Delgosh, Raaz Dwivedi, Koulik Khamaru, Sang Min Han, Jichan Chung, Kuan-Yun Lee, Amirali Aghazadeh, Wenglong Mou, Efe Aras, Banghua Zhu and Yaoqing Yang. They were great sources of support, knowledge, and friendship. Engaging coffee breaks and delightful discussions with them always lightened my mood. I will miss going to impromptu hiking trips and dinners with my friends at BLISS. I would also like to thank my friends outside BLISS—Sourav Ghosh, Hari Prasanna Das, Santanu Maity, Priyanka Roy and Milind Hedge, with whom I had an awesome time. Furthermore, I would also like to take this opportunity to have a special shout out to the friendly and extremely helpful staff of EECS at Berkeley. In particular, Shirley Salanio has made my (and everybody's) life at EECS better; whenever I needed help, I knew I could get timely support Shirley that would solve any bureaucratic issue I had, and also hear her nice words, which would brighten any day.

Last but not the least, I would like to thank my parents Dipika Ghosh and Asit Ghosh, to whom I dedicate this thesis. Their love and support always stayed with me, and worked as a driving force in my PhD journey. They taught me the importance of education and always motivated me to pursue a career in academics. I owe all my achievements to them and their unconditional love and support.

Chapter 1

Introduction

In recent years, we are witnessing an unprecedented growth in the amount of high dimensional data being sensed, collected and processed to drive modern applications across many domains. This growth has challenged classical algorithms to scale up to big-data regimes. Furthermore, in time sensitive applications such as control and decision-making, collaborative learning, and medical imaging, efficient implementations of such *scaled-up* algorithms are necessary. Moreover, in applications like personalized recommendation and advertisement placement, adapting to the structure of the problem is of fundamental interest. This *adaptation* can reduce both the problem complexity and compute time. In this thesis, we focus on three fundamental aspects of machine learning: *statistical guarantees*, *efficient implementations* and *adaptation to problem complexity*. We consider three different frameworks in which data is presented to the learner: (a) Federated (Distributed) Learning (FL), (b) supervised batch learning and (c) online (streaming) setup with Bandit/Reinforcement Learning (RL).

1.1 Algorithms for Federated (Distributed) Learning (FL)

In many real-world applications, the size of training datasets has grown significantly over the years to the point that it is becoming crucial to implement learning algorithms in a distributed fashion. A commonly used distributed learning framework is data parallelism, in which large-scale datasets are distributed over multiple *worker machines* for parallel processing in order to speed up computation. In many applications, data are stored in end users' own devices such as mobile phones and personal computers, and in these applications, fully utilizing the on-device machine intelligence is an important direction for next-generation distributed learning. Federated Learning (FL) [1–3] is a recently proposed distributed computing paradigm that is designed towards this goal, and has received significant attention. Many statistical and computational challenges arise in Federated Learning, due to the highly decentralized system architecture.

In a standard Federated (or distributed) learning framework, a set of worker machines store the data, perform local computations, and communicate local updates (ex. gradients, Hessians) to the center machine (e.g., a parameter server). The center machine processes the results from workers

to update the model parameters. Such distributed frameworks need to address the following three fundamental challenges:

1. *Data Heterogeneity*: First, the heterogeneity of data across machines can often dampen the learning rate. This is a major concern in Federated Learning, since the worker machines are end users' personal devices and the data across users are different (in distribution) from one another.
2. *Communication Cost*: Second, the gains due to parallelization are often bottlenecked in practice by heavy communication overheads between workers and the central machine. This is especially the case for large clusters of worker machines or for modern deep learning applications using models with millions of parameters (for example, NLP models, such as BERT [4], may have well over 100 million parameters). Moreover, in Federated Learning, communication from a user device to the central server is directly tied to the user's upload bandwidth costs. Therefore, it is of paramount importance to reduce communication overhead in distributed learning algorithms.
3. *Adversarial Robustness*: Third, messages from workers are susceptible to errors due to hardware faults or software bugs, stalled computations, data crashes, and unpredictable communication channels. In scenarios such as Federated Learning, users may as well be malicious and act adversarially. The inherent unpredictable (and potentially adversarial) nature of compute units is typically modeled as *Byzantine failures*. Even if a single worker is Byzantine, it can be fatal to most learning algorithms ([5]).

In Chapters 2-5, we address these fundamental issues and propose algorithms that are provable and computation friendly. In Chapter 2, we propose an iterative clustering algorithm that partitions the worker machines and learns the optimal model for each clusters simultaneously, yielding an optimal statistical rate in certain settings. In Chapter 3 we propose first order optimization algorithms to deal with communication efficiency and Byzantine attacks, and show the optimality of our proposed scheme in certain parameter regime. Furthermore, in Chapters 4 and 5, we extend this to second order and beyond second order algorithms.

1.1.1 Data Heterogeneity: Clustered FL

Exploiting data heterogeneity is particularly crucial in applications such as recommendation systems and personalized advertisement placement, and it benefits both the users' and the enterprises. For example, mobile phone users who read news articles may be interested in different categories of news like politics, sports or fashion; advertisement platforms might need to send different categories of ads to different groups of customers. These indicate that leveraging the heterogeneity among the users is of potential interest—on the one hand, each machine itself may not have enough data and thus we need to better utilize the similarity among the users; on the other hand, if we treat the data from all the users as i.i.d. samples, we may not be able to provide personalized predictions. This problem has recently received much attention [6–8].




Figure 1.1: An overview of IFCA with 2 clusters, orange and green: (a) The server broadcast models. (b) Worker machines identify their cluster memberships and run local updates. (c) The worker machines send back the local models to server. (d) Average the models within the same estimated cluster S_j .

In Chapter 2, we study one of the formulations of FL with non-i.i.d. data, i.e., the *clustered Federated Learning* [7, 9]. We assume that the users are partitioned into different clusters; for example, the clusters may represent groups of users interested in politics, sports, etc, and our goal is to train models for every cluster of users. We note that cluster structure is very common in applications such as recommender systems [10, 11]. The main challenge of our problem is that the *cluster identities of the users are unknown*, and we have to simultaneously solve two problems: identifying the cluster membership of each user and optimizing each of the cluster models in a distributed setting.

We assume that there are k different data distributions, $\mathcal{D}_1, \dots, \mathcal{D}_k$, and that the m machines are partitioned into k disjoint clusters, S_1^*, \dots, S_k^* . We assume no knowledge of the cluster identity of each machine, i.e., the partition S_1^*, \dots, S_k^* is not revealed to the learning algorithm. We assume that every worker machine $i \in S_j^*$ contains n i.i.d. data points $z^{i,1}, \dots, z^{i,n}$ drawn from \mathcal{D}_j , where each data point $z^{i,j}$ consists of a pair of feature and response denoted by $z^{i,\ell} = (x^{i,\ell}, y^{i,\ell})$.

Let $f(w; z) : \mathcal{W} \rightarrow \mathbb{R}$ be the loss function associated with data point z , where $\mathcal{W} \subseteq \mathbb{R}^d$ is the parameter space. We choose $\mathcal{W} = \mathbb{R}^d$. Our goal is to minimize the population loss function

$$F^j(w) := \mathbb{E}_{z \sim \mathcal{D}_j} [f(w; z)]$$

for all $j \in [k]$. In particular, we try to find solutions $\{\hat{w}_j\}_{j=1}^k$ that are close to $w_j^* = \operatorname{argmin}_{w \in \Theta} F^j(w)$, $j \in [k]$.

In order to achieve this goal, we propose a framework and analyze a distributed method, named the *Iterative Federated Clustering Algorithm (IFCA)* for clustered FL. The basic idea of our algorithm is a strategy that alternates between estimating the cluster identities and minimizing the loss functions, and thus can be seen as an Alternating Minimization algorithm in a distributed

setting. The details of the algorithm is given in Figure 1.1 for two clusters, green and orange. With the broadcasted models, the workers perform an objective based clustering to estimate their cluster identity, and correspondingly update the model by taking gradient steps. Furthermore, when the cluster structure is ambiguous, we propose to leverage the weight sharing technique in multi-task learning [12] and combine it with IFCA. More specifically, we learn the shared representation layers using data from all the users, and use IFCA to train separate final layers for each individual cluster.

We further establish convergence rates of our algorithm, for strongly convex losses under the assumption of good initialization. We prove exponential convergence speed, and for both settings, we can obtain *near optimal* statistical error rates in certain regimes. With quadratic loss, we show that the statistical error rate of our problem is $\tilde{O}(\sqrt{\frac{d}{mn}} + \text{minor-term})$, which is near optimal (see [13]). For strongly convex loss, the error rate obtained is $\tilde{O}(\sqrt{\frac{d}{mn}} + \frac{1}{n})$, which is sub-optimal by a factor of $\frac{1}{n}$. Furthermore, extensive experiments on MNIST, CIFAR-10 and Federated EMNIST data-sets validate our theoretical findings.

1.1.2 Communication Efficiency and Byzantine Resilience

In Chapters 3-5, we address the challenges of communication efficiency and Byzantine resilience (against adversarial machines) simultaneously. As mentioned earlier, Chapter 3 introduces first order gradient based methods, and Chapters 4 and 5, extends those to second order and beyond second order methods respectively.

Gradient Based (First Order) Methods

Both the challenges of communication efficiency and Byzantine-robustness, have recently attracted significant research attention, albeit mostly separately. In particular, several recent works have proposed various quantization or sparsification techniques to reduce the communication overhead ([14–16]). The goal of these quantization schemes is to compute an unbiased estimate of the gradient with bounded second moment in order to achieve good convergence guarantees. The problem of developing Byzantine-robust distributed algorithms has been considered in [17–19].

Once again, in Chapter 3, we consider m worker machines, each storing n data points. The data points are generated from some unknown distribution \mathcal{D} . The objective is to learn a parametric model that minimizes a non-convex population loss function $F : \mathcal{W} \rightarrow \mathbb{R}$, where F is defined as an expectation over \mathcal{D} , and $\mathcal{W} \subseteq \mathbb{R}^d$ denotes the parameter space. For gradient compression at workers, we consider the notion of a δ -approximate compressor from [20] defined below:

Definition 1 (δ -Approximate Compressor). *An operator $\mathcal{Q}(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is defined as δ -approximate compressor on a set $\mathcal{S} \subseteq \mathbb{R}^d$ if, $\forall x \in \mathcal{S}$,*

$$\|\mathcal{Q}(x) - x\|^2 \leq (1 - \delta)\|x\|^2,$$

where $\delta \in (0, 1]$ is the compression factor.

Examples of such compressors include sign-based compressors like QSGD ([21]), ℓ_1 -QSGD ([20]) and top- k sparsification ([15]). Furthermore, for Byzantine robustness, we assume that $0 \leq \alpha < 1/2$ fraction of the worker machines are Byzantine, and make no assumptions on the behavior of the Byzantine workers. Byzantine workers can send any arbitrary values to the center machine. In addition, they may completely know the learning algorithm and are allowed to collude with each other. In contrast to blind multiplicative adversaries assumed in literature (see [22]), we consider unrestricted adversaries.

We propose a communication-efficient and robust distributed gradient descent (GD) algorithm. The algorithm takes as input the gradients compressed using a δ -approximate compressor along with their norms, and performs a simple thresholding operation on based on gradient norms to discard $\beta > \alpha$ fraction of workers with the largest norm values. We show that our algorithm achieves the following statistical error rate¹ for the regime $\delta > 4\beta + 4\alpha - 8\alpha^2 + 4\alpha^3$:

$$\tilde{\mathcal{O}} \left(d^2 \left[\frac{\alpha^2}{n} + \frac{1-\delta}{n} + \frac{1}{mn} \right] \right). \quad (1.1)$$

We first note that when $\delta = 1$ (uncompressed), the error rate is $\tilde{\mathcal{O}}(d^2[\frac{\alpha^2}{n} + \frac{1}{mn}])$, which matches the optimal rate of [23]. Furthermore, for a fixed d and the compression factor δ satisfying $\delta \geq 1 - \alpha^2$, the rate $\tilde{\mathcal{O}}(\frac{\alpha^2}{n} + \frac{1}{mn})$, is optimal and order-wise identical to the case of no compression [23]. In other words, in this parameter regime, we get compression for free.

Furthermore, we strengthen our learning algorithm by using error feedback to correct the direction of the local gradient (Algorithm 3). We show (both theoretically and via experiments) that using error-feedback with a δ -approximate compressor indeed speeds up the convergence rate and attains better (statistical) error rate.

Distributed Approximate Newton Method

An alternative way to reduce communication cost is to use second order optimization algorithms; which are known to converge much faster than their first order counterparts. At the expense of increasing computation cost at the worker machines (which is not an issue in applications like Federated Learning; see [3]), the second order algorithms reduce the number of iterations between the worker and center machines. Indeed, a handful of algorithms has been developed using this philosophy, such as DANE [24], DISCO [25], GIANT [26], DINGO [27], Newton-MR [28], INEXACT DANE and AIDE [29].

One major drawback of the above-mentioned algorithms is that they are prone to adversarial or Byzantine attacks. On the other hand, the Byzantine-robust distributed algorithms (see [18, 19, 23, 30–33]) are all variants of first order gradient based optimization algorithm.

In Chapter 4, we propose COMRADE, a distributed approximate Newton-type algorithm that communicates less and is resilient to Byzantine workers. Specifically, with m worker machines and one center machine, we minimize a regularized convex loss $f : \mathbb{R}^d \rightarrow \mathbb{R}$, which is additive over the available data points. Furthermore, similar to previous chapters, we assume that α fraction of the

¹Note that $\mathcal{O}(\cdot)$ hides multiplicative constants, while $\tilde{\mathcal{O}}(\cdot)$ further hides logarithmic factors.

worker machines are Byzantine, where $\alpha \in [0, 1/2)$. To the best of our knowledge, this is the first work that addresses the problem of Byzantine resilience in second order optimization.

In our proposed algorithm, the worker machines communicate *only once* per iteration with the center machine. This is in sharp contrast with the state-of-the-art distributed second order algorithms (like GIANT [26], DINGO [27], Determinantal Averaging [34]), which sequentially estimates functions of local gradients and Hessians and communicate them with the center machine. In this way, they end up communicating twice per iteration with the center machine. We show that this sequential estimation is redundant. Instead, in COMRADE, the worker machines only send a d dimensional vector, the product of the inverse of local Hessian and the local gradient. Hence, in this way, we save $\mathcal{O}(d)$ bits of communication per iteration. Furthermore, in Section 4.5, we argue that, in order to cut down further communication, the worker machines can even compress the local Hessian inverse and gradient product, using a δ -approximate compressor.

For Byzantine resilience, COMRADE employs a simple thresholding policy on the norms of the local Hessian inverse and local gradient product to discard $\beta > \alpha$ fraction of workers having largest norm values. Since the norm of the Hessian-inverse and gradient product determines the *amount* of movement for Newton-type algorithms, this norm corresponds to a natural metric for identifying and filtering out Byzantine workers.

We prove the linear-quadratic rate of convergence of our proposed algorithm for strongly convex loss functions. In particular, suppose each worker machines contain s data points; and let $\Delta_t = w_t - w^*$, where w_t is the t -th iterate of COMRADE, and w^* is the optimal model we want to estimate. In 4 (Theorem 2), we show that

$$\|\Delta_{t+1}\| \leq \max\{\Psi_t^{(1)}\|\Delta_t\|, \Psi_t^{(2)}\|\Delta_t\|^2\} + (\Psi_t^{(3)} + \alpha)\sqrt{\frac{1}{s}},$$

where $\{\Psi_t^{(i)}\}_{i=1}^3$ are quantities dependent on several problem parameters. Notice that the above implies a quadratic rate of convergence when $\|\Delta_t\| \geq \Psi_t^{(1)}/\Psi_t^{(2)}$. Subsequently, when $\|\Delta_t\|$ becomes sufficiently small, the above condition is violated and the convergence slows down to a linear rate. The error-floor, which is $\mathcal{O}(1/\sqrt{s})$ comes from the Byzantine resilience subroutine in conjunction with the simultaneous estimation of Hessian and gradient.

Cubic Regularized Newton Method

In Chapter 5, we address the problem of saddle point avoidance in non-convex optimization in a Federated Learning (FL) framework. In order to fit complex machine learning models, one often requires to find local minima of a non-convex loss $f(\cdot)$, instead of critical points only, which may include several saddle points. Training deep neural networks and other high-capacity learning architectures [35, 36] are some of the examples where finding local minima is crucial. [36, 37] shows that the stationary points of these problems are in fact saddle points and far away from any local minimum, and hence designing efficient algorithm that escapes saddle points is of interest. Moreover, in [38, 39], it is argued that saddle points can lead to highly sub-optimal solutions in many problems of interest. This issue is amplified in high dimension as shown in [40], and becomes

the main bottleneck in training deep neural nets. Furthermore, a line of recent work [39, 41, 42], shows that for many non-convex problems, it is sufficient to find a local minimum.

The issue of saddle point avoidance becomes non-trivial in the presence of Byzantine workers. Since we do not assume anything on the behavior of the Byzantine workers, it is certainly conceivable that by appropriately modifying their messages to the center, they can create *fake local minima* that are close to the saddle point of the loss function $f(\cdot)$, and these are far away from the true local minima of $f(\cdot)$. This is popularly known as the *saddle-point attack* (see [31]), and it can arbitrarily destroy the performance of any non-robust learning algorithm.

In Chapter 5, we propose a communication efficient distributed optimization algorithm that escapes the saddle points as well as resists the Byzantine attacks simultaneously. In particular, we consider a variation of the famous cubic-regularized Newton algorithm of Nesterov and Polyak [43]. Being a second order algorithm, it converges very fast compared to its first order counterparts (see [43]). Also, in [43–45], it is shown that cubic-regularized Newton can efficiently escape the saddle points of a non-convex function, by pushing the Hessian towards a positive semi-definite matrix.

A point w is said to satisfy the ϵ -second order stationary condition of the loss function $f(\cdot)$ if,

$$\|\nabla f(w)\| \leq \epsilon \quad \lambda_{\min}(\nabla^2 f(w)) \geq -\sqrt{\epsilon}.$$

$\nabla f(w)$ denotes the gradient of the function and $\lambda_{\min}(\nabla^2 f(w))$ denotes the minimum eigenvalue of the Hessian of the function. Hence, under the assumption (which is standard in the literature, see [31, 46]) that all saddle points are strict (i.e., $\lambda_{\min}(\nabla^2 f(w_s)) < 0$ for any saddle point w_s), all second order stationary points (with $\epsilon = 0$) are local minima, and hence converging to a stationary point is equivalent to converging to a local minima.

We consider a distributed variant of the cubic regularized Newton algorithm. In this scheme, the center machine asks the workers to solve an auxiliary function and return the result. The center machine aggregates the solution of the worker machines and takes a descent step. Furthermore, we use a simple norm-based thresholding approach to robustify the distributed cubic-regularized Newton method. We prove that the algorithm convergence at a rate of $\frac{1}{T^{2/3}}$, which is faster than the first order methods (which converge at $1/\sqrt{T}$ rate, see [31]). Hence, the number of iterations (and hence the communication cost) required to achieve a target accuracy is much fewer than the first order methods. Experimental results on LIBSVM ([47]) datasets validate our theoretical findings.

1.2 Learning in Supervised Batch Setup

In the batch setup, the learner has data and labels to begin with (standard supervised learning framework). Here, we work with a computationally efficient and statistically sound Alternating Minimization (AM) algorithm, typically used to solve non-convex problems. In particular, we apply AM to a non-convex problem, namely max-affine regression and obtain a near-optimal statistical rate. Max-affine regression refers to a model where the unknown regression function is modeled as a maximum of k unknown affine functions for a fixed $k \geq 1$. This generalizes linear regression and (real) phase retrieval, and is closely related to convex regression. Working within a non-asymptotic framework, we study this problem in the high-dimensional setting assuming that k is a fixed constant,

and focus on the estimation of the unknown coefficients of the affine functions underlying the model. In Chapters 6 and 7, we analyze max-affine regression under 2 designs respectively: Gaussian and Small-ball.

Formally, max-affine regression implies the following regression model:

$$Y = \max_{1 \leq j \leq k} (\langle X, \theta_j^* \rangle + b_j^*) + \epsilon \quad (1.2)$$

where Y is a univariate response, X is a d -dimensional vector of covariates and ϵ models zero-mean noise that is independent of X . We assume that $k \geq 1$ is a known integer and study the problem of estimating the unknown parameters $\theta_1^*, \dots, \theta_k^* \in \mathbb{R}^d$ and $b_1^*, \dots, b_k^* \in \mathbb{R}$ from independent observations $(x_1, y_1), \dots, (x_n, y_n)$ drawn according to the model (1.2).

Immediately, one may observe that when $k = 1$, equation (1.2) corresponds to the classical linear regression model. Also, when $k = 2$, the intercepts $b_2^* = b_1^* = 0$, and $\theta_2^* = -\theta_1^* = \theta^*$, we have

$$Y = |\langle X, \theta^* \rangle| + \epsilon. \quad (1.3)$$

The problem of recovering θ^* from observations drawn according to the above model is known as (real) phase retrieval—variants of which arise in a diverse array of science and engineering applications [48–51]. Furthermore, since $x \mapsto \max_{1 \leq j \leq k} (\langle x, \theta_j^* \rangle + b_j^*)$ is always a convex function, the model (1.2) serves as a parametric approximation to the non-parametric convex regression model

$$Y = \phi^*(X) + \epsilon, \quad (1.4)$$

where $\phi^* : \mathbb{R}^d \rightarrow \mathbb{R}$ is an unknown convex function. It is well known that convex regression suffers from the curse of dimensionality (see, e.g., [52–54]). Hence, one would need to work with more structured sub-classes of convex functions. Since convex functions can be approximated to arbitrary accuracy by maxima of affine functions, it is reasonable to enforce a regularization on the problem by considering only those convex functions that can be written as a maximum of a fixed number of affine functions.

With the augmented notation $\beta_j^* := (\theta_j^*, b_j^*) \in \mathbb{R}^{d+1}$ for $j = 1, \dots, k$ and (ξ_i, y_i) for $i = 1, \dots, n$, where $\xi_i := (x_i, 1) \in \mathbb{R}^{d+1}$, we minimize the least squares objective

$$(\widehat{\beta}_1^{(\text{ls})}, \dots, \widehat{\beta}_k^{(\text{ls})}) \in \underset{\beta_1, \dots, \beta_k \in \mathbb{R}^{d+1}}{\operatorname{argmin}} \sum_{i=1}^n \left(y_i - \max_{1 \leq j \leq k} \langle \xi_i, \beta_j \rangle \right)^2. \quad (1.5)$$

In order to do this, we use alternating minimization (AM) algorithm of [55], which is an iterative algorithm for estimating the parameters $\beta_1^*, \dots, \beta_k^*$. In the t -th iteration of the algorithm, the current estimates $\beta_1^{(t)}, \dots, \beta_k^{(t)}$ are used to partition the observation indices $1, \dots, n$ into k sets $S_1^{(t)}, \dots, S_k^{(t)}$ such that $j \in \operatorname{argmax}_{u \in [k]} \langle \xi_i, \beta_u^{(t)} \rangle$ for every $i \in S_j^{(t)}$. For each $1 \leq j \leq k$, the next estimate $\beta_j^{(t+1)}$ is then obtained by performing a least squares fit (or equivalently, linear regression) to the data $(\xi_i, y_i), i \in S_j^{(t)}$. Figure 1.2 pictorially represents the working of the AM algorithm.




Figure 1.2: Alternating Minimization in action: in the first phase, it disambiguates the max index and partitions the observations in k buckets. Within each bucket, AM solves a linear regression problem to obtain the next iterate.

We show (in Chapter 6) that for each $\epsilon > 0$, the parameter estimates $\beta_1^{(t)}, \dots, \beta_k^{(t)}$ returned by the AM algorithm at iteration t satisfy, with high probability, the inequality

$$\sum_{j=1}^k \|\beta_j^{(t)} - \beta_j^*\|^2 \leq \epsilon + C(\beta_1^*, \dots, \beta_k^*) \frac{\sigma^2 kd}{n} \log(kd) \log\left(\frac{n}{kd}\right) \quad (1.6)$$

for every $t \geq \log_{4/3} \left(\frac{\sum_{j=1}^k \|\beta_j^{(0)} - \beta_j^*\|^2}{\epsilon} \right)$, provided that the sample size n is sufficiently large and that the initial estimates satisfy the condition

$$\min_{c>0} \max_{1 \leq j \leq k} \|c\beta_j^{(0)} - \beta_j^*\|^2 \leq \frac{1}{k} c(\beta_1^*, \dots, \beta_k^*). \quad (1.7)$$

Here $C(\beta_1^*, \dots, \beta_k^*)$ and $c(\beta_1^*, \dots, \beta_k^*)$ are constants depending only on true parameters $\beta_1^*, \dots, \beta_k^*$.

In Chapter 7, we relax the assumption of Gaussian covariates and show that a similar phenomenon occurs under significantly weaker statistical assumptions. In particular, we allow the distribution of the covariates to come from the larger class of sub-Gaussian distributions that satisfy a *small-ball* condition. In addition, we also consider the scenario of *universal* parameter estimation, meaning that our guarantees hold uniformly over all $\beta_1^*, \dots, \beta_k^*$. This allows the parameters to be chosen with knowledge of the realized covariates, a robust setting that is commonly studied in signal processing applications like phase retrieval [56]. Our covariate assumption relies on the following definition.

Definition 2. (*Small-ball*) A distribution P_X satisfies a (ζ, c_s) -small-ball property if, for $X \sim P_X$ and each $\delta > 0$, we have

$$\sup_{u \in \mathbb{S}^{d-1}, w \in \mathbb{R}} \Pr \{ (\langle X, u \rangle + w)^2 \leq \delta \} \leq (c_s \delta)^\zeta. \quad (1.8)$$

The small-ball properties of various classes of distributions have been studied extensively in the probability literature [57, 58], and many natural distributions possess this property provided they are not too “peaky”. We now present our assumption on the covariate distribution;

Assumption 1. *The covariates are drawn from a distribution, P_X , which is isotropic, η -sub-Gaussian, and satisfies a (ζ, c_s) small-ball condition.*

We now provide a few examples, where covariates satisfy the above assumption:

(a) $\text{Unif}[-\sqrt{3}, \sqrt{3}]$, (b) the standard Gaussian and (c) any random variable with density $f(x) \propto e^{-\|x\|^c}$ for a positive constant c . In Chapter 7, we discuss these examples in detail and obtain the parameters (ζ, c_s) and η .

Applying the same AM algorithm, we obtain guarantees similar to that of Chapter 6. As a special instance, for phase retrieval problem, we obtain some new results. To the best of our knowledge, all previous results on the AM algorithm for phase retrieval [59, 60] only held under the assumptions of Gaussian covariates and noiseless observations, and/or required resampling of the measurements [61]. Ours is thus the first work to handle non-Gaussian covariates in the presence of noise, while also analyzing the algorithm without resampling.

1.3 Efficient and Adaptive Algorithms for Bandit and Reinforcement Learning (RL)

In the online learning framework, data is presented to the user in a streaming fashion, and the job of the learner is to balance between exploration and exploitation judiciously to enforce statistical learning. In this part of the thesis (Chapters 8-9), we provide statistically tractable and efficient algorithms for online learning. In particular, we are interested in the question of model selection or adaptation in online learning. Model selection has been studied in supervised offline setup to a great extent (via popular methods like k -fold cross validation, which has theoretical guarantees as well as practical effectiveness). However, in online learning, the question of model selection has received little interest. Our main contribution in this part of the thesis is to propose an efficient and provable algorithm that adapts to the problem complexity automatically. We study the model selection problem in 3 settings: stochastic linear bandits, contextual bandits (beyond linear structure) and finally Reinforcement learning with linear structure.

In Chapter 8, we consider the standard setup of stochastic linear bandits, where, at each round, having observed a context vector, the learner chooses an action and observes a noisy linearly parameterized reward. We propose a provable, efficient algorithm that captures the complexity of the problem and automatically adapts to it. Via regret analysis, we show optimality of our scheme,

showing that the cost of adaptation is a small additive constant. In Chapter 9, we extend these to the problem of generic contextual bandit beyond linear structure.

Model Selection

Model Selection for Bandits(or Reinforcement) Learning, refers to choosing the appropriate hypothesis class, to model the mapping from arms (actions) to expected rewards. Model selection plays an important role in applications such as personalized recommendations, (see the motivating example in Chapter 8). Moreover, in Reinforcement Learning with huge state and action spaces, it is empirically observed that exploring only a small subset of the state and action spaces are often sufficient to obtain high rewards. Hence, in these applications one needs an adaptive algorithm that identifies the implicit structure of the problem and explores or exploits judiciously.

Formally, a family of nested hypothesis classes $\mathcal{H}_f, f \in \mathcal{F}$ needs to be specified, where each class posits a plausible model for mapping arms (or actions) to expected rewards. The true model is assumed to be contained in the family \mathcal{F} which is totally ordered, where if $f_1 \leq f_2$, then $\mathcal{H}_{f_1} \subseteq \mathcal{H}_{f_2}$. Model selection guarantees then refers to algorithms whose regret scales in the complexity of the *smallest hypothesis class containing the true model*, even though the algorithm was not aware apriori.

1.3.1 Adaptation in Stochastic Linear Bandits

We consider the problem of *model selection* for two popular stochastic linear bandit settings, and propose algorithms that adapts to the *unknown* problem complexity. In the first setting, we consider the K armed mixture bandits, where the mean reward of arm $i \in [K]^2$, is

$$\mu_i + \langle \alpha_{i,t}, \theta^* \rangle$$

where $\alpha_{i,t} \in \mathbb{R}^d$ is the known context vector of arm i at time t , and $\mu_i \in \mathbb{R}, \theta^* \in \mathbb{R}^d$ are unknown and needs to be estimated. This setting also contains the standard Multi-Armed Bandit (MAB) problem [62, 63] when $\theta^* = 0$.

Popular linear bandit algorithms, like LinUCB, OFUL (see [64–66]) handle the case with no bias ($\mu_i = 0$), while OSOM [67], the recent improvement can handle arm-bias. Implicitly, all the above algorithms assume an upper bound on the norm of $\|\theta^*\| \leq L$, which is supplied as an input. Crucially however, the regret guarantees scale linearly in the upper bound L . In contrast, we choose $\|\theta^*\|$ as the problem complexity and consider a sequence of nested hypothesis classes, each positing a different upper bound on $\|\theta^*\|$. We propose Adaptive Linear Bandit (ALB), a novel phase based algorithm, that, without any upper bound on the norm $\|\theta^*\|$, *adapts to the true complexity* of the problem instance, and achieves a regret scaling linearly in the true norm $\|\theta^*\|$. The details of this algorithm is given in Figure 1.3. It consistently estimates upper-bounds on parameter norm over a confidence ellipsoid given by [67], and we show that the size of this confidence set shrinks with time and as a result, the norm estimates converge to the true norm estimates.

²By $[r]$, we denote the set of positive integers $\{1, 2, \dots, r\}$.




Figure 1.3: Model selection with norm refinement. At the end of each epoch, an over-estimate of norm is computed, and with shrinking confidence set, these norm estimates converge to the true norm of θ^* .

As a corollary, our algorithm’s performance matches the minimax regret of simple MAB when $\theta^* = 0$, even though the algorithm did not apriori know that $\theta^* = 0$. Intuitively, if θ^* is very small, the contribution of the contextual term will be low, and on the other hand, a large non zero θ^* , the contextual part dominates. Hence, $\|\theta^*\|$ can be chosen as a natural complexity parameter. In Chapter 8, we also provide motivations and examples behind the choice of $\|\theta^*\|$ as a problem complexity parameter, as it balances between the standard MAB and the linear bandit problems.

We show that ALB achieves the regret

$$R(T) = \tilde{O}(\|\theta^*\| \sqrt{T}),$$

where $\|\theta^*\|$ is apriori unknown, and T is the duration of the interaction between the player and the environment, known as the horizon. As a corollary, when $\theta^* = 0$, ALB recovers the minimax regret for the simple bandit algorithm without such knowledge of θ^* . ALB is the first algorithm that uses parameter norm as model section criteria for linear bandits. Prior state of art algorithms [67] achieve a regret of $\tilde{O}(L\sqrt{T})$, where L is the upper bound on $\|\theta^*\|$, fed as an input to the problem.

In the second setting, we consider is the standard linear stochastic bandit [66] with possibly an infinite number of arms, where the mean reward of any arm $x \in \mathbb{R}^d$ (arms are vectors in this case) given by $\langle x, \theta^* \rangle$, where $\theta^* \in \mathbb{R}^d$ is unknown. For this setting, we consider model selection from among a total of d different hypothesis classes, with each class positing a different cardinality for the support of θ^* . The sparsity of θ^* , denoted by $d^* \leq d$, is unknown to the algorithm, and we define d^* as the problem complexity We exhibit a novel algorithm, where the regret scales linearly in the unknown cardinality of the support of θ^* . The regret scaling of our algorithm matches that of an oracle that has knowledge of the optimal support cardinality [68],[69], thereby achieving model selection guarantees. Our algorithm is the first known algorithm to obtain regret scaling matching

that of an oracle that has knowledge of the true support. This is in contrast to standard linear bandit algorithms such as [66], where the regret scales linearly in d . We also extend this methodology to the case when the number of arms is finite (see [70]) and obtain similar regret rates matching the oracle. We further verify through synthetic and real-data experiments that the performance gains are fundamental and not artifacts of mathematical bounds. In particular, we show 1.5 – 3x drop in cumulative regret over non-adaptive algorithms.

1.3.2 Adaption for Contextual Bandits Beyond Linear Structure

Here, we focus on the main contribution of the paper—a provable model selection guarantee for the (generic) stochastic contextual bandit problem. The interaction between the learner and the environment can be modelled as following:

Let \mathcal{A} be the set of K actions, and let \mathcal{X} be the set of d dimensional contexts. At time t , Nature picks (x_t, r_t) in an i.i.d fashion, where $x_t \in \mathcal{X}$ and a context dependent $r_t \in [0, 1]^K$. Upon observing the context, the learner takes action $a_t \in \mathcal{A}$, and obtains the reward of $r_t(a_t)$. We restrict to the class of realizable rewards defined as:

Assumption 2. (Realizability:) *There exists a predictor $f^* \in \mathcal{F}$, such that $\mathbb{E}[r_t(a)|x] = f^*(x, a)$, for all x and a .*

Note that the realizability assumption is standard in the literature ([70, 71]).

In the contextual bandit literature ([71, 72]) it is generally assumed that the true regression function f^* is unknown, but we know the function class \mathcal{F} where it belongs. Hence, we pay some price (denoted by the regret) for this. To set up notation, for any $f \in \mathcal{F}$, we define a policy induced by the function, $\pi_f(x) = \operatorname{argmin}_{a \in \mathcal{A}} f(x, a)$. So, we need to compete with the policy induced by the true regressor π_{f^*} . We define the regret over T rounds as the following:

$$R(T) = \sum_{t=1}^T [r_t(\pi_{f^*}(x_t)) - r_t(a_t)].$$

However, in contrast to the standard setting, in the model selection framework does not assume the knowledge of \mathcal{F} . Instead, we are given a nested class of M function classes, $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \dots \subset \mathcal{F}_M$. The smallest function class where the true regressor lies is denoted by \mathcal{F}_{d^*} , where $d^* \in [M]$. Hence, the regret of the contextual bandit algorithm should depend on \mathcal{F}_{d^*} (and not the largest class \mathcal{F}_M). However, we do not know d^* , and our goal is to propose adaptive algorithms such that the regret depends on the *actual* problem complexity \mathcal{F}_{d^*} .

In order to achieve this, we propose an algorithm, namely Adaptive Contextual Bandits (ACB). Our algorithm selects the correct model via successive refinements over multiple (doubling) epochs. We use the FALCON algorithm of [71], and add a model selection phase at the beginning of each epoch. In other words, over multiple epochs, we successively refine our estimates of the *proper* model class where the true regressor function f^* lies. We show that ACB selects the correct model class with high probability and attains a regret of

$$R(T) \lesssim \sqrt{KT \log(|\mathcal{F}_{d^*}|T/\delta)} + \text{minor-term},$$

where K is the number of actions, and $|\cdot|$ denotes the cardinality. Although the above expression makes sense only with finite function classes, a simple extension (see Chapter 9) to infinite function classes is possible. Note that, this regret scaling is optimal (see [71]). The cost of model selection therefore is a minor additive term in the regret expression.

1.3.3 Model Selection for Reinforcement Learning

We study the problem of model selection in Reinforcement Learning with function approximation. We restrict our attention to the framework of model based episodic *linear* MDP. In particular, similar to Chapter 8, we study both the *norm* and *sparsity/dimension* of the problem parameter as problem complexity and obtain model selection guarantees, meaning that the performance (regret) of our algorithm depends on the complexity of the problem instance, as compared to some predefined upper-bound on the complexity parameters.

In this chapter, we consider an heterogeneous, episodic Markov decision process denoted by a tuple $M(\mathcal{S}, \mathcal{A}, H, \{\mathbb{P}_h\}_{h=1}^H, \{r_h\}_{h=1}^H)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, H is the length of each episode, \mathbb{P}_h is the transition probability at step h such that $\mathbb{P}_h(\cdot|s, a)$ is the distribution over states provided action a is taken for state s at step h , and $r_h : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the deterministic reward function at step h . A policy π is a collection of H functions $\{\pi_h : \mathcal{S} \rightarrow \mathcal{A}\}_{h \in [H]}$, where each of them maps a state s to an action a .

The RL agent interacts with the environment for K episodes to learn the unknown transition kernels $\{\mathbb{P}_h\}$ and hence to improve its policy. For each $k \geq 1$, the environment picks a starting state s_1^k , and the agent chooses a policy π^k that will be followed through the whole k -th episode. The objective is to design a learning algorithm that constructs a sequence $\{\pi^k\}$ that minimizes the total regret

$$R(K) = \sum_{k=1}^K \left[V_1^*(s_1^k) - V_1^{\pi^k}(s_1^k) \right].$$

where $V_1(\cdot)$ is the value function at the beginning of each episode, and $V_1^*(\cdot)$ is the optimal value function. In this chapter, we consider a special class of MDPs called *linear kernel MDPs* (a.k.a., linear mixture MDPs) [73]. Roughly speaking, it means that the transition kernel $\{\mathbb{P}_h\}$ can be represented as a linear function of a given feature map $\phi : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^d$. Formally, we have the following definition.

Definition 3 (Linear kernel MDP). $M(\mathcal{S}, \mathcal{A}, H, \{\mathbb{P}_h\}_{h=1}^H, \{r_h\}_{h=1}^H)$ is called an *heterogeneous, episodic B -bounded linear kernel MDP* if there exists a known feature mapping $\phi : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^d$ and an unknown vector $\theta_h \in \mathbb{R}^d$ with $\|\theta_h\|_2 \leq B$ such that

(i) For any state-action-state triplet $(s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$ and step $h \in [H]$, $\mathbb{P}(s'|s, a) = \langle \phi(s'|s, a), \theta_h \rangle$.

(ii) For any bounded function $V : \mathcal{S} \rightarrow [0, 1]$ and any tuple $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have $\|\phi_V(s, a)\|_2 \leq 1$, where $\phi_V(s, a) := \sum_{s' \in \mathcal{S}} \phi(s'|s, a)V(s')$.

In the learning problem, the vectors $\{\theta_h^*\}_{h=1}^H$ are unknown to the learner. The episodic MDP is parameterized by $\Theta^* = \{\theta_h^*\}_{h=1}^H$ and we denote the MDP by M_{Θ^*} . In what follows, we define two natural complexity measure of M_{Θ^*} : (a) $\|\theta_h^*\|$ for all $h \in [H]$, and (b) $\|\theta_h^*\|_0$ for all $h \in [H]$.

Similar to Chapter 8, we propose algorithms that use successive refinement strategy to estimate the norm and sparsity of $\{\theta_h^*\}_{h=1}^H$, and adapts to it. As a result, we obtain regret guarantees that depend on the problem complexity parameter alone. We characterize the cost of model selection, and in both cases we obtain a regret of

$$R(K) \leq \mathcal{O}(\sqrt{K}) + \text{minor-term}$$

with high probability. We observe that $R(K)$ retains the \sqrt{K} regret, which is optimal (see [74]). However, the cost of model selection is the minor term added to $R(K)$. Note that this term is K independent, and hence can be treated as constant. A notable related work, which considers similar model selection problem is [75]. However, the paper obtains a sub-optimal rate of $\mathcal{O}(K^{2/3})$. Our main contribution is keeping the optimal rate of $\mathcal{O}(\sqrt{K})$, while pushing the cost of model selection in an additive constant independent of K .

Part I

Learning in Federated Framework

In this part of the thesis, we study some problems in Federated Learning (FL). In particular we address (a) data heterogeneity, (b) communication bottleneck and (c) Byzantine robustness.

- **Chapter 2:** We study the FL framework where users are partitioned into clusters, and our proposed algorithm learn the cluster identities and the optimal model for each cluster simultaneously.
- **Chapter 3:** We address the communication bottleneck and robustness issues via proposing first order gradient based algorithms.
- **Chapter 4:** We propose a second order distributed Newton type algorithm to address the communication bottleneck and furthermore, make it Byzantine resilient.
- **Chapter 5:** We target the saddle point avoidance problem in non-convex optimization via proposing an efficient and robust cubic-regularized Newton method.

Chapter 2

Clustered Federated Learning

We address the problem of Federated Learning (FL) where users are distributed and partitioned into clusters. This setup captures settings where different groups of users have their own objectives (learning tasks) but by aggregating their data with others in the same cluster (same learning task), they can leverage the strength in numbers in order to perform more efficient Federated Learning. We propose a new framework dubbed the Iterative Federated Clustering Algorithm (IFCA), which alternately estimates the cluster identities of the users and optimizes model parameters for the user clusters via gradient descent. We analyze the convergence rate of this algorithm first in a linear model with squared loss and then for generic strongly convex and smooth loss functions. We show that in both settings, with good initialization, IFCA converges at an exponential rate, and discuss the optimality of the statistical error rate. When the clustering structure is ambiguous, we propose to train the models by combining IFCA with the weight sharing technique in multi-task learning. In the experiments, we show that our algorithm can succeed even if we relax the requirements on initialization with random initialization and multiple restarts. We also present experimental results showing that our algorithm is efficient in non-convex problems such as neural networks. We demonstrate the benefits of IFCA over the baselines on several clustered FL benchmarks.

2.1 Introduction

In this chapter, we study one of the formulations of FL with non-i.i.d. data, i.e., the *clustered Federated Learning* [7, 9]. We assume that the users are partitioned into different clusters; for example, the clusters may represent groups of users interested in politics, sports, etc, and our goal is to train models for every cluster of users. We note that cluster structure is very common in applications such as recommender systems [10, 11]. The main challenge of our problem is that the *cluster identities of the users are unknown*, and we have to simultaneously solve two problems: identifying the cluster membership of each user and optimizing each of the cluster models in a distributed setting. In order to achieve this goal, we propose a framework and analyze a distributed method, named the *Iterative Federated Clustering Algorithm (IFCA)* for clustered FL. The basic idea of our algorithm is a strategy that alternates between estimating the cluster identities and

minimizing the loss functions, and thus can be seen as an Alternating Minimization algorithm in a distributed setting. One of the major advantages of our algorithm is that it does not require a centralized clustering algorithm, and thus significantly reduces the computational cost at the center machine. When the cluster structure is ambiguous, we propose to leverage the weight sharing technique in multi-task learning [12] and combine it with IFCA. More specifically, we learn the shared representation layers using data from all the users, and use IFCA to train separate final layers for each individual cluster.

We also present experimental evidence of its performance in practical settings: We show that our algorithm can succeed even if we relax the initialization requirements with random initialization and multiple restarts; and we also present results showing that our algorithm is efficient on neural networks. We demonstrate the effectiveness of IFCA on two clustered FL benchmarks created based on the MNIST and CIFAR-10 datasets, respectively, as well as the Federated EMNIST dataset [76] which is a more realistic benchmark for FL and has ambiguous cluster structure.

Here, we emphasize that clustered Federated Learning is not the only approach to modeling the non-i.i.d. nature of the problem, and different algorithms may be more suitable for different application scenarios; see Section 2.2 for more discussions. That said, our approach to modeling and the resulting IFCA framework is certainly an important and relatively unexplored direction in Federated Learning. We would also like to note that our theoretical analysis makes contributions to statistical estimation problems with latent variables in distributed settings. In fact, both mixture of regressions [77] and mixture of classifiers [78] can be considered as special cases of our problem in the centralized setting. We discuss more about these algorithms in Section 2.2.

Notation: We use $[r]$ to denote the set of integers $\{1, 2, \dots, r\}$. We use $\|\cdot\|$ to denote the ℓ_2 norm of vectors. We use $x \gtrsim y$ if there exists a sufficiently large constant $c > 0$ such that $x \geq cy$, and define $x \lesssim y$ similarly. We use $\text{poly}(m)$ to denote a polynomial in m with arbitrarily large constant degree.

2.2 Related work

During the development of our clustered FL framework, we became aware of a concurrent and independent work by Mansour et al. [9], in which the authors propose clustered FL as one of the formulations for personalization in Federated Learning. The algorithms proposed here and by Mansour et al. are similar. However, we make an important contribution by establishing the *convergence rate* of the *population loss function* under good initialization, which simultaneously guarantees both convergence of the training loss and generalization to test data; whereas in [9], the authors provided only *generalization* guarantees. We discuss other related work in the following.

Federated Learning and non-i.i.d. data: Learning with a distributed computing framework has been studied extensively in various settings [79–81]. As mentioned in Section 2.1, Federated Learning [1–3, 82] is one of the modern distributed learning frameworks that aims to better utilize the data and computing power on edge devices. A central problem in FL is that the data on the users’ personal devices are usually non-i.i.d. Several formulations and solutions have been proposed to tackle this problem. A line of research focuses on learning a single global model from

non-i.i.d. data [83–88]. Other lines of research focus more on learning personalized models [6, 7, 89]. In particular, the MOCHA algorithm [6] considers a multi-task learning setting and forms a deterministic optimization problem with the correlation matrix of the users being a regularization term. Our work differs from MOCHA since we consider a statistical setting with cluster structure. Another approach is to formulate Federated Learning with non-i.i.d. data as a meta learning problem [8, 89, 90]. In this setup, the objective is to first obtain a single global model, and then each device fine-tunes the model using its local data. The underlying assumption of this formulation is that the data distributions among different users are similar, and the global model can serve as a good initialization. The formulation of clustered FL has been considered in two recent works [7, 32]. Both of the two works use *centralized* clustering algorithm such as K -means, in which the center machine has to identify the cluster identities of all the users, leading to high computational cost at the center. As a result, these algorithms may not be suitable for large models such as deep neural networks or applications with a large number of users. In contrast, our algorithm uses a decentralized approach to identify the cluster identities and thus is more suitable for large-scale applications.

Latent variable problems: As mentioned in Section 2.1, our formulation can be considered as a statistical estimation problem with latent variables in a distributed setting, and the latent variables are the cluster identities. Latent variable problem is a classical topic in statistics and non-convex optimization; examples include Gaussian mixture models (GMM) [91, 92], mixture of linear regressions [77, 93, 94], and phase retrieval [95, 96]. Expectation Maximization (EM) and Alternating Minimization (AM) are two popular approaches to solving these problems. Despite the wide applications, their convergence analyses in the finite sample setting are known to be hard, due to the non-convexity nature of their optimization landscape. In recent years, some progress has been made towards understanding the convergence of EM and AM in the centralized setting [61, 97–100]. For example, if started from a suitable point, they have fast convergence rate, and occasionally they enjoy super-linear speed of convergence [91, 101]. Here, we provide new insights to these algorithms in the FL setting.

2.3 Problem formulation

We begin with a standard statistical learning setting of empirical risk minimization (ERM). Our goal is to learn parametric models by minimizing some loss functions defined by the data. We consider a distributed learning setting where we have one center machine and m worker machines (i.e., each worker machine corresponds to a user in the Federated Learning framework). The center machine and worker machines can communicate with each other using some predefined communication protocol. We assume that there are k different data distributions, $\mathcal{D}_1, \dots, \mathcal{D}_k$, and that the m machines are partitioned into k disjoint clusters, S_1^*, \dots, S_k^* . We assume no knowledge of the cluster identity of each machine, i.e., the partition S_1^*, \dots, S_k^* is not revealed to the learning algorithm. We assume that every worker machine $i \in S_j^*$ contains n i.i.d. data points $z^{i,1}, \dots, z^{i,n}$ drawn from \mathcal{D}_j , where each data point $z^{i,j}$ consists of a pair of feature and response denoted by $z^{i,\ell} = (x^{i,\ell}, y^{i,\ell})$.




Figure 2.1: An overview of IFCA (model averaging). (a) The server broadcast models. (b) Worker machines identify their cluster memberships and run local updates. (c) The worker machines send back the local models to server. (d) Average the models within the same estimated cluster S_j .

Let $f(w; z) : \mathcal{W} \rightarrow \mathbb{R}$ be the loss function associated with data point z , where $\mathcal{W} \subseteq \mathbb{R}^d$ is the parameter space. In this paper, we choose $\mathcal{W} = \mathbb{R}^d$. Our goal is to minimize the population loss function $F^j(w) := \mathbb{E}_{z \sim \mathcal{D}_j}[f(w; z)]$ for all $j \in [k]$. For the purpose of theoretical analysis in Section 2.5, we focus on the strongly convex losses, in which case we can prove guarantees for estimating the unique solution that minimizes each population loss function. In particular, we try to find solutions $\{\hat{w}_j\}_{j=1}^k$ that are close to $w_j^* = \operatorname{argmin}_{w \in \Theta} F^j(w)$, $j \in [k]$. In our problem, since we only have access to finite data, we take advantage of the empirical loss functions. In particular, let $Z \subseteq \{z^{i,1}, \dots, z^{i,n}\}$ be a subset of the data points on the i -th machine. We define the empirical loss associated with Z as $F_i(w; Z) = \frac{1}{|Z|} \sum_{z \in Z} f(w; z)$. When it is clear from the context, we may also use the shorthand notation $F_i(w)$ to denote an empirical loss associated with some (or all) data on the i -th worker.

2.4 Algorithm

In this section, we provide details of our algorithm. We name this scheme *Iterative Federated Clustering Algorithm* (IFCA). The main idea is to alternatively minimize the loss functions while estimating the cluster identities. We discuss two variations of IFCA, namely gradient averaging and model averaging. The algorithm is formally presented in Algorithm 1 and illustrated in Figure 2.1.

The algorithm starts with k initial model parameters $w_j^{(0)}$, $j \in [k]$. In the t -th iteration of IFCA, the center machine selects a random subset of worker machines, $M_t \subseteq [m]$, and broadcasts the current model parameters $\{w_j^{(t)}\}_{j=1}^k$ to the worker machines in M_t . Here, we call M_t the set of *participating devices*. Recall that each worker machine is equipped with local empirical loss function $F_i(\cdot)$. Using the received parameter estimates and F_i , the i -th worker machine ($i \in M_t$) estimates

Algorithm 1: Iterative Federated Clustering Algorithm (IFCA)

-
- 1: **Input:** number of clusters k , step size γ , $j \in [k]$, initialization $w_j^{(0)}$, $j \in [k]$
number of parallel iterations T , number of local gradient steps τ (for model averaging).
 - 2: **for** $t = 0, 1, \dots, T - 1$ **do**
 - 3: center machine: broadcast $w_j^{(t)}$, $j \in [k]$
 - 4: $M_t \leftarrow$ random subset of worker machines (participating devices)
 - 5: **for** worker machine $i \in M_t$ in parallel **do**
 - 6: cluster identity estimate $\hat{j} = \operatorname{argmin}_{j \in [k]} F_i(w_j^{(t)})$
 - 7: define one-hot encoding vector $s_i = \{s_{i,j}\}_{j=1}^k$ with $s_{i,j} = \mathbf{1}\{j = \hat{j}\}$
 - 8: **option I** (gradient averaging):
 - 9: compute (stochastic) gradient: $g_i = \widehat{\nabla} F_i(w_j^{(t)})$, send back s_i, g_i
to the center machine
 - 10: **option II** (model averaging):
 - 11: $\tilde{w}_i = \text{LocalUpdate}(w_j^{(t)}, \gamma, \tau)$, send back s_i, \tilde{w}_i
to the center machine
 - 12: **end for**
 - 13: center machine:
 - 14: **option I** (gradient averaging): $w_j^{(t+1)} = w_j^{(t)} - \frac{\gamma}{m} \sum_{i \in M_t} s_{i,j} g_i, \forall j \in [k]$
 - 15: **option II** (model averaging): $w_j^{(t+1)} = \sum_{i \in M_t} s_{i,j} \tilde{w}_i / \sum_{i \in M_t} s_{i,j}, \forall j \in [k]$
 - 16: **end for**
 - 17: **return** $w_j^{(T)}$, $j \in [k]$
LocalUpdate $(\tilde{w}^{(0)}, \gamma, \tau)$ at the i -th worker machine
 - 18: **for** $q = 0, \dots, \tau - 1$ **do**
 - 19: (stochastic) gradient descent $\tilde{w}^{(q+1)} = \tilde{w}^{(q)} - \gamma \widehat{\nabla} F_i(\tilde{w}^{(q)})$
 - 20: **end for**
 - 21: **return** $\tilde{w}^{(\tau)}$
-

its cluster identity via finding the model parameter with lowest loss, i.e., $\hat{j} = \operatorname{argmin}_{j \in [k]} F_i(w_j^{(t)})$ (ties can be broken arbitrarily). If we choose the option of gradient averaging, the worker machine then computes a (stochastic) gradient of the local empirical loss F_i at $w_j^{(t)}$, and sends its cluster identity estimate and gradient back to the center machine. After receiving the gradients and cluster identity estimates from all the participating worker machines, the center machine then collects all the gradient updates from worker machines whose cluster identity estimates are the same and conducts gradient descent update on the model parameter of the corresponding cluster. If we choose the option of model averaging (similar to the Federated Averaging algorithm [3]), each participating device needs to run τ steps of local (stochastic) gradient descent updates, get the updated model, and send the new model and its cluster identity estimate to the center machine. The center machine then averages the new models from the worker machines whose cluster identity estimates are the same.

2.4.1 Practical implementation of IFCA

We clarify a few issues regarding the practical implementation of IFCA. In some real-world problems, the cluster structure may be ambiguous, which means that although the distributions of data from different clusters are different, there exists some common properties of the data from all the users that the model should leverage. For these problems, we propose to use the weight sharing technique in multi-task learning [12] and combine it with IFCA. More specifically, when we train neural network models, we can share the weights for the first a few layers among all the clusters so that we can learn a good representation using all the available data, and then run IFCA algorithm only on the last (or last few) layers to address the different distributions among different clusters. Using the notation in Algorithm 1, we run IFCA on a subset of the coordinates of $w_j^{(t)}$, and run vanilla gradient averaging or Federated Averaging on the remaining coordinates. Another benefit of this implementation is that we can reduce the communication cost: Instead of sending k models to all the worker machines, the center machine only needs to send k different versions of a subset of all the weights, and one single copy of the shared layers.

Another technique to reduce communication cost is that when the center machine observes that the cluster identities of all the worker machines are stable, i.e., the estimates of their cluster identities do not change for several parallel iterations, then the center machine can stop sending k models to each worker machine, and instead, it can simply send the model corresponding to each worker machine's cluster identity estimate.

2.5 Theoretical guarantees

In this section, we present convergence guarantees of IFCA. In order to streamline our theoretical analysis, we make several simplifications: we consider the IFCA with gradient averaging, and assume that all the worker machines participate in every rounds of IFCA, i.e., $M_t = [m]$ for all t . In addition, we also use the *re-sampling* technique for the purpose of theoretical analysis. In particular, suppose that we run a total of T parallel iterations. We partition the n data points on each machine into $2T$ disjoint subsets, each with $n' = \frac{n}{2T}$ data points. For the i -th machine, we denote the subsets as $\widehat{Z}_i^{(0)}, \dots, \widehat{Z}_i^{(T-1)}$ and $Z_i^{(0)}, \dots, Z_i^{(T-1)}$. In the t -th iteration, we use $\widehat{Z}_i^{(t)}$ to estimate the cluster identity, and use $Z_i^{(t)}$ to conduct gradient descent. As we can see, we use fresh data samples for each iteration of the algorithm. Furthermore, in each iteration, we use different set of data points for obtaining the cluster estimate and computing the gradient. This is done in order to remove the inter-dependence between the cluster estimation and the gradient computation, and ensure that in each iteration, we use fresh i.i.d. data that are independent of the current model parameter. We would like to emphasize that re-sampling is a standard tool used in statistics [61, 98, 101–103], and that it is for theoretical tractability only and is not required in practice as we show in Section 2.6.

Under these conditions, the update rule for the parameter vector of the j -th cluster can be written as

$$S_j^{(t)} = \{i \in [m] : j = \operatorname{argmin}_{j' \in [k]} F_i(w_{j'}^{(t)}; \widehat{Z}_i^{(t)})\}, \quad w_j^{(t+1)} = w_j^{(t)} - \frac{\gamma}{m} \sum_{i \in S_j^{(t)}} \nabla F_i(w_j^{(t)}; Z_i^{(t)}),$$

where $S_j^{(t)}$ denotes the set of worker machines whose cluster identity estimate is j in the t -th iteration. In the following, we discuss the convergence guarantee of IFCA under two models: in Section 2.5.1, we analyze the algorithm under a linear model with Gaussian features and squared loss, and in Section 2.5.2, we analyze the algorithm under a more general setting of strongly convex loss functions.

2.5.1 Linear models with squared loss

In this section, we analyze our algorithm in a concrete linear model. This model can be seen as a warm-up example for more general problems with strongly convex loss functions that we discuss in Section 2.5.2, as well as a distributed formulation of the widely studied mixture of linear regression problem [98, 103]. We assume that the data on the worker machines in the j -th cluster are generated in the following way: for $i \in S_j^*$, the feature-response pair of the i -th worker machine machine satisfies

$$y^{i,\ell} = \langle x^{i,\ell}, w_j^* \rangle + \epsilon^{i,\ell},$$

where $x^{i,\ell} \sim \mathcal{N}(0, I_d)$ and the additive noise $\epsilon^{i,\ell} \sim \mathcal{N}(0, \sigma^2)$ is independent of $x^{i,\ell}$. Furthermore, we use the squared loss function $f(w; x, y) = (y - \langle x, w \rangle)^2$. As we can see, this model is the mixture of linear regression model in the distributed setting. We observe that under the above setting, the parameters $\{w_j^*\}_{j=1}^k$ are the minimizers of the population loss function $F^j(\cdot)$.

We proceed to analyze our algorithm. We define $p_j := |S_j^*|/m$ as the fraction of worker machines belonging to the j -th cluster, and let $p := \min\{p_1, p_2, \dots, p_k\}$. We also define the minimum separation Δ as $\Delta := \min_{j \neq j'} \|w_j^* - w_{j'}^*\|$, and $\rho := \frac{\Delta^2}{\sigma^2}$ as the signal-to-noise ratio. Before we establish our convergence result, we state a few assumptions. Here, recall that n' denotes the number of data that each worker uses in each step.

Assumption 3. *The initialization of parameters $w_j^{(0)}$ satisfy $\|w_j^{(0)} - w_j^*\| \leq \frac{1}{4}\Delta, \forall j \in [k]$.*

Assumption 4. *Without loss of generality, we assume that $\max_{j \in [k]} \|w_j^*\| \lesssim 1$, and that $\sigma \lesssim 1$. We also assume that $n' \gtrsim \left(\frac{\rho+1}{\rho}\right)^2 \log m$, $d \gtrsim \log m$, $p \gtrsim \frac{\log m}{m}$, $p m n' \gtrsim d$, and $\Delta \gtrsim \frac{\sigma}{p} \sqrt{\frac{d}{m n'}} + \exp(-c(\frac{\rho}{\rho+1})^2 n')$ for some universal constant c .*

In Assumption 3, we assume that the initialization is close enough to w_j^* . We note that this is a standard assumption in the convergence analysis of mixture models [99, 104], due to the non-convex optimization landscape of mixture model problems. In Assumption 4, we put mild assumptions on n' , m , p , and d . The condition that $p m n' \gtrsim d$ simply assumes that the total number of data that we use in each iteration for each cluster is at least as large as the dimension of the parameter space. The condition that $\Delta \gtrsim \frac{\sigma}{p} \sqrt{\frac{d}{m n'}} + \exp(-c(\frac{\rho}{\rho+1})^2 n')$ ensures that the iterates stay close to w_j^* .

We first provide a single step analysis of our algorithm. We assume that at a certain iteration, we obtain parameter vectors w_j that are close to the ground truth parameters w_j^* , and show that w_j converges to w_j^* at an exponential rate with an error floor.

Theorem 1. *Consider the linear model and assume that Assumptions 3 and 4 hold. Suppose that in a certain iteration of the IFCA algorithm we obtain parameter vectors w_j with $\|w_j - w_j^*\| \leq \frac{1}{4}\Delta$. Let w_j^+ be iterate after this iteration. Then there exist universal constants $c_1, c_2, c_3, c_4 > 0$ such that when we choose step size $\gamma = c_1/p$, with probability at least $1 - 1/\text{poly}(m)$, we have for all $j \in [k]$,*

$$\|w_j^+ - w_j^*\| \leq \frac{1}{2}\|w_j - w_j^*\| + c_2 \frac{\sigma}{p} \sqrt{\frac{d}{mn'}} + c_3 \exp\left(-c_4 \left(\frac{\rho}{\rho+1}\right)^2 n'\right).$$

We prove Theorem 1 in Appendix 2.8. Here, we briefly summarize the proof idea. Using the initialization condition, we show that the set $\{S_j\}_{j=1}^k$ has a significant overlap with $\{S_j^*\}_{j=1}^k$. In the overlapped set, we then argue that the gradient step provides a contraction and error floor due to the basic properties of linear regression. We then bound the gradient norm of the miss-classified machines and add them to the error floor. We complete the proof by combining the contributions of properly classified and miss-classified worker machines. We can then iteratively apply Theorem 1 and obtain accuracy of the final solution \hat{w}_j in the following corollary.

Corollary 1. *Consider the linear model and assume that Assumptions 3 and 4 hold. By choosing step size $\gamma = c_1/p$, with probability at least $1 - \frac{\log(\Delta/4\varepsilon)}{\text{poly}(m)}$, after $T = \log \frac{\Delta}{4\varepsilon}$ parallel iterations, we have for all $j \in [k]$, $\|\hat{w}_j - w_j^*\| \leq \varepsilon$, where $\varepsilon = c_5 \frac{\sigma}{p} \sqrt{\frac{d}{mn'}} + c_6 \exp(-c_4 (\frac{\rho}{\rho+1})^2 n')$.*

Let us examine the final accuracy. Since the number of data points on each worker machine $n = 2n'T = 2n' \log(\Delta/4\varepsilon)$, we know that for the smallest cluster, there are a total of $2pmn' \log(\Delta/4\varepsilon)$ data points. According to the minimax estimation rate of linear regression [13], we know that even if we know the ground truth cluster identities, we cannot obtain an error rate better than $\mathcal{O}(\sigma \sqrt{\frac{d}{pmn' \log(\Delta/4\varepsilon)}})$. Comparing this rate with our statistical accuracy ε , we can see that the first term $\frac{\sigma}{p} \sqrt{\frac{d}{mn'}}$ in ε is equivalent to the minimax rate up to a logarithmic factor and a dependency on p , and the second term in ε decays exponentially fast in n' , and therefore, our final statistical error rate is near optimal.

2.5.2 Strongly convex loss functions

In this section, we study a more general scenario where the population loss functions of the k clusters are strongly convex and smooth. In contrast to the previous section, our analysis do not rely on any specific statistical model, and thus can be applied to more general machine learning problems. We start with reviewing the standard definitions of strongly convex and smooth functions $F : \mathbb{R}^d \mapsto \mathbb{R}$.

Definition 4. F is λ -strongly convex if $\forall w, w', F(w') \geq F(w) + \langle \nabla F(w), w' - w \rangle + \frac{\lambda}{2} \|w' - w\|^2$.

Definition 5. F is L -smooth if $\forall w, w', \|\nabla F(w) - \nabla F(w')\| \leq L \|w - w'\|$.

In this section, we assume that the population loss functions $F^j(w)$ are strongly convex and smooth.

Assumption 5. *The population loss function $F^j(w)$ is λ -strongly convex and L -smooth, $\forall j \in [k]$.*

We note that we do not make any convexity or smoothness assumptions on the individual loss function $f(w; z)$. Instead, we make the following distributional assumptions on $f(w; z)$ and $\nabla f(w; z)$.

Assumption 6. *For every w and every $j \in [k]$, the variance of $f(w; z)$ is upper bounded by η^2 , when z is sampled according to \mathcal{D}_j , i.e., $\mathbb{E}_{z \sim \mathcal{D}_j}[(f(w; z) - F^j(w))^2] \leq \eta^2$*

Assumption 7. *For every w and every $j \in [k]$, the variance of $\nabla f(w; z)$ is upper bounded by v^2 , when z is sampled according to \mathcal{D}_j , i.e., $\mathbb{E}_{z \sim \mathcal{D}_j}[\|\nabla f(w; z) - \nabla F^j(w)\|_2^2] \leq v^2$*

Bounded variance of gradient is very common in analyzing SGD [105]. In this paper we use loss function value to determine cluster identity, so we also need to have a probabilistic assumption on $f(w; z)$. We note that bounded variance is a relatively weak assumption on the tail behavior of probability distributions. In addition to the assumptions above, we still use some definitions from Section 2.5.1, i.e., $\Delta := \min_{j \neq j'} \|w_j^* - w_{j'}^*\|$, and $p = \min_{j \in [k]} p_j$ with $p_j = |S_j^*|/m$. We make the following assumptions on the initialization, n' , p , and Δ .

Assumption 8. *Without loss of generality, we assume that $\max_{j \in [k]} \|w_j^*\| \lesssim 1$. We also assume that $\|w_j^{(0)} - w_j^*\| \leq \frac{1}{4} \sqrt{\frac{\lambda}{L}} \Delta$, $\forall j \in [k]$, $n' \gtrsim \frac{k\eta^2}{\lambda^2 \Delta^4}$, $p \gtrsim \frac{\log(mn')}{m}$, and that*

$$\Delta \geq \tilde{\mathcal{O}}(\max\{(n')^{-1/5}, m^{-1/6}(n')^{-1/3}\}).$$

Here, for simplicity, the $\tilde{\mathcal{O}}$ notation omits any logarithmic factors and quantities that do not depend on m and n' . As we can see, again we need to assume good initialization, due to the nature of the mixture model, and the assumptions that we impose on n' , p , and Δ are relatively mild; in particular, the assumption on Δ ensures that the iterates stay close to an ℓ_2 ball around w_j^* .

Theorem 2. *Suppose Assumptions 5-8 hold. Choose step size $\gamma = 1/L$. Then, with probability at least $1 - \delta$, after $T = \frac{8L}{p\lambda} \log(\frac{\Delta}{2\varepsilon})$ parallel iterations, we have for all $j \in [k]$, $\|\hat{w}_j - w_j^*\| \leq \varepsilon$, where*

$$\varepsilon \lesssim \frac{vkL \log(mn')}{p^{5/2} \lambda^2 \delta \sqrt{mn'}} + \frac{\eta^2 L^2 k \log(mn')}{p^2 \lambda^4 \delta \Delta^4 n'} + \tilde{\mathcal{O}}\left(\frac{1}{n' \sqrt{m}}\right).$$

We prove Theorem 2 in the Appendix 2.9. Similar to Section 2.5.1, to prove this result, we first prove a per-iteration contraction

$$\|w_j^+ - w_j^*\| \leq \left(1 - \frac{p\lambda}{8L}\right) \|w_j - w_j^*\| + \tilde{\mathcal{O}}\left(\frac{1}{\sqrt{mn'}} + \frac{1}{n'} + \frac{1}{n' \sqrt{m}}\right), \forall j \in [k],$$

and then derive the convergence rate. To better interpret the result, we focus on the dependency on m and n and treat other quantities as constants. Then, since $n = 2n'T$, we know that n and n' are of the same scale up to a logarithmic factor. Therefore, the final statistical error rate that we obtain is $\varepsilon = \tilde{\mathcal{O}}\left(\frac{1}{\sqrt{mn}} + \frac{1}{n}\right)$. As discussed in Section 2.5.1, $\frac{1}{\sqrt{mn}}$ is the optimal rate even if we

know the cluster identities; thus our statistical rate is near optimal in the regime where $n \gtrsim m$. In comparison with the statistical rate in linear models $\tilde{\mathcal{O}}(\frac{1}{\sqrt{mn}} + \exp(-n))$, we note that the major difference is in the second term. The additional terms of the linear model and the strongly convex case are $\exp(-n)$ and $\frac{1}{n}$, respectively. We note that this is due to different statistical assumptions: in for the linear model, we assume Gaussian noise whereas here we only assume bounded variance.

2.6 Experiments

In this section, we present our experimental results, which not only validate the theoretical claims in Section 2.5, but also demonstrate that our algorithm can be efficiently applied beyond the regime we discussed in the theory. We emphasize that we *do not* re-sample fresh data points at each iteration. Furthermore, the requirement on the initialization can be relaxed. More specifically, for linear models, we observe that random initialization with a few restarts is sufficient to ensure convergence of Algorithm 1. In our experiments, we also show that our algorithm works efficiently for problems with non-convex loss functions such as neural networks.

2.6.1 Synthetic data

We begin with evaluation of Algorithm 1 with gradient averaging (option I) on linear models with squared loss, as described in Section 2.5.1. For all $j \in [k]$, we first generate $w_j^* \sim \text{Ber}(0.5)$ coordinate-wise, and then rescale their ℓ_2 norm to R . This ensures that the separation between the w_j^* 's is proportional to R in expectation, and thus, in this experiment, we use R to represent the *separation* between the ground truth parameter vectors. Moreover, we simulate the scenario where all the worker machines participate in all iterations, and all the clusters contain same number of worker machines. For each trial of the experiment, we first generate the parameter vectors w_j^* 's, fix them, and then randomly initialize $w_j^{(0)}$ according to an independent coordinate-wise Bernoulli distribution. We then run Algorithm 1 for 300 iterations, with a constant step size. For $k = 2$ and $k = 4$, we choose the step size in $\{0.01, 0.1, 1\}$, $\{0.5, 1.0, 2.0\}$, respectively. In order to determine whether we successfully learn the model or not, we sweep over the aforementioned step sizes and define the following distance metric: $\text{dist} = \frac{1}{k} \sum_{j=1}^k \|\hat{w}_j - w_j^*\|$, where $\{\hat{w}_j\}_{j=1}^k$ are the parameter estimates obtained from Algorithm 1. A trial is dubbed *successful* if for a fixed set of w_j^* , among 10 random initialization of $w_j^{(0)}$, at least in one scenario, we obtain $\text{dist} \leq 0.6\sigma$.

In Fig. 2.2 (a-b), we plot the empirical success probability over 40 trials, with respect to the separation parameter R . We set the problem parameters as (a) $(m, n, d) = (100, 100, 1000)$ with $k = 2$, and (b) $(m, n, d) = (400, 100, 1000)$ with $k = 4$. As we can see, when R becomes larger, i.e., the separation between parameters increases, and the problem becomes easier to solve, yielding in a higher success probability. This validates our theoretical result that higher signal-to-noise ratio produces smaller error floor. In Fig. 2.2 (c-d), we characterize the dependence on m and n , with fixing R and d with $(R, d) = (0.1, 1000)$ for (c) and $(R, d) = (0.5, 1000)$ for (d). We observe that when we increase m and/or n , the success probability improves. This validates our theoretical finding that more data and/or more worker machines help improve the performance of the algorithm.




Figure 2.2: Success probability with respect to: (a), (b) the separation scale R and the scale of additive noise σ ; (c), (d) the number of worker machines m and the sample size on each machine n . In (a) and (b), we see that the success probability gets better with increasing R , i.e., more separation between ground truth parameter vectors, and in (c) and (d), we note that the success probability improves with an increase of mn , i.e., more data on each machine and/or more machines.

2.6.2 Rotated MNIST and CIFAR

We also create clustered FL datasets based on the MNIST [106] and CIFAR-10 [107] datasets. In order to simulate an environment where the data on different worker machines are generated from different distributions, we augment the datasets using rotation, and create the Rotated MNIST [108] and Rotated CIFAR datasets. For **Rotated MNIST**, recall that the MNIST dataset has 60000 training images and 10000 test images with 10 classes. We first augment the dataset by applying 0, 90, 180, 270 degrees of rotation to the images, resulting in $k = 4$ clusters. For given m and n satisfying $mn = 60000k$, we randomly partition the images into m worker machines so that each machine holds n images *with the same rotation*. We also split the test data into $m_{test} = 10000k/n$ worker machines in the same way. The **Rotated CIFAR** dataset is also created in a similar way as

Table 2.1: Test accuracies(%) \pm std on Rotated MNIST ($k = 4$) and Rotated CIFAR ($k = 2$)

m, n	Rotated MNIST			Rotated CIFAR
	4800, 50	2400, 100	1200, 200	200, 500
IFCA (ours)	94.20 \pm 0.03	95.05 \pm 0.02	95.25 \pm 0.40	81.51 \pm 1.37
global model	86.74 \pm 0.04	88.65 \pm 0.08	89.73 \pm 0.13	77.87 \pm 0.39
local model	63.32 \pm 0.02	73.66 \pm 0.04	80.05 \pm 0.02	33.97 \pm 1.19

Rotated MNIST, with the main difference being that we create $k = 2$ clusters with 0 and 180 degrees of rotation. We note that creating different tasks by manipulating standard datasets such as MNIST and CIFAR-10 has been widely adopted in the continual learning research community [108–110]. For clustered FL, creating datasets using rotation helps us simulate a federated learning setup with clear cluster structure.

For our MNIST experiments, we use the fully connected neural network with ReLU activations, with a single hidden layer of size 200; and for our CIFAR experiments, we use a convolution neural network model which consists of 2 convolutional layers followed by 2 fully connected layers, and the images are preprocessed by standard data augmentation such as flipping and random cropping.

We compare our IFCA algorithm with two baseline algorithms, i.e., the *global model*, and *local model* schemes. For **IFCA**, we use model averaging (option II in Algorithm 1). For MNIST experiments, we use full worker machines participation ($M_t = [m]$ for all t). For LocalUpdate step in Algorithm 1, we choose $\tau = 10$ and step size $\gamma = 0.1$. For CIFAR experiments, we choose $|M_t| = 0.1m$, and apply step size decay 0.99, and we also set $\tau = 5$ and batch size 50 for LocalUpdate process, following prior works [1]. In the **global model** scheme, the algorithm tries to learn single global model that can make predictions from all the distributions. The algorithm does not consider cluster identities, so model averaging step in Algorithm 1 becomes $w^{(t+1)} = \sum_{i \in M_t} \tilde{w}_i / |M_t|$, i.e. averaged over parameters from all the participating machines. In the **local model** scheme, the model in each node performs gradient descent only on local data available, and model averaging is not performed.

For IFCA and the global model scheme, we perform inference in the following way. For every test worker machine, we run inference on all learned models (k models for IFCA and one model for global model scheme), and calculate the accuracy from the model that produces the smallest loss value. For testing the local model baselines, the models are tested by measuring the accuracy on the test data with the same distribution (i.e. those have the same rotation). We report the accuracy averaged over all the models in worker machines. For all algorithms, we run experiment with 5 different random seeds and report the average and standard deviation.

Our experimental results are shown in Table 2.1. We can observe that our algorithm performs better than the two baselines. As we run the IFCA algorithm, we observe that we can gradually find the underlying cluster identities of the worker machines, and after the correct cluster is found, each model is trained and tested using data with the same distribution, resulting in better accuracy. The global model baseline performs worse than ours since it tries to fit all the data from different distributions, and cannot provide personalized predictions. The local model baseline algorithm overfits to the local data easily, leading to worse performance than ours.

2.6.3 Federated EMNIST

We provide additional experimental results on the Federated EMNIST (FEMNIST) [76], which is a realistic FL dataset where the data points on every worker machine are the handwritten digits or letters from a specific writer. Although the data distribution among all the users are similar, there might be ambiguous cluster structure since the writing styles of different people may be clustered. We use the weight sharing technique mentioned in Section 2.4.1. We use a neural network with two convolutional layers, with a max pooling layer after each convolutional layer, followed by two fully connected layers. We share the weights of all the layers, except the last layer which is trained by IFCA. We treat the number of clusters k as a hyper parameter and run the experiments with different values of k . We compare IFCA with the global model and local model approaches, as well as the one-shot centralized clustering algorithm in [32]. The test accuracies are shown in Table 2.2, with mean and standard deviation computed over 5 independent runs. As we can see, IFCA shows clear advantage over the global model and local model approaches. The results of IFCA and the one-shot algorithm are similar. However, as we emphasized in Section 2.2, IFCA does not run a centralized clustering procedure, and thus reduces the computational cost at the center machine. As a final note, we observe that IFCA is robust to the choice of the number of clusters k . The results of the algorithm with $k = 2$ and $k = 3$ are similar, and we notice that when $k > 3$, IFCA automatically identifies 3 clusters, and the remaining clusters are empty. This indicates the applicability of IFCA in real-world problems where the cluster structure is ambiguous and the number of clusters is unknown.

Table 2.2: Test accuracies (%) \pm std on FEMNIST

IFCA ($k = 2$)	IFCA ($k = 3$)	one-shot ($k = 2$)	one-shot ($k = 3$)	global	local
87.99 \pm 0.35	87.89 \pm 0.52	87.41 \pm 0.39	87.38 \pm 0.37	84.45 \pm 0.51	75.85 \pm 0.72

2.7 Conclusion and Open Problems

In this paper, we address the clustered FL problem. We propose an iterative algorithm and obtain convergence guarantees for strongly convex and smooth functions. In experiments, we achieve this via random initialization with multiple restarts, and we show that our algorithm works efficiently beyond the convex regime. An immediate future work is to extend the analysis to weakly convex and non-convex functions. Also, the convergence guarantees are local, i.e., a good initialization is required. Obtaining a provable initialization for the clustered FL is an interesting future direction.

Our formulation is one of the problem setups for personalized Federated Learning. We expect that, overall our framework will better protect the users' privacy in a Federated Learning system while still provide personalized predictions. The reason is that our algorithm does not require the

users to send any of their own personal data to the central server, and the users can still learn a personalized model using their on-device computing power. One potential risk is that our algorithm still requires the users to send the estimates of their cluster identities to the central server. Thus there might still be privacy concerns in this step. We suggest that before applying our algorithm, or generally any FL algorithms, in a real-world system, we should first request the users' consent.

Appendix

In our proofs, we use c, c_1, c_2, \dots to denote positive universal constants, the value of which may differ across instances. For a matrix A , we write $\|A\|_{op}$ and $\|A\|_F$ as the operator norm and Frobenius norm, respectively. For a set S , we use \bar{S} to denote the complement of the set.

2.8 Proof of Theorem 1

Since we only analyze a single iteration, for simplicity we drop the superscript that indicates the iteration counter. Suppose that at a particular iteration, we have model parameters $w_j, j \in [k]$, for the k clusters. We denote the *estimation* of the set of worker machines that belongs to the j -th cluster by S_j , and recall that the true clusters are denoted by $S_j^*, j \in [k]$.

Probability of erroneous cluster identity estimation We begin with the analysis of the probability of incorrect cluster identity estimation. Suppose that a worker machine i belongs to S_j^* . We define the event $\mathcal{E}_i^{j,j'}$ as the event when the i -th machine is classified to the j' -th cluster, i.e., $i \in S_{j'}$. Thus the event that worker i is correctly classified is $\mathcal{E}_i^{j,j}$, and we use the shorthand notation $\mathcal{E}_i := \mathcal{E}_i^{j,j}$. We now provide the following lemma that bounds the probability of $\mathcal{E}_i^{j,j'}$ for $j' \neq j$.

Lemma 1. *Suppose that worker machine $i \in S_j^*$. Let $\rho := \frac{\Delta^2}{\sigma^2}$. Then there exist universal constants c_1 and c_2 such that for any $j' \neq j$,*

$$\mathbb{P}(\mathcal{E}_i^{j,j'}) \leq c_1 \exp\left(-c_2 n' \left(\frac{\rho}{\rho+1}\right)^2\right),$$

and by union bound

$$\mathbb{P}(\bar{\mathcal{E}}_i) \leq c_1 k \exp\left(-c_2 n' \left(\frac{\rho}{\rho+1}\right)^2\right).$$

We prove Lemma 1 in Appendix 2.8.1.

Now we proceed to analyze the gradient descent step. Without loss of generality, we only analyze the first cluster. The update rule of w_1 in this iteration can be written as

$$w_1^\dagger = w_1 - \frac{\gamma}{m} \sum_{i \in S_1} \nabla F_i(w_1; Z_i),$$

where Z_i is the set of the n' data points that we use to compute gradient in this iteration on a particular worker machine.

We use the shorthand notation $F_i(w) := F_i(w; Z_i)$, and note that $F_i(w)$ can be written in the matrix form as

$$F_i(w) = \frac{1}{n'} \|Y_i - X_i w\|^2,$$

where we have the feature matrix $X_i \in \mathbb{R}^{n' \times d}$ and response vector $Y_i = X_i w_1^* + \epsilon_i$. According to our model, all the entries of X_i are i.i.d. sampled according to $\mathcal{N}(0, 1)$, and $\epsilon_i \sim \mathcal{N}(0, \sigma^2 I)$.

We first notice that

$$\|w_1^+ - w_1^*\| = \underbrace{\|w_1 - w_1^* - \frac{\gamma}{m} \sum_{i \in S_1 \cap S_1^*} \nabla F_i(\theta_1)\|}_{T_1} + \underbrace{\frac{\gamma}{m} \sum_{i \in S_1 \cap \bar{S}_1^*} \nabla F_i(\theta_1)}_{T_2} \leq \|T_1\| + \|T_2\|.$$

We control the two terms separately. Let us first focus on $\|T_1\|$.

Bound $\|T_1\|$ To simplify notation, we concatenate all the feature matrices and response vectors of all the worker machines in $S_1 \cap S_1^*$ and get the new feature matrix $X \in \mathbb{R}^{N \times d}$, $Y \in \mathbb{R}^N$ with $Y = X w_1^* + \epsilon$, where $N := n' |S_1 \cap S_1^*|$. It is then easy to verify that

$$\begin{aligned} T_1 &= (I - \frac{2\gamma}{mn'} X^\top X)(w_1 - w_1^*) + \frac{2\gamma}{mn'} X^\top \epsilon \\ &= (I - \frac{2\gamma}{mn'} \mathbb{E}[X^\top X])(w_1 - w_1^*) + \frac{2\gamma}{mn'} (\mathbb{E}[X^\top X] - X^\top X)(w_1 - w_1^*) + \frac{2\gamma}{mn'} X^\top \epsilon \\ &= (1 - \frac{2\gamma N}{mn'}) (w_1 - w_1^*) + \frac{2\gamma}{mn'} (\mathbb{E}[X^\top X] - X^\top X)(w_1 - w_1^*) + \frac{2\gamma}{mn'} X^\top \epsilon. \end{aligned}$$

Therefore

$$\|T_1\| \leq (1 - \frac{2\gamma N}{mn'}) \|w_1 - w_1^*\| + \frac{2\gamma}{mn'} \|X^\top X - \mathbb{E}[X^\top X]\|_{op} \|w_1 - w_1^*\| + \frac{2\gamma}{mn'} \|X^\top \epsilon\|. \quad (2.1)$$

Thus in order to bound $\|T_1\|$, we need to analyze two terms, $\|X^\top X - \mathbb{E}[X^\top X]\|_{op}$ and $\|X^\top \epsilon\|$. To bound $\|X^\top X - \mathbb{E}[X^\top X]\|_{op}$, we first provide an analysis of N showing that it is large enough. Using Lemma 1 in conjunction with Assumption 4, we see that the probability of correctly classifying any worker machine i , given by $\mathbb{P}(\mathcal{E}_i)$, satisfies $\mathbb{P}(\mathcal{E}_i) \geq \frac{1}{2}$. Hence, we obtain

$$\mathbb{E}[|S_1 \cap S_1^*|] \geq \mathbb{E}[\frac{1}{2}|S_1^*|] = \frac{1}{2} p_1 m,$$

where we use the fact that $|S_1^*| = p_1 m$. Since $|S_1 \cap S_1^*|$ is a sum of Bernoulli random variables with success probability at least $\frac{1}{2}$, we obtain

$$\mathbb{P}\left(|S_1 \cap S_1^*| \leq \frac{1}{4} p_1 m\right) \leq \mathbb{P}\left(\left||S_1 \cap S_1^*| - \mathbb{E}[|S_1 \cap S_1^*|]\right| \geq \frac{1}{4} p_1 m\right) \leq 2 \exp(-cpm),$$

where $p = \min\{p_1, p_2, \dots, p_k\}$, and the second step follows from Hoeffding's inequality. Hence, we obtain $|S_1 \cap S_1^*| \geq \frac{1}{4}p_1 m$ with high probability, which yields

$$\mathbb{P}(N \geq \frac{1}{4}p_1 m n') \geq 1 - 2 \exp(-c p m). \quad (2.2)$$

By combining this fact with our assumption that $p m n' \gtrsim d$, we know that $N \gtrsim d$. Then, according to the concentration of the covariance of Gaussian random vectors [13], we know that with probability at least $1 - 2 \exp(-\frac{1}{2}d)$,

$$\|X^\top X - \mathbb{E}[X^\top X]\|_{op} \leq 6\sqrt{dN} \lesssim N. \quad (2.3)$$

We now proceed to bound $\|X^\top \epsilon\|$. In particular, we use the following lemma.

Lemma 2. *Consider a random matrix $X \in \mathbb{R}^{N \times d}$ with i.i.d. entries sampled according to $\mathcal{N}(0, 1)$, and $\epsilon \in \mathbb{R}^N$ be a random vector sampled according to $\mathcal{N}(0, \sigma^2 I)$, independently of X . Then we have with probability at least $1 - 2 \exp(-c_1 \max\{d, N\})$,*

$$\|X\|_{op} \leq c \max\{\sqrt{d}, \sqrt{N}\},$$

and with probability at least $1 - c_2 \exp(-c_3 \min\{d, N\})$,

$$\|X^\top \epsilon\| \leq c_4 \sigma \sqrt{dN}.$$

We prove Lemma 2 in Appendix 2.8.2. Now we can combine (2.1), (2.3), (2.2), and Lemma 2 and obtain with probability at least $1 - c_1 \exp(-c_2 p m) - c_3 \exp(-c_4 d)$,

$$\|T_1\| \leq (1 - c_5 \gamma p) \|w_1 - w_1^*\| + c_6 \gamma \sigma \sqrt{\frac{d}{m n'}}. \quad (2.4)$$

Since we assume that $p \gtrsim \frac{\log m}{m}$ and $d \gtrsim \log m$, the success probability can be simplified as $1 - 1/\text{poly}(m)$.

Bound $\|T_2\|$ We first condition on S_1 . We have the following:

$$\nabla F_i(w_1) = \frac{2}{n'} X_i^\top (Y_i - X_i w_1).$$

For $i \in S_1 \cap S_j^*$, with $j \neq 1$, we have $Y_i = X_i w_j^* + \epsilon_i$, and so we obtain

$$n' \nabla F_i(w_1) = 2X_i^\top X_i (w_j^* - w_1) + 2X_i^\top \epsilon_i,$$

which yields

$$n' \|\nabla F_i(w_1)\| \lesssim \|X_i\|_{op}^2 + \|X_i^\top \epsilon_i\|, \quad (2.5)$$

where we use the fact that $\|w_j^* - w_1\| \leq \|w_j^*\| + \|w_1^*\| + \|w_1^* - w_1\| \lesssim 1$. Then, we combine (2.5) and Lemma 2 and get with probability at least $1 - c_1 \exp(-c_2 \min\{d, n'\})$,

$$\|\nabla F_i(w_1)\| \leq \frac{1}{n'}(c_3 \max\{d, n'\} + c_4 \sigma \sqrt{dn'}) \leq c_5 \max\{1, \frac{d}{n'}\}, \quad (2.6)$$

where we use our assumption that $\sigma \lesssim 1$. By union bound, we know that with probability at least $1 - c_1 m \exp(-c_2 \min\{d, n'\})$, (2.6) holds for all $j \in \overline{S_1^*}$. In addition, since we assume that $n' \gtrsim \log m$, $d \gtrsim \log m$, this probability can be lower bounded by $1 - 1/\text{poly}(m)$. This implies that conditioned on S_1 , with probability at least $1 - 1/\text{poly}(m)$,

$$\|T_2\| \leq c_5 \frac{\gamma}{m} |S_1 \cap \overline{S_1^*}| \max\{1, \frac{d}{n'}\}. \quad (2.7)$$

Since we choose $\gamma = \frac{c}{p}$, we have $\frac{\gamma}{m} \max\{1, \frac{d}{n'}\} \lesssim 1$, where we use our assumption that $pmn' \gtrsim d$. This shows that with probability at least $1 - 1/\text{poly}(m)$,

$$\|T_2\| \leq c_5 |S_1 \cap \overline{S_1^*}|. \quad (2.8)$$

We then analyze $|S_1 \cap \overline{S_1^*}|$. By Lemma 1, we have

$$\mathbb{E}[|S_1 \cap \overline{S_1^*}|] \leq c_6 m \exp(-c_7 (\frac{\rho}{\rho+1})^2 n'). \quad (2.9)$$

According to Assumption 4, we know that $n' \geq c(\frac{\rho+1}{\rho})^2 \log m$, for some constant c that is large enough. Therefore, $m \leq \exp(\frac{1}{c}(\frac{\rho}{\rho+1})^2 n')$, and thus, as long as c is large enough such that $\frac{1}{c} < c_7$ where c_7 is defined in (2.9), we have

$$\mathbb{E}[|S_1 \cap \overline{S_1^*}|] \leq c_6 \exp(-c_8 (\frac{\rho}{\rho+1})^2 n'). \quad (2.10)$$

and then by Markov's inequality, we have

$$\mathbb{P}\left(|S_1 \cap \overline{S_1^*}| \leq c_6 \exp(-\frac{c_8}{2} (\frac{\rho}{\rho+1})^2 n')\right) \geq 1 - \exp(-\frac{c_8}{2} (\frac{\rho}{\rho+1})^2 n') \geq 1 - \text{poly}(m). \quad (2.11)$$

Combining (2.8) with (2.11), we know that with probability at least $1 - 1/\text{poly}(m)$,

$$\|T_2\| \leq c_1 \exp(-c_2 (\frac{\rho}{\rho+1})^2 n').$$

Using this fact and (2.4), we obtain that with probability at least $1 - 1/\text{poly}(m)$,

$$\|w_1^+ - w_1^*\| \leq (1 - c_1 \gamma p) \|w_1 - w_1^*\| + c_2 \gamma \sigma \sqrt{\frac{d}{mn'}} + c_3 \exp(-c_4 (\frac{\rho}{\rho+1})^2 n').$$

Then we can complete the proof for the first cluster by choosing $\gamma = \frac{1}{2c_1 p}$. To complete the proof for all the k clusters, we can use union bound, and the success probability is $1 - k/\text{poly}(m)$. However, since $k \leq m$ by definition, we still have success probability $1 - 1/\text{poly}(m)$.

2.8.1 Proof of Lemma 1

Without loss of generality, we analyze $\mathcal{E}_i^{1,j}$ for some $j \neq 1$. By definition, we have

$$\mathcal{E}_i^{1,j} = \{F_i(w_j; \widehat{Z}_i) \leq F_i(w_1; \widehat{Z}_i)\},$$

where \widehat{Z}_i is the set of n' data points that we use to estimate the cluster identity in this iteration. We write the data points in \widehat{Z}_i in matrix form with feature matrix $X_i \in \mathbb{R}^{n' \times d}$ and response vector $Y_i = X_i w_1^* + \epsilon_i$. According to our model, all the entries of X_i are i.i.d. sampled according to $\mathcal{N}(0, 1)$, and $\epsilon_i \sim \mathcal{N}(0, \sigma^2 I)$. Then, we have

$$\mathbb{P}\{\mathcal{E}_i^{1,j}\} = \mathbb{P}\{\|X_i(w_1^* - w_1) + \epsilon_i\|^2 \geq \|X_i(w_1^* - w_j) + \epsilon_i\|^2\}.$$

Consider the random vector $X_i(w_1^* - w_j) + \epsilon_i$, and in particular consider the ℓ -th coordinate of it. Since X_i and ϵ_i are independent and we resample (X_i, Y_i) at each iteration, the ℓ -th coordinate of $X_i(w_1^* - w_j) + \epsilon_i$ is a Gaussian random variable with mean 0 and variance $\|w_j - w_1^*\|^2 + \sigma^2$. Since X_i and ϵ_i contain independent rows, the distribution of $\|X_i(w_1^* - w_j) + \epsilon_i\|^2$ is given by $(\|w_j - w_1^*\|^2 + \sigma^2)u_j$, where u_j is a standard Chi-squared random variable n' degrees of freedom. We now calculate the an upper bound on the following probability:

$$\begin{aligned} & \mathbb{P}\{\|X_i(w_1^* - w_1) + \epsilon_i\|^2 \geq \|X_i(w_1^* - w_j) + \epsilon_i\|^2\} \\ & \stackrel{(i)}{\leq} \mathbb{P}\{\|X_i(w_1^* - w_j) + \epsilon_i\|^2 \leq t\} + \mathbb{P}\{\|X_i(w_1^* - w_1) + \epsilon_i\|^2 > t\} \\ & \leq \mathbb{P}\{(\|w_j - w_1^*\|^2 + \sigma^2)u_j \leq t\} + \mathbb{P}\{(\|w_1 - w_1^*\|^2 + \sigma^2)u_j > t\}, \end{aligned} \quad (2.12)$$

where (i) holds for all $t \geq 0$. For the first term, we use the concentration property of Chi-squared random variables. Using the fact that $\|w_j - w_1^*\| \geq \|w_j^* - w_1^*\| - \|w_j - w_j^*\| \geq \frac{3}{4}\Delta$, we have

$$\mathbb{P}\{(\|w_j - w_1^*\|^2 + \sigma^2)u_j \leq t\} \leq \mathbb{P}\left\{\left(\frac{9}{16}\Delta^2 + \sigma^2\right)u_j \leq t\right\}. \quad (2.13)$$

Similarly, using the initialization condition, $\|w_1 - w_1^*\| \leq \frac{1}{4}\Delta$, the second term of equation (2.12) can be simplified as

$$\mathbb{P}\{(\|w_1 - w_1^*\|^2 + \sigma^2)u_j > t\} \leq \mathbb{P}\left\{\left(\frac{1}{16}\Delta^2 + \sigma^2\right)u_j > t\right\}. \quad (2.14)$$

Based on the above observation, we now choose $t = n'(\frac{5}{16}\Delta^2 + \sigma^2)$. Recall that $\rho := \frac{\Delta^2}{\sigma^2}$. Then the inequality (2.13) can be rewritten as

$$\mathbb{P}\{(\|w_j - w_1^*\|^2 + \sigma^2)u_j \leq t\} \leq \mathbb{P}\left\{\frac{u_j}{n'} - 1 \leq -\frac{4\rho}{9\rho + 16}\right\}.$$

According to the concentration results for standard Chi-squared distribution [13], we know that there exists universal constants c_1 and c_2 such that

$$\mathbb{P}\{(\|w_j - w_1^*\|^2 + \sigma^2)u_j \leq t\} \leq c_1 \exp\left(-c_2 n' \left(\frac{\rho}{\rho + 1}\right)^2\right). \quad (2.15)$$

Similarly, the inequality (2.14) can be rewritten as

$$\mathbb{P} \left\{ (\|w_1 - w_1^*\|^2 + \sigma^2)u_j > t \right\} \leq \mathbb{P} \left\{ \frac{u_j}{n'} - 1 > \frac{4\rho}{\rho + 16} \right\}$$

and again according to the concentration of Chi-squared distribution, there exists universal constants c_3 and c_4 such that

$$\mathbb{P} \left\{ (\|w_1 - w_1^*\|^2 + \sigma^2)u_1 > t \right\} \leq c_3 \exp \left(-c_4 n' \left(\frac{\rho}{\rho + 1} \right)^2 \right). \quad (2.16)$$

The proof can be completed by combining (2.12), (2.15) and (2.16).

2.8.2 Proof of Lemma 2

According to Theorem 5.39 of [111], we have with probability at least $1 - 2 \exp(-c_1 \max\{d, N\})$,

$$\|X\|_{op} \leq c \max\{\sqrt{d}, \sqrt{N}\},$$

where c and c_1 are universal constants. As for $\|X^\top \epsilon\|$, we first condition on X . According to the Hanson-Wright inequality [112], we obtain for every $t \geq 0$

$$\mathbb{P} \left(\left| \|X^\top \epsilon\| - \sigma \|X^\top\|_F \right| > t \right) \leq 2 \exp \left(-c \frac{t^2}{\sigma^2 \|X^\top\|_{op}^2} \right). \quad (2.17)$$

Using Chi-squared concentration [13], we obtain with probability at least $1 - 2 \exp(-cdN)$,

$$\|X\|_F \leq c\sqrt{dN}.$$

Furthermore, using the fact that $\|X^\top\|_{op} = \|X\|_{op}$ and substituting $t = \sigma\sqrt{dN}$ in (2.17), we obtain with probability at least $1 - c_2 \exp(-c_3 \min\{d, N\})$,

$$\|X^\top \epsilon\| \leq c_4 \sigma \sqrt{dN}.$$

2.9 Proof of Theorem 2

The proof of this theorem is similar to that of the linear model. We begin with a single-step analysis.

2.9.1 Analysis for a single step

Suppose that at a certain step, we have model parameters $w_j, j \in [k]$ for the k clusters. Assume that $\|w_j - w_j^*\| \leq \frac{1}{4} \sqrt{\frac{\lambda}{L}} \Delta$, for all $j \in [k]$.

Probability of erroneous cluster identity estimation: We first calculate the probability of erroneous estimation of worker machines' cluster identity. We define the events $\mathcal{E}_i^{j,j'}$ in the same way as in Appendix 2.8, and have the following lemma.

Lemma 3. *Suppose that worker machine $i \in S_j^*$. Then there exists a universal constants c_1 such that for any $j' \neq j$,*

$$\mathbb{P}(\mathcal{E}_i^{j,j'}) \leq c_1 \frac{\eta^2}{\lambda^2 \Delta^4 n'},$$

and by union bound

$$\mathbb{P}(\overline{\mathcal{E}}_i) \leq c_1 \frac{k\eta^2}{\lambda^2 \Delta^4 n'}.$$

We prove Lemma 3 in Appendix 2.9.3. Now we proceed to analyze the gradient descent iteration. Without loss of generality, we focus on w_1 . We have

$$\|w_1^+ - w_1^*\| = \|w_1 - w_1^* - \frac{\gamma}{m} \sum_{i \in S_1} \nabla F_i(\theta_1)\|,$$

where $F_i(w) := F_i(w; Z_i)$ with Z_i being the set of data points on the i -th worker machine that we use to compute the gradient, and S_1 is the set of indices returned by Algorithm 1 corresponding to the first cluster. Since

$$S_1 = (S_1 \cap S_1^*) \cup (S_1 \cap \overline{S_1^*})$$

and the sets are disjoint, we have

$$\|w_1^+ - w_1^*\| = \underbrace{\|w_1 - w_1^* - \frac{\gamma}{m} \sum_{i \in S_1 \cap S_1^*} \nabla F_i(\theta_1)\|}_{T_1} - \underbrace{\frac{\gamma}{m} \sum_{i \in S_1 \cap \overline{S_1^*}} \nabla F_i(\theta_1)}_{T_2}.$$

Using triangle inequality, we obtain

$$\|w_1^+ - w_1^*\| \leq \|T_1\| + \|T_2\|,$$

and we control both the terms separately. Let us first focus on $\|T_1\|$.

Bound $\|T_1\|$ We first split T_1 in the following way:

$$T_1 = \underbrace{w_1 - w_1^* - \hat{\gamma} \nabla F^1(w_1)}_{T_{11}} + \underbrace{\hat{\gamma} \left(\nabla F^1(w_1) - \frac{1}{|S_1 \cap S_1^*|} \sum_{i \in S_1 \cap S_1^*} \nabla F_i(w_1) \right)}_{T_{12}}, \quad (2.18)$$

where $\hat{\gamma} := \gamma \frac{|S_1 \cap S_1^*|}{m}$. Let us condition on S_1 . According to standard analysis technique for gradient descent on strongly convex functions, we know that when $\hat{\gamma} \leq \frac{1}{L}$,

$$\|T_{11}\| = \|w_1 - w_1^* - \hat{\gamma} \nabla F^1(w_1)\| \leq \left(1 - \frac{\hat{\gamma} \lambda L}{\lambda + L}\right) \|w_1 - w_1^*\|. \quad (2.19)$$

Further, we have $\mathbb{E}[\|T_{12}\|^2] = \frac{v^2}{n'|S_1 \cap S_1^*|}$, which implies $\mathbb{E}[\|T_{12}\|] \leq \frac{v}{\sqrt{n'|S_1 \cap S_1^*|}}$, and thus by Markov's inequality, for any $\delta_0 > 0$, with probability at least $1 - \delta_0$,

$$\|T_{12}\| \leq \frac{v}{\delta_0 \sqrt{n'|S_1 \cap S_1^*|}}. \quad (2.20)$$

We then analyze $|S_1 \cap S_1^*|$. Similar to the proof of Theorem 1, we can show that $|S_1 \cap S_1^*|$ is large enough. From Lemma 3 and using our assumption, we see that the probability of correctly classifying any worker machine i , given by $\mathbb{P}(\mathcal{E}_i)$, satisfies $\mathbb{P}(\mathcal{E}_i) \geq \frac{1}{2}$. Recall $p = \min\{p_1, p_2, \dots, p_k\}$, and we obtain $|S_1 \cap S_1^*| \geq \frac{1}{4}p_1m$ with probability at least $1 - 2\exp(-cpm)$. Let us condition on $|S_1 \cap S_1^*| \geq \frac{1}{4}p_1m$ and choose $\gamma = 1/L$. Then $\hat{\gamma} \leq 1/L$ is satisfied, and on the other hand $\hat{\gamma} \geq \frac{p}{4L}$. Plug this fact in (2.19), we obtain

$$\|T_{11}\| \leq \left(1 - \frac{p\lambda}{8L}\right) \|w_1 - w_1^*\|. \quad (2.21)$$

We then combine (2.20) and (2.21) and have with probability at least $1 - \delta_0 - 2\exp(-cpm)$,

$$\|T_1\| \leq \left(1 - \frac{p\lambda}{8L}\right) \|w_1 - w_1^*\| + \frac{2v}{\delta_0 L \sqrt{pmn'}}. \quad (2.22)$$

Bound $\|T_2\|$ Let us define $T_{2j} := \sum_{S_1 \cap S_j^*} \nabla F_i(w_1)$, $j \geq 2$. We have $T_2 = \frac{\gamma}{m} \sum_{j=2}^k T_{2j}$. We condition on S_1 and first analyze T_{2j} . We have

$$T_{2j} = |S_1 \cap S_j^*| \nabla F^j(w_1) + \sum_{i \in S_1 \cap S_j^*} (\nabla F_i(w_1) - \nabla F^j(w_1)). \quad (2.23)$$

Due to the smoothness of $F^j(w)$, we know that

$$\|\nabla F^j(w_1)\| \leq L \|w_1 - w_j^*\| \leq 3L, \quad (2.24)$$

where we use the fact that $\|w_1 - w_j^*\| \leq \|w_j^*\| + \|w_1^*\| + \|w_1 - w_1^*\| \leq 1 + 1 + \frac{1}{4}\sqrt{\frac{\lambda}{L}}\Delta \leq 3$. In addition, we have

$$\mathbb{E} \left[\left\| \sum_{i \in S_1 \cap S_j^*} \nabla F_i(w_1) - \nabla F^j(w_1) \right\|^2 \right] = |S_1 \cap S_j^*| \frac{v^2}{n'},$$

which implies

$$\mathbb{E} \left[\left\| \sum_{i \in S_1 \cap S_j^*} \nabla F_i(w_1) - \nabla F^j(w_1) \right\| \right] \leq \sqrt{|S_1 \cap S_j^*|} \frac{v}{\sqrt{n'}},$$

and then according to Markov's inequality, for any $\delta_1 \in (0, 1)$, with probability at least $1 - \delta_1$,

$$\left\| \sum_{i \in S_1 \cap S_j^*} \nabla F_i(w_1) - \nabla F^j(w_1) \right\| \leq \sqrt{|S_1 \cap S_j^*|} \frac{v}{\delta_1 \sqrt{n'}}. \quad (2.25)$$

Then, by combining (2.24) and (2.25), we know that with probability at least $1 - \delta_1$,

$$\|T_{2j}\| \leq 3L|S_1 \cap S_j^*| + \sqrt{|S_1 \cap S_j^*|} \frac{v}{\delta_1 \sqrt{n'}}. \quad (2.26)$$

By union bound, we know that with probability at least $1 - k\delta_1$, (2.26) applies to all $j \neq 1$. Then, we have with probability at least $1 - k\delta_1$,

$$\|T_2\| \leq \frac{3\gamma L}{m} |S_1 \cap \overline{S_1^*}| + \frac{\gamma v \sqrt{k}}{\delta_1 m \sqrt{n'}} \sqrt{|S_1 \cap \overline{S_1^*}|}. \quad (2.27)$$

According to Lemma 3, we know that

$$\mathbb{E}[|S_1 \cap \overline{S_1^*}|] \leq c_1 \frac{\eta^2 m}{\lambda^2 \Delta^4 n'}.$$

Then by Markov's inequality, we know that with probability at least $1 - \delta_2$,

$$|S_1 \cap \overline{S_1^*}| \leq c_1 \frac{\eta^2 m}{\delta_2 \lambda^2 \Delta^4 n'}. \quad (2.28)$$

Now we combine (2.27) with (2.28) and obtain with probability at least $1 - k\delta_1 - \delta_2$,

$$\|T_2\| \leq c_1 \frac{\eta^2}{\delta_2 \lambda^2 \Delta^4 n'} + c_2 \frac{v \eta \sqrt{k}}{\delta_1 \sqrt{\delta_2} \lambda L \Delta^2 \sqrt{m n'}}. \quad (2.29)$$

Combining (2.22) and (2.29), we know that with probability at least $1 - \delta_0 - k\delta_1 - \delta_2 - 2 \exp(-cpm)$,

$$\|w_1^+ - w_1^*\| \leq \left(1 - \frac{p\lambda}{8L}\right) \|w_1 - w_1^*\| + \frac{2v}{\delta_0 L \sqrt{p m n'}} + c_1 \frac{\eta^2}{\delta_2 \lambda^2 \Delta^4 n'} + c_2 \frac{v \eta \sqrt{k}}{\delta_1 \sqrt{\delta_2} \lambda L \Delta^2 \sqrt{m n'}}. \quad (2.30)$$

In the following, we let $\delta_3 := \delta_0 + k\delta_1 + \delta_2 + 2 \exp(-cpm)$, and

$$\varepsilon_0 = \frac{2v}{\delta_0 L \sqrt{p m n'}} + c_1 \frac{\eta^2}{\delta_2 \lambda^2 \Delta^4 n'} + c_2 \frac{v \eta \sqrt{k}}{\delta_1 \sqrt{\delta_2} \lambda L \Delta^2 \sqrt{m n'}}.$$

Let us simplify this expression. We first choose $\delta \in (0, 1)$ as the failure probability of the entire algorithm. Then, we choose

$$\delta_0 = \frac{p\lambda\delta}{CkL \log(mn')}, \quad \delta_1 = \frac{p\lambda\delta}{Ck^2L \log(mn')}, \quad \delta_2 = \frac{p\lambda\delta}{CkL \log(mn')},$$

for some constant $C > 0$ that is large enough. In addition, since we assume that $p \gtrsim \frac{\log(mn')}{m}$, we have $\exp(-cpm) \leq 1/\text{poly}(mn') \lesssim \frac{p\lambda\delta}{kL \log(mn')}$. Consider all these facts, we obtain

$$\delta_3 = \frac{4p\lambda\delta}{CkL \log(mn')}, \quad (2.31)$$

$$\varepsilon_0 \lesssim \frac{vk \log(mn')}{p^{3/2}\lambda\delta\sqrt{mn'}} + \frac{\eta^2 Lk \log(mn')}{p\lambda^3\delta\Delta^4 n'} + \frac{v\eta k^3 \sqrt{L} \log^{3/2}(mn')}{p^{3/2}\lambda^{5/2}\delta^{3/2}\Delta^2 \sqrt{mn'}}. \quad (2.32)$$

In addition, by union bound, we know that with probability at least $1 - k\delta_3$, for all $j \in [k]$,

$$\|w_j^+ - w_j^*\| \leq \left(1 - \frac{p\lambda}{8L}\right) \|w_j - w_j^*\| + \varepsilon_0. \quad (2.33)$$

2.9.2 Convergence of the algorithm

We now analyze the convergence of the entire algorithm. First, we can verify that as long as

$$\varepsilon_0 \leq \frac{p}{32} \left(\frac{\lambda}{L}\right)^{3/2} \Delta, \quad (2.34)$$

we can guarantee that $\|w_j^+ - w_j^*\| \leq \frac{1}{4} \sqrt{\frac{\lambda}{L}} \Delta$. We can also verify that as long as there is

$$\Delta \geq \tilde{\mathcal{O}}(\max\{(n')^{-1/5}, m^{-1/6}(n')^{-1/3}\}), \quad (2.35)$$

using the definition of ε_0 in (2.32), we know that (2.34) holds. Here, in the $\tilde{\mathcal{O}}$ notation, we omit the logarithmic factors and quantities that does not depend on m and n' . In this case, we can iteratively apply (2.33) for T iterations and obtain that with probability at least $1 - kT\delta_3$,

$$\|w_j^{(T)} - w_j^*\| \leq \left(1 - \frac{p\lambda}{8L}\right)^T \|w_j^{(0)} - w_j^*\| + \frac{8L}{p\lambda} \varepsilon_0.$$

Then, we know that when we choose

$$T = \frac{8L}{p\lambda} \log \left(\frac{p\lambda\Delta}{32\varepsilon_0 L} \right), \quad (2.36)$$

we have

$$\left(1 - \frac{p\lambda}{8L}\right)^T \|w_j^{(0)} - w_j^*\| \leq \exp\left(-\frac{p\lambda}{8L} T\right) \frac{1}{4} \sqrt{\frac{\lambda}{L}} \Delta \leq \frac{8}{p} \sqrt{\frac{L}{\lambda}} \varepsilon_0,$$

which implies $\|w_j^{(T)} - w_j^*\| \leq \frac{16L}{p\lambda} \varepsilon_0$. Finally, we check the failure probability. The failure probability is

$$kT\delta_3 \leq \frac{8kL}{p\lambda} \log \left(\frac{p\lambda\Delta}{32\varepsilon_0 L} \right) \frac{4p\lambda\delta}{CkL \log(mn')} = \frac{32\delta \log\left(\frac{p\lambda\Delta}{32\varepsilon_0 L}\right)}{C \log(mn')} \leq \delta \frac{\log\left(\frac{1}{\varepsilon_0}\right)}{\log((mn')^{C/32})}.$$

On the other hand, according to (2.32), we know that

$$\frac{1}{\varepsilon_0} \leq \tilde{\mathcal{O}}(\max\{\sqrt{mn'}, n'\}),$$

then, as long as C is large enough, we can guarantee that $(mn')^{C/32} > \frac{1}{\varepsilon_0}$, which implies that the failure probability is upper bounded by δ . Our final error floor can be obtained by redefining

$$\varepsilon := \frac{16L}{p\lambda} \varepsilon_0.$$

2.9.3 Proof of Lemma 3

Without loss of generality, we bound the probability of $\mathcal{E}_i^{1,j}$ for some $j \neq 1$. We know that

$$\mathcal{E}_i^{1,j} = \left\{ F_i(w_1; \widehat{Z}_i) \geq F_i(w_j; \widehat{Z}_i) \right\},$$

where \widehat{Z}_i is the set of n' data points that we use to estimate the cluster identity in this iteration. In the following, we use the shorthand notation $F_i(w) := F_i(w; \widehat{Z}_i)$. We have

$$\mathbb{P}(\mathcal{E}_i^{1,j}) \leq \mathbb{P}(F_i(w_1) > t) + \mathbb{P}(F_i(w_j) \leq t)$$

for all $t \geq 0$. We choose $t = \frac{F^1(w_1) + F^1(w_j)}{2}$. With this choice, we obtain

$$\mathbb{P}(F_i(w_1) > t) = \mathbb{P}\left(F_i(w_1) > \frac{F^1(w_1) + F^1(w_j)}{2}\right) \quad (2.37)$$

$$= \mathbb{P}\left(F_i(w_1) - F^1(w_1) > \frac{F^1(w_j) - F^1(w_1)}{2}\right). \quad (2.38)$$

Similarly, for the second term, we have

$$\mathbb{P}(F_i(w_j) \leq t) = \mathbb{P}\left(F_i(w_j) - F^1(w_j) \leq -\frac{F^1(w_j) - F^1(w_1)}{2}\right). \quad (2.39)$$

Based on our assumption, we know that $\|w_j - w_1\| \geq \Delta - \frac{1}{4}\sqrt{\frac{\lambda}{L}}\Delta \geq \frac{3}{4}\Delta$. According to the strong convexity of $F^1(\cdot)$,

$$F^1(w_j) \geq F^1(w_1^*) + \frac{\lambda}{2}\|w_j - w_1^*\|^2 \geq F^1(w_1^*) + \frac{9\lambda}{32}\Delta^2,$$

and according to the smoothness of $F^1(\cdot)$,

$$F^1(w_1) \leq F^1(w_1^*) + \frac{L}{2}\|w_1 - w_1^*\|^2 \leq F^1(w_1^*) + \frac{L}{2} \frac{\lambda}{16L} \Delta^2 = F^1(w_1^*) + \frac{\lambda}{32} \Delta^2.$$

Therefore, $F^1(w_j) - F^1(w_1) \geq \frac{\lambda}{4}\Delta^2$. Then, according to Chebyshev's inequality, we obtain that $\mathbb{P}(F_i(w_1) > t) \leq \frac{64\eta^2}{\lambda^2\Delta^4n'}$ and that $\mathbb{P}(F_i(w_j) \leq t) \leq \frac{64\eta^2}{\lambda^2\Delta^4n'}$, which complete the proof.

Chapter 3

Communication-Efficient and Byzantine-Robust Distributed First Order Learning

In this chapter, we develop a communication-efficient distributed learning algorithm that is robust against Byzantine worker machines. We propose and analyze a distributed gradient-descent algorithm that performs a simple thresholding based on gradient norms to mitigate Byzantine failures. Furthermore, for communication efficiency, we consider a generic class of δ -approximate compressors from Karimireddi et al. [20] that encompasses sign-based compressors and top- k sparsification. Our algorithm uses compressed gradients and gradient norms for aggregation and Byzantine removal respectively. We establish the statistical error rate for non-convex smooth loss functions. We show that, in certain range of the compression factor δ , the (order-wise) rate of convergence is not affected by the compression operation. Moreover, we analyze the compressed gradient descent algorithm with error feedback (proposed in [20]) in a distributed setting and in the presence of Byzantine worker machines. We show that exploiting error feedback improves the statistical error rate. Finally, we experimentally validate our results and show good performance in convergence for convex (least-square regression) and non-convex (neural network training) problems.

3.1 Introduction

We propose provable and efficient algorithms that address the issues of communication efficiency and Byzantine robustness simultaneously. Note that, both these challenges, have recently attracted significant research attention, albeit mostly separately. In particular, several recent works have proposed various quantization or sparsification techniques to reduce the communication overhead ([14–16, 21, 113–116]). The goal of these quantization schemes is to compute an unbiased estimate of the gradient with bounded second moment in order to achieve good convergence guarantees. The problem of developing Byzantine-robust distributed algorithms has been considered in [17–19, 23, 30–33].

A notable exception to considering communication overhead separately from Byzantine robustness is the recent work of [22]. In this work, a sign-based compression algorithm signSGD of [117] is shown to be Byzantine fault-tolerant. The main idea of signSGD is to communicate the coordinate-wise signs of the gradient vector to reduce communication and employ a majority vote during the aggregation to mitigate the effect of Byzantine units. However, signSGD suffers from two major drawbacks. First, sign-based algorithms do not converge in general ([20]). In particular, [20, Section 3] presents several convex counter examples where signSGD fails to converge even though [22, Theorem 2] shows convergence guarantee for non-convex objective under certain assumptions. Second, signSGD can handle only a limited class of adversaries, namely *blind multiplicative adversaries* ([22]). Such an adversary manipulates the gradients of the worker machines by multiplying it (element-wise) with a vector that can scale and randomize the sign of each coordinate of the gradient. However, the vector must be chosen before observing the gradient (hence ‘blind’).

In this work, we develop communication-efficient and robust learning algorithms that overcome both these drawbacks¹. Specifically, we consider the following distributed learning setup. There are m worker machines, each storing n data points. The data points are generated from some unknown distribution \mathcal{D} . The objective is to learn a parametric model that minimizes a population loss function $F : \mathcal{W} \rightarrow \mathbb{R}$, where F is defined as an expectation over \mathcal{D} , and $\mathcal{W} \subseteq \mathbb{R}^d$ denotes the parameter space. We choose the loss function F to be non-convex. With the rapid rise of the neural networks, the study of *local minima* in non-convex optimization framework has become imperative [35, 36]. For gradient compression at workers, we consider the notion of a δ -approximate compressor from [20] that encompasses sign-based compressors like QSGD ([21]), ℓ_1 -QSGD ([20]) and top- k sparsification ([15]). We assume that $0 \leq \alpha < 1/2$ fraction of the worker machines are Byzantine. In contrast to blind multiplicative adversaries, we consider unrestricted adversaries.

Our key idea is to use a simple threshold (on local gradient norms) based Byzantine resilience scheme in contrast with computationally complex robust aggregation methods such as coordinate wise median or trimmed mean of [23]. Our main result is to show that, for a wide range of compression factor δ , the statistical error rate of our proposed threshold-based scheme is (order-wise) identical to the case of no compression considered in [23]. In fact, our algorithm achieves order-wise optimal error-rate in parameters (α, n, m) . Furthermore, to alleviate convergence issues associated with sign-based compressors, we employ the technique of error-feedback from [20]. In this setup, the worker machines store the difference between the actual and compressed gradient and add it back to the next step so that the *correct* direction of the gradient is not forgotten. We show that using error feedback with our threshold based Byzantine resilience scheme not only achieves better statistical error rate but also improves the rate of convergence. We outline our specific contributions in the following.

Our Contributions: We propose a communication-efficient and robust distributed gradient descent (GD) algorithm. The algorithm takes as input the gradients compressed using a δ -approximate compressor along with the norms² (of either compressed or uncompressed gradients), and performs a simple thresholding operation on based on gradient norms to discard $\beta > \alpha$ fraction of workers with

¹We compare our algorithm with signSGD in Section 3.8.

²We can handle any convex norm.

the largest norm values. We establish the statistical error rate of the algorithm for arbitrary smooth population loss functions as a function of the number of worker machines m , the number of data points on each machine n , dimension d , and the compression factor δ . In particular, we show that our algorithm achieves the following statistical error rate³ for the regime $\delta > 4\beta + 4\alpha - 8\alpha^2 + 4\alpha^3$:

$$\tilde{\mathcal{O}} \left(d^2 \left[\frac{\alpha^2}{n} + \frac{1-\delta}{n} + \frac{1}{mn} \right] \right). \quad (3.1)$$

We first note that when $\delta = 1$ (uncompressed), the error rate is $\tilde{\mathcal{O}}(d^2[\frac{\alpha^2}{n} + \frac{1}{mn}])$, which matches [23]. Notice that we use a simple threshold (on local gradient norms) based Byzantine resilience scheme in contrast with the coordinate wise median or trimmed mean of [23]. We note that for a fixed d and the compression factor δ satisfying $\delta \geq 1 - \alpha^2$, the statistical error rate become $\tilde{\mathcal{O}}(\frac{\alpha^2}{n} + \frac{1}{mn})$, which is order-wise identical to the case of no compression [23]. In other words, in this parameter regime, the compression term does not contribute (order-wise) to the statistical error. Moreover, it is shown in [23] that, for strongly-convex loss functions and a fixed d , no algorithm can achieve an error lower than $\tilde{\Omega}(\frac{\alpha^2}{n} + \frac{1}{mn})$, implying that our algorithm is order-wise optimal in terms of the statistical error rate in the parameters (α, n, m) .

Furthermore, we strengthen our distributed learning algorithm by using error feedback to correct the direction of the local gradient (Algorithm 3). We show (both theoretically and via experiments) that using error-feedback with a δ -approximate compressor indeed speeds up the convergence rate and attains better (statistical) error rate. Under the assumption that the gradient norm of the local loss function is upper-bounded by σ , we obtain the following (statistical) error rate:

$$\tilde{\mathcal{O}} \left(d^2 \left[\frac{\alpha^2}{n} + \frac{(1-\delta)\sigma^2}{d^2 \delta} + \frac{1}{mn} \right] \right)$$

provided a similar (δ, α) trade-off⁴. We note that in the no-compression setting ($\delta = 1$), we recover the $\tilde{\mathcal{O}}(\frac{\alpha^2}{n} + \frac{1}{mn})$ rate. Furthermore, in Section 3.7.2, we argue that, when $\delta = \Theta(1)$, the error rate of Algorithm 3 is strictly better than that of Algorithm 2. In experiments (Section 3.8), we also see a reflection of this fact.

We experimentally evaluate our algorithm for convex and non-convex losses. For the convex case, we choose the linear regression problem, and for the non-convex case, we train a ReLU activated feed-forward fully connected neural net. We compare our algorithm with the non-Byzantine case and signSGD with majority vote, and observe that our algorithm converges faster using the standard MNIST dataset.

A major technical challenge is to handle compression and the Byzantine behavior of the worker machines simultaneously. We build up on the techniques of [23] to control the Byzantine machines. In particular, using certain distributional assumption on the partial derivative of the loss function and exploiting uniform bounds via careful covering arguments, we show that the local gradient on a non-Byzantine worker machine is close to the gradient of the population loss function.

³Throughout the paper $\mathcal{O}(\cdot)$ hides multiplicative constants, while $\tilde{\mathcal{O}}(\cdot)$ further hides logarithmic factors.

⁴See Theorem 5 for details.

Note that in some settings, our results may not have an optimal dependence on dimension d . This is due to the norm-based Byzantine removal schemes. Obtaining optimal dependence on d is an interesting future direction.

Organization: We describe the problem formulation in Sec. 3.2, and give a brief overview of δ -compressors in Sec. 3.3. Then, we present our proposed algorithm in Sec. 2.4. We analyze the algorithm, first, for a *restricted* (as described next) adversarial model in Section 3.5, and in the subsequent section, remove this restriction. In Section 3.5, we restrict our attention to an adversarial model in which Byzantine workers can provide arbitrary values as an input to the compression algorithm, but they correctly implement the same compression scheme as mandated. In Section 3.6, we remove this restriction on the Byzantine machines. As a consequence, we observe (in Theorem 4) that the modified algorithm works under a stricter assumption, and performs slightly worse than the one in restricted adversary setting. In Section 3.7, we strengthen our algorithm by including error-feedback at worker machines, and provide statistical guarantees for non-convex smooth loss functions. We show that error-feedback indeed improves the performance of our optimization algorithm in the presence of arbitrary adversaries.

3.1.1 Related Work

Gradient Compression: The foundation of gradient quantization was laid in [118] and [119]. In the work of [21, 114, 115] each co-ordinate of the gradient vector is represented with a small number of bits. Using this, an unbiased estimate of the gradient is computed. In these works, the communication cost is $\Omega(\sqrt{d})$ bits. In [113], a quantization scheme was proposed for distributed mean estimation. The tradeoff between communication and accuracy is studied in [120]. Variance reduction in communication efficient stochastic distributed learning has been studied in [121]. Sparsification techniques are also used instead of quantization to reduce communication cost. Gradient sparsification has been studied in [14–16] with provable guarantees. The main idea is to communicate top components of the d -dimensional local gradient to get good estimate of the true global gradient.

Byzantine Robust Optimization: In the distributed learning context, a generic framework of one shot median based robust learning has been proposed in [19]. In [30] the issue of byzantine failure is tackled by grouping the servers in batches and computing the median of batched servers. Later in [23, 31], co-ordinate wise median, trimmed mean and iterative filtering based algorithm have been proposed and optimal statistical error rate is obtained. Also, [122, 123] considers adversaries may steer convergence to bad local minimizers. In this work, we do not assume such adversaries.

Gradient compression and Byzantine robust optimization have simultaneously been addressed in a recent paper [22]. Here, the authors use signSGD as compressor and majority voting as robust aggregator. As explained in [20], signSGD can run into convergence issues. Also, [22] can handle a restricted class of adversaries that are *multiplicative* (i.e., multiply each coordinate of gradient by arbitrary scalar) and *blind* (i.e., determine how to corrupt the gradient before observing the true

gradient). Here, for compression, we use a generic δ approximate compressor. Also, we can handle arbitrary Byzantine worker machines.

Very recently, [20] uses error-feedback to remove some of the issues of sign based compression schemes. In this work, we extend the framework to a distributed setting and prove theoretical guarantees in the presence of Byzantine worker machines.

Notation: We assume $C, C_1, C_2, \dots, c, c_1, \dots$ as positive universal constants, the value of which may differ from instance to instance. $[r]$ denotes the set of natural numbers $\{1, 2, \dots, r\}$. Also, $\|\cdot\|$ denotes the ℓ_2 norm of a vector and the operator norm of a matrix unless otherwise specified.

3.2 Problem Formulation

In this section, we formally set up the problem. We consider a standard statistical problem of risk minimization. In a distributed setting, suppose we have one central and m worker nodes and the worker nodes communicate to the central node. Each worker node contains n data points. We assume that the mn data points are sampled independently from some unknown distribution \mathcal{D} . Also, let $f(w, z)$ be the non-convex loss function of a parameter vector $w \in \mathcal{W} \subseteq \mathbb{R}^d$ corresponding to data point z , where \mathcal{W} is the parameter space. Hence, the population loss function is $F(w) = \mathbb{E}_{z \sim \mathcal{D}}[f(w, z)]$. Our goal is to obtain the following:

$$w^* = \operatorname{argmin}_{w \in \mathcal{W}} F(w),$$

where we assume \mathcal{W} to be a convex and compact subset of \mathbb{R}^d with diameter D . In other words, we have $\|w_1 - w_2\| \leq D$ for all $w_1, w_2 \in \mathcal{W}$. Each worker node is associated with a local loss defined as $F_i(w) = \frac{1}{n} \sum_{j=1}^n f(w, z^{i,j})$, where $z^{i,j}$ denotes the j -th data point in the i -th machine. This is precisely the empirical risk function of the i -th worker node.

We assume a setup where worker i compresses the local gradient and sends to the central machine. The central machine aggregates the compressed gradients, takes a gradient step to update the model and broadcasts the updated model to be used in the subsequent iteration. Furthermore, we assume that α fraction of the total workers nodes are Byzantine, for some $\alpha < 1/2$. Byzantine workers can send any arbitrary values to the central machine. In addition, Byzantine workers may completely know the learning algorithm and are allowed to collude with each other.

Next, we define a few (standard) quantities that will be required in our analysis.

Definition 6. (*Sub-exponential random variable*) A zero mean random variable Y is called v -sub-exponential if $\mathbb{E}[e^{\lambda Y}] \leq e^{\frac{1}{2}\lambda^2 v^2}$, for all $|\lambda| \leq \frac{1}{v}$.

Definition 7. (*Smoothness*) A function $h(\cdot)$ is L_F -smooth if $h(w) \leq h(w') + \langle \nabla h(w'), w - w' \rangle + \frac{L_F}{2} \|w - w'\|^2 \forall w, w'$.

Definition 8. (*Lipschitz*) A function $h(\cdot)$ is L -Lipschitz if $\|h(w) - h(w')\| \leq L \|w - w'\| \forall w, w'$.

3.3 Compression At Worker Machines

In this section, we consider a generic class of compressors from [15] and [20] as described in the following.

Definition 9 (δ -Approximate Compressor). *An operator $\mathcal{Q}(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is defined as δ -approximate compressor on a set $\mathcal{S} \subseteq \mathbb{R}^d$ if, $\forall x \in \mathcal{S}$,*

$$\|\mathcal{Q}(x) - x\|^2 \leq (1 - \delta)\|x\|^2,$$

where $\delta \in (0, 1]$ is the compression factor.

Furthermore, a randomized operator $\mathcal{Q}(\cdot)$ is δ -approximate compressor on a set $\mathcal{S} \subseteq \mathbb{R}^d$ if,

$$\mathbb{E} (\|\mathcal{Q}(x) - x\|^2) \leq (1 - \delta)\|x\|^2$$

holds for all $x \in \mathcal{S}$, where the expectation is taken with respect to the randomness of $\mathcal{Q}(\cdot)$. For the clarity of exposition, we consider the deterministic form of the compressor (as in Definition 9). However, the results can be easily extended for randomized $\mathcal{Q}(\cdot)$.

Notice that $\delta = 1$ implies $\mathcal{Q}(x) = x$ (no compression). We list a few examples of δ -approximate compressors (including a few from [20]) here:

1. top_k operator, which selects k coordinates with largest absolute value; for $1 \leq k \leq d$, $(\mathcal{Q}(x))_i = (x)_{\pi(i)}$ if $i \leq k$, and 0 otherwise, where π is a permutation of $[d]$ with $(|x|)_{\pi(i)} \geq (|x|)_{\pi(i+1)}$ for $i \in [d - 1]$. This is a k/d -approximate compressor.
2. k -PCA that uses top k eigenvectors to approximate a matrix X ([114]).
3. Quantized SGD (QSGD) [21], where $\mathcal{Q}(x_i) = \|x\| \cdot \text{sgn}(x_i) \cdot \xi_i(x)$, where $\text{sgn}(x_i)$ is the coordinate-wise sign vector, and $\xi_i(x)$ is defined as following: let $0 \leq l \leq s$, be an integer such that $|x_i|/\|x\| \in [l/s, (l + 1)/s]$. Then, $\xi_i = l/s$ with probability $1 - \frac{|x_i|}{c\|x\|\sqrt{d}} + l$ and $(l + 1)/s$ otherwise. [21] shows that it is a $1 - \min(d/s^2, \sqrt{d}/s)$ -approximate compressor.
4. Quantized SGD with ℓ_1 norm [20], $\mathcal{Q}(x) = \frac{\|x\|_1}{d} \text{sgn}(x)$, which is $\frac{\|x\|_1^2}{d\|x\|^2}$ -approximate compressor. Here, we call this compression scheme as ℓ_1 -QSGD.

Apart from these examples, several randomized compressors are also discussed in [15]. Also, the *signSGD* compressor, $\mathcal{Q}(x) = \text{sgn}(x)$, where $\text{sgn}(x)$ is the (coordinate-wise) sign operator, was proposed in [22, 117]. Here the local machines send a d -dimensional vector containing coordinate-wise sign of the gradients.

Algorithm 2: Robust Compressed Gradient Descent

1: **Input:** Step size γ , Compressor $\mathcal{Q}(\cdot)$, $q > 1$, $\beta < 1$. Also define,

$$\mathcal{C}(x) = \begin{cases} \{\mathcal{Q}(x), \|x\|_q\} & \forall x \in \mathbb{R}^d & \text{Option I} \\ \{\mathcal{Q}(x), \|\mathcal{Q}(x)\|_q\} & \forall x \in \mathbb{R}^d & \text{Option II} \end{cases}$$

2: **Initialize:** Initial iterate $w_0 \in \mathcal{W}$

3: **for** $t = 0, 1, \dots, T - 1$ **do**

4: Central machine: broadcasts w_t

for $i \in [m]$ **do in parallel**

5: i -th worker machine:

- Non-Byzantine:
 - Computes $\nabla F_i(w_t)$; sends $\mathcal{C}(\nabla F_i(w_t))$ to the central machine,
- Byzantine:
 - Generates \star (arbitrary), and sends $\mathcal{C}(\star)$ to the central machine: Option I,
 - Sends \star to the central machine: Option II,

end for

6: Central Machine:

- Sort the worker machines in a non decreasing order according to
 - Local gradient norm: Option I,
 - Compressed local gradient norm: Option II,
- Return the indices of the first $1 - \beta$, fraction of elements as \mathcal{U}_t ,
- Update model parameter: $w_{t+1} = \Pi_{\mathcal{W}} \left(w_t - \frac{\gamma}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} \mathcal{Q}(\nabla F_i(w_t)) \right)$.

7: **end for**

3.4 Robust Compressed Gradient Descent

In this section, we describe a communication-efficient and robust distributed gradient descent algorithm for δ -approximate compressors. The optimization algorithm we use is formally given in Algorithm 2. Note that the algorithm uses a compression scheme $\mathcal{Q}(\cdot)$ to reduce communication cost and a norm based thresholding to remove Byzantine worker nodes. As seen in Algorithm 2,

robust compressed gradient descent operates under two different setting, namely *Option I* and *Option II*.

Option I and II are analyzed in Sections 3.5 and 3.6 respectively. For Option I, we use a δ -approximate compressor along with the norm information. In particular, we use: $\mathcal{C}(x) = \{\|x\|_q, \mathcal{Q}(x)\}$ where $q \geq 1$. $\mathcal{C}(x)$ is comprised of a scalar (norm of x) and a compressed vector $\mathcal{Q}(x)$. For compressors such as QSGD ([21]) and ℓ_1 -QSGD ([20]), the quantity $\mathcal{Q}(\cdot)$ has the norm information and hence sending the norm separately is not required.

As seen in Option I of Algorithm 2, worker node i compresses the local gradient $\nabla F_i(\cdot)$ sends $\mathcal{C}(\nabla F_i(\cdot))$ to the central machine. Adversary nodes can send arbitrary $\mathcal{C}(\star)$ to the central machine. The central machine aggregates the gradients, takes a gradient step and broadcasts the updated model for next iteration.

For Option I, we restrict to the setting where the Byzantine worker machines can send arbitrary values to the input of the compression algorithm, but they adhere to the compression algorithm. In particular, Byzantine workers can provide arbitrary values, \star to the input of the compression algorithm, $\mathcal{Q}(\cdot)$ but they correctly implement the same compression algorithm, i.e., computes $\mathcal{Q}(\star)$.

We now explain how Algorithm 2 tackles the Byzantine worker machines. The central machine receives the compressed gradients comprising a scalar ($\|x\|_q, q \geq 1$) and a quantized vector ($\mathcal{Q}(x)$) and outputs a set of indices \mathcal{U} with $|\mathcal{U}| = (1 - \beta)m$. Here we employ a simple thresholding scheme on the (local) gradient norm. Note that, if the Byzantine worker machines try to diverge the learning algorithm by increasing the norm of the local gradients; Algorithm 2 can identify them as outliers. Furthermore, when the Byzantine machines behave like inliers, they can not diverge the learning algorithm since $\alpha < 1/2$. In the subsequent sections, we show theoretical justification of this argument.

With Option II, we remove this restriction on Byzantine machines at the cost of slightly weakening the convergence guarantees. This is explained in Section 3.6. With Option II, the i -th local machine sends $\mathcal{C} = \{\mathcal{Q}(\nabla F_i(w_t)), \|\mathcal{Q}(\nabla F_i(w_t))\|_q\}$ to the central machine, where $q \geq 1$. Effectively, the i -th local machine just sends $\mathcal{Q}(\nabla F_i(w_t))$ since its norm can be computed at the central machine. Byzantine workers just send arbitrary (\star) vector instead of compressed local gradient. Note that the Byzantine workers here do not adhere to any compression rule.

The Byzantine resilience scheme with Option II is similar to Option I except the fact that the central machine sorts the worker machines according to the norm of the compressed gradients rather than the norm of the gradients.

3.5 Distributed Learning with Restricted Adversaries

In this section, we analyze the performance of Algorithm 4 with *Option I*. We restrict to an adversarial model in which Byzantine workers can provide arbitrary values to the input of the compression algorithm, but they adhere to the compression rule. Though this adversarial model is restricted, we argue that it is well-suited for applications wherein compression happens outside of worker machines. For example, Apache MXNet, a deep learning framework designed to be distributed on cloud infrastructures, uses NVIDIA Collective Communication Library (NCCL) that

employs gradient compression (see [124]). Also, in a Federated Learning setup the compression can be part of the communication protocol. Furthermore, this can happen when worker machines are divided into groups, and each group is associated with a *compression unit*. As an example, cores in a multi-core processor ([125]) acting as a group of worker machines with the compression carried out by a separate processor, or servers co-located on a rack ([126]) acting as a group with the compression carried out by the top-of-the-rack switch.

3.5.1 Main Results

We analyze Algorithm 4 (with Option I) and obtain the rate of the convergence under non-convex loss functions. We start with the following assumption.

Assumption 9. *For all z , the partial derivative of the loss function $f(\cdot, z)$ with respect to the k -th coordinate (denoted as $\partial_k f(\cdot, z)$) is L_k Lipschitz with respect to the first argument for each $k \in [d]$, and let $\hat{L} = \sqrt{\sum_{i=1}^d L_k^2}$. The population loss function $F(\cdot)$ is L_F smooth.*

We also make the following assumption on the tail behavior of the partial derivative of the loss function.

Assumption 10. *(Sub-exponential gradients) For all $k \in [d]$ and z , the quantity $\partial_k f(w, z)$ is v sub-exponential for all $w \in \mathcal{W}$.*

The assumption implies that the moments of the partial derivatives are bounded. We like to emphasize that the sub-exponential assumption on gradients is fairly common ([18, 23]). For instance, [23, Proposition 2] gives a concrete example of coordinate-wise sub-exponential gradients in the context of a regression problem. Furthermore, in [31], the gradients are assumed to be sub-gaussian, which is stronger than Assumption 10.

To simplify notation and for the clarity of exposition, we define the following three quantities which will be used throughout the chapter.

$$\epsilon_1 = v\sqrt{d} \left(\max \left\{ \frac{d}{n} \log(1 + 2nD\hat{L}d), \sqrt{\frac{d}{n} \log(1 + 2nD\hat{L}d)} \right\} \right) + \frac{1}{n}, \quad (3.2)$$

$$\epsilon_2 = v\sqrt{d} \left(\max \left\{ \frac{d}{(1-\alpha)mn} \log(1 + 2(1-\alpha)mnD\hat{L}d), \sqrt{\frac{d}{(1-\alpha)mn} \log(1 + 2(1-\alpha)mnD\hat{L}d)} \right\} \right), \quad (3.3)$$

$$\epsilon = 2 \left(1 + \frac{1}{\lambda_0} \right) \left[\left(\frac{1-\alpha}{1-\beta} \right)^2 \epsilon_2^2 + \left(\frac{\sqrt{1-\delta} + \alpha + \beta}{1-\beta} \right)^2 \epsilon_1^2 \right]. \quad (3.4)$$

where λ_0 is a positive constant. For intuition, one can think of $\epsilon_1 = \tilde{\mathcal{O}}\left(\frac{d}{\sqrt{n}}\right)$ and $\epsilon_2 = \tilde{\mathcal{O}}\left(\frac{d}{\sqrt{mn}}\right)$ as small problem dependent quantities. Assuming $\beta = c\alpha$ for a universal constant $c > 1$, we have

$$\epsilon = \tilde{\mathcal{O}} \left(d^2 \left[\frac{\alpha^2}{n} + \frac{1-\delta}{n} + \frac{1}{mn} \right] \right). \quad (3.5)$$

Assumption 11. (Size of parameter space \mathcal{W}) Suppose that $\|\nabla F(w)\| \leq M$ for all $w \in \mathcal{W}$. We assume that \mathcal{W} contains the ℓ_2 ball $\{w : \|w - w_0\| \leq c[(2 - \frac{c_0}{2})M + \sqrt{\epsilon}] \frac{F(w_0) - F(w^*)}{\epsilon}\}$, where c_0 is a constant, δ is the compression factor, w_0 is the initial parameter vector and ϵ is defined in equation (3.4).

We use the above assumption to ensure that the iterates of Algorithm 4 stays in \mathcal{W} . We emphasize that this is a standard assumption on the size of \mathcal{W} to control the iterates for non-convex loss function. Note that, similar assumptions have been used in prior works [23, Assumption 5], [31]. We point out that Assumption 11 is used for simplicity and is not a hard requirement. We show (in the proof of Theorem 3) that the iterates of Algorithm 4 stay in a bounded set around the initial iterate w_0 . Also, note that the dependence of M in the final statistical rate (implicit, via diameter D) is logarithmic (weak dependence), as will be seen in Theorem 3.

We provide the following rate of convergence to a critical point of the (non-convex) population loss function $F(\cdot)$.

Theorem 3. Suppose Assumptions 9, 10 and 11 hold, and $\alpha \leq \beta < 1/2$. For sufficiently small constant c , we choose the step size $\gamma = \frac{c}{L_F}$. Then, running Algorithm 4 for $T = C_3 \frac{L_F(F(w_0) - F(w^*))}{\epsilon}$ iterations yields

$$\min_{t=0, \dots, T} \|\nabla F(w_t)\|^2 \leq C \epsilon,$$

with probability greater than or equal to $1 - \frac{c_1(1-\alpha)md}{(1+n\bar{L}D)^d} - \frac{c_2d}{(1+(1-\alpha)mn\bar{L}D)^d}$, provided the compression factor satisfies $\delta > \delta_0 + 4\alpha - 9\alpha^2 + 4\alpha^3$, where $\delta_0 = \left(1 - \frac{(1-\beta)^2}{1+\lambda_0}\right)$ and λ_0 is a (sufficiently small) positive constant.

A few remarks are in order. In the following remarks, we fix the dimension d , and discuss the dependence of ϵ on (α, δ, n, m) .

Remark 1. We observe, from the definition of ϵ that the price for compression is $\tilde{O}(\frac{1-\delta}{n})$.

Remark 2. Substituting $\delta = 1$ (no compression) in ϵ , we get $\epsilon = \tilde{O}(\frac{\alpha^2}{n} + \frac{1}{mn})$, which matches the (statistical) rate of [23]. A simple *norm based thresholding* operation is computationally simple and efficient in the high dimensional settings compared to the coordinate wise median and trimmed mean to achieve robustness and obtain the the same statistical error and iteration complexity as [23]

Remark 3. When the compression factor δ is large enough, satisfying $\delta \geq 1 - \alpha^2$, we obtain $\epsilon = \tilde{O}(\frac{\alpha^2}{n} + \frac{1}{mn})$. In this regime, the iteration complexity and the final statistical error of Algorithm 4 is order-wise identical to the setting with no compression [23]. We emphasize here that a reasonable high δ is often observed in practical applications like training of neural nets [20, Figure 2].

Remark 4. (Optimality) For a distributed mean estimation problem, Observation 1 in [23] implies that any algorithm will yield an (statistical) error of $\Omega(\frac{\alpha^2}{n} + \frac{d}{mn})$. Hence, in the regime where $\delta \geq 1 - \alpha^2$, our error-rate is optimal.

Remark 5. For the convergence of Algorithm 4, we require $\delta > \delta_0 + 4\alpha - 9\alpha^2 + 4\alpha^3$, implying that our analysis will not work if δ is very close to 0. Note that a very small δ does not give good accuracy in practical applications [20, Figure 2]. Also, note that, from the definition of δ_0 , we can choose λ_0 sufficiently small at the expense of increasing the multiplicative constant in ϵ by a factor of $1/\lambda_0$. Since the error-rate considers asymptotic in m and n , increasing a constant factor is insignificant. A sufficiently small λ_0 implies $\delta_0 = \mathcal{O}(2\beta)$, and hence we require $\delta > 4\alpha + 2\beta$ (ignoring the higher order dependence).

Remark 6. The requirement $\delta > 4\alpha + 2\beta$ can be seen as a trade-off between the amount of compression and the fraction of adversaries in the system. As α increases, the amount of (tolerable) compression decreases and vice versa.

Remark 7. (Rate of Convergence) Algorithm 4 with T iterations yields

$$\min_{t=0,\dots,T} \|\nabla F(w_t)\|^2 \leq \frac{C_1 L_F (F(w_0) - F(w^*))}{T + 1} + C_2 \epsilon$$

with high probability. We see that Algorithm 4 converges at a rate of $\mathcal{O}(1/T)$, and finally plateaus at an error floor of ϵ . Note that the rate of convergence is same as [23]. Hence, even with compression, the (order-wise) convergence rate is unaffected.

3.6 Distributed Optimization with Arbitrary Adversaries

In this section we remove the assumption of restricted adversary (as in Section 3.5) and make the learning algorithm robust to the adversarial effects of both the computation and compression unit. In particular, here we consider Algorithm 4 with Option II. Hence, the Byzantine machines do not need to adhere to the mandated compression algorithm. However, in this setting, the statistical error-rate of our proposed algorithm is slightly weaker than that of Theorem 3. Furthermore, the (δ, α) trade-off is stricter compared to Theorem 3.

3.6.1 Main Results

We continue to assume that the population loss function $F(\cdot)$ is smooth and non-convex and analyze Algorithm 4 with Option II. We have the following result. For the clarity of exposition, we define the following quantity which will be used in the results of this section:

$$\tilde{\epsilon} = 2\left(1 + \frac{1}{\lambda_0}\right) \left(\left(\frac{(1+\beta)\sqrt{1-\delta} + \alpha + \beta}{1-\beta} \right)^2 \epsilon_1^2 + \left(\frac{1-\alpha}{1-\beta} \right)^2 \epsilon_2^2 \right).$$

Comparing $\tilde{\epsilon}$ with ϵ , we observe that $\tilde{\epsilon} > \epsilon$. Also, note that,

$$\tilde{\epsilon} = \tilde{\mathcal{O}} \left(d^2 \left[\frac{\alpha^2}{n} + \frac{1-\delta}{n} + \frac{1}{mn} \right] \right), \quad (3.6)$$

which suggests that $\tilde{\epsilon}$ and ϵ are order-wise similar. We have the following assumption, which parallels Assumption 11, with ϵ replaced by $\tilde{\epsilon}$.

Assumption 12. (Size of parameter space \mathcal{W}) Suppose that $\|\nabla F(w)\| \leq M$ for all $w \in \mathcal{W}$. We assume that \mathcal{W} contains the ℓ_2 ball $\{w : \|w - w_0\| \leq c[(2 - \frac{c_0}{2})M + \sqrt{\tilde{\epsilon}}] \frac{F(w_0) - F(w^*)}{\tilde{\epsilon}}\}$, where c_0 is a constant, δ is the compression factor and $\tilde{\epsilon}$ is defined in equation (3.6).

Theorem 4. Suppose Assumptions 9,10 and 12 hold, and $\alpha \leq \beta < 1/2$. For sufficiently small constant c , we choose the step size $\gamma = \frac{c}{L_F}$. Then, running Algorithm 4 for $T = C_3 \frac{L_F(F(w_0) - F(w^*))}{\tilde{\epsilon}}$ iterations yields

$$\min_{t=0, \dots, T} \|\nabla F(w_t)\|^2 \leq C\tilde{\epsilon},$$

with probability greater than or equal to $1 - \frac{c_1(1-\alpha)md}{(1+n\hat{L}D)^d} - \frac{c_2d}{(1+(1-\alpha)mn\hat{L}D)^d}$, provided the compression factor satisfies $\delta > \tilde{\delta}_0 + 4\alpha - 8\alpha^2 + 4\alpha^3$, where $\tilde{\delta}_0 = \left(1 - \frac{(1-\beta)^2}{(1+\beta)^2(1+\lambda_0)}\right)$ and λ_0 is a (sufficiently small) positive constant.

Remark 8. The above result and their consequences resemble that of Theorem 3. Since $\tilde{\epsilon} > \epsilon$, the statistical error-rate in Theorem 4 is strictly worse than that of Theorem 3 (although order-wise they are same).

Remark 9. Note that the definition of δ_0 is different than in Theorem 3. For a sufficiently small λ_0 , we see $\tilde{\delta}_0 = \mathcal{O}(4\beta)$, which implies we require $\delta > 4\beta + 4\alpha$ for the convergence of Theorem 4. Note that this is a slightly strict requirement compared to Theorem 3. In particular, for a given δ , Algorithm 4 with Option II can tolerate less number of Byzantine machines compared to Option I.

Remark 10. The result in Theorem 4 is applicable for arbitrary adversaries, whereas Theorem 3 relies on the adversary being restrictive. Hence, we can view the limitation of Theorem 4 (such as worse statistical error-rate and stricter (δ, α) trade-off) as a price of accommodating arbitrary adversaries.

3.7 Byzantine Robust Distributed Learning with Error Feedback

We now investigate the role of error feedback [20] in distributed learning with Byzantine worker machines. We stick to the formulation of Section 3.1.

In order to address the issues of convergence for sign based algorithms (like *signSGD*), [20] proposes a class of optimization algorithms that uses *error feedback*. In this setting, the worker machine locally stores the error between the actual local gradient and its compressed counterpart. Using this as feedback, the worker machine adds this error term to the compressed gradient in the subsequent iteration. Intuitively, this accounts for correcting the the direction of the local gradient. The error-feedback has its roots in some of the classical communication system like “delta-sigma” modulator and adaptive modulator ([127]).

We analyze the distributed error feedback algorithm in the presence of Byzantine machines. The algorithm is presented in Algorithm 3. We observe that here the central machine sorts the worker machines according to the norm of the compressed local gradients, and ignore the largest β fraction.

Algorithm 3: Distributed Compressed Gradient Descent with Error Feedback

- 1: **Input:** Step size γ , Compressor $\mathcal{Q}(\cdot)$, parameter $\beta (> \alpha)$.
 - 2: **Initialize:** Initial iterate w_0 , $e_i(0) = 0 \forall i \in [m]$
 - 3: **for** $t = 0, 1, \dots, T - 1$ **do**
 - 4: Central machine: sends w_t to all worker
 for $i \in [m]$ **do in parallel**
 - 5: i -th non-Byzantine worker machine:
 - computes $p_i(w_t) = \gamma \nabla F_i(w_t) + e_i(t)$
 - sends $\mathcal{Q}(p_i(w_t))$ to the central machine
 - computes $e_i(t + 1) = p_i(w_t) - \mathcal{Q}(p_i(w_t))$
 - 6: Byzantine worker machine:
 - sends \star to the central machine.
 - 7: At Central machine:
 - sorts the worker machines in non-decreasing order according to $\|\mathcal{Q}(p_i(w_t))\|$.
 - returns the indices of the first $1 - \beta$ fraction of elements as \mathcal{U}_t .
 - $w_{t+1} = \Pi_{\mathcal{W}} \left(w_t - \frac{\gamma}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} \mathcal{Q}(p_i(w_t)) \right)$
 - 8: **end for**
-

Note that, similar to Section 3.6, we handle arbitrary adversaries. In the subsequent section, we show (both theoretically and experimentally) that the statistical error rate of Algorithm 3 is smaller than Algorithm 4.

3.7.1 Main Results

In this section we analyze Algorithm 3 and obtain the rate of the convergence under non-convex smooth loss functions. Throughout the section, we select γ as the step size and assume that Algorithm 3 is run for T iterations. We start with the following assumption.

Assumption 13. For all non-Byzantine worker machine i , the local loss functions $F_i(\cdot)$ satisfy $\|\nabla F_i(x)\|^2 \leq \sigma^2$, where $x \in \{w_j\}_{j=0}^T$, and $\{w_0, \dots, w_T\}$ are the iterates of Algorithm 3.

Note that since $F_i(\cdot)$ can be written as loss over data points of machine i , we observe that the bounded gradient condition is equivalent to the bounded second moment condition for SGD, and have featured in several previous works, see, e.g., [128], [129]. Here, we are using all the data points and (hence no randomness over the choice of data points) perform gradient descent instead

of SGD. Also, note that Assumption 13 is weaker than the bounded second moment condition since we do not require $\|\nabla F_i(x)\|^2$ to be bounded for all x ; just when $x \in \{w_j\}_{j=0}^T$.

We also require the following assumption on the size of the parameter space \mathcal{W} , which parallels Assumption 11 and 12.

Assumption 14. (Size of parameter space \mathcal{W}) Suppose that $\|\nabla F(w)\| \leq M$ for all $w \in \mathcal{W}$. We assume that \mathcal{W} contains the ℓ_2 ball $\{w : \|w - w_0\| \leq \gamma r^* T\}$, where

$$r^* = \epsilon_2 + M + \frac{6\beta(1 + \sqrt{1 - \delta})}{(1 - \beta)} \left(\epsilon_1 + M + \sqrt{\frac{3(1 - \delta)}{\delta}} \sigma \right) + \sqrt{\frac{12(1 - \delta)}{\delta}} \sigma,$$

and (ϵ_1, ϵ_2) are defined in equations (3.2) and (3.3) respectively.

Similar to Assumption 11 and 12, we use the above assumption to ensure that the iterates of Algorithm 3 stays in \mathcal{W} , and we emphasize that this is a standard assumption to control the iterates for non-convex loss function (see [23, 31]).

To simplify notation and for the clarity of exposition, we define the following quantities which will be used in the main results of this section.

$$\Delta_1 = \frac{9(1 + \sqrt{1 - \delta})^2}{2c(1 - \beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \left(\epsilon_1^2 + \frac{3(1 - \delta)}{\delta} \sigma^2 \right) + \frac{50}{c} \epsilon_2^2, \quad (3.7)$$

$$\Delta_2 = \frac{L^2}{2} \frac{3(1 - \delta)\sigma^2}{c\delta} + \frac{2L\epsilon_2^2}{c} + \left(\frac{1}{2} + L \right) \frac{9(1 + \sqrt{1 - \delta})^2}{c(1 - \beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \left(\epsilon_1^2 + \frac{3(1 - \delta)}{\delta} \sigma^2 \right), \quad (3.8)$$

$$\Delta_3 = \left(\frac{L^2}{100} + 25L^2 \right) \frac{3(1 - \delta)\sigma^2}{c\delta}, \quad (3.9)$$

where c is a universal constant.

We show the following rate of convergence to a critical point of the population loss function $F(\cdot)$.

Theorem 5. Suppose Assumptions 9, 10, 13 and 14 hold, and $\alpha \leq \beta < 1/2$. Then, running Algorithm 4 for T iterations with step size γ yields

$$\min_{t=0, \dots, T} \|\nabla F(w_t)\|^2 \leq \frac{F(w_0) - F^*}{c\gamma(T + 1)} + \Delta_1 + \gamma\Delta_2 + \gamma^2\Delta_3,$$

with probability greater than or equal to $1 - \frac{c_1(1 - \alpha)md}{(1 + n\hat{L}D)^d} - \frac{c_2d}{(1 + (1 - \alpha)mn\hat{L}D)^d}$, provided the compression factor satisfies $\frac{(1 + \sqrt{1 - \delta})^2}{(1 - \beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] < 0.107$. Here Δ_1, Δ_2 and Δ_3 are defined in equations (3.7), (3.8) and (3.9) respectively.

Remark 11. (Choice of Step Size γ) Substituting $\gamma = \frac{1}{\sqrt{T+1}}$, we obtain

$$\min_{t=0, \dots, T} \|\nabla F(w_t)\|^2 \leq \frac{F(w_0) - F^*}{c\sqrt{T+1}} + \Delta_1 + \frac{\Delta_2}{\sqrt{T+1}} + \frac{\Delta_3}{T+1},$$

with high probability. Hence, we observe that the quantity associated with Δ_3 goes down at a considerably faster rate ($\mathcal{O}(1/T)$) than the other terms and hence can be ignored, when T is large.

Remark 12. Note that when no Byzantine worker machines are present, i.e., $\alpha = \beta = 0$, we obtain

$$\Delta_1 = \frac{50}{c}\epsilon_2^2, \quad \Delta_2 = \frac{L^2}{2} \frac{3(1-\delta)\sigma^2}{c\delta} + \frac{2L\epsilon_2^2}{c}, \quad \Delta_3 = \left(\frac{L^2}{100} + 25L^2\right) \frac{3(1-\delta)\sigma^2}{c\delta}.$$

Additionally, if $\delta = \Theta(1)$ (this is quite common in applications like training of neural nets, as mentioned earlier), we obtain $\Delta_2 = C(L^2\sigma^2 + L\epsilon_2^2)$, and $\Delta_3 = C_1L^2$. Substituting $\epsilon_2 = \mathcal{O}\left(\frac{d}{\sqrt{mn}}\right)$ and for a fixed d , the upper bound in the above theorem is order-wise identical to that of standard SGD in a population loss minimization problem under similar setting.

Remark 13. (No compression setting) In the setting, where $\delta = 1$ (no compression), we obtain

$$\Delta_1 = \mathcal{O} \left[d^2 \left(\frac{\alpha^2}{n} + \frac{1}{mn} \right) \right],$$

and

$$\Delta_2 = \mathcal{O} \left[d^2 L \left(\frac{\alpha^2}{n} + \frac{1}{mn} \right) \right],$$

and $\Delta_3 = 0$. The statistical rate (obtained by making T sufficiently large) of the problem is Δ_1 , and this rate matches exactly to that of [23]. Hence, we could recover the optimal rate without compression. Furthermore, this rate is optimal in (α, m, n) as shown in [23].

Remark 14. In the next section, we show that when $\alpha \neq 0$ and $\delta = \Theta(1)$, the statistical error rate of Algorithm 3 is order-wise identical to the no-compression setting. Hence, we get the compression for free. Furthermore, we argue that error feedback improves the statistical (error) rate.

3.7.2 Comparison with Algorithm 4

We now compare the statistical rate of Algorithm 3 with Algorithm 4, where no feedback is used. Recall from Sections 3.5 and 3.6 that the statistical error rate of is given by,

$$\Delta_{no-feedback} = \mathcal{O} \left(d^2 \left[\frac{\alpha^2}{n} + \frac{1-\delta}{n} + \frac{1}{mn} \right] \right).$$

Let us compare it with the statistical rate of our algorithm, given by

$$\Delta_1 = \mathcal{O} \left(\frac{\alpha^2 d^2}{n} + \frac{\alpha^2 (1-\delta)\sigma^2}{\delta} + \frac{d^2}{mn} \right).$$




Figure 3.1: Training (cross entropy) loss for MNIST image. Different types of attack (a) labels with deterministic shift (9 – label) (b) random labels. Plots show theresholding scheme with different type of byzantine attacks achieve similar convergence as ‘no byzantine’ setup.

Observe that in the setting with error feedback, we have an additional problem parameter σ^2 . Hence, for the purpose of comparison, we first argue what the scaling of σ^2 should be. Since, $\|\nabla F_i(w_t)\|^2 \leq 2\|\nabla F_i(w_t) - \nabla F(w_t)\|^2 + 2\|\nabla F(w_t)\|^2$, using Lemma 6 of Appendix 3.10, we have

$$\|\nabla F_i(w_t)\|^2 \leq 2\frac{d^2}{n} + 2\|\nabla F(w_t)\|^2,$$

with high probability. Since from Assumption 13, we obtain $\min_{t=0,\dots,T} \|\nabla F_i(w_t)\|^2 \leq \sigma^2$, we obtain

$$\sigma^2 = \Theta\left(\frac{2d^2}{n} + 2\min_{t=0,\dots,T} \|\nabla F(w_t)\|^2\right).$$

As seen by [20], $\delta = \Theta(1)$ is a reasonable and practical parameter regime, and in this setting, with the above σ^2 , we observe that

$$\Delta_1 < \Delta_{no-feedback}.$$

provided $\alpha^2 \leq c$, for a constant c . Note that this condition is equivalent to the trade-off between the amount of compression and the fraction of Byzantine worker machines, featured in Theorem 5, which was required to show convergence of Algorithm 3. Hence, in the above mentioned parameter regime, the error rate with error feedback is strictly better than no-feedback setting.

In numerical experiments, we observe that the convergence of Algorithm 3 with error feedback is faster than Algorithm 4, which is intuitive since error feedback helps in correcting the direction of the local gradient. We now have a theoretical justification for this fact.




Figure 3.2: Convergence for (a) regression problem (b) training (cross entropy) loss for MNIST image. Plots show convergence beyond the theoretical bound on the number of Byzantine machine.




Figure 3.3: Convergence for (a) regression problem (b) training (cross entropy) loss for MNIST image. Plots show convergence with a natural Byzantine attack of $-\epsilon$ times the local gradient with high number of Byzantine machines for $\epsilon = 0.9$.

3.8 Experiments

In this section we validate the correctness of our proposed algorithms for linear regression problem and training ReLU network. In all the experiments, we choose the following compression scheme: given any $x \in \mathbb{R}^d$, we report $\mathcal{C}(x) = \{\frac{\|x\|_1}{d}, \text{sgn}(x)\}$ where $\text{sgn}(x)$ serves as the quantized vector and $\frac{\|x\|_1}{d}$ is the scaling factor. All the reported results are averaged over 20 different runs.

First we consider a least square regression problem $w^* = \text{argmin}_w \|Aw - b\|_2$. For the regression problem we generate matrix $A \in \mathbb{R}^{N \times d}$, vector $w^* \in \mathbb{R}^d$ by sampling each item independently from standard normal distribution and set $b = Aw^*$. Here we choose $N = 4000$ and consider $d = 1000$. We partition the data set equally into $m = 200$ servers. We randomly choose αm ($= 10, 20$) workers to be Byzantine and apply norm based thresholding operation with parameter




Figure 3.4: Comparison of Robust Compressed Gradient Descent with and without thresholding scheme in a regression problem. The plots show better convergence with thresholding.




Figure 3.5: Comparison of Robust Compressed Gradient Descent with majority vote based signSGD [22] in regression Problem. The plots show better convergence with thresholding in comparison to the majority vote based robustness of [22]

βm ($= 12, 22$) respectively. We simulate the Byzantine workers by adding i.i.d $\mathcal{N}(0, 10I_d)$ entries to the gradient. In our experiments the gradient is the most pertinent information of the the worker server. So we choose to add noise to the gradient to make it a Byzantine worker. However, later on, we consider several kinds of attack models. We choose $\|w_t - w^*\|$ as the error metric for this problem.

Effectiveness of thresholding: We compare Algorithm 4 with compressed gradient descent (with vanilla aggregation). Our method is equipped with Byzantine tolerance steps and the vanilla compressed gradient just computes the average of the compressed gradient sent by the workers. From Figure 3.4 it is evident that the the application of norm based thresholding scheme provides better convergence result compared to the compressed gradient method without it.

Comparison with signSGD with majority vote: In [22], a communication efficient byzantine tolerant algorithm is proposed where communication efficiency is achieved by communicating sign of the gradient and robustness is attained by taking co-ordinate wise majority vote. The robustness in our algorithm comes from thresholding operation on the scaling factor. We show a comparison of both method in Figure 3.5 in the regression setup depicted above. Our method shows a better trend in convergence.

Error-feedback with thresholding scheme: We demonstrate the effectiveness of Byzantine resilience with error-feedback scheme as described in Algorithm 3. We compare our scheme with Algorithm 4 (which does not use error feedback) in Figure 3.6.

Feed-forward Neural Net with ReLU activation: Next, we show the effectiveness of our method in training a fully connected feed forward neural net. We implement the neural net in pytorch and use the digit recognition dataset MNIST ([106]). We partition 60,000 training data into 200 different worker nodes. The neural net is equipped with 1000 node hidden layer with ReLU activation function and we choose *cross-entropy-loss* as the loss function. We simulate the Byzantine workers by adding i.i.d $\mathcal{N}(0, 10I_d)$ entries to the gradient. In Figure 3.7 we compare our robust compressed gradient descent scheme with the trimmed mean scheme of [23] and majority vote based *signSGD* scheme of [22]. Compared to the majority vote based scheme, our scheme converges faster. Further, our method shows as good as performance of trimmed mean despite the fact the robust scheme of [23] is an uncompressed scheme and uses a more complicated aggregation rules.

Different Types of Attacks: In the previous paragraph we compared our scheme with existing scheme with additive Gaussian noise as a form of byzantine attack. We also show convergence results with the following type of attacks, which are quite common ([23]) in neural net training with digit recognition dataset [106]. (a) *Random label*: the byzantine worker machines randomly replaces the labels of the data, and (b) *Deterministic Shift*: byzantine workers in a deterministic manner replace the labels y with $9 - y$ (0 becomes 9, 9 becomes 0). In Figure 3.1 we show the convergence results with different numbers of byzantine worker nodes.

Large Number of Byzantine Workers: In Figures 3.2 and 3.3, we show the convergence results that holds beyond the theoretical limit (as shown in Theorem 3 and 4) of the number of Byzantine servers in the regression problem and neural net training. In Figure 3.2, for the regression problem, the Byzantine attack is additive Gaussian noise as described before and our algorithm is robust up to 40% ($\alpha = .4$) of the workers being Byzantine. While training of the feed-forward neural network, we apply a deterministic shift as the Byzantine attack, and the algorithm converges even for 40% ($\alpha = .4$) Byzantine workers.

Note that our robust algorithm is in essence a stochastic gradient descent algorithm. Thus, a ‘natural’ Byzantine attack would be when a Byzantine worker sends $-\epsilon g$ where $0 \leq \epsilon \leq 1$ and g is the local gradient making the algorithm ‘ascent’ type. We choose $\epsilon = 0.9$ and show convergence




Figure 3.6: Comparison of norm based thresholding with and without error feedback. The plots show that error feedback based scheme offers better convergence.




Figure 3.7: Training (cross entropy) loss for MNIST image. Comparison with (a) Uncompressed Trimmed mean [23] (b) majority based signSGD of [22]. In plot (a) show that Robust Gradient descent matches the convergence of the uncompressed trimmed mean [23]. Plot (b) show a faster convergence compared to the algorithm of [22].

for the regression problem for up to 40% byzantine workers, and for the neural network training for up to 33% Byzantine workers in Figure 3.3.

3.9 Conclusion and Open Problems

We address the problem of robust distributed optimization where the worker machines send the compressed gradient (as opposed to the full gradient) to the central machine. We propose a first order optimization algorithm and provide theoretical guarantees and experimental validation under different setup. In some settings, we assume a restricted adversary (that adheres to the compression

algorithm). An immediate future work would be to remove such assumption and obtain a learning algorithm with arbitrary adversaries uniformly for all δ -approximate compressors. It might also be interesting to study a second order distributed optimization algorithm with compressed gradients and Hessians.

Appendix

3.10 Analysis of Algorithm 4

In this section, we provide analysis of the Lemmas required for the proof of Theorem 3 and Theorem 4.

Notation: Let \mathcal{M} and \mathcal{B} denote the set of non-Byzantine and Byzantine worker machines. Furthermore, \mathcal{U}_t and \mathcal{T}_t denote untrimmed and trimmed worker machines. So evidently,

$$|\mathcal{M}| + |\mathcal{B}| = |\mathcal{U}_t| + |\mathcal{T}_t| = m.$$

3.10.1 Proof of Theorem 3

Let $g(w_t) = \frac{1}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} \mathcal{Q}(\nabla F_i(w_t))$ and $\Delta = g(w_t) - \nabla F(w_t)$. We have the following Lemma to control of $\|\Delta\|^2$.

Lemma 4. For any $\lambda > 0$, we have,

$$\|\Delta\|^2 \leq (1 + \lambda) \left(\frac{\sqrt{1 - \delta} + 2\alpha}{1 - \beta} \right)^2 \|\nabla F(w_t)\|^2 + \tilde{\epsilon}(\lambda)$$

with probability greater than or equal to $1 - \frac{c_1(1-\alpha)md}{(1+n\hat{L}D)^d} - \frac{c_2d}{(1+(1-\alpha)mn\hat{L}D)^d}$, where

$$\tilde{\epsilon}(\lambda) = 2\left(1 + \frac{1}{\lambda}\right) \left[\left(\frac{\sqrt{1 - \delta} + \alpha + \beta}{1 - \beta} \right)^2 \epsilon_1^2 + \left(\frac{1 - \alpha}{1 - \beta} \right)^2 \epsilon_2^2 \right].$$

with ϵ_1 and ϵ_2 as defined in equation (3.2) and (3.3) respectively.

The proof of the lemma is deferred to Section 3.10.3. We prove the theorem using the above lemma.

We first show that with Assumption 11 and with the choice of step size γ , we always stay in \mathcal{W} without projection. Recall that $g(w_t) = \frac{1}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} \mathcal{Q}(\nabla F_i(w_t))$ and $\Delta = g(w_t) - \nabla F(w_t)$. We have

$$\begin{aligned} \|w_{t+1} - w^*\| &\leq \|w_t - w^*\| + \gamma(\|\nabla F(w_t)\| + \|g(w_t) - \nabla F(w_t)\|) \\ &\leq \|w_t - w^*\| + \frac{c}{L_F}(\|\nabla F(w_t)\| + \|\Delta\|) \end{aligned}$$

We use Lemma 4 with $\lambda = \lambda_0$ for a sufficiently small positive constant λ_0 . Define $\delta_0 = \left(1 - \frac{(1-\beta)^2}{1+\lambda_0}\right)$. A little algebra shows that provided $\delta > \delta_0 + 4\alpha - 9\alpha^2 + 4\alpha^3$, we obtain

$$\|\Delta\|^2 \leq (1 - c_0)\|\nabla F(w_t)\|^2 + \epsilon$$

with probability greater than or equal to $1 - \frac{c_1(1-\alpha)md}{(1+n\bar{L}D)^d} - \frac{c_2d}{(1+(1-\alpha)mn\bar{L}D)^d}$, where c_0 is a positive constant and ϵ is defined in equation (3.4). Substituting, we obtain

$$\begin{aligned} \|w_{t+1} - w^*\| &\leq \|w_t - w^*\| + \frac{c_1}{L_F} \left((1 + \sqrt{1 - c_0})\|\nabla F(w_t)\| + \sqrt{\epsilon} \right) \\ &\leq \|w_t - w^*\| + \frac{c_1}{L_F} \left(\left(2 - \frac{c_0}{2}\right)\|\nabla F(w_t)\| + \sqrt{\epsilon} \right). \end{aligned}$$

where we use the fact that $\sqrt{1 - c_0} \leq 1 - c_0/2$. Now, running $T = C \frac{L_F(F(w_0) - F(w^*))}{\epsilon}$ iterations, we see that Assumption 11 ensures that the iterations of Algorithm 4 is always in \mathcal{W} . Hence, let us now analyze the algorithm without the projection step.

Using the smoothness of $F(\cdot)$, we have

$$F(w_{t+1}) \leq F(w_t) + \langle \nabla F(w_t), w_{t+1} - w_t \rangle + \frac{L_F}{2} \|w_{t+1} - w_t\|^2.$$

Using the iteration of Algorithm 4, we obtain

$$\begin{aligned} F(w_{t+1}) &\leq F(w_t) - \gamma \langle \nabla F(w_t), \nabla F(w_t) + \Delta \rangle + \frac{\gamma^2 L_F}{2} \|\nabla F(w_t) + \Delta\|^2 \\ &\leq F(w_t) - \gamma \|\nabla F(w_t)\|^2 - \gamma \langle \nabla F(w_t), \Delta \rangle + \frac{\gamma^2 L_F}{2} \|\nabla F(w_t)\|^2 \\ &\quad + \frac{\gamma^2 L_F}{2} \|\Delta\|^2 + \gamma^2 L_F \langle \nabla F(w_t), \Delta \rangle \\ &\leq F(w_t) - \left(\gamma - \frac{\gamma^2 L_F}{2}\right) \|\nabla F(w_t)\|^2 + (\gamma + \gamma^2 L_F) \left(\frac{\rho}{2} \|\nabla F(w_t)\|^2 + \frac{1}{2\rho} \|\Delta\|^2 \right) \\ &\quad + \frac{\gamma^2 L_F}{2} \|\Delta\|^2, \end{aligned}$$

where $\rho > 0$ and the last inequality follows from Young's inequality. Substituting $\rho = 1$, we obtain

$$(\gamma/2 - \gamma^2 L_F) \|\nabla F(w_t)\|^2 \leq F(w_t) - F(w_{t+1}) + (\gamma/2 + \gamma^2 L_F) \|\Delta\|^2.$$

We now use Lemma 4 to obtain

$$\begin{aligned} \left(\frac{\gamma}{2} - \gamma^2 L_F\right) \|\nabla F(w_t)\|^2 &\leq F(w_t) - F(w_{t+1}) \\ &\quad + (\gamma/2 + \gamma^2 L_F) \left((1 + \lambda) \left(\frac{\sqrt{1 - \delta} + 2\alpha}{1 - \beta} \right)^2 \|\nabla F(w_t)\|^2 + \tilde{\epsilon}(\lambda) \right). \end{aligned}$$

with high probability. Upon further simplification, we have

$$\begin{aligned} & \left(\frac{\gamma}{2} - \frac{\gamma}{2}(1 + \lambda) \left(\frac{\sqrt{1 - \delta} + 2\alpha}{1 - \beta} \right)^2 - (1 + \lambda) \left(\frac{\sqrt{1 - \delta} + 2\alpha}{1 - \beta} \right)^2 \gamma^2 L_F - \gamma^2 L_F \right) \|\nabla F(w_t)\|^2 \\ & \leq F(w_t) - F(w_{t+1}) + (\gamma/2 + \gamma^2 L_F) \tilde{\epsilon}(\lambda). \end{aligned}$$

We now substitute $\gamma = \frac{c}{L_F}$, for a small enough constant c , so that we can ignore the contributions of the terms with quadratic dependence on γ . We substitute $\lambda = \lambda_0$ for a sufficiently small positive constant λ_0 . Provided $\delta > \delta_0 + 4\alpha - 9\alpha^2 + 4\alpha^3$, where $\delta_0 = \left(1 - \frac{(1-\beta)^2}{1+\lambda_0}\right)^2$, we have

$$\left(\frac{\gamma}{2} - \frac{\gamma}{2}(1 + \lambda) \left(\frac{\sqrt{1 - \delta} + 2\alpha}{1 - \beta} \right)^2 - (1 + \lambda) \left(\frac{\sqrt{1 - \delta} + 2\alpha}{1 - \beta} \right)^2 \gamma^2 L_F - \gamma^2 L_F \right) = \frac{c_1}{L_F},$$

where c_1 is a constant. With this choice, we obtain

$$\frac{1}{T+1} \sum_{t=0}^T \|\nabla F(w_t)\|^2 \leq C_1 \frac{L_F(F(w_0) - F(w^*))}{T+1} + C_2 \epsilon$$

where the first term is obtained from a telescopic sum and ϵ is defined in equation (3.4). Finally, we obtain

$$\min_{t=0, \dots, T} \|\nabla F(w_t)\|^2 \leq C_1 \frac{L_F(F(w_0) - F(w^*))}{T+1} + C_2 \epsilon$$

with probability greater than or equal to $1 - \frac{c_1(1-\alpha)md}{(1+nLD)^d} - \frac{c_2d}{(1+(1-\alpha)mnLD)^d}$, proving Theorem 3.

3.10.2 Proof of Theorem 4

The proof of convergence for Theorem 4 follows the same steps as Theorem 3. Recall that the quantity of interest is

$$\tilde{\Delta} = g(w_t) - \nabla F(w_t)$$

for which we prove bound in the following lemma.

Lemma 5. *For any $\lambda > 0$, we have,*

$$\|\tilde{\Delta}\|^2 \leq ((1 + \lambda) \left(\frac{(1 + \beta)\sqrt{1 - \delta} + 2\alpha}{1 - \beta} \right)^2 \|\nabla F(w_t)\|^2 + \tilde{\epsilon}(\lambda))$$

with probability greater than or equal to $1 - \frac{c_1(1-\alpha)md}{(1+nLD)^d} - \frac{c_2d}{(1+(1-\alpha)mnLD)^d}$, where

$$\tilde{\epsilon}(\lambda) = 2\left(1 + \frac{1}{\lambda}\right) \left(\left(\frac{(1 + \beta)\sqrt{1 - \delta} + \alpha + \beta}{1 - \beta} \right)^2 \epsilon_1^2 + \left(\frac{1 - \alpha}{1 - \beta} \right)^2 \epsilon_2^2 \right).$$

with ϵ_1 and ϵ_2 as defined in equation (3.2) and (3.3) respectively.

Taking the above lemma for granted, we proceed to prove Theorem 4. The proof of Lemma 5 is deferred to Section 3.10.6.

The proof parallels the proof of 3, except the fact that we use Lemma 5 to upper bound $\|\tilde{\Delta}\|^2$. Correspondingly, a little algebra shows that we require $\delta > \tilde{\delta}_0 + 4\alpha - 8\alpha^2 + 4\alpha^3$, where $\tilde{\delta}_0 = \left(1 - \frac{(1-\beta)^2}{(1+\beta)^2(1+\lambda_0)}\right)$, where λ_0 is a sufficiently small positive constant. With the above requirement, the proof follows the same steps as Theorem 3 and hence we omit the details here.

3.10.3 Proof of Lemma 4:

We require the following result to prove Lemma 4. In the following result, we show that for non-Byzantine worker machine i , the local gradient $\nabla F_i(w_t)$ is concentrated around the global gradient $\nabla F(w_t)$.

Lemma 6. *We have*

$$\max_{i \in \mathcal{M}} \|\nabla F_i(w_t) - \nabla F(w_t)\| \leq \epsilon_1$$

with probability exceeding $1 - \frac{2(1-\alpha)md}{(1+n\hat{L}D)^d}$, where ϵ_1 is defined in equation (3.2).

Furthermore, we have the following Lemma which implies that the average of local gradients $\nabla F_i(w_t)$ over non-Byzantine worker machines is close to its expectation $\nabla F(w_t)$.

Lemma 7. *We have*

$$\left\| \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla F_i(w_t) - \nabla F(w_t) \right\| \leq \epsilon_2.$$

with probability exceeding $1 - \frac{2(1-\alpha)md}{(1+n\hat{L}D)^d} - \frac{2d}{(1+(1-\alpha)mn\hat{L}D)^d}$, where ϵ_2 is defined in equation (3.3).

Recall the definition of Δ . Using triangle inequality, we obtain

$$\|\Delta\| \leq \underbrace{\left\| \frac{1}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} \mathcal{Q}(\nabla F_i(w_t)) - \frac{1}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} \nabla F_i(w_t) \right\|}_{T_1} + \underbrace{\left\| \frac{1}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} \nabla F_i(w_t) - \nabla F(w_t) \right\|}_{T_2}$$

We first control T_1 . Using the compression scheme (Definition 9), we obtain

$$\begin{aligned} T_1 &= \left\| \frac{1}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} \mathcal{Q}(\nabla F_i(w_t)) - \frac{1}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} \nabla F_i(w_t) \right\| \leq \frac{\sqrt{1-\delta}}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} \|\nabla F_i(w_t)\| \\ &\leq \frac{\sqrt{1-\delta}}{|\mathcal{U}_t|} \left[\sum_{i \in \mathcal{M}} \|\nabla F_i(w_t)\| - \sum_{i \in \mathcal{M} \cap \mathcal{T}_t} \|\nabla F_i(w_t)\| + \sum_{i \in \mathcal{B} \cap \mathcal{U}_t} \|\nabla F_i(w_t)\| \right] \\ &\leq \frac{\sqrt{1-\delta}}{|\mathcal{U}_t|} \left[\sum_{i \in \mathcal{M}} \|\nabla F_i(w_t)\| + \sum_{i \in \mathcal{B} \cap \mathcal{U}_t} \|\nabla F_i(w_t)\| \right] \end{aligned}$$

Since $\beta \geq \alpha$, we ensure that $\mathcal{M} \cap \mathcal{T}_t \neq \emptyset$. We have,

$$\begin{aligned}
 T_1 &\leq \frac{\sqrt{1-\delta}}{|\mathcal{U}_t|} \left[\sum_{i \in \mathcal{M}} \|\nabla F_i(w_t)\| + \alpha m \max_{i \in \mathcal{M}} \|\nabla F_i(w_t)\| \right] \\
 &\leq \underbrace{\frac{\sqrt{1-\delta}}{|\mathcal{U}_t|} \left[\sum_{i \in \mathcal{M}} \|\nabla F_i(w_t) - \nabla F(w_t)\| + \sum_{i \in \mathcal{M}} \|\nabla F(w_t)\| \right]}_{T_3} \\
 &\quad + \underbrace{\frac{\alpha m \sqrt{1-\delta}}{|\mathcal{U}_t|} \max_{i \in \mathcal{M}} [\|\nabla F_i(w_t) - \nabla F(w_t)\| + \|\nabla F(w_t)\|]}_{T_4}
 \end{aligned}$$

We now upper-bound T_3 . We have

$$\begin{aligned}
 T_3 &\leq \frac{\sqrt{1-\delta}|\mathcal{M}|}{|\mathcal{U}_t|} \max_{i \in \mathcal{M}} \|\nabla F_i(w_t) - \nabla F(w_t)\| + \frac{\sqrt{1-\delta}|\mathcal{M}|}{|\mathcal{U}_t|} \|\nabla F(w_t)\| \\
 &\leq \frac{\sqrt{1-\delta}(1-\alpha)}{(1-\beta)} \max_{i \in \mathcal{M}} \|\nabla F_i(w_t) - \nabla F(w_t)\| + \frac{\sqrt{1-\delta}(1-\alpha)}{(1-\beta)} \|\nabla F(w_t)\| \\
 &\leq \frac{\sqrt{1-\delta}(1-\alpha)}{(1-\beta)} \epsilon_1 + \frac{\sqrt{1-\delta}(1-\alpha)}{(1-\beta)} \|\nabla F(w_t)\|
 \end{aligned}$$

with probability exceeding $1 - \frac{2(1-\alpha)md}{(1+nLD)^d}$, where we use Lemma 6. Similarly, for T_4 , we have

$$T_4 \leq \frac{\sqrt{1-\delta}\alpha}{1-\beta} \epsilon_1 + \frac{\sqrt{1-\delta}\alpha}{1-\beta} \|\nabla F(w_t)\|.$$

We now control the terms in T_2 . We obtain the following:

$$\begin{aligned}
 T_2 &\leq \frac{1}{|\mathcal{U}_t|} \left\| \sum_{i \in \mathcal{U}_t} \nabla F_i(w_t) - \nabla F(w_t) \right\| \\
 &\leq \frac{1}{|\mathcal{U}_t|} \left\| \sum_{i \in \mathcal{M}} (\nabla F_i(w_t) - \nabla F(w_t)) - \sum_{i \in \mathcal{M} \cap \mathcal{T}_t} (\nabla F_i(w_t) - \nabla F(w_t)) \right\| \\
 &\quad + \sum_{i \in \mathcal{B} \cap \mathcal{T}_t} \|\nabla F_i(w_t) - \nabla F(w_t)\| \\
 &\leq \frac{1}{|\mathcal{U}_t|} \left\| \sum_{i \in \mathcal{M}} (\nabla F_i(w_t) - \nabla F(w_t)) \right\| + \frac{1}{|\mathcal{U}_t|} \left\| \sum_{i \in \mathcal{M} \cap \mathcal{T}_t} (\nabla F_i(w_t) - \nabla F(w_t)) \right\| \\
 &\quad + \frac{1}{|\mathcal{U}_t|} \left\| \sum_{i \in \mathcal{B} \cap \mathcal{T}_t} (\nabla F_i(w_t) - \nabla F(w_t)) \right\|.
 \end{aligned}$$

Using Lemma 7, we have

$$\frac{1}{|\mathcal{U}_t|} \left\| \sum_{i \in \mathcal{M}} (\nabla F_i(w_t) - \nabla F(w_t)) \right\| \leq \frac{1 - \alpha}{1 - \beta} \epsilon_2.$$

with probability exceeding $1 - \frac{2(1-\alpha)md}{(1+n\hat{L}D)^d} - \frac{2d}{(1+(1-\alpha)mn\hat{L}D)^d}$. Also, we obtain

$$\frac{1}{|\mathcal{U}_t|} \left\| \sum_{i \in \mathcal{M} \cap \mathcal{T}_t} (\nabla F_i(w_t) - \nabla F(w_t)) \right\| \leq \frac{\beta}{1 - \alpha} \max_{i \in \mathcal{M}} \|\nabla F_i(w_t) - \nabla F(w_t)\| \leq \frac{\beta}{1 - \alpha} \epsilon_1,$$

with probability at least $1 - \frac{2(1-\alpha)md}{(1+n\hat{L}D)^d}$, where the last inequality is derived from Lemma 6. Finally, for the Byzantine term, we have

$$\begin{aligned} \frac{1}{|\mathcal{U}_t|} \left\| \sum_{i \in \mathcal{B} \cap \mathcal{T}_t} (\nabla F_i(w_t) - \nabla F(w_t)) \right\| &\leq \frac{\alpha}{1 - \beta} \max_{i \in \mathcal{B} \cap \mathcal{T}_t} \|\nabla F_i(w_t)\| + \frac{\alpha}{1 - \beta} \|\nabla F(w_t)\| \\ &\leq \frac{\alpha}{1 - \beta} \max_{i \in \mathcal{M}} \|\nabla F_i(w_t)\| + \frac{\alpha}{1 - \beta} \|\nabla F(w_t)\| \\ &\leq \frac{\alpha}{1 - \beta} \max_{i \in \mathcal{M}} \|\nabla F_i(w_t) - \nabla F(w_t)\| + \frac{2\alpha}{1 - \beta} \|\nabla F(w_t)\| \\ &\leq \frac{\alpha}{1 - \beta} \epsilon_1 + \frac{2\alpha}{1 - \beta} \|\nabla F(w_t)\|, \end{aligned}$$

with high probability, where the last inequality follows from Lemma 6.

Combining all the terms of T_1 and T_2 , we obtain,

$$\|\Delta\| \leq \frac{\sqrt{1 - \delta} + 2\alpha}{1 - \beta} \|\nabla F(w_t)\| + \frac{\sqrt{1 - \delta} + \alpha + \beta}{1 - \beta} \epsilon_1 + \frac{1 - \alpha}{1 - \beta} \epsilon_2.$$

Now, using Young's inequality, for any $\lambda > 0$, we obtain

$$\|\Delta\|^2 \leq (1 + \lambda) \left(\frac{\sqrt{1 - \delta} + 2\alpha}{1 - \beta} \right)^2 \|\nabla F(w_t)\|^2 + \tilde{\epsilon}(\lambda)$$

where

$$\tilde{\epsilon}(\lambda) = 2\left(1 + \frac{1}{\lambda}\right) \left[\left(\frac{\sqrt{1 - \delta} + \alpha + \beta}{1 - \beta} \right)^2 \epsilon_1^2 + \left(\frac{1 - \alpha}{1 - \beta} \right)^2 \epsilon_2^2 \right].$$

3.10.4 Proof of Lemma 6:

For a fixed $i \in \mathcal{M}$, we first analyze the quantity $\|\nabla F_i(w_t) - \nabla F(w_t)\|$. Notice that i is non-Byzantine. Recall that machine i has n independent data points. We use the sub-exponential concentration to control this term. Let us rewrite the concentration inequality.

Univariate sub-exponential concentration: Suppose Y is univariate random variable with $\mathbb{E}Y = \mu$ and y_1, \dots, y_n are i.i.d draws of Y . Also, Y is v sub-exponential. From sub-exponential concentration (Hoeffding's inequality), we obtain

$$\mathbb{P}\left(\left|\frac{1}{n}\sum_{i=1}^n y_i - \mu\right| > t\right) \leq 2 \exp\left\{-n \min\left(\frac{t}{v}, \frac{t^2}{v^2}\right)\right\}.$$

We directly use this to the k -th partial derivative of F_i . Let $\partial_k f(w_t, z^{i,j})$ be the partial derivative of the loss function with respect to k -th coordinate on i -th machine with j -th data point. From Assumption 10, we obtain

$$\mathbb{P}\left(\left|\frac{1}{n}\sum_{j=1}^n \partial_k f(w_t, z^{i,j}) - \partial_k F(w_t)\right| \geq t\right) \leq 2 \exp\left\{-n \min\left(\frac{t}{v}, \frac{t^2}{v^2}\right)\right\}.$$

Since $\nabla F_i(w_t) = \frac{1}{n} \sum_{j=1}^n \nabla f(w_t, z^{i,j})$, denoting $\nabla F_i^{(k)}(w_t)$ as the k -th coordinate of $\nabla F_i(w_t)$, we have

$$|\nabla F_i^{(k)}(w_t) - \partial_k F(w_t)| \leq t$$

with probability at least $1 - 2 \exp\left\{-n \min\left(\frac{t}{v}, \frac{t^2}{v^2}\right)\right\}$.

This result holds for a particular w_t . To extend this for all $w \in \mathcal{W}$, we exploit the covering net argument and the Lipschitz continuity of the partial derivative of the loss function (Assumption 9). Let $\{w_1, \dots, w_N\}$ be a δ covering of \mathcal{W} . Since \mathcal{W} has diameter D , from Vershynin, we obtain $N \leq (1 + \frac{D}{\delta})^d$. Hence with probability at least

$$1 - 2Nd \exp\left\{-n \min\left(\frac{t}{v}, \frac{t^2}{v^2}\right)\right\},$$

we have

$$|\nabla F_i^{(k)}(w) - \partial_k F(w)| \leq t$$

for all $w \in \{w_1, \dots, w_N\}$ and $k \in [d]$. This implies

$$\|\nabla F_i(w_t) - \nabla F(w_t)\| \leq t\sqrt{d},$$

with probability greater than or equal to $1 - 2Nd \exp\left\{-n \min\left(\frac{t}{v}, \frac{t^2}{v^2}\right)\right\}$.

We now reason about $w \in \mathcal{W} \setminus \{w_1, \dots, w_N\}$ via Lipschitzness (Assumption 9). From the definition of δ cover, for any $w \in \mathcal{W}$, there exists w_ℓ , an element of the cover such that $\|w - w_\ell\| \leq \delta$. Hence, we obtain

$$|\nabla F_i^{(k)}(w) - \partial_k F(w)| \leq t + 2L_k \delta$$

for all $w \in \mathcal{W}$ and consequently

$$\|\nabla F_i(w_t) - \nabla F(w_t)\| \leq \sqrt{d}t + 2\delta\hat{L}$$

with probability at least $1 - 2Nd \exp\{-n \min(\frac{t}{v}, \frac{t^2}{v^2})\}$, where $\hat{L} = \sqrt{\sum_{k=1}^d L_k^2}$.

Choosing $\delta = \frac{1}{2n\hat{L}}$ and

$$t = v \max\left\{\frac{d}{n} \log(1 + 2n\hat{L}d), \sqrt{\frac{d}{n} \log(1 + 2n\hat{L}d)}\right\},$$

we obtain

$$\|\nabla F_i(w_t) - \nabla F(w_t)\| \leq v\sqrt{d} \left(\max\left\{\frac{d}{n} \log(1 + 2n\hat{L}d), \sqrt{\frac{d}{n} \log(1 + 2n\hat{L}d)}\right\} \right) + \frac{1}{n} = \epsilon_1, \quad (3.10)$$

with probability greater than $1 - \frac{d}{(1+n\hat{L}D)^d}$. Taking union bound on all non-Byzantine machines yields the theorem.

3.10.5 Proof of Lemma 7

We need to upper bound the following quantity:

$$\left\| \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} (\nabla F_i(w_t) - \nabla F(w_t)) \right\|$$

We now use similar argument (sub-exponential concentration) like Lemma 6. The only difference is that in this case, we also consider *averaging* over worker nodes. We obtain the following:

$$\left\| \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} (\nabla F_i(w_t) - \nabla F(w_t)) \right\| \leq \epsilon_2$$

where

$$\epsilon_2 = v\sqrt{d} \left(\max\left\{\frac{d}{(1-\alpha)mn} \log(1 + 2(1-\alpha)mn\hat{L}d), \sqrt{\frac{d}{(1-\alpha)mn} \log(1 + 2(1-\alpha)mn\hat{L}d)}\right\}\right),$$

with probability $1 - \frac{2d}{(1+(1-\alpha)mn\hat{L}D)^d}$.

3.10.6 Proof of Lemma 5

Here we prove an upper bound on the norm of

$$\tilde{\Delta} = g(w_t) - \nabla F(w_t)$$

where $g(w_t) = \frac{1}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} Q(\nabla F_i(w_t))$.

We have

$$\begin{aligned} \|\tilde{\Delta}\| &= \left\| \frac{1}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} Q(\nabla F_i(w_t)) - \nabla F(w_t) \right\| \\ &= \frac{1}{|\mathcal{U}_t|} \left\| \sum_{i \in \mathcal{M}} [Q(\nabla F_i(w_t)) - \nabla F(w_t)] - \sum_{i \in (\mathcal{M} \cap \mathcal{T}_t)} [Q(\nabla F_i(w_t)) - \nabla F(w_t)] \right. \\ &\quad \left. + \sum_{i \in (\mathcal{B} \cap \mathcal{U}_t)} [Q(\nabla F_i(w_t)) - \nabla F(w_t)] \right\| \\ &\leq \frac{1}{|\mathcal{U}_t|} \left(\underbrace{\left\| \sum_{i \in \mathcal{M}} Q(\nabla F_i(w_t)) - \nabla F(w_t) \right\|}_{T_1} + \underbrace{\left\| \sum_{i \in (\mathcal{M} \cap \mathcal{T}_t)} Q(\nabla F_i(w_t)) - \nabla F(w_t) \right\|}_{T_2} \right. \\ &\quad \left. + \underbrace{\left\| \sum_{i \in (\mathcal{B} \cap \mathcal{U}_t)} Q(\nabla F_i(w_t)) - \nabla F(w_t) \right\|}_{T_3} \right) \end{aligned}$$

Now we bound each term separately. For the first term, we have

$$\begin{aligned} \frac{1}{|\mathcal{U}_t|} T_1 &= \frac{1}{|\mathcal{U}_t|} \left\| \sum_{i \in \mathcal{M}} Q(\nabla F_i(w_t)) - \nabla F(w_t) \right\| \\ &= \frac{1}{|\mathcal{U}_t|} \left\| \sum_{i \in \mathcal{M}} Q(\nabla F_i(w_t)) - \nabla F_i(w_t) \right\| + \frac{1}{|\mathcal{U}_t|} \left\| \sum_{i \in \mathcal{M}} \nabla F_i(w_t) - \nabla F(w_t) \right\| \\ &\leq \frac{1}{|\mathcal{U}_t|} \sum_{i \in \mathcal{M}} \left(\left\| Q(\nabla F_i(w_t)) - \nabla F_i(w_t) \right\| \right) + \frac{1 - \alpha}{1 - \beta} \epsilon_2 \\ &\leq \frac{1}{|\mathcal{U}_t|} \sum_{i \in \mathcal{M}} \left(\sqrt{1 - \delta} \|\nabla F_i(w_t)\| \right) + \frac{1 - \alpha}{1 - \beta} \epsilon_2 \\ &\leq \frac{\sqrt{1 - \delta}}{|\mathcal{U}_t|} \sum_{i \in \mathcal{M}} \left(\|\nabla F(w_t)\| + \|\nabla F_i(w_t) - \nabla F(w_t)\| \right) + \frac{1 - \alpha}{1 - \beta} \epsilon_2 \\ &\leq \frac{\sqrt{1 - \delta}(1 - \alpha)}{1 - \beta} \|\nabla F(w_t)\| + \frac{\sqrt{1 - \delta}(1 - \alpha)}{1 - \beta} \epsilon_1 + \frac{1 - \alpha}{1 - \beta} \epsilon_2 \end{aligned}$$

where we use the definition of a δ -approximate compressor, Lemma 6 and Lemma 7. Similarly, we can bound T_2 as

$$\begin{aligned}
 T_2 &\leq \sum_{i \in (\mathcal{M} \cap \mathcal{T}_t)} \|Q(\nabla F_i(w_t)) - \nabla F(w_t)\| \\
 &\leq \beta m \max_{i \in \mathcal{M}} \|Q(\nabla F_i(w_t)) - \nabla F(w_t)\| \\
 &\leq \beta m \max_{i \in \mathcal{M}} \left(\sqrt{1 - \delta} \|\nabla F_i(w_t)\| + \|\nabla F_i(w_t) - \nabla F(w_t)\| \right) \\
 &\leq \beta m \max_{i \in \mathcal{M}} \left(\sqrt{1 - \delta} \|\nabla F(w_t)\| + (1 + \sqrt{1 - \delta}) \|\nabla F_i(w_t) - \nabla F(w_t)\| \right)
 \end{aligned}$$

where we use the definition of δ -approximate compressor. Hence invoking Lemma 6, we obtain

$$\frac{1}{|\mathcal{U}_t|} T_2 \leq \frac{\beta \sqrt{1 - \delta}}{1 - \beta} \|\nabla F(w_t)\| + \frac{\beta(1 + \sqrt{1 - \delta})}{1 - \beta} \epsilon_1$$

Also, owing to the trimming with $\beta > \alpha$, we have at least one good machine in the set \mathcal{T}_t for all t . Now each term in the set $\mathcal{B} \cap \mathcal{U}_t$, we have

$$\begin{aligned}
 T_3 &= \sum_{i \in (\mathcal{B} \cap \mathcal{U}_t)} \|Q(\nabla F_i(w_t)) - \nabla F(w_t)\| \\
 &\leq \alpha m (\max_{i \in \mathcal{M}} \|Q(\nabla F_i(w_t))\| + \|\nabla F(w_t)\|) \\
 &\leq \alpha m (\max_{i \in \mathcal{M}} \sqrt{1 - \delta} \|\nabla F_i(w_t)\| + \|\nabla F_i(w_t)\| + \|\nabla F(w_t)\|) \\
 &\leq \alpha m \left((1 + \sqrt{1 - \delta}) \epsilon_1 + (2 + \sqrt{1 - \delta}) \|\nabla F(w_t)\| \right) \\
 \frac{1}{|\mathcal{U}_t|} T_3 &\leq \frac{\alpha(2 + \sqrt{1 - \delta})}{1 - \beta} \|\nabla F(w_t)\| + \frac{\alpha(1 + \sqrt{1 - \delta})}{1 - \beta} \epsilon_1
 \end{aligned}$$

where we use Lemma 6. Putting T_1, T_2, T_3 we get

$$\begin{aligned}
 \|\tilde{\Delta}\| &\leq \left(\frac{\sqrt{1 - \delta}(1 - \alpha)}{1 - \beta} + \frac{\beta \sqrt{1 - \delta}}{1 - \beta} + \frac{\alpha(2 + \sqrt{1 - \delta})}{1 - \beta} \right) \|\nabla F(w_t)\| \\
 &\quad + \left(\frac{\sqrt{1 - \delta}(1 - \alpha)}{1 - \beta} + \frac{\beta(1 + \sqrt{1 - \delta})}{1 - \beta} + \frac{\alpha(1 + \sqrt{1 - \delta})}{1 - \beta} \right) \epsilon_1 + \frac{1 - \alpha}{1 - \beta} \epsilon_2 \\
 &= \left(\frac{(1 + \beta)\sqrt{1 - \delta} + 2\alpha}{1 - \beta} \right) \|\nabla F(w_t)\| + \left(\frac{(1 + \beta)\sqrt{1 - \delta} + \alpha + \beta}{1 - \beta} \right) \epsilon_1 + \frac{1 - \alpha}{1 - \beta} \epsilon_2 \\
 \|\tilde{\Delta}\|^2 &\leq (1 + \lambda) \left(\frac{(1 + \beta)\sqrt{1 - \delta} + 2\alpha}{1 - \beta} \right)^2 \|\nabla F(w_t)\|^2 + \tilde{\epsilon}(\lambda)
 \end{aligned}$$

where $\tilde{\epsilon}(\lambda) = 2(1 + \frac{1}{\lambda}) \left(\left(\frac{(1 + \beta)\sqrt{1 - \delta} + \alpha + \beta}{1 - \beta} \right)^2 \epsilon_1^2 + \left(\frac{1 - \alpha}{1 - \beta} \right)^2 \epsilon_2^2 \right)$. Hence, the lemma follows.

3.11 Proof of Theorem 5

We first define an auxiliary sequence defined as:

$$\tilde{w}_t = w_t - \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} e_i(t)$$

Hence, we obtain

$$\tilde{w}_{t+1} = w_{t+1} - \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} e_i(t+1).$$

For notational simplicity, let us drop the subscript t from \mathcal{U}_t and \mathcal{T}_t and denote them as \mathcal{U} and \mathcal{T} .

Since (we will ensure that the iterates remain in the parameter space and hence we can ignore the projection step),

$$w_{t+1} = w_t - \frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{U}} p_i(w_t),$$

we get

$$\begin{aligned} \tilde{w}_{t+1} &= w_t - \frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{U}} \mathcal{C}(p_i(w_t)) - \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} e_i(t+1) \\ &= w_t - \frac{1}{|\mathcal{U}|} \left(\sum_{i \in \mathcal{M}} \mathcal{C}(p_i(w_t)) + \sum_{i \in \mathcal{B} \cap \mathcal{U}} \mathcal{C}(p_i(w_t)) - \sum_{i \in \mathcal{M} \cap \mathcal{T}} \mathcal{C}(p_i(w_t)) \right) - \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} e_i(t+1) \\ &= w_t - \left(\frac{1-\alpha}{1-\beta} \right) \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \mathcal{C}(p_i(w_t)) - \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} e_i(t+1) \\ &\quad - \frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{B} \cap \mathcal{U}} \mathcal{C}(p_i(w_t)) + \frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{M} \cap \mathcal{T}} \mathcal{C}(p_i(w_t)) \end{aligned}$$

Since $\mathcal{C}(p_i(w_t)) + e_i(t+1) = p_i(w_t)$ for all $i \in \mathcal{M}$, we obtain

$$\begin{aligned} \left(\frac{1-\alpha}{1-\beta} \right) \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \mathcal{C}(p_i(w_t)) + \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} e_i(t+1) &= \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} p_i(w_t) \\ &\quad + \frac{\beta-\alpha}{1-\beta} \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \mathcal{C}(p_i(w_t)) \end{aligned}$$

Let us denote $T_1 = \frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{B} \cap \mathcal{U}} \mathcal{C}(p_i(w_t))$, $T_2 = \frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{M} \cap \mathcal{T}} \mathcal{C}(p_i(w_t))$

and $T_3 = \frac{\beta-\alpha}{1-\beta} \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \mathcal{C}(p_i(w_t))$. With this, we obtain

$$\begin{aligned} \tilde{w}_{t+1} &= w_t - \frac{1}{|\mathcal{M}|} p_i(w_t) - T_1 + T_2 - T_3 \\ &= \tilde{w}_t + \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} e_i(t) - \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} p_i(w_t) - \tilde{T} \\ &= \tilde{w}_t - \gamma \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla F_i(w_t) - \tilde{T} \end{aligned}$$

where $\tilde{T} = T_1 - T_2 + T_3$. Observe that the auxiliary sequence looks similar to a distributed gradient step with a presence of \tilde{T} . For the convergence analysis, we will use this relation along with an upper bound on $\|\tilde{T}\|$.

Using this auxiliary sequence, we first ensure that the iterates of our algorithm remains close to one another. To that end, we have

$$\begin{aligned} w_{t+1} - w_t &= \tilde{w}_{t+1} - \tilde{w}_t + \frac{1}{|\mathcal{M}|} e_i(t+1) - \frac{1}{|\mathcal{M}|} e_i(t) \\ &= -\gamma \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla F_i(w_t) - \tilde{T} + \frac{1}{|\mathcal{M}|} e_i(t+1) - \frac{1}{|\mathcal{M}|} e_i(t). \end{aligned}$$

Hence, we obtain

$$\begin{aligned} \|w_{t+1} - w_t\| &\leq \left\| \gamma \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla F_i(w_t) \right\| + \|\tilde{T}\| + \left\| \frac{1}{|\mathcal{M}|} e_i(t+1) \right\| + \left\| \frac{1}{|\mathcal{M}|} e_i(t) \right\| \\ &\leq \gamma \left\| \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla F_i(w_t) - \nabla F(w_t) \right\| + \gamma \|\nabla F(w_t)\| + \|\tilde{T}\| + \left\| \frac{1}{|\mathcal{M}|} e_i(t+1) \right\| \\ &\quad + \left\| \frac{1}{|\mathcal{M}|} e_i(t) \right\| \\ &\leq \gamma \epsilon_2 + \gamma \|\nabla F(w_t)\| + \|\tilde{T}\| + \left\| \frac{1}{|\mathcal{M}|} e_i(t+1) \right\| + \left\| \frac{1}{|\mathcal{M}|} e_i(t) \right\|. \end{aligned}$$

Now, using Lemma 8 and Lemma 9 in conjunction with Assumption 11 ensures the iterates of Algorithm 3 stays in the parameter space \mathcal{W} .

We assume that the global loss function $F(\cdot)$ is L_F smooth. We get

$$F(\tilde{w}_{t+1}) \leq F(\tilde{w}_t) + \langle \nabla F(\tilde{w}_t), \tilde{w}_{t+1} - \tilde{w}_t \rangle + \frac{L_F}{2} \|\tilde{w}_{t+1} - \tilde{w}_t\|^2.$$

Now, we use the above recursive equation

$$\tilde{w}_{t+1} = \tilde{w}_t - \gamma \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla F_i(w_t) - \tilde{T}.$$

Substituting, we obtain

$$\begin{aligned}
 F(\tilde{w}_{t+1}) &\leq F(\tilde{w}_t) - \gamma \langle \nabla F(\tilde{w}_t), \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla F_i(w_t) \rangle - \langle \nabla F(\tilde{w}_t), \tilde{T} \rangle \\
 &\quad + \frac{L_F}{2} \left\| \frac{\gamma}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla F_i(w_t) + \tilde{T} \right\|^2 \\
 &\leq F(\tilde{w}_t) - \gamma \langle \nabla F(\tilde{w}_t), \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla F_i(w_t) \rangle - \langle \nabla F(\tilde{w}_t), \tilde{T} \rangle \\
 &\quad + L_F \gamma^2 \left\| \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla F_i(w_t) \right\|^2 + L_F \|\tilde{T}\|^2
 \end{aligned} \tag{3.11}$$

In the subsequent calculation, we use the following definition of smoothness:

$$\|\nabla F(y_1) - \nabla F(y_2)\| \leq L_F \|y_1 - y_2\|$$

for all y_1 and $y_2 \in \mathbb{R}^d$.

Rewriting the right hand side (R.H.S) of equation (3.11), we obtain

$$\begin{aligned}
 R.H.S &= \underbrace{F(\tilde{w}_t) - \gamma \langle \nabla F(\tilde{w}_t), \nabla F(w_t) \rangle}_{\text{Term-I}} + \underbrace{\gamma \langle \nabla F(\tilde{w}_t), \nabla F(w_t) - \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla F_i(w_t) \rangle}_{\text{Term-II}} \\
 &\quad + \underbrace{\langle \nabla F(w_t), -\tilde{T} \rangle + \langle \nabla F(\tilde{w}_t) - \nabla F(w_t), -\tilde{T} \rangle}_{\text{Term-III}} \\
 &\quad + \underbrace{2L_F \gamma^2 \left\| \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla F_i(w_t) - \nabla F(w_t) \right\|^2 + 2L_F \gamma^2 \|\nabla F(w_t)\|^2 + L_F \|\tilde{T}\|^2}_{\text{Term-IV}}.
 \end{aligned}$$

We now control the 4 terms separately. We start with Term-I.

Control of Term-I: We obtain

$$\begin{aligned}
 \text{Term-I} &= F(\tilde{w}_t) - \gamma \langle \nabla F(w_t), \nabla F(w_t) \rangle - \gamma \langle \nabla F(\tilde{w}_t) - \nabla F(w_t), \nabla F(w_t) \rangle \\
 &\leq F(\tilde{w}_t) - \gamma \|\nabla F(w_t)\|^2 + 25\gamma \|\nabla F(\tilde{w}_t) - \nabla F(w_t)\|^2 + \frac{\gamma}{100} \|\nabla F(w_t)\|^2,
 \end{aligned}$$

where we use Young's inequality ($\langle a, b \rangle \leq \frac{\rho}{2} \|a\|^2 + \frac{1}{2\rho} \|b\|^2$ with $\rho = 50$) in the last inequality. Using the smoothness of $F(\cdot)$, we obtain

$$\text{Term-I} \leq F(\tilde{w}_t) - \gamma \|\nabla F(w_t)\|^2 + \frac{\gamma}{100} \|\nabla F(w_t)\|^2 + 25\gamma L_F^2 \left\| \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} e_i(t) \right\|^2. \tag{3.12}$$

Control of Term-II: Similarly, for Term-II, we have

$$\begin{aligned} \text{Term-II} &= \gamma \langle \nabla F(\tilde{w}_t), \nabla F(w_t) - \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla F_i(w_t) \rangle \leq 50\gamma\epsilon_2^2 + \frac{\gamma}{200} \|\nabla F(\tilde{w}_t)\|^2 \\ &\leq 50\gamma\epsilon_2^2 + \frac{\gamma}{100} \|\nabla F(w_t)\|^2 + \frac{\gamma L_F^2}{100} \left\| \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} e_i(t) \right\|^2. \end{aligned} \quad (3.13)$$

Control of Term-III: We obtain

$$\begin{aligned} \text{Term-III} &= \langle \nabla F(w_t), -\tilde{T} \rangle + \langle \nabla F(\tilde{w}_t) - \nabla F(w_t), -\tilde{T} \rangle \\ &\leq \frac{\gamma}{2} \|\nabla F(w_t)\|^2 + \frac{1}{2\gamma} \|\tilde{T}\|^2 + \frac{L_F^2}{2} \left\| \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} e_i(t) \right\|^2 + \frac{1}{2} \|\tilde{T}\|^2. \end{aligned} \quad (3.14)$$

Control of Term-IV:

$$\begin{aligned} \text{Term-IV} &= 2L_F\gamma^2 \left\| \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla F_i(w_t) - \nabla F(w_t) \right\|^2 + 2L_F\gamma^2 \|\nabla F(w_t)\|^2 + L_F \|\tilde{T}\|^2 \\ &\leq 2L_F\gamma^2\epsilon_2^2 + 2L_F\gamma^2 \|\nabla F(w_t)\|^2 + L_F \|\tilde{T}\|^2 \end{aligned} \quad (3.15)$$

Combining all 4 terms, we obtain

$$\begin{aligned} F(\tilde{w}_{t+1}) &\leq F(\tilde{w}_t) - \left(\frac{\gamma}{2} - \frac{\gamma}{50} - 2L_F\gamma^2 \right) \|\nabla F(w_t)\|^2 \\ &\quad + \left(25\gamma L_F^2 + \frac{\gamma L_F^2}{100} + \frac{L_F^2}{2} \right) \left\| \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} e_i(t) \right\|^2 \\ &\quad + 50\gamma\epsilon_2^2 + 2L_F\gamma^2\epsilon_2^2 + \left(\frac{1}{2\gamma} + \frac{1}{2} + L_F \right) \|\tilde{T}\|^2 \end{aligned} \quad (3.16)$$

We now control the error sequence and $\|\tilde{T}\|^2$. These will be separate lemmas, but here we write it as a whole.

Control of error sequence:

Lemma 8. For all $i \in \mathcal{M}$, we have

$$\|e_i(t)\|^2 \leq \frac{3(1-\delta)}{\delta} \gamma^2 \sigma^2$$

for all $t \geq 0$.

Proof. For machine $i \in \mathcal{M}$, we have

$$\|e_i(t+1)\|^2 = \|\mathcal{C}(p_i(w_t)) - p_i(w_t)\|^2 \leq (1-\delta) \|p_i(w_t)\|^2 = (1-\delta) \|\gamma \nabla F_i(w_t) + e_i(t)\|^2$$

Using technique similar to the proof of [20, Lemma 3] and using $\|\nabla F_i(w_t)\|^2 \leq \sigma^2$, we obtain

$$\|e_i(t+1)\|^2 \leq \frac{2(1-\delta)(1+1/\eta)}{\delta} \gamma^2 \sigma^2$$

where $\eta > 0$. Substituting $\eta = 2$ implies

$$\|e_i(t+1)\|^2 \leq \frac{3(1-\delta)}{\delta} \gamma^2 \sigma^2 \tag{3.17}$$

for all $i \in \mathcal{M}$. This also implies

$$\max_{i \in \mathcal{M}} \|e_i(t+1)\|^2 \leq \frac{3(1-\delta)}{\delta} \gamma^2 \sigma^2.$$

□

Control of $\|\tilde{T}\|^2$:

Lemma 9. *We obtain*

$$\|\tilde{T}\|^2 \leq \frac{9(1 + \sqrt{1-\delta})^2 \gamma^2}{(1-\beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \left(\epsilon_1^2 + \|\nabla F(w_t)\|^2 + \frac{3(1-\delta)}{\delta} \sigma^2 \right)$$

with probability exceeding $1 - \frac{2(1-\alpha)md}{(1+nLD)^d}$.

Proof. We have

$$\|\tilde{T}\| = \|T_1 - T_2 + T_3\| \leq \|T_1\| + \|T_2\| + \|T_3\|.$$

We control these 3 terms separately. We obtain

$$\|T_1\| = \left\| \frac{1}{|\mathcal{U}|} \sum_{i \in B \cap \mathcal{U}} \mathcal{C}(p_i(w_t)) \right\| \leq \frac{1}{(1-\beta)m} \sum_{i \in B \cap \mathcal{U}} \|\mathcal{C}(p_i(w_t))\|.$$

Since the worker machines are sorted according to $\|\mathcal{C}(p_i(w_t))\|$ (the central machine only gets to

see $\mathcal{C}(p_i(w_t))$, and so the most natural metric to sort is $\|\mathcal{C}(p_i(w_t))\|$, we obtain

$$\begin{aligned}
 \|T_1\| &\leq \frac{\alpha m}{(1-\beta)m} \max_{i \in \mathcal{M}} \|\mathcal{C}(p_i(w_t))\| \\
 &\leq (1 + \sqrt{1-\delta}) \frac{\alpha m}{(1-\beta)m} \max_{i \in \mathcal{M}} \|p_i(w_t)\| \\
 &\leq (1 + \sqrt{1-\delta}) \frac{\alpha m}{(1-\beta)m} \max_{i \in \mathcal{M}} \|\gamma \nabla F_i(w_t) + e_i(t)\| \\
 &\leq (1 + \sqrt{1-\delta}) \frac{\alpha}{(1-\beta)} \gamma \max_{i \in \mathcal{M}} \|\nabla F_i(w_t) - \nabla F(w_t)\| \\
 &\quad + (1 + \sqrt{1-\delta}) \frac{\alpha}{(1-\beta)} \gamma \|\nabla F(w_t)\| + (1 + \sqrt{1-\delta}) \frac{\alpha}{(1-\beta)} \max_{i \in \mathcal{M}} \|e_i(t)\| \\
 &\leq (1 + \sqrt{1-\delta}) \frac{\alpha \gamma \epsilon_1}{(1-\beta)} + (1 + \sqrt{1-\delta}) \frac{\alpha \gamma}{(1-\beta)} \|\nabla F(w_t)\| \\
 &\quad + (1 + \sqrt{1-\delta}) \frac{\alpha \gamma \sigma}{(1-\beta)} \sqrt{\frac{3(1-\delta)}{\delta}}.
 \end{aligned}$$

Hence,

$$\|T_1\|^2 \leq 3 \frac{(1 + \sqrt{1-\delta})^2}{(1-\beta)^2} \alpha^2 \gamma^2 \left(\epsilon_1^2 + \|\nabla F(w_t)\|^2 + \frac{3(1-\delta)}{\delta} \sigma^2 \right).$$

Similarly, we obtain,

$$\|T_2\|^2 \leq 3 \frac{(1 + \sqrt{1-\delta})^2}{(1-\beta)^2} \beta^2 \gamma^2 \left(\epsilon_1^2 + \|\nabla F(w_t)\|^2 + \frac{3(1-\delta)}{\delta} \sigma^2 \right).$$

For T_3 , we have

$$\begin{aligned}
 \|T_3\| &= \frac{\beta - \alpha}{1 - \beta} \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \|\mathcal{C}(p_i(w_t))\| \leq \frac{\beta - \alpha}{1 - \beta} \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} (1 + \sqrt{1-\delta}) \|p_i(w_t)\| \\
 &\leq (1 + \sqrt{1-\delta}) \frac{\beta - \alpha}{1 - \beta} \max_{i \in \mathcal{M}} \|p_i(w_t)\|
 \end{aligned}$$

Using the previous calculation, we obtain

$$\begin{aligned}
 \|T_3\| &\leq (1 + \sqrt{1-\delta}) \frac{(\beta - \alpha) \gamma \epsilon_1}{(1-\beta)} + (1 + \sqrt{1-\delta}) \frac{(\beta - \alpha) \gamma}{(1-\beta)} \|\nabla F(w_t)\| \\
 &\quad + (1 + \sqrt{1-\delta}) \frac{(\beta - \alpha) \gamma \sigma}{(1-\beta)} \sqrt{\frac{3(1-\delta)}{\delta}},
 \end{aligned}$$

and as a result,

$$\|T_3\|^2 \leq 3 \frac{(1 + \sqrt{1-\delta})^2}{(1-\beta)^2} (\beta - \alpha)^2 \gamma^2 \left(\epsilon_1^2 + \|\nabla F(w_t)\|^2 + \frac{3(1-\delta)}{\delta} \sigma^2 \right).$$

Combining the above 3 terms, we obtain

$$\begin{aligned} \|\tilde{T}\|^2 &\leq 3\|T_1\|^2 + 3\|T_2\|^2 + 3\|T_3\|^2 \\ &\leq \frac{9(1 + \sqrt{1 - \delta})^2 \gamma^2}{(1 - \beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \left(\epsilon_1^2 + \|\nabla F(w_t)\|^2 + \frac{3(1 - \delta)}{\delta} \sigma^2 \right). \end{aligned}$$

□

Back to the convergence of $F(\cdot)$: We use the above bound on $\|\tilde{T}\|^2$ and Lemma 8 to conclude the proof of the main convergence result. Recall equation (3.16):

$$\begin{aligned} F(\tilde{w}_{t+1}) &\leq F(\tilde{w}_t) - \left(\frac{\gamma}{2} - \frac{\gamma}{50} - 2L_F\gamma^2 \right) \|\nabla F(w_t)\|^2 \\ &\quad + \left(25\gamma L_F^2 + \frac{\gamma L_F^2}{100} + \frac{L_F^2}{2} \right) \left\| \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} e_i(t) \right\|^2 \\ &\quad + 50\gamma\epsilon_2^2 + 2L_F\gamma^2\epsilon_2^2 + \left(\frac{1}{2\gamma} + \frac{1}{2} + L_F \right) \|\tilde{T}\|^2 \end{aligned} \tag{3.18}$$

First, let us compute the term associated with the error sequence. Note that (from Cauchy-Schwartz inequality)

$$\left\| \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} e_i(t) \right\|^2 \leq \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \|e_i(t)\|^2,$$

and from equation (3.17), we obtain

$$\left\| \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} e_i(t) \right\|^2 \leq \frac{3(1 - \delta)}{\delta} \gamma^2 \sigma^2,$$

and so the error term is upper bounded by

$$\left(\frac{\gamma^2 L_F^2}{2} + \frac{\gamma^3 L_F^2}{100} + 25\gamma^3 L_F^2 \right) \frac{3(1 - \delta)\sigma^2}{\delta}.$$

We now substitute the expression for $\|\tilde{T}\|^2$. We obtain

$$\left(\frac{1}{2\gamma} + \frac{1}{2} + L_F \right) \|\tilde{T}\|^2 = \frac{1}{2\gamma} \|\tilde{T}\|^2 + \left(\frac{1}{2} + L_F \right) \|\tilde{T}\|^2.$$

The first term in the above equation is

$$\begin{aligned} \frac{1}{2\gamma} \|\tilde{T}\|^2 &\leq \frac{9\gamma(1 + \sqrt{1 - \delta})^2}{2(1 - \beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \left(\epsilon_1^2 + \|\nabla F(w_t)\|^2 + \frac{3(1 - \delta)}{\delta} \sigma^2 \right) \\ &\leq \frac{9\gamma(1 + \sqrt{1 - \delta})^2}{2(1 - \beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \|\nabla F(w_t)\|^2 \\ &\quad + \frac{9\gamma(1 + \sqrt{1 - \delta})^2}{2(1 - \beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \left(\epsilon_1^2 + \frac{3(1 - \delta)}{\delta} \sigma^2 \right), \end{aligned}$$

and the second term is

$$\begin{aligned}
 \left(\frac{1}{2} + L_F\right) \|\tilde{T}\|^2 &\leq \left(\frac{1}{2} + L_F\right) \frac{9\gamma^2(1 + \sqrt{1 - \delta})^2}{(1 - \beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \\
 &\quad \times \left(\epsilon_1^2 + \|\nabla F(w_t)\|^2 + \frac{3(1 - \delta)}{\delta} \sigma^2\right) \\
 &\leq \left(\frac{1}{2} + L_F\right) \frac{9\gamma^2(1 + \sqrt{1 - \delta})^2}{(1 - \beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \|\nabla F(w_t)\|^2 \\
 &\quad + \left(\frac{1}{2} + L_F\right) \frac{9\gamma^2(1 + \sqrt{1 - \delta})^2}{(1 - \beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \left(\epsilon_1^2 + \frac{3(1 - \delta)}{\delta} \sigma^2\right)
 \end{aligned}$$

Collecting all the above terms, the coefficient of $-\gamma\|\nabla F(w_t)\|^2$ is given by

$$\begin{aligned}
 &\frac{1}{2} - \frac{1}{50} - 2L_F\gamma - \frac{9(1 + \sqrt{1 - \delta})^2}{2(1 - \beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \\
 &\quad - \left(\frac{1}{2} + L_F\right) \frac{9\gamma(1 + \sqrt{1 - \delta})^2}{(1 - \beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2].
 \end{aligned}$$

Provided we select a sufficiently small γ , a little algebra shows that if

$$\frac{9(1 + \sqrt{1 - \delta})^2}{2(1 - \beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] < \left(\frac{1}{2} - \frac{1}{50}\right),$$

the coefficient of $\|\nabla F(w_t)\|^2$ becomes $-c\gamma$, where $c > 0$ is a universal constant. Considering the other terms and rewriting equation (3.16), we obtain

$$\begin{aligned}
 &F(\tilde{w}_{t+1}) \\
 &\leq F(\tilde{w}_t) - c\gamma\|\nabla F(w_t)\|^2 + \left(\frac{\gamma^2 L_F^2}{2} + \frac{\gamma^3 L_F^2}{100} + 25\gamma^3 L_F^2\right) \frac{3(1 - \delta)\sigma^2}{\delta} + 50\gamma\epsilon_2^2 \\
 &\quad + 2L_F\gamma^2\epsilon_2^2 + \frac{9\gamma(1 + \sqrt{1 - \delta})^2}{2(1 - \beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \left(\epsilon_1^2 + \frac{3(1 - \delta)}{\delta} \sigma^2\right) \\
 &\quad + \left(\frac{1}{2} + L_F\right) \frac{9\gamma^2(1 + \sqrt{1 - \delta})^2}{(1 - \beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \left(\epsilon_1^2 + \frac{3(1 - \delta)}{\delta} \sigma^2\right).
 \end{aligned}$$

Continuing, we get

$$\begin{aligned}
& \frac{1}{T+1} \sum_{t=0}^T \|\nabla F(w_t)\|^2 \\
& \leq \frac{1}{c\gamma(T+1)} \sum_{t=0}^T (F(\tilde{w}_t) - F(\tilde{w}_{t+1})) \\
& \quad + \left(\frac{\gamma L_F^2}{2} + \frac{\gamma^2 L_F^2}{100} + 25\gamma^2 L_F^2 \right) \frac{3(1-\delta)\sigma^2}{c\delta} + \frac{50}{c} \epsilon_2^2 \\
& \quad + \frac{2L_F\gamma\epsilon_2^2}{c} + \frac{9(1+\sqrt{1-\delta})^2}{2c(1-\beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \left(\epsilon_1^2 + \frac{3(1-\delta)}{\delta} \sigma^2 \right) \\
& \quad + \left(\frac{1}{2} + L_F \right) \frac{9\gamma(1+\sqrt{1-\delta})^2}{c(1-\beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \left(\epsilon_1^2 + \frac{3(1-\delta)}{\delta} \sigma^2 \right).
\end{aligned}$$

Using the telescoping sum, we obtain

$$\begin{aligned}
\min_{t=0,\dots,T} \|\nabla F(w_t)\|^2 & \leq \frac{F(w_0) - F^*}{c\gamma(T+1)} \\
& \quad + \left[\frac{9(1+\sqrt{1-\delta})^2}{2c(1-\beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \left(\epsilon_1^2 + \frac{3(1-\delta)}{\delta} \sigma^2 \right) + \frac{50}{c} \epsilon_2^2 \right] \\
& \quad + \gamma \left[\frac{L_F^2}{2} \frac{3(1-\delta)\sigma^2}{c\delta} + \frac{2L_F\epsilon_2^2}{c} \right. \\
& \quad \left. + \left(\frac{1}{2} + L_F \right) \frac{9(1+\sqrt{1-\delta})^2}{c(1-\beta)^2} [\alpha^2 + \beta^2 + (\beta - \alpha)^2] \left(\epsilon_1^2 + \frac{3(1-\delta)}{\delta} \sigma^2 \right) \right] \\
& \quad + \gamma^2 \left[\left(\frac{L_F^2}{100} + 25L_F^2 \right) \frac{3(1-\delta)\sigma^2}{c\delta} \right]
\end{aligned}$$

Simplifying the above expression, we write

$$\min_{t=0,\dots,T} \|\nabla F(w_t)\|^2 \leq \frac{F(w_0) - F^*}{c\gamma(T+1)} + \Delta_1 + \gamma\Delta_2 + \gamma^2\Delta_3,$$

where the definition of Δ_1 , Δ_2 and Δ_3 are immediate from the above expression.

Chapter 4

Efficient and Robust FL with Newton Algorithm

We develop a distributed second order optimization algorithm that is communication-efficient as well as robust against Byzantine failures of the worker machines. We propose COMRADE (COMunication-efficient and Robust Approximate Distributed nEwton), an iterative second order algorithm, where the worker machines communicate *only once* per iteration with the center machine. This is in sharp contrast with the state-of-the-art distributed second order algorithms like GIANT [26] and DINGO[27], where the worker machines send (functions of) local gradient and Hessian sequentially; thus ending up communicating twice with the center machine per iteration. Moreover, we show that the worker machines can further compress the local information before sending it to the center. In addition, we employ a simple norm based thresholding rule to filter-out the Byzantine worker machines. We establish the linear-quadratic rate of convergence of COMRADE and establish that the communication savings and Byzantine resilience result in only a small statistical error rate for arbitrary convex loss functions. To the best of our knowledge, this is the first work that addresses the issue of Byzantine resilience in second order distributed optimization. Furthermore, we validate our theoretical results with extensive experiments on synthetic and benchmark LIBSVM [47] data-sets and demonstrate convergence guarantees.

4.1 Introduction

In this chapter, we propose COMRADE, a distributed approximate Newton-type algorithm that communicates less and is resilient to Byzantine workers. Specifically, we consider a distributed setup with m worker machines and one center machine. The goal is to minimize a regularized convex loss $f : \mathbb{R}^d \rightarrow \mathbb{R}$, which is additive over the available data points. Furthermore, we assume that α fraction of the worker machines are Byzantine, where $\alpha \in [0, 1/2)$. We assume that Byzantine workers can send any arbitrary values to the center machine. In addition, they may completely know the learning algorithm and are allowed to collude with each other. To the best of our knowledge, this is the first work that addresses the problem of Byzantine resilience in second order optimization.

In our proposed algorithm, the worker machines communicate *only once* per iteration with the center machine. This is in sharp contrast with the state-of-the-art distributed second order algorithms (like GIANT [26], DINGO [27], Determinantal Averaging [34]), which sequentially estimates functions of local gradients and Hessians and communicate them with the center machine. In this way, they end up communicating twice per iteration with the center machine. We show that this sequential estimation is redundant. Instead, in COMRADE, the worker machines only send a d dimensional vector, the product of the inverse of local Hessian and the local gradient. Via sketching arguments, we show that the empirical mean of the product of local Hessian inverse and local gradient is close to the global Hessian inverse and gradient product, and thus just sending the above-mentioned product is sufficient to ensure convergence. Hence, in this way, we save $\mathcal{O}(d)$ bits of communication per iteration. Furthermore, in Section 4.5, we argue that, in order to cut down further communication, the worker machines can even compress the local Hessian inverse and gradient product. Specifically, we use a (generic) δ -approximate compressor ([20]) for this, that encompasses sign-based compressors like QSGD [21] and top_k sparsification [15].

For Byzantine resilience, COMRADE employs a simple thresholding policy on the norms of the local Hessian inverse and local gradient product. Note that norm-based thresholding is computationally much simpler in comparison to existing co-ordinate wise median or trimmed mean ([23]) algorithms. Since the norm of the Hessian-inverse and gradient product determines the *amount* of movement for Newton-type algorithms, this norm corresponds to a natural metric for identifying and filtering out Byzantine workers.

Our Contributions: We propose a communication efficient Newton-type algorithm that is robust to Byzantine worker machines. Our proposed algorithm, COMRADE takes as input the local Hessian inverse and gradient product (or a compressed version of it) from the worker machines, and performs a simple thresholding operation on the norm of the said vector to discard $\beta > \alpha$ fraction of workers having largest norm values. We prove the linear-quadratic rate of convergence of our proposed algorithm for strongly convex loss functions. In particular, suppose there are m worker machines, each containing s data points; and let $\Delta_t = w_t - w^*$, where w_t is the t -th iterate of COMRADE, and w^* is the optimal model we want to estimate. In Theorem 2, we show that

$$\|\Delta_{t+1}\| \leq \max\{\Psi_t^{(1)}\|\Delta_t\|, \Psi_t^{(2)}\|\Delta_t\|^2\} + (\Psi_t^{(3)} + \alpha)\sqrt{\frac{1}{s}},$$

where $\{\Psi_t^{(i)}\}_{i=1}^3$ are quantities dependent on several problem parameters. Notice that the above implies a quadratic rate of convergence when $\|\Delta_t\| \geq \Psi_t^{(1)}/\Psi_t^{(2)}$. Subsequently, when $\|\Delta_t\|$ becomes sufficiently small, the above condition is violated and the convergence slows down to a linear rate. The error-floor, which is $\mathcal{O}(1/\sqrt{s})$ comes from the Byzantine resilience subroutine in conjunction with the simultaneous estimation of Hessian and gradient. Furthermore, in Section 4.5, we consider worker machines compressing the local Hessian inverse and gradient product via a δ -approximate compressor [20], and show that the (order-wise) rate of convergence remain unchanged, and the compression factor, δ affects the constants only.

We experimentally validate our proposed algorithm, COMRADE, with several benchmark data-sets. We consider several types of Byzantine attacks and observe that COMRADE is robust against Byzantine worker machines, yielding better classification accuracy compared to the existing state-of-the-art second order algorithms.

A major technical challenge of this work is to approximate local gradient and Hessian simultaneously in the presence of Byzantine workers. We use sketching, similar to [26], along with the norm based Byzantine resilience technique. Using *incoherence* (defined shortly) of the local Hessian along with concentration results originating from uniform sampling, we obtain the simultaneous gradient and Hessian approximation. Furthermore, ensuring at least one non-Byzantine machine gets trimmed at every iteration of COMRADE, we control the influence of Byzantine workers.

Related Work: Second order optimization has received a lot of attention in the recent years in the distributed setting owing to its attractive convergence speed. The fundamentals of second order optimization is laid out in [24], and an extension with better convergence rates is presented in [29]. Recently, in GIANT [26] algorithm, each worker machine computes an approximate Newton direction in each iteration and the center machine averages them to obtain a *globally improved* approximate Newton direction. Furthermore, DINGO [27] generalizes second order optimization beyond convex functions by extending the Newton-MR [28] algorithm in a distributed setting. Very recently, [34] proposes Determinantal averaging to correct the inversion bias of the second order optimization. A slightly different line of work ([130], [131], [132]) uses Hessian sketching to solve a large-scale distributed learning problems.

Furthermore, as mentioned in Chapter 3, Byzantine robust optimization has also received significant interest over the years, in papers like [19, 30]. Later in [23, 31], co-ordinate wise median, trimmed mean and iterative filtering based algorithm have been proposed and optimal statistical error rate is obtained. Also, [122, 123] consider adversaries may steer convergence to bad local minimizers for non-convex optimization problems.

Organization: In Section 4.3, we first analyze COMRADE with *one round* of communication per iteration. We assume $\alpha = 0$, and focus on the communication efficiency aspect only. Subsequently, in Section 4.4, we make $\alpha \neq 0$, thereby addressing communication efficiency and Byzantine resilience simultaneously. Further, in Section 4.5 we augment a compression scheme along with the setting of Section 4.4. Finally, in Section 4.6, we validate our theoretical findings with experiments. Proofs of all theoretical results can be found in the supplementary material.

Notation: For a positive integer r , $[r]$ denotes the set $\{1, 2, \dots, r\}$. For a vector v , we use $\|v\|$ to denote the ℓ_2 norm unless otherwise specified. For a matrix X , we denote $\|X\|_2$ denotes the operator norm, $\sigma_{max}(X)$ and $\sigma_{min}(X)$ denote the maximum and minimum singular value. Throughout the chapter, we use C, C_1, c, c_1 to denote positive universal constants, whose value changes with instances.

4.2 Problem Formulation

We begin with the standard statistical learning framework for empirical risk minimization, where the objective is to minimize the following loss function:

$$f(w) = \frac{1}{n} \sum_{j=1}^n \ell_j(w^T \mathbf{x}_j) + \frac{\lambda}{2} \|w\|^2, \quad (4.1)$$

where, the loss functions $\ell_j : \mathbb{R} \rightarrow \mathbb{R}$, $j \in [n]$ are *convex, twice differentiable and smooth*. Moreover, $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n \in \mathbb{R}^d$ denote the input feature vectors and $y_1, y_2, \dots, y_n \in \mathbb{R}$ denote the corresponding responses. Furthermore, we assume that the function f is *strongly convex*, implying the existence of a unique minimizer of (4.1). We denote this minimizer by w^* . Note that the response $\{y_j\}_{j=1}^n$ is captured by the corresponding loss function $\{\ell_j\}_{j=1}^n$. Some examples of ℓ_j are

$$\text{logistic loss: } \ell_j(z_j) = \log(1 - \exp(-z_j y_j)), \quad \text{squared loss: } \ell_j(z_j) = \frac{1}{2}(z_j - y_j)^2$$

We consider the framework of distributed optimization with m worker machines, where the feature vectors and the loss functions $(\mathbf{x}_1, \ell_1), \dots, (\mathbf{x}_n, \ell_n)$ are partitioned homogeneously among them. Furthermore, we assume that α fraction of the worker machines are Byzantine for some $\alpha < \frac{1}{2}$. The Byzantine machines, by nature, may send any arbitrary values to the center machine. Moreover, they can even collude with each other and plan malicious attacks with complete information of the learning algorithm.

4.3 COMRADE Can Communicate Less

We first present the Newton-type learning algorithm, namely COMRADE without any Byzantine workers, i.e., $\alpha = 0$. It is formally given in Algorithm 4 (with $\beta = 0$). In each iteration of our algorithm, every worker machine computes the local Hessian and local gradient and sends the local second order update (which is the product of the inverse of the local Hessian and local gradient) to the center machine. The center machine aggregates the updates from the worker machines by averaging them and updates the model parameter w . Later the center machine broadcast the parameter w to all the worker machines.

In any iteration t , a standard Newton algorithm requires the computation of exact Hessian (\mathbf{H}_t) and gradient (\mathbf{g}_t) of the loss function which can be written as

$$\mathbf{g}_t = \frac{1}{n} \sum_{i=1}^n \ell'_j(w_t^\top \mathbf{x}_i) \mathbf{x}_i + \lambda w_t, \quad \mathbf{H}_t = \frac{1}{n} \sum_{i=1}^n \ell''_j(w_t^\top \mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top + \lambda \mathbf{I}. \quad (4.2)$$

In a distributed set up, the exact Hessian (\mathbf{H}_t) and gradient (\mathbf{g}_t) can be computed in parallel in the following manner. In each iteration, the center machine ‘broadcasts’ the model parameter w_t to the worker machines and each worker machine computes its own local gradient and Hessian. Then the

center machine can compute the exact gradient and exact Hessian by averaging the the local gradient vectors and local Hessian matrices. But for each worker machine the per iteration communication complexity is $\mathcal{O}(d)$ for the gradient computation and $\mathcal{O}(d^2)$ for the Hessian computation. Using Algorithm 4, we reduce the communication cost to only $\mathcal{O}(d)$ per iteration, which is the same as the first order methods.

Each worker machine possess s samples drawn uniformly from $\{(\mathbf{x}_1, \ell_1), (\mathbf{x}_2, \ell_2), \dots, (\mathbf{x}_n, \ell_n)\}$. By S_i , we denote the indices of the samples held by worker machine i . At any iteration t , the worker machine computes the local Hessian $\mathbf{H}_{i,t}$ and local gradient $\mathbf{g}_{i,t}$ as

$$\mathbf{g}_{i,t} = \frac{1}{s} \sum_{i \in S_i} \ell'_j(w_t^\top \mathbf{x}_i) \mathbf{x}_i + \lambda w_t, \quad \mathbf{H}_{i,t} = \frac{1}{s} \sum_{i \in S_i} \ell''_j(w_t^\top \mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top + \lambda \mathbf{I}. \quad (4.3)$$

It is evident from the uniform sampling that $\mathbb{E}[\mathbf{g}_{i,t}] = \mathbf{g}_t$ and $\mathbb{E}[\mathbf{H}_{i,t}] = \mathbf{H}_t$. The update direction from the worker machine is defined as $\hat{\mathbf{p}}_{i,t} = (\mathbf{H}_{i,t})^{-1} \mathbf{g}_{i,t}$. Each worker machine requires $\mathcal{O}(sd^2)$ operations to compute the Hessian matrix $\mathbf{H}_{i,t}$ and $\mathcal{O}(d^3)$ operations to invert the matrix. In practice, the computational cost can be reduced by employing conjugate gradient method. The center machine computes the parameter update direction $\hat{\mathbf{p}}_t = \frac{1}{m} \sum_{i=1}^m \hat{\mathbf{p}}_{i,t}$.

We show that given large enough sample in each worker machine (s is large) and with incoherent data points (the information is spread out and not concentrated to a small number of sample data points), the local Hessian $\mathbf{H}_{i,t}$ is close to the global Hessian \mathbf{H}_t in spectral norm, and the local gradient $\mathbf{g}_{i,t}$ is close to the global gradient \mathbf{g}_t . Subsequently, we prove that the empirical average of the local updates acts as a good proxy for the global Newton update and achieves good convergence guarantee.

4.3.1 Theoretical Guarantee

We define the matrix $\mathbf{A}_t^\top = [\mathbf{a}_1^\top, \dots, \mathbf{a}_n^\top] \in \mathbb{R}^{d \times n}$ where $\mathbf{a}_j = \sqrt{\ell''_j(w^\top \mathbf{x}_j)} \mathbf{x}_j$. So the exact Hessian in equation (4.2) is $\mathbf{H}_t = \frac{1}{n} \mathbf{A}_t^\top \mathbf{A}_t + \lambda \mathbf{I}$. Also we define $\mathbf{B}_t = [\mathbf{b}_1, \dots, \mathbf{b}_n] \in \mathbb{R}^{d \times n}$ where $\mathbf{b}_i = \ell'_i(w^\top \mathbf{x}_i) \mathbf{x}_i$. So the exact gradient in equation (4.2) is $\mathbf{g}_t = \frac{1}{n} \mathbf{B}_t \mathbf{1} + \lambda w_t$

Definition 10 (Coherence of a Matrix). *Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ be any matrix with $\mathbf{U} \in \mathbb{R}^{n \times d}$ being its orthonormal basis (the left singular vectors). The row coherence of the matrix \mathbf{A} is defined as $\mu(\mathbf{A}) = \frac{n}{d} \max_i \|\mathbf{u}_i\|^2 \in [1, \frac{n}{d}]$, where \mathbf{u}_i is the i th row of \mathbf{U} .*

Remark 15. If the coherence of \mathbf{A}_t is small, it can be shown that the Hessian matrix can be approximated well via selecting a subset of rows. Note that this is a fairly common to use coherence condition as an approximation tool (see [133–135])

In the following, we assume that the Hessian matrix is L -Lipschitz (see definition below), which is a standard assumption for the analysis of the second order method for general smooth loss function (as seen in [26],[34]).

Assumption 15. *The Hessian matrix of the loss function f is L -Lipschitz continuous i.e.,*

$$\|\nabla^2 f(w) - \nabla^2 f(w')\|_2 \leq L \|w - w'\|.$$

Algorithm 4: COMMunication-efficient and Robust Approximate Distributed nEwton (COMRADE)

- 1: **Input:** Step size γ , parameter $\beta \geq 0$
 - 2: **Initialize:** Initial iterate $w_0 \in \mathbb{R}^d$
 - 3: **for** $t = 0, 1, \dots, T - 1$ **do**
 - 4: Central machine: broadcasts w_t
for $i \in [m]$ **do in parallel**
 - 5: i -th worker machine:
 - Non-Byzantine: Computes local gradient $\mathbf{g}_{i,t}$ and local Hessian $\mathbf{H}_{i,t}$; sends $\hat{\mathbf{p}}_{i,t} = (\mathbf{H}_{i,t})^{-1} \mathbf{g}_{i,t}$ to the central machine,
 - Byzantine: Generates \star (arbitrary), and sends it to the center machine
 - end for**
 - 6: Center Machine:
 - Sort the worker machines in a non decreasing order according to norm of updates $\{\hat{\mathbf{p}}_{i,t}\}_{i=1}^m$ from the local machines
 - Return the indices of the first $1 - \beta$ fraction of machines as \mathcal{U}_t ,
 - Approximate Newton Update direction : $\hat{\mathbf{p}}_t = \frac{1}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} \hat{\mathbf{p}}_{i,t}$
 - Update model parameter: $w_{t+1} = w_t - \gamma \hat{\mathbf{p}}_t$.
 - 7: **end for**
-

In the following theorem, we provide the convergence rate of COMRADE (with $\alpha = \beta = 0$) in the terms of $\Delta_t = w_t - w^*$. Also, we define $\kappa_t = \sigma_{\max}(\mathbf{H}_t) / \sigma_{\min}(\mathbf{H}_t)$ as the condition number of \mathbf{H}_t , and hence $\kappa_t \geq 1$.

Theorem 6. Let $\mu \in [1, \frac{n}{d}]$ be the coherence of \mathbf{A}_t . Suppose $\gamma = 1$ and $s \geq \frac{3\mu d}{\eta^2} \log \frac{md}{\rho}$ for some $\eta, \rho \in (0, 1)$. Under Assumption 15, with probability exceeding $1 - \rho$, we obtain

$$\|\Delta_{t+1}\| \leq \max\left\{\sqrt{\kappa_t \left(\frac{\zeta^2}{1 - \zeta^2}\right)} \|\Delta_t\|, \frac{L}{\sigma_{\min}(\mathbf{H}_t)} \|\Delta_t\|^2\right\} + \frac{2\epsilon}{\sqrt{\sigma_{\min}(\mathbf{H}_t)}},$$

where $\zeta = \nu \left(\frac{\eta}{\sqrt{m}} + \frac{\eta^2}{1-\eta}\right)$, $\nu = \frac{\sigma_{\max}(\mathbf{A}^\top \mathbf{A})}{\sigma_{\max}(\mathbf{A}^\top \mathbf{A}) + n\lambda} \leq 1$, and

$$\epsilon = \frac{1}{1 - \eta} \frac{1}{\sqrt{\sigma_{\min}(\mathbf{H}_t)}} \left(1 + \sqrt{2 \ln\left(\frac{m}{\rho}\right)}\right) \sqrt{\frac{1}{s}} \max_i \|\mathbf{b}_i\|. \quad (4.4)$$

Remark 16. It is well known that a distributed Newton method has linear-quadratic convergence rate. In Theorem 6 the quadratic term comes from the standard analysis of Newton method. The

linear term (which is small) arises owing to Hessian approximation. It gets smaller with better Hessian approximation (smaller η), and thus the above rate becomes quadratic one. The small error floor arises due to the gradient approximation in the worker machines, which is essential for the one round of communication per iteration. The error floor is $\propto \frac{1}{\sqrt{s}}$ where s is the number of samples in each worker machine. So for a sufficiently large s , the error floor becomes negligible.

Remark 17. The sample size in each worker machine is dependent on the coherence of the matrix \mathbf{A}_t and the dimension d of the problem. Theoretically, the analysis is feasible for the case of $s \geq d$ (since we work with $\mathbf{H}_{i,t}^{-1}$). However, when $s < d$, one can replace the inverse by a pseudo-inverse (modulo some changes in convergence rate).

4.4 COMRADE Can Resist Byzantine Workers

In this section, we analyze COMRADE with Byzantine workers. We assume that $\alpha (< 1/2)$ fraction of worker machines are Byzantine. We define the set of Byzantine worker machines by \mathcal{B} and the set of the good (non-Byzantine) machines by \mathcal{M} . COMRADE employs a ‘norm based thresholding’ scheme on the local Hessian inverse and gradient product to tackle the Byzantine workers.

In the t -th iteration, the center machine outputs a set \mathcal{U}_t with $|\mathcal{U}_t| = (1 - \beta)m$, consisting the indices of the worker machines with smallest norm. Hence, we ‘trim’ the worker machines that may try to diverge the learning algorithm. We denote the set of trimmed machines as \mathcal{T}_t . Moreover, we take $\beta > \alpha$ to ensure at least one good machine falls in \mathcal{T}_t . This condition helps us to control the Byzantine worker machines. Finally, the update is given by $\hat{\mathbf{p}}_t = \frac{1}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} \hat{\mathbf{p}}_{i,t}$. We define:

$$\epsilon_{byz}^2 = [3(\frac{1-\alpha}{1-\beta})^2 + 4\kappa_t(\frac{\alpha}{1-\beta})^2]\epsilon^2, \quad (4.5)$$

$$\zeta_{byz}^2 = 2(\frac{1-\alpha}{1-\beta})^2(\frac{\nu}{1-\eta})^2 + \nu^2(\frac{1-\alpha}{1-\beta})^2(\frac{\eta}{\sqrt{(1-\alpha)m}} + \frac{\eta^2}{1-\eta})^2 + 4\kappa_t(\frac{\alpha}{1-\beta})^2[2 + (\frac{\nu}{1-\eta})^2]. \quad (4.6)$$

ϵ is defined in (4.4), $\nu = \frac{\sigma_{\max}(\mathbf{A}^T \mathbf{A})}{\sigma_{\max}(\mathbf{A}^T \mathbf{A}) + n\lambda}$ and κ_t is the condition number of the exact Hessian \mathbf{H}_t .

Theorem 7. *Let $\mu \in [1, \frac{n}{d}]$ be the coherence of \mathbf{A}_t . Suppose $\gamma = 1$ and $s \geq \frac{3\mu d}{\eta^2} \log \frac{md}{\rho}$ for some $\eta, \rho \in (0, 1)$. For $0 \leq \alpha < \beta < 1/2$, under Assumption 15, with probability exceeding $1 - \rho$, Algorithm 4 yields*

$$\|\Delta_{t+1}\| \leq \max\left\{\sqrt{\kappa_t\left(\frac{\zeta_{byz}^2}{1-\zeta_{byz}^2}\right)}\|\Delta_t\|, \frac{L}{\sigma_{\min}(\mathbf{H}_t)}\|\Delta_t\|^2\right\} + \frac{2\epsilon_{byz}}{\sqrt{\sigma_{\min}(\mathbf{H}_t)}},$$

where ζ_{byz} and ϵ_{byz} are defined in equations (4.5) and (4.6) respectively.

The remarks of Section 4.3 is also applicable here. On top of that, we have the following remarks:

Remark 18. Compared to the convergence rate of Theorem 6, the rate here remains order-wise same even with Byzantine robustness. The coefficient of the quadratic term remains unchanged but the linear rate and the error floor suffers a little bit (by a small constant factor).

Remark 19. Note that for Theorem 7 to hold, we require $\alpha \sim 1/\sqrt{\kappa_t}$ for all t . In cases where κ_t is large, this can impose a stricter condition on α . However, we conjecture that this dependence can be improved via applying a more intricate (and perhaps computation heavy) Byzantine resilience algorithm. In this work, we kept the Byzantine resilience scheme simple at the expense of this condition on α .

4.5 COMRADE Can Communicate Even Less and Resist Byzantine Workers

In Section 4.3 we analyze COMRADE with an additional feature. We let the worker machines further reduce the communication cost by applying a generic class of δ -approximate compressor [20] on the parameter update of Algorithm 4. We first define the class of δ -approximate compressor:

Definition 11. An operator $\mathcal{Q} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is defined as δ -approximate compressor on a set $S \subset \mathbb{R}^d$ if, $\forall x \in S$, $\|\mathcal{Q}(x) - x\|^2 \leq (1 - \delta) \|x\|^2$, where $\delta \in [0, 1]$ is the compression factor.

The above definition can be extended for any randomized operator \mathcal{Q} satisfying

$$\mathbb{E}(\|\mathcal{Q}(x) - x\|^2) \leq (1 - \delta) \|x\|^2,$$

for all $\forall x \in S$. The expectation is taken over the randomization of the operator. Notice that $\delta = 1$ implies that $\mathcal{Q}(x) = x$ (no compression). Examples of δ -approximate compressor include QSGD [21], ℓ_1 -QSGD [20], top_k sparsification and rand_k [15].

Worker machine i computes the product of local Hessian inverse inverse and local gradient and then apply δ -approximate compressor to obtain $\mathcal{Q}(\mathbf{H}_{i,t}^{-1} \mathbf{g}_{i,t})$; and finally sends this compressed vector to the center. The Byzantine resilience subroutine remains the same—except, instead of sorting with respect to $\|\mathbf{H}_{i,t}^{-1} \mathbf{g}_{i,t}\|$, the center machine now sorts according to $\|\mathcal{Q}(\mathbf{H}_{i,t}^{-1} \mathbf{g}_{i,t})\|$. The center machine aggregates the compressed updates by averaging $\mathcal{Q}(\hat{\mathbf{p}}) = \frac{1}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} \mathcal{Q}(\hat{\mathbf{p}}_{i,t})$, and take the next step as $w_{t+1} = w_t - \gamma \mathcal{Q}(\hat{\mathbf{p}})$.

Recall the definition of ϵ from (4.4). We also use the following notation : $\zeta_{\mathcal{M}}^2 = \nu \left(\frac{\eta}{\sqrt{(1-\alpha)m}} + \frac{\eta^2}{1-\eta} \right)$, $\zeta_1 = \frac{\nu}{1-\eta}$ and $\nu = \frac{\sigma_{\max}(\mathbf{A}^T \mathbf{A})}{\sigma_{\max}(\mathbf{A}^T \mathbf{A}) + n\lambda}$. Furthermore, we define the following:

$$\epsilon_{\text{comp,byz}}^2 = \left[3 \left(\frac{1-\alpha}{1-\beta} \right)^2 + 4\kappa_t \left(\frac{\alpha}{1-\beta} \right)^2 \right] (1 + \kappa(1-\delta)) \epsilon^2 \quad (4.7)$$

$$\begin{aligned} \zeta_{\text{comp,byz}}^2 &= 2 \left(\frac{1-\alpha}{1-\beta} \right)^2 (\zeta_1^2 + \kappa_t(1-\rho)((1 + \zeta_1^2))) + \left(\frac{1-\alpha}{1-\beta} \right)^2 (\zeta_{\mathcal{M}}^2 + \kappa_t(1-\delta)((1 + \zeta_1^2))) \\ &\quad + 4\kappa_t \left(\frac{\alpha}{1-\beta} \right)^2 (2 + (\zeta_1^2 + \kappa_t(1-\delta)((1 + \zeta_1^2)))) \end{aligned} \quad (4.8)$$

Theorem 8. Let $\mu \in [1, \frac{n}{d}]$ be the coherence of \mathbf{A}_t . Let $\gamma = 1$ and $s \geq \frac{3\mu d}{\eta^2} \log \frac{md}{\rho}$ for some $\eta, \rho \in (0, 1)$. For $0 \leq \alpha < \beta < 1/2$, under Assumption 15 and with \mathcal{Q} being the δ -approximate




Figure 4.1: (First row) Comparison of training accuracy between COMRADE(Algorithm 4) and GIANT [26] with (a) w5a (b) a9a (c) Epsilon (d) Covtype dataset. (Second row) Training accuracy of (e) GIANT for ‘flipped label’ and (f) ‘negative update’ attack; and comparison of Robust GIANT and COMRADE with a9a dataset for (g) ‘flipped label’ and (h) ‘negative update’ attack.

compressor, with probability exceeding $1 - \rho$, we obtain

$$\|\Delta_{t+1}\| \leq \max\left\{\sqrt{\kappa_t \left(\frac{\zeta_{comp,byz}^2}{1 - \zeta_{comp,byz}^2}\right)} \|\Delta_t\|, \frac{L}{\sigma_{min}(\mathbf{H}_t)} \|\Delta_t\|^2\right\} + \frac{\epsilon_{comp,byz}}{\sqrt{\sigma_{min}(\mathbf{H}_t)}}$$

where $\epsilon_{comp,byz}$ and $\zeta_{comp,byz}$ are given in equations (4.7) and (4.8) respectively.

Remark 20. With no compression ($\delta = 1$) we get back the convergence guarantee of Theorem 7.

Remark 21. Note that even with compression, we retain the linear-quadratic rate of convergence of COMRADE. The constants are affected by a δ -dependent term.

4.6 Experimental Results

In this section we validate our algorithm, COMRADE in Byzantine and non-Byzantine setup on synthetically generated and benchmark LIBSVM [47] data-set. The experiments focus on the standard logistic regression problem. The logistic regression objective is defined as

$$\frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-y_i \mathbf{x}_i^\top w)) + \frac{\lambda}{2n} \|w\|^2,$$

where $w \in \mathbb{R}^d$ is the parameter, $\{\mathbf{x}_i\}_{i=1}^n \in \mathbb{R}^d$ are the feature data and $\{y_i\}_{i=1}^n \in \{0, 1\}$ are the corresponding labels. We use ‘mpi4py’ package for distributed framework (swarm2) at the University of Massachusetts Amherst [136] using mpi4py Python package. We choose ‘a9a’ ($d = 123, n \approx 32\text{K}$), ‘w5a’ ($d = 300, n \approx 10\text{k}$), ‘Epsilon’ ($d = 2000, n = 0.4\text{M}$) and ‘covtype.binary’ ($d = 54, n \approx 0.5\text{M}$) classification datasets and partition the data in 20 different worker machines. In the experiments, we choose two types of Byzantine attacks : (1). ‘flipped label’-attack where (for binary classification) the Byzantine worker machines flip the labels of the data, thus making the model learn with wrong labels, and (2). ‘negative update attack’ where the Byzantine worker machines compute the local update ($\hat{\mathbf{p}}_i$) and communicate $-c \times \hat{\mathbf{p}}_i$ with $c \in (0, 1)$ making the updates to be opposite of actual direction. We choose $\beta = \alpha + \frac{2}{m}$. We choose the regularization parameter $\lambda = 1$ and fixed step size. We ran the algorithms sufficient number of steps to ensure convergence.

In Figure 4.1(first row) we compare COMRADE in non-Byzantine setup ($\alpha = \beta = 0$) with the state-of the art algorithm GIANT [26]. It is evident from the plot that despite the fact that COMRADE requires less communication, the algorithm is able to achieve similar accuracy. Also, we show the ineffectiveness of GIANT in the presence of Byzantine attacks. In Figure 4.2((e),(f)) we show the accuracy for flipped label and negative update attacks. These plots are an indicator of the requirement of robustness in the learning algorithm. So we devise ‘Robust GIANT’, which is GIANT algorithm with added ‘norm based thresholding’ for robustness. In particular, we trim the worker machines based on the local gradient norm in the first round of communication of GIANT. Subsequently, in the second round of communication, the non-trimmed worker machines send the updates (product of local Hessian inverse and the local gradient) to the center machine. We compare COMRADE with ‘Robust GIANT’ in Figure 4.1((g),(h)) with 10% Byzantine worker machines for ‘a9a’ dataset. It is evident plot that COMRADE performs better than the ‘Robust GIANT’.

Next we show the accuracy of COMRADE with different numbers of Byzantine worker machines. Here we choose $c = 0.9$. We show the accuracy for ‘negative update’ attack in Figure 4.2(first row) and ‘flipped label’ attack in Figure 4.2 (second row). Furthermore, we show that COMRADE works even when δ -approximate compressor is applied to the updates. In Figure 4.2(Third row) we plot the training accuracies. For compression we apply the scheme known as QSGD [21]. Further experiments can be found in the supplementary material.

4.7 Conclusion

In this paper, we address the issue of communication efficiency and Byzantine robustness via second order optimization and norm based thresholding respectively for strongly convex loss. Extending our setting to handle weakly convex and non-convex loss is of immediate interest. We would also like to exploit local averaging with second order optimization. Moreover, an import aspect, privacy, is not addressed in this work. We keep this as our future research direction.




Figure 4.2: (First row) Accuracy of COMRADE with 10%, 15%, 20% Byzantine workers with ‘negative update’ attack for (a). w5a (b). a9a (c). covtype (d). Epsilon. (Second row) COMRADE accuracy with 10%, 15%, 20% Byzantine workers with ‘flipped label’ attack for (e) w5a (f) a9a (g) covtype (h) Epsilon. (Third row) Accuracy of COMRADE with ρ -approximate compressor (Section 4.5) with 10%, 15%, 20% Byzantine workers; (i) ‘flipped label’ attack for w5a (j) ‘negative update’ attack for w5a. (k) ‘flipped label’ attack for a9a . (l) ‘negative update’ attack for a9a dataset.

Appendix

4.8 Analysis of Section 4.3

Matrix Sketching

Here we briefly discuss the matrix sketching that is broadly used in the context of *randomized linear algebra*. For any matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ the sketched matrix $\mathbf{Z} \in \mathbb{R}^{s \times d}$ is defined as $\mathbf{S}^T \mathbf{A}$ where $\mathbf{S} \in \mathbb{R}^{n \times s}$ is the sketching matrix (typically $s < n$). Based on the scope and basis of the application, the sketched matrix is constructed by taking linear combination of the rows of matrix which is known as *random projection* or by sampling and scaling a subset of the rows of the matrix which is known as *random sampling*. The sketching is done to get a smaller representation of the original

matrix to reduce computational cost.

Here we consider a uniform row sampling scheme. The matrix \mathbf{Z} is formed by sampling and scaling rows of the matrix \mathbf{A} . Each row of the matrix \mathbf{A} is sampled with probability $p = \frac{1}{n}$ and scaled by multiplying with $\frac{1}{\sqrt{sp}}$.

$$\mathbb{P}\left(\mathbf{z}_i = \frac{\mathbf{a}_j}{\sqrt{sp}}\right) = p,$$

where \mathbf{z}_i is the i -th row matrix \mathbf{Z} and \mathbf{a}_j is the j th row of the matrix \mathbf{A} . Consequently the sketching matrix \mathbf{S} has one non-zero entry in each column.

We define the matrix $\mathbf{A}_t^\top = [\mathbf{a}_1^\top, \dots, \mathbf{a}_n^\top] \in \mathbb{R}^{d \times n}$ where $\mathbf{a}_j = \sqrt{\ell_j''(w^\top \mathbf{x}_j)} \mathbf{x}_j$. So the exact Hessian in equation (4.2) is $\mathbf{H}_t = \frac{1}{n} \mathbf{A}_t^\top \mathbf{A}_t + \lambda \mathbf{I}$. Assume that S_i is the set of features that are held by the i th worker machine. So the local Hessian is

$$\mathbf{H}_{i,t} = \frac{1}{s} \sum_{j \in S_i} \ell_j''(w^\top \mathbf{x}_j) \mathbf{x}_j \mathbf{x}_j^\top + \lambda \mathbf{I} = \frac{1}{s} \mathbf{A}_{i,t}^\top \mathbf{A}_{i,t} + \lambda \mathbf{I},$$

where $\mathbf{A}_{i,t} \in \mathbb{R}^{s \times d}$ and the row of the matrix $\mathbf{A}_{i,t}$ is indexed by S_i . Also we define $\mathbf{B}_t = [\mathbf{b}_1, \dots, \mathbf{b}_n] \in \mathbb{R}^{d \times n}$ where $\mathbf{b}_i = \ell_i'(w^\top \mathbf{x}_i) \mathbf{x}_i$. So the exact gradient in equation (4.2) is $\mathbf{g}_t = \frac{1}{n} \mathbf{B}_t \mathbf{1} + \lambda w_t$ and the local gradient is

$$\mathbf{g}_{i,t} = \frac{1}{s} \sum_{i \in S_i} \ell_j'(w_t^\top \mathbf{x}_i) \mathbf{x}_i + \lambda w_t = \frac{1}{s} \mathbf{B}_{i,t} \mathbf{1} + \lambda w_t,$$

where $\mathbf{B}_{i,t}$ is the matrix with column indexed by S_i . If $\{\mathbf{S}_i\}_{i=1}^m$ are the sketching matrices then the local Hessian and gradient can be expressed as

$$\mathbf{H}_{i,t} = \mathbf{A}_t^\top \mathbf{S}_i \mathbf{S}_i^\top \mathbf{A}_t + \lambda \mathbf{I} \quad \mathbf{g}_{i,t} = \frac{1}{n} \mathbf{B} \mathbf{S}_i \mathbf{S}_i^\top \mathbf{1} + \lambda w. \quad (4.9)$$

With the help of sketching idea later we show that the local hessian and gradient are close to the exact hessian and gradient.

The Quadratic function For the purpose of analysis we define an auxiliary quadratic function

$$\phi(\mathbf{p}) = \frac{1}{2} \mathbf{p}^\top \mathbf{H}_t \mathbf{p} - \mathbf{g}_t^\top \mathbf{p} = \frac{1}{2} \mathbf{p}^\top (\mathbf{A}_t^\top \mathbf{A}_t + \lambda \mathbf{I}) \mathbf{p} - \mathbf{g}_t^\top \mathbf{p}. \quad (4.10)$$

The optimal solution to the above function is

$$\mathbf{p}^* = \arg \min \phi(\mathbf{p}) = \mathbf{H}_t^{-1} \mathbf{g}_t = (\mathbf{A}_t^\top \mathbf{A}_t + \lambda \mathbf{I})^{-1} \mathbf{g}_t,$$

which is also the optimal direction of the global Newton update. In this work we consider the local and global (approximate) Newton direction to be

$$\hat{\mathbf{p}}_{i,t} = (\mathbf{A}^\top \mathbf{S}_i \mathbf{S}_i^\top \mathbf{A} + \lambda \mathbf{I})^{-1} \mathbf{g}_{i,t}, \quad \hat{\mathbf{p}}_t = \frac{1}{m} \sum_{i=1}^m \hat{\mathbf{p}}_{i,t}.$$

respectively. And it can be easily verified that each local update $\hat{\mathbf{p}}_{i,t}$ is optimal solution to the following quadratic function

$$\hat{\phi}_{i,t}(p) = \frac{1}{2} \mathbf{p}^\top (\mathbf{A}^\top \mathbf{S}_i \mathbf{S}_i^\top \mathbf{A} + \lambda \mathbf{I}) \mathbf{p} - \mathbf{g}_i^\top \mathbf{p}. \quad (4.11)$$

In our convergence analysis we show that value of the quadratic function in (4.10) with value $\hat{\mathbf{p}}_t$ is close to the optimal value.

Singular Value Decomposition (SVD) For any matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ with rank r , the singular value decomposition is defined as $\mathbf{A} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top$ where \mathbf{U} , \mathbf{V} are $n \times r$ and $d \times r$ column orthogonal matrices respectively and $\mathbf{\Sigma}$ is a $r \times r$ diagonal matrix with diagonal entries $\{\sigma_1, \dots, \sigma_r\}$. If \mathbf{A} is a symmetric positive semi-definite matrix then $\mathbf{U} = \mathbf{V}$.

4.8.1 Analysis

Lemma 10 (McDiarmid's Inequality). *Let $X = X_1, \dots, X_m$ be m independent random variables taking values from some set A , and assume that $f : A^m \rightarrow \mathbb{R}$ satisfies the following condition (bounded differences):*

$$\sup_{x_1, \dots, x_m, \hat{x}_i} |f(x_1, \dots, x_i, \dots, x_m) - f(x_1, \dots, \hat{x}_i, \dots, x_m)| \leq c_i,$$

for all $i \in \{1, \dots, m\}$. Then for any $\epsilon > 0$ we have

$$P[f(X_1, \dots, X_m) - \mathbb{E}[f(X_1, \dots, X_m)] \geq \epsilon] \leq \exp\left(-\frac{2\epsilon^2}{\sum_{i=1}^m c_i^2}\right).$$

The property described in the following Lemma 11 is a very useful result for uniform row sampling sketching matrix.

Lemma 11 (Lemma 8 [26]). *Let $\eta, \rho \in (0, 1)$ be a fixed parameter and $r = \text{rank}(\mathbf{A}_t)$ and $\mathbf{U} \in \mathbb{R}^{n \times r}$ be the orthonormal bases of the matrix \mathbf{A}_t . Let $\{\mathbf{S}_i\}_{i=1}^m$ be sketching matrices and $\mathbf{S} = \frac{1}{\sqrt{m}}[\mathbf{S}_1, \dots, \mathbf{S}_m] \in \mathbb{R}^{n \times ms}$. With probability $1 - \rho$ the following holds*

$$\|\mathbf{U}^\top \mathbf{S}_i \mathbf{S}_i^\top \mathbf{U} - \mathbf{I}\|_2 \leq \eta \quad \forall i \in [m] \quad \text{and} \quad \|\mathbf{U}^\top \mathbf{S} \mathbf{S}^\top \mathbf{U} - \mathbf{I}\|_2 \leq \frac{\eta}{\sqrt{m}}.$$

Lemma 12. *Let $\mathbf{S} \in \mathbb{R}^{n \times s}$ be any uniform sampling sketching matrix, then for any matrix $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_n] \in \mathbb{R}^{d \times n}$ with probability $1 - \rho$ for any $\rho > 0$ we have,*

$$\left\| \frac{1}{n} \mathbf{B} \mathbf{S} \mathbf{S}^\top \mathbf{1} - \frac{1}{n} \mathbf{B} \mathbf{1} \right\| \leq \left(1 + \sqrt{2 \ln\left(\frac{1}{\rho}\right)}\right) \sqrt{\frac{1}{s}} \max_i \|\mathbf{b}_i\|,$$

where $\mathbf{1}$ is all ones vector.

Proof. The vector $\mathbf{B}\mathbf{1}$ is the sum of column of the matrix \mathbf{B} and $\mathbf{BSS}^\top \mathbf{1}$ is the sum of uniformly sampled and scaled column of the matrix \mathbf{B} where the scaling factor is $\frac{1}{\sqrt{sp}}$ with $p = \frac{1}{n}$. If (i_1, \dots, i_s) is the set of sampled indices then $\mathbf{BSS}^\top \mathbf{1} = \sum_{k \in (i_1, \dots, i_s)} \frac{1}{sp} \mathbf{b}_k$.

Define the function $f(i_1, \dots, i_s) = \|\frac{1}{n} \mathbf{BSS}^\top \mathbf{1} - \frac{1}{n} \mathbf{B}\mathbf{1}\|$. Now consider a sampled set $(i_1, \dots, i_{j'}, \dots, i_s)$ with only one item (column) replaced then the bounded difference is

$$\begin{aligned} \Delta &= |f(i_1, \dots, i_j, \dots, i_s) - f(i_1, \dots, i_{j'}, \dots, i_s)| \\ &= \left| \frac{1}{n} \left\| \frac{1}{sp} \mathbf{b}_{i_j'} - \frac{1}{sp} \mathbf{b}_{i_j} \right\| \right| \leq \frac{2}{s} \max_i \|\mathbf{b}_i\|. \end{aligned}$$

Now we have the expectation

$$\begin{aligned} \mathbb{E}[\|\frac{1}{n} \mathbf{BSS}^\top \mathbf{1} - \frac{1}{n} \mathbf{B}\mathbf{1}\|^2] &\leq \frac{n}{sn^2} \sum_{i=1}^n \|\mathbf{b}_i\|^2 = \frac{1}{s} \max_i \|\mathbf{b}_i\|^2 \\ \Rightarrow \mathbb{E}[\|\frac{1}{n} \mathbf{BSS}^\top \mathbf{1} - \frac{1}{n} \mathbf{B}\mathbf{1}\|] &\leq \sqrt{\frac{1}{s}} \max_i \|\mathbf{b}_i\|. \end{aligned}$$

Using McDiarmid inequality (Lemma 10) we have

$$P\left[\left\|\frac{1}{n} \mathbf{BSS}^\top \mathbf{1} - \frac{1}{n} \mathbf{B}\mathbf{1}\right\| \geq \sqrt{\frac{1}{s}} \max_i \|\mathbf{b}_i\| + t\right] \leq \exp\left(-\frac{2t^2}{s\Delta^2}\right).$$

Equating the probability with ρ we have

$$\begin{aligned} \exp\left(-\frac{2t^2}{s\Delta^2}\right) &= \rho \\ \Rightarrow t &= \Delta \sqrt{\frac{s}{2} \ln\left(\frac{1}{\rho}\right)} = \max_i \|\mathbf{b}_i\| \sqrt{\frac{2}{s} \ln\left(\frac{1}{\rho}\right)}. \end{aligned}$$

Finally we have with probability $1 - \rho$

$$\left\|\frac{1}{n} \mathbf{BSS}^\top \mathbf{1} - \frac{1}{n} \mathbf{B}\mathbf{1}\right\| \leq \left(1 + \sqrt{2 \ln\left(\frac{1}{\rho}\right)}\right) \sqrt{\frac{1}{s}} \max_i \|\mathbf{b}_i\|.$$

□

Remark 22. For m sketching matrix $\{\mathbf{S}_i\}_{i=1}^m$, the bound in the Lemma 12 is

$$\left\|\frac{1}{n} \mathbf{B} \mathbf{S}_i \mathbf{S}_i^\top \mathbf{1} - \frac{1}{n} \mathbf{B}\mathbf{1}\right\| \leq \left(1 + \sqrt{2 \ln\left(\frac{m}{\rho}\right)}\right) \sqrt{\frac{1}{s}} \max_i \|\mathbf{b}_i\|,$$

with probability $1 - \rho$ for any $\rho > 0$ for all $i \in \{1, 2, \dots, m\}$. In the case that each worker machine holds data based on the uniform sketching matrix the local gradient is close to the exact gradient. Thus the local second order update acts as a good approximate to the exact Newton update.

Now we consider the update rule of GIANT [26] where the update is done in two rounds in each iteration. In the first round each worker machine computes and send the local gradient and the center machine computes the exact gradient \mathbf{g}_t in iteration t . Next the center machine broadcasts the exact gradient and each worker machine computes the local Hessian and send $\tilde{\mathbf{p}}_{i,t} = (\mathbf{H}_{i,t})^{-1}\mathbf{g}_t$ to the center machine and the center machine computes the approximate Newton direction $\tilde{\mathbf{p}}_t = \frac{1}{m} \sum_{i=1}^m \tilde{\mathbf{p}}_{i,t}$. Now based on this we restate the following lemma (Lemma 6 [26]).

Lemma 13. *Let $\{\mathbf{S}_i\}_{i=1}^m \in \mathbb{R}^{n \times s}$ be sketching matrices based on Lemma 11. Let ϕ_t be defined in (4.10) and $\tilde{\mathbf{p}}_t$ be the update. It holds that*

$$\min_{\mathbf{p}} \phi_t(\mathbf{p}) \leq \phi_t(\tilde{\mathbf{p}}_t) \leq (1 - \zeta^2) \min_{\mathbf{p}} \phi_t(\mathbf{p}),$$

where $\zeta = \nu(\frac{\eta}{\sqrt{m}} + \frac{\eta^2}{1-\eta})$ and $\nu = \frac{\sigma_{max}(\mathbf{A}^\top \mathbf{A})}{\sigma_{max}(\mathbf{A}^\top \mathbf{A}) + n\lambda} \leq 1$.

Now we prove similar guarantee for the update according to COMRADE in Algorithm 4.

Lemma 14. *Let $\{\mathbf{S}_i\}_{i=1}^m \in \mathbb{R}^{n \times s}$ be sketching matrices based on Lemma 11. Let ϕ_t be defined in (4.10) and $\hat{\mathbf{p}}_t$ be defined in Algorithm 4($\beta = 0$)*

$$\min_{\mathbf{p}} \phi_t(\mathbf{p}) \leq \phi_t(\hat{\mathbf{p}}_t) \leq \epsilon^2 + (1 - \zeta^2) \min_{\mathbf{p}} \phi_t(\mathbf{p}),$$

where $\epsilon = \frac{1}{1-\eta} \frac{1}{\sqrt{\sigma_{min}(\mathbf{H}_t)}} (1 + \sqrt{2 \ln(\frac{m}{\rho})}) \sqrt{\frac{1}{s} \max_i \|\mathbf{b}_i\|}$ and $\zeta = \nu(\frac{\eta}{\sqrt{m}} + \frac{\eta^2}{1-\eta})$ and $\nu = \frac{\sigma_{max}(\mathbf{A}^\top \mathbf{A})}{\sigma_{max}(\mathbf{A}^\top \mathbf{A}) + n\lambda}$.

Proof. First consider the quadratic function (4.10)

$$\begin{aligned} \phi_t(\hat{\mathbf{p}}_t) - \phi_t(\mathbf{p}^*) &= \frac{1}{2} \|\mathbf{H}_t^{\frac{1}{2}}(\hat{\mathbf{p}}_t - \mathbf{p}^*)\|^2 \\ &\leq \underbrace{(\|\mathbf{H}_t^{\frac{1}{2}}(\hat{\mathbf{p}}_t - \tilde{\mathbf{p}}_t)\|^2)}_{Term1} + \underbrace{(\|\mathbf{H}_t^{\frac{1}{2}}(\tilde{\mathbf{p}}_t - \mathbf{p}^*)\|^2)}_{Term2}, \end{aligned} \quad (4.12)$$

where $\tilde{\mathbf{p}}_t = \frac{1}{m} \sum_{i=1}^m (\mathbf{H}_{i,t})^{-1}\mathbf{g}_t$. First we bound the Term 2 of (4.12) using the quadratic function and Lemma 13

$$\begin{aligned} \frac{1}{2} \left\| \mathbf{H}_t^{\frac{1}{2}}(\tilde{\mathbf{p}}_t - \mathbf{p}^*) \right\|^2 &\leq \zeta^2 \left\| \mathbf{H}_t^{\frac{1}{2}} \mathbf{p}^* \right\|^2 \quad (\text{Using Lemma 13}) \\ &= -\zeta^2 \phi_t(\mathbf{p}^*). \end{aligned} \quad (4.13)$$

The step in equation (4.13) is from the definition of the function ϕ_t and \mathbf{p}^* . It can be shown that

$$\phi_t(\mathbf{p}^*) = - \left\| \mathbf{H}_t^{\frac{1}{2}} \mathbf{p}^* \right\|^2.$$

Now we bound the Term 1 in (4.12). By Lemma 11, we have $(1 - \eta)\mathbf{A}_t^\top \mathbf{A}_t \preceq \mathbf{A}_t^\top \mathbf{S}_i \mathbf{S}_i^\top \mathbf{A}_t \preceq (1 + \eta)\mathbf{A}_t^\top \mathbf{A}_t$. Following we have $(1 - \eta)\mathbf{H}_t \preceq \mathbf{H}_{i,t} \preceq (1 + \eta)\mathbf{H}_t$. Thus there exists matrix ξ_i satisfying

$$\mathbf{H}_t^{\frac{1}{2}} \mathbf{H}_{i,t}^{-1} \mathbf{H}_t^{\frac{1}{2}} = \mathbf{I} + \xi_i \quad \text{and} \quad -\frac{\eta}{1 + \eta} \preceq \xi_i \preceq \frac{\eta}{1 - \eta},$$

So we have,

$$\left\| \mathbf{H}_t^{\frac{1}{2}} \mathbf{H}_{i,t}^{-1} \mathbf{H}_t^{\frac{1}{2}} \right\| \leq 1 + \frac{\eta}{1 - \eta} = \frac{1}{1 - \eta}. \quad (4.14)$$

Now we have

$$\begin{aligned} \left\| \mathbf{H}_t^{\frac{1}{2}} (\hat{\mathbf{p}}_t - \tilde{\mathbf{p}}_t) \right\| &= \left\| \mathbf{H}_t^{\frac{1}{2}} \frac{1}{m} \sum_{i=1}^m (\hat{\mathbf{p}}_{i,t} - \tilde{\mathbf{p}}_{i,t}) \right\| \\ &\leq \frac{1}{m} \sum_{i=1}^m \left\| \mathbf{H}_t^{\frac{1}{2}} (\hat{\mathbf{p}}_{i,t} - \tilde{\mathbf{p}}_{i,t}) \right\| \\ &= \frac{1}{m} \sum_{i=1}^m \left\| \mathbf{H}_t^{\frac{1}{2}} \mathbf{H}_{i,t}^{-1} (\mathbf{g}_{i,t} - \mathbf{g}_t) \right\| \\ &= \frac{1}{m} \sum_{i=1}^m \left\| \mathbf{H}_t^{\frac{1}{2}} \mathbf{H}_{i,t}^{-1} \mathbf{H}_t^{\frac{1}{2}} \mathbf{H}_t^{-\frac{1}{2}} (\mathbf{g}_{i,t} - \mathbf{g}_t) \right\| \\ &\leq \frac{1}{m} \sum_{i=1}^m \left\| \mathbf{H}_t^{\frac{1}{2}} \mathbf{H}_{i,t}^{-1} \mathbf{H}_t^{\frac{1}{2}} \right\| \left\| \mathbf{H}_t^{-\frac{1}{2}} (\mathbf{g}_{i,t} - \mathbf{g}_t) \right\| \\ &\leq \frac{1}{1 - \eta} \frac{1}{m} \sum_{i=1}^m \left\| \mathbf{H}_t^{-\frac{1}{2}} (\mathbf{g}_{i,t} - \mathbf{g}_t) \right\| \quad (\text{Using (4.14)}) \\ &\leq \frac{1}{1 - \eta} \frac{1}{\sqrt{\sigma_{\min}(\mathbf{H}_t)}} \frac{1}{m} \sum_{i=1}^m \|\mathbf{g}_{i,t} - \mathbf{g}_t\|. \end{aligned} \quad (4.15)$$

Now we bound $\|\mathbf{g}_{i,t} - \mathbf{g}_t\|$ using Lemma 12,

$$\|\mathbf{g}_{i,t} - \mathbf{g}_t\| = \left\| \frac{1}{n} \mathbf{B} \mathbf{S} \mathbf{S}^\top \mathbf{1} - \frac{1}{n} \mathbf{B} \mathbf{1} \right\| \leq (1 + \sqrt{2 \ln(\frac{m}{\rho})}) \sqrt{\frac{1}{s}} \max_i \|\mathbf{b}_i\|.$$

Plugging it into equation (4.15) we get,

$$\begin{aligned} \left\| \mathbf{H}_t^{\frac{1}{2}} (\hat{\mathbf{p}}_t - \tilde{\mathbf{p}}_t) \right\| &\leq \frac{1}{1 - \eta} \frac{1}{\sqrt{\sigma_{\min}(\mathbf{H}_t)}} \frac{1}{m} \sum_{i=1}^m \|\mathbf{g}_{i,t} - \mathbf{g}_t\| \\ &\leq \frac{1}{1 - \eta} \frac{1}{\sqrt{\sigma_{\min}(\mathbf{H}_t)}} (1 + \sqrt{2 \ln(\frac{m}{\rho})}) \sqrt{\frac{1}{s}} \max_i \|\mathbf{b}_i\|. \end{aligned} \quad (4.16)$$

Now collecting the terms of (4.16) and (4.13) and plugging them into (4.12) we have

$$\begin{aligned}\phi_t(\hat{\mathbf{p}}_t) - \phi_t(\mathbf{p}^*) &\leq \epsilon^2 - \zeta^2 \phi_t(\mathbf{p}^*) \\ \Rightarrow \phi_t(\hat{\mathbf{p}}_t) &\leq \epsilon^2 + (1 - \zeta^2) \phi_t(\mathbf{p}^*),\end{aligned}$$

where ϵ is as defined in (4.4). □

Lemma 15. *Let $\zeta \in (0, 1)$, ϵ be any fixed parameter. And $\hat{\mathbf{p}}_t$ satisfies $\phi_t(\hat{\mathbf{p}}_t) \leq \epsilon^2 + (1 - \zeta^2) \min_{\mathbf{p}} \phi_t(\mathbf{p})$. Under the Assumption 15 (Hessian L -Lipschitz) and $\Delta_t = w_t - w^*$ satisfies*

$$\Delta_{t+1}^\top \mathbf{H}_t \Delta_{t+1} \leq L \|\Delta_{t+1}\| \|\Delta_t\|^2 + \frac{\zeta^2}{1 - \zeta^2} \Delta_t^\top \mathbf{H}_t \Delta_t + 2\epsilon^2.$$

Proof. We have $w_{t+1} = w_t - \hat{\mathbf{p}}_t$, $\Delta_t = w_t - w^*$ and $\Delta_{t+1} = w_{t+1} - w^*$. Also $\hat{\mathbf{p}}_t = w_t - w_{t+1} = \Delta_t - \Delta_{t+1}$. From the definition of ϕ we have,

$$\begin{aligned}\phi_t(\hat{\mathbf{p}}_t) &= \frac{1}{2} (\Delta_t - \Delta_{t+1})^\top \mathbf{H}_t (\Delta_t - \Delta_{t+1}) - (\Delta_t - \Delta_{t+1})^\top \mathbf{g}_t, \\ (1 - \zeta^2) \phi_t\left(\frac{1}{1 - \zeta^2} \Delta_t\right) &= \frac{1}{2(1 - \zeta^2)} \Delta_t^\top \mathbf{H}_t \Delta_t - \Delta_t^\top \mathbf{g}_t.\end{aligned}$$

From the above two equation we have

$$\begin{aligned}\phi_t(\hat{\mathbf{p}}_t) - (1 - \zeta^2) \phi_t\left(\frac{1}{1 - \zeta^2} \Delta_t\right) &= \frac{1}{2} \Delta_{t+1}^\top \mathbf{H}_t \Delta_{t+1} - \frac{1}{2} \Delta_t^\top \mathbf{H}_t \Delta_{t+1} + \frac{1}{2} \Delta_{t+1}^\top \mathbf{g}_t - \frac{\zeta^2}{2(1 - \zeta^2)} \Delta_t^\top \mathbf{H}_t \Delta_t.\end{aligned}$$

From Lemma 14 the following holds

$$\begin{aligned}\phi_t(\hat{\mathbf{p}}_t) &\leq \epsilon^2 + (1 - \zeta^2) \min_{\mathbf{p}} \phi_t(\mathbf{p}) \\ &\leq \epsilon^2 + (1 - \zeta^2) \phi_t\left(\frac{1}{1 - \zeta^2} \Delta_t\right).\end{aligned}$$

So we have

$$\frac{1}{2} \Delta_{t+1}^\top \mathbf{H}_t \Delta_{t+1} - \Delta_t^\top \mathbf{H}_t \Delta_{t+1} + \Delta_{t+1}^\top \mathbf{g}_t - \frac{\zeta^2}{2(1 - \zeta^2)} \Delta_t^\top \mathbf{H}_t \Delta_t \leq \epsilon^2. \quad (4.17)$$

Consider $\mathbf{g}_t = \mathbf{g}(w_t)$

$$\begin{aligned}\mathbf{g}(w_t) &= \mathbf{g}(w^*) + \left(\int_0^1 \nabla^2 f(w^* + z(w_t - w^*)) dz \right) (w_t - w^*) \\ &= \left(\int_0^1 \nabla^2 f(w^* + z(w_t - w^*)) dz \right) \Delta_t \quad (\text{as } \mathbf{g}(w^*) = 0).\end{aligned}$$

Now we bound the following

$$\begin{aligned}
\|\mathbf{H}_t \Delta_t - \mathbf{g}(w_t)\| &\leq \|\Delta_t\| \left\| \int_0^1 [\nabla^2 f(w_t) - \nabla^2 f(w^* + z(w_t - w^*))] dz \right\| \\
&\leq \|\Delta_t\| \int_0^1 \|\nabla^2 f(w_t) - \nabla^2 f(w^* + z(w_t - w^*))\| dz \quad (\text{By Jensen}) \\
&\leq \|\Delta_t\| \int_0^1 (1-z)L \|w_t - w^*\| dz \quad (\text{by } L\text{-Lipschitz assumption}) \\
&= \frac{L}{2} \|\Delta_t\|^2.
\end{aligned}$$

Plugging it into (4.17) we have

$$\begin{aligned}
\Delta_{t+1}^\top \mathbf{H}_t \Delta_{t+1} &\leq 2\Delta_{t+1}^\top (\mathbf{H}_t \Delta_t - \mathbf{g}_t) + \frac{\zeta^2}{(1-\zeta^2)} \Delta_t^\top \mathbf{H}_t \Delta_t + 2\epsilon^2 \\
&\leq 2\|\Delta_{t+1}\| \|\mathbf{H}_t \Delta_t - \mathbf{g}_t\| + \frac{\zeta^2}{(1-\zeta^2)} \Delta_t^\top \mathbf{H}_t \Delta_t + 2\epsilon^2 \\
&\leq L\|\Delta_{t+1}\| \|\Delta_t\|^2 + \frac{\zeta^2}{(1-\zeta^2)} \Delta_t^\top \mathbf{H}_t \Delta_t + 2\epsilon^2.
\end{aligned}$$

□

Proof of Theorem 6

Proof. From the Lemma 15 with probability $1 - \rho$

$$\begin{aligned}
\Delta_{t+1}^\top \mathbf{H}_t \Delta_{t+1} &\leq L\|\Delta_{t+1}\| \|\Delta_t\|^2 + \frac{\zeta^2}{(1-\zeta^2)} \Delta_t^\top \mathbf{H}_t \Delta_t + 2\epsilon^2 \\
&\leq L\|\Delta_{t+1}\| \|\Delta_t\|^2 + \left(\frac{\zeta^2}{1-\zeta^2} \sigma_{\max}(\mathbf{H}_t)\right) \|\Delta_t\|^2 + 2\epsilon^2.
\end{aligned}$$

So we have,

$$\|\Delta_{t+1}\| \leq \max\left\{ \sqrt{\frac{\sigma_{\max}(\mathbf{H}_t)}{\sigma_{\min}(\mathbf{H}_t)} \left(\frac{\zeta^2}{1-\zeta^2}\right)} \|\Delta_t\|, \frac{L}{\sigma_{\min}(\mathbf{H}_t)} \|\Delta_t\|^2 \right\} + \frac{2\epsilon}{\sqrt{\sigma_{\min}(\mathbf{H}_t)}}.$$

□

4.9 Analysis of Section 4.4

In this section we provide the theoretical analysis of the Byzantine robust method explained in Section 4.4 and prove the statistical guarantee. In any iteration t the following holds

$$\begin{aligned}
|\mathcal{U}_t| &= |(\mathcal{U}_t \cap \mathcal{M}_t)| + |(\mathcal{U}_t \cap \mathcal{B}_t)| \\
|\mathcal{M}_t| &= |(\mathcal{U}_t \cap \mathcal{M}_t)| + |(\mathcal{M}_t \cap \mathcal{T}_t)|.
\end{aligned}$$

Combining both we have

$$|\mathcal{U}_t| = |\mathcal{M}_t| - |(\mathcal{M}_t \cap \mathcal{T}_t)| + |(\mathcal{U}_t \cap \mathcal{B}_t)|.$$

Lemma 16. *Let $\{\mathbf{S}_i\}_{i=1}^m \in \mathbb{R}^{n \times s}$ be sketching matrices based on Lemma 11. Let ϕ_t be defined in (4.10) and $\hat{\mathbf{p}}_t$ be defined in Algorithm 4. It holds that*

$$\min_{\mathbf{p}} \phi_t(\mathbf{p}) \leq \phi_t(\hat{\mathbf{p}}_t) \leq \epsilon_{byz}^2 + (1 - \zeta_{byz}^2)\phi(\mathbf{p}^*),$$

where ϵ_{byz} and ζ_{byz} is defined in (4.5) and (4.6) respectively.

Proof. In the following analysis we omit the subscript 't'. From the definition of the quadratic function (4.10) we know that

$$\phi(\hat{\mathbf{p}}) - \phi(\mathbf{p}^*) = \frac{1}{2} \|\mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}} - \mathbf{p}^*)\|^2.$$

Now we consider

$$\begin{aligned} \frac{1}{2} \|\mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}} - \mathbf{p}^*)\|^2 &= \frac{1}{2} \|\mathbf{H}^{\frac{1}{2}}\left(\frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{U}} \hat{\mathbf{p}}_i - \mathbf{p}^*\right)\|^2 \\ &= \frac{1}{2} \|\mathbf{H}^{\frac{1}{2}} \frac{1}{|\mathcal{U}|} \left(\sum_{i \in \mathcal{M}} (\hat{\mathbf{p}}_i - \mathbf{p}^*) - \sum_{i \in (\mathcal{M} \cap \mathcal{T})} (\hat{\mathbf{p}}_i - \mathbf{p}^*) + \sum_{i \in (\mathcal{U} \cap \mathcal{B})} (\hat{\mathbf{p}}_i - \mathbf{p}^*) \right)\|^2 \\ &\leq \underbrace{\|\mathbf{H}^{\frac{1}{2}} \frac{1}{|\mathcal{U}|} \left(\sum_{i \in \mathcal{M}} (\hat{\mathbf{p}}_i - \mathbf{p}^*) \right)\|^2}_{Term1} + 2 \underbrace{\|\mathbf{H}^{\frac{1}{2}} \frac{1}{|\mathcal{U}|} \sum_{i \in (\mathcal{M} \cap \mathcal{T})} (\hat{\mathbf{p}}_i - \mathbf{p}^*)\|^2}_{Term2} \\ &\quad + 2 \underbrace{\|\mathbf{H}^{\frac{1}{2}} \frac{1}{|\mathcal{U}|} \sum_{i \in (\mathcal{U} \cap \mathcal{B})} (\hat{\mathbf{p}}_i - \mathbf{p}^*)\|^2}_{Term3}. \end{aligned}$$

Now we bound each term separately and use the result of the Lemma 14 to bound each term.

$$\begin{aligned} Term1 &= \|\mathbf{H}^{\frac{1}{2}} \frac{1}{|\mathcal{U}|} \left(\sum_{i \in \mathcal{M}} (\hat{\mathbf{p}}_i - \mathbf{p}^*) \right)\|^2 \\ &= \left(\frac{1 - \alpha}{1 - \beta} \right)^2 \|\mathbf{H}^{\frac{1}{2}} \frac{1}{|\mathcal{M}|} \left(\sum_{i \in \mathcal{M}} (\hat{\mathbf{p}}_i - \mathbf{p}^*) \right)\|^2 \\ &\leq \left(\frac{1 - \alpha}{1 - \beta} \right)^2 [\epsilon^2 + \zeta_{\mathcal{M}}^2 \|\mathbf{H}^{\frac{1}{2}} \mathbf{p}^*\|^2], \end{aligned}$$

where $\zeta_{\mathcal{M}} = \nu\left(\frac{\eta}{\sqrt{|\mathcal{M}|}} + \frac{\eta^2}{1 - \eta}\right) = \nu\left(\frac{\eta}{\sqrt{(1 - \alpha)m}} + \frac{\eta^2}{1 - \eta}\right)$.

Similarly the Term 2 can be bonded as it is a bound on good machines

$$\begin{aligned}
Term2 &= 2\|\mathbf{H}^{\frac{1}{2}}\frac{1}{|\mathcal{U}|}\sum_{i\in(\mathcal{M}\cap\mathcal{T})}(\hat{\mathbf{p}}_i - \mathbf{p}^*)\|^2 \\
&= 2\left(\frac{1-\alpha}{1-\beta}\right)^2\|\mathbf{H}^{\frac{1}{2}}\frac{1}{|\mathcal{M}\cap\mathcal{T}|}\sum_{i\in(\mathcal{M}\cap\mathcal{T})}(\hat{\mathbf{p}}_i - \mathbf{p}^*)\|^2 \\
&\leq 2\left(\frac{1-\alpha}{1-\beta}\right)^2[\epsilon^2 + \zeta_{\mathcal{M}\cap\mathcal{T}}^2\|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2],
\end{aligned}$$

where $\zeta_{\mathcal{M}\cap\mathcal{T}} = \nu\left(\frac{\eta}{\sqrt{|\mathcal{M}\cap\mathcal{T}|}} + \frac{\eta^2}{1-\eta}\right) \leq \nu\left(\frac{\eta}{\sqrt{(1-\beta)m}} + \frac{\eta^2}{1-\eta}\right)$.

For the Term 3 we know that $\beta > \alpha$ so all the untrimmed worker norm is bounded by a good machine as at least one good machine gets trimmed.

$$\begin{aligned}
Term3 &= 2\|\mathbf{H}^{\frac{1}{2}}\frac{1}{|\mathcal{U}|}\sum_{i\in(\mathcal{U}\cap\mathcal{B})}(\hat{\mathbf{p}}_i - \mathbf{p}^*)\|^2 \\
&\leq 2\sigma_{max}(\mathbf{H})\left(\frac{|\mathcal{U}\cap\mathcal{B}|}{|\mathcal{U}|}\right)^2\|\frac{1}{|\mathcal{U}\cap\mathcal{B}|}\sum_{i\in(\mathcal{U}\cap\mathcal{B})}(\hat{\mathbf{p}}_i - \mathbf{p}^*)\|^2 \\
&\leq 2\sigma_{max}(\mathbf{H})\left(\frac{|\mathcal{U}\cap\mathcal{B}|}{|\mathcal{U}|}\right)^2\frac{1}{|\mathcal{U}\cap\mathcal{B}|}\sum_{i\in(\mathcal{U}\cap\mathcal{B})}\|(\hat{\mathbf{p}}_i - \mathbf{p}^*)\|^2 \\
&\leq 4\sigma_{max}(\mathbf{H})\left(\frac{|\mathcal{U}\cap\mathcal{B}|}{|\mathcal{U}|}\right)^2\frac{1}{|\mathcal{U}\cap\mathcal{B}|}\sum_{i\in(\mathcal{U}\cap\mathcal{B})}(\|\hat{\mathbf{p}}_i\|^2 + \|\mathbf{p}^*\|^2) \\
&\leq 4\sigma_{max}(\mathbf{H})\left(\frac{|\mathcal{U}\cap\mathcal{B}|}{|\mathcal{U}|}\right)^2\max_{i\in\mathcal{M}}(\|\hat{\mathbf{p}}_i\|^2 + \|\mathbf{p}^*\|^2) \\
&\leq 4\sigma_{max}(\mathbf{H})\left(\frac{|\mathcal{U}\cap\mathcal{B}|}{|\mathcal{U}|}\right)^2\max_{i\in\mathcal{M}}(\|\hat{\mathbf{p}}_i - \mathbf{p}^*\|^2 + 2\|\mathbf{p}^*\|^2) \\
&\leq 4\kappa\left(\frac{|\mathcal{U}\cap\mathcal{B}|}{|\mathcal{U}|}\right)^2\max_{i\in\mathcal{M}}(\|\mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}}_i - \mathbf{p}^*)\|^2 + 2\|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2) \\
&\leq 4\kappa\left(\frac{|\mathcal{U}\cap\mathcal{B}|}{|\mathcal{U}|}\right)^2(\epsilon^2 + (2 + \zeta_1^2)\|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2) \\
&\leq 4\kappa\left(\frac{\alpha}{1-\beta}\right)^2(\epsilon^2 + (2 + \zeta_1^2)\|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2),
\end{aligned}$$

where $\zeta_1 = \nu\left(\eta + \frac{\eta^2}{1-\eta}\right) = \frac{\nu}{1-\eta}$ and $\kappa = \frac{\sigma_{max}(\mathbf{H})}{\sigma_{min}(\mathbf{H})}$.

Combining all the bounds on Term1 , Term2 and Term3 we have

$$\frac{1}{2}\|\mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}} - \mathbf{p}^*)\|^2 \leq \epsilon_{byz}^2 + \zeta_{byz}^2\|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2,$$

where

$$\begin{aligned}\epsilon_{byz}^2 &= \left(3 \left(\frac{1-\alpha}{1-\beta} \right)^2 + 4\kappa \left(\frac{\alpha}{1-\beta} \right)^2 \right) \epsilon^2, \\ \zeta_{byz}^2 &= 2 \left(\frac{1-\alpha}{1-\beta} \right)^2 \zeta_{\mathcal{M} \cap \mathcal{T}}^2 + \left(\frac{1-\alpha}{1-\beta} \right)^2 \zeta_{\mathcal{M}}^2 + 4\kappa \left(\frac{\alpha}{1-\beta} \right)^2 (2 + \zeta_1^2).\end{aligned}$$

Finally we have

$$\begin{aligned}\phi(\hat{\mathbf{p}}) - \phi(\mathbf{p}^*) &\leq \epsilon_{byz}^2 - \zeta_{byz}^2 \phi(\mathbf{p}^*) \\ \Rightarrow \phi(\hat{\mathbf{p}}) &\leq \epsilon_{byz}^2 + (1 - \zeta_{byz}^2) \phi(\mathbf{p}^*).\end{aligned}$$

□

Lemma 17. Let $\zeta_{byz} \in (0, 1)$, ϵ_{byz} be any fixed parameter. And $\hat{\mathbf{p}}_t$ satisfies $\phi_t(\hat{\mathbf{p}}_t) \leq \epsilon_{byz}^2 + (1 - \zeta_{byz}^2) \min_{\mathbf{p}} \phi_t(\mathbf{p})$. Under the Assumption 15 (Hessian L-Lipschitz) and $\Delta_t = w_t - w^*$ satisfies

$$\Delta_{t+1}^T \mathbf{H}_t \Delta_{t+1} \leq L \|\Delta_{t+1}\| \|\Delta_t\|^2 + \frac{\zeta_{byz}^2}{1 - \zeta_{byz}^2} \Delta_t^T \mathbf{H}_t \Delta_t + 2\epsilon_{byz}^2.$$

Proof. We choose $\zeta = \zeta_{byz}$ and $\epsilon = \epsilon_{byz}$ from the Lemma 16 and follow the proof of Lemma 15 to obtain the desired bound. □

Proof of Theorem 7

Proof. We get the desired bound by developing from the result of the Lemma 17 and following the proof of Theorem 6 □

4.10 Auxiliary Lemmas

Lemma 18. Let $\{\mathbf{S}_i\}_{i=1}^m \in \mathbb{R}^{n \times s}$ be sketching matrices satisfies the Lemma 11. Let ϕ_t be defined in (4.10) and $\hat{\mathbf{p}}_t$ be defined in Algorithm 4. It holds that

$$\min_{\mathbf{p}} \phi_t(\mathbf{p}) \leq \phi_t(\hat{\mathbf{p}}_t) \leq \epsilon_{byz}^2 + (1 - \alpha_{byz}^2) \phi(\mathbf{p}^*)$$

where

Proof. In the following analysis we omit the subscript 't'. From the definition of the quadratic function (4.10) we know that

$$\begin{aligned}\phi(\mathcal{Q}(\hat{\mathbf{p}})) - \phi(\mathbf{p}^*) &= \frac{1}{2} \|\mathbf{H}^{\frac{1}{2}}(\mathcal{Q}(\hat{\mathbf{p}}) - \mathbf{p}^*)\|^2 \\ &\leq \underbrace{\|\mathbf{H}^{\frac{1}{2}}(\mathcal{Q}(\hat{\mathbf{p}}) - \hat{\mathbf{p}})\|^2}_{\text{Term1}} + \underbrace{\|\mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}} - \mathbf{p}^*)\|^2}_{\text{Term2}}\end{aligned}$$

First we bound the Term 1.

$$\begin{aligned}
\|\mathbf{H}^{\frac{1}{2}}(\mathcal{Q}(\hat{\mathbf{p}}) - \hat{\mathbf{p}})\|^2 &\leq \sigma_{max}(\mathbf{H}_t)(1 - \rho)\|\hat{\mathbf{p}}\|^2 \\
&\leq \sigma_{max}(\mathbf{H}_t)(1 - \rho)[\|\hat{\mathbf{p}} - \mathbf{p}^*\|^2 + \|\mathbf{p}^*\|^2] \\
&\leq \frac{\sigma_{max}(\mathbf{H}_t)}{\sigma_{min}(\mathbf{H}_t)}(1 - \rho)[\|\mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}} - \mathbf{p}^*)\|^2 + \|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2] \\
&= \kappa(1 - \rho)[\|\mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}} - \mathbf{p}^*)\|^2 + \|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2]
\end{aligned}$$

Now plugging back the bound of Term 1, we get

$$\begin{aligned}
\phi(\mathcal{Q}(\hat{\mathbf{p}})) - \phi(\mathbf{p}^*) &\leq \kappa(1 - \rho)[\|\mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}} - \mathbf{p}^*)\|^2 + \|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2] + \|\mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}} - \mathbf{p}^*)\|^2 \\
&= (1 + \kappa(1 - \rho))[\|\mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}} - \mathbf{p}^*)\|^2] + \kappa(1 - \rho)\|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2
\end{aligned}$$

Now we use Lemma 14 to bound the term $\|\mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}} - \mathbf{p}^*)\|^2$ and we get,

$$\begin{aligned}
\phi(\mathcal{Q}(\hat{\mathbf{p}})) - \phi(\mathbf{p}^*) &\leq (1 + \kappa(1 - \rho))(\epsilon^2 + \alpha^2\|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2) + \kappa(1 - \rho)\|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2 \\
&= (1 + \kappa(1 - \rho))\epsilon^2 + [(1 + \kappa(1 - \rho))\alpha^2 + \kappa(1 - \rho)]\|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2 \\
\Rightarrow \phi(\mathcal{Q}(\hat{\mathbf{p}})) &\leq (1 + \kappa(1 - \rho))\epsilon^2 + (1 - [(1 + \kappa(1 - \rho))\alpha^2 + \kappa(1 - \rho)])\phi(\mathbf{p}^*) \\
&= \epsilon_b^2 + (1 - \alpha_b^2)\phi(\mathbf{p}^*)
\end{aligned}$$

where

$$\begin{aligned}
\epsilon_b^2 &= (1 + \kappa(1 - \rho))\epsilon^2 \\
\alpha_b^2 &= (1 + \kappa(1 - \rho))\alpha^2 + \kappa(1 - \rho)
\end{aligned}$$

□

Lemma 19. $\Delta_t = w_t - w^*$ satisfies

$$\Delta_{t+1}^T \mathbf{H}_t \Delta_{t+1} \leq L\|\Delta_{t+1}\|\|\Delta_t\|^2 + \frac{\alpha_b^2}{1 - \alpha_b^2} \Delta_t^T \mathbf{H}_t \Delta_t + 2\epsilon_b^2$$

Theorem 9. Let $\mu_t \in [1, \frac{n}{d}]$ be the coherence of \mathbf{A}_t and m be the number of partitions. For some $\eta, \rho \in (0, 1)$, with probability $1 - \rho$

$$\|\Delta_{t+1}\| \leq \max\left\{\sqrt{\frac{\sigma_{max}(\mathbf{H}_t)}{\sigma_{min}(\mathbf{H}_t)}}\left(\frac{\alpha_b^2}{1 - \alpha_b^2}\|\Delta_t\|, \frac{L}{\sigma_{min}(\mathbf{H}_t)}\|\Delta_t\|^2\right)\right\} + \frac{\epsilon_b}{\sqrt{\sigma_{min}(\mathbf{H}_t)}}$$

4.11 Analysis of Section 4.5

First we prove the following lemma that will be useful in our subsequent calculations. Consider that $\mathcal{Q}(\hat{\mathbf{p}}) = \frac{1}{|B|} \sum_{i \in B} \mathcal{Q}(\hat{\mathbf{p}}_i)$. And also we use the following notation $\zeta_B = \nu \left(\frac{\eta}{\sqrt{|B|}} + \frac{\eta^2}{1-\eta} \right)$, $\nu = \frac{\sigma_{\max}(\mathbf{A}^\top \mathbf{A})}{\sigma_{\max}(\mathbf{A}^\top \mathbf{A}) + n\lambda} \leq 1$.

Lemma 20. *If $\mathcal{Q}(\hat{\mathbf{p}}_i)$ is the local update direction and \mathbf{p}^* is the optimal solution to the quadratic function ϕ then*

$$\left\| \mathbf{H}^{\frac{1}{2}}(\mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*) \right\|^2 \leq 1 + \kappa(1 - \delta)\epsilon^2 + (\zeta_B^2 + \kappa(1 - \delta)((1 + \zeta_1^2))) \left\| \mathbf{H}^{\frac{1}{2}}\mathbf{p}^* \right\|^2,$$

where \mathbf{H} is the exact Hessian and

$$\begin{aligned} \epsilon_1 &= \sqrt{(1 + \kappa(1 - \delta))\epsilon}, \\ \zeta_{\text{comp}, B}^2 &= (\zeta_B^2 + \kappa(1 - \delta)((1 + \zeta_1^2))). \end{aligned}$$

ϵ is defined in equation (4.4) and

Proof.

$$\begin{aligned} \left\| \mathbf{H}^{\frac{1}{2}}(\mathcal{Q}(\hat{\mathbf{p}}) - \mathbf{p}^*) \right\|^2 &= \left\| \mathbf{H}^{\frac{1}{2}}(\mathcal{Q}(\hat{\mathbf{p}}) - \hat{\mathbf{p}} + \hat{\mathbf{p}} - \mathbf{p}^*) \right\|^2 \\ &\leq 2 \left(\underbrace{\left\| \mathbf{H}^{\frac{1}{2}}(\mathcal{Q}(\hat{\mathbf{p}}) - \hat{\mathbf{p}}) \right\|^2}_{\text{Term1}} + \underbrace{\left\| \mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}} - \mathbf{p}^*) \right\|^2}_{\text{Term2}} \right). \end{aligned} \quad (4.18)$$

Following the proof of Lemma 14 we get

$$\left\| \mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}}_i - \mathbf{p}^*) \right\|^2 \leq \epsilon^2 + \zeta_1 \left\| \mathbf{H}^{\frac{1}{2}}\mathbf{p}^* \right\|^2, \quad (4.19)$$

where ϵ is as defined in (4.4). Now we consider the term

$$\begin{aligned} \left\| \mathbf{H}^{\frac{1}{2}}(\mathcal{Q}(\hat{\mathbf{p}}_i) - \hat{\mathbf{p}}_i) \right\|^2 &\leq \sigma_{\max}(\mathbf{H})(1 - \delta) \|\hat{\mathbf{p}}_i\|^2 \\ &\leq \sigma_{\max}(\mathbf{H})(1 - \delta) (\|\hat{\mathbf{p}}_i - \mathbf{p}^*\|^2 + \|\mathbf{p}^*\|^2) \\ &\leq \frac{\sigma_{\max}}{\sigma_{\min}}(1 - \delta) \left(\left\| \mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}}_i - \mathbf{p}^*) \right\|^2 + \left\| \mathbf{H}^{\frac{1}{2}}\mathbf{p}^* \right\|^2 \right) \\ &= \kappa(1 - \delta) \left(\left\| \mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}}_i - \mathbf{p}^*) \right\|^2 + \left\| \mathbf{H}^{\frac{1}{2}}\mathbf{p}^* \right\|^2 \right) \\ &\leq \kappa(1 - \delta) \left(\epsilon^2 + (1 + \zeta_1^2) \left\| \mathbf{H}^{\frac{1}{2}}\mathbf{p}^* \right\|^2 \right) \quad \text{Using (4.19)}. \end{aligned}$$

Now we use the above calculation and bound Term1

$$\begin{aligned} \left\| \mathbf{H}^{\frac{1}{2}}(\mathcal{Q}(\hat{\mathbf{p}}) - \hat{\mathbf{p}}) \right\|^2 &\leq \frac{1}{|B|} \sum_{i \in B} \left\| \mathbf{H}^{\frac{1}{2}}(\mathcal{Q}(\hat{\mathbf{p}}_i) - \hat{\mathbf{p}}_i) \right\|^2 \\ &\leq \kappa(1 - \delta) \left(\epsilon^2 + (1 + \zeta_1^2) \left\| \mathbf{H}^{\frac{1}{2}} \mathbf{p}^* \right\|^2 \right). \end{aligned} \quad (4.20)$$

We can bound the Term2 directly using the proof of Lemma 14

$$\left\| \mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}} - \mathbf{p}^*) \right\|^2 \leq \epsilon^2 + \zeta_B^2 \left\| \mathbf{H}^{\frac{1}{2}} \mathbf{p}^* \right\|^2. \quad (4.21)$$

Now we use (4.20) and (4.21) and plug them in (4.18)

$$\left\| \mathbf{H}^{\frac{1}{2}}(\mathcal{Q}(\hat{\mathbf{p}}) - \mathbf{p}^*) \right\|^2 \leq (1 + \kappa(1 - \delta))\epsilon^2 + (\zeta_B^2 + \kappa(1 - \delta)((1 + \zeta_1^2)) \left\| \mathbf{H}^{\frac{1}{2}} \mathbf{p}^* \right\|^2.$$

Now we define

$$\begin{aligned} \epsilon_1 &= \sqrt{(1 + \kappa(1 - \delta))}\epsilon \\ \zeta_{comp,B}^2 &= (\zeta_B^2 + \kappa(1 - \delta)((1 + \zeta_1^2))). \end{aligned}$$

□

Now we have the robust update in iteration t to be $\mathcal{Q}(\hat{\mathbf{p}}) = \frac{1}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} \mathcal{Q}(\hat{\mathbf{p}}_{i,t})$.

Lemma 21. *Let $\{\mathbf{S}_i\}_{i=1}^m \in \mathbb{R}^{n \times s}$ be sketching matrices based on Lemma 11. Let ϕ_t be defined in (4.10) and $\mathcal{Q}(\hat{\mathbf{p}}_t)$ be the update with \mathcal{Q} being δ -approximate compressor. It holds that*

$$\min_{\mathbf{p}} \phi_t(\mathbf{p}) \leq \phi_t(\mathcal{Q}(\hat{\mathbf{p}}_t)) \leq \epsilon_{comp,byz}^2 + (1 - \zeta_{comp,byz}^2)\phi_t(\mathbf{p}^*),$$

where $\epsilon_{comp,byz}$ and $\zeta_{comp,byz}^2$ is as defined in (4.7) and (4.8) respectively.

Proof. In the following analysis we omit the subscript ' t '. From the definition of the quadratic function (4.10) we know that

$$\phi(\mathcal{Q}(\hat{\mathbf{p}})) - \phi(\mathbf{p}^*) = \frac{1}{2} \left\| \mathbf{H}^{\frac{1}{2}}(\mathcal{Q}(\hat{\mathbf{p}}) - \mathbf{p}^*) \right\|^2.$$

Now we consider

$$\begin{aligned}
\frac{1}{2}\|\mathbf{H}^{\frac{1}{2}}(\mathcal{Q}(\hat{\mathbf{p}}) - \mathbf{p}^*)\|^2 &= \frac{1}{2}\|\mathbf{H}^{\frac{1}{2}}\left(\frac{1}{|\mathcal{U}|}\sum_{i \in \mathcal{U}} \mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*\right)\|^2 \\
&= \frac{1}{2}\|\mathbf{H}^{\frac{1}{2}}\frac{1}{|\mathcal{U}|}\left(\sum_{i \in \mathcal{M}} (\mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*) - \sum_{i \in (\mathcal{M} \cap \mathcal{T})} (\mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*)\right. \\
&\quad \left. + \sum_{i \in (\mathcal{U} \cap \mathcal{B})} (\mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*)\right)\|^2 \\
&\leq \underbrace{\|\mathbf{H}^{\frac{1}{2}}\frac{1}{|\mathcal{U}|}\left(\sum_{i \in \mathcal{M}} (\mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*)\right)\|^2}_{Term1} + 2\underbrace{\|\mathbf{H}^{\frac{1}{2}}\frac{1}{|\mathcal{U}|}\sum_{i \in (\mathcal{M} \cap \mathcal{T})} (\mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*)\|^2}_{Term2} \\
&\quad + 2\underbrace{\|\mathbf{H}^{\frac{1}{2}}\frac{1}{|\mathcal{U}|}\sum_{i \in (\mathcal{U} \cap \mathcal{B})} (\mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*)\|^2}_{Term3}.
\end{aligned}$$

Now we bound each term separately and use the Lemma 20

$$\begin{aligned}
Term1 &= \|\mathbf{H}^{\frac{1}{2}}\frac{1}{|\mathcal{U}|}\left(\sum_{i \in \mathcal{M}} (\mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*)\right)\|^2 \\
&= \left(\frac{1-\alpha}{1-\beta}\right)^2 \|\mathbf{H}^{\frac{1}{2}}\frac{1}{|\mathcal{M}|}\left(\sum_{i \in \mathcal{M}} (\mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*)\right)\|^2 \\
&\leq \left(\frac{1-\alpha}{1-\beta}\right)^2 [\epsilon_1^2 + \zeta_{comp, \mathcal{M}}^2 \|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2],
\end{aligned}$$

where $\zeta_{comp, \mathcal{M}}^2 = (\zeta_{\mathcal{M}}^2 + \kappa(1-\delta)((1+\zeta_1^2))$. Similarly the Term 2 can be bonded as it is a bound on good machines

$$\begin{aligned}
Term2 &= 2\|\mathbf{H}^{\frac{1}{2}}\frac{1}{|\mathcal{U}|}\sum_{i \in (\mathcal{M} \cap \mathcal{T})} (\mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*)\|^2 \\
&= 2\left(\frac{1-\alpha}{1-\beta}\right)^2 \|\mathbf{H}^{\frac{1}{2}}\frac{1}{|\mathcal{M} \cap \mathcal{T}|}\sum_{i \in (\mathcal{M} \cap \mathcal{T})} (\mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*)\|^2 \\
&\leq 2\left(\frac{1-\alpha}{1-\beta}\right)^2 [\epsilon_1^2 + \zeta_{comp, \mathcal{M} \cap \mathcal{T}}^2 \|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2].
\end{aligned}$$

For the Term 3 we know that $\beta > \alpha$ so all the untrimmed worker norm is bounded by a good

machine as at least one good machine gets trimmed.

$$\begin{aligned}
Term3 &= 2\|\mathbf{H}^{\frac{1}{2}}\frac{1}{|\mathcal{U}|}\sum_{i\in(\mathcal{U}\cap\mathcal{B})}(\mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*)\|^2 \\
&\leq 2\sigma_{max}(\mathbf{H})\left(\frac{|\mathcal{U}\cap\mathcal{B}|}{|\mathcal{U}|}\right)^2\left\|\frac{1}{|\mathcal{U}\cap\mathcal{B}|}\sum_{i\in(\mathcal{U}\cap\mathcal{B})}(\mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*)\right\|^2 \\
&\leq 2\sigma_{max}(\mathbf{H})\left(\frac{|\mathcal{U}\cap\mathcal{B}|}{|\mathcal{U}|}\right)^2\frac{1}{|\mathcal{U}\cap\mathcal{B}|}\sum_{i\in(\mathcal{U}\cap\mathcal{B})}\|(\mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*)\|^2 \\
&\leq 4\sigma_{max}(\mathbf{H})\left(\frac{|\mathcal{U}\cap\mathcal{B}|}{|\mathcal{U}|}\right)^2\frac{1}{|\mathcal{U}\cap\mathcal{B}|}\sum_{i\in(\mathcal{U}\cap\mathcal{B})}(\|\mathcal{Q}(\hat{\mathbf{p}}_i)\|^2 + \|\mathbf{p}^*\|^2) \\
&\leq 4\sigma_{max}(\mathbf{H})\left(\frac{|\mathcal{U}\cap\mathcal{B}|}{|\mathcal{U}|}\right)^2\max_{i\in\mathcal{M}}(\|\mathcal{Q}(\hat{\mathbf{p}}_i)\|^2 + \|\mathbf{p}^*\|^2) \\
&\leq 4\sigma_{max}(\mathbf{H})\left(\frac{|\mathcal{U}\cap\mathcal{B}|}{|\mathcal{U}|}\right)^2\max_{i\in\mathcal{M}}(\|\mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*\|^2 + 2\|\mathbf{p}^*\|^2) \\
&\leq 4\kappa\left(\frac{|\mathcal{U}\cap\mathcal{B}|}{|\mathcal{U}|}\right)^2\max_{i\in\mathcal{M}}(\|\mathbf{H}^{\frac{1}{2}}(\mathcal{Q}(\hat{\mathbf{p}}_i) - \mathbf{p}^*)\|^2 + 2\|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2) \\
&\leq 4\kappa\left(\frac{|\mathcal{U}\cap\mathcal{B}|}{|\mathcal{U}|}\right)^2(\epsilon_1^2 + (2 + \zeta_1^2)\|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2) \\
&\leq 4\kappa\left(\frac{\alpha}{1-\beta}\right)^2(\epsilon_1^2 + (2 + \zeta_1^2)\|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2).
\end{aligned}$$

Combining all the bounds on Term1 , Term2 and Term3 we have

$$\frac{1}{2}\|\mathbf{H}^{\frac{1}{2}}(\hat{\mathbf{p}} - \mathbf{p}^*)\|^2 \leq \epsilon_{byz}^2 + \zeta_{byz}^2\|\mathbf{H}^{\frac{1}{2}}\mathbf{p}^*\|^2,$$

where

$$\begin{aligned}
\epsilon_{comp,byz}^2 &= \left(3\left(\frac{1-\alpha}{1-\beta}\right)^2 + 4\kappa\left(\frac{\alpha}{1-\beta}\right)^2\right)\epsilon_1^2 \\
\zeta_{comp,byz}^2 &= 2\left(\frac{1-\alpha}{1-\beta}\right)^2\zeta_{comp,\mathcal{M}\cap\mathcal{T}}^2 + \left(\frac{1-\alpha}{1-\beta}\right)^2\zeta_{comp,\mathcal{M}}^2 + 4\kappa\left(\frac{\alpha}{1-\beta}\right)^2(2 + \zeta_{comp,1}^2).
\end{aligned}$$

Finally we have

$$\begin{aligned}
\phi(\hat{\mathbf{p}}) - \phi(\mathbf{p}^*) &\leq \epsilon_{comp,byz}^2 - \zeta_{comp,byz}^2\phi(\mathbf{p}^*) \\
\Rightarrow \phi(\hat{\mathbf{p}}) &\leq \epsilon_{comp,byz}^2 + (1 - \zeta_{comp,byz}^2)\phi(\mathbf{p}^*).
\end{aligned}$$

□

Lemma 22. Let $\zeta_{comp,byz} \in (0, 1)$, $\epsilon_{comp,byz}$ be any fixed parameter. And $\mathcal{Q}(\hat{p}_t)$ satisfies $\phi_t(\mathcal{Q}(\hat{p}_t)) \leq \epsilon_{byz}^2 + (1 - \zeta_{byz}^2) \min_{\mathbf{p}} \phi_t(\mathbf{p})$. Under the Assumption 15 (Hessian L -Lipschitz) and $\Delta_t = w_t - w^*$ satisfies

$$\Delta_{t+1}^T \mathbf{H}_t \Delta_{t+1} \leq L \|\Delta_{t+1}\| \|\Delta_t\|^2 + \frac{\zeta_{comp,byz}^2}{1 - \zeta_{comp,byz}^2} \Delta_t^T \mathbf{H}_t \Delta_t + 2\epsilon_{comp,byz}^2.$$

Proof. We choose $\zeta = \zeta_{comp,byz}$ and $\epsilon = \epsilon_{comp,byz}$ from the Lemma 21 and follow the proof of Lemma 15 to obtain the desired bound. \square

Proof of Theorem 8

Proof. We get the desired bound by developing from the result of the Lemma 22 and following the proof of Theorem 6 \square

Additional Experiment

In addition to the experimental results in Section 4.6, we provide some more experiments supporting the robustness of the COMRADE in two different types of attacks : 1. ‘Gaussian attack’: where the Byzantine workers add Gaussian Noise ($\mathcal{N}(\mu, \sigma^2)$) to the update and 2. ‘random label attack’: where the Byzantine worker machines learns based on random labels instead of proper labels.




Figure 4.3: (First row) Accuracy of COMRADE with 10%, 15%, 20% Byzantine workers with ‘Gaussian ’ attack for (a). w5a (b). a9a and ‘random label’ attack for (c). w5a (d).a9a. (Second row) Accuracy of COMRADE with ρ -approximate compressor (Section 4.5) with 10%, 15%, 20% Byzantine workers with ‘Gaussian ’ attack for (a). w5a (b). a9a and ‘random label’ attack for (c). w5a (d).a9a.

Chapter 5

Escaping Saddle Points in FL: Distributed Cubic Regularized Newton Method

We study the problem of optimizing a non-convex loss function (with saddle points) in a Federated Learning framework in the presence of Byzantine machines. Our proposed algorithm is a variant of the celebrated cubic-regularized Newton method of Nesterov and Polyak [43], which avoids saddle points efficiently and converges to local minima. Furthermore, our algorithm resists the presence of Byzantine machines, which may create *fake local minima* near the saddle points of the loss function, also known as saddle-point attack. We robustify the cubic-regularized Newton algorithm such that it avoids the saddle points and the fake local minimas efficiently. Furthermore, being a second order algorithm, the iteration complexity is much lower than its first order counterparts, and thus our algorithm communicates little with the parameter server. We obtain theoretical guarantees for our proposed scheme under several settings including approximate (sub-sampled) gradients and Hessians. Moreover, we validate our theoretical findings with experiments using standard datasets and several types of Byzantine attacks.

5.1 Introduction

In this chapter we propose provable and efficient algorithms for distributed learning that are communication efficient and Byzantine resilient. In particular, we focus on optimizing a non-convex loss function $f(\cdot)$ in a distributed optimization framework. We have m worker machines, out of which α fraction may behave in a Byzantine fashion, where $\alpha < \frac{1}{2}$. Optimizing a loss function in a distributed setup has gained a lot of attention in recent years [17, 19, 30, 33]. However, most of these approaches either work when $f(\cdot)$ is convex, or provide weak guarantees in the non-convex case (ex: zero gradient points, maybe a saddle point).

On the other hand, in order to fit complex machine learning models, one often requires to find local minima of a non-convex loss $f(\cdot)$, instead of critical points only, which may include several saddle points. Training deep neural networks and other high-capacity learning architectures [35, 36] are some of the examples where finding local minima is crucial. [36, 37] shows that the stationary

points of these problems are in fact saddle points and far away from any local minimum, and hence designing efficient algorithm that escapes saddle points is of interest. Moreover, in [38, 39], it is argued that saddle points can lead to highly sub-optimal solutions in many problems of interest. This issue is amplified in high dimension as shown in [40], and becomes the main bottleneck in training deep neural nets.

Furthermore, a line of recent work [39, 41, 42], shows that for many non-convex problems, it is sufficient to find a local minimum. In fact, in many problems of interest, all local minima are global minima (e.g., dictionary learning [42], phase retrieval [39], matrix sensing and completion [36, 41], and some of neural nets [37]). Also, in [137], it is argued that for more general neural nets, the local minima are as good as global minima.

The issue of saddle point avoidance becomes non-trivial in the presence of Byzantine workers. Since we do not assume anything on the behavior of the Byzantine workers, it is certainly conceivable that by appropriately modifying their messages to the center, they can create *fake local minima* that are close to the saddle point of the loss function $f(\cdot)$, and these are far away from the true local minima of $f(\cdot)$. This is popularly known as the *saddle-point attack* (see [31]), and it can arbitrarily destroy the performance of any non-robust learning algorithm. Hence, our goal is to design an algorithm that escapes saddle points of $f(\cdot)$ in an efficient manner as well as resists the saddle-point attack simultaneously. The complexity of such an algorithm emerges from the the interplay between non-convexity of the loss function and the behavior of the Byzantine machines.

The problem of saddle point avoidance in the context of non-convex optimization has received considerable attention in the past few years. In the seminal paper of Jin et al. [46], a gradient descent based approach is proposed. By defining a certain *perturbation condition* and adding Gaussian noise to the iterates of gradient descent, the algorithm of [46] provably escapes the saddle points of the non-convex loss function. A few papers [138, 139] following the above use various modifications to obtain saddle point avoidance guarantees. However, these algorithms are non-robust. A Byzantine robust saddle point avoidance algorithm is proposed by Yin et al. [31], and probably is the closest to this work. In [31], the authors propose a repeated check-and-escape type of first order gradient descent based algorithm. First of all, being a first order algorithm, the convergence rate is quite slow (the rate for gradient decay is $1/\sqrt{T}$, where T is the number of iterations). Moreover, implementation-wise, the algorithm presented in [31] is computation heavy, and takes potentially many iterations between the center and the worker machines. Hence, this algorithm is not efficient in terms of the communication cost.

In this work, we consider a variation of the famous cubic-regularized Newton algorithm of Nesterov and Polyak [43]. It is theoretically proved in [43] that a cubic-regularized Newton method with proper choice of parameters like step size always outperforms the gradient based first order schemes (like [31]) in all situations under consideration. Indeed, in Theorem 10, we observe that the rate of gradient decay is $\frac{1}{T^{2/3}}$, which is strictly better than the first order gradient based methods. In Section 5.6, we experimentally show that our scheme outperforms that of [31], in terms of iteration complexity and hence communication cost. Also, our algorithm is easy to implement whereas a range of hyper-parameter choice and tuning makes the implementation of ByzantinePGD [31] difficult.

In [43–45], it is shown that cubic-regularized Newton can efficiently escape the saddle points

of a non-convex function. Assuming the loss function has a Lipschitz continuous Hessian (see Assumption 19), the cubic-regularized Newton optimizes an auxiliary function (detailed in Section 5.3), which is an upper second order approximation of the original loss function. It is shown in [43] that the cubic regularized term in the auxiliary function pushes the Hessian towards a positive semi-definite matrix.

A point w is said to satisfy the ϵ -second order stationary condition of the loss function $f(\cdot)$ if,

$$\|\nabla f(w)\| \leq \epsilon \quad \lambda_{\min}(\nabla^2 f(w)) \geq -\sqrt{\epsilon}.$$

$\nabla f(w)$ denotes the gradient of the function and $\lambda_{\min}(\nabla^2 f(w))$ denotes the minimum eigenvalue of the Hessian of the function. Hence, under the assumption (which is standard in the literature, see [31, 46]) that all saddle points are strict (i.e., $\lambda_{\min}(\nabla^2 f(w_s)) < 0$ for any saddle point w_s), all second order stationary points (with $\epsilon = 0$) are local minima, and hence converging to a stationary point is equivalent to converging to a local minima.

We consider a distributed variant of the cubic regularized Newton algorithm. In this scheme, the center machine asks the workers to solve an auxiliary function and return the result. Note that the complexity of the problem is partially transferred to the worker machines. It is worth mentioning that in most distributed optimization paradigm, including Federated Learning, the workers possess sufficient compute power to handle this partial transfer of compute load, and in most cases, this is desirable [3]. The center machine aggregates the solution of the worker machines and takes a descent step. Note that, unlike gradient aggregation, the aggregation of the solutions of the local optimization problems is a highly non-linear operation. Hence, it is quite non-trivial to extend the centralized cubic regularized algorithm to a distributed one. The solution to the cubic regularization even lacks a closed form solution unlike the second order Hessian based update or the first order gradient based update. The analysis is carried out by leveraging the first order and second order stationary conditions of the auxiliary function solved in each worker machines.

In addition to this, we use a simple norm-based thresholding approach to robustify the distributed cubic-regularized Newton method. In [31], the authors use computation-heavy schemes like coordinate-wise median, trimmed mean and iterative filtering. In contrast to these approaches, our scheme is computationally efficient. Norm based thresholding is a standard trick for Byzantine resilience as featured in [140, 141]. However, since the local optimization problem lacks a closed form solution, using norm-based trimming is also technical challenging in this case. Handling the Byzantine worker machines becomes a bit more complicated as those stationary conditions of the good machines (non-Byzantine machine) do not hold for the Byzantine worker machines.

5.1.1 Our Contributions

1. We propose a novel distributed and robust cubic regularized Newton algorithm, that escapes saddle point efficiently. We prove that the algorithm convergence at a rate of $\frac{1}{T^{2/3}}$, which is faster than the first order methods (which converge at $1/\sqrt{T}$ rate, see [31]). Hence, the number of iterations (and hence the communication cost) required to achieve a target accuracy is much fewer than the first order methods. A simple simulation in Section 5.6, shows that the algorithm of [31] requires 36x more steps than ours, showing a huge communication gain.

2. The computation complexity of our algorithm is also much less than the existing schemes [31]. Part of computation is deferred to the worker machines, which is desirable in paradigm like Federated Learning.
3. We use norm-based thresholding to resist Byzantine workers. In previous works, computation heavy techniques like coordinate-wise median, coordinate-wise trimmed mean and spectral filtering are used to resist Byzantine workers. In contrast, our norm based thresholding in computation friendly.
4. We work with inexact gradients and Hessians, which is quite common in distributed setup like Federated Learning.
5. In Section 5.6, we verify our theoretical findings via experiments. We use benchmark LIBSVM ([47]) datasets for logistic regression and non-convex robust regression and show convergence results for both non-Byzantine and several different Byzantine attacks.

5.2 Related Work

In the recent years, there are handful first order algorithms [142–144] that focus on the escaping saddle points and convergence to local minima. The critical algorithmic aspect is running gradient based algorithm and adding perturbation to the iterates when the gradient is small. ByzantinePGD [31], PGD [46], Neon+GD[138], Neon2+GD [139] are examples of such algorithms. For faster convergence rate, second order Hessian based algorithms are developed for saddle point avoidance. The work of Nesterov and Polyak [43] first proposes the cubic regularized Newton method and provides analysis for the second order stationary condition. An algorithm called Adaptive Regularization with Cubics (ARC) was developed by [145, 146] where cubic regularized Newton method with access to inexact Hessian was studied. The inexactness of Hessians for the ARC algorithm is adaptive over iterations. Cubic regularization with both the gradient and Hessian being inexact was studied in [147]. In [44], a cubic regularized Newton with sub-sampled Hessian and gradient was proposed, but for analysis, the batch size of the sample changes in adaptive manner to provide guarantees for the inexactness of the Hessian and gradient. In this work, we also take a similar approach as [44], but we relax the adaptive nature of the sample size. Momentum based cubic regularized algorithm was studied in [45]. A variance reduced cubic regularized algorithm was proposed in [148, 149]. In terms of solving the cubic sub-problem, [150] proposes a gradient based algorithm and [151] provides a Hessian-vector product technique.

Furthermore, as mentioned in Chapters 3,4, the Byzantine resilience algorithms in FL have received significant interest (see [19, 22, 23, 30, 140]. Also, in Chapters 3 and 4, a norm based thresholding approach for Byzantine resilience for distributed first and second order algorithms respectively was proposed. However, none of these works provide saddle point avoidance guarantees (they only show convergence to a small gradient point). The work [31] is the only one that provides second order guarantee (Hessian positive semi-definite) under Byzantine attack, and we compare our algorithm with that of [31].

5.3 Problem Formulation

In this section, we formally set up the problem. Similar to Chapter 4, we minimize a loss function additive over losses of worker machines. The loss function takes the following form:

$$f(w) = \frac{1}{m} \sum_{i=1}^m f_i(w), \quad (5.1)$$

where the function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is twice differentiable and non-convex. Continuing with the notation from the previous chapters, we consider distributed optimization framework with m worker machines and one center machine where the worker machines communicate to the center machine. Each worker machine is associated with a local loss function f_i minimized over i.i.d. data points drawn from some unknown distribution. In addition to that, we also consider the case where α fraction of the worker machines are Byzantine for some $\alpha < \frac{1}{2}$. The Byzantine machines can send arbitrary updates to the central machine which can disrupt the learning. Furthermore, the Byzantine machines can collude with each other, create *fake local minima* or attack maliciously by gaining information about the learning algorithm and other workers.

In the rest of the chapter, the norm $\|\cdot\|$ will refer to ℓ_2 norm or spectral norm when the argument is a vector or a matrix respectively.

Algorithm 5: Byzantine Robust Distributed Cubic Regularized Newton Algorithm

- 1: **Input:** Step size η_k , parameter $\beta \geq 0, \gamma > 0, M > 0$
 - 2: **Initialize:** Initial iterate $w_0 \in \mathbb{R}^d$
 - 3: **for** $k = 0, 1, \dots, T - 1$ **do**
 - 4: **Central machine:** broadcasts w_k
 for $i \in [m]$ **do in parallel**
 - 5: i -th worker machine:
 - Non-Byzantine: Computes local gradient $\mathbf{g}_{i,k}$ and local Hessian $\mathbf{H}_{i,k}$; locally solves the problem described in equation (5.2) and sends $\mathbf{s}_{i,k+1}$ to the central machine,
 - Byzantine: Generates \star (arbitrary), and sends it to the center machine
 - end for**
 - 6: Center Machine:
 - Sort the worker machines in a non decreasing order according to norm of updates $\{\mathbf{s}_{i,k+1}\}_{i=1}^m$ from the local machines
 - Return the indices of the first $1 - \beta$ fraction of machines as \mathcal{U}_t ,
 - Update parameter: $w_{k+1} = w_k + \eta_k \frac{1}{|\mathcal{U}_t|} \sum_{i \in \mathcal{U}_t} \mathbf{s}_{i,k+1}$
 - 7: **end for**
-

5.4 Distributed Cubic Regularized Newton

We first focus on the non-Byzantine setup ($\alpha = 0; \beta = 0$ in Algorithm 5) of distributed cubic regularized Newton algorithm. Byzantine resilience attribute of Algorithm 5 is deferred to Section 5.5. As mentioned before, the data is drawn independently across worker machines from some unknown distribution. The local data at the i th machine is denoted by S_i . Starting with initialization w_0 , the central machine broadcasts the parameter to the worker machines. At k -th iteration, the i -th worker machine solves a cubic-regularized auxiliary loss function based on its local data:

$$\mathbf{s}_{i,k+1} = \underset{\mathbf{s}}{\operatorname{argmin}} \mathbf{g}_{i,k}^T \mathbf{s} + \frac{\gamma}{2} \mathbf{s}^T \mathbf{H}_{i,k} \mathbf{s} + \frac{M}{6} \gamma^2 \|\mathbf{s}\|^3, \quad (5.2)$$

where $M > 0, \gamma > 0$ are parameter and $\mathbf{g}_{i,t}, \mathbf{H}_{i,t}$ are the gradient and Hessian of the local loss function f_i computed on the independently sampled data (S_i) stored in the worker machine.

$$\begin{aligned} \mathbf{g}_{i,k} &= \nabla f_i(w_k) = \frac{1}{|S_i|} \sum_{z_i \in S_i} \nabla f_i(w_k, z_i), \\ \mathbf{H}_{i,k} &= \nabla^2 f_i(w_k) = \frac{1}{|S_i|} \sum_{z_i \in S_i} \nabla^2 f_i(w_k, z_i). \end{aligned}$$

After receiving the update $\mathbf{s}_{i,k+1}$, the central machine updates the parameter in the following way

$$\mathbf{s}_{k+1} = \mathbf{s}_k + \frac{\eta_k}{m} \sum_{i=1}^m \mathbf{s}_{i,k+1}, \quad (5.3)$$

where η_k is the step-size.

Remark 23. Note that, we introduce the parameter γ in the cubic regularized sub-problem. The parameter γ emphasizes the effect of the second and third order terms in the sub-problem. The choice of γ plays an important role in our analysis in handling the non-linear update from different worker machines. Such non-linearity vanishes if we choose $\gamma = 0$ and the distributed cubic Newton becomes distributed gradient descent algorithm.

5.4.1 Theoretical Guarantees

We have the following standard assumptions:

Assumption 16. *The non-convex loss function $f(\cdot)$ is twice continuously-differentiable and bounded below, i.e., $f^* = \inf_{w \in \mathbb{R}^d} f(x) > -\infty$.*

Assumption 17. *The loss $f(\cdot)$ is L -Lipschitz continuous ($\forall \mathbf{z}, \mathbf{y}, |f(\mathbf{z}) - f(\mathbf{y})| \leq L \|\mathbf{z} - \mathbf{y}\|$), has L_1 -Lipschitz gradients ($\|\nabla f(\mathbf{z}) - \nabla f(\mathbf{y})\| \leq L_1 \|\mathbf{z} - \mathbf{y}\|$) and L_2 -Lipschitz Hessian ($\|\nabla^2 f(\mathbf{z}) - \nabla^2 f(\mathbf{y})\| \leq L_2 \|\mathbf{z} - \mathbf{y}\|$).*

The above assumption states that the loss and the gradient and Hessian of the loss do not drastically change in the local neighborhood. These assumptions are standard in the analysis of the saddle point escape for cubic regularization (see [44, 147]).

In this work, each worker machine solves the cubic sub-problem as described in the equation (5.2). The gradient and Hessian used in the equation are inexact in nature as they are computed using the sub-sampled data.

Assumption 18. For a given $\epsilon_g > 0$ and for all k, i ,

$$\|\nabla f(w_k) - \mathbf{g}_{i,k}\| \leq \epsilon_g. \quad (5.4)$$

Assumption 19. For a given $\epsilon_H > 0$ and for all k, i ,

$$\|\nabla^2 f(w_k) - \mathbf{H}_{i,k}\| \leq \epsilon_H. \quad (5.5)$$

In the following section, we provide an exact characterization of ϵ_g and ϵ_H to justify the assumptions.

5.4.2 Inexact Gradient and Hessian

In this work, we assume that each worker machine solve the sub-problem described in equation (5.2) with the data sampled independently from some unknown distribution. So, in each iteration, the gradient and Hessian computed by each worker machine are actually sub-sampled gradient and Hessian of the objective function f and inexact in nature. In the following lemmas, we described the deviation conditions given the size of the sampled data and provide probabilistic deviation bound that ensure the deviation defined in Assumption 18 and 19.

Lemma 23. (Gradient deviation bound) Given S iid data sample, under the Assumption 17, we have $\|\mathbf{g}_i - \nabla f\| \leq c \left(\frac{L\sqrt{\log(2d/\delta)}}{\sqrt{|S|}} \right)$, with probability exceeding $1 - \delta$, where c is a constant and \mathbf{g}_i is the gradient computed in i -th worker machine.

Lemma 24. (Hessian deviation bound) Given S iid data sample, under the Assumption 17, we have $\|\mathbf{H}_i - \nabla^2 f\| \leq c_1 \left(\frac{L_1\sqrt{\log(2d/\delta)}}{\sqrt{|S|}} \right)$, with probability at least $1 - \delta$, where c_1 is a constant and \mathbf{H}_i is the Hessian computed in i -th worker machine.

Remark 24. In the Assumptions 18 and 19, the deviation bounds are in ℓ_2 norm for the gradient and in spectral norm for the Hessian. In the previous works, in centralized model, [44, 45, 147] study cubic regularization with sub-sampled and inexact gradient and Hessian. In the central model, the motivation of the sub-sampled Hessian and gradient is for the ease of the computation. Here, we use the deviation bounds as each of the worker machine only have access to a fraction of the data.

Also, in contrast to the sub-sampled analysis of [44], we choose the deviation of both the gradient and Hessian in the Assumptions 18 and 19 to be independent of the update s which is the solution of the sub-problem (5.2).

The analysis of the deviation bounds follows from the vector and matrix Bernstein inequalities. The assumption of the independent data in each worker can be relaxed. The analysis can be easily extended for data partition (non iid data), following an analysis of [141] the bound of $\sqrt{\frac{1}{|S|}}$ holds.

Theorem 10. *Under the Assumptions 16,17,18, 19, and $\alpha = 0$, after T iterations, the sequence $\{w_i\}_{i=1}^T$ generated by the Algorithm 5 with $\beta = 0$, contains a point \tilde{w} such that*

$$\begin{aligned} \|\nabla f(\tilde{w})\| &\leq \frac{\Psi_1}{T^{\frac{2}{3}}} + \frac{\Psi_2}{T} + (\epsilon_g + \epsilon_H), \\ \lambda_{\min}(\nabla^2 f(\tilde{x})) &\geq -\frac{\Psi_3}{T^{\frac{1}{3}}} - \epsilon_H, \end{aligned} \tag{5.6}$$

where, $\lambda_{\min}(\cdot)$ denotes the minimum eigenvalue and

$$\begin{aligned} \Psi_1 &= \left(\frac{L_2}{2} + \frac{M}{2}\right) \Psi^2, \Psi_2 = \epsilon_H \Psi^3, \Psi_3 = \left(\frac{M}{2} + L_2\right) \Psi \\ \Psi &= \left[\frac{f(w_0) - f^*}{\lambda} + \frac{\sum_{k=0}^{T-1} \eta_k^2}{2m^2 \lambda} ((m-1)L_1 + \epsilon_H) + \frac{\sum_{k=0}^{T-1} \eta_k \epsilon_g}{m\lambda} \right]^{\frac{1}{3}} \\ \lambda &= \left(\frac{M}{4m} - \frac{L_2}{6} - \frac{1}{2\eta_k m} ((m-1)L_1 + \epsilon_H) - \frac{\epsilon_g}{\eta_k^2} \right) \end{aligned}$$

Remark 25. We choose the step sizes $\{\eta_k\}_{k=0}^{T-1}$ such way that $\sum_{i=0}^T \eta_k$ and $\sum_{i=0}^T \eta_k^2$ is bounded. For the ease of choice, we can choose $\eta_k = \frac{c}{T}$, for some constant $c > 0$. Also, we choose $\eta_k = \gamma$.

Remark 26. Both the gradient and the minimum eigenvalue of the Hessian in the Theorem 10 have two parts. The first part decreases with the number iterations T . The gradient and the minimum eigenvalue of the Hessian have the rate of $O\left(\frac{1}{T^{\frac{2}{3}}}\right)$ and $O\left(\frac{1}{T^{\frac{1}{3}}}\right)$, respectively. Both of these rates match the rates of the centralized version of the cubic regularized Newton. In the second parts of the gradient bound and the minimum eigenvalue of the Hessian have the error floor of $\epsilon_g + \epsilon_H$ and ϵ_H , respectively. Both the terms ϵ_g and ϵ_H decrease at the rate of $\frac{1}{\sqrt{|S|}}$, where $|S|$ is the number of data in each of the worker machines.

Remark 27 (Two rounds of communication $\epsilon_g = 0$). We can improve the bound in the Theorem 10, with the calculation of the actual gradient which requires one more round of communication in each iteration. In the first iteration, all the worker machines compute the gradient based on the stored data and send it to the center machine. The center machine averages them and then broadcast the global gradient $\nabla f(w_k) = \frac{1}{m} \sum_{i=1}^m \mathbf{g}_{i,k}$ at iteration k . In this manner, the worker machines solve the sub-problem (5.2) with the actual gradient. The analysis follows same as that of the Theorem 10 with $\epsilon_g = 0$. This improves the gradient bound while the communication remains $O(d)$ in each iteration.

5.4.3 Solving cubic sub-problem

We use a gradient based approach for solving the cubic sub-problem (5.2) in each worker machines (see [147]). The worker machines computes the gradient and Hessian based on the local data stored in the machines and perform the following gradient descent algorithm to yields a solution withing certain tolerance.

Algorithm 6: Gradient based Cubic solver

- 1: **Input:** Step size ξ , local gradient g and Hessian \mathbf{H} and tolerance $\tau > 0$ and M, γ .
 - 2: **Initialize:** $\mathbf{s} \leftarrow 0; G \leftarrow g$
 - 3: While $\|G\| > \tau$;
 - $\mathbf{s} \leftarrow \mathbf{s} - \xi G$
 - $G = g + \gamma \mathbf{H}\mathbf{s} + \frac{M\gamma^2}{2} \|\mathbf{s}\|\mathbf{s}$
 - 4: Return \mathbf{s}
-

5.5 Byzantine Resilience

In this section, we analyze our algorithm’s resilience against Byzantine workers. We consider that $\alpha (< \frac{1}{2})$ fraction of the worker machines are Byzantine in nature. We denote the set of Byzantine worker machines by \mathcal{B} and the set of the rest of the good machines as \mathcal{M} . In each iteration, the good machines send the solution of the sub-cubic problem described in equation (5.2) and the Byzantine machines can send any arbitrary values or intentionally disrupt the learning algorithm with malicious updates. Moreover, in the non-convex optimization problems, one of the more complicated and important issue is to avoid saddle points which can yield highly sub-optimal results. In the presence of Byzantine worker machines, they can be in cohort to create a *fake local minima* and drive the algorithm into sub-optimal region. Lack of any robust measure towards these type of intentional and unintentional attacks can be catastrophic to the learning procedure as the learning algorithm can get stuck in such sub-optimal point. To tackle such Byzantine worker machines, we employ a simple process called *norm based thresholding*.

After receiving all the updates from the worker machines, the central machine outputs a set \mathcal{U} which consists of the indexes of the worker machines with smallest norm. We choose the size of the set \mathcal{U} to be $(1 - \beta)m$. Hence, we ‘trim’ β fraction of the worker machine so that we can control the iterated update by not letting the worker machines with large norm participate and diverge the learning process. We denote the set of trimmed machine as \mathcal{T} . We choose $\beta > \alpha$ so that at least one of the good machines gets trimmed. In this way, the norm of the all the updates in the set \mathcal{U} is bounded by at least the largest norm of the good machines.

Theorem 11. For $0 \leq \alpha \leq \beta \leq \frac{1}{2}$ and under the Assumptions 16,17,18, 19, after T iterations, the sequence $\{w_i\}_{i=1}^T$ generated by the Algorithm 5 contains a point \tilde{w} such that

$$\begin{aligned} \|\nabla f(\tilde{w})\| &\leq \frac{\Psi_{1,byz}}{T^{\frac{2}{3}}} + \frac{\Psi_{2,byz}}{T} + \epsilon_g + \frac{(1-\alpha)}{(1-\beta)}(1+\alpha m)\epsilon_H \\ \lambda_{\min}(\nabla^2 f(\tilde{w})) &\geq -\frac{\Psi_{3,byz}}{T^{\frac{1}{3}}} - \epsilon_H \quad \text{where,} \end{aligned} \quad (5.7)$$

$$\begin{aligned} \Psi_{1,byz} &= \left(\frac{L_2(1+\alpha m)(1-\alpha)}{2(1-\beta)} + \frac{M(1+\alpha m)^2}{2} \left(\frac{(1-\alpha)^2}{(1-\beta)} \right) \right) \Psi_{byz}^2 \\ \Psi_{2,byz} &= \frac{(1-\alpha)}{(1-\beta)}(1+\alpha m)\epsilon_H \Psi_{byz}^3 \\ \Psi_{3,byz} &= \left(\frac{M(1-\alpha)}{2(1-\beta)}(1+\alpha m) + L_2 \frac{(1+\alpha m)(1-\alpha)}{(1-\beta)} \right) \Psi_{byz} \\ \Psi_{byz} &= \left[\frac{f(w_0) - f^*}{\lambda_{byz}} + \frac{\sum_{k=0}^{T-1} \lambda_{floor}}{\lambda_{byz}} \right]^{\frac{1}{3}} \\ \lambda_{byz} &= \frac{M\gamma(1-\alpha)}{4\eta_k(1-\beta)^2 m} - \frac{(L(1+\alpha m) + \epsilon_g)(1-\alpha)}{\eta_k^2(1-\beta)} - \frac{(1-\alpha)\alpha m L_2}{6(1-\beta)} \\ &\quad - \frac{(1-\alpha)}{2\eta_k(1-\beta)^2 m} (1-\beta)mL_1 + \alpha(1-\beta)m^2 + \epsilon_H \\ \lambda_{floor} &= \frac{\eta_k \epsilon_g (1-\alpha)}{(1-\beta)} + \frac{\eta_k L}{(1-\beta)} + \frac{\eta_k^2}{2(1-\beta)^2 m^2} \times \\ &\quad [(2(1-\beta)m + 1)L_1 + \epsilon_H](1-\alpha)m + \alpha(1-\beta)m^2 L_1 \end{aligned}$$

Remark 28. Compared to the non-Byzantine part described in Theorem 10, the rate of remains same except for the error floor of the gradient bound suffering a small constant factor.

Remark 29. The condition for the step-size η_k remains same as described in the Remark 26 and we choose $\gamma = \frac{\eta_k(1-\alpha)}{(1-\beta)}(1+\alpha m)$.

Remark 30. In the previous work, [31] first provides a *perturbed gradient based algorithm* to escape the saddle point in non-convex optimization in the presence of Byzantine worker machines. Also, in that paper, the Byzantine resilience is achieved using techniques such as trimmed mean, median and collaborative filtering. These methods require additional assumptions (coordinate of the gradient being sub-exponential etc.) for the purpose of analysis. In this work, we perform a simple *norm based thresholding* that provides robustness towards any sorts of adversarial attacks. Also the perturbed gradient descent (PGD) actually requires multiple rounds of communications between the central machine and the worker machines whenever the norm of the gradient is small as this is an indication of either a local minima or a saddle point. In contrast to that, our method does not require any additional communication for *escaping* the saddle points. Our method provides such ability by virtue of cubic regularization.

Remark 31. Since our algorithm is second order in nature, it requires less number of iterations compared to the first order gradient based algorithms. Our algorithm achieves a superior rate of $O\left(\frac{1}{T^{\frac{2}{3}}}\right)$ compared to the gradient based approach of rate $O\left(\frac{1}{\sqrt{T}}\right)$. Our algorithm dominates ByzantinePGD [23] in terms of convergence, communication rounds and simplicity and efficiency of Byzantine resilience.

5.6 Experimental Results

In this section, we validate our algorithm in both Byzantine and non-Byzantine setup on benchmark LIBSVM ([47]) dataset in both convex and non-convex problems. We choose the following problems for our experiment.

1. Logistic regression:

$$\min_{w \in \mathbb{R}^d} \frac{1}{n} \sum_i \log(1 + \exp(-y_i \mathbf{x}_i^T w)) + \frac{\lambda}{2n} \|w\|^2. \quad (5.8)$$

2. Non-convex robust linear regression:

$$\min_{w \in \mathbb{R}^d} \frac{1}{n} \sum_i \log\left(\frac{(y_i - w^T \mathbf{x}_i)^2}{2} + 1\right), \quad (5.9)$$

where $w \in \mathbb{R}^d$ is the parameter, $\{\mathbf{x}_i\}_{i=1}^n \in \mathbb{R}^d$ are the feature vectors and $\{y_i\}_{i=1}^n \in \{0, 1\}$ are the corresponding labels. We choose ‘a9a’ ($d = 123, n \approx 32K$, we split the data into 70/30 and use as training/testing purpose) and ‘w8a’ (training data $d = 300, n \approx 50K$ and testing data $d = 300, n \approx 15K$) classification datasets and partition the data in 20 different worker machines.

In Figure 5.3, we show the performance of our algorithm in non-Byzantine setup ($\alpha = \beta = 0$). In the top row of Figure 5.3, we plot the classification accuracy on test data of both ‘a9a’ and ‘w8a’ datasets for logistic regression problem and in the bottom row of Figure 5.3, we plot the function value of the non-convex robust linear regression problem defined in equation (5.9) for training data of ‘a9a’ and ‘w8a’ datasets. We choose the learning rate $\eta_k = 1$ and the parameter $\lambda = 1$ and $M = \{10, 15, 20\}$.

Next, we show the effectiveness of our algorithm in Byzantine setup. In this work, we consider the following four Byzantine attacks: (1) ‘Gaussian Noise attack’: where the Byzantine worker machines add Gaussian noise to the update. (2) ‘Random label attack’: where the Byzantine worker machines train and learn based on random labels instead of the proper labels. (3) ‘Flipped label attack’: where (for Binary classification) the Byzantine worker machines flip the labels of the data and learn based on wrong labels. (4) ‘Negative update attack’: where the Byzantine workers computes the update s (here solves the sub-problem in Eq. (5.2)) and communicates $-c * s$ with $c \in (0, 1)$ making the direction of the update opposite of the actual one.




Figure 5.1: Function loss of the training data ‘a9a’ dataset (first row) and ‘w8a’ dataset (second row) with 10%, 15%, 20% Byzantine worker machines for (a,e). Flipped label attack.(b,f). Negative Update attack (c,g). Gaussian noise attack and (d,h). Random label attack for non-convex robust linear regression problem.

We show the classification accuracy on testing data of ‘a9a’ and ‘w8a’ dataset for logistic regression problem in Figure 5.2 and training function loss of ‘a9a’ and ‘w8a’ dataset for robust linear regression problem in the Figure 5.1. It is evident from the plots that a simple *norm based thresholding* makes the learning algorithm robust. We choose the parameters $\lambda = 1$, $M = 10$, learning rate $\eta_k = 1$, fraction of the Byzantine machines $\alpha = \{.1, .15, .2\}$ and $\beta = \alpha + \frac{2}{m}$.

We also compare our algorithm with ByzantinePGD [23]. For both the algorithms, we choose ℓ_2 norm of the gradient as a stopping criteria and compute the number of times the worker machines communicate with the center machine as a measure of communication cost and convergence rate. We choose $R = 10$, $r = 5$, $Q = 10$, $T_{th} = 10$ and ‘co-ordinate wise Trimmed mean’ for Byzantine resilience (see the algorithm in [23]). For our algorithm, we choose $M = 10$, $\lambda = 1$ in the non-Byzantine setup. The ByzantinePGD algorithm requires 257 rounds of communications (157 rounds for reaching the stopping criteria and 100 rounds for the ‘Escape’ sub-routine to check whether it is a ‘saddle point’) where our algorithm requires only 7 rounds of communication. We choose ‘w8a’ dataset for non-convex robust linear regression problem. Eventhough ByzantinePGD does not require the computation of Hessian and cubic sub-problem solving, our algorithm outperforms by a lot (36x) in terms of communication rounds. In the Appendix, we provide results in Byzantine settings.




Figure 5.2: Classification accuracy of the testing data ‘a9a’ dataset (first row) and ‘w8a’ dataset (second row) with 10%, 15%, 20% Byzantine worker machines for (a,e). Flipped label attack.(b,f). Negative Update attack (c,g). Gaussian noise attack and (d,h). Random label attack for logistic regression problem.




Figure 5.3: (First row) Accuracy of the algorithm for logistic regression on test data of (a). a9a and (b). w8a dataset. (Second row). Function value of the non-convex robust linear regression on the training data of (a). a9a and (b). w8a dataset.

5.7 Proofs

In this part, first, we establish some useful facts and lemmas. Next, we provide the missing proofs and analysis of Theorems 10 and 11 and additional experiments comparing ByzantinePGD [31] in Byzantine setup.

5.7.1 Some useful facts

For the purpose of analysis we use the following sets of inequalities.

Fact 1. For a_1, \dots, a_n we have the following inequality

$$\left\| \left(\sum_{i=1}^n a_i \right) \right\|^3 \leq \left(\sum_{i=1}^n \|a_i\| \right)^3 \leq n^2 \sum_{i=1}^n \|a_i\|^3 \quad (5.10)$$

$$\left\| \left(\sum_{i=1}^n a_i \right) \right\|^2 \leq \left(\sum_{i=1}^n \|a_i\| \right)^2 \leq n \sum_{i=1}^n \|a_i\|^2 \quad (5.11)$$

Fact 2. For $a_1, \dots, a_n > 0$ and $r < s$

$$\left(\frac{1}{n} \sum_{i=1}^n a_i^r \right)^{1/r} \leq \left(\frac{1}{n} \sum_{i=1}^n a_i^s \right)^{1/s} \quad (5.12)$$

Lemma 25 ([43]). *Under Assumption 17, i.e., the Hessian of the function is L_2 -Lipschitz continuous, for any $w, \mathbf{y} \in \mathbb{R}^d$, we have*

$$\|\nabla f(w) - \nabla f(\mathbf{y}) - \nabla^2 f(w)(\mathbf{y} - w)\| \leq \frac{L_2}{2} \|\mathbf{y} - w\|^2 \quad (5.13)$$

$$\left| f(\mathbf{y}) - f(w) - \nabla f(w)^T(\mathbf{y} - w) - \frac{1}{2}(\mathbf{y} - w)^T \nabla^2 f(w)(\mathbf{y} - w) \right| \leq \frac{L_2}{6} \|\mathbf{y} - w\|^3 \quad (5.14)$$

Next, we establish the following Lemma that provides some nice properties of the cubic subproblem.

Lemma 26. *Let $M > 0, \gamma > 0, \mathbf{g} \in \mathbb{R}^d, \mathbf{H} \in \mathbb{R}^{d \times d}$, and*

$$\mathbf{s} = \underset{w}{\operatorname{argmin}} \mathbf{g}^T w + \frac{\gamma}{2} w^T \mathbf{H} w + \frac{M\gamma^2}{6} \|w\|^3. \quad (5.15)$$

The following holds

$$\mathbf{g} + \gamma \mathbf{H} \mathbf{s} + \frac{M\gamma^2}{2} \|\mathbf{s}\| \mathbf{s} = \mathbf{0}, \quad (5.16)$$

$$\mathbf{H} + \frac{M\gamma}{2} \|\mathbf{s}\| \mathbf{I} \succeq \mathbf{0}, \quad (5.17)$$

$$\mathbf{g}^T \mathbf{s} + \frac{\gamma}{2} \mathbf{s}^T \mathbf{H} \mathbf{s} \leq -\frac{M}{4} \gamma^2 \|\mathbf{s}\|^3. \quad (5.18)$$

Proof. The equations (5.16) and (5.17) are from the first and second order optimal condition. We proof (5.18), by using the conditions of (5.16) and (5.17).

$$\mathbf{g}^T \mathbf{s} + \frac{\gamma}{2} \gamma \mathbf{s}^T \mathbf{H} \mathbf{s} = - \left(\gamma \mathbf{H} \mathbf{s} + \frac{M}{2} \gamma^2 \|\mathbf{s}\| \mathbf{s} \right)^T \mathbf{s} + \frac{\gamma}{2} \gamma \mathbf{s}^T \mathbf{H} \mathbf{s} \quad (5.19)$$

$$\begin{aligned} &= -\gamma \mathbf{s}^T \mathbf{H} \mathbf{s} - \frac{M}{2} \gamma^2 \|\mathbf{s}\|^3 + \frac{\gamma}{2} \gamma \mathbf{s}^T \mathbf{H} \mathbf{s} \\ &\leq \frac{M}{4} \gamma^2 \|\mathbf{s}\|^3 - \frac{M}{2} \gamma^2 \|\mathbf{s}\|^3 \\ &= -\frac{M}{4} \gamma^2 \|\mathbf{s}\|^3. \end{aligned} \quad (5.20)$$

In (5.19), we substitute the expression \mathbf{g} from the equation (5.16). In (5.20), we use the fact that $\mathbf{s}^T \mathbf{H} \mathbf{s} + \frac{M\gamma}{2} \|\mathbf{s}\|^3 > 0$ from the equation (5.17). □

5.7.2 Proof of Theorem 10

First we state the results of Lemma 26 for each worker node in iteration k ,

$$\mathbf{g}_{i,k} + \gamma \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} + \frac{M}{2} \gamma^2 \|\mathbf{s}_{i,k+1}\| \mathbf{s}_{i,k+1} = \mathbf{0} \quad (5.21)$$

$$\gamma \mathbf{H}_{i,k} + \frac{M}{2} \gamma^2 \|\mathbf{s}_{i,k+1}\| \mathbf{I} \succeq \mathbf{0} \quad (5.22)$$

$$\mathbf{g}_{i,k}^T \mathbf{s}_{i,k+1} + \frac{\gamma}{2} \mathbf{s}_{i,k+1}^T \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \leq -\frac{M}{4} \gamma^2 \|\mathbf{s}_{i,k+1}\|^3 \quad (5.23)$$

At iteration k , we have,

$$\begin{aligned}
 & f(w_{k+1}) - f(w_k) \\
 & \leq \nabla f(w_k)^T (w_{k+1} - w_k) + \frac{1}{2} (w_{k+1} - w_k)^T \nabla^2 f(w_k) (w_{k+1} - w_k) + \frac{L_2}{6} \|w_{k+1} - w_k\|^3 \\
 & \leq \eta_k \nabla f(w_k)^T \mathbf{s}_{k+1} + \eta_k^2 \frac{1}{2} \mathbf{s}_{k+1}^T \nabla^2 f(w_k) \mathbf{s}_{k+1} + \frac{L_2}{6} \|\eta_k \mathbf{s}_{k+1}\|^3 \\
 & = \frac{\eta_k}{m} \nabla f(w_k)^T \left(\sum_i \mathbf{s}_{i,k+1} \right) + \frac{\eta_k^2}{2} \left(\frac{1}{m} \sum_i \mathbf{s}_{i,k+1} \right)^T \nabla^2 f(w_k) \left(\frac{1}{m} \sum_i \mathbf{s}_{i,k+1} \right) + \frac{L_2 \eta_k^3}{6} \|\mathbf{s}_{k+1}\|^3 \quad (5.24)
 \end{aligned}$$

$$\begin{aligned}
 & \leq \frac{\eta_k}{m} \left(\sum_i \nabla f(w_k)^T \mathbf{s}_{i,k+1} \right) + \frac{\eta_k^2}{2m^2} \left(\sum_i \mathbf{s}_{i,k+1}^T \nabla^2 f(w_k) \mathbf{s}_{i,k+1} + \sum_{i \neq j} \mathbf{s}_{i,k+1}^T \nabla^2 f(w_k) \mathbf{s}_{j,k+1} \right) \\
 & + \frac{L_2 \eta_k^3}{6m} \sum_i \|\mathbf{s}_{i,k+1}\|^3 \\
 & \leq \frac{\eta_k}{m} \sum_i \left(\mathbf{g}_{i,k}^T \mathbf{s}_{i,k+1} + \frac{\gamma}{2} \mathbf{s}_{i,k+1}^T \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \right) + \frac{\eta_k}{m} \left(\sum_i (\nabla f(w_k) - \mathbf{g}_{i,k+1})^T \mathbf{s}_{i,k+1} \right) \\
 & - \frac{\eta_k \gamma}{2m} \sum_i \mathbf{s}_{i,k+1}^T \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \\
 & + \frac{\eta_k^2}{2m^2} \left(\sum_i \mathbf{s}_{i,k+1}^T \nabla^2 f(w_k) \mathbf{s}_{i,k+1} + \sum_{i \neq j} \mathbf{s}_{i,k+1}^T \nabla^2 f(w_k) \mathbf{s}_{j,k+1} \right) + \frac{L_2 \eta_k^3}{6m} \sum_i \|\mathbf{s}_{i,k+1}\|^3 \\
 & \leq \frac{\eta_k}{m} \sum_i -\frac{M}{4} \gamma^2 \|\mathbf{s}_{i,k+1}\|^3 + \frac{\eta_k}{m} \left(\sum_i \epsilon_g \|\mathbf{s}_{i,k+1}\| \right) - \left(\frac{\eta_k \gamma}{2m} - \frac{\eta_k^2}{2m^2} \right) \sum_i \mathbf{s}_{i,k+1}^T \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \\
 & + \frac{\eta_k^2}{2m^2} \left[\sum_i \mathbf{s}_{i,k+1}^T (\nabla^2 f(w_k) - \mathbf{H}_{i,k}) \mathbf{s}_{i,k+1} + \sum_{i \neq j} \mathbf{s}_{i,k+1}^T \nabla^2 f(w_k) \mathbf{s}_{j,k+1} \right] + \frac{L_2 \eta_k^3}{6m} \sum_i \|\mathbf{s}_{i,k+1}\|^3 \quad (5.25)
 \end{aligned}$$

$$\begin{aligned}
 & \leq \left(-\frac{M\gamma^2 \eta_k}{4m} + \frac{L_2 \eta_k^3}{6m} \right) \sum_i \|\mathbf{s}_{i,k+1}\|^3 + \frac{\eta_k \epsilon_g}{m} \left(\sum_i \|\mathbf{s}_{i,k+1}\| \right) \\
 & + \left(\frac{\eta_k \gamma}{2m} - \frac{\eta_k^2}{2m^2} \right) \sum_i \frac{M}{2} \gamma \|\mathbf{s}_{i,k+1}\|^3 \\
 & + \frac{\eta_k^2}{2m^2} \left[\epsilon_H \sum_i \|\mathbf{s}_{i,k+1}\|^2 + L_1 \sum_{i \neq j} \|\mathbf{s}_{i,k+1}\| \|\mathbf{s}_{j,k+1}\| \right] \quad (5.26)
 \end{aligned}$$

$$\begin{aligned}
 & = \left(-\frac{M\gamma^2 \eta_k}{4m^2} + \frac{L_2 \eta_k^3}{6m} \right) \sum_i \|\mathbf{s}_{i,k+1}\|^3 + \underbrace{\frac{\eta_k \epsilon_g}{m} \left(\sum_i \|\mathbf{s}_{i,k+1}\| \right)}_{\text{Term1}} \\
 & + \underbrace{\frac{\eta_k^2}{2m^2} \left[\epsilon_H \sum_i \|\mathbf{s}_{i,k+1}\|^2 + L_1 \sum_{i \neq j} \|\mathbf{s}_{i,k+1}\| \|\mathbf{s}_{j,k+1}\| \right]}_{\text{Term2}} \quad (5.27)
 \end{aligned}$$

In (5.24), we apply the inequality (5.20) on $\|\frac{1}{m} \sum_i \mathbf{s}_{i,k+1}\|^3$. In line (5.25), we use the gradient approximation from the Assumption 18. In line (5.25), we apply the fact that $\mathbf{s}_{i,k+1}^T \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} + \frac{M\gamma}{2} \|\mathbf{s}_{i,k+1}\|^3 > 0$ from the equation (5.22) and assume that $\gamma > \frac{\eta_k}{m}$. In line (5.26), we use the Assumption 19 and the fact that the Hessian of the objective function is bounded as the gradient is L_1 -Lipschitz continuous.

Now we bound the Term 1 and Term 2 in equation (5.27).

$$\text{Term 1} \leq \frac{\eta_k \epsilon_g}{m} \sum_i \|\mathbf{s}_{i,k+1}\| \leq \frac{\eta_k \epsilon_g}{m} \sum_i (1 + \|\mathbf{s}_{i,k+1}\|^3) \quad (5.28)$$

In line (5.28), we use the fact $a^2 b \leq a^3 + b^3$ where $a, b > 0$. We choose $a = 1$ and $b = \|\mathbf{s}_{i,k+1}\|$.

$$\begin{aligned} \text{Term 2} &\leq \frac{\eta_k^2}{2m^2} \left[\epsilon_H \sum_i \|\mathbf{s}_{i,k+1}\|^2 + L_1 \left(\left\| \sum_i \mathbf{s}_{i,k+1} \right\|^2 - \sum_i \|\mathbf{s}_{i,k+1}\|^2 \right) \right] \\ &\leq \frac{\eta_k^2}{2m^2} ((m-1)L_1 + \epsilon_H) \sum_i \|\mathbf{s}_{i,k+1}\|^2 \end{aligned} \quad (5.29)$$

$$\leq \frac{\eta_k^2}{2m^2} ((m-1)L_1 + \epsilon_H) \sum_i (\|\mathbf{s}_{i,k+1}\|^3 + 1) \quad (5.30)$$

In line (5.29), we use the inequality (5.11) to bound $\|\sum_i \mathbf{s}_{i,k+1}\|^2$ and in line (5.30), we choose $a = \|\mathbf{s}_{i,k+1}\|^2$ and $b = 1$ and use $a^2 b \leq a^3 + b^3$.

Combining the result of (5.28) and (5.30) in equation (5.27), we have

$$\begin{aligned} f(w_{k+1}) - f(w_k) &\leq \left[-\frac{M\gamma^2 \eta_k}{4m^2} + \frac{L_2 \eta_k^3}{6m} + \frac{\eta_k \epsilon_g}{m} + \frac{\eta_k^2}{2m^2} ((m-1)L_1 + \epsilon_H) \right] \sum_i \|\mathbf{s}_{i,k+1}\|^3 \\ &\quad + \left[\frac{\eta_k \epsilon_g}{m} + \frac{\eta_k^2}{2m^2} ((m-1)L_1 + \epsilon_H) \right] \end{aligned}$$

Now we consider that $\lambda = \left(\frac{M\gamma}{4\eta_k m} - \frac{L_2}{6} - \frac{1}{2\eta_k m} ((m-1)L_1 + \epsilon_H) - \frac{\epsilon_g}{\eta_k^2} \right)$. We can assure $\lambda > 0$ by choosing $M \geq \frac{4\eta_k m}{\gamma} \left(\frac{L_2}{6} + \frac{1}{2\eta_k m} ((m-1)L_1 + \epsilon_H) + \frac{\epsilon_g}{\eta_k^2} \right)$.

Now we have

$$\frac{1}{m} \sum_i \|\eta_k \mathbf{s}_{i,k+1}\|^3 \leq \frac{1}{\lambda} \left[f(w_k) - f(w_{k+1}) + \frac{\eta_k \epsilon_g}{m} + \frac{\eta_k^2}{2m^2} ((m-1)L_1 + \epsilon_H) \right]$$

Now we consider the step k_0 , where $k_0 = \arg \min_{0 \leq k \leq T-1} \|w_{k+1} - w_k\| = \arg \min_{0 \leq k \leq T-1} \|\eta_k \mathbf{s}_{k+1}\|$.

$$\begin{aligned}
 \min_{0 \leq k \leq T} \|w_{k+1} - w_k\|^3 &= \min_{0 \leq k \leq T} \|\eta_k \mathbf{s}_{k+1}\|^3 \\
 &\leq \frac{1}{m} \sum_{i=1}^m \|\eta_{k_0} \mathbf{s}_{i, k_0+1}\|^3 \\
 &\leq \frac{1}{T} \sum_{k=0}^{T-1} \frac{1}{m} \sum_{i=1}^m \|\eta_k \mathbf{s}_{i, k+1}\|^3 \\
 &\leq \frac{1}{T} \sum_{k=0}^{T-1} \frac{1}{\lambda} \left[f(w_k) - f(w_{k+1}) + \frac{\eta_k^2}{2m} ((m-1)L_1 + \epsilon_H) + \frac{\eta_k \epsilon_g}{m} \right] \\
 &= \frac{1}{T} \left[\frac{f(w_0) - f(w_T)}{\lambda} + \frac{\sum_{k=0}^{T-1} \eta_k^2}{2m\lambda} ((m-1)L_1 + \epsilon_H) + \frac{\sum_{k=0}^{T-1} \eta_k \epsilon_g}{m\lambda} \right] \\
 &\leq \frac{1}{T} \left[\frac{f(w_0) - f^*}{\lambda} + \frac{\sum_{k=0}^{T-1} \eta_k^2}{2m^2\lambda} ((m-1)L_1 + \epsilon_H) + \frac{\sum_{k=0}^{T-1} \eta_k \epsilon_g}{m\lambda} \right]
 \end{aligned}$$

Based on the calculation above, we have

$$\begin{aligned}
 \|w_{k_0+1} - w_{k_0}\| &\leq \left(\frac{1}{m} \sum_{i=1}^m \|\eta_{k_0} \mathbf{s}_{i, k_0+1}\|^3 \right)^{\frac{1}{3}} \\
 &\leq \frac{1}{T^{\frac{1}{3}}} \left[\frac{f(w_0) - f^*}{\lambda} + \frac{\sum_{k=0}^{T-1} \eta_k^2}{2m^2\lambda} ((m-1)L_1 + \epsilon_H) + \frac{\sum_{k=0}^{T-1} \eta_k \epsilon_g}{m\lambda} \right]^{\frac{1}{3}}
 \end{aligned}$$

With the proper choice step-size η_k , we choose

$$\Psi = \left[\frac{f(w_0) - f^*}{\lambda} + \frac{\sum_{k=0}^{T-1} \eta_k^2}{2m^2\lambda} ((m-1)L_1 + \epsilon_H) + \frac{\sum_{k=0}^{T-1} \eta_k \epsilon_g}{m\lambda} \right]^{\frac{1}{3}} \text{ and have}$$

$$\|w_{k_0+1} - w_{k_0}\| \leq \left(\frac{1}{m} \sum_{i=1}^m \|\eta_{k_0} \mathbf{s}_{i, k_0+1}\|^3 \right)^{\frac{1}{3}} \leq \frac{\Psi}{T^{\frac{1}{3}}} \tag{5.31}$$

Now the gradient condition

$$\begin{aligned}
 & \|\nabla f(w_{k+1})\| \\
 = & \left\| \nabla f(w_{k+1}) - \frac{1}{m} \sum_{i=1}^m \mathbf{g}_{i,k} - \frac{1}{m} \sum_{i=1}^m (\gamma \mathbf{H}_{i,k+1} \mathbf{s}_{i,k+1} - \frac{M\gamma^2}{2} \|\mathbf{s}_{i,k+1}\| \mathbf{s}_{i,k+1}) \right\| \quad (5.32) \\
 \leq & \left\| \nabla f(w_{k+1}) - \nabla f(w_k) - \nabla^2 f(w_k)(x_{k+1} - x_k) \right\| + \left\| \frac{1}{m} \sum_{i=1}^m (\mathbf{g}_{i,k} - \nabla f(w_k)) \right\| \\
 & + \left\| \nabla^2 f(w_k)(x_{k+1} - x_k) - \gamma \frac{1}{m} \sum_{i=1}^m \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \right\| + \left\| \frac{1}{m} \sum_{i=1}^m \frac{M\gamma^2}{2} \|\mathbf{s}_{i,k+1}\| \mathbf{s}_{i,k+1} \right\| \\
 \leq & \frac{L_2}{2} \|\eta_k \mathbf{s}_{k+1}\|^2 + \left\| \frac{\eta_k}{m} \sum_{i=1}^m \nabla^2 f(w_k) \mathbf{s}_{i,k+1} - \frac{\gamma}{m} \sum_{i=1}^m \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \right\| + \frac{M\gamma^2}{2m} \sum_i \|\mathbf{s}_{i,k+1}\|^2 + \epsilon_g \quad (5.33)
 \end{aligned}$$

$$\begin{aligned}
 \leq & \left(\frac{L_2 \eta_k^2}{2m} + \frac{M\gamma^2}{2m} \right) \sum_i \|\mathbf{s}_{i,k+1}\|^2 + \left\| \frac{\eta_k}{m} \sum_{i=1}^m \nabla^2 f(w_k) \mathbf{s}_{i,k+1} - \frac{\gamma}{m} \sum_{i=1}^m \nabla^2 f(w_k) \mathbf{s}_{i,k+1} \right\| \\
 & + \left\| \frac{\gamma}{m} \sum_{i=1}^m \nabla^2 f(w_k) \mathbf{s}_{i,k+1} - \frac{\gamma}{m} \sum_{i=1}^m \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \right\| + \epsilon_g \\
 \leq & \left(\frac{L_2 \eta_k^2}{2m} + \frac{M\gamma^2}{2m} \right) \sum_i \|\mathbf{s}_{i,k+1}\|^2 + \frac{(\eta_k - \gamma)L_1}{m} \sum_i \|\mathbf{s}_{i,k+1}\| + \frac{\gamma \epsilon_H}{m} \sum_i \|\mathbf{s}_{i,k+1}\| + \epsilon_g \quad (5.34)
 \end{aligned}$$

$$\begin{aligned}
 \leq & \left(\frac{L_2 \eta_k^2}{2} + \frac{M\gamma^2}{2} \right) \frac{1}{m} \sum_i \|\mathbf{s}_{i,k+1}\|^2 + (|\eta_k - \gamma|L_1 + \gamma \epsilon_H) \frac{1}{m} \sum_i \|\mathbf{s}_{i,k+1}\| + \epsilon_g \\
 \leq & \left(\frac{L_2}{2} + \frac{M\gamma^2}{2\eta_k^2} \right) \frac{1}{m} \sum_i \|\eta_k \mathbf{s}_{i,k+1}\|^2 + \left(|1 - \frac{\gamma}{\eta_k}|L_1 + \frac{\gamma}{\eta_k} \epsilon_H \right) \frac{1}{m} \sum_i \|\eta_k \mathbf{s}_{i,k+1}\| + \epsilon_g \\
 \leq & \left(\frac{L_2}{2} + \frac{M\gamma^2}{2\eta_k^2} \right) \frac{1}{m} \sum_i \|\eta_k \mathbf{s}_{i,k+1}\|^2 + \left(|1 - \frac{\gamma}{\eta_k}|L_1 + \frac{\gamma}{\eta_k} \epsilon_H \right) \frac{1}{m} \sum_i \|\eta_k \mathbf{s}_{i,k+1}\|^3 \quad (5.35)
 \end{aligned}$$

$$\begin{aligned}
 & + \left(\epsilon_g + \left(|1 - \frac{\gamma}{\eta_k}|L_1 + \frac{\gamma}{\eta_k} \epsilon_H \right) \right) \\
 \leq & \left(\frac{L_2}{2} + \frac{M\gamma^2}{2\eta_k^2} \right) \left[\frac{1}{m} \sum_i \|\eta_k \mathbf{s}_{i,k+1}\|^3 \right]^{\frac{2}{3}} + \left(|1 - \frac{\gamma}{\eta_k}|L_1 + \frac{\gamma}{\eta_k} \epsilon_H \right) \frac{1}{m} \sum_i \|\eta_k \mathbf{s}_{i,k+1}\|^3 \quad (5.36) \\
 & + \left(\epsilon_g + \left(|1 - \frac{\gamma}{\eta_k}|L_1 + \frac{\gamma}{\eta_k} \epsilon_H \right) \right) \quad (5.37)
 \end{aligned}$$

In line (5.32), we use first order condition described in equation (5.21). In line (5.33), we apply the result (5.13) from Lemma 25 and the approximate gradient condition from Assumption 18. In line (5.34), we apply the approximate Hessian condition from Assumption 19. We apply the

inequality of (5.12) in line (5.37). At step k_0 , by choosing $\eta_k = \gamma$, we have

$$\begin{aligned} \|\nabla f(w_{k_0+1})\| &\leq \left(\frac{L_2}{2} + \frac{M}{2}\right) \frac{\Psi^2}{T^{\frac{2}{3}}} + \epsilon_H \frac{\Psi^3}{T} + (\epsilon_g + \epsilon_H) \\ &= \frac{\Psi_1}{T^{\frac{2}{3}}} + \frac{\Psi_2}{T} + (\epsilon_g + \epsilon_H) \end{aligned} \quad (5.38)$$

where, $\Psi_1 = \left(\frac{L_2}{2} + \frac{M}{2}\right) \Psi^2$ and $\Psi_2 = \epsilon_H \Psi^3$. The Hessian bound

$$\begin{aligned} \lambda_{\min}(\nabla^2 f(w_{k+1})) &= \frac{1}{m} \sum_{i=1}^m \lambda_{\min} [\nabla^2 f(w_{k+1})] \\ &= \frac{1}{m} \sum_{i=1}^m \lambda_{\min} [\mathbf{H}_{i,k} - (\mathbf{H}_{i,k} - \nabla^2 f(w_{k+1}))] \\ &\geq \frac{1}{m} \sum_{i=1}^m [\lambda_{\min}(\mathbf{H}_{i,k}) - \|\mathbf{H}_{i,k} - \nabla^2 f(w_{k+1})\|] \\ &\geq \frac{1}{m} \sum_{i=1}^m \lambda_{\min}(\mathbf{H}_{i,k}) - \frac{1}{m} \sum_{i=1}^m \|\mathbf{H}_{i,k} - \nabla^2 f(w_{k+1})\| \end{aligned} \quad (5.39)$$

$$\begin{aligned} &\geq \frac{1}{m} \sum_{i=1}^m -\frac{M\gamma}{2} \|\mathbf{s}_{i,k+1}\| - \frac{1}{m} \sum_{i=1}^m \|\mathbf{H}_{i,k} - \nabla^2 f(w_k)\| \\ &\quad - \frac{1}{m} \sum_{i=1}^m \|\nabla^2 f(w_k) - \nabla^2 f(w_{k+1})\| \\ &\geq \frac{1}{m} \sum_{i=1}^m -\frac{M\gamma}{2} \|\mathbf{s}_{i,k+1}\| - \epsilon_H - \frac{1}{m} \sum_{i=1}^m \|\nabla^2 f(w_k) - \nabla^2 f(w_{k+1})\| \end{aligned} \quad (5.40)$$

$$\begin{aligned} &\geq \frac{1}{m} \sum_{i=1}^m -\frac{M\gamma}{2} \|\mathbf{s}_{i,k+1}\| - \epsilon_H - \frac{1}{m} \sum_{i=1}^m L_2 \|w_k - w_{k+1}\| \\ &\geq \left(-\frac{M\gamma}{2\eta_k} - L_2\right) \frac{1}{m} \sum_{i=1}^m \|\eta_k \mathbf{s}_{i,k+1}\| - \epsilon_H \\ &\geq -\left(\frac{M\gamma}{2\eta_k} + L_2\right) \left(\frac{1}{m} \sum_{i=1}^m \|\eta_k \mathbf{s}_{i,k+1}\|^3\right)^{1/3} - \epsilon_H \end{aligned} \quad (5.41)$$

Equation (5.39) follows from Weyl's inequality. We apply the Hessian approximation from the Assumption 19 in equation (5.40). In equation (5.41), we apply the result described in (5.12).

At step k_0 , by choosing $\eta_k = \gamma$, we have

$$\begin{aligned} \lambda_{\min}(\nabla^2 f(w_{k_0+1})) &\leq -\left(\frac{M}{2} + L_2\right) \frac{\Psi}{T^{\frac{1}{3}}} - \epsilon_H \\ &= -\frac{\Psi_3}{T^{\frac{1}{3}}} - \epsilon_H \end{aligned} \quad (5.42)$$

where, $\Psi_3 = \left(\frac{M}{2} + L_2\right) \Psi$.

5.7.3 Proof of Theorem 11

We consider the following

$$\begin{aligned}
 & f(w_{k+1}) - f(w_k) \\
 & \leq \nabla f(w_k)^T (w_{k+1} - w_k) + \frac{1}{2} (w_{k+1} - w_k)^T \nabla^2 f(w_k) (w_{k+1} - w_k) + \frac{L_2}{6} \|w_{k+1} - w_k\|^3 \\
 & = \underbrace{\frac{\eta_k}{|\mathcal{U}|} \nabla f(w_k)^T \sum_{i \in \mathcal{U}} \mathbf{s}_{i,k+1}}_{\text{Term1}} + \underbrace{\frac{\eta_k^2}{2|\mathcal{U}|^2} \left(\sum_{i \in \mathcal{U}} \mathbf{s}_{i,k+1} \right)^T \nabla^2 f(w_k) \left(\sum_{i \in \mathcal{U}} \mathbf{s}_{i,k+1} \right)}_{\text{Term2}} + \underbrace{\frac{L_2}{6} \left\| \frac{\eta_k}{|\mathcal{U}|} \sum_{i \in \mathcal{U}} \mathbf{s}_{i,k+1} \right\|^3}_{\text{Term3}}
 \end{aligned} \tag{5.43}$$

In line (5.43), we expand the update $w_{k+1} - w_k = \frac{\eta_k}{|\mathcal{U}|} \sum_{i \in \mathcal{U}} \mathbf{s}_{i,k+1}$. Also we use the following fact.

$$|\mathcal{U}| = |\mathcal{U} \cap \mathcal{M}| + |\mathcal{U} \cap \mathcal{B}| \tag{5.44}$$

$$|\mathcal{M}| = |\mathcal{U} \cap \mathcal{M}| + |\mathcal{T} \cap \mathcal{M}| \tag{5.45}$$

Combining both the equations (5.44) and (5.45), we have

$$|\mathcal{U}| = |\mathcal{M}| - |\mathcal{T} \cap \mathcal{M}| + |\mathcal{U} \cap \mathcal{B}| \tag{5.46}$$

We use the fact of (5.46) to bound each term in equation (5.43). First, consider the Term 1,

$$\begin{aligned}
 & \frac{\eta_k}{|\mathcal{U}|} \nabla f(w_k)^T \sum_{i \in \mathcal{U}} \mathbf{s}_{i,k+1} \\
 & = \frac{\eta_k}{(1-\beta)m} \nabla f(w_k)^T \left[\sum_{i \in \mathcal{M}} \mathbf{s}_{i,k+1} - \sum_{i \in \mathcal{M} \cap \mathcal{T}} \mathbf{s}_{i,k+1} + \sum_{i \in \mathcal{U} \cap \mathcal{B}} \mathbf{s}_{i,k+1} \right] \\
 & = \frac{\eta_k}{(1-\beta)m} \left[\sum_{i \in \mathcal{M}} \mathbf{g}_{i,k}^T \mathbf{s}_{i,k+1} - \sum_{i \in \mathcal{M} \cap \mathcal{T}} \nabla f(w_k)^T \mathbf{s}_{i,k+1} + \sum_{i \in \mathcal{U} \cap \mathcal{B}} \nabla f(w_k)^T \mathbf{s}_{i,k+1} \right. \\
 & \quad \left. + \frac{\gamma}{2} \sum_{i \in \mathcal{M}} \mathbf{s}_{i,k+1}^T \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \right] - \frac{\eta_k \gamma}{2(1-\beta)m} \sum_{i \in \mathcal{M}} \mathbf{s}_{i,k+1}^T \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \\
 & \quad + \frac{\eta_k}{(1-\beta)m} \sum_{i \in \mathcal{M}} (\nabla f(w_k) - \mathbf{g}_{i,k})^T \mathbf{s}_{i,k+1}
 \end{aligned} \tag{5.47}$$

$$\begin{aligned}
 & \leq -\frac{M\gamma^2\eta_k}{4(1-\beta)m} \left[\sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\|^3 \right] + \frac{\eta_k \epsilon_g}{(1-\beta)m} \sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\| - \frac{\eta_k \gamma}{2(1-\beta)m} \sum_{i \in \mathcal{M}} \mathbf{s}_{i,k+1}^T \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \\
 & \quad + \frac{\eta_k L}{(1-\beta)m} \left(\sum_{i \in \mathcal{M} \cap \mathcal{T}} \|\mathbf{s}_{i,k+1}\| + \sum_{i \in \mathcal{U} \cap \mathcal{B}} \|\mathbf{s}_{i,k+1}\| \right)
 \end{aligned} \tag{5.48}$$

We use the following facts in (5.48).

- $\mathbf{g}_{i,k}^T \mathbf{s}_{i,k+1} + \frac{\gamma}{2} \mathbf{s}_{i,k+1}^T \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \leq -\frac{M}{4} \gamma^2 \|\mathbf{s}_{i,k+1}\|^3$ and sum over the set \mathcal{M} .
- The gradient approximation described in Assumption 18.
- As the function f is L -Lipschitz, the gradient is bounded.

Now we bound Term 3 as follows,

$$\begin{aligned} \frac{L_2}{6} \left\| \frac{\eta_k}{|\mathcal{U}|} \sum_{i \in \mathcal{U}} \mathbf{s}_{i,k+1} \right\|^3 &\leq \frac{L_2 \eta_k^3}{6(1-\beta)m} \sum_{i \in \mathcal{U}} \|\mathbf{s}_{i,k+1}\|^3 \quad (5.49) \\ &\leq \frac{L_2 \eta_k^3}{6(1-\beta)m} \left[\sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\|^3 - \sum_{i \in \mathcal{M} \cap \mathcal{T}} \|\mathbf{s}_{i,k+1}\|^3 + \sum_{i \in \mathcal{U} \cap \mathcal{B}} \|\mathbf{s}_{i,k+1}\|^3 \right] \quad (5.50) \end{aligned}$$

In line (5.49), we use the inequality describe in (5.10) and in line (5.50), we split the sum using (5.46).

Finally, we bound Term 2

$$\begin{aligned} &\frac{\eta_k^2}{2|\mathcal{U}|^2} \left(\sum_{i \in \mathcal{U}} \mathbf{s}_{i,k+1} \right)^T \nabla^2 f(w_k) \left(\sum_{i \in \mathcal{U}} \mathbf{s}_{i,k+1} \right) \\ &= \frac{\eta_k^2}{2(1-\beta)^2 m^2} \left(\sum_{i \in \mathcal{U}} \mathbf{s}_{i,k+1}^T \nabla f(w_k) \mathbf{s}_{i,k+1} + \sum_{i \neq j \in \mathcal{U}} \mathbf{s}_{i,k+1}^T \nabla f(w_k) \mathbf{s}_{j,k+1} \right) \\ &= \frac{\eta_k^2}{2(1-\beta)^2 m^2} \left(\sum_{i \in \mathcal{M}} \mathbf{s}_{i,k+1}^T (\nabla^2 f(w_k) - \mathbf{H}_{i,k}) \mathbf{s}_{i,k+1} - \sum_{i \in \mathcal{M} \cap \mathcal{T}} \mathbf{s}_{i,k+1}^T \nabla^2 f(w_k) \mathbf{s}_{i,k+1} \right. \\ &\quad \left. + \sum_{i \in \mathcal{U} \cap \mathcal{B}} \mathbf{s}_{i,k+1}^T \nabla^2 f(w_k) \mathbf{s}_{i,k+1} \right) \\ &\quad + \frac{\eta_k^2}{2(1-\beta)^2 m^2} \sum_{i \neq j \in \mathcal{U}} \mathbf{s}_{i,k+1}^T \nabla^2 f(w_k) \mathbf{s}_{j,k+1} + \frac{\eta_k^2}{2(1-\beta)^2 m^2} \sum_{i \in \mathcal{M}} \mathbf{s}_{i,k+1}^T \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \quad (5.51) \\ &\leq \frac{\eta_k^2}{2(1-\beta)^2 m^2} \left[\epsilon_H \sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\|^2 + L_1 \sum_{i \in \mathcal{M} \cap \mathcal{T}} \|\mathbf{s}_{i,k+1}\|^2 + L_1 \sum_{i \in \mathcal{U} \cap \mathcal{B}} \|\mathbf{s}_{i,k+1}\|^2 \right] \\ &\quad + \frac{\eta_k^2}{2(1-\beta)^2 m^2} \left[L_1 \left\| \sum_{i \in \mathcal{U}} \mathbf{s}_{i,k+1} \right\|^2 - L_1 \sum_{i \in \mathcal{U}} \|\mathbf{s}_{i,k+1}\|^2 \right] + \frac{\eta_k^2}{2(1-\beta)^2 m^2} \sum_{i \in \mathcal{M}} \mathbf{s}_{i,k+1}^T \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \quad (5.52) \end{aligned}$$

$$\begin{aligned} &\leq \frac{\eta_k^2}{2(1-\beta)^2 m^2} \left[((1-\beta)m - 1)L_1 + \epsilon_H \right] \sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\|^2 + (1-\beta)m + 2) L_1 \sum_{i \in \mathcal{M} \cap \mathcal{T}} \|\mathbf{s}_{i,k+1}\|^2 \\ &\quad + (1-\beta)m L_1 \sum_{i \in \mathcal{U} \cap \mathcal{B}} \|\mathbf{s}_{i,k+1}\|^2 + \frac{\eta_k^2}{2(1-\beta)^2 m^2} \sum_{i \in \mathcal{M}} \mathbf{s}_{i,k+1}^T \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \quad (5.53) \end{aligned}$$

Now we collect the terms from equations (5.48), (5.50) and (5.53). First we focus on the terms that are summed over the set \mathcal{M} .

$$\begin{aligned}
 & \left(-\frac{M\gamma^2\eta_k}{4(1-\beta)m} + \frac{L_2\eta_k^3}{6(1-\beta)m} \right) \left[\sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\|^3 \right] - \frac{\eta_k\gamma}{2(1-\beta)m} \left(\gamma - \frac{\eta_k}{(1-\beta)m} \right) \sum_{i \in \mathcal{M}} \mathbf{s}_{i,k+1}^T \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \\
 & \quad + \frac{\eta_k\epsilon_g}{(1-\beta)m} \sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\| + \frac{\eta_k^2}{2(1-\beta)^2m^2} \left((1-\beta)m - 1 \right) L_1 + \epsilon_H \sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\|^2 \\
 \leq & \left(-\frac{M\gamma^2\eta_k}{4(1-\beta)m} + \frac{L_2\eta_k^3}{6(1-\beta)m} \right) \left[\sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\|^3 \right] \\
 & + \left[\frac{\eta_k\epsilon_g}{(1-\beta)m} + \frac{\eta_k^2}{2(1-\beta)^2m^2} \left((1-\beta)m - 1 \right) L_1 + \epsilon_H \right] \sum_{i \in \mathcal{M}} \left(\|\mathbf{s}_{i,k+1}\|^3 + 1 \right) \tag{5.54}
 \end{aligned}$$

$$\begin{aligned}
 = & \left(-\frac{M\gamma^2\eta_k}{4(1-\beta)m} + \frac{L_2\eta_k^3}{6(1-\beta)m} + \frac{\eta_k\epsilon_g}{(1-\beta)m} + \frac{\eta_k^2}{2(1-\beta)^2m^2} \left((1-\beta)m - 1 \right) L_1 + \epsilon_H \right) \\
 & \times \left[\sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\|^3 \right] + \frac{\eta_k\epsilon_g(1-\alpha)}{(1-\beta)} + \frac{\eta_k^2}{2(1-\beta)^2m} \left((1-\beta)m - 1 \right) L_1 + \epsilon_H (1-\alpha) \tag{5.55}
 \end{aligned}$$

In line (5.54), we use the fact of 5.22 and $\|\mathbf{s}_{i,k+1}\| \leq \|\mathbf{s}_{i,k+1}\|^3 + 1$ and $\|\mathbf{s}_{i,k+1}\|^2 \leq \|\mathbf{s}_{i,k+1}\|^3 + 1$ following the inequality $a^b \leq a^3 + b^3$. Also, we use the fact that $|\mathcal{M}| \leq (1-\alpha)m$.

Now we consider the terms with the set $\mathcal{M} \cap \mathcal{T}$ and $\mathcal{U} \cap \mathcal{B}$.

$$\begin{aligned}
 & \frac{\eta_k L}{(1-\beta)m} \left(\sum_{i \in \mathcal{M} \cap \mathcal{T}} \|\mathbf{s}_{i,k+1}\| + \sum_{i \in \mathcal{U} \cap \mathcal{B}} \|\mathbf{s}_{i,k+1}\| \right) \tag{5.56} \\
 & + \frac{L_2\eta_k^3}{6(1-\beta)m} \left[- \sum_{i \in \mathcal{M} \cap \mathcal{T}} \|\mathbf{s}_{i,k+1}\|^3 + \sum_{i \in \mathcal{U} \cap \mathcal{B}} \|\mathbf{s}_{i,k+1}\|^3 \right] \\
 & + \frac{\eta_k^2}{2(1-\beta)^2m^2} \left[(1-\beta)m + 2 \right) L_1 \sum_{i \in \mathcal{M} \cap \mathcal{T}} \|\mathbf{s}_{i,k+1}\|^2 + (1-\beta)m L_1 \sum_{i \in \mathcal{U} \cap \mathcal{B}} \|\mathbf{s}_{i,k+1}\|^2 \right] \\
 \leq & \left(\frac{\eta_k L}{(1-\beta)m} - \frac{L_2\eta_k^3}{6(1-\beta)m} + \frac{\eta_k^2}{2(1-\beta)^2m^2} \left((1-\beta)m + 2 \right) L_1 \right) \sum_{i \in \mathcal{M} \cap \mathcal{T}} \|\mathbf{s}_{i,k+1}\|^3 \\
 & + \left(\frac{\eta_k L}{(1-\beta)m} + \frac{L_2\eta_k^3}{6(1-\beta)m} + \frac{\eta_k^2}{2(1-\beta)^2m^2} (1-\beta)m L_1 \right) \sum_{i \in \mathcal{U} \cap \mathcal{B}} \|\mathbf{s}_{i,k+1}\|^3 \\
 & + \frac{\eta_k L}{(1-\beta)m} + \frac{\eta_k^2}{2(1-\beta)^2m} \left[(1-\beta)m + 2(1-\alpha) \right] L_1 \tag{5.57}
 \end{aligned}$$

Now as $\beta > \alpha$, at least one good machine is trimmed. So the norm of all the machine update in the set \mathcal{U} is upper bounded by the maximum norm of the good machine. We upper bound the terms as

follows,

$$\sum_{i \in \mathcal{U} \cap \mathcal{B}} \|\mathbf{s}_{i,k+1}\|^3 \leq \alpha m \max_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\|^3 \leq \alpha m \sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\|^3$$

and
$$\sum_{i \in \mathcal{M} \cap \mathcal{T}} \|\mathbf{s}_{i,k+1}\|^3 \leq \sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\|^3$$

We have

$$\begin{aligned} & \left(\frac{\eta_k L(1 + \alpha m)}{(1 - \beta)m} + \frac{(\alpha m - 1)L_2 \eta_k^3}{6(1 - \beta)m} + \frac{\eta_k^2}{2(1 - \beta)^2 m^2} ((1 - \beta)m + \alpha(1 - \beta)m^2 + 2)L_1 \right) \\ & \times \sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\|^3 + \frac{\eta_k L}{(1 - \beta)m} + \frac{\eta_k^2}{2(1 - \beta)^2 m} [(1 - \beta)m + 2(1 - \alpha)L_1] \end{aligned} \quad (5.58)$$

We combine the results (5.58) and (5.55). We have

$$f(w_{k+1}) - f(w_k) \leq -\lambda_{byz} \frac{1}{(1 - \alpha)m} \sum_{i \in \mathcal{M}} \|\eta_k \mathbf{s}_{i,k+1}\|^3 + \lambda_{floor} \quad (5.59)$$

where

$$\begin{aligned} \lambda_{byz} &= \left[\frac{M\gamma}{4\eta_k(1 - \beta)^2 m^2} - \frac{(L(1 + \alpha m) + \epsilon_g)}{\eta_k^2(1 - \beta)m} - \frac{\alpha m L_2}{6(1 - \beta)m} - \frac{1}{2\eta_k(1 - \beta)^2 m^2} (1 - \beta)m L_1 \right. \\ & \quad \left. + \alpha(1 - \beta)m^2 + \epsilon_H \right] (1 - \alpha)m \\ \lambda_{floor} &= \frac{\eta_k \epsilon_g (1 - \alpha)}{(1 - \beta)} + \frac{\eta_k L}{(1 - \beta)} \\ & \quad + \frac{\eta_k^2}{2(1 - \beta)^2 m^2} [(2(1 - \beta)m + 1)L_1 + \epsilon_H](1 - \alpha)m + \alpha(1 - \beta)m^2 L_1 \end{aligned}$$

We maintain $\lambda_{byz} > 0$ by choosing

$$\begin{aligned} M &> \frac{4\eta_k(1 - \beta)m}{\gamma} \left[\frac{(L(1 + \alpha m) + \epsilon_g)}{\eta_k^2} + \frac{\alpha m L_2}{6} \right. \\ & \quad \left. + \frac{1}{2\eta_k(1 - \beta)m} ((1 - \beta)m L_1 + \alpha(1 - \beta)m^2 + \epsilon_H) \right] \end{aligned}$$

Now we can have the following results from the proof of Theorem 10 for step

$$k_0 = \arg \min_{0 \leq k \leq T-1} \|w_{k+1} - w_k\|$$

$$\begin{aligned} \|w_{k_0+1} - w_{k_0}\| &\leq \left[\frac{1}{(1 - \alpha)m} \sum_{i \in \mathcal{M}} \|\eta_{k_0} \mathbf{s}_{i,k_0+1}\|^3 \right]^{\frac{1}{3}} \\ &\leq \frac{1}{T^{\frac{1}{3}}} \left[\frac{f(w_0) - f^*}{\lambda_{byz}} + \frac{\sum_{k=0}^{T-1} \lambda_{floor}}{\lambda_{byz}} \right]^{\frac{1}{3}} \end{aligned} \quad (5.60)$$

$$= \frac{\Psi_{byz}}{T^{\frac{1}{3}}} \quad (5.61)$$

where, $\Psi_{byz} = \left[\frac{f(w_0) - f^*}{\lambda_{byz}} + \frac{\sum_{k=0}^{T-1} \lambda_{f\text{loor}}}{\lambda_{byz}} \right]^{\frac{1}{3}}$.

The gradient condition

$$\|\nabla f(w_{k+1})\| = \left\| \nabla f(w_{k+1}) - \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \mathbf{g}_{i,k} - \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \gamma \mathbf{H}_{i,k+1} \mathbf{s}_{i,k+1} - \frac{M\gamma^2}{2} \|\mathbf{s}_{i,k+1}\| \mathbf{s}_{i,k+1} \right\| \quad (5.62)$$

$$\begin{aligned} &\leq \left\| \nabla f(w_{k+1}) - \nabla f(w_k) - \nabla^2 f(w_k)(x_{k+1} - x_k) \right\| + \left\| \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} (\mathbf{g}_{i,k} - \nabla f(w_k)) \right\| \\ &+ \left\| \nabla^2 f(w_k)(x_{k+1} - x_k) - \gamma \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \right\| + \left\| \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \frac{M\gamma^2}{2} \|\mathbf{s}_{i,k+1}\| \mathbf{s}_{i,k+1} \right\| \\ &\leq \frac{L_2 \eta_k^2}{2} \left\| \frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{U}} \mathbf{s}_{i,k+1} \right\|^2 + \epsilon_g + \frac{M\gamma^2}{2} \frac{1}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\|^2 \\ &+ \left\| \frac{\eta_k}{|\mathcal{U}|} \sum_{i \in \mathcal{U}} \nabla^2 f(w_k) \mathbf{s}_{i,k+1} - \frac{\gamma}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla^2 f(w_k) \mathbf{s}_{i,k+1} \right\| \\ &+ \left\| \frac{\gamma}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla^2 f(w_k) \mathbf{s}_{i,k+1} - \frac{\gamma}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \mathbf{H}_{i,k} \mathbf{s}_{i,k+1} \right\| \quad (5.63) \end{aligned}$$

In line (5.62), we use the fact the first order optimal condition (5.21) holds for the good machines in the set \mathcal{M} . And in (5.63), we use the in exact gradient condition from Assumption 18 and the condition (5.23). Consider the term

$$\begin{aligned} &\left\| \frac{\eta_k}{|\mathcal{U}|} \sum_{i \in \mathcal{U}} \nabla^2 f(w_k) \mathbf{s}_{i,k+1} - \frac{\gamma}{|\mathcal{M}|} \sum_{i \in \mathcal{M}} \nabla^2 f(w_k) \mathbf{s}_{i,k+1} \right\| \\ &= \left\| \frac{\eta_k}{(1-\beta)m} \left[\sum_{i \in \mathcal{M}} \nabla^2 f(w_k) \mathbf{s}_{i,k+1} - \sum_{i \in \mathcal{M} \cap \mathcal{T}} \nabla^2 f(w_k) \mathbf{s}_{i,k+1} + \sum_{i \in \mathcal{B} \cap \mathcal{U}} \nabla^2 f(w_k) \mathbf{s}_{i,k+1} \right] \right. \\ &\quad \left. - \frac{\gamma}{(1-\alpha)m} \sum_{i \in \mathcal{M}} \nabla^2 f(w_k) \mathbf{s}_{i,k+1} \right\| \\ &\leq \left\| \left(\frac{\eta_k(1+\alpha m)}{(1-\beta)m} - \frac{\gamma}{(1-\alpha)m} \right) \sum_{i \in \mathcal{M}} \nabla^2 f(w_k) \mathbf{s}_{i,k+1} \right\| \\ &\leq \left(\frac{\eta_k(1+\alpha m)}{(1-\beta)m} - \frac{\gamma}{(1-\alpha)m} \right) L_1 \sum_{i \in \mathcal{M}} \nabla^2 \|\mathbf{s}_{i,k+1}\| \quad (5.64) \end{aligned}$$

We choose $\gamma = \frac{\eta_k(1-\alpha)}{(1-\beta)}(1 + \alpha m)$ making the term in equation (5.64) equals to 0. Now we have,

$$\frac{L_2\eta_k^2}{2} \left\| \frac{1}{|\mathcal{U}|} \sum_{i \in \mathcal{U}} \mathbf{s}_{i,k+1} \right\|^2 \leq \frac{L_2\eta_k^2}{2(1-\beta)m} \sum_{i \in \mathcal{U}} \|\mathbf{s}_{i,k+1}\|^2 \leq \frac{L_2(1+\alpha m)\eta_k^2}{2(1-\beta)m} \sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\|^2 \quad (5.65)$$

Putting the calculation of (5.64) and (5.65), in (5.63), we have,

$$\|\nabla f(w_{k+1})\| \quad (5.66)$$

$$\leq \left(\frac{L_2(1+\alpha m)\eta_k^2}{2(1-\beta)m} + \frac{M\gamma^2}{2(1-\alpha)m} \right) \sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\|^2 + \epsilon_g + \frac{\gamma\epsilon_H}{(1-\alpha)m} \sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\| \quad (5.67)$$

$$\begin{aligned} &\leq \left(\frac{L_2(1+\alpha m)(1-\alpha)}{2(1-\beta)} + \frac{M\gamma^2}{2\eta_k^2} \right) \frac{1}{(1-\alpha)m} \sum_{i \in \mathcal{M}} \|\eta_k \mathbf{s}_{i,k+1}\|^2 + \epsilon_g \\ &\quad + \frac{\gamma\epsilon_H}{\eta_k} \frac{1}{(1-\alpha)m} \sum_{i \in \mathcal{M}} \|\eta_k \mathbf{s}_{i,k+1}\| \\ &\leq \left(\frac{L_2(1+\alpha m)(1-\alpha)}{2(1-\beta)} + \frac{M(1+\alpha m)^2}{2} \left(\frac{(1-\alpha)^2}{(1-\beta)} \right) \right) \left[\frac{1}{(1-\alpha)m} \sum_{i \in \mathcal{M}} \|\eta_k \mathbf{s}_{i,k+1}\|^3 \right]^{\frac{2}{3}} \\ &\quad + \frac{(1-\alpha)}{(1-\beta)}(1+\alpha m)\epsilon_H \left[\frac{1}{(1-\alpha)m} \sum_{i \in \mathcal{M}} \|\eta_k \mathbf{s}_{i,k+1}\|^3 \right] + \epsilon_g + \frac{(1-\alpha)}{(1-\beta)}(1+\alpha m)\epsilon_H \quad (5.68) \end{aligned}$$

We use the power mean inequality described in (5.12) in line (5.68). Then at step k_0 , we have,

$$\|\nabla f(w_{k_0+1})\| \leq \frac{\Psi_{1,byz}}{T^{\frac{2}{3}}} + \frac{\Psi_{2,byz}}{T} + \epsilon_g + \frac{(1-\alpha)}{(1-\beta)}(1+\alpha m)\epsilon_H, \quad (5.69)$$

where

$$\begin{aligned} \Psi_{1,byz} &= \left(\frac{L_2(1+\alpha m)(1-\alpha)}{2(1-\beta)} + \frac{M(1+\alpha m)^2}{2} \left(\frac{(1-\alpha)^2}{(1-\beta)} \right) \right) \Psi_{byz}^2 \\ \Psi_{2,byz} &= \frac{(1-\alpha)}{(1-\beta)}(1+\alpha m)\epsilon_H \Psi_{byz}^3 \end{aligned}$$

The Hessian bound is

$$\begin{aligned}
 & \lambda_{\min}(\nabla^2 f(w_{k+1})) \\
 &= \frac{1}{(1-\alpha)m} \sum_{i \in \mathcal{M}} \lambda_{\min} [\nabla^2 f(w_{k+1})] \\
 &= \frac{1}{(1-\alpha)m} \sum_{i \in \mathcal{M}} \lambda_{\min} [\mathbf{H}_{i,k} - (\mathbf{H}_{i,k} - \nabla^2 f(w_{k+1}))] \\
 &\geq \frac{1}{(1-\alpha)m} \sum_{i \in \mathcal{M}} [\lambda_{\min}(\mathbf{H}_{i,k}) - \|\mathbf{H}_{i,k} - \nabla^2 f(w_{k+1})\|] \tag{5.70}
 \end{aligned}$$

$$\begin{aligned}
 &\geq \frac{1}{(1-\alpha)m} \sum_{i \in \mathcal{M}} \lambda_{\min}(\mathbf{H}_{i,k}) - \frac{1}{(1-\alpha)m} \sum_{i \in \mathcal{M}} \|\mathbf{H}_{i,k} - \nabla^2 f(w_{k+1})\| \\
 &\geq \frac{1}{(1-\alpha)m} \sum_{i \in \mathcal{M}} -\frac{M\gamma}{2} \|\mathbf{s}_{i,k+1}\| - \frac{1}{(1-\alpha)m} \sum_{i \in \mathcal{M}} \|\mathbf{H}_{i,k} - \nabla^2 f(w_k)\| \\
 &\quad - \frac{1}{(1-\alpha)m} \sum_{i \in \mathcal{M}} \|\nabla^2 f(w_k) - \nabla^2 f(w_{k+1})\| \tag{5.71}
 \end{aligned}$$

$$\geq \frac{1}{(1-\alpha)m} \sum_{i \in \mathcal{M}} -\frac{M\gamma}{2} \|\mathbf{s}_{i,k+1}\| - \epsilon_H - \frac{1}{(1-\alpha)m} \sum_{i \in \mathcal{M}} L_2 \|w_k - w_{k+1}\| \tag{5.72}$$

$$\begin{aligned}
 &\geq -\frac{M\gamma}{2(1-\alpha)m} \sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\| - \epsilon_H - L_2 \left\| \frac{\eta_k}{(1-\beta)m} \sum_{i \in \mathcal{U}} \mathbf{s}_{i,k+1} \right\| \\
 &\geq -\frac{M\gamma}{2(1-\alpha)m} \sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\| - \epsilon_H - L_2 \frac{1}{(1-\beta)m} \sum_{i \in \mathcal{U}} \|\eta_k \mathbf{s}_{i,k+1}\| \\
 &\geq -\frac{M\gamma}{2(1-\alpha)m} \sum_{i \in \mathcal{M}} \|\mathbf{s}_{i,k+1}\| - \epsilon_H - L_2 \frac{(1+\alpha m)}{(1-\beta)m} \sum_{i \in \mathcal{M}} \|\eta_k \mathbf{s}_{i,k+1}\| \\
 &\geq -\left(\frac{M\gamma}{2\eta_k(1-\alpha)m} + L_2 \frac{(1+\alpha m)}{(1-\beta)m} \right) \sum_{i \in \mathcal{M}} \|\eta_k \mathbf{s}_{i,k+1}\| - \epsilon_H \\
 &\geq -\left(\frac{M\gamma}{2\eta_k} + L_2 \frac{(1+\alpha m)(1-\alpha)}{(1-\beta)} \right) \left[\frac{1}{(1-\alpha)m} \sum_{i \in \mathcal{M}} \|\eta_k \mathbf{s}_{i,k+1}\|^3 \right]^{\frac{1}{3}} - \epsilon_H \tag{5.73}
 \end{aligned}$$

In (5.70), we use the Weyl's inequality. In (5.72), we use the fact that Hessian is Lipschitz continuous. In (5.73), we use the power mean inequality described in (5.12). At step k_0 , we have

$$\lambda_{\min}(\nabla^2 f(w_{k_0+1})) \geq -\frac{\Psi_{3,byz}}{T^{\frac{1}{3}}} - \epsilon_H \tag{5.74}$$

where

$$\Psi_{3,byz} = \left(\frac{M(1-\alpha)}{2(1-\beta)}(1+\alpha m) + L_2 \frac{(1+\alpha m)(1-\alpha)}{(1-\beta)} \right) \Psi_{byz}$$

5.7.4 Additional Experiments

Here we compare our algorithm with the ByzantinePGD [31] in the byzantine setup. In the Table 5.1, we show the number iterations that is required by each algorithm to reach the stopping criteria based on the gradient norm. We choose the four adversarial attacks and the experimental setup that are described in Section 5.6.

	Gaussian Noise			Flipped Label		
	10%	15%	20%	10%	15%	20%
[31]	199.2	198.3	211.2	199.7	198.6	211.6
Our method	2.0	10.1	13.5	10.3	14.1	16.0

Table 5.1: Number of iterations required by ByzantinePGD [23] and our method when 10%, 15%, and 20% of the worker machines are Byzantine in nature.

	Negative Update			Random Label		
	10%	15%	20%	10%	15%	20%
[31]	200.7	197.6	210.5	199.9	198	209.7
Our method	8.7	10.1	13.5	8.8	10.0	10.7

Table 5.2: Number of iterations required by ByzantinePGD [23] and our method when 10%, 15%, and 20% of the worker machines are Byzantine in nature.

Part II

Efficient and Tractable Algorithms for Non-Convex Batch Learning

In this part of the thesis, we study some statistical and algorithmic aspects of non-convex optimization in the standard supervised batch setting. In particular we study max-affine regression, which is a generalization of several classical non-convex problems including the real phase retrieval.

- **Chapter 6:** We study the max affine regression with Gaussian covariates
- **Chapter 7:** We study the max-affine regression problem for small-ball covariate design.

Chapter 6

Max Affine Regression with Gaussian Design

Max-affine regression refers to a model where the unknown regression function is modeled as a maximum of k unknown affine functions for a fixed $k \geq 1$. This generalizes linear regression and (real) phase retrieval, and is closely related to convex regression. Working within a non-asymptotic framework, we study this problem in the high-dimensional setting assuming that k is a fixed constant, and focus on the estimation of the unknown coefficients of the affine functions underlying the model. We analyze a natural alternating minimization (AM) algorithm for the non-convex least squares objective when the design is Gaussian. We show that the AM algorithm, when initialized suitably, converges with high probability and at a geometric rate to a small ball around the optimal coefficients. In order to initialize the algorithm, we propose and analyze a combination of a spectral method and a random search algorithm in a low-dimensional space, which may be of independent interest. The final rate that we obtain is near-parametric and minimax optimal (up to a polylogarithmic factor) as a function of the dimension, sample size, and noise variance. In that sense, our approach should be viewed as a *direct* and implementable method of enforcing regularization to alleviate the curse of dimensionality in problems of the convex regression type. Numerical experiments illustrate the sharpness of our bounds in the various problem parameters.

6.1 Introduction

Max-affine regression refers to the regression model

$$Y = \max_{1 \leq j \leq k} (\langle X, \theta_j^* \rangle + b_j^*) + \epsilon \quad (6.1)$$

where Y is a univariate response, X is a d -dimensional vector of covariates and ϵ models zero-mean noise that is independent of X . We assume that $k \geq 1$ is a known integer and study the problem of estimating the unknown parameters $\theta_1^*, \dots, \theta_k^* \in \mathbb{R}^d$ and $b_1^*, \dots, b_k^* \in \mathbb{R}$ from independent observations $(x_1, y_1), \dots, (x_n, y_n)$ drawn according to the model (6.1). Furthermore,

we assume for concreteness¹ in this chapter that the covariate distribution is standard Gaussian, with $x_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, I_d)$.

Let us provide some motivation for studying the model (6.1). When $k = 1$, equation (6.1) corresponds to the classical linear regression model. When $k = 2$, the intercepts $b_2^* = b_1^* = 0$, and $\theta_2^* = -\theta_1^* = \theta^*$, the model (6.1) reduces to

$$Y = |\langle X, \theta^* \rangle| + \epsilon. \quad (6.2)$$

The problem of recovering θ^* from observations drawn according to the above model is known as (real) phase retrieval—variants of which arise in a diverse array of science and engineering applications [48–51]—and has associated with it an extensive statistical and algorithmic literature.

To motivate the model (6.1) for general k , note that the function $x \mapsto \max_{1 \leq j \leq k} (\langle x, \theta_j^* \rangle + b_j^*)$ is always a convex function and, thus, estimation under the model (6.1) can be used to fit convex functions to the observed data. Indeed, the model (6.1) serves as a parametric approximation to the non-parametric convex regression model

$$Y = \phi^*(X) + \epsilon, \quad (6.3)$$

where $\phi^* : \mathbb{R}^d \rightarrow \mathbb{R}$ is an unknown convex function. It is well-known that convex regression suffers from the curse of dimensionality unless d is small, which is basically a consequence of the fact that the metric entropy of natural totally bounded sub-classes of convex functions grows exponentially in d (see, e.g., [52–54]). To overcome this curse of dimensionality, one would need to work with more structured sub-classes of convex functions. Since convex functions can be approximated to arbitrary accuracy by maxima of affine functions, it is reasonable to regularize the problem by considering only those convex functions that can be written as a maximum of a fixed number of affine functions. Constraining the number of affine pieces in the function therefore presents a simple method to enforce structure, and such function classes have been introduced and studied in the convex regression literature (see e.g., [152]). This assumption directly leads to our model (6.1), and it has been argued by [55, 153, 154] that the parametric model (6.1) is a tractable alternative to the full non-parametric convex regression model (6.3) in common applications of convex regression to data arising in economics, finance and operations research where d is often moderate to large.

Another motivation for the model (6.1) comes from the problem of estimating convex sets from support function measurements. The support function of a compact convex set $K \subseteq \mathbb{R}^d$ is defined by $h_K(x) := \sup_{u \in K} \langle x, u \rangle$ for d -dimensional unit vectors x . The problem of estimating an unknown compact, convex set K^* from noisy measurements of $h_{K^*}(\cdot)$ arises in certain engineering applications such as robotic tactile sensing and projection magnetic resonance imaging (see, e.g., [155–157]). Specifically, the model considered here is

$$Y = h_{K^*}(X) + \epsilon,$$

and the goal is to estimate the set $K^* \subseteq \mathbb{R}^d$. As in convex regression, this problem suffers from a curse of dimensionality unless d is small, as is evident from known minimax lower bounds [158].

¹In Chapter 7 weakens distributional assumptions on the covariates, but this requires significantly more technical effort.

To alleviate this curse, it is natural to restrict K^* to the class of all polytopes with at most k extreme points for a fixed k ; such a restriction has been studied as a special case of enforcing structure in these problems by Soh and Chandrasekharan [159]. Under this restriction, one is led to the model (6.1) with $b_1^* = \dots = b_k^* = 0$, since if K^* is the polytope given by the convex hull of $\theta_1^*, \dots, \theta_k^* \in \mathbb{R}^d$, then its support function is equal to $x \mapsto \max_{1 \leq j \leq k} \langle x, \theta_j^* \rangle$.

Equipped with these motivating examples, our goal is to study a computationally efficient estimation methodology for the unknown parameters of the model (6.1) from i.i.d observations $(x_i, y_i)_{i=1}^n$. Before presenting our contributions, let us first rewrite the observation model (6.1) by using more convenient notation, and use it to describe existing estimation procedures for this model. Denote the unknown parameters by $\beta_j^* := (\theta_j^*, b_j^*) \in \mathbb{R}^{d+1}$ for $j = 1, \dots, k$ and the observations by (ξ_i, y_i) for $i = 1, \dots, n$, where $\xi_i := (x_i, 1) \in \mathbb{R}^{d+1}$. In this notation, the observation model takes the form

$$y_i = \max_{1 \leq j \leq k} \langle \xi_i, \beta_j^* \rangle + \epsilon_i, \quad \text{for } i = 1, 2, \dots, n. \quad (6.4)$$

We assume that in addition to the covariates being i.i.d. standard Gaussian, the noise variables $\epsilon_1, \dots, \epsilon_n$ are independent random variables drawn from a (univariate) distribution that is zero-mean and sub-Gaussian, with unknown sub-Gaussian parameter σ .

Let us now describe existing estimation procedures for max-affine regression. The most obvious approach is the global least squares estimator, defined as any minimizer of the least squares criterion

$$(\widehat{\beta}_1^{(\text{ls})}, \dots, \widehat{\beta}_k^{(\text{ls})}) \in \underset{\beta_1, \dots, \beta_k \in \mathbb{R}^{d+1}}{\operatorname{argmin}} \sum_{i=1}^n \left(y_i - \max_{1 \leq j \leq k} \langle \xi_i, \beta_j \rangle \right)^2. \quad (6.5)$$

It is easy to see (see Lemma 27 to follow) that a global minimizer of the least squares criterion above always exists but it will not—at least in general—be unique, since any relabeling of the indices of a minimizer will also be a minimizer. While the least squares estimator has appealing statistical properties (see, e.g. [158–160]), the optimization problem (6.5) is non-convex and, in general, NP-hard [103, 161] without further assumptions on the covariate matrix and response vector.²

It is interesting to compare (6.5) to the optimization problem used to compute the least squares estimator in the more general convex regression model (6.3), given by

$$\widehat{\phi}^{(\text{ls})} \in \underset{\phi}{\operatorname{argmin}} \sum_{i=1}^n (y_i - \phi(x_i))^2, \quad (6.6)$$

where the minimization is over all convex functions ϕ . In sharp contrast to the problem (6.5), the optimization problem (6.6) is convex [162, 163] and can be solved efficiently for fairly large values of the pair (d, n) [164]. Unfortunately however, the utility of $\widehat{\phi}^{(\text{ls})}$ in estimating the parameters of the max-affine model is debatable, as it is unclear how one may obtain estimates of the true parameters $\beta_1^*, \dots, \beta_k^*$ from $\widehat{\phi}^{(\text{ls})}$, which typically will *not* be a maximum of only k affine functions.

²We provide a proof of this in Appendix 6.9.14 for completeness.

Three heuristic techniques for solving the non-convex optimization problem (6.5) were empirically evaluated by Balázs [154, Chapters 6 and 7], who compared running times and performance of these techniques on a wide variety of real and synthetic datasets for convex regression. The first technique is the alternating minimization algorithm of Magnani and Boyd [55], the second technique is the convex adaptive partitioning (or CAP) algorithm of Hannah and Dunson [153], and the third is the adaptive max-affine partitioning algorithm proposed by Balázs himself [154]. The simplest and most intuitive of these three methods is the first alternating minimization (AM) algorithm, which is an iterative algorithm for estimating the parameters $\beta_1^*, \dots, \beta_k^*$ and forms the focus of our study. In the t -th iteration of the algorithm, the current estimates $\beta_1^{(t)}, \dots, \beta_k^{(t)}$ are used to partition the observation indices $1, \dots, n$ into k sets $S_1^{(t)}, \dots, S_k^{(t)}$ such that $j \in \operatorname{argmax}_{u \in [k]} \langle \xi_i, \beta_u^{(t)} \rangle$ for every $i \in S_j^{(t)}$. For each $1 \leq j \leq k$, the next estimate $\beta_j^{(t+1)}$ is then obtained by performing a least squares fit (or equivalently, linear regression) to the data $(\xi_i, y_i), i \in S_j^{(t)}$. More intuition and a formal description of the algorithm are provided in Section 6.2. Balázs found that when this algorithm was run on a variety of datasets with multiple random initializations, it compared favorably with the state of the art in terms of its final predictive performance—see, for example, Figures 7.4 and 7.5 in the thesis [154], which show encouraging results when the algorithm is used to fit convex functions to datasets of average wages and aircraft profile drag data, respectively. In the context of fitting convex sets to support function measurements, Soh and Chandrasekaran [159] recently proposed and empirically evaluated a similar algorithm in the case of isotropic covariates. However, to the best of our knowledge, no theoretical results exist to support the performance of such a technique.

In this chapter, we present a theoretical analysis of the AM algorithm for recovering the parameters of the max-affine regression model when the covariate distribution is Gaussian. This assumption forms a natural starting point for the study of many iterative algorithms in related problems [59–61, 99], and is also quite standard in theoretical investigations of multidimensional regression problems. Note that the AM algorithm described above can be seen as a generalization of classical AM algorithms for (real) phase retrieval [95, 165], which have recently been theoretically analyzed in a series of papers [59–61] for Gaussian designs. The AM—and the closely related expectation maximization³, or EM—methodology is widely used for parameter estimation in missing data problems [166, 167] and mixture models [91], including those with covariates such as mixtures-of-experts [168] and mixtures-of-regressions [169] models. Theoretical guarantees for such algorithms have been established in multiple statistical contexts [99, 170–172]; in the case when the likelihood is not unimodal, these are typically of the local convergence type. In particular, algorithms of the EM type return, for many such latent variable models, minimax-optimal parameter estimates when initialized in a neighborhood of the optimal solution (e.g., [169, 173, 174]); conversely, these algorithms can get stuck at spurious fixed points when initialized at random [175]. In some specific applications of EM to mixtures of two Gaussians [176, 177] and mixtures of two regressions [178], however, it has been shown that randomly initializing the EM algorithm suffices in order to obtain consistent parameter estimates. Here, we establish guarantees on the AM algorithm for max-affine regression that are of the former type: we prove local geometric convergence of the

³Indeed, for many problems, the EM algorithm reduces to AM in the noiseless limit, and AM should thus be viewed as a variant of EM that uses hard-thresholding to determine values of the latent variables.

AM iterates when initialized in a neighborhood of the optimal solution. We analyze the practical variant of the algorithm in which the steps are performed without sample-splitting. As in the case of mixture models [169, 179], we use spectral methods to obtain such an initialization.

Contributions Let us now describe our results in more detail. To simplify the exposition, we state simplified corollaries of our theorems; for precise statements, see Section 6.3. We prove in Theorem 12 that for each $\epsilon > 0$, the parameter estimates $\beta_1^{(t)}, \dots, \beta_k^{(t)}$ returned by the AM algorithm at iteration t satisfy, with high probability, the inequality

$$\sum_{j=1}^k \|\beta_j^{(t)} - \beta_j^*\|^2 \leq \epsilon + C(\beta_1^*, \dots, \beta_k^*) \frac{\sigma^2 kd}{n} \log(kd) \log\left(\frac{n}{kd}\right) \quad (6.7)$$

for every $t \geq \log_{4/3} \left(\frac{\sum_{j=1}^k \|\beta_j^{(0)} - \beta_j^*\|^2}{\epsilon} \right)$, provided that the sample size n is sufficiently large and that the initial estimates satisfy the condition

$$\min_{c>0} \max_{1 \leq j \leq k} \|c\beta_j^{(0)} - \beta_j^*\|^2 \leq \frac{1}{k} c(\beta_1^*, \dots, \beta_k^*). \quad (6.8)$$

Here $C(\beta_1^*, \dots, \beta_k^*)$ and $c(\beta_1^*, \dots, \beta_k^*)$ are constants depending only on the true parameters $\beta_1^*, \dots, \beta_k^*$, and their explicit values are given in Theorem 12. The constant c in equation (6.8) endows the initialization with a scale-invariance property: indeed, scaling all parameters $\beta_1^{(0)}, \dots, \beta_k^{(0)}$ by the same positive constant c produces the same initial partition of subsets $S_1^{(0)}, \dots, S_k^{(0)}$, from which the algorithm proceeds identically.

Treating k as a fixed constant, inequality (6.7) implies, under the initialization condition (6.8), that the parameter estimates returned by AM converge geometrically to within a small ball of the true parameters, and that this error term is nearly the parametric risk $\frac{\sigma^2 d}{n}$ up to a logarithmic factor. The initialization condition (6.8) requires the distance between the initial estimates and the true parameters to be at most a specific (k -dependent) constant. It has been empirically observed that there exist bad initializations under which the AM algorithm behaves poorly (see, e.g., [55, 154]) and assumption (6.8) is one way to rule these out.

A natural question based on our Theorem 12 is whether it is possible to produce preliminary estimates $\beta_1^{(0)}, \dots, \beta_k^{(0)}$ satisfying the initialization condition (6.8). Indeed, one such method is to repeatedly initialize parameters (uniformly) at random within the unit ball \mathbb{B}^{d+1} ; Balázs empirically observed in a close relative of such a scheme (see Figure 6.6 in his thesis [154]) that increasing the number of random initializations is often sufficient to get the AM algorithm to succeed. However, reasoning heuristically, the number of repetitions required to ensure that one such random initialization generates parameters that satisfy condition (6.8) increases exponentially in the ambient dimension d , and so it is reasonable to ask if, in large dimensions, there is some natural form of dimensionality reduction that allows us to perform this step in a lower-dimensional space.

When⁴ $k < d$, we show that a natural spectral method (described formally in Algorithm 8) is able to reduce the dimensionality of our problem from d to k . In particular, this method returns an orthonormal basis of vectors $\widehat{U}_1, \dots, \widehat{U}_k$ such that the k -dimensional linear subspace spanned by these vectors accurately estimates the subspace spanned by the vectors $\theta_1^*, \dots, \theta_k^*$. We form the matrix $\widehat{U} := [\widehat{U}_1 : \dots : \widehat{U}_k]$ by collecting these vectors as its columns, and in order to account for the intercepts, further append such a matrix to form the matrix $\widehat{V} := \begin{bmatrix} \widehat{U} & 0 \\ 0 & 1 \end{bmatrix} \in \mathbb{R}^{(d+1) \times (k+1)}$.

We then choose M random initializations in $(k+1)$ dimensions—the ℓ -th such initialization is given by a set of vectors $\nu_1^\ell, \dots, \nu_k^\ell \in \mathbb{R}^{k+1}$ each chosen uniformly at random from the $(k+1)$ -dimensional unit ball—so that the collection of k vectors $\{\widehat{V}\nu_j^\ell\}_{j=1}^k$ serves as our ℓ -th guess of the true parameters. In order to decide which of these random points to choose, we evaluate (on an independent set of samples) the goodness-of-fit statistic $\min_{c \geq 0} \sum_i (y_i - c \max_{1 \leq j \leq k} \langle \xi_i, \widehat{V}\nu_j^\ell \rangle)^2$ for each $1 \leq \ell \leq M$, where the minimization over the constant c accounts for the scale-invariance property alluded to above. Letting ℓ^* denote the index with the smallest loss, we then return the initialization $\beta_j^{(0)} = \widehat{V}\nu_j^{\ell^*}$ for $j = 1, \dots, k$.

Our algorithm can thus be viewed as a variant of the repeated random initialization evaluated by Balázs [154], but incurs significantly smaller computational cost, since we only run the full-blown iterative AM algorithm once. Note that our algorithm treats the number of initializations M as a tuning parameter to be chosen by the statistician, similar to Balázs [154], but we show a concrete upper bound on M that is sufficient to guarantee convergence. In particular, we show that in order to produce an initialization satisfying condition (6.8) with high probability, it suffices to choose M as a function *only* of the number of affine pieces k and other geometric parameters of the problem (and independently of the sample size n and ambient dimension d).

To produce our overall guarantee, we combine the initialization with the AM algorithm in Corollary 2, showing that provided the sample size scales linearly in the dimension (with a multiplicative pre-factor that depends polynomially on k and other problem-dependent parameters), we obtain estimates that are accurate up to the parametric risk. Our algorithm is also computationally efficient when k is treated as a fixed constant.

From a technical standpoint, our results for the AM algorithm are significantly more challenging to establish than related results in the literature [59, 99, 180, 181]. First, it is technically very challenging to compute the *population operator* [99]—corresponding to running the AM update in the infinite sample limit—in this setting, since the max function introduces intricate geometry in the problem that is difficult to reason about in closed form. Second, we are interested in analyzing the AM update without sample-splitting, and so cannot assume that the iterates are independent of the covariates; the latter assumption has been used fruitfully in the literature to simplify analyses of such algorithms [60, 61, 180]. Third, and unlike algorithms for phase retrieval [59, 181], our algorithm performs least squares using sub-matrices of the covariate matrix that are chosen *depending* on our random iterates. Accordingly, a key technical difficulty of the proof, which may be of independent interest, is to control the spectrum of these random matrices, rows of which are

⁴If $k \geq d$, then this dimensionality reduction step can be done away with and one can implement the random search routine directly.

drawn from (randomly) truncated variants of the Gaussian distribution.

Our spectral initialization algorithm is also a natural estimator based on the method-of-moments, and has been used in a variety of non-convex problems [169, 173, 174]. However, our guarantees for this step are once again non-trivial to establish. In particular, the eigengap of the population moment (on which the rates of the estimator depend) is difficult to compute in our case since the \max function is not differentiable, and so it is not clear that higher order moments return reasonable estimates even in the infinite sample limit (see Section 6.2). However, since we operate exclusively with Gaussian covariates, we are able to use some classical moment calculations for truncated Gaussian distributions [182] in order to bound the eigengap. Translating these calculations into an eigengap is quite technical, and involves the isolation of many properties of the population moments that may be of independent interest.

Finally, it is important to note that owing to the scale invariance of our initialization condition (6.8) and goodness-of-fit statistic, our random search scheme does not require a bound on the size of the parameters; it suffices to initialize parameters uniformly within the *unit* ball. This is in contrast to other search procedures employed for similar problems [103, 183], which are based on covering arguments and require a bound on the maximum norm of the unknown parameters.

Notation For a positive integer n , let $[n] := \{1, 2, \dots, n\}$. For a finite set S , we use $|S|$ to denote its cardinality. All logarithms are to the natural base unless otherwise mentioned. For two sequences $\{a_n\}_{n=1}^\infty$ and $\{b_n\}_{n=1}^\infty$, we write $a_n \lesssim b_n$ if there is a universal constant C such that $a_n \leq Cb_n$ for all $n \geq 1$. The relation $a_n \gtrsim b_n$ is defined analogously, and we use $a_n \sim b_n$ to indicate that both $a_n \gtrsim b_n$ and $a_n \lesssim b_n$ hold simultaneously. We use c, C, c_1, c_2, \dots to denote universal constants that may change from line to line. For a pair of vectors (u, v) , we let $u \otimes v := uv^\top$ denote their outer product. We use $\|\cdot\|$ to denote the ℓ_2 norm unless otherwise stated. Denote by I_d the $d \times d$ identity matrix. We let $\mathbf{1}\{\mathcal{E}\}$ denote the indicator of an event \mathcal{E} . Let $\text{sgn}(t)$ denote the sign of a scalar t , with the convention that $\text{sgn}(0) = 1$. Let $\lambda_i(\Gamma)$ denote the i -th largest eigenvalue of a symmetric matrix Γ . Let $\mathbb{S}^{d-1} := \{v \in \mathbb{R}^d : \|v\| = 1\}$ denote the unit sphere in d -dimensions, and use $\mathbb{B}^d := \{v \in \mathbb{R}^d : \|v\| \leq 1\}$ to denote the d -dimensional unit ball.

6.2 Background and problem formulation

In this section, we formally introduce the geometric parameters underlying the max-affine regression model, as well as the methodology we use to perform parameter estimation.

6.2.1 Model and Geometric Parameters

We work throughout with the observation model defined in equation (6.4); recall that our covariates are drawn i.i.d. from a standard Gaussian distribution, and that our noise is σ -sub-Gaussian. We let $X \in \mathbb{R}^{n \times d}$ denote the covariate matrix with row i given by the vector x_i , and collect the responses in a vector $y \in \mathbb{R}^n$.

Recall that $\xi_i = (x_i, 1) \in \mathbb{R}^{d+1}$ for each $i \in [n]$; the matrix of appended covariates $\Xi \in \mathbb{R}^{n \times (d+1)}$ is defined by appending a vector of ones to the right of the matrix X . Our primary goal is to use the data (X, y) —or equivalently, the pair (Ξ, y) —to estimate the underlying parameters $\{\beta_j^*\}_{j=1}^k$.

An important consideration in achieving such a goal is the “effective” sample size with which we observe the parameter β_j^* . Toward that end, for $X \sim \mathcal{N}(0, I_d)$, let

$$\pi_j(\beta_1^*, \dots, \beta_k^*) := \Pr \left\{ \langle X, \theta_j^* \rangle + b_j^* = \max_{j' \in [k]} (\langle X, \theta_{j'}^* \rangle + b_{j'}^*) \right\} \quad (6.9)$$

denote the probability with which the j -th parameter $\beta_j^* = (\theta_j^*, b_j^*)$ attains the maximum. Note that the event on which more than one of the parameters attains the maximum has measure zero, except in the case where $\beta_i^* = \beta_j^*$ for some $i \neq j$. We explicitly disallow this case and assume that the parameters $\beta_1^*, \dots, \beta_k^*$ are distinct. Let

$$\pi_{\min}(\beta_1^*, \dots, \beta_k^*) := \min_{j \in [k]} \pi_j(\beta_1^*, \dots, \beta_k^*), \quad (6.10)$$

and assume that we have $\pi_{\min}(\beta_1^*, \dots, \beta_k^*) > 0$; in other words, we ignore vacuous cases in which some parameter is never observed. Roughly speaking, the sample size of the parameter that is observed most rarely is given by $\min_{j \in [k]} \pi_j n \sim n \cdot \pi_{\min}(\beta_1^*, \dots, \beta_k^*)$, and so the error in estimating this parameter should naturally depend on $\pi_{\min}(\beta_1^*, \dots, \beta_k^*)$. By definition, we always have $\pi_{\min}(\beta_1^*, \dots, \beta_k^*) \leq 1/k$.

Since we are interested in performing parameter estimation under the max-affine regression model, a few geometric quantities also appear in our bounds, and serve as natural notions of “signal strength” and “condition number” of the estimation problem. The signal strength is given by the minimum separation

$$\Delta(\beta_1^*, \dots, \beta_k^*) = \min_{j, j': j \neq j'} \|\theta_j^* - \theta_{j'}^*\|^2;$$

we also assume that Δ is strictly positive, since otherwise, a particular parameter is never observed. To denote a natural form of conditioning, define the quantities

$$\kappa_j(\beta_1^*, \dots, \beta_k^*) = \frac{\max_{j' \neq j} \|\theta_j^* - \theta_{j'}^*\|^2}{\min_{j' \neq j} \|\theta_j^* - \theta_{j'}^*\|^2}, \quad \text{with} \quad \kappa(\beta_1^*, \dots, \beta_k^*) = \max_{j \in [k]} \kappa_j(\beta_1^*, \dots, \beta_k^*).$$

Finally, let $B_{\max}(\beta_1^*, \dots, \beta_k^*) := \max_{j \in [k]} \|\beta_j^*\|$ denote the maximum norm of any unknown parameter. We often use the shorthand

$$\begin{aligned} \pi_{\min} &= \pi_{\min}(\beta_1^*, \dots, \beta_k^*), & \Delta &= \Delta(\beta_1^*, \dots, \beta_k^*), \\ \kappa &= \kappa(\beta_1^*, \dots, \beta_k^*), & \text{and} & \quad B_{\max} = B_{\max}(\beta_1^*, \dots, \beta_k^*) \end{aligned}$$

when the true parameters $\beta_1^*, \dots, \beta_k^*$ are clear from context.

6.2.2 Methodology

As discussed in the introduction, the most natural estimation procedure from i.i.d. samples $(\xi_i, y_i)_{i=1}^n$ of the model (6.4) is the least squares estimator (6.5). The following lemma (which does not seem to have been explicitly stated previously in the literature, except in the case $k = 2$ [160, 184]) proves that the least squares estimator $(\widehat{\beta}_1^{(\text{ls})}, \dots, \widehat{\beta}_k^{(\text{ls})})$ always exists. Note, however, that it will not be unique in general since any relabeling of a minimizer is also a minimizer.

Lemma 27. *The least squares estimator $(\widehat{\beta}_1^{(\text{ls})}, \dots, \widehat{\beta}_k^{(\text{ls})})$ exists for every dataset (Ξ, y) .*

We postpone the proof of Lemma 27 to Appendix 6.5.1. In spite of the fact that the least squares estimator always exists, the problem (6.5) is non-convex and NP-hard in general. The AM algorithm presents a tractable approach towards solving it in the statistical setting that we consider.

Alternating Minimization

We now formally describe the AM algorithm proposed by Magnani and Boyd [55]. For each β_1, \dots, β_k , define the sets

$$S_j(\beta_1, \dots, \beta_k) := \left\{ i \in [n] : j = \min_{1 \leq u \leq k} \operatorname{argmax} (\langle \xi_i, \beta_u \rangle) \right\} \quad (6.11)$$

for $j = 1, \dots, k$. In words, the set $S_j(\beta_1, \dots, \beta_k)$ contains the indices of samples on which parameter β_j attains the maximum; in the case of a tie, samples having multiple parameters attaining the maximum are assigned to the set with the smallest corresponding index (i.e., ties are broken in the lexicographic order⁵). Thus, the sets $\{S_j(\beta_1, \dots, \beta_k)\}_{j=1}^k$ define a partition of $[n]$. The AM algorithm employs an iterative scheme where one first constructs the partition $S_j(\beta_1^{(t)}, \dots, \beta_k^{(t)})$ based on the current iterates $\beta_1^{(t)}, \dots, \beta_k^{(t)}$ and then calculates the next parameter estimate $\beta_j^{(t+1)}$ by a least squares fit to the dataset $\{(\xi_i, y_i), i \in S_j(\beta_1^{(t)}, \dots, \beta_k^{(t)})\}$. The algorithm (also described below as Algorithm 7) is, clearly, quite intuitive and presents a natural approach to solving (6.5).

As a sanity check, Lemma 28 (stated and proved in Appendix 6.5.1) shows that the global least squares estimator (6.5) is a fixed-point of this iterative scheme under a mild technical assumption.

We also note that the AM algorithm was proposed by Soh [185] in the context of estimating structured convex sets from support function measurements. It should be viewed as a generalization of a classical algorithm for (real) phase retrieval due to Fienup [95], which has been more recently analyzed in a series of papers [59, 61] for Gaussian designs. While some analyses of AM algorithms assume sample-splitting across iterations (e.g. [60, 61, 180]), we consider the more practical variant of AM run without sample-splitting, since the update (6.12a)-(6.12b) is run on the full data (Ξ, y) in every iteration.

⁵In principle, it is sufficient to define the sets $S_j(\beta_1, \dots, \beta_k), j \in [k]$ as any partition of $[n]$ having the property that $\langle \xi_i, \beta_j \rangle = \max_{u \in [k]} \langle \xi_i, \beta_u \rangle$ for every $j \in [k]$ and $i \in S_j(\beta_1, \dots, \beta_k)$; here “any” means that ties can be broken according to an arbitrary rule, and we have chosen this rule to be the lexicographic order in equation (6.11).

Algorithm 7: Alternating minimization for estimating maximum of k affine functions**Input:** Data $\{\xi_i, y_i\}_{i=1}^n$; initial parameter estimates $\beta_1^{(0)}, \dots, \beta_k^{(0)}$; number of iterations T .**Output:** Final estimator of parameters $\hat{\beta}_1, \dots, \hat{\beta}_k$.2 Initialize $t \leftarrow 0$.3 **repeat**

5 Compute maximizing index sets

$$S_j^{(t)} = S_j(\beta_1^{(t)}, \dots, \beta_k^{(t)}), \quad (6.12a)$$

for each $j \in [k]$, according to equation (6.11).

7 Update

$$\beta_j^{(t+1)} \in \operatorname{argmin}_{\beta \in \mathbb{R}^{d+1}} \sum_{i \in S_j^{(t)}} (y_i - \langle \xi_i, \beta \rangle)^2, \quad (6.12b)$$

for each $j \in [k]$.8 **until** $t = T$;10 Return $\hat{\beta}_j = \beta_j^{(T)}$ for each $j \in [k]$.**Initialization**

The alternating minimization algorithm described above requires an initialization. While the algorithm was proposed to be run from a random initialization with restarts [55, 159], we propose to initialize the algorithm from parameter estimates that are sufficiently close to the optimal parameters. This is similar to multiple procedures to solve non-convex optimization problems in statistical settings (e.g., [99, 186]), that are based on iterative algorithms that exhibit *local* convergence to the unknown parameters. Such algorithms are typically initialized by using a moment method, which (under various covariate assumptions) returns useful parameter estimates.

Our approach to the initialization problem is similar, in that we combine a moment method with random search in a lower-dimensional space. For convenience of analysis, we split the n samples into two equal parts—assume that n is even without loss of generality—and perform each of the above steps on different samples so as to maintain independence between the two steps. The formal algorithm is presented in two parts as Algorithms 8 and 9.

In related problems [98, 169, 173, 187], a combination of a second order and third order method (involving tensor decomposition) is employed to obtain parameter estimates in one shot. Take the problem of learning generalized linear models [187] as an example; here, the analysis of the moment method relies on the link function being (at least) three times differentiable so that the population moment quantities can be explicitly computed. After showing that these expectations are closed form functions of the unknown parameters, matrix/tensor perturbation tools are then applied to show that the empirical moments concentrate about their population counterparts. However, in our setting, the max function is not differentiable, and so it is not clear that higher order moments

Algorithm 8: PCA for k -dimensional subspace initialization**Input:** Data $\{\xi_i, y_i\}_{i=1}^n$.**Output:** Matrix $\widehat{U} \in \mathbb{R}^{d \times k}$ having orthonormal columns that (approximately) span the k -dimensional subspace spanned by the vectors $\theta_1^*, \dots, \theta_k^*$.

2 Compute the quantities

$$\widehat{M}_1 = \frac{2}{n} \sum_{i=1}^{n/2} y_i x_i \quad \text{and} \quad \widehat{M}_2 = \frac{2}{n} \sum_{i=1}^{n/2} y_i (x_i x_i^\top - I_d), \quad (6.13)$$

and let $\widehat{M} = \widehat{M}_1 \otimes \widehat{M}_1 + \widehat{M}_2$; here, I_d denotes the $d \times d$ identity matrix and \otimes denotes the outer product.

4 Perform the eigendecomposition $\widehat{M} = \widehat{P} \widehat{\Lambda} \widehat{P}^\top$, and use the first k columns of \widehat{P} (corresponding to the k largest eigenvalues) to form the matrix $\widehat{U} \in \mathbb{R}^{d \times k}$. Return \widehat{U} .

return reasonable estimates even in expectation since Stein's lemma (on which many of these results rely) is not applicable⁶ in this setting. Nevertheless, we show that the second order moment returns a k -dimensional subspace that is close to the true span of the parameters $\{\theta_j^*\}_{j=1}^k$; the degree of closeness depends only on the geometric properties of these parameters.

Let us also briefly discuss Algorithm 9, which corresponds to performing random search in $(k+1) \cdot k$ dimensional space to obtain the final initialization. In addition to the random initialization employed in step 1 of this algorithm, we also use the mean squared error on a holdout set (corresponding to samples $n/2+1$ through n) to select the final parameter estimates. In particular, we evaluate the error in a scale-invariant fashion; the computation of the optimal constant c in step 2 of the algorithm can be performed in closed form for each fixed index ℓ , since for a pair of vectors (u, v) having equal dimension, we have

$$\operatorname{argmin}_{c \geq 0} \|u - cv\|^2 = \max \left\{ \frac{\langle u, v \rangle}{\|v\|^2}, 0 \right\}.$$

A key parameter that governs the performance of our search procedure is the number of initializations M ; we show in the sequel that it suffices to take M to be a quantity that depends only on the number of affine pieces k , and on other geometric parameters in the problem.

Our overall algorithm should be viewed as a slight variation of the AM algorithm with random restarts. It inherits similar empirical performance (see panel (b) of Figure 6.2 to follow), while significantly reducing the computational cost, since operations are now performed in ambient dimension $k+1$, and the iterative AM algorithm is run only once overall. It also produces *provable* parameter estimates, and as we show in the sequel, the number of random initializations M can be set independently of the pair (n, d) . Having stated the necessary background and described our methodology, we now proceed to statements and discussions of our main results.

⁶A natural workaround is to use Stein's lemma on the infinitely differentiable "softmax" surrogate function, but our approach to this involved balancing the estimation error (which, in turn, involves derivatives of the softmax function) and approximation error terms, and led to suboptimal dependence on the dimension.

Algorithm 9: Low-dimensional random search

Input: Data $\{\xi_i, y_i\}_{i=1}^n$, subspace estimate $\widehat{U} \in \mathbb{R}^{d \times k}$ having orthonormal columns that (approximately) span the k dimensional subspace spanned by the vectors $\theta_1^*, \dots, \theta_k^*$, and number of random initializations $M \in \mathbb{N}$.

Output: Initial estimator of parameters $\beta_1^{(0)}, \dots, \beta_k^{(0)}$.

- 2 Choose $M \cdot k$ random points ν_j^ℓ i.i.d. for $\ell \in [M]$ and $j \in [k]$, each uniformly from the $(k+1)$ -dimensional unit ball \mathbb{B}^{k+1} . Let

$$\widehat{V} = \begin{bmatrix} \widehat{U} & 0 \\ 0 & 1 \end{bmatrix}$$

be a matrix in $\mathbb{R}^{(d+1) \times (k+1)}$ having orthonormal columns.

- 4 Compute the index

$$\ell^* \in \operatorname{argmin}_{\ell \in [M]} \frac{2}{n} \left\{ \min_{c \geq 0} \sum_{i=n/2+1}^n \left(y_i - c \max_{j \in [k]} \langle \xi_i, \widehat{V} \nu_j^\ell \rangle \right)^2 \right\}.$$

- 6 Return the $(d+1)$ -dimensional parameters

$$\beta_j^{(0)} = \widehat{V} \nu_j^{\ell^*} \quad \text{for each } j \in [k].$$

6.3 Main results

In this section, we present our main theoretical results for the methodology introduced in Section 6.2.

6.3.1 Local geometric convergence of alternating minimization

We now establish local convergence results for the AM algorithm. Recall the definition of the parameters $(\pi_{\min}, \Delta, \kappa)$ introduced in Section 6.2, and the assumption that the covariates $\{x_i\}_{i=1}^n$ are drawn i.i.d. from the standard Gaussian distribution $\mathcal{N}(0, I_d)$. Throughout the paper, we assume that the true parameters $\beta_1^*, \dots, \beta_k^*$ are fixed.

Theorem 12. *There exists a tuple of universal constants (c_1, c_2) such that if the sample size satisfies the bound*

$$n \geq c_1 \max \{d, 10 \log n\} \max \left\{ \frac{k\kappa}{\pi_{\min}^3}, \sigma^2 \frac{k^5 \kappa^2}{\Delta \pi_{\min}^{15}} \right\},$$

then for all initializations $\beta_1^{(0)}, \dots, \beta_k^{(0)}$ satisfying the bound

$$\min_{c>0} \max_{1 \leq j \neq j' \leq k} \frac{\left\| c \left(\beta_j^{(0)} - \beta_{j'}^{(0)} \right) - \left(\beta_j^* - \beta_{j'}^* \right) \right\|}{\left\| \theta_j^* - \theta_{j'}^* \right\|} \leq c_2 \frac{\pi_{\min}^6}{k^2 \kappa} \log^{-3/2} \left(\frac{k^2 \kappa}{\pi_{\min}^6} \right), \quad (6.14a)$$

the estimation error at all iterations $t \geq 1$ is simultaneously bounded as

$$\sum_{j=1}^k \left\| \beta_j^{(t)} - \beta_j^* \right\|^2 \leq \left(\frac{3}{4} \right)^t \left(\sum_{j=1}^k \left\| c^* \beta_j^{(0)} - \beta_j^* \right\|^2 \right) + c_1 \sigma^2 \frac{kd}{\pi_{\min}^3 n} \log(kd) \log(n/kd) \quad (6.14b)$$

with probability exceeding $1 - c_2 \left(k \exp \left(-c_1 n \frac{\pi_{\min}^6}{k^2} \right) + \frac{k^2}{n^7} \right)$. Here, the positive scalar c^* minimizes the LHS of inequality (6.14a).

See Appendix 6.6 for a concise mathematical statement of the probability bound.

Let us interpret the various facets of Theorem 12. As mentioned before, it is a local convergence result, which requires the initialization $\beta_1^{(0)}, \dots, \beta_k^{(0)}$ to satisfy condition (6.14a). In the well-balanced case (with $\pi_{\min} \sim 1/k$) and treating k as a fixed constant, the initialization condition (6.14a) posits that the parameters are a constant “distance” from the true parameters. Notably, closeness is measured in a relative sense, and between pairwise differences of the parameter estimates as opposed to the parameters themselves; the intuition for this is that the initialization $\beta_1^{(0)}, \dots, \beta_k^{(0)}$ induces the initial partition of samples $S_1(\beta_1^{(0)}, \dots, \beta_k^{(0)}), \dots, S_1(\beta_1^{(0)}, \dots, \beta_k^{(0)})$, whose closeness to the true partition depends only on the relative pairwise differences between parameters, and is also invariant to a global scaling of the parameters. It is also worth noting that local geometric convergence of the AM algorithm is guaranteed uniformly from *all* initializations satisfying condition (6.14a). In particular, the initialization parameters are not additionally required to be independent of the covariates or noise, and this allows us to use the same n samples for initialization of the parameters.

Let us now turn our attention to the bound (6.14b), which consists of two terms. In the limit $t \rightarrow \infty$, the final parameters provide an estimate of the true parameters that is accurate to within the second term of the bound (6.14b). Up to a constant, this is the statistical error term

$$\delta_{n,\sigma}(d, k, \pi_{\min}) = \sigma^2 \frac{kd}{\pi_{\min}^3 n} \log(kd) \log(n/kd) \quad (6.15)$$

that converges to 0 as $n \rightarrow \infty$, thereby providing a consistent estimate in the large sample limit. Notice that the dependence of $\delta_{n,\sigma}(d, k, \pi_{\min})$ on the tuple (σ, d, n) is minimax-optimal up to the logarithmic factor $\log(n/d)$, since a matching lower bound can be proved for the linear regression problem when $k = 1$. In Proposition 16, (see Appendix 6.9.1) we also show a parametric lower bound on the minimax estimation error for general k , of the order $\sigma^2 kd/n$. Panel (c) of Figure 6.1 verifies in a simulation that the statistical error depends linearly on d/n . The dependence of the statistical error on the pair (k, π_{\min}) is more involved, and we do not yet know if these are optimal. As discussed before, a linear dependence of π_{\min} is immediate from a sample-size argument; the




Figure 6.1: Convergence of the AM with Gaussian covariates— in panel (a), we plot the noiseless sample complexity of AM; we fix $\|\beta_i^*\| = 1$ for all $i \in [k]$, $\sigma = 0$ and $\pi_{\min} = 1/k$. We say β_i^* is recovered if $\|\beta_i^{(t)} - \beta_i^*\| \leq 0.01$. For a fixed dimension d , we run a linear search on the number of samples n , such that the empirical probability of success over 100 trials is more than 0.95, and output the least such n . In panel (b), we plot the optimization error (in blue) $\sum_{j=1}^k \|\beta_j^{(t)} - \beta_j^{(T)}\|^2$ and the deviation from the true parameters (in red) $\sum_{j=1}^k \|\beta_j^{(t)} - \beta_j^*\|^2 / \sigma^2$ over iterations t for different σ (0.15, 0.25, 0.4, 0.5), with $k = 5$, $d = 100$, $T = 50$ and $n = 5d$, and averaged over 50 trials. Panel (c) shows that the estimation error at $T = 50$ scales at the parametric rate d/n , where we have chosen a fixed $k = 5$ and $\sigma = 0.25$. Panel (d) shows the variation of this error as a function of π_{\min} where we fix $k = 3$, $d = 2$, $n = 10^3$, $\sigma = 0.4$.

cubic dependence arises because the sub-matrices of Ξ chosen over the course of the algorithm are not always well-conditioned, and their condition number scales (at most) as π_{\min}^2 . In Appendix 6.9.3, we show a low-dimensional example (with $d = 2$ and $k = 3$) in which the least squares estimator incurs a parameter estimation error of the order $\frac{1}{\pi_{\min}^3 n}$ even when provided with the *true* partition of covariates $\{S_j(\beta_1^*, \dots, \beta_k^*)\}_{j=1}^3$. While this does not constitute an information theoretic lower bound, it provides strong evidence to suggest that our dependence on π_{\min} is optimal at least when viewed in isolation. We verify this intuition via simulation: in panel (d) of Figure 6.1, we observe

that on this example, the error of the final AM iterate varies linearly with the quantity $1/\pi_{\min}^3$.

The first term of the bound (6.14b) is an optimization error that is best interpreted in the noiseless case $\sigma = 0$, wherein the parameters $\beta_1^{(t)}, \dots, \beta_k^{(t)}$ converge at a geometric rate to the true parameters $\beta_1^*, \dots, \beta_k^*$, as verified in panel (a) of Figure 6.1. In particular, in the noiseless case, we obtain exact recovery of the parameters provided $n \geq C \frac{kd}{\pi_{\min}^3} \log(n/d)$. Thus, the “sample complexity” of parameter recovery is linear in the dimension d , which is optimal (panel (a) verifies this fact). In the well-balanced case, the dependence on k is quartic, but lower bounds based on parameter counting suggest that the true dependence ought to be linear. Again, we are not aware of whether the dependence on π_{\min} in the noiseless case is optimal; our simulations shown in panel (a) suggests that the sample complexity depends inversely on π_{\min} , and so closing this gap is an interesting open problem. When $\sigma > 0$, we have an overall sample size requirement

$$n \geq c \max \{d, 10 \log n\} \cdot \max \left\{ \frac{k\kappa}{\pi_{\min}^3}, \sigma^2 \frac{k^5 \kappa^2}{\Delta \pi_{\min}^{15}} \right\} := n_{\text{AM}}(c). \quad (6.16)$$

As a final remark, note that Theorem 12 holds under Gaussian covariates and when the true parameters $\beta_1^*, \dots, \beta_k^*$ are fixed independently of the covariates. In Chapter 7, it is shown that both of these features of the result can be relaxed, i.e., AM converges geometrically even under a milder covariate assumption, and this convergence occurs *for all* true parameters that are geometrically similar.

6.3.2 Initialization

In this section, we provide guarantees on the initialization method described in Algorithms 8 and 9 in Theorems 13 and 14, respectively.

Consider the matrices \widehat{U} and \widehat{M} defined in Algorithm 8. Algorithm 8 is a moment method: we extract the top k principal components of a carefully chosen moment statistic of the data to obtain a subspace estimate \widehat{U} . Spectral algorithms such as these have been used to obtain initializations in a wide variety of non-convex problems [56, 98, 188] to obtain an accurate estimate of the subspace spanned by the unknown parameters. It is well-known that the performance of the algorithm in recovering a k -dimensional subspace depends on $\lambda_k(\mathbb{E}[\widehat{M}])$, which is the k -th largest eigenvalue of the population moment $\mathbb{E}[\widehat{M}]$. We show in the proof (see the discussion following Lemma 33) that there is a strictly positive scalar γ such that

$$\lambda_k(\mathbb{E}[\widehat{M}]) \geq \gamma. \quad (6.17)$$

It should be stressed that we obtain an explicit expression for γ as a function of the various problem parameters (in equation (6.48) of the proof) that is, a priori⁷, independent of the ambient dimension d .

⁷While this may seem surprising—after all, the unknown parameters $\theta_1^*, \dots, \theta_k^*$ live in dimension d —all the interesting action is confined to the k dimensional subspace spanned by these parameters and γ is a function of the geometry induced by the parameters on this subspace.

This characterization is the main novelty of our contribution, and allows us to establish the following guarantee on the PCA algorithm. We let $U^* \in \mathbb{R}^{d \times k}$ denote a matrix whose orthonormal columns span the linear subspace spanned by the vectors $\theta_1^*, \dots, \theta_k^*$, and define the quantity

$$\varsigma := \max_{j \in [k]} \{ \|\theta_j^*\|_1 + b_j^* \}. \quad (6.18)$$

Theorem 13. *There is a universal constant C such that \widehat{U} satisfies the bound*

$$\|\widehat{U}\widehat{U}^\top - U^*(U^*)^\top\|_F^2 \leq C \left(\frac{\sigma^2 + \varsigma^2}{\gamma^2} \right) \frac{kd \log^3(nk)}{n}$$

with probability greater than $1 - Cn^{-10}$.

The proof of Theorem 13 is provided in Appendix 6.7. We have thus shown that the projection matrix $U^*(U^*)^\top$ onto the true subspace spanned by the vectors $\theta_1^*, \dots, \theta_k^*$ can be estimated at the parametric rate via our PCA procedure. This guarantee is illustrated via simulation in panel (a) of Figure 6.2.

Let us now turn to establishing a guarantee on Algorithm 9 when it is given a (generic) subspace estimate \widehat{U} as input. Since the model (6.1) is only identifiable up to a relabeling of the individual parameters, we can only hope to show that a suitably permuted set of the initial parameters is close to the true parameters. Toward that end, let \mathcal{P}_k denote the set of all permutations from $[k] \rightarrow [k]$, and let

$$\text{dist} \left(\left\{ \beta_j^{(0)} \right\}_{j=1}^k, \left\{ \beta_j^* \right\}_{j=1}^k \right) := \min_{P \in \mathcal{P}_k} \sum_{j=1}^k \|\beta_{P(j)}^{(0)} - \beta_j^*\|^2 \quad (6.19)$$

denote the minimum distance attainable via a relabeling of the parameters. With this notation in place, we are now ready to state our result for parameter initialization. In it, we assume that the input matrix \widehat{U} is fixed independently of the samples used to carry out the random search.

Theorem 14. *Let $\delta \in (0, 1)$ and $0 \leq r \leq \frac{\Delta \pi_{\min}^{5/2} \log^{-1/2}(k/\pi_{\min})}{8k^3}$ denote two positive scalars. Further suppose that*

$$\|\widehat{U}\widehat{U}^\top - U^*(U^*)^\top\|_{op} \leq \frac{\Delta \pi_{\min}^{3/2}}{8B_{\max} k^2}, \quad \text{and that} \quad M \geq \left(1 + \frac{B_{\max}}{r} \right)^{k^2} \log(1/\delta).$$

Then there is a tuple of universal constants (c_1, c_2) such that if

$$n \geq c_1 \max \left\{ d \frac{k^3}{\pi_{\min}^3} \log^2(\pi_{\min}/k), \sigma^2 \frac{k^3}{\pi_{\min}^3 \Delta^2} \log M \right\},$$

then

$$\begin{aligned} & \min_{c > 0} \text{dist} \left(\left\{ c \beta_j^{(0)} \right\}_{j=1}^k, \left\{ \beta_j^* \right\}_{j=1}^k \right) \\ & \leq c_1 \left(\frac{k}{\pi_{\min}} \right)^3 \left\{ 4k \left(r^2 + B_{\max}^2 \|\widehat{U}\widehat{U}^\top - U^*(U^*)^\top\|_{op}^2 \right) + \frac{\sigma^2 \log M}{n} \right\} \end{aligned}$$




Figure 6.2: Simulation of the PCA and overall guarantees. We assume that the true parameter matrix $\Theta^* = A^*(U^*)^\top$ for a $\mathbb{R}^{d \times k}$ matrix U^* and an invertible $A^* \in \mathbb{R}^{k \times k}$, and that Algorithm 8 returns a subspace estimate \widehat{U} . Panel (a) reveals the subspace estimation error as a function of d/n , which is corroborated by Theorem 13. In panel (b), we compare the performance of our overall algorithm (in red) with that of AM with repeated random initialization [154] (in blue) averaged over 50 trials. We fix $k = 3$, $d = 50$, $n = 35kd$ and $\sigma = 0.1$. For a sufficiently large M , both schemes perform in a similar fashion.

with probability exceeding $1 - \delta - c_1 k M \exp\left(-c_2 n \frac{\pi_{\min}^4}{k^4 \log^2(k/\pi_{\min})}\right)$.

We prove Theorem 14 in Appendix 6.8. Combining Theorems 13 and 14 with some algebra then allows us to prove a guarantee for the initialization procedure that combines Algorithms 8 and 9 in sequence. In particular, fix a pair of positive scalars $\epsilon \leq \Delta$ and $\delta \in (0, 1)$. Then combining the theorems shows that if (for an appropriately large universal constant c), we have

$M \geq \left(1 + c \frac{B_{\max} k^3 \log^{1/2}(k/\pi_{\min})}{\epsilon \pi_{\min}^{5/2}}\right)^{k^2} \log(1/\delta)$, and the sample size n is greater than

$$n_{\text{init}}(\epsilon, M, c) := c \max\left\{d \frac{k}{\pi_{\min}}, \sigma^2 \frac{k^5}{\pi_{\min}^5 \epsilon^2} \log(k/\pi_{\min}) \log(M/\delta), d \log^3(nk) \log(k/\pi_{\min}) \frac{k^7 B_{\max}^2}{\gamma^2 \pi_{\min}^5 \epsilon^2} (\sigma^2 + \varsigma^2)\right\}, \quad (6.20)$$

then $\min_{c>0} \text{dist}\left(\left\{c\beta_j^{(0)}\right\}_{j=1}^k, \left\{\beta_j^*\right\}_{j=1}^k\right) \leq \epsilon^2$ with probability greater than $1 - \delta - cn^{-10}$.

Equipped with this guarantee on our initialization step, we are now in a position to state an end-to-end guarantee on our overall methodology in the next section.

6.3.3 Overall algorithmic guarantee

Assume without loss of generality that the identity permutation minimizes the distance measure dist , so that $\beta_j^{(0)}$ is the estimate of the parameter β_j^* for each $j \in [k]$. Recall the statistical error

$\delta_{n,\sigma}(d, k, \pi_{\min})$ defined in equation (6.15), which is, up to a constant factor, the final (squared) radius of the ball to which the AM update converges when initialized suitably, and the notation $n_{\text{AM}}(c)$ and $n_{\text{init}}(\epsilon, M, c)$ from equations (6.16) and (6.20), respectively. We now state a guarantee for our overall procedure that runs Algorithms 8, 9, and 7 in that sequence; we omit the proof since it follows by simply putting together the pieces from Theorem 12 and the discussion above.

Corollary 2. *There exist universal constants c_1 and c_2 such that for each $\delta \in (0, 1)$, if*

$$M \geq \left(1 + c_1 \frac{\mathbf{B}_{\max} k^4 \log^{1/2}(k/\pi_{\min})}{\pi_{\min}^{11/2}} \right)^{k^2} \log(1/\delta), n \geq \max \left\{ n_{\text{init}} \left(c_2 \frac{\pi_{\min}^6}{k^2}, c_1, M \right), n_{\text{AM}}(c_1) \right\}$$

and $T_0 = c_1 \log \left(\frac{1}{\delta_{n,\sigma}(d, k, \pi_{\min})} \right),$

then the combined algorithm satisfies, simultaneously for all $T \geq T_0$, the bound

$$\Pr \left\{ \sum_{j=1}^k \|\beta_j^{(T)} - \beta_j^*\|^2 \geq c_1 \delta_{n,\sigma}(d, k, \pi_{\min}) \right\} \leq \delta + c_1 \left(n^{-10} + k \exp \left(-c_2 n \frac{\pi_{\min}^6}{k^2} \right) + \frac{k^2}{n^7} \right).$$

We thus obtain, an algorithm that when given a number of samples that is near-linear in the ambient dimension, achieves the rate $\delta_{n,\sigma}(d, k, \pi_{\min}) = \frac{\sigma^2 kd}{\pi_{\min}^3 n} \log(kd) \log(n/kd)$ of estimation of all kd parameters in squared ℓ_2 norm. This convergence is illustrated in simulation in Figure 6.2, in which we choose $k = 3$, $d = 50$ and $n = 35kd$. Interestingly, panel (b) of this figure shows that our provable multi-step algorithm has performance similar to the algorithm that runs AM with repeated random initializations.

The computational complexity of our overall algorithm (with exact matrix inversions) is given by $\mathcal{O} \left(knd^2 \log \left(\frac{1}{\delta_{n,\sigma}(d,k,\pi_{\min})} \right) + Mnd \right)$, where we also assume that the k top eigenvectors of the matrix \widehat{M} are computed exactly in Algorithm 8. This guarantee can also be extended to the case where the linear system is solved up to some numerical precision by (say) a conjugate gradient method and the eigenvectors of \widehat{M} are computed using the power method, thereby reducing the computational complexity. Such an extension is standard and we do not detail it here.

6.3.4 Proof sketch and technical challenges

Let us first sketch, at a high level, the ideas required to establish guarantees on the AM algorithm. We need to control the iterates of the AM algorithm without sample-splitting across iterations, and so the iterates themselves are random and depend on the sequence of random variables $(\xi_i, \epsilon_i)_{i=1}^n$. A popular and recent approach to handling this issue in related iterative algorithms (e.g., [99]) goes through two steps: first, the population update, corresponding to running (6.12a)-(6.12b) in the case $n \rightarrow \infty$, is analyzed, after which the random iterates in the finite-sample case are shown to be close to their (non-random) population counterparts by using concentration bounds for the associated empirical process. The main challenge in our setting is that the population update is quite non-trivial

to write down since it involves a delicate understanding of the geometry of the covariate distribution induced by the maxima of affine functions. We thus resort to handling the random iterates directly, thereby sidestepping the calculation of the population operator entirely.

In order to convey the principal difficulties associated with our approach, let us present a bound on the error obtained after running a single step of the algorithm, starting at the parameters β_1, \dots, β_k and obtaining, as a result of one step of the algorithm, the parameters $\beta_1^+, \dots, \beta_k^+$. We use the shorthand notation $S_j := S_j(\beta_1, \dots, \beta_k)$, and let $P_{\Xi^j(\beta_1, \dots, \beta_k)}$ denote the projection matrix onto the range of the matrix Ξ_{S_j} .

Let y^* denote the vector with entry i given by $\max_{\ell \in [k]} \langle \xi_i, \beta_\ell^* \rangle$. We have

$$\begin{aligned} \|\Xi_{S_j}(\beta_j^+ - \beta_j^*)\|^2 &= \|P_{\Xi^j(\beta_1, \dots, \beta_k)} y_{S_j} - \Xi_{S_j} \beta_j^*\|^2 \\ &= \|P_{\Xi^j(\beta_1, \dots, \beta_k)} y_{S_j}^* + P_{\Xi^j(\beta_1, \dots, \beta_k)} \epsilon_{S_j} - \Xi_{S_j} \beta_j^*\|^2 \\ &\leq 2\|P_{\Xi^j(\beta_1, \dots, \beta_k)} (y_{S_j}^* - \Xi_{S_j} \beta_j^*)\|^2 + 2\|P_{\Xi^j(\beta_1, \dots, \beta_k)} \epsilon_{S_j}\|^2 \\ &\leq 2\|y_{S_j}^* - \Xi_{S_j} \beta_j^*\|^2 + 2\|P_{\Xi^j(\beta_1, \dots, \beta_k)} \epsilon_{S_j}\|^2, \end{aligned} \quad (6.21)$$

where we have used the fact that the projection operator is non-expansive on a convex set.

Let

$$\{\langle \xi_i, \beta_\ell \rangle = \max\} := \left\{ \langle \xi_i, \beta_\ell \rangle = \max_{u \in [k]} \langle \xi_i, \beta_u \rangle \right\}, \text{ for each } i \in [n], \ell \in [k]$$

denote a convenient shorthand for these events. The first term on the RHS of inequality (6.21) can be written as

$$\sum_{i \in S_j} (y_i^* - \langle \xi_i, \beta_j^* \rangle)^2 \leq \sum_{i=1}^n \sum_{j': j' \neq j} \mathbf{1} \{ \langle \xi_i, \beta_j \rangle = \max \text{ and } \langle \xi_i, \beta_{j'}^* \rangle = \max \} \langle \xi_i, \beta_{j'}^* - \beta_j^* \rangle^2,$$

where the inequality accounts for ties. Each indicator random variable is bounded, in turn, as

$$\begin{aligned} \mathbf{1} \{ \langle \xi_i, \beta_j \rangle = \max \text{ and } \langle \xi_i, \beta_{j'}^* \rangle = \max \} &\leq \mathbf{1} \{ \langle \xi_i, \beta_j \rangle \geq \langle \xi_i, \beta_{j'} \rangle \text{ and } \langle \xi_i, \beta_{j'}^* \rangle \geq \langle \xi_i, \beta_j^* \rangle \} \\ &= \mathbf{1} \{ \langle \xi_i, \beta_j - \beta_{j'} \rangle \cdot \langle \xi_i, \beta_j^* - \beta_{j'}^* \rangle \leq 0 \}. \end{aligned}$$

Switching the order of summation yields the bound

$$\sum_{i \in S_j} (y_i^* - \langle \xi_i, \beta_j^* \rangle)^2 \leq \sum_{j': j' \neq j} \sum_{i=1}^n \mathbf{1} \{ \langle \xi_i, \beta_j - \beta_{j'} \rangle \cdot \langle \xi_i, \beta_j^* - \beta_{j'}^* \rangle \leq 0 \} \langle \xi_i, \beta_j^* - \beta_{j'}^* \rangle^2.$$

Recalling our notation for the minimum eigenvalue of a symmetric matrix, the LHS of inequality (6.21) can be bounded as

$$\|\Xi_{S_j}(\beta_j^+ - \beta_j^*)\|^2 \geq \lambda_{\min} \left(\Xi_{S_j}^\top \Xi_{S_j} \right) \cdot \|\beta_j^+ - \beta_j^*\|^2.$$

Putting together the pieces yields, for each $j \in [k]$, the pointwise bound

$$\begin{aligned} & \frac{1}{2} \lambda_{\min} \left(\Xi_{S_j}^\top \Xi_{S_j} \right) \cdot \|\beta_j^+ - \beta_j^*\|^2 \\ & \leq \sum_{j': j' \neq j} \sum_{i=1}^n \mathbf{1} \{ \langle \xi_i, \beta_j - \beta_{j'} \rangle \cdot \langle \xi_i, \beta_j^* - \beta_{j'}^* \rangle \leq 0 \} \langle \xi_i, \beta_j^* - \beta_{j'}^* \rangle^2 + \|P_{\Xi^j(\beta_1, \dots, \beta_k)} \epsilon_{S_j}\|^2. \end{aligned} \tag{6.22}$$

Up to this point, note that all steps of the proof were deterministic. Observe from equation (6.22) that in order to obtain an error bound on the next parameter, we need to control three distinct quantities: (a) the noise term $\|P_{\Xi^j(\beta_1, \dots, \beta_k)} \epsilon_{S_j}\|^2$, (b) the prediction error of the noiseless problem, given by a pairwise sum of terms of the form $\mathbf{1} \{ \langle \xi_i, \beta_j - \beta_{j'} \rangle \cdot \langle \xi_i, \beta_j^* - \beta_{j'}^* \rangle \leq 0 \} \langle \xi_i, \beta_j^* - \beta_{j'}^* \rangle^2$, and (c) the minimum eigen value of the covariate matrix restricted to the set S_j , denoted by $\lambda_{\min} \left(\Xi_{S_j}^\top \Xi_{S_j} \right)$. Since the set S_j is in itself random and depends on the pair (Ξ, ϵ) (since the current parameters were obtained over the course of running the algorithm), obtaining such a bound is especially challenging.

For step (a)—handled by Lemma 29—we apply standard concentration bounds for quadratic forms of sub-Gaussian random variables in conjunction with bounds on the *growth functions* of multi-class classifiers [189]. Crucially, this affords a uniform bound on the noise irrespective of which iterate the alternating minimization update is run from. To show step (b)—in Lemma 30—we generalize a result of Waldspurger [59]. Finally, the key difficulty in step (c) is to control the spectrum of random matrices, rows of which are drawn from (randomly) truncated variants of the Gaussian distribution. The expectation of such a random matrix can be characterized by appealing to tail bounds on the non-central χ^2 distribution, and the Gaussian covariate assumption additionally allows us to show that an analogous result holds for the random matrix with high probability (see Lemma 31). Here, our initialization condition is crucial: the aforementioned singular value control suffices for the sub-matrices formed by the *true* parameters, and we translate these bounds to the sub-matrices generated by random parameters by appealing to the fact that the initialization is sufficiently close to the truth.

Having discussed our proof of the AM update in some detail, let us now turn to a brief discussion of the techniques used to prove Theorems 13 and 14. As mentioned before, our proof of Theorem 13 relies on a lower bound on the eigengap of the population moment. We obtain such a lower bound by appealing to classical moment calculations for suitably truncated Gaussian distributions [182]. Translating these calculations into an eigengap is quite technical, and involves the isolation of many properties of the population moments that may be of independent interest. As briefly alluded to in Section 6.2, the heart of the technical difficulty is due to the fact that that max function is not differentiable, and so moments cannot be calculated by repeated applications of Stein’s lemma like in related problems [98, 188, 190].

In order to establish Theorem 14, we crucially use the scale-invariance property of the initialization along with some arguments involving empirical process theory to show that the goodness-of-fit statistic employed in the algorithm is able to isolate a good initialization. Establishing these bounds requires us to relate the prediction and estimation errors in the problem (in Lemma 43), which may be of independent interest.

6.4 Discussion

We conclude with short discussions of prediction error guarantees, a comparison with adaptivity in convex regression, related models, and future directions.

6.4.1 Guarantees on prediction error

While our principal focus in this paper was on estimation of the unknown parameters $\{\theta_j^*, b_j^*\}_{j=1}^k$, the complementary question of prediction error is also interesting and important. In particular, suppose that we produce the max-affine function estimate $\widehat{\phi}^{(\text{MA})}$ given by $\widehat{\phi}^{(\text{MA})}(x) := \max_{j \in [k]} (\langle x, \widehat{\theta}_j \rangle + \widehat{b}_j)$ for each $x \in \mathbb{R}^d$, and measure its performance via the prediction error

$$\frac{1}{n} \sum_{i=1}^n (\widehat{\phi}^{(\text{ls})}(x_i) - \phi^*(x_i))^2, \quad (6.23)$$

where $\phi^*(x) = \max_{j \in [k]} (\langle x, \theta_j^* \rangle + b_j^*)$ denotes the “true” function. When ϕ^* belongs to the subclass of k -piece affine functions induced by parameters in the set $B_{\text{vol}}(\pi_{\min}, \Delta, \kappa)$ and the covariates are drawn from a Gaussian distribution, our results imply (via Theorem 12, and by using Lemma 43 to translate our estimation error guarantee into a prediction error guarantee) the rate

$$\frac{1}{n} \sum_{i=1}^n (\widehat{\phi}^{(\text{MA})}(x_i) - \phi^*(x_i))^2 \leq C(\pi_{\min}, \Delta, \kappa) \frac{kd}{n} \log(kd) \log(n/kd). \quad (6.24)$$

At least in principle, an explicit dependence on π_{\min} should not be expected in the prediction error, since if a particular pair of parameters (θ, b) attains the maximum extremely rarely (resulting in a small value of π_{\min}), then we may simply drop these parameters from the estimate (and estimate the function with the maximum of the remaining $k - 1$ pieces) without affecting the prediction error significantly. Indeed the minimax risk of prediction (without any requirements of computational efficiency) is known to be independent of the geometry of the problem instance (see, e.g., [158]).

We also note that polynomial-time algorithms with small prediction error are known, without any dependence on π_{\min} . In particular, [191, Theorem 1.8] shows that the sample complexity for obtaining ϵ -accurate estimates in prediction error is bounded by $n \leq \exp\{c_1(k/\epsilon)^{\log k}\} d^{c_2}$ for absolute constants c_1 and c_2 . While the dependence on both ϵ and d can likely be improved⁸, these results provide additional evidence that the prediction error is much less sensitive to the geometry of the instance than the estimation error considered in this paper.

6.4.2 Comparison with algorithms for convex regression

As mentioned earlier, the most standard estimator in convex regression is the convex least squares estimator $\widehat{\phi}^{(\text{ls})}$ defined as in (6.6) which can be computed efficiently as shown in [162, 164]. The

⁸Note that unlike our paper, this work makes only boundedness assumptions on the covariates, and their focus is not on achieving the optimal dimension/samoke size dependence.

performance of $\widehat{\phi}^{(\text{ls})}$ in the max-affine regression model (6.1) has been the subject of some interest in the literature on adaptivity of shape-constrained estimators (see [192] for an overview of results of this type). These results mainly focus on the prediction error:

$$\frac{1}{n} \sum_{i=1}^n (\widehat{\phi}^{(\text{ls})}(x_i) - \phi^*(x_i))^2$$

as opposed to estimation of the parameters $\theta_j^*, b_j^*, j = 1, \dots, k$ which is our main focus. There is actually no natural way of obtaining parameter estimates from $\widehat{\phi}^{(\text{ls})}$ as $\widehat{\phi}^{(\text{ls})}$ will typically be a maximum of a strictly larger than k number of affine functions. Let us now compare our results with the existing results on the prediction error of the convex least squares estimator. When $d = 1$, it has been showed by [193, 194] that

$$\frac{1}{n} \sum_{i=1}^n (\widehat{\phi}^{(\text{ls})}(x_i) - \phi^*(x_i))^2 \leq \frac{k \log n}{n}$$

with high probability assuming that $\{x_i\}_{i=1}^n$ are uniformly spaced on the interval $[0, 1]$. For $d \geq 2$, Han and Wellner [152] studied the adaptivity properties of $\widehat{\phi}^{(\text{bls})}$ which is the least squares estimator over the class of *bounded* convex functions which is different from $\widehat{\phi}^{(\text{ls})}$ and computationally tricky to compute. However, [152] showed that unless $d \leq 4$,

$$\frac{1}{n} \sum_{i=1}^n (\widehat{\phi}^{(\text{bls})}(x_i) - \phi^*(x_i))^2 \leq C_m n^{-4/d} (\log n)^{d+4}$$

with high probability assuming that the covariates are drawn from a distribution supported on a convex polytope with m simplices (the constant pre-factor C_m depends on m). Comparing these two results with our result on the prediction error (6.24), we see that when $d = 1$, our results in the prediction error are strictly weaker than those of prior work [193, 194], but as soon as $d \geq 2$, they are significantly stronger than existing adaptivity results [152], at least for a sub-class of k -affine functions. We emphasize once again that the focus of the body of work differs from ours, and so the comparison presented above is necessarily incomplete.

A parallel line of work (including our own) eschews the c-LSE (and its variants) entirely and pursues a different avenue to alleviate the curse of dimensionality⁹, by directly fitting convex functions consisting of a certain number of affine pieces [55], or more broadly, by treating the number of affine functions as a tuning parameter to be chosen in a data-dependent fashion via cross-validation [153]. Hannah and Dunson [153] showed that performing estimation under a carefully chosen sequence of models of the form (6.1) via their “convex adaptive partitioning”, or CAP estimator is able to obtain consistent prediction rates for general convex regression problems. However, it is unclear if the CAP estimator is able to avoid the curse of dimensionality in the special case when the true function is k -piece affine.

⁹Note that setting $k \sim n^{d/(4+d)}$, we can (essentially) recover the entire class of convex functions from the maxima of k affine functions (see, e.g., Balázs [154]), so interesting parametric structure is only expected to emerge when k is essentially constant, or grows very slowly with n .

6.4.3 Related models

Models closely related to (6.1) also appear in second price auctions, where an item having d features is bid on and sold to the highest bidder at the second highest bid [195, 196]. Assuming that each of k user groups bids on an item and that each bid is a linear function of the features, one can use a variant of the model (6.1) with the \max function replaced by the second order statistic to estimate the individual bids of the user groups based on historical data. Another related problem is that of multi-class classification [189], in which one of k labels is assigned to each sample based on the argmax function, i.e., for a class of functions \mathcal{F} , we have the model $Y = \operatorname{argmax}_{1 \leq j \leq k} f_j(X)$ for j distinct functions $f_1, \dots, f_k \in \mathcal{F}$. When \mathcal{F} is the class of linear functions based on d features, this can be viewed as the “classification” variant of our regression problem.

The model (6.1) can also be seen as a special case of multi-index models [197, 198] as well as mixture-of-experts models [199, 200]. Multi-index models are of the form $Y = g(\langle \theta_1^*, X \rangle, \dots, \langle \theta_k^*, X \rangle) + \epsilon$ for an *unknown* function g and this function g is taken to be the $\max(\cdot)$ function in the model (6.1). In the mixture-of-experts model, the covariate space is partitioned into k regions via certain *gating functions*, and the observation model is given by k distinct regression functions: one on each region. The model (6.1) is clearly a member of this class, since the $\max(\cdot)$ function implicitly defines a partition of \mathbb{R}^d depending on which of the k linear functions of X attains the maximum, and on each of these partitions, the regression function is linear in X .

6.4.4 Open Problems

In this paper, we analyzed a natural alternating minimization algorithm for estimating the maximum of unknown affine functions, and established that it enjoys local linear convergence to a ball around the optimal parameters. We also proposed an initialization based on PCA followed by random search in a lower-dimensional space. An interesting open question is if there are other efficient methods besides random search that work just as well post dimensionality reduction. Another interesting question has to do with the necessity of dimensionality reduction: in simulations (see, e.g., Figure 6.2), we have observed that if the AM algorithm is repeatedly initialized in $(d + 1)$ -dimensional space without dimensionality reduction, then the number of repetitions required to obtain an initialization from which it succeeds (with high probability) is similar to the number of repetitions required after dimensionality reduction. This suggests that our (sufficient) initialization condition (6.14a) may be too stringent, and that the necessary conditions on the initialization to ensure convergence of the AM algorithm are actually much weaker. We leave such a characterization for future work, but note that some such conditions must exist: the AM algorithm when run from a single random initialization, for instance, fails with constant probability when $k \geq 3$. Understanding the behavior of the randomly initialized AM algorithm is also an open problem in the context of phase retrieval [59, 201].

We note that once again that the Gaussian assumption made in this paper for convenience of analysis can be relaxed to allow (for instance) log-concave covariate distributions, which includes the uniform distribution on $[-1, 1]^d$ common in nonparametric statistics. Such an extension requires

significant technical effort and the structure of the proof also changes slightly; simultaneously, the dependence of the eventual error bounds on the parameter π_{\min} is also different in the more general setting. In particular, Lemmas 30-32 in the current paper must be extended, and this requires, among other things, an analysis of random matrices whose rows are drawn from a (truncated) small-ball distribution. Our companion paper [202] is also concerned with the *universal* setting in which guarantees are proved uniformly over all choices of parameters once the covariates have been drawn, in contrast to the setting of the current paper in which parameters are fixed in advance. Universal guarantees are commonly sought out in statistical signal processing applications, including phase retrieval [56].

In the broader context of max-affine estimation, it is also interesting to analyze other non-convex procedures (e.g. gradient descent) to obtain conditions under which they obtain accurate parameter estimates. The CAP estimator of Hannah and Dunson [153] and the adaptive max-affine partitioning algorithm of Balázs [154] are also interesting procedures for estimation under these models, and it would be interesting to analyze their performance when the number of affine pieces k is fixed and known. For applications in which the dimension d is very large, it is also interesting to study the model with additional restrictions of sparsity on the unknown parameters—such problems are known to exhibit interesting statistical-computational gaps even in the special case of sparse phase retrieval (see, e.g., Cai et al. [203]). We also note that in practice, and especially for convex regression, the parameter k would be unknown and must be estimated (say) via cross-validation. Understanding the behavior of such a two-stage estimator is an important direction of future work.

Appendix

We now present proofs of our main results. We assume throughout that the sample size n is larger than some universal constant. Values of constants c, c_1, c', \dots may change from line to line. Statements of our theorems, for instance, minimize the number of constants by typically using one of these to denote a large enough constant, and another to denote a small enough constant.

6.5 Proofs

6.5.1 Technical results concerning the global LSE

In this section, we provide a proof of the existence of the global least squares estimator that was stated in the main text. We also state and prove a lemma that shows that the global LSE is a fixed point of the AM update under a mild technical condition.

6.5.2 Proof of Lemma 27

Fix data $(x_1, y_1), \dots, (x_n, y_n)$ and let

$$L(\gamma_1, \dots, \gamma_k) := \sum_{i=1}^n \left(y_i - \max_{j \in [k]} \langle \xi_i, \gamma_j \rangle \right)^2$$

denote the objective function in (6.5) with $\xi_i := (x_i, 1)$. The goal is to show that a global minimizer of $L(\gamma_1, \dots, \gamma_k)$ over $\gamma_1, \dots, \gamma_k \in \mathbb{R}^{d+1}$ exists. For $\gamma_1, \dots, \gamma_k \in \mathbb{R}^{d+1}$, let $S_1^\gamma, \dots, S_k^\gamma$ denote a fixed partition of $[n]$ having the property that

$$\langle \xi_i, \gamma_j \rangle = \max_{u \in [k]} \langle \xi_i, \gamma_u \rangle \quad \text{for every } j \in [k] \text{ and } i \in S_j^\gamma.$$

Also, let $\widehat{\beta}_1^\gamma, \dots, \widehat{\beta}_k^\gamma$ denote the solution to the following constrained least squares problem:

$$\begin{aligned} & \underset{\beta_1, \dots, \beta_k}{\text{minimize}} && \sum_{j=1}^k \sum_{i \in S_j^\gamma} (y_i - \langle \xi_i, \beta_j \rangle)^2 \\ & \text{subject to} && \langle \xi_i, \beta_j \rangle \geq \langle \xi_i, \beta_u \rangle, u, j \in [k], i \in S_j^\gamma. \end{aligned}$$

Note that the above quadratic problem is feasible as $\gamma_1, \dots, \gamma_k$ satisfies the constraint and, consequently, $\widehat{\beta}_1^\gamma, \dots, \widehat{\beta}_k^\gamma$ exists uniquely for every $\gamma_1, \dots, \gamma_k \in \mathbb{R}^{d+1}$. Note further that, by construction,

$$L(\widehat{\beta}_1^\gamma, \dots, \widehat{\beta}_k^\gamma) \leq L(\gamma_1, \dots, \gamma_k).$$

and that the set

$$\Delta := \left\{ (\widehat{\beta}_1^\gamma, \dots, \widehat{\beta}_k^\gamma) : \gamma_1, \dots, \gamma_k \in \mathbb{R}^{d+1} \right\}$$

is finite because $\widehat{\beta}_1^\gamma, \dots, \widehat{\beta}_k^\gamma$ depends on $\gamma_1, \dots, \gamma_k$ only through the partition $S_1^\gamma, \dots, S_k^\gamma$ and the number of possible such partitions of $[n]$ is obviously finite. Finally, it is evident that

$$(\widehat{\beta}_1^{(\text{ls})}, \dots, \widehat{\beta}_k^{(\text{ls})}) = \underset{(\beta_1, \dots, \beta_k) \in \Delta}{\text{argmin}} L(\beta_1, \dots, \beta_k)$$

is a global minimizer of $L(\gamma_1, \dots, \gamma_k)$ as

$$L(\widehat{\beta}_1^{(\text{ls})}, \dots, \widehat{\beta}_k^{(\text{ls})}) \leq L(\widehat{\beta}_1^\gamma, \dots, \widehat{\beta}_k^\gamma) \leq L(\gamma_1, \dots, \gamma_k)$$

for every $\gamma_1, \dots, \gamma_k$. This concludes the proof of Lemma 27.

6.5.3 Fixed point of AM update

The following lemma establishes that the global LSE is a fixed point of the AM update under a mild technical condition.

Lemma 28. *Consider the global least squares estimator (6.5). Suppose that the k values $\langle \xi_i, \widehat{\beta}_j^{(ls)} \rangle$ for $j = 1, \dots, k$ are distinct for each $i \in [n]$. Then*

$$\widehat{\beta}_j^{(ls)} \in \operatorname{argmin}_{\beta \in \mathbb{R}^{d+1}} \sum_{i \in S_j(\widehat{\beta}_1^{(ls)}, \dots, \widehat{\beta}_k^{(ls)})} (y_i - \langle \xi_i, \beta \rangle)^2 \quad \text{for every } j \in [k]. \quad (6.25)$$

Proof. It is clearly enough to prove (6.25) for $j = 1$. Suppose that $\widehat{\beta}_1^{(ls)}$ does not minimize the least squares criterion over $S_1(\widehat{\beta}_1^{(ls)}, \dots, \widehat{\beta}_k^{(ls)})$. Let

$$\widehat{\gamma}_1^{(ls)} \in \operatorname{argmin}_{\beta \in \mathbb{R}^{d+1}} \sum_{i \in S_1(\widehat{\beta}_1^{(ls)}, \dots, \widehat{\beta}_k^{(ls)})} (y_i - \langle \xi_i, \beta \rangle)^2$$

be any other least squares minimizer over $S_1(\widehat{\beta}_1^{(ls)}, \dots, \widehat{\beta}_k^{(ls)})$ and let, for $\epsilon > 0$,

$$\widetilde{\beta}_1 := \widehat{\beta}_1^{(ls)} + \epsilon \left(\widehat{\gamma}_1^{(ls)} - \widehat{\beta}_1^{(ls)} \right).$$

When $\epsilon > 0$ is sufficiently small, we have

$$S_j(\widetilde{\beta}_1, \widehat{\beta}_2^{(ls)}, \dots, \widehat{\beta}_k^{(ls)}) = S_j(\widehat{\beta}_1^{(ls)}, \dots, \widehat{\beta}_k^{(ls)}) \quad \text{for every } j \in [k]$$

due to the no ties assumption and the fact that $\widetilde{\beta}_1$ and $\widehat{\beta}_1^{(ls)}$ can be made arbitrarily close as ϵ becomes small. Thus, if

$$U(\beta_1, \dots, \beta_k) := \sum_{i=1}^n \left(y_i - \max_{j \in [k]} \langle \xi_i, \beta_j \rangle \right)^2 = \sum_{j \in [k]} \sum_{i \in S_j(\beta_1, \dots, \beta_k)} (y_i - \langle \xi_i, \beta_j \rangle)^2,$$

then

$$\begin{aligned} & U(\widetilde{\beta}_1, \widehat{\beta}_2^{(ls)}, \dots, \widehat{\beta}_k^{(ls)}) \\ &= \sum_{i \in S_1(\widetilde{\beta}_1, \widehat{\beta}_2^{(ls)}, \dots, \widehat{\beta}_k^{(ls)})} (y_i - \langle \xi_i, \widetilde{\beta}_1 \rangle)^2 + \sum_{j \geq 2} \sum_{i \in S_j(\widetilde{\beta}_1, \widehat{\beta}_2^{(ls)}, \dots, \widehat{\beta}_k^{(ls)})} (y_i - \langle \xi_i, \widehat{\beta}_j^{(ls)} \rangle)^2 \\ &= \sum_{i \in S_1(\widehat{\beta}_1^{(ls)}, \widehat{\beta}_2^{(ls)}, \dots, \widehat{\beta}_k^{(ls)})} (y_i - \langle \xi_i, \widetilde{\beta}_1 \rangle)^2 + \sum_{j \geq 2} \sum_{i \in S_j(\widehat{\beta}_1^{(ls)}, \widehat{\beta}_2^{(ls)}, \dots, \widehat{\beta}_k^{(ls)})} (y_i - \langle \xi_i, \widehat{\beta}_j^{(ls)} \rangle)^2 \\ &< \sum_{i \in S_1(\widehat{\beta}_1^{(ls)}, \widehat{\beta}_2^{(ls)}, \dots, \widehat{\beta}_k^{(ls)})} (y_i - \langle \xi_i, \widehat{\beta}_1^{(ls)} \rangle)^2 + \sum_{j \geq 2} \sum_{i \in S_j(\widehat{\beta}_1^{(ls)}, \widehat{\beta}_2^{(ls)}, \dots, \widehat{\beta}_k^{(ls)})} (y_i - \langle \xi_i, \widehat{\beta}_j^{(ls)} \rangle)^2 \\ &= U(\widehat{\beta}_1^{(ls)}, \widehat{\beta}_2^{(ls)}, \dots, \widehat{\beta}_k^{(ls)}) \end{aligned}$$

where the strict inequality above comes from the fact that $\widetilde{\beta}_1$ is closer to the least squares solution $\widehat{\gamma}_1^{(ls)}$ compared to $\widehat{\beta}_1^{(ls)}$. This leads to a contradiction as the criterion function is smaller than its value at a global minimizer, thereby concluding the proof. \square

6.6 Proof of Theorem 12

Let us begin by introducing some shorthand notation, and providing a formal statement of the probability bound guaranteed by the theorem. For a scalar w^* , vectors $u^* \in \mathbb{R}^d$ and $v^* = (u^*, w^*) \in \mathbb{R}^{d+1}$, and a positive scalar r , let $\mathcal{B}_{v^*}(r) = \left\{ v \in \mathbb{R}^{d+1} : \frac{\|v-v^*\|}{\|u^*\|} \leq r \right\}$, and let

$$\mathcal{I}(r; \{\beta_j^*\}_{j=1}^k) = \left\{ \beta_1, \dots, \beta_k \in \mathbb{R}^{d+1} : \exists c > 0 : c(\beta_i - \beta_j) \in \mathcal{B}_{\beta_i^* - \beta_j^*}(r) \text{ for all } 1 \leq i \neq j \leq k \right\}.$$

Also, use the shorthand

$$\vartheta_t \left(r; \{\beta_j^*\}_{j=1}^k \right) := \sup_{\beta_1^{(0)}, \dots, \beta_k^{(0)} \in \mathcal{I}(r)} \sum_{j=1}^k \|\beta_j^{(t)} - \beta_j^*\|^2 - \left(\frac{3}{4} \right)^t \left(\sum_{j=1}^k \|c^* \beta_j^{(0)} - \beta_j^*\|^2 \right), \text{ and}$$

$$\delta_{n,\sigma}^{\mathcal{N}}(d, k, \pi_{\min}) := \sigma^2 \frac{kd}{\pi_{\min}^3 n} \log(kd) \log(n/kd)$$

to denote the error tracked over iterations (with c^* denoting the smallest $c > 0$ such that $c(\beta_i - \beta_j) \in \mathcal{B}_{\beta_i^* - \beta_j^*}(r)$ for all $1 \leq i \neq j \leq k$), and a proxy for the final statistical rate, respectively.

Theorem 12 states that there are universal constants c_1 and c_2 such that if the sample size obeys the condition $n \geq n_{\text{AM}}(c_1)$, then we have

$$\Pr \left\{ \max_{t \geq 1} \vartheta_t \left(c_2 \frac{\pi_{\min}^6}{k^2 \kappa}; \{\beta_j^*\}_{j=1}^k \right) \geq c_1 \delta_{n,\sigma}^{\mathcal{N}}(d, k, \pi_{\min}) \right\} \leq c_2 \left(k \exp \left(-c_1 n \frac{\pi_{\min}^6}{k^2} \right) + \frac{k^2}{n^7} \right). \quad (6.26)$$

Let us assume, without loss of generality, that the scalar c^* above is equal to 1. It is convenient to state and prove another result that guarantees a one-step contraction, from which Theorem 12 follows as a corollary. In order to state this result, we assume that one step of the alternating minimization update (6.12a)-(6.12b) is run starting from the parameters $\{\beta_j\}_{j=1}^k$ to produce the next iterate $\{\beta_j^+\}_{j=1}^k$. We use the shorthand

$$\begin{aligned} v_{i,j}^* &= \beta_i^* - \beta_j^*, \\ v_{i,j} &= \beta_i - \beta_j, \text{ and} \\ v_{i,j}^+ &= \beta_i^+ - \beta_j^+. \end{aligned}$$

Also recall the definitions of the geometric quantities (Δ, κ) . The following proposition guarantees the one step contraction bound.

Proposition 15. *There exist universal constants c_1 and c_2 such that*

(a) *If the sample size satisfies the bound $n \geq c_1 \max \left\{ \frac{dk}{\pi_{\min}^3}, \log n \cdot \frac{k^2}{\pi_{\min}^6} \right\}$, then for all parameters $\{\beta_j\}_{j=1}^k$ satisfying*

$$\max_{1 \leq j \neq j' \leq k} \frac{\|v_{j,j'} - v_{j,j'}^*\|}{\|\theta_j^* - \theta_{j'}^*\|} \log^{3/2} \left(\frac{\|\theta_j^* - \theta_{j'}^*\|}{\|v_{j,j'} - v_{j,j'}^*\|} \right) \leq c_2 \frac{\pi_{\min}^6}{k^2 \kappa}, \quad (6.27a)$$

we have, simultaneously for all pairs $1 \leq j \neq \ell \leq k$, the bound

$$\begin{aligned} \frac{\|v_{j,\ell}^+ - v_{j,\ell}^*\|^2}{\|\theta_j^* - \theta_\ell^*\|^2} &\leq \max \left\{ \frac{d\kappa}{\pi_{\min}^3 n}, \frac{1}{4k} \right\} \left(\sum_{j'=1}^k \frac{\|v_{j,j'} - v_{j,j'}^*\|^2}{\|\theta_j^* - \theta_{j'}^*\|^2} + \frac{\|v_{\ell,j'} - v_{\ell,j'}^*\|^2}{\|\theta_\ell^* - \theta_{j'}^*\|^2} \right) \\ &\quad + c_1 \frac{\sigma^2}{\Delta} \frac{kd}{\pi_{\min}^3 n} \log(n/d) \end{aligned} \quad (6.27b)$$

with probability exceeding $1 - c_1 \left(k \exp \left(-c_2 n \frac{k^2}{\pi_{\min}^6} \right) + \frac{k^2}{n^7} \right)$.

(b) If the sample size satisfies the bound $n \geq c_1 \max \left\{ \frac{dk}{\pi_{\min}^3}, \log n \cdot \frac{k^2}{\pi_{\min}^6} \right\}$, then for all parameters $\{\beta_j\}_{j=1}^k$ satisfying

$$\max_{1 \leq j \neq j' \leq k} \frac{\|v_{j,j'} - v_{j,j'}^*\|}{\|\theta_j^* - \theta_{j'}^*\|} \log^{3/2} \left(\frac{\|\theta_j^* - \theta_{j'}^*\|}{\|v_{j,j'} - v_{j,j'}^*\|} \right) \leq c_2 \frac{\pi_{\min}^6}{k^2}, \quad (6.28a)$$

we have the overall estimation error bound

$$\sum_{i=1}^k \|\beta_j^+ - \beta_j^*\|^2 \leq \frac{3}{4} \cdot \left(\sum_{i=1}^k \|\beta_j - \beta_j^*\|^2 \right) + c_1 \sigma^2 \frac{kd}{\pi_{\min}^3 n} \log(k) \log(n/dk) \quad (6.28b)$$

with probability exceeding $1 - c_1 \left(k \exp \left(-c_2 n \frac{k^2}{\pi_{\min}^6} \right) + \frac{k^2}{n^7} \right)$.

Let us briefly comment on why Proposition 15 implies Theorem 12 as a corollary. Clearly, equations (6.28a) and (6.28b) in conjunction show that the estimation error decays geometrically after running one step of the algorithm. The only remaining detail to be verified is that the next iterates $\{\beta_j^+\}_{j=1}^k$ also satisfy condition (6.27a) provided the sample size is large enough; in that case, the one step estimation bound (6.28b) can be applied recursively to obtain the final bound (6.14b).

With the constant c_2 from the proposition, let r_b be the largest value in the interval $[0, e^{-3/2}]$ such that $r_b \log^{3/2}(1/r_b) \leq c_2 \frac{\pi_{\min}^6}{k^2}$. Similarly, let r_a be the largest value in the interval $[0, e^{-3/2}]$ such that $r_a \log^{3/2}(1/r_a) \leq c_2 \frac{\pi_{\min}^6}{k^2 \kappa}$. Bounds on both of these values will be used repeatedly later on.

Assume that the current parameters satisfy the bound (6.27a). Choosing $n \geq 4\kappa d / \pi_{\min}^3$ and applying inequality (6.27b), we have, for each pair $1 \leq j \neq \ell \leq k$, the bound

$$\begin{aligned} \frac{\|v_{j,\ell}^+ - v_{j,\ell}^*\|^2}{\|\theta_j^* - \theta_\ell^*\|^2} &\leq \frac{1}{4k} \left(\sum_{j'=1}^k \frac{\|v_{j,j'} - v_{j,j'}^*\|^2}{\|\theta_j^* - \theta_{j'}^*\|^2} + \frac{\|v_{\ell,j'} - v_{\ell,j'}^*\|^2}{\|\theta_\ell^* - \theta_{j'}^*\|^2} \right) + c_1 \frac{\sigma^2}{\Delta} \frac{kd}{\pi_{\min}^3 n} \log(n/d) \\ &\leq \frac{1}{2} r_a^2 + c_1 \frac{\sigma^2}{\Delta} \frac{kd}{\pi_{\min}^3 n} \log(n/d). \end{aligned}$$

Further, if $n \geq C \sigma^2 \frac{kd}{\pi_{\min}^3 \Delta r_a^2}$ for a sufficiently large constant C , we have

$$\frac{\|v_{j,\ell}^+ - v_{j,\ell}^*\|^2}{\|\theta_j^* - \theta_\ell^*\|^2} \leq r_a^2.$$

Thus, the parameters $\{\beta_j^+\}_{j=1}^k$ satisfy inequality (6.27a) for the sample size choice required by Theorem 12. Finally, noting, for a pair of small enough scalars (a, b) , the implication

$$a \leq \frac{b}{2} \log^{-3/2}(1/b) \implies a \log^{3/2}(1/a) \leq b,$$

and adjusting the constants appropriately to simplify the probability statement completes the proof of the theorem. It now remains to establish Proposition 15.

6.6.1 Proof of Proposition 15

Recall that we denote by

$$S_j(\beta_1, \dots, \beta_k) := \left\{ 1 \leq i \leq n : \langle \xi_i, \beta_j \rangle = \max_{1 \leq u \leq k} (\langle \xi_i, \beta_u \rangle) \right\},$$

the indices of the rows for which β_j attains the maximum, and we additionally keep this sets disjoint by breaking ties lexicographically. To lighten notation, we use the shorthand

$$\Xi^j(\beta_1, \dots, \beta_k) := \Xi_{S_j(\beta_1, \dots, \beta_k)}.$$

Recall the notation

$$\mathcal{B}_{v^*}(r) = \left\{ v \in \mathbb{R}^{d+1} : \frac{\|v - v^*\|}{\|u^*\|} \leq r \right\}$$

introduced before, and the definitions of the pair of scalars (r_a, r_b) . To be agnostic to the scale invariance of the problem, we set $c^* = 1$ and define the set of parameters

$$\mathcal{I}(r) = \left\{ \beta_1, \dots, \beta_k : v_{i,j} \in \mathcal{B}_{v_{i,j}^*}(r) \text{ for all } 1 \leq i \neq j \leq k \right\},$$

and use the shorthand $\mathcal{I}_a := \mathcal{I}(r_a)$ and $\mathcal{I}_b := \mathcal{I}(r_b)$, to denote the set of parameters satisfying conditions (6.27a) and (6.28a), respectively,

Finally, recall the deterministic bound (6.22) established in Section 6.3.4, restated below for convenience.

$$\begin{aligned} \frac{1}{2} \lambda_{\min} \left(\Xi_{S_j}^\top \Xi_{S_j} \right) \cdot \|\beta_j^+ - \beta_j^*\|^2 &\leq \sum_{j': j' \neq j} \sum_{i=1}^n \mathbf{1} \{ \langle \xi_i, v_{j,j'} \rangle \cdot \langle \xi_i, v_{j,j'}^* \rangle \leq 0 \} \langle \xi_i, v_{j,j'}^* \rangle^2 \\ &\quad + \|P_{\Xi^j(\beta_1, \dots, \beta_k) \in S_j}\|^2. \end{aligned}$$

It suffices to show high probability bounds on the various quantities appearing in this bound. First, we claim that the noise terms are uniformly bounded as

$$\Pr \left\{ \sup_{\beta_1, \dots, \beta_k \in \mathbb{R}^{d+1}} \sum_{j=1}^k \|P_{\Xi^j(\beta_1, \dots, \beta_k)} \epsilon_{S_j(\beta_1, \dots, \beta_k)}\|^2 \geq 2\sigma^2 k(d+1) \log(kd) \log(n/kd) \right\} \leq \binom{n}{kd}^{-1}, \quad (6.29a.I)$$

$$\Pr \left\{ \sup_{\beta_1, \dots, \beta_k \in \mathbb{R}^{d+1}} \|P_{\Xi^j(\beta_1, \dots, \beta_k)} \epsilon_{S_j(\beta_1, \dots, \beta_k)}\|^2 \geq 2\sigma^2 k(d+1) \log(n/d) \right\} \leq \binom{n}{d}^{-1} \text{ for each } j \in [k]. \quad (6.29a.II)$$

Second, we show that the indicator quantities are simultaneously bounded for all j, j' pairs. In particular, we claim that there exists a tuple of universal constants (C, c_1, c_2, c') such that for each positive scalar $r \leq 1/24$, we have

$$\begin{aligned} & \Pr \left\{ \exists 1 \leq j \neq j' \leq k, v_{j,j'} \in \mathcal{B}_{v_{j,j'}^*}(r) : \sum_{j':j' \neq j} \sum_{i=1}^n \mathbf{1} \{ \langle \xi_i, v_{j,j'} \rangle \cdot \langle \xi_i, v_{j,j'}^* \rangle \leq 0 \} \langle \xi_i, v_{j,j'}^* \rangle^2 \right. \\ & \left. \geq C \max\{d, nr \log^{3/2}(1/r)\} \sum_{j':j' \neq j} \|v_{j,j'} - v_{j,j'}^*\|^2 \right\} \leq c_1 \binom{k}{2} \left\{ ne^{-c_2 n} + e^{-c' \max\{d, 10 \log n\}} \right\}. \end{aligned} \quad (6.29b)$$

Finally, we show a bound on the LHS of the bound (6.22) by handling the singular values of (random) sub-matrices of Ξ with a uniform bound. In particular, we claim that there are universal constants (C, c, c') such that if $n \geq C \max \left\{ \frac{dk}{\pi_{\min}^3}, \log n \cdot \frac{k^2}{\pi_{\min}^6} \right\}$, then for each $j \in [k]$, we have

$$\begin{aligned} & \Pr \left\{ \inf_{\beta_1, \dots, \beta_k \in \mathcal{I}_b} \lambda_{\min} \left(\Xi^j(\beta_1, \dots, \beta_k)^\top \cdot \Xi^j(\beta_1, \dots, \beta_k) \right) \leq C \pi_{\min}^3 n \right\} \\ & \leq c \exp \left(-cn \frac{\pi_{\min}^6}{k^2} \right) + c' n^{-10}. \end{aligned} \quad (6.29c)$$

Notice that claim (6.29a.I) implicitly defines a high probability event $\mathcal{E}^{(a.I)}$, claim (6.29a.II) defines high probability events $\mathcal{E}_j^{(a.II)}$, claim (6.29b) defines a high probability event $\mathcal{E}^{(b)}(r)$, and claim (6.29c) defines high probability events $\mathcal{E}_j^{(c)}$. Define the intersection of these events as

$$\mathcal{E}(r) := \mathcal{E}^{(a.I)} \cap \left(\bigcap_{j \in [k]} \mathcal{E}_j^{(a.II)} \right) \cap \mathcal{E}^{(b)}(r) \cap \left(\bigcap_{j \in [k]} \mathcal{E}_j^{(c)} \right),$$

and note that the claims in conjunction with the union bound guarantee that if the condition on the sample size $n \geq c_1 \max \left\{ \frac{dk}{\pi_{\min}^3}, \log n \cdot \frac{k^2}{\pi_{\min}^6} \right\}$ holds, then for all $r \leq r_b$, we have

$$\Pr \{ \mathcal{E}(r) \} \geq 1 - c_1 \left(k \exp \left(-c_2 n \frac{\pi_{\min}^6}{k^2} \right) + \frac{k^2}{n^7} \right),$$

where we have adjusted constants appropriately in stating the bound. We are now ready to prove the two parts of the proposition.

Proof of part (a): Work on the event $\mathcal{E}(r_a)$. Normalizing inequality (6.22) by n and using claims (6.29a.II), (6.29b), and (6.29c) with $r = r_a$ then yields, simultaneously for all $j \in [k]$, the bound

$$\begin{aligned} \|\beta_j^+ - \beta_j^*\|^2 &\leq C \max \left\{ \frac{d}{\pi_{\min}^3 n}, \frac{r_a}{\pi_{\min}^3} \log^{3/2}(1/r_a) \right\} \sum_{j': j' \neq j} \|v_{j,j'} - v_{j,j'}^*\|^2 + C' \sigma^2 \frac{kd}{\pi_{\min}^3 n} \log(n/d) \\ &\stackrel{(i)}{\leq} \max \left\{ \frac{Cd}{\pi_{\min}^3 n}, \frac{1}{4k\kappa} \right\} \sum_{j': j' \neq j} \|v_{j,j'} - v_{j,j'}^*\|^2 + C' \sigma^2 \frac{kd}{\pi_{\min}^3 n} \log(n/d), \end{aligned}$$

where in step (i), we have used the definition of the quantity r_a . Using this bound for the indices j, ℓ in conjunction with the definition of the quantity κ proves inequality (6.27b). \square

Proof of part (b): We now work on the event $\mathcal{E}(r_b)$ and proceed again (see equation (6.22)) from the bound

$$\begin{aligned} \|\beta_j^+ - \beta_j^*\|^2 &\leq C \max \left\{ \frac{d}{\pi_{\min}^3 n}, \frac{r_b}{\pi_{\min}^3} \log^{3/2}(1/r_b) \right\} \sum_{j': j' \neq j} \|v_{j,j'} - v_{j,j'}^*\|^2 \\ &\quad + \frac{C}{\pi_{\min}^3 n} \|P_{\Xi^j(\beta_1, \dots, \beta_k)} \epsilon_{S_j}\|^2. \end{aligned}$$

Summing over $j \in [k]$ and using the fact that $\|a + b\|^2 \leq 2\|a\|^2 + 2\|b\|^2$, we obtain

$$\begin{aligned} \sum_{j=1}^k \|\beta_j^+ - \beta_j^*\|^2 &\leq C \max \left\{ \frac{kd}{\pi_{\min}^3 n}, \frac{kr_b}{\pi_{\min}^3} \log^{3/2}(1/r_b) \right\} \left(\sum_{j=1}^k \|\beta_j - \beta_j^*\|^2 \right) \\ &\quad + \frac{C}{\pi_{\min}^3 n} \sum_{j \in [k]} \|P_{\Xi^j(\beta_1, \dots, \beta_k)} \epsilon_{S_j}\|^2 \\ &\stackrel{(ii)}{\leq} \frac{3}{4} \left(\sum_{j=1}^k \|\beta_j - \beta_j^*\|^2 \right) + C' \sigma^2 \frac{kd}{\pi_{\min}^3 n} \log(k) \log(n/kd), \end{aligned}$$

where in step (ii), we have used the definition of the quantity r_b , the bound $n \geq Ckd/\pi_{\min}^3$, and claim (6.29a.I). This completes the proof. \square

We now prove each of the claims in turn. This constitutes the technical meat of our proof, and involves multiple technical lemmas whose proofs are postponed to the end of the section.

Proof of claims (6.29a.I) and (6.29a.II): We begin by stating a general lemma about concentration properties of the noise.

Lemma 29. Consider a random variable $z \in \mathbb{R}^n$ with i.i.d. σ -sub-Gaussian entries, and a fixed matrix $\Xi \in \mathbb{R}^{n \times (d+1)}$. Then, we have

$$\sup_{\beta_1, \dots, \beta_k \in \mathbb{R}^{d+1}} \sum_{j=1}^k \|P_{\Xi^j(\beta_1, \dots, \beta_k)} z\|^2 \leq 2\sigma^2 k(d+1) \log(kd) \log(n/kd) \quad (6.30a)$$

with probability greater than $1 - \binom{n}{kd}^{-1}$ and

$$\sup_{\beta_1, \dots, \beta_k \in \mathbb{R}^{d+1}} \max_{j \in [k]} \|P_{\Xi^j(\beta_1, \dots, \beta_k)} z_{S_j(\beta_1, \dots, \beta_k)}\|^2 \leq 2\sigma^2 k(d+1) \log(n/d) \quad (6.30b)$$

with probability greater than $1 - \binom{n}{d}^{-1}$.

The proof of the claims follows directly from Lemma 29, since the noise vector ϵ is independent of the matrix Ξ , and $\mathcal{I}_b \subseteq (\mathbb{R}^{d+1})^{\otimes k}$. \square

Proof of claim (6.29b): We now state a lemma that directly handles indicator functions as they appear in the claim.

Lemma 30. Let $u^* \in \mathbb{R}^d$ and $w^* \in \mathbb{R}$, and consider a fixed parameter $v^* = (u^*, w^*) \in \mathbb{R}^{d+1}$. Then there are universal constants (c_1, c_2, c_3, c_4) such that for all positive scalars $r \leq 1/24$, we have

$$\sup_{v \in \mathcal{B}_{v^*}(r)} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{\langle \xi_i, v \rangle \cdot \langle \xi_i, v^* \rangle \leq 0\}} \langle \xi_i, v^* \rangle^2 \right) / \|v - v^*\|^2 \leq c_1 \cdot \max \left\{ \frac{d}{n}, r \log^{3/2} \left(\frac{1}{r} \right) \right\}$$

with probability exceeding $1 - c_1 e^{-c_2 \max\{d, 10 \log n\}} - c_3 n e^{-c_4 n}$. Here, we adopt the convention that $0/0 = 0$.

Applying Lemma 30 with $v = v_{j,j'}$ and $v^* = v_{j',j'}$ for all pairs (j, j') and using a union bound directly yields the claim. \square

Proof of claim (6.29c): For this claim, we state three technical lemmas pertaining to the singular values of random matrices whose rows are formed by truncated Gaussian random vectors. We let $\text{vol}(K)$ denote the volume of a set $K \subseteq \mathbb{R}^d$ with respect to d -dimensional standard Gaussian measure, i.e., with $\text{vol}(K) = \Pr\{Z \in K\}$ for $Z \sim \mathcal{N}(0, I_d)$.

Lemma 31. Suppose n vectors $\{x_i\}_{i=1}^n$ are drawn i.i.d. from $\mathcal{N}(0, I_d)$, and $K \subseteq \mathbb{R}^d$ is a fixed convex set. Then there exists a tuple of universal constants (c_1, c_2) such that if $\text{vol}^3(K)n \geq c_1 d \log^2(1/\text{vol}(K))$, then

$$\lambda_{\min} \left(\sum_{i: x_i \in K} \xi_i \xi_i^\top \right) \geq c_2 \text{vol}^3(K) \cdot n$$

with probability greater than $1 - c_1 \exp\left(-c_2 n \frac{\text{vol}^4(K)}{\log^2(1/\text{vol}(K))}\right) - c_1 \exp(-c_2 n \cdot \text{vol}(K))$.

For a pair of scalars (w, w') and d -dimensional vectors (u, u') , define the *wedge* formed by the $d + 1$ -dimensional vectors $v = (u, w)$ and $v' = (u', w')$ as the region

$$W(v, v') = \{x \in \mathbb{R}^d : (\langle x, u \rangle + w) \cdot (\langle x, u' \rangle + w') \leq 0\},$$

and let $\mathcal{W}_\delta = \{W = W(v, v') : \text{vol}(W) \leq \delta\}$ denote the set of all wedges with Gaussian volume less than δ . The next lemma bounds the maximum singular value of a sub-matrix formed by any such wedge.

Lemma 32. *There is a tuple of universal constants (c_1, c_2) such that if $n \geq c_1 \max\{d, \frac{\log n}{\delta^2}\}$, then*

$$\sup_{W \in \mathcal{W}_\delta} \lambda_{\max} \left(\sum_{i: x_i \in W} \xi_i \xi_i^\top \right) \leq c_1 n \sqrt{\delta}$$

with probability greater than $1 - \exp(-c_2 n \delta^2) - n^{-10}$.

We are now ready to proceed to a proof of claim (6.29c). For convenience, introduce the shorthand notation

$$S_j^* := S_j(\beta_1^*, \dots, \beta_k^*)$$

to denote the set of indices corresponding to observations generated by the true parameter β_j^* . Letting $A \Delta B := (A \setminus B) \cup (B \setminus A)$ denote the symmetric difference between two sets A and B , we have

$$\lambda_{\min} \left(\Xi_{S_j}^\top \Xi_{S_j} \right) \geq \lambda_{\min} \left(\Xi_{S_j^*}^\top \Xi_{S_j^*} \right) - \lambda_{\max} \left(\Xi_{S_j^* \Delta S_j}^\top \Xi_{S_j^* \Delta S_j} \right).$$

Recall that by definition, we have

$$\begin{aligned} S_j^* \Delta S_j &= \{i : \langle \xi_i, \beta_j^* \rangle\} = \max \text{ and } \langle \xi_i, \beta_j \rangle \neq \max \} \cup \{i : \langle \xi_i, \beta_j^* \rangle \neq \max \text{ and } \langle \xi_i, \beta_j \rangle = \max \} \\ &\subseteq \bigcup_{j' \in [k] \setminus j} \{i : \langle \xi_i, v_{j,j'}^* \rangle \cdot \langle \xi_i, v_{j,j'} \rangle < 0\} \\ &= \bigcup_{j' \in [k] \setminus j} \{i : x_i \in W(v_{j,j'}^*, v_{j,j'})\}. \end{aligned} \quad (6.31)$$

Putting together the pieces, we have

$$\lambda_{\min} \left(\Xi_{S_j}^\top \Xi_{S_j} \right) \geq \lambda_{\min} \left(\Xi_{S_j^*}^\top \Xi_{S_j^*} \right) - \sum_{j' \neq j} \lambda_{\max} \left(\sum_{i: x_i \in W(v_{j,j'}^*, v_{j,j'})} \xi_i \xi_i^\top \right). \quad (6.32)$$

Now by Lemma 36, the definition of the set \mathcal{I}_b , and the definition of r_b , we have

$$\text{vol} \left(W(v_{j,j'}^*, v_{j,j'}) \right) \leq n r_b \log^{1/2}(1/r_b) \leq C \frac{\pi_{\min}^6}{k^2}.$$

Owing to the sample size assumption $n \geq C \max \left\{ d, k^2 \frac{\log n}{\pi_{\min}^6} \right\}$, the conditions of Lemma 32 are satisfied, and applying it yields

$$\sup_{v_{j,j'} \in \mathcal{B}_{v_{j,j'}^*}(r_a)} \lambda_{\max} \left(\sum_{i: x_i \in W(v_{j,j'}^*, v_{j,j'})} \xi_i \xi_i^\top \right) \leq nC \frac{\pi_{\min}^3}{k}$$

with probability exceeding $1 - n^{-10} - \exp \left(-cn \frac{\pi_{\min}^6}{k^2} \right)$. Moreover, Lemma 31 guarantees the bound $\lambda_{\min} \left(\Xi_{S_j^*}^\top \Xi_{S_j^*} \right) \geq c_2 n \cdot \pi_{\min}^3$, so that putting together the pieces, we have

$$\begin{aligned} \inf_{\beta_1, \dots, \beta_k \in \mathcal{I}_b} \lambda_{\min} \left(\Xi_{S_j}^\top \Xi_{S_j} \right) &\geq c_2 n \pi_{\min}^3 - Cnk \frac{\pi_{\min}^3}{k} \\ &\geq C\pi_{\min}^3 n, \end{aligned} \tag{6.33}$$

with probability greater than $1 - c \exp \left(-cn \frac{\pi_{\min}^6}{k^2} \right) - n^{-10}$. These assertions hold provided

$$n \geq C \max \left\{ d \cdot \frac{k}{\pi_{\min}^3}, \frac{k^2 \log n}{\pi_{\min}^6} \right\},$$

and this completes the proof. \square

Having proved the claims, we turn to proofs of our technical lemmas.

Proof of Lemma 29

In this proof, we assume that $\sigma = 1$; our bounds can finally be scaled by σ^2 .

It is natural to prove the bound (6.30b) first followed by bound (6.30a). First, consider a fixed set of parameters $\{\beta_1, \dots, \beta_k\}$. Then, we have

$$\|P_{\Xi^j(\beta_1, \dots, \beta_k)} z_{S^j}\|^2 = \|UU^\top z_{S^j}\|^2,$$

where $U \in \mathbb{R}^{|\Xi^j| \times (d+1)}$ denotes a matrix with orthonormal columns that span the range of $\Xi^j(\beta_1, \dots, \beta_k)$.

Applying the Hanson-Wright inequality for independent sub-Gaussians (see [112, Theorem 2.1]) and noting that $\|UU^\top\|_F \leq \sqrt{d+1}$ we obtain

$$\Pr \left\{ \|UU^\top z_{S^j}\|^2 \geq (d+1) + t \right\} \leq e^{-ct},$$

for each $t \geq 0$. In particular, this implies that the random variable $\|UU^\top z_{S^j}\|^2$ is sub-exponential.

This tail bound holds for a fixed partition of the rows of Ξ ; we now take a union bound over all possible partitions. Toward that end, define the sets

$$\mathcal{S}^j = \left\{ S_j(\beta_1, \dots, \beta_k) : \beta_1, \dots, \beta_k \in \mathbb{R}^{d+1} \right\}, \text{ for each } j \in [k].$$

From Lemma 38, we have the bound $|\mathcal{S}^j| \leq 2^{ckd \log(en/d)}$. Thus, applying the union bound, we obtain

$$\Pr \left\{ \sup_{\beta_1, \dots, \beta_k \in \mathbb{R}^{d+1}} \|P_{\Xi^j(\beta_1, \dots, \beta_k)} z_{S^j}\|^2 \geq (d+1) + t \right\} \leq |\mathcal{S}^j| e^{-ct},$$

and substituting $t = ck(d+1) \log(n/d)$ and performing some algebra establishes bound (6.30b).

In order to establish bound (6.30a), we once again consider the random variable $\sum_{j=1}^k \|P_{\Xi^j(\beta_1, \dots, \beta_k)} z_{S^j}\|^2$ for a fixed set of parameters $\{\beta_1, \dots, \beta_k\}$. Note that this is the sum of k independent sub-exponential random variables and can be thought of as a quadratic form of the entire vector z . So once again from the Hanson-Wright inequality, we have

$$\Pr \left\{ \sup_{\beta_1, \dots, \beta_k \in \mathbb{R}^{d+1}} \sum_{j=1}^k \|P_{\Xi^j(\beta_1, \dots, \beta_k)} z_{S^j}\|^2 \geq k(d+1) + t \right\} \leq e^{-ct/k}$$

for all $t \geq 0$.

Also define the set of all possible partitions of the n points via the max-affine function; we have the set

$$\mathcal{S} = \{S_1(\beta_1, \dots, \beta_k), \dots, S_k(\beta_1, \dots, \beta_k) : \beta_1, \dots, \beta_k \in \mathbb{R}^{d+1}\}.$$

Lemma 39 yields the bound $|\mathcal{S}| \leq 2^{ckd \log(kd) \log(n/kd)}$, and combining a union bound with the high probability bound above establishes bound (6.30a) after some algebraic manipulation. \square

Proof of Lemma 30

Let $\gamma_v = v - v^*$; we have

$$\begin{aligned} \mathbf{1} \{ \langle \xi_i, v \rangle \cdot \langle \xi_i, v^* \rangle \leq 0 \} \langle \xi_i, v^* \rangle^2 &\leq \mathbf{1} \{ \langle \xi_i, v \rangle \cdot \langle \xi_i, v^* \rangle \leq 0 \} \langle \xi_i, \gamma_v \rangle^2 \\ &\leq \mathbf{1} \{ \langle \xi_i, \gamma_v \rangle^2 \geq \langle \xi_i, v^* \rangle^2 \} \langle \xi_i, \gamma_v \rangle^2. \end{aligned}$$

Define the (random) set $K_v = \{i : \langle \xi_i, \gamma_v \rangle^2 > \langle \xi_i, v^* \rangle^2\}$; we have the bound

$$\frac{1}{n} \sum_{i=1}^n \mathbf{1} \{ \langle \xi_i, v \rangle \cdot \langle \xi_i, v^* \rangle \leq 0 \} \langle \xi_i, v^* \rangle^2 \leq \frac{1}{n} \|\Xi_{K_v} \gamma_v\|^2.$$

We now show that the quantity $\|\Xi_{K_v} \gamma_v\|^2$ is bounded uniformly for all $v \in \mathcal{B}_{v^*}(r)$ for small enough r . Recall that u^* is the “linear” portion of v^* , and let $m = \max\{d, 10 \log n, n \cdot (16r \cdot \sqrt{\log(1/r)})\}$

(note that m depends implicitly on r). We claim that for all $r \in (0, 1/24]$, we have

$$\Pr \left\{ \sup_{v \in \mathcal{B}_{v^*}(r)} |K_v| > m \right\} \leq 4e^{-c \max\{d, 10 \log n\}} + cne^{-c'n}, \text{ and} \quad (6.34a)$$

$$\Pr \left\{ \bigcup_{\substack{T \subseteq [n]: \\ |T| \leq m}} \sup_{\substack{\omega \in \mathbb{R}^{d+1} \\ \omega \neq 0}} \frac{\|\Xi_T \omega\|^2}{\|\omega\|^2} \geq (2d + 20m \log(n/m)) \right\} \leq e^{-c \max\{d, 10 \log n\}}. \quad (6.34b)$$

Taking these claims as given, the proof of the lemma is immediate, since $\frac{n}{m} \leq \frac{1}{16r \log(1/r)}$, so that $\log(n/m) \leq C \log(1/r)$.

Proof of claim (6.34a): By definition of the set K_v , we have

$$\begin{aligned} \Pr \left\{ \sup_{v \in \mathcal{B}_{v^*}(r)} |K_v| > m \right\} &\leq \sum_{\substack{T \subseteq [n]: \\ |T| > m}} \Pr \left\{ \exists v \in \mathcal{B}_{v^*}(r) : \|\Xi_T \gamma_v\|^2 \geq \|\Xi_T v^*\|^2 \right\} \\ &= \sum_{\substack{T \subseteq [n]: \\ |T| > m}} \Pr \left\{ \exists v \in \mathcal{B}_{v^*}(r) : \frac{\|\gamma_v\|^2 \|\Xi_T \gamma_v\|^2}{\|u^*\|^2 \|\gamma_v\|^2} \geq \frac{\|\Xi_T v^*\|^2}{\|u^*\|^2} \right\} \\ &\leq \sum_{\substack{T \subseteq [n]: \\ |T| > m}} \Pr \left\{ \exists v \in \mathcal{B}_{v^*}(r) : r^2 \frac{\|\Xi_T \gamma_v\|^2}{\|\gamma_v\|^2} \geq \frac{\|\Xi_T v^*\|^2}{\|u^*\|^2} \right\} \\ &\leq \sum_{\substack{T \subseteq [n]: \\ |T| > m}} \left(\Pr \left\{ \exists v \in \mathcal{B}_{v^*}(r) : \frac{\|\Xi_T \gamma_v\|^2}{\|\gamma_v\|^2} \geq (\sqrt{d} + \sqrt{|T|} + t_T)^2 \right\} \right. \\ &\quad \left. + \Pr \left\{ \frac{\|\Xi_T v^*\|^2}{\|u^*\|^2} \leq r^2 (\sqrt{d} + \sqrt{|T|} + t_T)^2 \right\} \right), \end{aligned}$$

where the final step follows by the union bound and holds for all positive scalars $\{t_T\}_{T \subseteq [n]}$. For some fixed subset T of size ℓ , we have the tail bounds

$$\Pr \left\{ \sup_{\substack{\omega \in \mathbb{R}^{d+1} \\ \omega \neq 0}} \frac{\|\Xi_T \omega\|^2}{\|\omega\|^2} \geq (\sqrt{d} + \sqrt{\ell} + t)^2 \right\} \stackrel{(i)}{\leq} 2e^{-t^2/2}, \text{ for all } t \geq 0, \text{ and} \quad (6.35a)$$

$$\Pr \left\{ \frac{\|\Xi_T v^*\|^2}{\|u^*\|^2} \leq \delta \ell \right\} \stackrel{(ii)}{\leq} (e\delta)^{\ell/2} \text{ for all } \delta \geq 0, \quad (6.35b)$$

where step (i) follows from the sub-Gaussianity of the covariate matrix (see Lemma 40), and step (ii) from a tail bound for the non-central χ^2 distribution (see Lemma 41).

Substituting these bounds yields

$$\begin{aligned} \Pr\left\{\sup_{v \in \mathcal{B}_{v^*}(r)} |K_v| > m\right\} &\leq \sum_{\ell=m+1}^n \binom{n}{\ell} \left[2e^{-t_\ell^2/2} + \left(er^2 \cdot \frac{(\sqrt{d} + \sqrt{\ell} + t_\ell)^2}{\ell} \right)^{\ell/2} \right] \\ &\leq \sum_{\ell=m+1}^n \binom{n}{\ell} \left[2e^{-t_\ell^2/2} + \left(2r \cdot \frac{\sqrt{d} + \sqrt{\ell} + t_\ell}{\sqrt{\ell}} \right)^\ell \right]. \end{aligned}$$

Recall that t_ℓ was a free (non-negative) variable to be chosen. We now split the proof into two cases and choose this parameter differently for the two cases.

Case 1, $m \leq \ell < n/e$: Substituting the choice $t_\ell = 4\sqrt{\ell \log(n/\ell)}$, we obtain

$$\begin{aligned} &\binom{n}{\ell} \left[2e^{-t_\ell^2/2} + \left(2r \cdot \frac{\sqrt{d} + \sqrt{\ell} + t_\ell}{\sqrt{\ell}} \right)^\ell \right] \\ &\leq \left(\frac{n}{\ell}\right)^{-c\ell} + \binom{n}{\ell} \cdot \left(2r \cdot \frac{\sqrt{d} + 5\sqrt{\ell \log(n/\ell)}}{\sqrt{\ell}} \right)^\ell \\ &\stackrel{(i)}{\leq} \left(\frac{n}{\ell}\right)^{-c\ell} + \binom{n}{\ell} \cdot \left(2r \cdot (1 + 5\sqrt{\log(n/\ell)}) \right)^\ell \\ &\stackrel{(ii)}{\leq} \left(\frac{n}{\ell}\right)^{-c\ell} + \binom{n}{\ell} \cdot \left(12r \cdot \sqrt{\log(n/\ell)} \right)^\ell \\ &\leq \left(\frac{n}{\ell}\right)^{-c\ell} + \left(12 \left(\frac{en}{\ell}\right) r \cdot \sqrt{\log(n/\ell)} \right)^\ell, \end{aligned}$$

where step (i) follows from the bound $m \geq d$, and step (ii) from the bound $\ell \leq n/e$.

Now note that the second term is only problematic for small ℓ . For all $\ell \geq m = n \cdot (16r \cdot \sqrt{\log(1/r)})$, we have

$$\left(12 \left(\frac{en}{\ell}\right) r \cdot \sqrt{\log(n/\ell)} \right)^\ell \leq (3/4)^\ell.$$

The first term, on the other hand, satisfies the bound $\left(\frac{n}{\ell}\right)^{-c\ell} \leq (3/4)^\ell$ for sufficiently large n .

Case 2, $\ell \geq n/e$: In this case, setting $t_\ell = 2\sqrt{n}$ for each ℓ yields the bound

$$\begin{aligned} \binom{n}{\ell} \left[2e^{-t_\ell^2/2} + \left(2r \cdot \frac{\sqrt{d} + \sqrt{\ell} + t_\ell}{\sqrt{\ell}} \right)^\ell \right] &\leq 2 \binom{n}{n/2} e^{-2n} + (12r)^\ell \\ &\leq ce^{-c'n}, \end{aligned}$$

where we have used the fact that $d \leq n/2$ and $r \leq 1/24$.

Putting together the pieces from both cases, we have shown that for all $r \in (0, 1/24]$, we have

$$\begin{aligned} \Pr\left\{ \sup_{v \in \mathcal{B}_{v^*}(r)} |K_v| > m \right\} &\leq cne^{-c'n} + \sum_{\ell=m+1}^{n/e} (3/4)^\ell \\ &\leq cne^{-c'n} + 4(3/4)^{\max\{d, 10 \log n\}}, \end{aligned}$$

thus completing the proof of the claim.

Proof of claim (6.34b): The proof of this claim follows immediately from the steps used to establish the previous claim. In particular, writing

$$\begin{aligned} &\Pr \left\{ \bigcup_{\substack{T \subseteq [n]: \omega: \|\omega\|=1 \\ |T| \leq m}} \|\Xi_T \omega\|^2 \geq 2d + 20m \log(n/m) \right\} \\ &\leq \Pr \left\{ \bigcup_{\substack{T \subseteq [n]: \omega: \|\omega\|=1 \\ |T| \leq m}} \|\Xi_T \omega\|^2 \geq \left(\sqrt{d} + \sqrt{m} + \sqrt{4m \log(n/m)} \right)^2 \right\} \\ &\leq \sum_{\ell=1}^m \Pr \left\{ \bigcup_{\substack{T \subseteq [n]: \omega: \|\omega\|=1 \\ |T|=\ell}} \|\Xi_T \omega\|^2 \geq \left(\sqrt{d} + \sqrt{m} + \sqrt{4m \log(n/m)} \right)^2 \right\} \\ &\stackrel{(iv)}{\leq} 2 \sum_{\ell=1}^m \binom{n}{\ell} \exp\{-2m \log(n/m)\} \\ &\leq 2 \left(\frac{n}{m} \right)^{-cm} \leq 2e^{-c \max\{d, 10 \log n\}}, \end{aligned}$$

where step (iv) follows from the tail bound (6.35a). \square

Proof of Lemma 31

The lemma follows from some structural results on the truncated Gaussian distribution. Using the shorthand $\mathfrak{vol} := \mathfrak{vol}(K)$ and letting ψ denote the d -dimensional Gaussian density, consider a random vector τ drawn from the distribution having density $h(y) = \frac{1}{\mathfrak{vol}} \psi(y) \mathbf{1}\{y \in K\}$, and denote its mean and second moment matrix by μ_τ and Σ_τ , respectively. Also denote the recentered random variable by $\tilde{\tau} = \tau - \mu_\tau$. We claim that

$$\|\mu_\tau\|^2 \leq C \log(1/\mathfrak{vol}), \quad (6.36a)$$

$$C \mathfrak{vol}^2 \cdot I \preceq \Sigma_\tau \preceq (1 + C \log(1/\mathfrak{vol})) I, \quad \text{and} \quad (6.36b)$$

$$\tilde{\tau} \text{ is } c\text{-sub-Gaussian for a universal constant } c. \quad (6.36c)$$

Taking these claims as given for the moment, let us prove the lemma.

The claims (6.36a) and (6.36c) taken together imply that the random variable τ is sub-Gaussian with parameter $\zeta^2 \leq 2c^2 + 2C \log(1/\text{vol})$. Now consider m i.i.d. draws of τ given by $\{\tau_i\}_{i=1}^m$; standard results (see, e.g., Vershynin [111, Remark 5.40], or Wainwright [13, Theorem 6.2]) yield the bound

$$\Pr \left\{ \left\| \frac{1}{m} \sum_{i=1}^m \tau_i \tau_i^\top - \Sigma_\tau \right\|_{\text{op}} \geq \zeta^2 \left(\frac{d}{m} + \sqrt{\frac{d}{m}} + \delta \right) \right\} \leq 2 \exp(-cn \min\{\delta, \delta^2\}).$$

Using this bound along with claim (6.36b) and Weyl's inequality yields

$$\lambda_{\min} \left(\frac{1}{m} \sum_{i=1}^m \tau_i \tau_i^\top \right) \geq C \text{vol}^2 - \zeta^2 \left(\frac{d}{m} + \sqrt{\frac{d}{m}} + \delta \right) \quad (6.37)$$

with probability greater than $1 - 2 \exp(-cn \min\{\delta, \delta^2\})$.

Furthermore, when n samples are drawn from a standard Gaussian distribution, the number m of them that fall in the set K satisfies $m \geq \frac{1}{2}n \cdot \text{vol}$ with high probability. In particular, this follows from a straightforward binomial tail bound, which yields

$$\Pr \left\{ m \leq \frac{n \cdot \text{vol}}{2} \right\} \leq \exp(-cn \cdot \text{vol}). \quad (6.38)$$

Recall our choice $n \geq Cd \frac{\log^2(1/\text{vol})}{\text{vol}^3}$, which in conjunction with the bound (6.38) ensures that $C \text{vol}^2 \geq \frac{1}{8}\sigma^2 \sqrt{\frac{d}{m}}$ with high probability. Setting $\delta = C \text{vol}^2 / \sigma^2$ in inequality (6.37), we have

$$\lambda_{\min} \left(\frac{1}{m} \sum_{i=1}^m \tau_i \tau_i^\top \right) \geq \frac{C}{2} \text{vol}^2$$

with probability greater than $1 - 2 \exp(-cn \text{vol}^4 / \sigma^4)$. Putting together the pieces thus proves the lemma. It remains to show the various claims. \square

Proof of claim (6.36a) Let $\tau_{\mathcal{A}}$ denote a random variable formed as a result of truncating the Gaussian distribution to a (general) set \mathcal{A} with volume vol . Letting $\mu_{\mathcal{A}}$ denote its mean, the dual norm definition of the ℓ_2 norm yields

$$\begin{aligned} \|\mu_{\mathcal{A}}\| &= \sup_{v \in \mathbb{S}^{d-1}} \langle v, \mu_{\mathcal{A}} \rangle \\ &\leq \sup_{v \in \mathbb{S}^{d-1}} \mathbb{E} |\langle v, \tau_{\mathcal{A}} \rangle|. \end{aligned}$$

Let us now evaluate an upper bound on the quantity $\mathbb{E} |\langle v, \tau_{\mathcal{A}} \rangle|$. In the calculation, for any d -dimensional vector y , we use the shorthand $y_v := v^\top y$ and $y_{\setminus v} := U_{\setminus v}^\top y$ for a matrix $U_{\setminus v} \in \mathbb{R}^{d \times (d-1)}$

having orthonormal columns that span the subspace orthogonal to v . Letting $\mathcal{A}_v \subseteq \mathbb{R}$ denote the projection of \mathcal{A} onto the direction v , define the set $\mathcal{A}_{\setminus v}(w) \subseteq \mathbb{R}^{d-1}$ via

$$\mathcal{A}_{\setminus v}(w) = \{y_{\setminus v} \in \mathbb{R}^{d-1} : y \in \mathcal{A} \text{ and } y_v = w\}.$$

Letting ψ_d denote the d -dimensional standard Gaussian pdf, we have

$$\begin{aligned} \mathbb{E}|\langle v, \tau_{\mathcal{A}} \rangle| &= \frac{1}{\text{vol}} \int_{y \in \mathcal{A}} |y^\top v| \psi_d(y) dy \\ &= \frac{1}{\text{vol}} \int_{y \in \mathcal{A}} |y_v| \psi(y_v) \psi_{d-1}(y_{\setminus v}) dy \\ &= \frac{1}{\text{vol}} \int_{y_v \in \mathcal{A}_v} |y_v| \psi(y_v) \underbrace{\left(\int_{y_{\setminus v} \in \mathcal{A}_{\setminus v}(y_v)} \psi_{d-1}(y_{\setminus v}) dy_{\setminus v} \right)}_{f(y_v)} dy_v \\ &\stackrel{(i)}{\leq} \frac{1}{\text{vol}} \int_{y_v \in \mathcal{A}_v} |y_v| \psi(y_v) dy_v, \end{aligned} \tag{6.39}$$

where step (i) follows since $f(y_v) \leq 1$ point-wise. On the other hand, we have

$$\text{vol} = \int_{y_v \in \mathcal{A}_v} \psi(y_v) \left(\int_{y_{\setminus v} \in \mathcal{A}_{\setminus v}(y_v)} \psi_{d-1} dy_{\setminus v} \right) dy_v \leq \int_{y_v \in \mathcal{A}_v} \psi(y_v) dy_v. \tag{6.40}$$

Combining inequalities (6.39) and (6.40) and letting $w = y_v$, an upper bound on $\|\mu_\tau\|$ can be obtained by solving the one-dimensional problem given by

$$\begin{aligned} \|\mu_\tau\| &\leq \sup_{\mathcal{S} \subseteq \mathbb{R}} \frac{1}{\text{vol}} \int_{w \in \mathcal{S}} |w| \psi(w) dw \\ &\text{s.t. } \int_{w \in \mathcal{S}} \psi(w) dw \geq \text{vol}. \end{aligned}$$

It can be verified that the optimal solution to the problem above is given by choosing the truncation set $\mathcal{S} = (\infty, -\beta) \cup [\beta, \infty)$ for some threshold $\beta > 0$. With this choice, the constraint can be written as

$$\text{vol} \leq \int_{|w| \geq \beta} \psi(w) dw \leq 2\sqrt{\frac{2}{\pi}} \frac{1}{\beta} e^{-\beta^2/2},$$

where we have used a standard Gaussian tail bound. Simplifying yields the bound

$$\beta \leq 2\sqrt{\log(C/\text{vol})}.$$

Furthermore, we have

$$\begin{aligned} \frac{1}{\mathbf{vol}} \int_{|w| \geq \beta} |w| \psi(w) dw &= \frac{C}{\mathbf{vol}} e^{-\beta^2/2} \\ &\stackrel{(ii)}{\lesssim} \frac{\beta^3}{\beta^2 - 1} \\ &\leq c \sqrt{\log(1/\mathbf{vol})}, \end{aligned}$$

where step (ii) follows from the bound $\Pr\{Z \geq z\} \geq \psi(z) \left(\frac{1}{z} - \frac{1}{z^3}\right)$ valid for a standard Gaussian variate Z . Putting together the pieces, we have

$$\|\mu_\tau\|^2 \leq c \log(1/\mathbf{vol}).$$

□

Proof of claim (6.36b) Let us first show the upper bound. Writing $\text{cov}(\tau)$ for the covariance matrix, we have

$$\begin{aligned} \|\Sigma_\tau\|_{\text{op}} &\leq \|\text{cov}(\tau)\|_{\text{op}} + \|\mu_\tau\|^2 \\ &\stackrel{(iii)}{\leq} \|I\|_{\text{op}} + C \log(1/\mathbf{vol}), \end{aligned}$$

where step (iii) follows from the fact that $\text{cov}(\tau) \preceq \text{cov}(Z)$, since truncating a Gaussian to a convex set reduces its variance along all directions [204, 205].

We now proceed to the lower bound. Let \mathbb{P}_K denote the Gaussian distribution truncated to the set K . Recall that we denoted the probability that a Gaussian random variable falls in the set K by $\mathbf{vol}(K)$; use the shorthand $\mathbf{vol} = \mathbf{vol}(K)$. Define the polynomial

$$p_u(x) = \langle x - \mathbb{E}_{X \sim \mathbb{P}_K}[X], u \rangle^2;$$

note that we are interested in a lower bound on $\inf_{u \in \mathbb{S}^{d-1}} \mathbb{E}_{X \sim \mathbb{P}_K}[p_u(X)]$.

For $\delta > 0$, define the set

$$S_\delta := \{x \in \mathbb{R}^d : p_u(x) \leq \delta\} \subseteq \mathbb{R}^d.$$

Letting Z denote a d -dimensional standard Gaussian random vector and using the shorthand $\alpha := \mathbb{E}_{X \sim \mathbb{P}_K}[X]$, we have

$$\Pr\{Z \in S_\delta\} = \Pr\{\langle Z - \alpha, u \rangle^2 \leq \delta\} \tag{6.41}$$

$$= \Pr\left\{\langle \alpha, u \rangle - \sqrt{\delta} \leq \langle Z, u \rangle \leq \langle \alpha, u \rangle + \sqrt{\delta}\right\} \tag{6.42}$$

$$= \int_{\langle \alpha, u \rangle - \sqrt{\delta}}^{\langle \alpha, u \rangle + \sqrt{\delta}} \psi(x) dx \leq \sqrt{\frac{2}{\pi}} \delta, \tag{6.43}$$

where in the final step, we have used the fact that $\psi(x) \leq 1/\sqrt{2\pi}$ for all $x \in \mathbb{R}$.

Consequently, we have

$$\begin{aligned}
 \mathbb{E}_{X \sim \mathbb{P}_K}[p_u(X)] &= \frac{1}{\text{vol}} \mathbb{E}_Z [p_u(Z) \mathbf{1}\{Z \in K\}] \\
 &\geq \frac{1}{\text{vol}} \mathbb{E}_Z [p_u(Z) \mathbf{1}\{Z \in K \cap S_\delta^c\}] \\
 &\stackrel{(iv)}{\geq} \frac{1}{\text{vol}} \mathbb{E}_Z [\delta \mathbf{1}\{Z \in K \cap S_\delta^c\}] \\
 &= \frac{\delta}{\text{vol}} \Pr\{Z \in K \cap S_\delta^c\} \\
 &\stackrel{(v)}{\geq} \delta \frac{\text{vol} - \sqrt{\frac{2}{\pi}}\delta}{\text{vol}}.
 \end{aligned}$$

Here, step (iv) follows from the definition of the set S_δ , which ensures that $p_u(x) \geq \delta$ for all $x \in S_\delta^c$. Step (v) follows as a consequence of equation (6.43), since

$$\Pr\{Z \in K \cap S_\delta^c\} = \Pr\{Z \in K\} - \Pr\{Z \in S_\delta\} \geq \text{vol} - \sqrt{\frac{2}{\pi}}\delta.$$

Finally, choosing $\delta = c \text{vol}^2$ for a suitably small constant c , we have $\mathbb{E}_{X \sim \mathbb{P}_K}[p_u(X)] \geq C \text{vol}^2$ for a fixed $u \in \mathbb{S}^{d-1}$. Since u was chosen arbitrarily, this proves the claim. \square

Proof of claim (6.36c) Since the random variable ξ is obtained by truncating a Gaussian random variable to a convex set, it is 1-strongly log-concave. Thus, standard results [206, Theorem 2.15] show that the random variable $\tilde{\xi}$ is c -sub-Gaussian. \square

Proof of Lemma 32

For a pair of $d + 1$ -dimensional vectors (v, v') , denote by

$$n_{W(v, v')} = \#\{i : x_i \in W(v, v')\} \tag{6.44}$$

the random variable that counts the number of points that fall within the wedge $W(v, v')$; recall our notation W_δ for the set of all wedges with Gaussian volume less than δ . Since each wedge is formed by the intersection of two hyperplanes, applying Lemmas 37 and 38 in conjunction yields that there are universal constants (c, c', C) such that

$$\sup_{W \in \mathcal{W}_\delta} n_W \leq c\delta n \tag{6.45}$$

with probability exceeding $1 - \exp(-c'n\delta^2)$, provided $n \geq \frac{C}{\delta^2}$. In words, the maximum number of points that fall in *any* wedge of volume δ is linear in δn with high probability.

It thus suffices to bound, simultaneously, the maximum singular value of every sub-matrix of Ξ having (at most) $c\delta n$ rows. Applying [181, Theorem 5.7] yields the bound¹⁰

$$\Pr \left\{ \max_{S: |S| \leq c\delta n} \lambda_{\max} \left(\sum_{i \in S} \xi_i \xi_i^\top \right) \geq c_1 n \sqrt{\delta} \right\} \leq n^{-10},$$

where we have used the lower bound $n \geq c \max\{d, \log n / \delta\}$ on the sample size.

Putting together the pieces, we have that if $n \geq c \max\{d, \frac{\log n}{\delta^2}\}$, then

$$\sup_{W \in \mathcal{W}_\delta} \lambda_{\max} \left(\sum_{i: x_i \in W} \xi_i \xi_i^\top \right) \leq c_1 n \sqrt{\delta}$$

with probability exceeding $1 - n^{-10} - \exp(-c'n\delta^2)$. \square

6.7 Proof of Theorem 13

We dedicate the first portion of the proof to a precise definition of the quantity γ .

Let $\Theta^* \in \mathbb{R}^{k \times d}$ denote a matrix with rows $(\theta_j^*)^\top, j = 1, \dots, k$ and let $\Sigma = \Theta^* (\Theta^*)^\top \in \mathbb{R}^{k \times k}$. We employ the decomposition $\Theta^* = A^* (U^*)^\top$, where $A^* \in \mathbb{R}^{k \times k}$ is the invertible matrix of coefficients and $U^* \in \mathbb{R}^{d \times k}$ is a matrix of orthonormal columns. Note that for $X \sim N(0, I_d)$, the vector in \mathbb{R}^k with j -th component $\langle X, \theta_j^* \rangle + b_j^*$ is distributed as $Z + b^*$ where $Z \sim N(0, \Sigma)$ and the vector $b^* \in \mathbb{R}^k$ collects the scalars $\{b_j^*\}_{j=1}^k$ in its entries. For $Z \sim \mathcal{N}(0, \Sigma)$, let

$$\rho = \frac{\mathbb{E} [\max(Z + b^*) Z^\top \Sigma^{-1} \mathbf{1}]}{\sqrt{\mathbb{E} [(\max(Z + b^*))^2] \cdot \mathbb{E} [(Z^\top \Sigma^{-1} \mathbf{1})^2]}} \quad (6.46)$$

denote the correlation coefficient between the maximum and a particular linear combination of a multivariate Gaussian distribution. Variants of such quantities have been studied extensively in the statistical literature (see, e.g., James [207]). For our purposes, the fact that $\max(Z + b^*) Z \neq 0$ for any finite b^* , coupled with a full-rank Σ , ensure that $\rho \neq 0$ for any fixed k . Also define the positive scalar $\varrho := \sqrt{\mathbb{E}[(\max(Z + b^*))^2]}$, which tracks the average size of our observations. Also recall the quantity ς defined in the main section.

For each $j \in [k]$ consider the zero-mean Gaussian random vector with covariance $(\mathbf{1} \cdot e_j^\top - I) A^* (A^*)^\top (\mathbf{1} \cdot e_j^\top - I)^\top$. This is effectively a Gaussian that lives in $k - 1$ dimensions, with density that we denote by $\tilde{\psi}_j(x_1, x_2, \dots, x_{j-1}, 0, x_{j+1}, \dots, x_k)$ at point $(x_1, x_2, \dots, x_{j-1}, 0, x_{j+1}, \dots, x_k)$ (the

¹⁰Strictly speaking, [181, Theorem 5.7] applies to Gaussian random matrices, i.e., without the appended column of ones. By multiplying each row of Ξ with an independent Rademacher RV (see the proof of Lemma 40) to obtain a sub-Gaussian random matrix with the same singular values, and noting also that the proof technique of [181, Theorem 5.7] relies on chaining and holds for a sub-Gaussian random matrix, one can show that the same result also holds for the matrix Ξ .

density is not defined elsewhere). Truncate this random vector to the region $\{x_i \geq b_i^* - b_j^* : i \in [k]\}$; this results in the *truncated* Gaussian density $\psi_j(x_1, x_2, \dots, x_{j-1}, 0, x_{j+1}, \dots, x_k)$ for each $j \in [k]$.

For any $x \in \mathbb{R}^k$ such that $x_j = 0$, define

$$F_i^j(x) = \int_{b_1^* - b_j^*}^{\infty} \dots \int_{b_{i-1}^* - b_j^*}^{\infty} \int_{b_{i+1}^* - b_j^*}^{\infty} \dots \int_{b_k^* - b_j^*}^{\infty} \psi_j(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_k) dx_k \dots dx_{i+1} dx_{i-1} \dots dx_1 \quad (6.47)$$

to be the i -th marginal density of this truncated Gaussian evaluated at the point x , with the convention that $F_j^j(\cdot) = 0$ everywhere. Also define the vector F^j by setting its i -th entry to $(F^j)_i = F_i^j(b_i^* - b_j^*)$.

Now let P denote the matrix with entries

$$P_{i,j} = \begin{cases} (F^j)_i / \sum_{k \neq j} (F^j)_k & \text{if } i \neq j \\ 0 & \text{otherwise.} \end{cases}$$

Note that the matrix P is the transition matrix of an irreducible, aperiodic Markov chain, with one eigenvalue equal to 1. Consequently, the matrix $I - P$ is rank $k - 1$. With this setup in place, let

$$\gamma := \min \left\{ \rho^2 \varrho^2, \min_{j \in [k]} \left(\sum_{k \neq j} (F^j)_k \right) \lambda_k(\Sigma) \cdot \sqrt{\lambda_{k-1}((I - P^\top)(I - P))} \right\} \quad (6.48)$$

denote a positive scalar that will serve as a bound on our eigengap.

Let $M_1 = \mathbb{E}[\max(\Theta^* X + b^*) X]$ and $M_2 = \mathbb{E}[\max(\Theta^* X + b^*)(X X^\top - I_d)]$ denote the expectations of the first and second moment estimators, respectively.

For a random variable $W \sim \mathcal{N}(b^*, \Sigma)$, we often use the shorthand

$$\{W_j = \max\} := \{W_j \geq W_i \text{ for all } 1 \leq i \leq k\}.$$

Finally, collect the probabilities $\{\pi_j\}_{j=1}^k$ defined in equation (6.9) in a vector $\pi \in \mathbb{R}^k$. We use $\mathbf{1}$ to denote the all-ones vector in k dimensions.

We are ready to state our two main lemmas.

Lemma 33. (a) *The first moment satisfies*

$$M_1 = (\Theta^*)^\top \pi \quad \text{and} \quad \langle M_1, (\Theta^*)^\top \Sigma^{-1} \mathbf{1} \rangle = \rho \varrho \left\| (\Theta^*)^\top \Sigma^{-1} \mathbf{1} \right\|.$$

(b) *The second moment satisfies*

$$M_2 \succeq 0, \quad M_2 (\Theta^*)^\top \Sigma^{-1} \mathbf{1} = 0, \quad \text{rank}(M_2) = k - 1 \quad \text{and}$$

$$\lambda_{k-1}(M_2) \geq \min_{j \in [k]} \left(\sum_{k \neq j} (F^j)_k \right) \lambda_k(\Sigma) \cdot \sqrt{\lambda_{k-1}((I - P^\top)(I - P))}.$$

We combine this lemma with a result that shows that the empirical moments concentrate about their expectations.

Lemma 34. *For an absolute constant C , we have*

$$\Pr \left\{ \left\| \widehat{M}_1 - M_1 \right\|^2 \geq C_1 (\sigma^2 + \varsigma^2) \frac{d \log^2(nk)}{n} \right\} \leq 5dn^{-12}, \text{ and} \quad (6.49a)$$

$$\Pr \left\{ \left\| \widehat{M}_2 - M_2 \right\|_{op}^2 \geq C (\sigma^2 + \varsigma^2) \frac{d \log^3(nk)}{n} \right\} \leq 5dn^{-12}. \quad (6.49b)$$

Lemma 33 is proved at the end of this section, and Lemma 34 is proved in Appendix 6.9.10. For now, we take both lemmas as given and proceed to a proof of Theorem 13.

Recall the matrix $\widehat{M} = \widehat{M}_1 \otimes \widehat{M}_1 + \widehat{M}_2$ and let $M = M_1 \otimes M_1 + M_2$. By Lemma 33, the matrix M is positive semidefinite with k non-zero eigenvalues. In particular, using the shorthand $\bar{\theta} := (\Theta^*)^\top \Sigma^{-1} \mathbf{1}$, we have $\bar{\theta} \in \text{nullspace}(M_2)$, and so

$$\bar{\theta}^\top M \bar{\theta} = \langle \bar{\theta}, M_1 \rangle^2 = \rho^2 \varrho^2 \|\bar{\theta}\|^2,$$

where the final inequality follows by part (a) of Lemma 33.

Thus, there is a k -dimensional subspace orthogonal to the nullspace of M (and so the range of M is k dimensional). For any unit vector v in this subspace, we have

$$v^\top M v \geq \min\{\rho^2 \varrho^2, \lambda_{k-1}(M_2)\}.$$

Thus, the k th eigenvalue of M satisfies

$$\lambda_k(M) \geq \min \left\{ \rho^2 \varrho^2, \min_{j \in k} \left(\sum_{k \neq j} (F^j)_k \right) \lambda_k(\Sigma) \cdot \sqrt{\lambda_{k-1}((I - P^\top)(I - P))} \right\} = \gamma,$$

where the equality follows by definition (6.48). By Lemma 34, we have

$$\begin{aligned} \left\| \widehat{M} - M \right\|_{op}^2 &\leq 2 \left\| \widehat{M}_2 - M_2 \right\|_{op}^2 + 2 \left\| \widehat{M}_1 \otimes \widehat{M}_1 - M_1 \otimes M_1 \right\|_{op}^2 \\ &\leq 2C (\sigma^2 + \varsigma^2 \log^2(nk)) \frac{d \log(nk)}{n} + 16 \left\| \widehat{M}_1 - M_1 \right\|^2 \|M_1\|^2 + 4 \left\| \widehat{M}_1 - M_1 \right\|^4 \\ &\leq C' (\sigma^2 + \varsigma^2 \log^2(nk)) \frac{d \log(nk)}{n}, \end{aligned}$$

where the last two inequalities each hold with probability greater than $1 - 2n^{-10}$.

We denote the estimated and true eigenspaces by \widehat{U} and U^* , respectively. Applying [208, Theorem 2] yields the bound

$$\|U^*(U^*)^\top - \widehat{U}\widehat{U}^\top\|_{\text{F}}^2 \leq C \left(\frac{\sigma^2 + \varsigma^2}{\gamma^2} \right) \frac{kd \log^3(nk)}{n},$$

thereby proving the required result. □

We now proceed to a proof of Lemma 33.

6.7.1 Proof of Lemma 33

Recall our decomposition $\Theta^* = A^*(U^*)^\top$, where $U^* \in \mathbb{R}^{d \times k}$ is a matrix of orthonormal columns, and $A^* \in \mathbb{R}^{k \times k}$ is an invertible matrix of coefficients. Since we are always concerned with random variables of the form Θ^*X with X Gaussian, we may assume without loss of generality by the rotation invariance of the Gaussian distribution that $U^* = [e_1^d \ e_2^d \ \dots \ e_k^d]$, where e_i^d denotes the i th standard basis vector in \mathbb{R}^d .

We let $X_i^j = (X_i, X_{i+1}, \dots, X_j)$ denote a sub-vector of the random vector X , so that by the above argument, we have $\Theta^*X \stackrel{d}{=} A^*X_1^k$.

Calculating M_1 : Using the shorthand $Z = A^*X_1^k$, we have

$$\begin{aligned} M_1 &= \mathbb{E}[\max(\Theta^*X + b^*)X] \\ &= U^*\mathbb{E}[\max(A^*X_1^k + b^*)X] \\ &= U^*(A^*)^{-1}\mathbb{E}[\max(Z + b^*)Z]. \end{aligned}$$

Now using Stein's lemma¹¹, by a calculation similar to the one performed also in Seigel [209] and Liu [210], we have

$$\mathbb{E}[\max(Z + b^*)Z] = \Sigma\pi,$$

where $\pi \in \mathbb{R}^k$ is the vector of probabilities, the j -th of which is given by equation (6.9), and we have used $\Sigma = A^*(A^*)^\top = (\Theta^*)(\Theta^*)^\top$ to denote the covariance matrix of Z .

Therefore, we have the first moment

$$M_1 = U^*(A^*)^{-1}A^*(A^*)^\top\pi = (\Theta^*)^\top\pi.$$

Correlation bound: By computation, we have

$$\begin{aligned} \langle M_1, (\Theta^*)^\top\Sigma^{-1}\mathbf{1} \rangle &= \mathbb{E}[\max(Z + b^*)\langle Z, \Sigma^{-1}\mathbf{1} \rangle] \stackrel{(i)}{=} \rho \cdot \sqrt{\mathbb{E}[(\max(Z + b^*))^2] \cdot \mathbb{E}[\langle Z, \Sigma^{-1}\mathbf{1} \rangle^2]} \\ &\stackrel{(ii)}{=} \rho\varrho \cdot \|(\Theta^*)^\top\Sigma^{-1}\mathbf{1}\|, \end{aligned}$$

where step (i) follows from the definition (6.46) of the quantity ρ , and step (ii) from explicitly calculating the expectation and recalling the definition of ϱ .

Positive semidefiniteness of M_2 : For some $u \in \mathbb{R}^d$, let $f(X) = \max(\Theta^*X + b)$ and $g_u(X) = \langle u, X \rangle^2$. Since g_u is an even function, we have $\mathbb{E}[g_u(X)X] = 0$. Furthermore, since both f and g_u are convex, applying Lemma 42 (see Appendix 6.9.9) yields the bound

$$\mathbb{E}[f(X)g_u(X)] \geq \mathbb{E}[f(X)]\mathbb{E}[g_u(X)],$$

¹¹One can also derive $M_1 = (\Theta^*)^\top\pi$ directly applying Stein's lemma $\mathbb{E}Xf(X) = \mathbb{E}\nabla f(X)$ to $f(x) := \max(\Theta^*X + b^*)$ so that $\nabla f(x)$ equals θ_j^* whenever x belongs to the region when j is maximized.

so that substituting yields the bound

$$u^\top \mathbb{E}[\max(\Theta^* X + b) X X^\top] u \geq u^\top \mathbb{E}[\max(\Theta^* X + b) I] u.$$

Since this holds for all $u \in \mathbb{R}^d$, we have shown that the matrix $\mathbb{E}[\max(\Theta^* X + b)(X X^\top - I)]$ is positive semidefinite.

Calculating M_2 : We now use Stein's lemma to compute an explicit expression for the moment M_2 . By the preceding substitution, we have

$$\begin{aligned} M_2 &= \mathbb{E} \left[\max(A^* X_1^k + b^*) \begin{bmatrix} X_1^k (X_1^k)^\top - I_k & X_1^k (X_{k+1}^d)^\top \\ X_{k+1}^d (X_1^k)^\top & X_{k+1}^d (X_{k+1}^d)^\top - I_{d-k} \end{bmatrix} \right] \\ &= \begin{bmatrix} \mathbb{E} [\max(A^* X_1^k + b^*) (X_1^k (X_1^k)^\top - I_k)] & 0 \\ 0 & 0 \end{bmatrix} \end{aligned}$$

Once again using the substitution $Z = A^* X_1^k$ and $\Sigma = A^* (A^*)^\top$, we have

$$M_2 = U^* (A^*)^{-1} \mathbb{E} [\max(Z + b^*) (Z Z^\top - \Sigma)] (A^*)^{-\top} (U^*)^\top,$$

and applying Stein's lemma yields

$$\mathbb{E} [\max(Z + b^*) (Z Z^\top - \Sigma)] = \Sigma \Pi^\top = \Pi \Sigma,$$

where $\Pi \in \mathbb{R}^{k \times k}$ denotes a matrix with entry i, j given by $\Pi_{i,j} = \mathbb{E}[Z_i \mathbf{1} \{Z_j + b_j^* = \max\}]$, and the final equality follows by symmetry of the matrix.

Simplifying further, we have

$$M_2 = U^* (A^*)^{-1} \Pi A^* (U^*)^\top.$$

Nullspace of M_2 : Notice that $\Pi \mathbf{1} = \mathbb{E}[Z] = 0$, so that

$$M_2 (\Theta^*)^\top \Sigma^{-1} \mathbf{1} = U^* (A^*)^{-1} \Pi A^* (U^*)^\top U^* (A^*)^\top \Sigma^{-1} \mathbf{1} = 0.$$

Rank of M_2 and bound on $\lambda_{k-1}(M_2)$: By the previous claim, we have $\text{rank}(M_2) \leq k - 1$. Furthermore, the matrix M_2 has $d - k$ eigenvalues equal to zero, and the other k of its eigenvalues equal to those of Π , all of which are positive (by the PSD property of M_2), and at least one of which is zero. Thus, it suffices to work with the eigenvalues of Π ; in particular, a lower bound on $\lambda_{k-1}(\Pi)$ directly implies a lower bound on $\lambda_{k-1}(M_2)$.

Let us first show that $\lambda_{k-1}(\Pi) > 0$. Since we know that a zero-eigenvector of Π is the all-ones vector $\mathbf{1}$, it suffices to show that $x^\top \Pi x \neq 0$ when $\langle x, \mathbf{1} \rangle = 0$. We use the shorthand $x \perp \mathbf{1}$ to denote any such vector.

We now explicitly evaluate the entries of the matrix Π . We denote the j th column of this matrix by Π_j . We have

$$\begin{aligned}\Pi_j &= \mathbb{E}[Z\mathbf{1} \{Z_j + b_j^* = \max\}] \\ &= \mathbb{E}[\mathbf{1} \cdot Z_j \mathbf{1} \{Z_j + b_j^* = \max\}] - \mathbb{E}[(\mathbf{1} \cdot Z_j - Z)\mathbf{1} \{Z_j + b_j^* = \max\}] \\ &= \mathbf{1} \cdot \mathbb{E}[Z_j \mathbf{1} \{Z_j + b_j^* = \max\}] - \mathbb{E}[(\mathbf{1} \cdot Z_j - Z)\mathbf{1} \{Z_j + b_j^* = \max\}].\end{aligned}\quad (6.50)$$

For any $x \perp \mathbf{1}$, we have $x^\top \mathbf{1} \mathbb{E}[Z\mathbf{1} \{Z + b^* = \max\}]^\top \mathbf{1} = 0$, so that in order to show that $x^\top \Pi x \neq 0$, it suffices to consider just the second term in the expression (6.50).

In order to focus on this term, consider the matrix Φ with column j given by

$$\Phi_j = \mathbb{E}[(\mathbf{1} \cdot Z_j - Z)\mathbf{1} \{Z_j - Z \geq b^* - b_j^*\}].$$

where the indicator random variable above is computed element-wise. We are interested in evaluating the eigenvalues of the matrix $-\Phi$.

The quantity Φ_j can be viewed as the first moment of a (lower) truncated, multivariate Gaussian with (original) covariance matrix

$$\kappa_j = (\mathbf{1} \cdot e_j^\top - I)A^*(A^*)^\top (\mathbf{1} \cdot e_j^\top - I)^\top.$$

Recalling the column vectors F^j defined (in equation (6.47)) for each $j \in [k]$ and applying [211, (11)] (see also Tallis [182] for a similar classical result), we may explicitly evaluate the vector Φ_j , as

$$\begin{aligned}\Phi_j &= \kappa_j F^j \\ &\stackrel{(iii)}{=} (\mathbf{1} \cdot e_j^\top - I)A^*(A^*)^\top G_j\end{aligned}$$

where in step (iii), we have let G_j denote a vector in \mathbb{R}^k with entry i given by

$$(G_j)_i = \begin{cases} -(F^j)_i & \text{if } j \neq i \\ \sum_{k \neq j} (F^j)_k & \text{otherwise.} \end{cases}$$

Letting $G \in \mathbb{R}^{k \times k}$ denote the matrix with G_j as its j th column, and for $x \perp \mathbf{1}$, we have

$$x^\top (-\Phi)x = x^\top \Sigma Gx,$$

since once again, for each $x \perp \mathbf{1}$, we have $x^\top \mathbf{1} \cdot e_j^\top A^*(A^*)^\top (\mathbf{1} \cdot e_j^\top - I)^\top x = 0$.

Now consider the matrix ΣG . In order to show the claimed bound, it suffices to show that $x^\top \Sigma Gx \neq 0$ if $x \perp \mathbf{1}$. We show this by combining two claims:

Claim 1: The nullspace of G is one-dimensional.

Claim 2: Both the left and right eigenvectors of ΣG that correspond to this nullspace are not orthogonal to the $\mathbf{1}$ vector.

We show both claims concurrently. The nullspace of G is clearly non-trivial, since $\mathbf{1}^\top G = 0$. Let us first show, by contradiction, that the left eigenvector corresponding to this nullspace dimension is

not orthogonal to the all-ones vector. Toward that x_ℓ denote the aforementioned left eigenvector which also satisfies $\langle x_\ell, \mathbf{1} \rangle = 0$. By virtue of being a left eigenvector, x_ℓ satisfies $\Sigma x_\ell = \mathbf{1}$, or in other words, $x_\ell = \Sigma^{-1} \mathbf{1}$. Since $x_\ell \perp \mathbf{1}$, we have $\mathbf{1}^\top \Sigma \mathbf{1} = 0$, but this contradicts the positive definiteness of Σ .

It remains to establish that the null-space of G is in fact only one-dimensional, and that its right eigenvector is not orthogonal to the all-ones vector. Notice that we may write the matrix as

$$G = (I - P^\top) \text{diag}(G),$$

where we recall that the matrix P is defined with entries

$$P_{i,j} = \begin{cases} (F^j)_i / \sum_{k \neq j} (F^j)_k & \text{if } i \neq j \\ 0 & \text{otherwise.} \end{cases}$$

Since all of the entries of P are positive and sum to 1 along the rows, the matrix P matrix can be viewed as the transition matrix of a Markov chain. Furthermore, since this Markov chain communicates, it is irreducible and aperiodic, with only one eigenvalue equal to 1. Thus, the matrix $I - P^\top$ is rank $k - 1$, thereby establishing that the nullspace of G is one-dimensional. Furthermore, the right eigenvector x_r of G is a non-negative vector by the Perron-Frobenius theorem, so that it cannot satisfy $\langle x_r, \mathbf{1} \rangle = 0$.

We have thus established both claims, which together show that $\lambda_{k-1}(M_2) \neq 0$. Further noting that the matrix M_2 is positive semi-definite, we have

$$\lambda_{k-1}(M_2) \geq \min_{j \in [k]} G_{j,j} \cdot \lambda_{\min}(\Sigma) \sqrt{\lambda_{k-1}[(I - P^\top)(I - P)]},$$

and this completes the proof of the claim, and consequently, the lemma. \square

6.8 Proof of Theorem 14

Recall the matrix \widehat{V} formed by appending a standard basis vector to \widehat{U} . First, we show that there is a point among the randomly chosen initializations that is sufficiently close to the true parameters. Toward that end, let $c_0 := r + B_{\max}$ and define $\beta_j^\ell = \widehat{V} \nu_j^\ell$ for each $j \in [k]$ and $\ell \in [M]$. Let

$$\ell^\# := \operatorname{argmin}_{\ell \in [M]} \left\{ \max_{j \in [k]} \|c_0 \beta_j^\ell - \beta_j^*\| \right\},$$

and define the event

$$\mathcal{E}_1(M, r) := \left\{ \max_{j \in [k]} \|c_0 \beta_j^{\ell^\#} - \beta_j^*\| \geq r + B_{\max} \|\widehat{U} \widehat{U}^\top - U^*(U^*)^\top\|_{\text{op}} \right\};$$

in words $\mathcal{E}_1(M, r)$ is the event that *none* of the randomly initialized points (when scaled by a fixed constant c_0) is close to the true parameters. The following lemma bounds the probability of such an event provided M is sufficiently large.

Lemma 35. *If $M \geq (1 + \frac{B_{\max}}{r})^{k^2} \log(1/\delta)$, then $\Pr \{ \mathcal{E}_1(M, r) \} \leq \delta$.*

Taking the lemma as given, let us now proceed to the proof of the theorem. Define the shorthand

$$\mathcal{P}(\beta_1, \dots, \beta_k) := \frac{2}{n} \sum_{i=n/2+1}^n \left(\max_{j \in [k]} \langle \xi_i, \beta_j \rangle - \max_{j \in [k]} \langle \xi_i, \beta_j^* \rangle \right)^2$$

for each set of parameters $\beta_1, \dots, \beta_k \in \mathbb{R}^{d+1}$.

For each $\ell \in [M]$, let

$$c_\ell := \operatorname{argmin}_{c \geq 0} \frac{2}{n} \sum_{i=n/2+1}^n \left(y_i - c \max_{j \in [k]} \langle \xi_i, \beta_j^\ell \rangle \right)^2,$$

and recall that ℓ^* is the index returned by the algorithm. Also note that trivially, we have $c_{\ell^*} > 0$ with probability tending to 1 exponentially in n , so that this pathological case in which the initial partition is random can be ignored.

Due to sample splitting, the parameters β_j^ℓ are independent of the noise sequence $\{\epsilon_i\}_{i=n/2+1}^n$. Thus, applying Lemma 44 from Appendix 6.9.13 yields the bound

$$\Pr \left\{ \mathcal{P}(c_{\ell^*} \beta_1^{\ell^*}, \dots, c_{\ell^*} \beta_k^{\ell^*}) \geq c_1 \left\{ \min_{\substack{c \geq 0 \\ \ell \in [M]}} \mathcal{P}(c \beta_1^\ell, \dots, c \beta_k^\ell) + \frac{\sigma^2 t (\sqrt{\log M} + c_1)}{n} \right\} \right\} \\ \leq e^{-c_2 n t (\sqrt{\log M} + c_1)},$$

valid for all $t \geq \sqrt{\log M} + c_1$ and suitable universal constants c_1 and c_2 . Setting $t = \sqrt{\log M} + c_1$ and on this event, we have

$$\mathcal{P}(c_{\ell^*} \beta_1^{\ell^*}, \dots, c_{\ell^*} \beta_k^{\ell^*}) \leq c_1 \mathcal{P}(c_0 \beta_1^{\ell^\#}, \dots, c_0 \beta_k^{\ell^\#}) + c_1 \frac{\sigma^2 \log M}{n}$$

with probability greater than $1 - e^{-c_2 n}$.

To complete the proof, let $C(\pi_{\min}, k) := c_2 \left(\frac{k}{\pi_{\min}} \right)^3$ for a suitable constant c_2 and apply

Lemma 43 twice (note that here we use the assumption $n \geq Cd \frac{k^3}{\pi_{\min}^3} \log^2(k/\pi_{\min})$) in order to obtain

$$\begin{aligned}
 \sum_{j \in [k]} \min_{j' \in [k]} \|\beta_j^* - c_{\ell^*} \beta_{j'}^{\ell^*}\|^2 &\leq C(\pi_{\min}, k) \cdot \mathcal{P}(c_{\ell^*} \beta_1^{\ell^*}, \dots, \beta_k^{\ell^*}) \\
 &\leq c_1 \cdot C(\pi_{\min}, k) \cdot \left\{ \mathcal{P}(c_0 \beta_1^{\ell^\#}, \dots, c_0 \beta_k^{\ell^\#}) + \frac{\sigma^2 \log M}{n} \right\} \\
 &\leq c_1 \cdot C(\pi_{\min}, k) \cdot \left\{ 2 \sum_{j=1}^k \|c_0 \beta_j^{\ell^\#} - \beta_j^*\|^2 + \frac{\sigma^2 \log M}{n} \right\} \\
 &\leq c_1 \cdot C(\pi_{\min}, k) \cdot \left\{ 2k \max_{j \in [k]} \|c_0 \beta_j^{\ell^\#} - \beta_j^*\|^2 + \frac{\sigma^2 \log M}{n} \right\} \\
 &\stackrel{(ii)}{\leq} c_1 \cdot C(\pi_{\min}, k) \left\{ 4k \left(r^2 + \mathbf{B}_{\max}^2 \|\widehat{U} \widehat{U}^\top - U^* (U^*)^\top\|_{\text{op}}^2 \right) + \frac{\sigma^2 \log M}{n} \right\}
 \end{aligned}$$

on an event of suitably high probability, where step (ii) follows from Lemma 35 and on the event $\mathcal{E}_1^c(M, r)$. Note that implicitly, we have also used a union bound over all M choices of our parameters, which leads to the overall probability bound of $1 - c_1 k M \exp\left(-c_2 \frac{\pi_{\min}^4}{k^4 \log^2(k/\pi_{\min})}\right)$.

Finally, note that provided the RHS above is less than $\Delta^2/4$, each minimum on the LHS is attained for a unique index j' . This condition is ensured by the sample size assumption of the theorem; thus, we have

$$\begin{aligned}
 \min_{c > 0} \text{dist} \left(\left\{ c \beta_j^{(0)} \right\}_{j=1}^k, \left\{ \beta_j^* \right\}_{j=1}^k \right) &\leq c_1 \cdot C(\pi_{\min}, k) \\
 &\quad \times \left\{ 4k \left(r^2 + \mathbf{B}_{\max}^2 \|\widehat{U} \widehat{U}^\top - U^* (U^*)^\top\|_{\text{op}}^2 \right) + \frac{\sigma^2 \log M}{n} \right\}.
 \end{aligned}$$

Combining the various probability bounds then completes the proof. \square

6.8.1 Proof of Lemma 35

Recall that U^* is a matrix of orthonormal columns spanning the k -dimensional subspace spanned by the vectors $\{\theta_1^*, \dots, \theta_k^*\}$. Define the matrix

$$V^* = \begin{bmatrix} U^* & 0 \\ 0 & 1 \end{bmatrix};$$

for each $j \in [k]$, we have $\beta_j^* = V^* \nu_j^*$ for some vector $\nu_j^* \in \mathbb{R}^{k+1}$. Also define the rotation matrix

$$O = \begin{bmatrix} \widehat{U}^\top U^* & 0 \\ 0 & 1 \end{bmatrix},$$

so that $\widehat{V}O - V^* = \begin{bmatrix} \widehat{U}\widehat{U}^\top U^* - U^* & 0 \\ 0 & 0 \end{bmatrix}$ and we have $\|\widehat{V}O - V^*\| = \|\widehat{U}\widehat{U}^\top - U^*(U^*)^\top\|$ for any unitarily invariant norm $\|\cdot\|$.

Now for each $j \in [k]$ and $\ell \in [M]$, applying the triangle inequality yields

$$\begin{aligned} \|c_0\beta_j^\ell - \beta_j^*\| &\leq \|c_0\widehat{V}O\nu_j^\ell - \widehat{V}O\nu_j^*\| + \|\widehat{V}O\nu_j^* - V^*\nu_j^*\| \\ &\leq \|c_0\nu_j^\ell - \nu_j^*\| + \|\nu_j^*\| \|\widehat{V}O - V^*\|_{\text{op}} \\ &\leq \|c_0\nu_j^\ell - \nu_j^*\| + \mathbf{B}_{\max} \|\widehat{U}\widehat{U}^\top - U^*(U^*)^\top\|_{\text{op}}. \end{aligned}$$

For each pair (ℓ, j) , define the event

$$\mathcal{E}_j^\ell(r) := \{\|c_0\nu_j^\ell - \nu_j^*\| \leq r\}.$$

We claim that if $M \geq \left(1 + \frac{\mathbf{B}_{\max}}{r}\right)^{k^2} \log(1/\delta)$, we have

$$\Pr \left\{ \bigcup_{\ell \in [M]} \bigcap_{j \in [k]} \mathcal{E}_j^\ell(r) \right\} \geq 1 - \delta. \quad (6.51)$$

Indeed, such a claim suffices, since it implies that

$$\min_{\ell \in [M]} \max_{j \in [k]} \|c_0\beta_j^\ell - \beta_j^*\| \leq r + \mathbf{B}_{\max} \|\widehat{U}\widehat{U}^\top - U^*(U^*)^\top\|_{\text{op}}$$

with probability exceeding $1 - \delta$, thereby proving the theorem. It remains to establish claim (6.51).

Denote by p the probability with which for a fixed pair (ℓ, j) , we have $\|c_0\nu_j^\ell - \nu_j^*\| \leq r$. This is the ratio of the volume of the ℓ_2 -ball of radius r and the ℓ_2 -ball of radius c_0 , and so we have $p = \left(\frac{r}{r + \mathbf{B}_{\max}}\right)^k$. Thus, we have

$$\begin{aligned} \Pr \left\{ \bigcap_{\ell \in [M]} \left(\bigcap_{j \in [k]} \mathcal{E}_j^\ell(r) \right)^c \right\} &\leq (1 - p^k)^M \\ &\leq e^{-p^k M} \\ &\stackrel{(i)}{\leq} \delta, \end{aligned}$$

where step (i) holds provided $M \geq \frac{1}{p^k} \log(1/\delta)$. Putting together the pieces completes the proof. \square

6.9 Technical Lemmas and Background

6.9.1 Fundamental limits

In this section, we present two lower bounds: one on the minimax risk of parameter estimation, and another on the risk of the least squares estimator with side-information.

6.9.2 Minimax lower bounds

Recall our notation Θ^* for the matrix whose columns consist of the parameters $\theta_1^*, \dots, \theta_k^*$. Assume that the intercepts b_1^*, \dots, b_k^* are identically zero, so that $\xi_i = x_i$ and $\Xi = X$. For a fixed matrix X , consider the observation model

$$y = \max(X\Theta^*) + \epsilon, \quad (6.52)$$

where $y \in \mathbb{R}^n$, the noise $\epsilon \sim \mathcal{N}(0, \sigma^2 I_n)$ is chosen independently of X , and the max function is computed row-wise.

Proposition 16. *There is an absolute constant C such that the minimax risk of estimation satisfies*

$$\inf_{\hat{\Theta}} \sup_{\Theta^* \in \mathbb{R}^{k \times d}} \mathbb{E} \left[\frac{1}{n} \|X(\hat{\Theta} - \Theta^*)\|_F^2 \right] \geq C \frac{\sigma^2 kd}{n}.$$

Here, the expectation is taken over the noise ϵ , and infimum is over all measurable functions of the observations (X, y) . Indeed, when X is a random Gaussian matrix, it is well conditioned and has singular values of the order \sqrt{n} , so that this bound immediately yields

$$\inf_{\hat{\Theta}} \sup_{\Theta^* \in \mathbb{R}^{k \times d}} \mathbb{E} \left[\frac{1}{n} \|\hat{\Theta} - \Theta^*\|_F^2 \right] \geq C \frac{\sigma^2 kd}{n}.$$

Let us now provide a proof of the proposition.

Proof. The proof is based on a standard application of Fano's inequality (see, e.g., Wainwright [13, Chapter 15] and Tsybakov [212, Chapter 2]). For a tolerance level $\delta > 0$ to be chosen, we choose the local set

$$F = \left\{ X\Theta \in \mathbb{R}^{n \times k} \mid \|X\Theta^\top\|_F \leq 4\delta\sqrt{kn} \right\}$$

and let $\{X\Theta^1, \dots, X\Theta^M\}$ be a $2\delta\sqrt{kn}$ -packing of the set in the Frobenius norm. This can be achieved by packing the j -th column $Q_j := \{X\theta_j \mid \|X\theta_j\|_2 \leq 4\delta\sqrt{n}\}$ at level $2\delta\sqrt{n}$ in ℓ_2 norm for all $j \in [k]$. Standard results yield the bound $\log M \geq C_1 \cdot kd \log 2$.

For each $i \neq j$, we have

$$2\delta\sqrt{k} \leq \frac{\|X(\Theta^i - \Theta^j)\|_F}{\sqrt{n}} \leq 8\delta\sqrt{k}. \quad (6.53)$$

Let $\mathbb{P}_j = \mathcal{N}(\max(X(\Theta^j)), \sigma^2 I_n)$ denote the distribution of the observation vector y when the true parameter is Θ^j . We thus obtain

$$D_{\text{KL}}(\mathbb{P}_j \parallel \mathbb{P}_i) = \frac{1}{2\sigma^2} \left\| \max(X(\Theta^j)) - \max(X(\Theta^i)) \right\|_2^2 \leq \frac{1}{2\sigma^2} \|X(\Theta^j - \Theta^i)\|_F^2,$$

where the inequality follows since the max function is 1-Lipschitz in ℓ_2 norm. Putting together the pieces yields

$$D_{\text{KL}}(\mathbb{P}_j \parallel \mathbb{P}_i) \leq \frac{32k\delta^2 n}{\sigma^2},$$

so that the condition

$$\frac{\frac{1}{M^2} \sum_{i,j} D_{\text{KL}}(\mathbb{P}_{\Theta^j} \parallel \mathbb{P}_{\Theta^k}) + \log 2}{\log M} \leq \frac{1}{2}$$

is satisfied with the choice $\delta^2 = C \frac{\sigma^2 d}{n}$. Finally, applying Fano's inequality (see, e.g., [13, Proposition 15.2]) yields the minimax lower bound

$$\inf_{\hat{\Theta}} \sup_{\Theta^*} \mathbb{E} \left[\frac{1}{n} \|X(\hat{\Theta} - \Theta^*)\|_{\text{F}}^2 \right] \geq C \frac{\sigma^2 k d}{n}. \quad (6.54)$$

□

6.9.3 Performance of unconstrained least squares with side-information

In this section, we perform an explicit computation when $k = 3$ and $d = 2$ to illustrate the cubic π_{\min} dependence of the error incurred by the unconstrained least squares estimator, even when provided access to the true partition $\{S_j(\beta_1^*, \dots, \beta_3^*)\}_{j=1}^3$.

We begin by defining our unknown parameters. For a scalar $\alpha \in (0, \pi/4)$, let

$$\theta_1^* = \sin(\alpha) \cdot e_1, \quad \theta_2^* = \cos(\alpha) \cdot e_2, \quad \text{and } \theta_3^* = -\cos(\alpha) \cdot e_2,$$

and set $b_j^* = 0$ for $j = 1, 2, 3$.

Now an explicit computation yields that the cone on which θ_1^* attains the maximum is given by

$$\mathcal{C}_1 := \left\{ x \in \mathbb{R}^2 : \langle x, \theta_1^* \rangle \geq \max_{j \in [k]} \langle x, \theta_j^* \rangle \right\} = \{x \in \mathbb{R}^2 : x_1 \geq 0, |x_2| \leq x_1 \tan(\alpha)\}.$$

Now consider a Gaussian random vector in \mathbb{R}^2 truncated to that cone. In particular, consider a two-dimensional random variable W with density $\psi(x) \mathbf{1}\{x \in \mathcal{C}_1\} / \text{vol}(\mathcal{C}_1)$, where ψ is the two-dimensional standard Gaussian density and $\text{vol}(S)$ denotes the Gaussian volume of a set S . Note that we have $\text{vol}(\mathcal{C}_1) = \alpha/\pi$ by construction.

Let us now compute the second order statistics of W , using polar coordinates with R^2 denoting a χ_2^2 random variable. The individual second moments take the form

$$\mathbb{E}[W_1^2] = \frac{\pi}{\alpha} \mathbb{E}[R^2] \left(\frac{1}{2\pi} \int_{-\alpha}^{\alpha} \cos^2 \phi d\phi \right) = 1,$$

and

$$\mathbb{E}[W_2^2] = \frac{\pi}{\alpha} \mathbb{E}[R^2] \left(\frac{1}{2\pi} \int_{-\alpha}^{\alpha} \sin^2 \phi d\phi \right) = \frac{1}{\alpha} (\alpha - \sin(2\alpha)/2) \sim \alpha^2.$$

On the other hand, the cross terms are given by

$$\mathbb{E}[W_1 W_2] = \frac{\pi}{\alpha} \mathbb{E}[R^2] \left(\frac{1}{2\pi} \int_{-\alpha}^{\alpha} \sin(\phi) \cos(\phi) d\phi \right) = 0.$$

Thus, it can be verified that for all $\alpha \in [0, \pi/4]$, the second moment matrix of W has a tuple of singular values $(1, c\alpha^2)$ for an absolute constant c .

Let us now use this calculation to reason about the least squares estimator. Drawing n samples from the Gaussian distribution on \mathbb{R}^2 , we expect $n_1 \sim \frac{\alpha}{\pi} n$ of them to fall in the set \mathcal{C}_1 with high probability. Collect these samples as rows of a matrix X_1 . When n is large enough, i.e., on the order of α^{-3} , standard bounds (as in Section 6.6.1) can be applied to explicitly evaluate the singular values of the matrix $\frac{1}{n_1} X_1^\top X_1$. In particular, we have

$$\lambda_1 \left(\frac{1}{n_1} X_1^\top X_1 \right) = c' \quad \text{and} \quad \lambda_2 \left(\frac{1}{n_1} X_1^\top X_1 \right) = c\alpha^2.$$

We now provide the $n_1 \times 2$ matrix X_1 as side information to a procedure whose goal is to estimate the unknown parameters. Clearly, given this matrix, a natural procedure to run in order to estimate θ_1^* is the (unconstrained) least squares estimator on these samples, which we denote by $\hat{\theta}_1$. As is well known, the rate obtained (in the fixed design setting) by this estimator with σ -sub-Gaussian noise is given by

$$\begin{aligned} \mathbb{E} \left[\|\hat{\theta}_1 - \theta_1^*\|^2 \right] &= \sigma^2 \text{tr}(X_1^\top X_1)^{-1} \\ &= \sigma^2 \frac{1}{n_1} (c\alpha^{-2} + c') \\ &\sim \sigma^2 \frac{1}{\alpha^3 n}, \end{aligned}$$

where the last two relations hold with exponentially high probability in n . We have thus shown that the unconstrained least squares estimator (even when provided with additional side information) attains an error having cubic dependence on $\alpha \sim \pi_{\min}$. While this does not constitute an information theoretic lower bound, our calculation provides some evidence for the fact that, at least when viewed in isolation, the dependence of our statistical error bound (6.15) on π_{\min} is optimal for Gaussian covariates.

6.9.4 Background and technical lemmas used in the proof of Theorem 12

In this section, we collect statements and proofs of some technical lemmas used in the proofs of our results concerning the AM algorithm.

6.9.5 Bounds on the “volumes” of wedges in \mathbb{R}^d

For a pair of scalars (w, w') and d -dimensional vectors (u, u') , recall that we define the *wedge* formed by the $d + 1$ -dimensional vectors $v = (u, w)$ and $v' = (u', w')$ as the region

$$W(v, v') = \{x \in \mathbb{R}^d : (\langle x, u \rangle + w) \cdot (\langle x, u' \rangle + w') \leq 0\}.$$

Note that the wedge is a purely geometric object.

For any set $\mathcal{C} \subseteq \mathbb{R}^d$, let

$$\text{vol}(\mathcal{C}) = \Pr_{X \sim \mathcal{N}(0, I_d)} \{X \in \mathcal{C}\}$$

denote the volume of the set under the measure corresponding to the covariate distribution.

We now bound the volume of a wedge for the Gaussian distribution.

Lemma 36. *Suppose that for a pair of scalars (w, w') , d -dimensional vectors (u, u') , and $v = (u, w)$ and $v' = (u', w')$, we have $\frac{\|v - v'\|}{\|u\|} < 1/2$. Then, there is a positive constant C such that*

$$\text{vol}(W(v, v')) \leq C \frac{\|v - v'\|}{\|u\|} \log^{1/2} \left(\frac{2\|u\|}{\|v - v'\|} \right).$$

Proof of Lemma 36

Using the notation $\xi = (x, 1) \in \mathbb{R}^{d+1}$ to denote the appended covariate, we have

$$\text{vol}(W(v, v')) = \Pr \{ \langle \xi, v \rangle \cdot \langle \xi, v' \rangle \leq 0 \},$$

where the probability is computed with respect to Gaussian measure.

In order to prove a bound on this probability, we begin by bounding the associated indicator random variable as

$$\begin{aligned} \mathbf{1} \{ \langle \xi, v \rangle \cdot \langle \xi, v' \rangle \leq 0 \} &\leq \mathbf{1} \{ \langle \xi, v' - v \rangle^2 \geq \langle \xi, v \rangle^2 \} \\ &\leq \mathbf{1} \{ \langle \xi, v' - v \rangle^2 \geq t \} + \mathbf{1} \{ \langle \xi, v \rangle^2 \leq t \}, \end{aligned} \quad (6.55)$$

where inequality (6.55) holds for all $t \geq 0$. In order to bound the expectation of the second term, we write

$$\begin{aligned} \Pr \{ \langle \xi, v \rangle^2 \leq t \} &= \Pr \{ \|u\|^2 \chi_{nc}^2 \leq t \} \\ &\stackrel{(i)}{\leq} \left(\frac{et}{\|u\|^2} \right)^{1/2} \end{aligned}$$

where χ_{nc}^2 is a non-central chi-square random variable centered at $\frac{w}{\|u\|}$, and step (i) follows from standard χ^2 tail bounds (see Lemma 41).

It remains to control the expectation of the first term on the RHS of inequality (6.55). We have

$$\begin{aligned} \Pr \{ \langle \xi, v' - v \rangle^2 \geq t \} &\leq \Pr \{ 2 \langle x, u' - u \rangle^2 + 2(w' - w)^2 \geq t \} \\ &\leq \Pr \left\{ \|u - u'\|^2 \chi^2 \geq \frac{t}{2} - \|v - v'\|^2 \right\}. \end{aligned}$$

Now, invoking a standard sub-exponential tail bound on the upper tail of a χ^2 random variable yields

$$\begin{aligned} \Pr \{ \langle \xi, v' - v \rangle^2 \geq t \} &\leq c_1 \exp \left(- \frac{c_2}{\|u - u'\|^2} \left\{ \frac{t}{2} - \|v - v'\|^2 \right\} \right) \\ &\leq c_1 \exp \left(- \frac{c_2}{\|v - v'\|^2} \left\{ \frac{t}{2} - \|v - v'\|^2 \right\} \right). \end{aligned}$$

Putting all the pieces together, we obtain

$$\text{vol}(W(v, v')) \leq c_1 \exp \left(- \frac{c_2}{\|v - v'\|^2} \left\{ \frac{t}{2} - \|v - v'\|^2 \right\} \right) + \left(\frac{et}{\|u\|^2} \right)^{1/2}.$$

Substituting $t = 2c \|v - v'\|^2 \log(2\|u\|/\|v - v'\|)$, which is a valid choice provided $\frac{\|v - v'\|}{\|u\|} < 1/2$, yields the desired result. \square

6.9.6 Growth Functions and Uniform Empirical Concentration

We now briefly introduce growth functions and uniform laws derived from them, and refer the interested reader to Mohri et al. [213] for a more in-depth exposition on these topics.

We define growth functions in the general multi-class setting [189]. Let \mathcal{X} denote a set, and let \mathcal{F} denote a family of functions mapping $\mathcal{X} \mapsto \{0, 1, \dots, k-1\}$. The *growth function* $\Pi_{\mathcal{F}} : \mathbb{N} \rightarrow \mathbb{R}$ of \mathcal{F} is defined via

$$\Pi_{\mathcal{F}}(n) := \max_{x_1, \dots, x_n \in \mathcal{X}} \text{abs} \{ \{f(x_1), f(x_2), \dots, f(x_n)\} : f \in \mathcal{F} \}.$$

In words, it is the cardinality of all possible labelings of n points in the set \mathcal{X} by functions in the family \mathcal{F} .

A widely studied special case arises in the case $k = 2$, with the class of binary functions. In this case, a natural function class \mathcal{F} is formed by defining \mathcal{C} to be a family of subsets of \mathcal{X} , and identifying each set $C \in \mathcal{C}$ with its indicator function $f_C := 1_C : \mathcal{X} \rightarrow \{0, 1\}$. In this case, define $\mathcal{F}_{\mathcal{C}} = \{f_C : C \in \mathcal{C}\}$. A bound on the growth function for such binary function provides following guarantee for the uniform convergence for the empirical measures of sets belonging to \mathcal{C} .

Lemma 37 (Theorem 2 in [214]). *Let \mathcal{C} be a family of subsets of a set \mathcal{X} . Let μ be a probability measure on \mathcal{X} , and let $\hat{\mu}_m := \frac{1}{m} \sum_{i=1}^m \delta_{X_i}$ be the empirical measure obtained from m independent copies of a random variable X with distribution μ . For every u such that $m \geq 2/u^2$, we have*

$$\Pr \left\{ \sup_{C \in \mathcal{C}} \text{abs} \hat{\mu}_m(C) - \sigma(C) \geq u \right\} \leq 4\Pi_{\mathcal{F}_{\mathcal{C}}}(2m) \exp(-mu^2/16). \quad (6.56)$$

We conclude this section by collecting some results on the growth functions of various function classes. For our development, it will be specialized to the case $\mathcal{X} = \mathbb{R}^d$.

Define the class of binary functions $\mathcal{F}_{\mathcal{H}}$ as the set of all functions of the form

$$f_{\theta,b}(x) := \frac{\text{sgn}(\langle x, \theta \rangle + b) + 1}{2};$$

specifically, let $\mathcal{F}_{\mathcal{H}} := \{f_{\theta,b} : \theta \in \mathbb{R}^d, b \in \mathbb{R}\}$. In particular, these are all functions that can be formed by a d -dimensional hyperplane.

Using the shorthand $B_1^k = \{B_1, \dots, B_k\}$, define the binary function

$$g_{\theta_1^k, b_1^k}(x) := \prod_{i=1}^k f_{\theta_i, b_i}(x),$$

and the binary function class corresponding to the intersection of k hyperplanes

$$\mathcal{G}_{\mathcal{H}^k} := \left\{ g_{\theta_1^k, b_1^k} : \theta_1, \dots, \theta_k \in \mathbb{R}^d, b_1, \dots, b_k \in \mathbb{R} \right\}.$$

Finally, we are interested in the argmax function over hyperplanes. Here, define the function

$$m_{\theta_1^k, b_1^k}(x) := \operatorname{argmax}_{j \in [k]} (\langle \theta_j, x \rangle + b_j) - 1,$$

mapping $\mathbb{R}^d \mapsto \{0, \dots, k-1\}$. The function class that collects all such functions is given by

$$\mathcal{M}_k := \left\{ m_{\theta_1^k, b_1^k} : \theta_1, \dots, \theta_k \in \mathbb{R}^d, b_1, \dots, b_k \in \mathbb{R} \right\}.$$

The following results bound the growth functions of each of these function classes. We first consider the function classes $\mathcal{F}_{\mathcal{H}}$ and $\mathcal{G}_{\mathcal{H}^k}$, for which bounds on the VC dimension directly yield bounds on the growth function.

Lemma 38 (Sauer-Shelah (e.g. Section 3 of Mohri et al. [213])). *We have*

$$\Pi_{\mathcal{F}_{\mathcal{H}}}(n) \leq \left(\frac{en}{d+1} \right)^{d+1}, \text{ and} \tag{6.57}$$

$$\Pi_{\mathcal{G}_{\mathcal{H}^k}}(n) \leq \left(\frac{en}{d+1} \right)^{k(d+1)}. \tag{6.58}$$

The second bound can be improved (see, e.g. [215]), but we state the version obtained by a trivial composition of individual halfspaces.

The following bound on the growth function of the class \mathcal{M}_k is also known.

Lemma 39 (Theorem 3.1 of Daniely et al. [189]). *For an absolute constant C , we have*

$$\Pi_{\mathcal{M}_k}(n) \leq \left(\frac{en}{Ck(d+1)\log(kd)} \right)^{Ck(d+1)\log(kd)}.$$

6.9.7 Singular value bound

We now state and prove a technical lemma that bound the maximum singular value of a matrix whose rows are drawn from a sub-Gaussian distribution.

Lemma 40. *Suppose that the covariates are drawn i.i.d. from a η -sub-Gaussian distribution. Then for a fixed subset $S \in [n]$ of size ℓ and each $t \geq 0$, we have*

$$\Pr \left\{ \lambda_{\max} (\Xi_S^\top \Xi_S) \geq \ell + \tilde{\eta}^2 (\sqrt{\ell d} + d + \ell t) \right\} \leq 2e^{-\ell \min\{t, t^2\}},$$

where $\tilde{\eta} = \max\{\eta, 1\}$.

Proof of Lemma 40

Let $\{z_i\}_{i=1}^\ell$ denote i.i.d. Rademacher variables, and collect these in an ℓ -dimensional vector z . Let $D = \text{diag}(z)$ denote a diagonal matrix, and note that by unitary invariance of the singular values, the singular values of the matrix $\tilde{\Xi}_S = D\Xi_S$ are the same as those of Ξ_S .

By construction, the matrix $\tilde{\Xi}_S$ has i.i.d. rows, and the i -th row is given by $z_i(x_i, 1)$. For a $d+1$ dimensional vector $\tilde{\lambda} = (\lambda, w)$ with $\lambda \in \mathbb{R}^d$ and $w \in \mathbb{R}$, we have

$$\begin{aligned} \mathbb{E} \left[\exp(\langle \tilde{\lambda}, z_i(x_i, 1) \rangle) \right] &= \frac{e^w}{2} \cdot \mathbb{E} [\exp(\langle \lambda, x_i \rangle)] + \frac{e^{-w}}{2} \cdot \mathbb{E} [\exp(-\langle \lambda, x_i \rangle)] \\ &= \exp(\|\lambda\|^2 \eta^2 / 2) \cdot \frac{1}{2} (e^w + e^{-w}) \\ &\leq \exp(\|\lambda\|^2 \eta^2 / 2) \cdot \exp(w^2 / 2) \leq \exp(\|\tilde{\lambda}\|^2 \tilde{\eta}^2 / 2). \end{aligned}$$

where we have used the fact that x_i is zero-mean and η sub-Gaussian.

Since the rows of $\tilde{\Xi}_S$ are i.i.d., zero-mean, and $\tilde{\eta}$ -sub-Gaussian, applying [13, Theorem 6.2] immediately yields the lemma. \square

6.9.8 Anti-concentration of χ^2 random variable

The following lemma shows the anti-concentration of the central and non-central χ^2 random variable.

Lemma 41. *Let Z_ℓ and Z'_ℓ denote central and non-central χ^2 random variables with ℓ degrees of freedom, respectively. Then for all $p \in [0, \ell]$, we have*

$$\Pr\{Z'_\ell \leq p\} \leq \Pr\{Z_\ell \leq p\} \leq \left(\frac{p}{\ell} \exp\left(1 - \frac{p}{\ell}\right) \right)^{\ell/2} = \exp\left(-\frac{\ell}{2} \left[\log \frac{\ell}{p} + \frac{p}{\ell} - 1 \right]\right) \quad (6.59)$$

Proof of Lemma 41

The fact that $Z'_\ell \stackrel{st.}{\leq} Z_\ell$ follows from standard results that guarantee that central χ^2 random variables stochastically dominate their non-central counterparts.

The tail bound is a simple consequence of the Chernoff bound. In particular, we have for all $\lambda > 0$ that

$$\begin{aligned} \Pr\{Z_\ell \leq p\} &= \Pr\{\exp(-\lambda Z_\ell) \geq \exp(-\lambda p)\} \\ &\leq \exp(\lambda p) \mathbb{E}[\exp(-\lambda Z_\ell)] \\ &= \exp(\lambda p)(1 + 2\lambda)^{-\frac{\ell}{2}}. \end{aligned} \tag{6.60}$$

where in the last step, we have used $\mathbb{E}[\exp(-\lambda Z_\ell)] = (1 + 2\lambda)^{-\frac{\ell}{2}}$, which is valid for all $\lambda > -1/2$. Minimizing the last expression over $\lambda > 0$ then yields the choice $\lambda^* = \frac{1}{2} \left(\frac{\ell}{p} - 1 \right)$, which is greater than 0 for all $0 \leq p \leq \ell$. Substituting this choice back into equation (6.60) proves the lemma. \square

6.9.9 Background and technical lemmas used in the proof of Theorem 13

We begin by stating a result of Hargé [216, Theorem 1.2] (see also Hu [217]) that guarantees that convex functions of a Gaussian random vector are positively correlated. We state it below in the notation of the current paper.

Lemma 42 ([216]). *Let f and g be two convex functions on \mathbb{R}^d , and let X be a standard d -dimensional Gaussian vector. Then*

$$\mathbb{E}[f(X)g(X)] \geq (1 + \langle m(g), m(f) \rangle) \mathbb{E}[f(X)] \mathbb{E}[g(X)], \tag{6.61}$$

where for any d -variate function h , we have $m(h) = \frac{\mathbb{E}[Xh(X)]}{\mathbb{E}[h(X)]}$.

We also prove Lemma 34, which was used in the proof of Theorem 13.

6.9.10 Proof of Lemma 34

We prove each bound separately. First, by the rotation invariance of the Gaussian distribution, we may assume that $U^* = [e_1^d \dots e_k^d]$, so that the max is computed as a function of the k coordinates X_1, \dots, X_k .

We also define some events that we make use of repeatedly in the proofs. For each $i \in [n]$, define the events

$$\begin{aligned} \mathcal{E}_i &= \{|x_{i,j}| \leq 5\sqrt{\log(2nk)} \text{ for all } 1 \leq j \leq k\}, \text{ and} \\ \mathcal{F}_i &= \{|\epsilon_i| \leq 5\sigma\sqrt{\log(2n)}\}. \end{aligned}$$

Note that by standard sub-Gaussian tail bounds, we have $\Pr\{\mathcal{E}_i^c\} \leq 2n^{-12}$ and $\Pr\{\mathcal{F}_i^c\} \leq 2n^{-12}$ for each $i \in [n]$. For notational convenience, define for each i the modified covariate $z_i = x_i \cdot \mathbf{1}\{\mathcal{E}_i\}$.

We have

$$|\max(\Theta^* z_i + b^*)| \leq C \max_{j \in [k]} \|\theta_j^*\|_1 \sqrt{\log(nk)} + |b_j^*| \leq \left(C \sqrt{\log(nk)} \right) \varsigma$$

almost surely, where in the second bound, we have used the shorthand $\varsigma = \max_j (\|\theta_j^*\|_1 + \|b_j^*\|_1)$ as defined in equation (6.18). With this setup in place, we are now ready to prove both deviation bounds.

Proof of bound (6.49a)

Let us first bound the deviation of the first moment. We work with the decomposition

$$\widehat{M}_1 - M_1 = \frac{2}{n} \sum_{i=1}^{n/2} \underbrace{\max(\Theta^* x_i + b^*) x_i - \mathbb{E}[\max(\Theta^* X + b^*) X]}_{T_i^1} + \frac{2}{n} \sum_{i=1}^{n/2} \underbrace{\epsilon_i x_i}_{T_i^2}.$$

By triangle inequality, it suffices to bound the norms of each of the two sums separately. We now use the further decomposition

$$\begin{aligned} T_i^1 &= \underbrace{\max(\Theta^* x_i + b^*) x_i - \max(\Theta^* z_i + b^*) z_i}_{P_i} + \underbrace{\max(\Theta^* z_i + b^*) z_i - \mathbb{E}[\max(\Theta^* z_i + b^*) z_i]}_{Q_i} \\ &\quad + \underbrace{\mathbb{E}[\max(\Theta^* z_i + b^*) z_i] - \mathbb{E}[\max(\Theta^* x_i + b^*) x_i]}_{R_i}. \end{aligned}$$

Since $z_i = x_i$ with probability greater than $1 - 2n^{-12}$, the term $P_i = 0$ on this event.

Also, for each fixed $j \in [k]$, applying the Hoeffding inequality yields the bound

$$\Pr \left\{ \left| \frac{2}{n} \sum_{i=1}^{n/2} Q_{i,j} \right| \geq t \right\} \leq 2 \exp \left\{ - \frac{nt^2}{8C^2\varsigma^2(\log(nk))^2} \right\}.$$

On the other hand, for $j \in [d] \setminus [k]$, we have

$$\begin{aligned} \left| \frac{2}{n} \sum_{i=1}^{n/2} Q_{i,j} \right| &\leq \varsigma \frac{2}{n} \sum_{i=1}^{n/2} z_{i,j} \\ &= \varsigma \left| \frac{2}{n} \sum_{i=1}^{n/2} x_{i,j} \right|. \end{aligned}$$

Standard Gaussian tail bounds then yield

$$\Pr \left\{ \left| \frac{2}{n} \sum_{i=1}^{n/2} Q_{i,j} \right| \geq st \sqrt{\log(nk)} \right\} \leq 2 \exp \left\{ - \frac{nt^2}{8} \right\}$$

for each $t \geq 0$. Putting together the pieces with a union bound and choosing constants appropriately, we then have

$$\Pr \left\{ \left\| \frac{2}{n} \sum_{i=1}^{n/2} Q_i \right\|^2 \geq \frac{1}{n} \cdot Ck\varsigma^2 (\log(nk))^2 + \frac{1}{n} \cdot C'(d-k)\varsigma^2 \log(nk) \right\} \leq 2dn^{-12}.$$

It remains to handle the final terms $\{R_i\}_{i=1}^n$. Note that when $j \notin [k]$, we have $R_{i,j} = 0$. It therefore suffices to bound the various $R_{i,j}$ terms when $j \in [k]$. We have

$$\begin{aligned} |R_{i,j}| &= |\mathbb{E}[\max(\Theta^* z_i + b^*) z_{i,j}] - \mathbb{E}[\max(\Theta^* x_i + b^*) x_{i,j} \mathbf{1}\{\mathcal{E}_i\}] - \mathbb{E}[\max(\Theta^* x_i + b^*) x_{i,j} \mathbf{1}\{\mathcal{E}_i^c\}]| \\ &= |\mathbb{E}[\max(\Theta^* x_i + b^*) x_{i,j} \mathbf{1}\{\mathcal{E}_i^c\}]| \end{aligned}$$

Expanding this further, we have

$$\begin{aligned} |R_{i,j}| &\leq \mathbb{E}[\max_{\ell \in [k]} (|\langle \theta_\ell^*, x_i \rangle| + |b_\ell^*|) |x_{i,j}| \mathbf{1}\{\mathcal{E}_i^c\}] \\ &\leq \mathbb{E}[|x_{i,j}| \|x_i\|_\infty (\|\Theta^*\|_{1,\infty} + \|b^*\|_\infty) \mathbf{1}\{\mathcal{E}_i^c\}] \\ &= \varsigma \mathbb{E}[|x_{i,j}| \|x_i\|_\infty \mathbf{1}\{\mathcal{E}_i^c\}] \\ &\leq \varsigma \sum_{\ell=1}^k \mathbb{E}[|x_{i,j}| |x_{i,\ell}| \mathbf{1}\{\mathcal{E}_i^c\}]. \end{aligned}$$

Note that for a pair (X_1, X_2) of i.i.d. random variables, Jensen's inequality yields the bounds

$$\begin{aligned} \mathbb{E}[|X_1 X_2| \mathbf{1}\{X_1, X_2 \geq \lambda\}] &\leq \mathbb{E}[X_1^2 \mathbf{1}\{|X_1| \geq \lambda\}] \text{ for all } \lambda \geq 0, \text{ and} \\ \mathbb{E}[|X_1| \mathbf{1}\{|X_1| \geq \lambda\}] &\leq \mathbb{E}[X_1^2 \mathbf{1}\{|X_1| \geq \lambda\}] \text{ for all } \lambda \geq 1. \end{aligned}$$

Furthermore, if X is a standard Gaussian random variable, then a simple calculation (see also Burkardt [218]) yields the bound

$$\mathbb{E}[X^2 \mid |X| \geq \lambda] \leq \frac{1}{2\sqrt{2\pi}} \lambda e^{-\lambda^2/2}, \text{ for all } \lambda \geq \sqrt{2}.$$

Putting together the pieces with with $\lambda = 5\sqrt{\log(2nk)}$, we have

$$|R_{i,j}|^2 \leq Ck^2\varsigma^2 \log(nk)(nk)^{-24},$$

and summing over $j \in [k]$ yields the bound

$$\left\| \frac{2}{n} \sum_{i=1}^{n/2} R_i \right\|^2 \leq Ck^2\varsigma^2 \log(nk)(nk)^{-24}.$$

Finally, putting together the pieces with a union bound yields the desired bound on the random variable $\left\| \frac{2}{n} \sum_{i=1}^{n/2} T_i^1 \right\|^2$.

The second term can be bounded more easily; in particular, on the intersection of the events $\{\mathcal{F}_i\}_{i=1}^n$, we have

$$\left\| \frac{2}{n} \sum_{i=1}^{n/2} T_i^2 \right\|^2 \leq C\sigma^2 \log n \left\| \frac{2}{n} \sum_{i=1}^{n/2} x_i \right\|^2 \leq C\sigma^2 \frac{(d + \log n) \log n}{n},$$

where the final bound holds with probability greater than $1 - cn^{-10}$. Finally, putting the bounds together yields the result.

Proof of bound (6.49b)

Once again, we decompose the required term as

$$\widehat{M}_2 - M_2 = \frac{2}{n} \sum_{i=1}^{n/2} \underbrace{\max(\Theta^* x_i + b^*) (x_i x_i^\top - I_d)}_{\tau_i^1} + \frac{2}{n} \sum_{i=1}^{n/2} \underbrace{\epsilon_i (x_i x_i^\top - I_d)}_{\tau_i^2}.$$

We use the further decomposition

$$\begin{aligned} \tau_i^1 &= \underbrace{\max(\Theta^* x_i + b^*) (x_i x_i^\top - I_d) - \max(\Theta^* z_i + b^*) (z_i z_i^\top - I_d)}_{\phi_i} \\ &\quad + \underbrace{\max(\Theta^* z_i + b^*) (z_i z_i^\top - I_d) - \mathbb{E}[\max(\Theta^* z_i + b^*) (z_i z_i^\top - I_d)]}_{\kappa_i} \\ &\quad + \underbrace{\mathbb{E}[\max(\Theta^* z_i + b^*) (z_i z_i^\top - I_d)] - \mathbb{E}[\max(\Theta^* x_i + b^*) (x_i x_i^\top - I_d)]}_{\rho_i}. \end{aligned}$$

As before, since $z_i = x_i$ with probability greater than $1 - 2n^{-12}$, the term $\phi_i = 0$ on this event.

Let us further decompose κ_i as

$$\begin{aligned} \kappa_i &= \underbrace{\left(\max(\Theta^* z_i + b^*) + \varsigma \sqrt{\log(nk)} \right) z_i z_i^\top - \mathbb{E} \left[\left(\max(\Theta^* z_i + b^*) + \varsigma \sqrt{\log(nk)} \right) z_i z_i^\top \right]}_{\kappa_i^{(1)}} \\ &\quad + \underbrace{\varsigma \sqrt{\log(nk)} \mathbb{E} [z_i z_i^\top] - I_d + I_d}_{\kappa_i^{(2)}} \cdot \underbrace{\left(\mathbb{E}[\max(\Theta^* z_i + b^*)] - \max(\Theta^* z_i + b^*) \right)}_{\kappa_i^{(3)}}, \end{aligned}$$

so that

$$\left\| \frac{2}{n} \sum_{i=1}^n \kappa_i \right\|_{\text{op}} \leq \left\| \frac{2}{n} \sum_{i=1}^n \kappa_i^{(1)} \right\|_{\text{op}} + \left\| \frac{2}{n} \sum_{i=1}^n \kappa_i^{(2)} \right\|_{\text{op}} + \left| \frac{2}{n} \sum_{i=1}^{n/2} \kappa_i^{(3)} \right|.$$

Since $|\max(\Theta^* z_i + b^*)| \leq C\varsigma\sqrt{\log(nk)}$, the random vector $\sqrt{\max(\Theta^* z_i + b^*) + C\varsigma\sqrt{\log(nk)}}z_i$ is well-defined and bounded; sub-Gaussian concentration bounds [13] can therefore be applied to obtain

$$\mathbb{P}\left[\left\|\frac{1}{n}\sum_{i=1}^n \kappa_i^{(1)}\right\|_{\text{op}} \geq c_1\varsigma^2(\log(nk))^2 \left\{\sqrt{\frac{d}{n}} + \frac{d}{n} + \delta\right\}\right] \leq c_2 \exp(-n \min(\delta, \delta^2))$$

where $c_1 \log(nk) = \max(\Theta^* z_i + b^*) + \varsigma\sqrt{\log(nk)} \leq 2\varsigma \log(nk)$. Reasoning similarly for the second term, we have

$$\mathbb{P}\left[\left\|\frac{1}{n}\sum_{i=1}^n \kappa_i^{(2)}\right\|_{\text{op}} \geq c_1\varsigma^2(\log(nk))^2 \left\{\sqrt{\frac{d}{n}} + \frac{d}{n} + \delta\right\}\right] \leq c_2 \exp(-n \min(\delta, \delta^2)).$$

Combining these bounds setting $\delta = c_1\sqrt{\frac{d}{n}}$, we have

$$\left\|\frac{2}{n}\sum_{i=1}^n \kappa_i^{(1)}\right\|_{\text{op}} + \left\|\frac{2}{n}\sum_{i=1}^n \kappa_i^{(2)}\right\|_{\text{op}} \leq C\varsigma^2(\log(nk))^2 \left\{\sqrt{\frac{d}{n}} + \frac{d}{n}\right\}$$

with probability at least $1 - c \exp(-c'd)$.

The term $\kappa_i^{(3)}$, on the other hand, can be controlled directly via Hoeffding's inequality. Since $\max(\Theta^* z_i + b^*)$ is $C\varsigma\sqrt{\log(nk)}$ sub-Gaussian, we obtain

$$\mathbb{P}\left[\left|\frac{2}{n}\sum_{i=1}^{n/2} \kappa_i^{(3)}\right| \geq \varsigma\sqrt{\log(nk)}t\right] \leq 2 \exp\left\{-\frac{nt^2}{32}\right\}.$$

Choosing $t = c\sqrt{\frac{d+\log n}{n}}$ and putting together all the pieces, we obtain

$$\left\|\frac{2}{n}\sum_{i=1}^n \kappa_i\right\|_{\text{op}} \leq C\varsigma^2(\log(nk))^2 \left\{\sqrt{\frac{d+\log n}{n}} + \frac{d+\log n}{n}\right\} + c\varsigma\sqrt{\log(nk)}\sqrt{\frac{d}{n}}$$

with probability at least $1 - cn^{-12}$.

It remains to handle the terms $\{\rho_i\}_{i=1}^{n/2}$, and to do so, we use a similar argument to before. We first bound the absolute value of the (p, q) th entry of each matrix as

$$\begin{aligned} |\rho_i(p, q)| &= |\mathbb{E}[\max(\Theta^* z_i + b^*)z_i z_i^\top(p, q)] - \mathbb{E}[\max(\Theta^* x_i + b^*)x_i x_i^\top(p, q)\mathbf{1}\{\mathcal{E}_i\}] \\ &\quad + \mathbb{E}[\max(\Theta^* x_i + b^*)x_i x_i^\top(p, q)\mathbf{1}\{\mathcal{E}_i^c\}]| = |\mathbb{E}[\max(\Theta^* x_i + b^*)x_i x_i^\top(p, q)\mathbf{1}\{\mathcal{E}_i^c\}]| \end{aligned}$$

Expanding this further, we have

$$\begin{aligned} |\rho_i(p, q)| &\leq \mathbb{E}[\max_{\ell \in [k]}(|\theta_\ell^* \cdot x_i| + |b_\ell^*|)|x_{i,p}x_{i,q}|\mathbf{1}\{\mathcal{E}_i^c\}] \\ &\leq \varsigma\mathbb{E}[|x_{i,p}x_{i,q}|\|x_i\|_\infty\mathbf{1}\{\mathcal{E}_i^c\}] \\ &\leq \mathbb{E}\left[|x_{i,p}x_{i,q}|\sum_{\ell \in [k]}|x_{i,\ell}|\mathbf{1}\{\mathcal{E}_i^c\}\right]. \end{aligned}$$

Also note that $\rho_{p,q} = 0$ unless $p \in [k], q \in [k]$. Hence we finally need to control the terms of the form $\mathbb{E}[|X|^3 \mathbf{1}\{|X| \geq \lambda\}]$ for a standard Gaussian X . Substituting $\lambda = 5\sqrt{\log(nk)}$, a simple calculation of truncated third moment of standard Gaussian ([218]) yields

$$|\rho_i(p, q)| \leq \varsigma \log^2(nk)(nk)^{-10},$$

and proceeding as before provides a strictly lower order bound on $\|\rho_i\|_{\text{op}}$ than the remaining terms.

The term τ_i^2 can be bounded more easily. Specifically, on the intersection of the events $\{\mathcal{F}_i\}_{i=1}^{n/2}$, applying [13, Lemma 6.2], we have

$$\left\| \frac{2}{n} \sum_{i=1}^{n/2} \tau_i^2 \right\|_{\text{op}}^2 \leq C\sigma^2 \log n \left\| \frac{2}{n} \sum_{i=1}^{n/2} x_i x_i^\top - I \right\|_{\text{op}}^2 \leq C\sigma^2 \log n \left\{ \frac{d + \log n}{n} + \frac{(d + \log n)^2}{n^2} \right\}$$

where the final bound holds with probability greater than $1 - cn^{-12}$. Finally combining all the terms yield the desired result. \square

6.9.11 Background and technical lemmas used in the proof of Theorem 14

In this section, we collect two technical lemmas that were used to prove Theorem 14.

6.9.12 Prediction and estimation error

Here, we connect the prediction error to the estimation error, which may be of independent interest. Recall our notation dist for the minimum distance between parameters obtainable after relabeling.

Lemma 43. *There exists a tuple of universal constants (c_1, c_2) such that for each set of parameters $\beta_1, \dots, \beta_k \in \mathbb{R}^{d+1}$:*

1. *If $n \geq c_1 d$, then we have*

$$\frac{1}{n} \sum_{i=1}^n \left(\max_{j \in [k]} \langle \xi_i, \beta_j \rangle - \max_{j \in [k]} \langle \xi_i, \beta_j^* \rangle \right)^2 \leq c_1 \text{dist}(\{\beta_j\}_{j=1}^k, \{\beta_j^*\}_{j=1}^k)$$

with probability exceeding $1 - c_1 \exp(-c_2 n)$.

2. *If $n \geq c_1 d \frac{k^3}{\pi_{\min}} \log^2(k/\pi_{\min})$, then we have*

$$c_2 \left(\frac{\pi_{\min}}{k} \right)^3 \sum_{j \in [k]} \min_{j' \in [k]} \|\beta_j^* - \beta_{j'}\|^2 \leq \frac{1}{n} \sum_{i=1}^n \left(\max_{j \in [k]} \langle \xi_i, \beta_j \rangle - \max_{j \in [k]} \langle \xi_i, \beta_j^* \rangle \right)^2$$

with probability exceeding $1 - c_1 k \exp\left(-c_2 n \frac{\pi_{\min}^4}{k^4 \log^2(k/\pi_{\min})}\right)$.

Proof. To prove the part 1 of the lemma, we leverage the fact that the max function is 1-Lipschitz with respect to the ℓ_2 -norm. Consequently, we obtain

$$\frac{1}{n} \sum_{i=1}^n \left(\max_{j \in [k]} \langle \xi_i, \beta_j \rangle - \max_{j \in [k]} \langle \xi_i, \beta_j^* \rangle \right)^2 \leq \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^k (\xi_i^\top (\beta_j - \beta_j^*))^2,$$

where we have ordered the parameters such that $\text{dist} \left(\{\beta_j\}_{j=1}^k, \{\beta_j^*\}_{j=1}^k \right)$ is minimized. We now use the fact that the rows of Ξ are 1-sub-Gaussian (this is restatement of the conclusion of Lemma 40) to complete the proof.

We now proceed to a proof of part 2 of the lemma. Recall the setup of Appendix 6.7 along with notation $(\{x_i\}_{i=1}^n, \Theta^*, b^*, \beta^*)$. Specifically, we have $\beta_j^* = (\theta_j^*, b_j^*)$ and $(\Theta^*)^\top = [\theta_1^* \ \theta_2^* \ \dots \ \theta_k^*]$. Similarly let $\beta_j = (\theta_j, b_j) \in \mathbb{R}^{d+1}$ and $\Theta^\top = [\theta_1 \ \theta_2 \ \dots \ \theta_k]$. In the notation of Section 6.6, we define for each pair (Θ, b) , the sets

$$S_j(\Theta, b) = \left\{ i \in [n] : \langle x_i, \theta_j \rangle + b_j = \max_{j' \in [k]} (\langle x_i, \theta_{j'} \rangle + b_{j'}) \right\}, \quad j \in [k].$$

We use the shorthand $S_j^* = S_j(\Theta^*, b^*)$ and $\widehat{S}_j = S_j(\Theta, b)$ for the rest of the proof. By definition, we have

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n (\max(\Theta x_i + b) - \max(\Theta^* x_i + b^*))^2 \\ &= \frac{1}{n} \sum_{\substack{\ell \in [k] \\ m \in [k]}} \sum_{i \in S_\ell^* \cap \widehat{S}_m} \left((\langle \theta_\ell^*, x_i \rangle + b_\ell^*) - (\langle \theta_m, x_i \rangle + b_m) \right)^2 \\ &= \frac{1}{n} \sum_{\substack{\ell \in [k] \\ m \in [k]}} \sum_{i \in S_\ell^* \cap \widehat{S}_m} \left(\langle \beta_\ell^*, \xi_i \rangle - \langle \beta_m, \xi_i \rangle \right)^2 \\ &= \frac{1}{n} \sum_{\substack{\ell \in [k] \\ m \in [k]}} \|\widetilde{\Xi}_{\ell, m}(\beta_\ell^* - \beta_m)\|^2, \end{aligned}$$

where we have let $\widetilde{\Xi}_{\ell, m}$ denote the sub-matrix of Ξ with rows indexed by the set $S_\ell^* \cap \widehat{S}_m$. It is also useful to define the convex sets

$$\begin{aligned} K_\ell^* &:= \left\{ x \in \mathbb{R}^d : \langle x, \theta_\ell^* \rangle + b_\ell^* = \max_{j' \in [k]} (\langle x, \theta_{j'}^* \rangle + b_{j'}^*) \right\}, \quad \text{and} \\ K_m &:= \left\{ x : \langle x, \theta_m \rangle + b_m = \max_{j' \in [k]} (\langle x, \theta_{j'} \rangle + b_{j'}) \right\} \end{aligned}$$

for each pair $(\ell, m) \in k \times k$. By definition, for each $\ell \in [k]$, there exists a corresponding index m_ℓ such that $\text{vol}(K_\ell^* \cap K_{m_\ell}) \geq \frac{\pi_{\min}}{k}$. Proceeding from above, we have

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n (\max(\Theta x_i + b) - \max(\Theta^* x_i + b^*))^2 &\geq \frac{1}{n} \sum_{\ell \in [k]} \left\| \tilde{\Xi}_{\ell, m_\ell} (\beta_\ell^* - \beta_{m_\ell}) \right\|^2 \\ &\geq \frac{1}{n} \sum_{\ell \in [k]} \lambda_{\min} \left(\tilde{\Xi}_{\ell, m_\ell}^\top \tilde{\Xi}_{\ell, m_\ell} \right) \|\beta_\ell^* - \beta_{m_\ell}\|^2. \end{aligned}$$

Finally, applying Lemma 31 in conjunction with the bound $\text{vol}(K_\ell^* \cap K_{m_\ell}) \geq \frac{\pi_{\min}}{k}$, we obtain that provided $n \geq c_1 d \cdot \frac{k^3}{\pi_{\min}^3} \log^2(k/\pi_{\min})$, we have

$$\lambda_{\min} \left(\tilde{\Xi}_{\ell, m_\ell}^\top \tilde{\Xi}_{\ell, m_\ell} \right) \geq \left(\frac{\pi_{\min}}{k} \right)^3 \cdot n$$

with probability exceeding $1 - c_1 \exp \left(-c_2 n \frac{\pi_{\min}^4}{k^4 \log^2(k/\pi_{\min})} \right)$ for each index $\ell \in [k]$. Taking a union bound over the k indices and combining the pieces completes the proof. \square

6.9.13 Projection onto a finite collection of rays

Consider a vector $\theta^* \in \mathbb{R}^n$ observed via the observation model

$$y = \theta^* + \epsilon,$$

where ϵ has independent, zero-mean, σ -sub-Gaussian entries. For a fixed set of M vectors $\{\theta_1, \dots, \theta_M\}$, denote by $\mathbb{C} := \{c\theta_\ell : c \geq 0, \ell \in [M]\}$ the set of all one-sided rays obtainable with these vectors.

Now consider the projection estimate

$$P_{\mathbb{C}}(y) = \operatorname{argmin}_{\theta \in \mathbb{C}} \|y - \theta\|^2,$$

which exists since the projection onto each ray exists. The following lemma proves an oracle inequality on the error of such an estimate.

Lemma 44. *There are universal constants c, C, c_1 and c_2 such that*

$$\Pr \left\{ \|P_{\mathbb{C}}(y) - \theta^*\|^2 \geq c \left(\min_{\theta \in \mathbb{C}} \|\theta - \theta^*\|^2 + \sigma^2 t (\log M + c_1) \right) \right\} \leq c_2 e^{-nt(\sqrt{\log M} + c_1)},$$

for all $t \geq C\sigma(\sqrt{\log M} + c_1)$.

Proof. We follow the standard technique for bounding the error for non-parametric least squares estimators. From the definition, we have

$$P_{\mathbb{C}}(y) = \operatorname{argmin}_{\theta \in \mathbb{C}} \|y - \theta\|^2.$$

We substitute the expression for y and obtain

$$P_{\mathbb{C}}(y) = \operatorname{argmax}_{\theta \in \mathbb{C}} [2\langle \epsilon, \theta - \theta^* \rangle - \|\theta - \theta^*\|^2].$$

To obtain an upper bound on $\|P_{\mathbb{C}}(y) - \theta^*\|^2$, it is sufficient to control the following quantity (e.g. see [219, Chapter 3], [13, Chapter 13]):

$$\mathbb{E} \left[\sup_{\theta \in \mathbb{C}: \|\theta - \theta^*\| \leq \delta} \langle \epsilon, \theta - \theta^* \rangle \right]$$

for some $\delta > 0$ to be chosen later. Since ϵ is σ -sub-Gaussian, we use Dudley's entropy integral to control the term above. We obtain

$$\mathbb{E} \left[\sup_{\theta \in \mathbb{C}: \|\theta - \theta^*\| \leq \delta} \langle \epsilon, \theta - \theta^* \rangle \right] \leq C\sigma \int_0^\delta \sqrt{\log N(\epsilon, \{\theta \in \mathbb{C}, \|\theta - \theta^*\| \leq \delta\}, \ell_2)} d\epsilon,$$

where $N(\epsilon, S, \ell_2)$ is the ϵ -covering number of a compact set S in ℓ_2 norm. Note that \mathbb{C} contains scaled versions of M fixed vectors $\{\theta_1, \dots, \theta_M\}$. For a fixed θ_i , with $i \in [M]$, the covering number $N(\epsilon, \{c\theta_i : c \in \mathbb{R}, \|\theta_i - \theta^*\| \leq \delta\}, \ell_2)$ is equivalent to the covering number of a bounded interval (in 1 dimension). Using [220], this is $(1 + \frac{2\delta}{\epsilon})$. Since there are M such fixed vectors, we obtain

$$N(\epsilon, \{\theta \in \mathbb{C}, \|\theta - \theta^*\| \leq \delta\}, \ell_2) \leq C_1 M (1 + \frac{\delta}{\epsilon}).$$

Substituting, we obtain

$$\mathbb{E} \left[\sup_{\theta \in \mathbb{C}: \|\theta - \theta^*\| \leq \delta} \langle \epsilon, \theta - \theta^* \rangle \right] \leq C\sigma \left(\delta \sqrt{\log M} + C_1 \delta \right).$$

Now, the critical inequality ([13, Chapter 13]) takes the form

$$\delta\sigma(\sqrt{\log M} + C_1) \lesssim \delta^2.$$

Hence we can choose $\delta = C_2\sigma(\sqrt{\log M} + C_1)$. Now, for any $t \geq \delta$, invoking [13, Theorem 13.2] yields the oracle inequality

$$\begin{aligned} \|P_{\mathbb{C}}(y) - \theta^*\|^2 &\leq c (\|\theta^* - P_{\mathbb{C}}(\theta^*)\|^2 + \sigma^2 t (\log M + c_1)) \\ &= c \left(\min_{\theta \in \mathbb{C}} \|\theta - \theta^*\|^2 + \sigma^2 t (\log M + c_1) \right), \end{aligned}$$

with probability exceeding $1 - c_2 e^{-nt(\sqrt{\log M} + c_1)}$, which proves the lemma. \square

6.9.14 NP-hardness of real phase retrieval

Our discussion borrows from a similar proof established in [103, Proposition 1] for mixtures of linear regressions. Recall that with n i.i.d observations $\{(x_i, y_i)\}_{i=1}^n$, the max-affine model takes the form

$$y_i = \max_{1 \leq j \leq k} (\langle x_i, \theta_j^* \rangle + b_j^*) + \epsilon_i,$$

where $\{\epsilon_i\}_{i=1}^n$ is a sequence of i.i.d zero mean sub-Gaussian noise.

We now consider a special case, where $k = 2$, $b_1^* = b_2^* = 0$, and $\theta_1^* = -\theta_2^*$, corresponding to the real phase retrieval problem. Furthermore, we consider the noiseless case $\epsilon = 0$. Our covariate matrix is given by $X \in \mathbb{R}^{n \times d}$ and the response vector by $y \in \mathbb{R}^n$. We now show that even in this special case, there is family of instances (X, y) such that solving the least squares problem (6.5) is NP-hard. In particular, we say that a “solution” to the noiseless phase retrieval problem exists on an instance (X, y) if the least squares objective in equation (6.5) has minimum value zero.

Proposition 17. *Deciding whether a problem instance (X, y) has a solution to the noiseless phase retrieval problem is NP-hard.*

Proof. The proof follows from a reduction to the subset-sum problem, the decision version of which is stated as follows: given p numbers $a_1, \dots, a_p \in \mathbb{R}$, we must decide if there exists a partition $S \subseteq [p]$ such that

$$\sum_{i \in S} a_i = \sum_{j \in S^c} a_j.$$

For each p -dimensional vector a , we design a problem instance (X, y) such that solving the noiseless (real) phase retrieval problem on (X, y) implies deciding on the subset sum problem specified by a .

To accomplish this, take $n = 2p + 1$ and $d = p$, and define the instance

$$X = \begin{bmatrix} I_p \\ I_p \\ 1 \dots 1 \end{bmatrix} \quad \text{and} \quad y = \begin{bmatrix} a \\ -a \\ 0 \end{bmatrix},$$

where I_p denotes the $p \times p$ identity matrix. By construction, finding a solution to the noiseless (real) phase retrieval problem on this instance corresponds to finding a subset $S \subseteq [2p + 1]$ and a pair of vectors (θ_1^*, θ_2^*) with $\theta_1^* = -\theta_2^*$, such that $X_S \theta_1^* = y_S$, and $X_{S^c} \theta_2^* = y_{S^c}$. Here X_S and y_S are the sub-matrix and sub-vector of X and y respectively restricted to the set S . Note that in general, the set S cannot contain the index i and $p + i$, since they correspond to two mutually exclusive equations. From this observation, we have $\theta_1^*(i) = \{a_i, -a_i\}$, and $\theta_1^*(i) = -\theta_2^*(i)$, where $\theta_1^*(i)$ and $\theta_2^*(i)$ denote the i -th coordinate of θ_1^* and θ_2^* , respectively.

As a consequence, if θ_1^* (and $\theta_2^* = -\theta_1^*$) satisfies the first $2p$ equations in this system, then the final equation demands that

$$\sum_{i \in S} \theta_1^*(i) = 0 = \sum_{j \in S^c} \theta_2^*(j).$$

By construction, note that this is accomplished if and only if

$$\sum_{i \in S} a_i = \sum_{j \in S^c} a_j,$$

and so a solution to the noiseless (real) phase retrieval problem on (X, y) yields a solution to the subset-sum problem, as desired. \square

Chapter 7

Max-affine Regression Beyond Gaussian Design—Small Ball covariates

In this chapter, we show that the Alternating Minimization (AM) algorithm is significantly robust. In particular, it converges locally under a weaker, “small-ball” design assumption (accommodating, for instance, and bounded log-concave distribution), and even when the underlying parameters are chosen with knowledge of the realized covariates. Once again, the final rate obtained by the procedure is near-parametric and minimax optimal (up to a polylogarithmic factor) as a function of the dimension, sample size, and noise variance. As a by-product of our analysis, we obtain convergence guarantees on a classical algorithm for the (real) phase retrieval problem in the presence of noise, and under considerably weaker assumptions on the design distribution than was previously known.

7.1 Introduction

Similar to Chapter 6, we show that AM converges under significantly weaker statistical assumptions. In particular, we allow the distribution of the covariates to come from the larger class of sub-Gaussian distributions that satisfy a *small-ball* condition. In addition, we also consider the scenario of *universal* parameter estimation, meaning that our guarantees hold uniformly over all $\beta_1^*, \dots, \beta_k^*$. This allows the parameters to be chosen with knowledge of the realized covariates, a robust setting that is commonly studied in signal processing applications like phase retrieval [56]. In contrast, our prior work [221] only handled the case where the parameters are fixed.

More precisely, our covariate assumption relies on the following definition.

Definition 12. (*Small-ball*) A distribution P_X satisfies a (ζ, c_s) -small-ball property if, for $X \sim P_X$ and each $\delta > 0$, we have

$$\sup_{u \in \mathbb{S}^{d-1}, w \in \mathbb{R}} \Pr \{ (\langle X, u \rangle + w)^2 \leq \delta \} \leq (c_s \delta)^\zeta. \quad (7.1)$$

The small-ball properties of various classes of distributions have been studied extensively in the probability literature [57, 58], and many natural distributions possess this property provided they are not too “peaky”. For instance, a simple calculation yields that provided the density of $\langle X, u \rangle$ is bounded by \sqrt{c} for each $u \in \mathbb{S}^{d-1}$, the distribution P_X satisfies the $(1/2, c)$ -small ball property. We now present our assumption on the covariate distribution; recall that $X \in \mathbb{R}^d$ is said to be η -sub-Gaussian if

$$\sup_{u \in \mathbb{S}^{d-1}} \mathbb{E}[\exp(\lambda \langle X, u \rangle)] \leq \exp\left(\frac{\lambda^2 \eta^2}{2}\right) \text{ for each } \lambda \in \mathbb{R}.$$

Assumption 20. *The distribution P_X is isotropic, η -sub-Gaussian, and satisfies a (ζ, c_s) small-ball condition.*

Let us briefly state a few examples where Assumption 20 is satisfied with particular values of the tuple (η, ζ, c_s) . The first is the class of compactly supported log-concave random vectors, which satisfy the small ball conditions with $(\zeta, c_s) = (1/2, C)$ for an absolute constant C (see [222, Appendix G.1]). Boundedness further implies sub-Gaussianity. As a specific example, consider X with each entry drawn i.i.d. according to the distribution $\text{Unif}[-\sqrt{3}, \sqrt{3}]$, which is commonly used as a random design in investigations of non-parametric regression problems [223]. The associated distribution P_X is isotropic by definition, and has $(\eta, \zeta, c_s) = (12, 1/2, C)$. Similarly, any other uniform distribution on a bounded, isotropic convex set would also satisfy Assumption 20. The second (canonical) example for which Assumption 20 is satisfied is the standard Gaussian distribution. As we verify in Appendix 7.6.1 with χ^2 tail bounds, the standard Gaussian satisfies $(\eta, \zeta, c_s) = (1, 1/2, e)$. Also, while we have only presented examples in which $\zeta = 1/2$, there are distributions that satisfy the small ball condition for other values of ζ : for example, any random variable with density $f(x) \propto e^{-\|x\|^c}$ for a positive constant c .

In order to make our guarantees on universal parameter estimation more clear, let us define a few geometric quantities induced by the max affine regression model (6.1). For $X \sim P_X$, let

$$\pi_j(\beta_1^*, \dots, \beta_k^*) := \Pr\{\langle X, \theta_j^* \rangle + b_j^* = \max_{j' \in [k]} (\langle X, \theta_{j'}^* \rangle + b_{j'}^*)\},$$

and define

$$\pi_{\min}(\beta_1^*, \dots, \beta_k^*) := \min_{j \in [k]} \pi_j(\beta_1^*, \dots, \beta_k^*).$$

For a fixed set of true parameters, the quantity $\pi_{\min}(\beta_1^*, \dots, \beta_k^*) \cdot n$ is the expected number of samples that are noisy linear combinations of one of these parameters. Thus, even if the underlying parameters are fixed, we can only hope to estimate these parameters when $\pi_{\min}(\beta_1^*, \dots, \beta_k^*)$ is sufficiently large.

The signal strength of the problem is the minimum separation

$$\Delta(\beta_1^*, \dots, \beta_k^*) = \min_{j, j': j \neq j'} \|\theta_j^* - \theta_{j'}^*\|^2.$$

We also define a notion of condition number, given by

$$\kappa(\beta_1^*, \dots, \beta_k^*) = \max_{j \in [k]} \frac{\max_{j' \neq j} \|\theta_j^* - \theta_{j'}^*\|^2}{\min_{j' \neq j} \|\theta_j^* - \theta_{j'}^*\|^2}.$$

We often use the shorthand

$$\pi_{\min} = \pi_{\min}(\beta_1^*, \dots, \beta_k^*), \quad \Delta = \Delta(\beta_1^*, \dots, \beta_k^*), \quad \text{and} \quad \kappa = \kappa(\beta_1^*, \dots, \beta_k^*).$$

when the true parameters $\beta_1^*, \dots, \beta_k^*$ are clear from context.

Recall that our goal was to prove a result that holds uniformly for all true parameters $\{\beta_j^*\}_{j=1}^k$. However, this is clearly impossible in a general sense, since we cannot hope to obtain consistent estimates if some parameters are never observed in the sample. A workaround is to hold certain geometric quantities fixed while sweeping over all possible allowable parameters $\beta_j^*, j = 1, \dots, k$. Accordingly, for each triple of positive scalars (π, Δ, κ) , we define the set of “admissible” true parameters as

$$\text{B}_{\text{vol}}(\pi, \Delta, \kappa) = \{\beta_1, \dots, \beta_k : \pi_{\min}(\beta_1, \dots, \beta_k) \geq \pi, \Delta(\beta_1, \dots, \beta_k) \geq \Delta, \kappa(\beta_1, \dots, \beta_k) \leq \kappa\}.$$

We let the true parameters $\beta_1^*, \dots, \beta_k^*$ take values in the set $\text{B}_{\text{vol}}(\pi, \Delta, \kappa)$, and prove guarantees uniformly over all such $\beta_1^*, \dots, \beta_k^*$.

With these definitions at hand, we are now ready to discuss our contributions.

Contributions: Suppose Assumption 20 holds, and $\{\beta_j^{(t)}\}_{j=1}^k$ are the parameter estimates returned by the AM algorithm at the t -th iteration. In Theorem 18, we show that for any $\epsilon > 0$, there is a sufficiently large iteration t such that

$$\sum_{i=1}^k \|\beta_i^{(t)} - \beta_i^*\|^2 \leq \epsilon + C_{\eta, \zeta, c_s}^{(1)} \frac{\sigma^2 k d}{n \pi_{\min}^{1+2\zeta-1}} \log(kd) \log\left(\frac{n}{kd}\right).$$

Such a result holds *simultaneously* for all $(\beta_1^*, \dots, \beta_k^*) \in \text{B}_{\text{vol}}(\pi_{\min}, \Delta, \kappa)$ with high probability provided the sample size is large enough and the initialization is chosen close enough to the true parameters. All of these aspects of the result are quantitative, and clarified in the statements of the formal theorems to follow.

In Corollary 3, we specialize Theorem 18 to the phase retrieval problem, showing that a slight variant of the same AM algorithm [95] exhibits linear convergence provided the covariates (called “measurements” in the signal processing literature) satisfy Assumption 20. Since our result holds for universal parameter estimation, we allow the underlying “signal” to be chosen adversarially, with knowledge of the measurements. Such a robust setting is common for phase retrieval problems. However, to the best of our knowledge, all previous results on the AM algorithm for phase retrieval [59, 60] only held under the assumptions of Gaussian covariates and noiseless observations, and/or required resampling of the measurements [61]. Ours is thus the first work to

handle non-Gaussian covariates in the presence of noise, while also analyzing the algorithm without resampling.

In Section 6.3, we validate the local convergence of AM with $\text{Unif}[-\sqrt{3}, \sqrt{3}]^{\otimes d}$ covariates. We observe that, when $\sigma = 0$, AM recovers the true parameters $\{\beta_i^*\}_{i=1}^k$ within 4 – 5 iterations. Furthermore, when $\sigma \neq 0$, the iterates of AM plateaus at the statistical error of the problem, which has a linear dependence on d/n . Moreover, in Section 7.2.4, we empirically show the performance of two initialization techniques; (a) random initialization with multiple restarts and (b) heuristic Principal Component Analysis (PCA) based initialization. For (a), we randomly initialize AM with multiple seeds, and observe that around 80 – 100 restarts is sufficient for a moderate dimensional problem. Furthermore, for (b), we use an appropriately defined second moment matrix (a function of $\{y_i, x_i\}_{i=1}^n$), and running PCA in conjunction with a random search method (Algorithm 3 of [221]), we obtain the initialization required for AM. Hence, the requirement of ‘good’ initialization is not a deterrent in the practical usage of AM.

7.2 Main results

Let us now state and discuss our results in precise terms.

7.2.1 Local geometric convergence of alternating minimization

For each pair $1 \leq i \neq j \leq k$ and $t \geq 0$, we use the shorthand $v_{i,j}^* = \beta_i^* - \beta_j^*$ and $v_{i,j}^{(t)} = \beta_i^{(t)} - \beta_j^{(t)}$ to denote the pairwise differences between parameters.

Theorem 18. *Suppose that Assumption 20 holds. Then there exists a pair of universal positive constants (c_1, c_2) and positive constants $(C_{\eta, \zeta, c_s}^{(1)}, C_{\eta, \zeta, c_s}^{(2)})$ depending only on the triple (η, ζ, c_s) such that if the sample size satisfies the bound*

$$n \geq C_{\eta, \zeta, c_s}^{(1)} \max \{d, 10 \log n\} \times \max \left\{ \frac{k\kappa}{\pi_{\min}^{1+2\zeta-1}} \log(n/d), \sigma^2 \frac{k^{1+2\zeta-1}}{\Delta^{\zeta-1} \pi_{\min}^{(1+2\zeta-1)^2}} \log^{2+2\zeta-1} \left(\frac{k}{\pi_{\min}^{1+2\zeta-1}} \right) \log(kd) \right\} \quad (7.2)$$

then simultaneously for all true parameters $\beta_1^*, \dots, \beta_k^* \in \mathcal{B}_{\text{Vol}}(\pi_{\min}, \Delta, \kappa)$ and all initializations satisfying

$$\min_{c>0} \max_{1 \leq j \neq j' \leq k} \frac{\|c v_{j,j'}^{(0)} - v_{j,j'}^*\|}{\|\theta_j^* - \theta_{j'}^*\|} \leq C_{\eta, \zeta, c_s}^{(2)} \left(\frac{\pi_{\min}^{1+2\zeta-1}}{k\kappa} \right)^{\zeta-1} \left[\log^{-(1+\zeta-1)} \left(\frac{k\kappa}{\pi_{\min}^{1+2\zeta-1}} \right) \right], \quad (7.3a)$$

the estimation error at all iterations $t \geq 1$ satisfies

$$\sum_{j=1}^k \|\beta_j^{(t)} - \beta_j^*\|^2 \leq \left(\frac{3}{4} \right)^t \left(\sum_{j=1}^k \|c^* \beta_j^{(0)} - \beta_j^*\|^2 \right) + C_{\eta, \zeta, c_s}^{(1)} \cdot \sigma^2 \frac{kd}{n\pi_{\min}^{1+2\zeta-1}} \log(kd) \log(n/kd) \quad (7.3b)$$

with probability exceeding $1 - c_1 \left\{ \frac{k^2}{n^7} + \exp(-c_2 n \pi_{\min}^2) \right\}$. Here, $c^* (> 0)$ minimizes the LHS of inequality (7.3a).

Let us discuss the initialization conditions of the theorem in more detail. We require the initialization $\beta_1^{(0)}, \dots, \beta_k^{(0)}$ to satisfy condition (7.3a). In the well-balanced case (with $\pi_{\min} \sim 1/k$) and treating k as constant, the initialization condition (7.3a) posits that the parameters are a constant “distance” from the true parameters. Closeness here is measured in a relative sense, i.e., between pairwise differences of the parameter estimates as opposed to the parameters themselves. The intuition for this is that $\beta_1^{(0)}, \dots, \beta_k^{(0)}$ induces a partition of samples $S_1(\beta_1^{(0)}, \dots, \beta_k^{(0)}), \dots, S_1(\beta_1^{(0)}, \dots, \beta_k^{(0)})$, and the closeness to this to the true partition depends only on the relative pairwise differences between parameters. Furthermore, the initialization condition is also invariant to a global scaling of the parameters, since scaling does not change the initial partition of samples. Also note that the geometric convergence guarantee (7.3b) holds uniformly for all initializations satisfying condition (7.3a). Hence, the initialization parameters are not additionally required to be independent of the covariates or noise.

Let us now turn our attention to the bound (7.3b), which consists of two terms. When $t \rightarrow \infty$, the second term of the bound (7.3b) provides an estimate of the closeness of the final parameters to the true parameters. Up to a constant, this is the statistical error term

$$\delta_{n,\sigma}^{\text{sb}}(d, k, \pi_{\min}) = \sigma^2 \frac{kd}{\pi_{\min}^{1+2\zeta^{-1}} n} \log(kd) \log(n/kd) \quad (7.4)$$

that converges to 0 as $n \rightarrow \infty$, thereby providing a consistent estimate in the large sample limit. The dependence on π_{\min} is discussed shortly.

The first term of (7.3b) is an optimization error that is best interpreted in the noiseless case $\sigma = 0$, wherein the parameters $\beta_1^{(t)}, \dots, \beta_k^{(t)}$ converge at a geometric rate to the true parameters $\beta_1^*, \dots, \beta_k^*$. In the noiseless case, we obtain exact recovery of the parameters provided

$$n \geq C \frac{kd}{\pi_{\min}^{1+2\zeta^{-1}}} \log(n/d).$$

Thus, the “sample complexity” of parameter recovery is linear in the dimension d , which is optimal. In the well-balanced case, we require $n \sim k^{2+2\zeta^{-1}} d$, but lower bounds based on parameter counting suggest that the true dependence ought to be linear. We are not aware of whether the dependence on π_{\min} in the noiseless case is optimal; our simulations in panel (a) suggest that the sample complexity depends inversely on π_{\min} , and so closing this gap is an interesting open problem.

In Figure 7.1, we verify that for independent, isotropic covariates chosen uniformly from a symmetric interval¹, initializing the AM algorithm in a neighborhood of the true parameters suffices to ensure that it converges to the true parameters. Furthermore, both the sample size requirement and final error of the algorithm exhibit the behaviors predicted by Theorem 18.

We now compare Theorem 18 with Chapter 6 (Theorem 1) for the special case of Gaussian covariates, where $\eta = 1$ and $\zeta = 1/2$. In this case, all terms of the form π_{\min}^3 in [221, Theorem 1]

¹Such a distribution is compactly supported and log-concave, and therefore satisfies Assumption 20.




Figure 7.1: Convergence of AM when the covariates are drawn i.i.d. from the distribution $\text{Unif}[-\sqrt{3}, \sqrt{3}]^{\otimes d}$. In panel (a), we plot the noiseless sample complexity of AM; we fix $\|\beta_i^*\| = 1$ for all $i \in [k]$, $\sigma = 0$ and $\pi_{\min} = 1/k$. We say β_i^* is recovered if $\|\beta_i^{(t)} - \beta_i^*\| \leq 0.01$. For a fixed dimension d , we run a linear search on the number of samples n , such that the empirical probability of success over 100 trials is more than 0.95, and output the least such n . In panel (b), we plot the optimization error (in blue) $\sum_{j=1}^k \|\beta_j^{(t)} - \beta_j^{(T)}\|^2$ and the deviation from the true parameters (in red) $\sum_{j=1}^k \|\beta_j^{(t)} - \beta_j^*\|^2 / \sigma^2$ over iterations t for various values of σ in the tuple $(0.15, 0.25, 0.4, 0.5)$, with $k = 5$, $d = 100$, $T = 50$ and $n = 5d$. the resulting error curves are averaged over 50 trials. Panel (c) shows that the estimation error at $T = 50$ scales at the parametric rate d/n , where we have chosen a fixed $k = 5$ and $\sigma = 0.25$.

are replaced by terms of the form π_{\min}^5 . In particular, we see that the initialization condition (7.3a) is more stringent and the final statistical rate of the estimate (corresponding to the limit $t \rightarrow \infty$) now attains an estimation error that is a factor π_{\min}^{-2} higher than the corresponding rate of Chapter 6 (Theorem 1). The sample size requirement is similarly affected. On the other hand, the geometric convergence result (7.3b) now holds uniformly for all true parameters $\beta_1^*, \dots, \beta_k^* \in \text{B}_{\text{vol}}(\pi_{\min}, \Delta, \kappa)$, as opposed to Chapter 6 (Theorem 1), which holds only when the true parameters are held fixed. The more stringent initialization condition and sample size requirements can be viewed as the price

Algorithm 10: Alternating minimization for real phase retrieval

Input: Data $\{x_i, y_i\}_{i=1}^n$; initial parameter estimate $\theta^{(0)} \in \mathbb{R}^d$; number of iterations T .

Output: Final estimator $\hat{\theta}$.

2 Initialize $t \leftarrow 0$.

3 **repeat**

5 Compute sign vector $s^{(t)}$ with i -th entry as

$$s_i^{(t)} = \text{sgn}(\langle x_i, \theta^{(t)} \rangle) \quad \text{for each } i \in [n]. \quad (7.5a)$$

6 Update

$$\theta^{(t+1)} = \arg \min_{\theta \in \mathbb{R}^d} \sum_{i=1}^n (y_i - s_i^{(t)} \langle x_i, \theta \rangle)^2. \quad (7.5b)$$

7 **until** $t = T$;

9 Return $\hat{\theta} = \theta^{(T)}$.

to pay for the more robust convergence of the AM algorithm. Notably, the dependence on σ, k, n, d remains unchanged.

7.2.2 Consequences for phase retrieval

A notable consequence of Theorem 18 is that it can be applied to the phase retrieval model—in which results are usually proved uniformly over all true parameters [56, 224]—to yield a convergence result under general distributional assumptions on the covariates. In particular, setting $\pi_{\min} = 1/2$ and $k = 2$ yields a local linear convergence result for the AM algorithm of the Gershberg-Saxton-Fienup type (presented for completeness as Algorithm 10) uniformly *for all* θ^* provided the covariates satisfy a small-ball condition. We note that other algorithms for phase retrieval have also been shown to succeed under such small-ball assumptions [225, 226].

Corollary 3. *Suppose that Assumption 20 holds, and that $\theta^{(t)}$ is the t -th iterate of Algorithm 10. There exists a universal positive constant c_1 and a pair of constants $(C_{\eta, \zeta, c_s}^{(1)}, C_{\eta, \zeta, c_s}^{(2)})$ depending only on (η, ζ, c_s) such that if*

$$n \geq C_{\eta, \zeta, c_s}^{(1)} \max\{d, 10 \log n\} \log(n/d),$$

then simultaneously for all true parameters $\theta^ \in \mathbb{R}^d$ and all initializations $\theta^{(0)}$ satisfying*

$$\min_{c > 0} \min_{s \in \{-1, 1\}} \frac{\|c\theta^{(0)} - s\theta^*\|}{\|\theta^*\|} \leq C_{\eta, \zeta, c_s}^{(2)}, \quad (7.6a)$$

the estimation error for all $t \geq 1$ satisfies

$$\min_{s \in \{-1,1\}} \|\theta^{(t)} - s\theta_j^*\|^2 \leq \left(\frac{3}{4}\right)^t \min_{s \in \{-1,1\}} \|\theta^{(0)} - s\theta^*\|^2 + \frac{c_1 \sigma^2 d \log(n/d)}{n}$$

with probability exceeding $1 - c_1 n^{-7}$.

Let us now compare this with the sharpest existing local convergence result of AM for phase retrieval due to Waldspurger [59], which holds for Gaussian covariates and in the noiseless setting². Specializing Corollary 3 to the noiseless setting, we observe that provided the ratio n/d is larger than a fixed constant (that depends only on (η, ζ, c_s)), we obtain exact recovery of the underlying parameter with high probability, up to a global sign, provided the measurement vectors are sub-Gaussian and satisfy a small-ball condition. To the best of our knowledge, prior work on the AM algorithm had not established provable guarantees for non-Gaussian covariates even in the noiseless setting. In the noisy case, Corollary 3 guarantees convergence of the iterates to a small neighborhood around either θ^* or $-\theta^*$, and the size of this neighborhood is within a logarithmic factor of being minimax optimal [56, 184]. Once again, to the best of our knowledge, guarantees for the AM algorithm as applied to noisy phase retrieval did not exist in the literature.

7.2.3 Proof ideas and technical challenges

We now sketch the high level proof ideas required to establish guarantees on the AM algorithm. Some parts of the proof here is similar to that of Chapter 6. For example, Similar to Chapter 6, here also, the proof consists of 3 steps: (a) controlling the behavior of noise, (b) analyzing the prediction error of the noiseless problem and (c) translating the prediction error to estimation error by inverting a specifically chosen sub-matrix of the covariate matrix. The analysis of (a) parallels to that of Chapter 6. However, for part (b) and (c), the weakened assumption on covariates poses non-trivial technical difficulties.

In order to control the prediction error of the noiseless problem, we crucially use the anti concentration property of the small-ball distribution. Using the good initialization condition in conjunction with the above mentioned lemmas, we upper-bound the noiseless prediction error of the AM iterates. Finally, in order to translate our bounds on the prediction error into bounds on parameter estimation, we invert specifically chosen sub-matrices of the covariate matrix over the course of the algorithm. Our bounds naturally depend on how these sub-matrices are conditioned. A key technical difficulty of the proof is to control the spectrum of these random matrices, rows of which are drawn from (randomly) truncated variants of the covariate distribution. Our techniques for controlling the spectrum of these matrices is more broadly applicable, and we expect this result to be of broader interest.

The technical meat of the paper crucially depends on the two technical lemmas we obtain that bound the maximum and minimum singular values of a matrix whose rows are drawn from a

²Waldspurger [59] deals with the complex phase retrieval, whose analysis is significantly more complicated than real phase retrieval considered here.




Figure 7.2: Initialization and overall guarantee with bounded uniform covariates drawn according to the distribution $\text{Unif}[-\sqrt{3}, \sqrt{3}]^{\otimes d}$. In panel (a) we fix $k = 3, d = 10, n = 30kd, \sigma = 0.1$ and $\pi_{\min} = 1/3$. We evaluate the performance of AM with repeated random initialization [154] and plot the estimation error with M , the number of restarts. In panel (b), we run Algorithm 2 from our companion paper [221] in order to initialize the AM algorithm. We fix $k = 3, d = 50, n = 30kd, \sigma = 0.1$ and $\pi_{\min} = 1/3$. All the data points are obtained as a result of averaging 25 trials of the experiment.

sub-Gaussian distribution obeying the small-ball property. Note that these results hold uniformly for any matrix that has sufficiently many rows. Our results on the minimum singular value are similar in spirit to those of Rudelson and Vershynin [58], but proved under a slightly different set of assumptions. We apply these results on carefully chosen sub-matrices of the covariate matrix Ξ .

7.2.4 Initialization and overall guarantee

We conclude our discussion of the main results with some simulations demonstrating that the ‘good’ initialization condition required in Theorem 18 for the convergence of the AM algorithm is indeed achievable in practice. In particular, we consider two approaches. We first evaluate the random initialization scheme with multiple restarts proposed by Balasz [154]. We also simulate the PCA-based approach proposed in Chapter 6 for Gaussian covariates, wherein we use a moment method to convert the regression problem in d -dimensions into one in $k + 1$ dimensions, followed by a random initialization scheme. This method is particularly useful in the high dimensional regime $d \gg k$, and is described formally in Chapter 6 (Algorithms 2 and 3).

In Figure 7.2 (a), we initialize AM randomly, with multiple restarts, following [154]. We observe that with sufficiently many restarts, AM converges. Hence, the requirement of ‘good’ initialization can be replaced by random initialization with multiple restarts. This is how AM is used in practice as seen in [154]. However, we observe that with increasing dimension, the number of restarts M increases rapidly, and hence we get hit by the ‘curse of dimensionality’. For this reason, we consider a low dimensional setting with $d = 10$ and observe that around 110 restarts suffices. Understanding this repeated random initialization algorithm and obtaining explicit dependence on the number of restarts (as a function of dimension d , and the number of affine pieces k) is an interesting open

problem.

We attempt to resolve the curse of dimensionality via a heuristic PCA algorithm. Our algorithm matches exactly Algorithm 2 of Chapter 6, except the definition of \widehat{M}_2 (the second moment used in the spectral method, see [221] for details). We define $\widehat{M}_2 = \frac{2}{n} \sum_{i=1}^{n/2} (y_i^2/2) (x_i x_i^\top - I_d)$. After obtaining the subspace spanned the true parameters, we run the random search method (Algorithm 3 of Chapter 6) to obtain the initialization, and finally run AM with it. We consider a moderate dimensional problem ($d = 50$), and show (in Figure 7.2(b)) that 100 restarts are sufficient. Note that as expected, this algorithm can handle a high dimensional problem; thanks to the PCA algorithm. This implies that, with a slight variation in the definition of \widehat{M}_2 , the PCA algorithm is indeed robust with respect to the covariate distribution: it works (at least empirically) even when the covariates are sub-Gaussian satisfying a small-ball condition. Obtaining provable guarantees for PCA with non-Gaussian covariates is kept as an interesting future work.

7.3 Discussion and Open Problems

We analyze a natural alternating minimization algorithm for estimating the maximum of unknown affine functions, when the covariates come from sub-Gaussian distribution with small-ball condition. Additionally, we also consider the setting of universal parameter estimation, meaning that our results hold uniformly for all true parameters. In particular, this also allows the parameters be chosen with the knowledge of the covariates. We establish that AM enjoys local linear convergence to a ball around the optimal parameters. We characterize the statistical error of AM, and empirically verify the dependence on problem parameters like d and n . As a corollary, specializing to the (real) phase retrieval setting, we obtain several new results; in particular we obtain provable guarantees for the phase retrieval problem with non-Gaussian covariates and in the presence of noise.

Moreover, we experimentally show two techniques for the initializing AM using random initialization (with restarts) and PCA. However, we do not have a provable initialization guarantee. An immediate future work would to obtain a provable initialization. An important question that remains to be addressed is whether the heuristic of random initialization with multiple restarts, that has been employed both by us and in prior work can be shown to provably converge with an explicit bound on the number of repetitions. Understanding the behavior of the randomly initialized AM algorithm is also an open problem in the context of phase retrieval [59, 201].

Also, it would be interesting to study this problem in very large dimension with a sparsity assumption on the true parameters. Furthermore, another interesting problem is to obtain a minimax lower bound on estimation error in order to understand the (actual) dependence on several problem parameters. We keep these as our future endeavors.

Appendix

We now present proofs of our main results. We assume throughout that the sample size n is larger than some universal constant. Also, we assume that the values of constants c, c_1, c', \dots may change

from line to line.

7.4 Proof of Theorem 18

Throughout the proof, we make references to the proof of [221, Theorem 1]. Recall our observation model (6.1). We use the notation $\Xi \in \mathbb{R}^{n \times (d+1)}$ for the covariate matrix, $y \in \mathbb{R}^n$ for the response vector. Recall that the noise is σ sub-Gaussian. Let us begin by introducing some shorthand notation. For a scalar w^* , vectors $u^* \in \mathbb{R}^d$ and $v^* = (u^*, w^*) \in \mathbb{R}^{d+1}$, and a positive scalar r , let $\mathcal{B}_{v^*}(r) = \left\{ v \in \mathbb{R}^{d+1} : \frac{\|v-v^*\|}{\|u^*\|} \leq r \right\}$. Define the set

$$\mathcal{I} \left(r; \{\beta_j^*\}_{j=1}^k \right) = \left\{ \beta_1, \dots, \beta_k \in \mathbb{R}^{d+1} : \exists c > 0 : c(\beta_i - \beta_j) \in \mathcal{B}_{\beta_i^* - \beta_j^*}(r) \text{ for all } 1 \leq i \neq j \leq k \right\}.$$

Also, use

$$\vartheta_t \left(r; \{\beta_j^*\}_{j=1}^k \right) := \sup_{\beta_1^{(0)}, \dots, \beta_k^{(0)} \in \mathcal{I}(r)} \sum_{j=1}^k \|\beta_j^{(t)} - \beta_j^*\|^2 - \left(\frac{3}{4} \right)^t \left(\sum_{j=1}^k \|c^* \beta_j^{(0)} - \beta_j^*\|^2 \right)$$

to denote the error tracked over iterations (with c^* denoting the smallest $c > 0$ such that $c(\beta_i - \beta_j) \in \mathcal{B}_{\beta_i^* - \beta_j^*}(r)$ for all $1 \leq i \neq j \leq k$). Finally, we use the shorthand

$$\delta_{n,\sigma}^{\text{sb}}(d, k, \pi_{\min}) := \sigma^2 \frac{kd}{n\pi_{\min}^{1+2\zeta-1}} \log(kd) \log(n/kd).$$

to denote the final statistical rate. With the rest of the notation remaining the same as before, the theorem claims that there exist constants such that if condition (7.2) is satisfied, then we have

$$\begin{aligned} & \Pr \left\{ \max_{t \geq 1} \sup_{\beta_1^*, \dots, \beta_k^* \in \mathcal{B}_{\text{vol}}(\pi_{\min}, \Delta, \kappa)} \vartheta_t \left(C_{\eta, \zeta, c_s}^{(2)} \left(\frac{\pi_{\min}^{1+2\zeta-1}}{\kappa k} \right)^{\zeta-1} ; \{\beta_j^*\}_{j=1}^k \right) \geq C_{\eta, \zeta, c_s}^{(1)} \delta_{n,\sigma}^{\text{sb}}(d, k, \pi_{\min}) \right\} \\ & \leq 1 - c_1 \left\{ \frac{k^2}{n^7} + \exp(-c_2 n \pi_{\min}^2) \right\} \end{aligned} \quad (7.7)$$

We dedicate the rest of the proof to establishing (7.7). Without loss of generality, we assume that the scalar c^* is equal to 1. It is convenient to state and prove another result that guarantees a one-step contraction, from which Theorem 18 follows as a corollary. In order to state this result, we assume that one step of the alternating minimization update is run starting from the parameters $\{\beta_j\}_{j=1}^k$ to produce the next iterate $\{\beta_j^+\}_{j=1}^k$. In the statement of the proposition, we use the shorthand

$$\begin{aligned} v_{i,j}^* &= \beta_i^* - \beta_j^*, \\ v_{i,j} &= \beta_i - \beta_j, \text{ and} \\ v_{i,j}^+ &= \beta_i^+ - \beta_j^+. \end{aligned}$$

Also recall the definitions of the geometric quantities (Δ, κ) . The following proposition guarantees the one step contraction bound.

Proposition 19. *Suppose that Assumption 20 holds. Then there exists a tuple of universal constants (c_1, c_2) and another tuple of constants $(C_{\eta, \zeta, c_s}^{(1)}, C_{\eta, \zeta, c_s}^{(2)})$ depending only on the tuple (η, ζ, c_s) such that*

(a) *If the sample size satisfies the bound $n \geq c_1 \max \{d, 10 \log n\} \max \left\{ \frac{1+\zeta^{-1}}{\pi_{\min}}, \frac{k}{\pi_{\min}^2} \log(n/d) \right\}$, then for any set of parameters $\{\beta_j^*\}_{j=1}^k \in \mathcal{B}_{\text{vol}}(\pi_{\min}, \Delta, \kappa)$ and $\{\beta_j\}_{j=1}^k$ satisfying*

$$\max_{1 \leq j \neq j' \leq k} \left(\frac{\|v_{j,j'} - v_{j,j'}^*\|}{\|\theta_j^* - \theta_{j'}^*\|} \right)^\zeta \log^{1+\zeta} \left(\frac{\|\theta_j^* - \theta_{j'}^*\|}{\|v_{j,j'} - v_{j,j'}^*\|} \right) \leq C_{\eta, \zeta, c_s}^{(2)} \left(\frac{\pi_{\min}^{1+2\zeta^{-1}}}{k\kappa} \right), \quad (7.8a)$$

we have, simultaneously for all pairs $1 \leq j \neq \ell \leq k$, the bound

$$\begin{aligned} \frac{\|v_{j,\ell}^+ - v_{j,\ell}^*\|^2}{\|\theta_j^* - \theta_\ell^*\|^2} &\leq C_{\eta, \zeta, c_s}^{(1)} \max \left\{ \frac{d\zeta\kappa}{\pi_{\min}^{1+2\zeta^{-1}} n}, \frac{1}{4k} \right\} \left(\sum_{j'=1}^k \frac{\|v_{j,j'} - v_{j,j'}^*\|^2}{\|\theta_j^* - \theta_{j'}^*\|^2} + \frac{\|v_{\ell,j'} - v_{\ell,j'}^*\|^2}{\|\theta_\ell^* - \theta_{j'}^*\|^2} \right) \\ &\quad + C_{\eta, \zeta, c_s}^{(1)} \sigma^2 \frac{kd}{\pi_{\min}^{1+2\zeta^{-1}} n} \log(n/d) \end{aligned} \quad (7.8b)$$

with probability exceeding $1 - c_1 \left\{ \frac{k^2}{n^7} + \exp(-c_2 n \pi_{\min}^2) \right\}$.

(b) *If the sample size satisfies the bound*

$$n \geq c_1 \max \left\{ \max \{d, 10 \log n\} \max \left\{ \frac{1 + \zeta^{-1}}{\pi_{\min}}, \frac{k}{\pi_{\min}^2} \log(n/d) \right\}, C_{\eta, \zeta, c_s}^{(1)} \frac{kd}{\pi_{\min}^{1+2\zeta^{-1}}} \right\},$$

then for any set of parameters $\{\beta_j^\}_{j=1}^k \in \mathcal{B}_{\text{vol}}(\pi_{\min}, \Delta, \kappa)$ and $\{\beta_j\}_{j=1}^k$ satisfying*

$$\max_{1 \leq j \neq j' \leq k} \left(\frac{\|v_{j,j'} - v_{j,j'}^*\|}{\|\theta_j^* - \theta_{j'}^*\|} \right)^\zeta \log^{1+\zeta} \left(\frac{\|\theta_j^* - \theta_{j'}^*\|}{\|v_{j,j'} - v_{j,j'}^*\|} \right) \leq C_{\eta, \zeta, c_s}^{(2)} \left(\frac{\pi_{\min}^{1+2\zeta^{-1}}}{k} \right), \quad (7.9a)$$

we have the overall estimation error bound

$$\sum_{i=1}^k \|\beta_j^+ - \beta_j^*\|^2 \leq \frac{3}{4} \cdot \left(\sum_{i=1}^k \|\beta_j - \beta_j^*\|^2 \right) + C_{\eta, \zeta, c_s}^{(1)} \sigma^2 \frac{kd}{\pi_{\min}^{1+2\zeta^{-1}} n} \log(kd) \log(n/kd) \quad (7.9b)$$

with probability exceeding $1 - c_1 \left\{ \frac{k^2}{n^7} + \exp(-c_2 n \pi_{\min}^2) \right\}$.

Let us briefly comment on how Theorem 18 follows from Proposition 19. Note that the equations (7.9a) and (7.9b) in conjunction clearly show that geometric decay of the estimation error after running one step of the algorithm. We only need to verify that the next iterates $\{\beta_j^+\}_{j=1}^k$ also satisfy condition (7.8a) provided the sample size n is large enough. Hence, the one step estimation bound (7.9b) can be applied recursively to obtain the final bound (7.3b).

For the constant $C_{\eta, \zeta, c_s}^{(2)}$ in the proposition, let r_a be the largest scalar in the interval $[0, e^{-(1+\zeta^{-1})}]$ such that $r_a^\zeta \log^{1+\zeta}(1/r_a) \leq C_{\eta, \zeta, c_s}^{(2)} \left(\frac{\pi_{\min}^{1+2\zeta^{-1}}}{k\kappa} \right)$, and let r_b be the largest scalar in the interval $[0, e^{-(1+\zeta^{-1})}]$ with $r_b^\zeta \log^{1+\zeta}(1/r_b) \leq C_{\eta, \zeta, c_s}^{(2)} \left(\frac{\pi_{\min}^{1+2\zeta^{-1}}}{k} \right)$.

Assume that the current parameters satisfy the bound (7.8a). Choosing $n \geq 4\kappa\zeta d / \pi_{\min}^{(1+2\zeta^{-1})\zeta^{-1}}$ and applying inequality (7.8b), we have, for each pair $1 \leq j \neq \ell \leq k$, the bound

$$\begin{aligned} \frac{\|v_{j,\ell}^+ - v_{j,\ell}^*\|^2}{\|\theta_j^* - \theta_\ell^*\|^2} &\leq C_{\eta, \zeta, c_s}^{(1)} \frac{1}{4k} \left(\sum_{j'=1}^k \frac{\|v_{j,j'} - v_{j,j'}^*\|^2}{\|\theta_j^* - \theta_{j'}^*\|^2} + \frac{\|v_{\ell,j'} - v_{\ell,j'}^*\|^2}{\|\theta_\ell^* - \theta_{j'}^*\|^2} \right) \\ &\quad + C_{\eta, \zeta, c_s}^{(2)} \sigma^2 \frac{kd}{\pi_{\min}^{1+2\zeta^{-1}} n} \log(n/d) \\ &\leq \frac{1}{2} r_a^2 + C_{\eta, \zeta, c_s}^{(2)} \sigma^2 \frac{kd}{\pi_{\min}^{1+2\zeta^{-1}} n} \log(n/d) \end{aligned}$$

Further, if

$$n \geq C \sigma^2 \frac{k^{1+2\zeta^{-1}}}{r_a^2 \Delta^{\zeta^{-1}} \pi_{\min}^{(1+2\zeta^{-1})^2}} \log^{2+2\zeta^{-1}} \left(\frac{k}{\pi_{\min}^{1+2\zeta^{-1}}} \right) \log(kd),$$

for a sufficiently large constant C , we have

$$\frac{\|v_{j,\ell}^+ - v_{j,\ell}^*\|^2}{\|\theta_j^* - \theta_\ell^*\|^2} \leq r_a^2.$$

Thus, the parameters $\{\beta_j^+\}_{j=1}^k$ satisfy inequality (7.8a) for the sample size choice required by Theorem 18. Finally, adjusting constants and simplifying the probability statement completes the proof of the theorem.

7.4.1 Proof of Proposition 19

We use the shorthand notation $S_j := S_j(\beta_1, \dots, \beta_k)$, and let P_{S_j} denote the projection matrix onto the range of the matrix Ξ_{S_j} . Recall our notation for the difference vectors.

Let y^* denote the vector with entry i given by $\max_{\ell \in [k]} \langle \xi_i, \beta_\ell^* \rangle$. We have

$$\begin{aligned} \|\Xi_{S_j}(\beta_j^+ - \beta_j^*)\|^2 &= \|P_{S_j} y_{S_j} - \Xi_{S_j} \beta_j^*\|^2 \\ &= \|P_{S_j} y_{S_j}^* + P_{S_j} \epsilon_{S_j} - \Xi_{S_j} \beta_j^*\|^2 \\ &\leq 2\|P_{S_j}(y_{S_j}^* - \Xi_{S_j} \beta_j^*)\|^2 + 2\|P_{S_j} \epsilon_{S_j}\|^2 \\ &\leq 2\|y_{S_j}^* - \Xi_{S_j} \beta_j^*\|^2 + 2\|P_{S_j} \epsilon_{S_j}\|^2, \end{aligned} \tag{7.10}$$

where we have used the fact that the projection operator is non-expansive on a convex set.

Let

$$\{\langle \xi_i, \beta_\ell \rangle = \max\} := \left\{ \langle \xi_i, \beta_\ell \rangle = \max_{u \in [k]} \langle \xi_i, \beta_u \rangle \right\}, \text{ for each } i \in [n], \ell \in [k]$$

denote a convenient shorthand for these events. The first term on the RHS of inequality (7.10) can be written as

$$\sum_{i \in S_j} (y_i^* - \langle \xi_i, \beta_j^* \rangle)^2 \leq \sum_{i=1}^n \sum_{j': j' \neq j} \mathbf{1} \{ \langle \xi_i, \beta_j \rangle = \max \text{ and } \langle \xi_i, \beta_{j'} \rangle = \max \} \langle \xi_i, \beta_{j'}^* - \beta_j^* \rangle^2,$$

where the inequality accounts for ties. Each indicator random variable is bounded, in turn, as

$$\begin{aligned} \mathbf{1} \{ \langle \xi_i, \beta_j \rangle = \max \text{ and } \langle \xi_i, \beta_{j'} \rangle = \max \} &\leq \mathbf{1} \{ \langle \xi_i, \beta_j \rangle \geq \langle \xi_i, \beta_{j'} \rangle \text{ and } \langle \xi_i, \beta_{j'}^* \rangle \geq \langle \xi_i, \beta_j^* \rangle \} \\ &= \mathbf{1} \{ \langle \xi_i, v_{j,j'} \rangle \cdot \langle \xi_i, v_{j,j'}^* \rangle \leq 0 \}. \end{aligned}$$

Switching the order of summation yields the bound

$$\sum_{i \in S_j} (y_i^* - \langle \xi_i, \beta_j^* \rangle)^2 \leq \sum_{j': j' \neq j} \sum_{i=1}^n \mathbf{1} \{ \langle \xi_i, v_{j,j'} \rangle \cdot \langle \xi_i, v_{j,j'}^* \rangle \leq 0 \} \langle \xi_i, v_{j,j'}^* \rangle^2.$$

Recalling our notation for the minimum eigenvalue of a symmetric matrix, the LHS of inequality (7.10) can be bounded as

$$\|\Xi_{S_j}(\beta_j^+ - \beta_j^*)\|^2 \geq \lambda_{\min} \left(\Xi_{S_j}^\top \Xi_{S_j} \right) \cdot \|\beta_j^+ - \beta_j^*\|^2.$$

Hence, we arrive at the bound

$$\lambda_{\min} \left(\Xi_{S_j}^\top \Xi_{S_j} \right) \cdot \|\beta_j^+ - \beta_j^*\|^2 \lesssim \sum_{j': j' \neq j} \sum_{i=1}^n \mathbf{1} \{ \langle \xi_i, v_{j,j'} \rangle \cdot \langle \xi_i, v_{j,j'}^* \rangle \leq 0 \} \langle \xi_i, v_{j,j'}^* \rangle^2 + \|P_{S_j} \epsilon_{S_j}\|^2.$$

For the second term (noise) on the RHS, we apply [221, Lemma 3]. We obtain

$$\sup_{\beta_1, \dots, \beta_k \in \mathbb{R}^{d+1}} \max_{j \in [k]} \|P_{S_j} \epsilon_{S_j}\|^2 \leq 2\sigma^2 k(d+1) \log(n/d)$$

with probability greater than $1 - \binom{n}{d}^{-1}$.

We now make two claims to bound the remaining two terms in the bound. Recall that Assumption 20 holds, and use the shorthand $\tilde{\eta} := \max\{1, \eta\}$. Our first claim bounds the indicator quantities under this assumption. In particular, we claim that there exists a tuple of universal

constants (C, c_1, C', c') such that for any $r \leq 1/24$, we have

$$\begin{aligned} & \Pr \left\{ \exists \beta_1^*, \dots, \beta_k^* \in \mathbb{R}^{d+1} \text{ and } \beta_1, \dots, \beta_k \text{ such that } v_{j,j'} \in \mathcal{B}_{v^*}(r) \text{ for all } 1 \leq j \neq j' \leq k : \right. \\ & \quad \sum_{j':j' \neq j} \sum_{i=1}^n \mathbf{1} \left\{ \langle \xi_i, v_{j,j'} \rangle \cdot \langle \xi_i, v_{j,j'}^* \rangle \leq 0 \right\} \langle x_i, v_{j,j'}^* \rangle^2 \\ & \quad \geq C\tilde{\eta}^2 \sum_{j':j' \neq j} \|v_{j,j'} - v_{j,j'}^*\|^2 \max \left\{ (1 + \zeta^{-1})d, n \cdot (1 + \zeta) \left\{ (\tilde{\eta}r)^\zeta \log^{\zeta+1}(1/\tilde{\eta}r) \right\} \right\} \left. \right\} \\ & \leq c_1 \binom{k}{2} \left\{ n e^{-C'n} + c e^{-c' \max\{d, 10 \log n\}} \right\}. \end{aligned} \quad (7.11a)$$

The singular values of sub-matrices of Ξ are then handled by the following claim: that there exist universal constants (C, c, c') and a constant C'_{η, ζ, c_s} that depends only the tuple (η, ζ, c_s) , such that if $n \geq Cd \max \left\{ \frac{1+\zeta^{-1}}{\pi_{\min}}, \frac{k}{\pi_{\min}^2} \log(n/kd) \right\}$, then

$$\begin{aligned} & \Pr \left\{ \inf_{\beta_1^*, \dots, \beta_k^* \in \mathcal{B}_{\text{vol}}(\pi_{\min}, \Delta, \kappa)} \inf_{\substack{\beta_1, \dots, \beta_k \in \mathbb{R}^{d+1}: \\ v_{j,j'} \in \mathcal{B}_{v_{j,j'}^*}(r_b)}} \min_{j \in [k]} \lambda_{\min} \left(\Xi^j(\beta_1, \dots, \beta_k)^\top \cdot \Xi^j(\beta_1, \dots, \beta_k) \right) \right. \\ & \quad \left. \leq C'_{\eta, \zeta, c_s} \pi_{\min}^{1+2\zeta^{-1}} n \right\} \leq c \exp(-c'n \cdot \pi_{\min}^2). \end{aligned} \quad (7.11b)$$

The proof of the proposition follows directly from these claims. \square

We now prove both the given claims in turn.

Proof of claim (7.11a): The proof of the claim hinges on the following lemma.

Lemma 45. *Suppose that Assumption 20 holds. Then, there exist universal constants (C, c_1, c_2) such for all positive scalars $r \leq 1/24$, we have*

$$\begin{aligned} & \sup_{\substack{v^* \in \mathbb{R}^{d+1} \\ v \in \mathcal{B}_{v^*}(r)}} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{1} \left\{ \langle \xi_i, v \rangle \cdot \langle \xi_i, v^* \rangle \leq 0 \right\} \langle \xi_i, v^* \rangle^2 \right) \Big/ \|v - v^*\|^2 \\ & \leq C\tilde{\eta}^2 \cdot \max \left\{ (1 + \zeta^{-1}) \frac{\max\{d, 10 \log n\}}{n}, (1 + \zeta) (\tilde{\eta}r)^\zeta \cdot \log^{\zeta+1} \left(\frac{1}{\tilde{\eta}r} \right) \right\} \end{aligned}$$

with probability exceeding $1 - c_1 n^{-8}$, where again, we adopt the convention that $0/0 = 0$.

This lemma immediately establishes the claim in conjunction with a union bound over all $\binom{k}{2}$ pairs of parameters. \square

Proof of claim (7.11b): For this claim, we require a uniform bound on singular values uniformly over all choices of the true parameters $\{\beta^*\}_{j=1}^k \in \mathcal{B}_{\text{vol}}(\pi_{\min}, \Delta, \kappa)$. We proceed by bounding the minimum singular values of *all* sub-matrices of Ξ of a certain size, and then show that each sub-matrix Ξ^j encountered over the course of the algorithm has a certain size with high probability.

In the following lemma, we use the shorthand $C_{\eta, \zeta, c_s} := 4c_s^2 \max\{9(\tilde{\eta})^2 \zeta, 1\}$.

Lemma 46. *Suppose that Assumption 20 holds, and that for a scalar $\alpha \in (0, 1)$, the sample size obeys the lower bound $n \geq \alpha^{-1} \max \left\{ 4d, \frac{d+1}{\zeta} \right\}$. Then we have*

$$\min_{S:|S|=\alpha n} \lambda_{\min} (\Xi_S^\top \Xi_S) \geq \frac{1}{C_{\eta,\zeta,c_s} \log C_{\eta,\zeta,c_s} + 2C\zeta^{-1} \log(e^2/\alpha)} \left(\frac{\alpha}{e^2} \right)^{2/\zeta} \cdot \alpha n$$

with probability greater than $1 - 3 \exp(-\alpha n)$.

We combine this lemma with a lower bound on the size of each subset.

Lemma 47. *Suppose that $n \geq 4 \frac{kd}{\pi_{\min}^2} \log(n/kd)$. Then we have*

$$\inf_{\beta_1^*, \dots, \beta_k^* \in \mathcal{B}_{\text{vol}}(\pi_{\min}, \Delta, \kappa)} \inf_{\substack{\beta_1, \dots, \beta_k \in \mathbb{R}^{d+1}: \\ v_{j,j'} \in \mathcal{B}_{v^*}^{j,j'}(r_b)}} \min_{j \in [k]} |S_j(\beta_1, \dots, \beta_k)| \geq n \cdot \frac{\pi_{\min}}{4}$$

with probability exceeding $1 - 2 \exp(-cn\pi_{\min}^2)$.

We are now ready to proceed to a proof of claim (7.11b). Note that the condition of the theorem guarantees that we have $n \geq \max \left\{ 4\alpha^{-1}d, \alpha^{-1} \frac{d+1}{\zeta}, 4 \frac{kd}{\pi_{\min}^2} \log(n/kd) \right\}$ provided $\alpha \geq c\pi_{\min}^2$.

Choosing $\alpha = \pi_{\min}/4$ and conditioning on the intersection of the pair of events guaranteed by Lemmas 46 and 47, we have

$$\begin{aligned} & \inf_{\beta_1^*, \dots, \beta_k^* \in \mathcal{B}_{\text{vol}}(\pi_{\min}, \Delta, \kappa)} \inf_{\substack{\beta_1, \dots, \beta_k \in \mathbb{R}^{d+1}: \\ v_{j,j'} \in \mathcal{B}_{v^*}^{j,j'}(r_b)}} \min_{j \in [k]} \lambda_{\min} (\Xi^j(\beta_1, \dots, \beta_k)^\top \Xi^j(\beta_1, \dots, \beta_k)) \\ & \geq \frac{1}{C_{\eta,\zeta,c_s} \log C_{\eta,\zeta,c_s} + 2C\zeta^{-1} \log(4e^2/\pi_{\min})} \left(\frac{\pi_{\min}}{4e^2} \right)^{2/\zeta} \cdot \frac{\pi_{\min} n}{4}. \end{aligned}$$

Combining this bound with the various high probability statements completes the proof of the claim. \square

Having proved the claims, we turn to proofs of the three technical lemmas.

Proof of Lemma 45

Consider a fixed pair (v, v^*) , and as before, let $\gamma_v = v - v^*$; we have

$$\begin{aligned} \mathbf{1} \{ \langle \xi_i, v \rangle \cdot \langle \xi_i, v^* \rangle \leq 0 \} \langle \xi_i, v^* \rangle^2 & \leq \mathbf{1} \{ \langle \xi_i, v \rangle \cdot \langle \xi_i, v^* \rangle \leq 0 \} \langle \xi_i, \gamma_v \rangle^2 \\ & \leq \mathbf{1} \{ \langle \xi_i, \gamma_v \rangle^2 \geq \langle \xi_i, v^* \rangle^2 \} \langle \xi_i, \gamma_v \rangle^2. \end{aligned}$$

Define the (random) set $K_v = \{i : \langle \xi_i, \gamma_v \rangle^2 \geq \langle \xi_i, v^* \rangle^2\}$, and note that this quantity implicitly depends on v^* as well. We have the bound

$$\frac{1}{n} \sum_{i=1}^n \mathbf{1} \{ \langle \xi_i, v \rangle \cdot \langle \xi_i, v^* \rangle \leq 0 \} \langle \xi_i, v^* \rangle^2 \leq \frac{1}{n} \|\Xi_{K_v} \gamma_v\|^2.$$

We now show that the quantity $\|\Xi_K \gamma_v\|^2$ is bounded as desired for all pairs (v^*, v) such that $v \in \mathcal{B}_{v^*}(r)$. Recall that $v^* = (u^* w^*)$, and let

$$m = \max \left\{ 4d, (d+1)/\zeta, C(1+\zeta^{-1}) \log n, n \left(Cr\tilde{\eta} \cdot \log \left(\frac{1}{r\tilde{\eta}} \right) \right)^\zeta \right\},$$

which implicitly depends on r . We claim that for all $0 \leq r \leq 1/24$, we have

$$\Pr \left\{ \sup_{\substack{v^* \in \mathbb{R}^{d+1} \\ v \in \mathcal{B}_{v^*}(r)}} |K_v| > m \right\} \leq \frac{3}{1 - (3/4)^\zeta} (3/4)^{\max\{d, 10 \log n\}} + cne^{-c'n}, \text{ and} \quad (7.12a)$$

$$\Pr \left\{ \bigcup_{\substack{T \subseteq [n]: \\ |T| \leq m}} \sup_{\omega \in \mathbb{R}^{d+1}} \frac{\|\Xi_T \omega\|^2}{\|\omega\|^2} \geq d + 4m\tilde{\eta}^2 \log(n/m) \right\} \leq \binom{n}{m}^{-1}. \quad (7.12b)$$

Taking these claims as given, the proof of the lemma is immediate, since $\frac{n}{m} \leq (Cr \log(1/r\tilde{\eta}))^{-\zeta}$, so that $\log(en/m) \leq C(1+\zeta) \log(1/r)$.

We now proceed to the proofs of the two claims, (7.12a) and (7.12b). The proof of claim (7.12b) follows from claim (7.12a) by using Lemma 40 in Appendix 6.9.5, so we dedicate the rest of the proof to establishing claim (7.12a).

Proof of claim (7.12a): By definition of the set K_v , we have

$$\begin{aligned} \Pr \left\{ \sup_{\substack{v^* \in \mathbb{R}^{d+1} \\ v \in \mathcal{B}_{v^*}(r)}} |K_v| > m \right\} &\leq \sum_{\substack{T \subseteq [n]: \\ |T| > m}} \Pr \left\{ \exists v^* \in \mathbb{R}^{d+1}, v \in \mathcal{B}_{v^*}(r) : \|\Xi_T \gamma_v\|^2 \geq \|\Xi_T v^*\|^2 \right\} \\ &= \sum_{\substack{T \subseteq [n]: \\ |T| > m}} \Pr \left\{ \exists v^* \in \mathbb{R}^{d+1}, v \in \mathcal{B}_{v^*}(r) : \frac{\|\gamma_v\|^2}{\|u^*\|^2} \frac{\|\Xi_T \gamma_v\|^2}{\|\gamma_v\|^2} \geq \frac{\|\Xi_T v^*\|^2}{\|u^*\|^2} \right\} \\ &\leq \sum_{\substack{T \subseteq [n]: \\ |T| > m}} \Pr \left\{ \exists v^* \in \mathbb{R}^{d+1}, v \in \mathcal{B}_{v^*}(r) : r^2 \frac{\|\Xi_T \gamma_v\|^2}{\|\gamma_v\|^2} \geq \frac{\|\Xi_T v^*\|^2}{\|u^*\|^2} \right\} \\ &\leq \sum_{\substack{T \subseteq [n]: \\ |T| > m}} \left(\Pr \left\{ \exists v^* \in \mathbb{R}^{d+1}, v \in \mathcal{B}_{v^*}(r) : \frac{\|\Xi_T \gamma_v\|^2}{\|\gamma_v\|^2} \geq |T| + \tilde{\eta}^2(\sqrt{d|T|} + d + |T|t_T) \right\} \right. \\ &\quad \left. + \Pr \left\{ \frac{\|\Xi_T v^*\|^2}{\|u^*\|^2} \leq r^2 \left\{ |T| + \tilde{\eta}^2(\sqrt{d|T|} + d + |T|t_T) \right\} \right\} \right) \end{aligned}$$

where the final step follows from the union bound and holds for all positive scalars $\{t_T\}_{T \subseteq n}$. For a fixed subset T of size $\ell \geq m$, we have the tail bounds

$$\Pr \left\{ \sup_{\omega \in \mathbb{R}^{d+1}} \frac{\|\Xi_T \omega\|^2}{\|\omega\|^2} \geq \ell + \tilde{\eta}^2(\sqrt{\ell d} + d + \ell t) \right\} \leq 2e^{-\ell \min\{t, t^2\}}, \text{ and} \quad (7.13a)$$

$$\Pr \left\{ \inf_{v^* \in \mathbb{R}^{d+1}} \frac{\|\Xi_T v^*\|^2}{\|u^*\|^2} \leq \delta \ell \right\} \leq 3(4c_s^2 \max\{9(\tilde{\eta})^2 \zeta, 1\} \delta \log(1/\delta))^{\ell \zeta/2}, \quad (7.13b)$$

where inequality (7.13a) follows from Lemma 40, and inequality (7.13b) from Lemma 53 since we have $m \geq \max\left\{\frac{d+1}{\zeta}, 4d\right\}$.

Now use the shorthand $C_{\eta, \zeta, c_s} := 4c_s^2 \max\{9(\tilde{\eta})^2 \zeta, 1\}$ as before. Substituting these bounds and letting $t_\ell \geq 1$ be a parameter to be chosen, we have

$$\begin{aligned} & \Pr\{|K_v| > m\} \\ & \leq \sum_{\ell=m+1}^n \binom{n}{\ell} \left[2e^{-\ell t_\ell} + 3(C_{\eta, \zeta, c_s} r^2 \cdot \frac{\ell + \tilde{\eta}^2(\sqrt{\ell d} + d + \ell t_\ell)}{\ell} \log(\frac{\ell}{r^2 \ell + r^2 \tilde{\eta}^2(\sqrt{\ell d} + d + \ell t_\ell)})) \right]^{\zeta \ell/2}. \end{aligned}$$

We now split the proof into two cases and choose t_ℓ differently for the two cases.

Case 1, $m \leq \ell < n/e$: Substituting the choice $t_\ell = 4 \log(n/\ell)$, we obtain

$$\begin{aligned} \left(\frac{\ell + \tilde{\eta}^2(\sqrt{\ell d} + d + \ell t_\ell)}{\ell} \right)^{1/2} & \leq \frac{\sqrt{\ell} + 5\tilde{\eta}(\sqrt{\ell \log(n/\ell)} + \sqrt{d})}{\sqrt{\ell}} \\ & \stackrel{(i)}{\leq} 1 + 6\tilde{\eta}\sqrt{\log(n/\ell)} \\ & \stackrel{(ii)}{\leq} 7\tilde{\eta} \cdot \sqrt{\log(n/\ell)}. \end{aligned}$$

where step (i) follows from the bound $m \geq d$, and step (ii) from the bound $\tilde{\eta} \geq 1$. Putting together the pieces and noting that the map $x \mapsto x \log(1/x)$ is an increasing on the interval $(0, e^{-1})$, we have, for each ℓ in this set, the bound

$$\Pr\{|K_v| = \ell\} \leq 2 \left(\frac{en}{\ell}\right)^{-c\ell} + 3 \binom{n}{\ell} \left(14C_{\eta, \zeta, c_s}^{1/2} r \tilde{\eta} \sqrt{\log(n/\ell) \cdot \log \frac{1}{7r\tilde{\eta}\sqrt{\log(n/\ell)}}} \right)^{\zeta \ell}.$$

For a sufficiently large constant C , and $\ell \geq n(Cr\tilde{\eta}\log(1/r\tilde{\eta}))^\zeta$, the second term is bounded by $3(3/4)^{\zeta \ell}$.

Case 2, $\ell \geq n/e$: The argument for this case is identical to before (with the choice $t_\ell = 2n/\ell$). Setting C'_{η,ζ,c_s} to be a sufficiently small η -dependent constant we obtain, for each ℓ in this range and for all $r \leq C'_{\eta,\zeta,c_s}$,

$$\left(3C_{\eta,\zeta,c_s} r^2 \frac{\ell + \tilde{\eta}^2(\sqrt{\ell d} + d + \ell t_\ell)}{\ell} \right) \leq 1/3,$$

where we also use the fact that $d \leq n/2$. Once again, using the non-decreasing nature of the map $x \mapsto x \log(1/x)$ on the interval $(0, e^{-1})$, we have

$$\begin{aligned} \binom{n}{\ell} & \left[2e^{-\ell t_\ell} + 3 \left(C_{\eta,\zeta,c_s} r^2 \cdot \frac{\ell + \tilde{\eta}^2(\sqrt{\ell d} + d + \ell t_\ell)}{\ell} \log \left(\frac{\ell}{r^2 \ell + r^2 \tilde{\eta}^2(\sqrt{\ell d} + d + \ell t_\ell)} \right) \right)^{\zeta \ell/2} \right] \\ & \leq 2e^{-cn} + 3 \left(\sqrt{\frac{1}{3}} \right)^{\zeta \ell}. \end{aligned}$$

Putting together the pieces from both cases, we have shown that for all $r \leq C'_{\eta,\zeta,c_s}$, we obtain

$$\begin{aligned} \Pr\{|K_v| > m\} & \leq 2ne^{-cn} + 3 \sum_{\ell=m+1}^n (3/4)^{\zeta \ell} \\ & \leq 2ne^{-cn} + np^{10\zeta^{-1} \log n}, \end{aligned}$$

where $p = (3/4)^\zeta$. This completes the proof of the claim. □

Proof of Lemma 46

The proof of this lemma follows from Lemma 53 in Appendix 6.9.5 in conjunction with the union bound. In particular, we have

$$\begin{aligned} \Pr \left\{ \min_{S:|S|=\alpha n} \lambda_{\min}(\Xi_S^\top \Xi_S) \leq \psi \alpha \cdot n \right\} & \leq 3 \binom{n}{\alpha n} (C_{\eta,\zeta,c_s} \psi \log(1/\psi))^{\zeta \alpha n} \\ & \leq 3 \left(e \frac{(C_{\eta,\zeta,c_s} \psi \log(1/\psi))^\zeta}{\alpha} \right)^{\alpha n}. \end{aligned}$$

Finally, setting $\psi = \frac{1}{C_{\eta,\zeta,c_s} \log C_{\eta,\zeta,c_s} + 2C\zeta^{-1} \log(e/\alpha)} \left(\frac{\alpha}{e^2} \right)^{2/\zeta}$ and performing some algebraic manipulation yields the claimed bound. □

Proof of Lemma 47

Note that we have $n \geq 4 \frac{kd}{\pi_{\min}^2} \log(n/kd)$. Noting that the S_j can be thought of as the indicator vector corresponding to the intersection of k halfspaces, applying [221, Lemmas 11 and 12] in conjunction

yields the bound

$$\inf_{\beta_1^*, \dots, \beta_k^* \in \mathcal{B}_{\text{vol}}(\pi_{\min}, \Delta, \kappa)} |S_j(\beta_1^*, \dots, \beta_k^*)| \geq n \cdot \frac{\pi_{\min}}{2}$$

with probability exceeding $1 - ce^{-c'n\pi_{\min}^2}$. Furthermore, we have

$$|S_j(\beta_1, \dots, \beta_k)| \geq |S_j(\beta_1^*, \dots, \beta_k^*)| - |S_j(\beta_1^*, \dots, \beta_k^*) \Delta S_j(\beta_1, \dots, \beta_k)|.$$

Using the notation W for a wedge and the notation n_W for the number of points in the wedge W , we have

$$\begin{aligned} |S_j(\beta_1^*, \dots, \beta_k^*) \Delta S_j(\beta_1, \dots, \beta_k)| &\stackrel{(i)}{\leq} \sum_{j'=1}^k n_{W(v_{j,j'}, v_{j,j'}^*)} \\ &\stackrel{(ii)}{\leq} 2 \sum_{j'=1}^k n \text{vol}(W(v_{j,j'}, v_{j,j'}^*)), \end{aligned}$$

where step (i) follows from [221, equation (33)], and step (ii) follows once again from [221, Lemma 11]. The volume of a wedge under a distribution satisfying Assumption 20 is upper bounded in Appendix 6.9.5. Applying Lemma 51, we have

$$\begin{aligned} |S_j(\beta_1, \dots, \beta_k)| &\geq n \cdot \frac{\pi_{\min}}{2} - \sum_{j'=1}^k \left(C''_{\eta, \zeta, c_s} \left(\frac{\|v_{j,j'} - v_{j,j'}^*\|}{\|\theta_j^* - \theta_{j'}^*\|} \right)^2 \log \left(\frac{\|\theta_j^* - \theta_{j'}^*\|}{\|v_{j,j'} - v_{j,j'}^*\|} \right) \right)^\zeta \\ &\stackrel{(iii)}{\geq} n \cdot \frac{\pi_{\min}}{2} - k \cdot n \frac{\pi_{\min}}{4k} \\ &\geq \frac{\pi_{\min} n}{4}, \end{aligned}$$

where in step (iii), we have used condition (7.9a) satisfied by all the parameters $v_{j,j'}$, by which we have

$$\frac{\|v_{j,j'} - v_{j,j'}^*\|}{\|\theta_j^* - \theta_{j'}^*\|} \log \left(\frac{\|\theta_j^* - \theta_{j'}^*\|}{\|v_{j,j'} - v_{j,j'}^*\|} \right) \leq C_{\eta, \zeta, c_s}^{(1)} \left(\frac{\pi_{\min}^{1+2\zeta^{-1}}}{k} \right)^{\zeta^{-1}}.$$

This further implies, for a sufficiently small constant $C_{\eta, \zeta, c_s}^{(1)}$, that

$$\frac{\|v_{j,j'} - v_{j,j'}^*\|}{\|\theta_j^* - \theta_{j'}^*\|} \log^{1/2} \left(\frac{\|\theta_j^* - \theta_{j'}^*\|}{\|v_{j,j'} - v_{j,j'}^*\|} \right) \leq \left(\frac{\pi_{\min}}{4k} \right)^{\zeta^{-1/2}}.$$

Since these steps held for an arbitrary index j , the proof of the lemma is complete. □

7.5 Proof of Corollary 3

Since the proof is more or less subsumed by Theorem 18, we only sketch the details. Let θ denote the current iterate and θ^+ denote the next iterate. Proceeding as before, we have the deterministic bound

$$\|X(\theta^+ - \theta^*)\|^2 \lesssim \sum_{i=1}^n \mathbf{1} \{ \langle x_i, \theta \rangle \langle x_i, \theta^* \rangle \leq 0 \} \langle x_i, \theta^* \rangle^2 + \|P_X \epsilon\|^2,$$

and so this case is even simpler than the max-affine setting, since we no longer select specific sub-matrices of X on which to invert a linear system. Standard bounds on sub-Gaussian random matrices yield

$$\Pr \left\{ \|X^\top X - n\|_{\text{op}} \geq \eta^2(d + 4\sqrt{\log n}) \right\} \leq 2n^{-8}. \quad (7.14)$$

By assumption, we have $n/\eta^2 \geq 2(d + 4\sqrt{\log n})$. Thus, applying Lemmas 45 and 3 (of [221]) along with the bound (7.14) to obtain that simultaneously for all pairs (θ^*, θ) satisfying $\|\theta - \theta^*\|/\|\theta^*\| \leq r \leq 1/24$, we have

$$\begin{aligned} \|\theta^+ - \theta^*\|^2 &\leq c_1 \left\{ \eta^2 \cdot \max \left\{ (1 + \zeta^{-1}) \frac{\max\{d, 10 \log n\}}{n}, (1 + \zeta) (\eta r)^\zeta \cdot \log^{\zeta+1} \left(\frac{1}{\eta r} \right) \right\} \right\} \\ &\quad \times \|\theta - \theta^*\|^2 + c_1 \sigma^2 \frac{d}{n} \log \left(\frac{n}{d} \right), \end{aligned}$$

with probability exceeding $1 - c_1 n^{-7}$. Choosing a small enough scalar r (depending on the tuple (η, ζ, c_s)) and a large enough n to make the quantity within the braces less than $3/4$ completes the proof.

7.6 Technical Lemmas and Background

7.6.1 Small ball properties and examples

We begin with a technical lemma taken from Rudelson and Vershynin [58, Corollary 1.4] that shows that product measures of small-ball distributions also satisfy the small-ball condition.

Lemma 48 ([58]). *For P_X satisfying the (ζ, c_s) -small-ball property (7.1) and $x_1, \dots, x_m \stackrel{i.i.d.}{\sim} P_X$, there is a universal constant C such that we have*

$$\sup_{\substack{u \in \mathbb{S}^{d-1} \\ w \in \mathbb{R}}} \Pr \left\{ \sum_{i=1}^m (\langle x_i, u \rangle + w)^2 \leq \delta m \right\} \leq (C \cdot c_s \delta)^{m\zeta} \text{ for all } \delta > 0. \quad (7.15)$$

We also verify that log-concave distributions and the standard Gaussian distribution satisfy the small-ball condition (7.1).

7.6.2 Log-concave distribution

The following result, taken from Carbery and Wright [227, Theorem 8], provides a small-ball bound for log-concave distributions almost directly.

Lemma 49 ([227]). *Let $p : \mathbb{R}^d \mapsto \mathbb{R}$ denote polynomial of degree (at most) ℓ , let X denote a d -dimensional random vector drawn from a log-concave distribution, and let $S_\delta = \{x \in \mathbb{R}^d : |p(x)| \leq \delta\}$ for each $\delta > 0$. Then for each $q > 0$, we have*

$$\Pr\{X \in S_\delta\} \cdot \left(\mathbb{E}_X [|p(X)|^{q/\ell}]\right)^{1/q} \leq Cq\delta^{1/\ell}.$$

For a unit norm vector $u \in \mathbb{S}^{d-1}$, setting $p(x) = (\langle x, u \rangle + w)^2$ and $q = 2$, we obtain

$$\Pr(|\langle X, u \rangle + w|^2 \leq \delta) \leq C \frac{1}{1+w^2} \delta^{1/2} \leq C\delta^{1/2},$$

where we have used the fact that $\mathbb{E}[\langle X, u \rangle^2] = 1$ since P_X is isotropic. Since this holds for each pair $(u, w) \in \mathbb{S}^{d-1} \times \mathbb{R}$, we have verified that X satisfies small ball condition with $(\zeta, c_s) = (1/2, C^2)$.

7.6.3 Standard Gaussian distribution

This is the canonical example of a sub-Gaussian distribution satisfying a small-ball condition. Suppose X is a standard Gaussian random variable. For a unit vector $u \in \mathbb{S}^{d-1}$, this implies that $\langle X, u \rangle^2 + w$ is a central χ^2 random variable with 1 degree of freedom.

Lemma 50. *Let Z_ℓ and Z'_ℓ denote central and non-central χ^2 random variables with ℓ degrees of freedom, respectively. Then for all $p \in [0, \ell]$, we have*

$$\Pr\{Z'_\ell \leq p\} \leq \Pr\{Z_\ell \leq p\} \leq \left(\frac{p}{\ell} \exp\left(1 - \frac{p}{\ell}\right)\right)^{\ell/2} = \exp\left(-\frac{\ell}{2} \left[\log \frac{\ell}{p} + \frac{p}{\ell} - 1\right]\right) \quad (7.16)$$

Applying the above lemma for $\ell = 1$ and $t = \delta$, we obtain

$$\Pr\{|\langle X, u \rangle + w|^2 \leq \delta\} \leq (e\delta)^{1/2}.$$

Hence the standard Gaussian satisfies the small-ball condition with $(\zeta, c_s) = (1/2, e)$.

For completeness, we provide a proof of Lemma 50 below.

Proof of Lemma 50

The fact that $Z'_\ell \stackrel{st.}{\leq} Z_\ell$ follows from standard results that guarantee that central χ^2 random variables stochastically dominate their non-central counterparts.

The tail bound is a simple consequence of the Chernoff bound. In particular, we have for all $\lambda > 0$ that

$$\begin{aligned} \Pr\{Z_\ell \leq p\} &= \Pr\{\exp(-\lambda Z_\ell) \geq \exp(-\lambda p)\} \\ &\leq \exp(\lambda p) \mathbb{E}[\exp(-\lambda Z_\ell)] \\ &= \exp(\lambda p)(1 + 2\lambda)^{-\frac{\ell}{2}}. \end{aligned} \tag{7.17}$$

where in the last step, we have used $\mathbb{E}[\exp(-\lambda Z_\ell)] = (1 + 2\lambda)^{-\frac{\ell}{2}}$, which is valid for all $\lambda > -1/2$. Minimizing the last expression over $\lambda > 0$ then yields the choice $\lambda^* = \frac{1}{2} \left(\frac{\ell}{p} - 1 \right)$, which is greater than 0 for all $0 \leq p \leq \ell$. Substituting this choice back into equation (7.17) proves the lemma. \square

7.6.4 Background and technical lemmas used in the proofs of Theorem 18

In this section, we collect statements and proofs of some technical lemmas used in the proofs of our results concerning the AM algorithm.

7.6.5 Bounds on the “volumes” of wedges in \mathbb{R}^d

For a pair of scalars (w, w') and d -dimensional vectors (u, u') , recall that we define the *wedge* formed by the $d + 1$ -dimensional vectors $v = (u, w)$ and $v' = (u', w')$ as the region

$$W(v, v') = \{x \in \mathbb{R}^d : (\langle x, u \rangle + w) \cdot (\langle x, u' \rangle + w') \leq 0\}.$$

Note that the wedge is a purely geometric object.

For any set $\mathcal{C} \subseteq \mathbb{R}^d$, let

$$\text{vol}_{P_X}(\mathcal{C}) = \Pr_{X \sim P_X} \{X \in \mathcal{C}\}$$

denote the volume of the set under the measure corresponding to the covariate distribution.

We now bound the volume of a wedge for distributions satisfying the small-ball condition. When the distribution P_X is η sub-Gaussian and satisfies a (ζ, c_s) -small-ball condition, we have the following result.

Lemma 51. *Suppose that Assumption 20 holds, and that for a pair of scalars (w, w') , d -dimensional vectors (u, u') , and $v = (u, w)$ and $v' = (u', w')$, we have $\frac{\|v - v'\|}{\|u\|} < 1/2$. Then, there is a positive constant C such that*

$$\text{vol}_{P_X}(W(v, v')) \leq C \left(C'_{\eta, \zeta, c_s} \frac{\|v - v'\|^2}{\|u\|^2} \log \left(\frac{\|u\|}{\|v - v'\|} \right) \right)^\zeta,$$

where C'_{η, ζ, c_s} is a constant that depends only on the tuple (η, ζ, c_s) .

Proof of Lemma 51

Using the notation $\xi = (x, 1) \in \mathbb{R}^{d+1}$ to denote the appended covariate, we have

$$\text{vol}(W(v, v')) = \Pr \{ \langle \xi, v \rangle \cdot \langle \xi, v' \rangle \leq 0 \},$$

where the probability is computed with respect to Gaussian measure.

In order to prove a bound on this probability, we begin by bounding the associated indicator random variable as

$$\begin{aligned} \mathbf{1} \{ \langle \xi, v \rangle \cdot \langle \xi, v' \rangle \leq 0 \} &= \mathbf{1} \{ \langle \xi, v' - v \rangle^2 \geq \langle \xi, v \rangle^2 \} \\ &\leq \mathbf{1} \{ \langle \xi, v' - v \rangle^2 \geq t \} + \mathbf{1} \{ \langle \xi, v \rangle^2 \leq t \}, \end{aligned} \quad (7.18)$$

where inequality (7.18) holds for all $t \geq 0$.

We now have

$$\Pr \{ \langle \xi, v \rangle^2 \leq t \} \leq \left(\frac{c_5 t}{\|u\|^2} \right)^\zeta,$$

since P_X satisfies a small-ball condition.

Furthermore, the sub-Gaussianity of the covariate distribution yields the upper tail bound

$$\begin{aligned} \Pr \{ \langle \xi, v' - v \rangle^2 \geq t \} &\leq c_1 \exp \left(- \frac{c_2}{\eta^2 \cdot \|u - u'\|^2} \left\{ \frac{t}{2} - \|v - v'\|^2 \right\} \right) \\ &\leq c_1 \exp \left(- \frac{c_2}{\eta^2 \cdot \|v - v'\|^2} \left\{ \frac{t}{2} - \|v - v'\|^2 \right\} \right). \end{aligned}$$

Putting all the pieces together, we obtain

$$\text{vol}_{P_X}(W(v, v')) \leq c_1 \exp \left(- \frac{c_2}{\eta^2 \cdot \|v - v'\|^2} \left\{ \frac{t}{2} - \|v - v'\|^2 \right\} \right) + \left(\frac{c_5 t}{\|u\|^2} \right)^\zeta.$$

Substituting $t = 2c\zeta\eta^2 \|v - v'\|^2 (\log(\|u\|/\|v - v'\|))$ yields the desired result. □

7.6.6 Uniform bounds on singular values of (sub-)matrices

We now state and prove two technical lemmas that bound the maximum and minimum singular values of a matrix whose rows are drawn from a sub-Gaussian distribution obeying the small-ball property. Our results on the minimum singular value are similar in spirit to those of Rudelson and Vershynin [228], but proved under a slightly different set of assumptions.

Lemma 52. *Suppose that the covariates are drawn i.i.d. from a η -sub-Gaussian distribution. Then for a fixed subset $S \in [n]$ of size ℓ and each $t \geq 0$, we have*

$$\Pr \left\{ \lambda_{\max}(\Xi_S^\top \Xi_S) \geq \ell + \tilde{\eta}^2(\sqrt{\ell d} + d + \ell t) \right\} \leq 2e^{-\ell \min\{t, t^2\}},$$

where $\tilde{\eta} = \max\{\eta, 1\}$.

The second lemma controls the minimum singular value of any sub-matrix of Ξ that has sufficiently many rows. Recall that distribution P_X is isotropic and η -sub-Gaussian, and satisfies the (ζ, c_s) small-ball condition (7.1).

Lemma 53. *Suppose that Assumption 20 holds, and that $\ell \geq \max \left\{ 4d, \frac{d+1}{\zeta} \right\}$. Let $\tilde{\eta} = \max\{\eta, 1\}$. Then for a fixed subset $S \in [n]$ of size ℓ and for each positive $\epsilon < \min\{(\tilde{\eta})^2\zeta, e^{-4/\zeta}\}$, we have*

$$\Pr \left\{ \lambda_{\min} \left(\Xi_S^\top \Xi_S \right) \leq \ell \epsilon \right\} \leq 3 \left(4c_s^2 \max \left\{ 9\tilde{\eta}^2\zeta, 1 \right\} \epsilon \log(1/\epsilon) \right)^{\ell\zeta/2}.$$

Proof of Lemma 52

Let $\{z_i\}_{i=1}^\ell$ denote i.i.d. Rademacher variables, and collect these in an ℓ -dimensional vector z . Let $D = \text{diag}(z)$ denote a diagonal matrix, and note that by unitary invariance of the singular values, the singular values of the matrix $\tilde{\Xi}_S = D\Xi_S$ are the same as those of Ξ_S .

By construction, the matrix $\tilde{\Xi}_S$ has i.i.d. rows, and the i -th row is given by $z_i(x_i, 1)$. For a $d + 1$ dimensional vector $\tilde{\lambda} = (\lambda, w)$ with $\lambda \in \mathbb{R}^d$ and $w \in \mathbb{R}$, we have

$$\begin{aligned} \mathbb{E} \left[\exp(\langle \tilde{\lambda}, z_i(x_i, 1) \rangle) \right] &= \frac{e^w}{2} \cdot \mathbb{E} [\exp(\langle \lambda, x_i \rangle)] + \frac{e^{-w}}{2} \cdot \mathbb{E} [\exp(-\langle \lambda, x_i \rangle)] \\ &= \exp(\|\lambda\|^2\eta^2/2) \cdot \frac{1}{2} (e^w + e^{-w}) \\ &\leq \exp(\|\lambda\|^2\eta^2/2) \cdot \exp(w^2/2) \leq \exp(\|\tilde{\lambda}\|^2\tilde{\eta}^2/2). \end{aligned}$$

where we have used the fact that x_i is zero-mean and η sub-Gaussian.

Since the rows of $\tilde{\Xi}_S$ are i.i.d., zero-mean, and $\tilde{\eta}$ -sub-Gaussian, applying [13, Theorem 6.2] immediately yields the lemma. □

Proof of Lemma 53

We let M denote the $\ell \times (d + 1)$ matrix Ξ_S . By the variational characterization of the minimum eigenvalue, we have

$$\lambda_{\min} (M^\top M) = \inf_{v \in \mathbb{S}^d} \|Mv\|^2 = \left(\inf_{v \in \mathbb{S}^d} \|Mv\| \right)^2.$$

Let $Z_v = \|Mv\|$ denote a random process indexed by v , and let $Z = \inf_{v \in \mathbb{S}^d} Z_v$; we are interested bounding the lower tail of Z . Consider a ρ -covering $\{v^1, \dots, v^N\}$ of the set \mathbb{S}^d in ℓ_2 norm, with $N \leq (1 + 2/\rho)^{d+1}$. Letting v^j be the closest element of the cover to v , we have

$$Z_v \geq Z_{v^j} - |Z_v - Z_{v^j}| \geq Z_{v^j} - \|M\|_{\text{op}} \cdot \rho,$$

so that we have the bound $Z \geq \min_{j \in [N]} Z_{v^j} - \|M\|_{\text{op}} \cdot \rho$. We have thus reduced the infimum over the unit shell to a finite minimum.

We thus have

$$\begin{aligned}
 \Pr \{ \lambda_{\min} (M^\top M) \leq \ell \epsilon \} &= \Pr \{ Z^2 \leq \ell \epsilon \} \\
 &= \Pr \{ Z \leq \sqrt{\ell \epsilon} \} \\
 &\leq \Pr \left\{ \min_{j \in [N]} Z_{vj} \leq 2\sqrt{\ell \epsilon} \right\} + \Pr \left\{ \|M\|_{\text{op}} \rho \geq \sqrt{\ell \epsilon} \right\} \\
 &\leq N \Pr \left\{ Z_{vj} \leq 2\sqrt{\ell \epsilon} \right\} + \Pr \left\{ \|M\|_{\text{op}} \geq \sqrt{\ell \epsilon} / \rho \right\},
 \end{aligned}$$

where we have used the union bound in each of the last two steps.

Now note that for each $j \in [N]$, the small ball condition yields the bound

$$\Pr \left\{ Z_{vj} \leq 2\sqrt{\ell \epsilon} \right\} = \Pr \left\{ \|Mv^j\|^2 \leq 2\ell \epsilon \right\} \leq (2C_{c_s} \epsilon)^{\ell \zeta},$$

where we have used Lemma 48 to reason about the product measure of small-ball distributions.

Furthermore, since M has $\tilde{\eta}$ sub-Gaussian rows, we may apply Lemma 52 to obtain

$$\Pr \left\{ \|M\|_{\text{op}} \geq \sqrt{\ell \epsilon} / \rho \right\} \leq 2 \exp \left(-\ell \left\{ \frac{\epsilon}{\rho^2 \tilde{\eta}^2} - \tilde{\eta}^{-2} - \frac{d + \sqrt{\ell d}}{\ell} \right\} \right),$$

which holds provided $\left\{ \frac{\epsilon}{\rho^2 \tilde{\eta}^2} - \tilde{\eta}^{-2} - \frac{d + \sqrt{\ell d}}{\ell} \right\} \geq 1$. When $\epsilon \leq e^{-4/\zeta}$, the choice $\rho = \rho_0 := \sqrt{\frac{\epsilon}{\zeta \tilde{\eta}^2}} \log^{-1}(1/\epsilon)$ ensures that

$$\frac{\epsilon}{\rho_0^2 \tilde{\eta}^2} - \tilde{\eta}^{-2} - \frac{d + \sqrt{\ell d}}{\ell} \stackrel{(i)}{\geq} \zeta \log(1/\epsilon) - 2 \geq \frac{\zeta}{2} \log(1/\epsilon),$$

where in step (i) we have used the properties $\tilde{\eta} \geq 1$ and $\ell \geq 4d$. This yields the bound

$$\Pr \left\{ \|M\|_{\text{op}} \geq \sqrt{\ell \epsilon} / \rho \right\} \leq 2\epsilon^{\ell \zeta / 2},$$

and putting together the pieces by substituting $N = \left(1 + \frac{2}{\rho_0}\right)^{d+1}$, we have

$$\Pr \left\{ \lambda_{\min} (M^\top M) \leq \ell \epsilon \right\} \leq \left(1 + \frac{2}{\rho_0}\right)^{d+1} (2c_s \epsilon)^{\ell \zeta} + 2\epsilon^{\ell \zeta / 2}.$$

Now note that we have $\epsilon < \min\{1, \tilde{\eta}^2 \zeta\}$, so that $\rho_0 \leq 1$, so that $1 + 2/\rho_0 \leq 3/\rho_0$. Therefore, we have

$$\begin{aligned}
 \Pr \left\{ \lambda_{\min} (M^\top M) \leq \ell \epsilon \right\} &\leq (9\tilde{\eta}^2 \zeta)^{(d+1)/2} (2c_s)^{\ell \zeta} \epsilon^{\ell \zeta - (d+1)/2} \log^{d+1}(1/\epsilon) + 2\epsilon^{\ell \zeta / 2} \\
 &\leq 3 (4c_s^2 \max\{9\tilde{\eta}^2 \zeta, 1\} \epsilon \log(1/\epsilon))^{\ell \zeta / 2},
 \end{aligned}$$

where we have used the fact that $\ell \zeta \geq d + 1$. This completes the proof. \square




Figure 7.3: AM with Rademacher and $\frac{1}{\sqrt{2.4}}$ (Bin(10, 0.4) – 4) covariates in log-log scale; the best-fit line to plot (a) has slope 1.94 and the best-fit line to plot (b) has slope 1.96, and hence the sample complexity scales linearly with k but quadratically with dimension d .

7.6.7 Numerical experiments: Noiseless Sample Complexity

In this section, we present some numerical experiments to illustrate the noiseless sample complexity of the AM algorithm under different covariate assumptions. We choose $\beta_i^* = e_i$ for $i \in [k]$, where e_i is the i -th canonical basis. Hence, $\|\beta_i^*\| = 1$ and $\pi_{\min} = 1/k$. We also choose the noiseless setting, i.e., $\sigma = 0$. We say β_i^* is recovered if $\|\beta_i^{(t)} - \beta_i^*\| \leq 0.01$. For a fixed dimension d , we run a linear search on the number of samples n , such that the empirical probability of success over 100 trials is more than 0.95, and output the smallest such n .

Here, we consider the distributions that do not satisfy the small ball condition. It is observed that for these distributions, the sample complexity for perfect parameter recovery has a quadratic dependence on dimension d . In particular, we draw the covariates i.i.d from Rademacher and $\frac{1}{\sqrt{2.4}}$ (Bin(10, 0.4) – 4) distributions. In Figure 7.3, we plot the noiseless sample complexity and dimension on a log scale. Since the slope of the (best-fit) lines in Figure 7.3 is very close to 2, we conclude that the number of samples required for perfect recovery of the parameters is linear in k , but quadratic in d .

Part III

Learning and Adaptation in Bandit and RL Framework

In this part of the thesis, we study the problem of adaptation (or model selection) in bandit and Reinforcement Learning framework.

- **Chapter 8:** We study the problem of adaptation/model selection for classical stochastic linear bandits.
- **Chapter 9:** We extend our study of model selection to generic contextual bandits beyond linear structure.
- **Chapter 10:** We study the problem of model selection for Reinforcement Learning with function approximation.

Chapter 8

Problem-Complexity Adaptive Model Selection for Stochastic Linear Bandits

We consider the problem of *model selection* for two popular stochastic linear bandit settings, and propose algorithms that adapt to the *unknown* problem complexity. In the first setting, we consider the K armed mixture bandits, where the mean reward of arm $i \in [K]$, is $\mu_i + \langle \alpha_{i,t}, \theta^* \rangle$, with $\alpha_{i,t} \in \mathbb{R}^d$ being the known context vector and $\mu_i \in [-1, 1]$ and θ^* are unknown parameters. We define $\|\theta^*\|$ as the problem complexity and we propose Adaptive Linear Bandit (ALB), a novel phase based algorithm that adapts to the true problem complexity, $\|\theta^*\|$. and achieves regret of $\tilde{O}(\|\theta^*\|\sqrt{T})$, where $\|\theta^*\|$ is a priori unknown. In the second setting, we consider the standard linear bandit problem (with possibly an infinite number of arms) where the sparsity of θ^* , denoted by $d^* \leq d$, is unknown to the algorithm. Defining d^* as the problem complexity (similar to [70]), we show that ALB achieves $\tilde{O}(d^*\sqrt{T})$ regret, matching that of an oracle who knew the true sparsity level. This methodology is then extended to the case of finitely many arms and similar results are proven. We further verify through synthetic and real-data experiments that the performance gains are fundamental and not artifacts of mathematical bounds. In particular, we show 1.5 – 3x drop in cumulative regret over non-adaptive algorithms.

8.1 Introduction

We consider two canonical settings for the stochastic MAB problem. The first is the K armed mixture MAB setting, in which the mean reward from any arm $i \in [K]$ is given by $\mu_i + \langle \theta^*, \alpha_{i,t} \rangle$, where $\alpha_{i,t} \in \mathbb{R}^d$ is the known context vector of arm i at time t , and $\mu_i \in \mathbb{R}$, $\theta^* \in \mathbb{R}^d$ are unknown and needs to be estimated. This setting also contains the standard MAB [62, 63] when $\theta^* = 0$. Popular linear bandit algorithms, like LinUCB, OFUL (see [64–66]) handle the case with no bias ($\mu_i = 0$), while OSOM [67], the recent improvement can handle arm-bias. Implicitly, all the above algorithms assume an upper bound on the norm of $\|\theta^*\| \leq L$, which is supplied as an input. Crucially however, the regret guarantees scale linearly in the upper bound L . In contrast, we choose $\|\theta^*\|$ as the problem complexity, and provide a novel phase based algorithm, that, without any upper

bound on the norm $\|\theta^*\|$, *adapts to the true complexity* of the problem instance, and achieves a regret scaling linearly in the true norm $\|\theta^*\|$. As a corollary, our algorithm’s performance matches the minimax regret of simple MAB when $\theta^* = 0$, even though the algorithm did not apriori know that $\theta^* = 0$. Formally, we consider a continuum of hypothesis classes, with each class positing a different upper bound on the norm $\|\theta^*\|$, where the complexity of a class is the upper bound posited. As our regret bound scales linearly in $\|\theta^*\|$ (the complexity of the smallest hypothesis class containing the instance) as opposed to an upper bound on $\|\theta^*\|$, our algorithm achieves model selection guarantees.

The second setting we consider is the standard linear stochastic bandit [66] with possibly an infinite number of arms, where the mean reward of any arm $x \in \mathbb{R}^d$ (arms are vectors in this case) given by $\langle x, \theta^* \rangle$, where $\theta^* \in \mathbb{R}^d$ is unknown. For this setting, we consider model selection from among a total of d different hypothesis classes, with each class positing a different cardinality for the support of θ^* . We exhibit a novel algorithm, where the regret scales linearly in the unknown cardinality of the support of θ^* . The regret scaling of our algorithm matches that of an oracle that has knowledge of the optimal support cardinality [68],[69], thereby achieving model selection guarantees. Our algorithm is the first known algorithm to obtain regret scaling matching that of an oracle that has knowledge of the true support. This is in contrast to standard linear bandit algorithms such as [66], where the regret scales linearly in d . We also extend this methodology to the case when the number of arms is finite and obtain similar regret rates matching the oracle. Model selection with dimension as a measure of complexity was also recently studied by [70], in which the classical contextual bandit [64] with a finite number of arms was considered. We clarify here that although our results for the finite arm setting yields a better (optimal) regret scaling with respect to the time horizon T and the support of θ^* (denoted by d^*), our guarantee depends on a problem dependent parameter and thus not uniform over all instances. In contrast, the results of [70], although sub-optimal in d^* and T , is uniform over all problem instances. Closing this gap is an interesting future direction.

8.1.1 Our Contributions

1. Successive Refinement Algorithms for Stochastic Linear Bandit - We present two novel epoch based algorithms, ALB (Adaptive Linear Bandit) - Norm and ALB - Dim, that achieve model selection guarantees for both families of hypothesis respectively. For the K armed mixture MAB setting, ALB-Norm, at the beginning of each phase, estimates an *upper bound* on the norm of $\|\theta^*\|$. Subsequently, the algorithm assumes this bound to be true during the phase, and the upper bound is re-estimated at the end of a phase. Similarly for the linear bandit setting, ALB-Dim estimates the support of θ^* at the beginning of each phase and subsequently only plays from this estimated support during the phase. In both settings, we show the estimates converge to the true underlying value —in the first case, the estimate of norm $\|\theta^*\|$ converges to the true norm, and in the second case, for all time after a random time with finite expectation, the estimated support equals the true support. Our algorithms are reminiscent of successive rejects algorithm [229] for standard MAB, with the crucial difference being that our algorithm is *non-monotone*. Once rejected, an arm is never pulled in the classical successive rejects. In contrast, our algorithm is successive

refinement and is not necessarily monotone —a hypothesis class discarded earlier can be considered at a later point of time.

2. Regret depending on the Complexity of the smallest Hypothesis Class - In the K armed mixture MAB setting, `ALB-NORM`'s regret scale as $\tilde{O}(\|\theta^*\|\sqrt{T})$, which is superior compared to state of art algorithms such as `OSOM` [67], whose regret scales as $\tilde{O}(L\sqrt{T})$, where L is an upper bound on $\|\theta^*\|$ that is supplied as an input. As a corollary, we get the ‘best of both worlds’ guarantee of [67], where if $\theta^* = 0$, our regret bound recovers known minimax regret guarantee of simple MAB. Similarly, for the linear bandit setting with unknown support, `ALB-DIM` achieves a regret of $\tilde{O}(d^*\sqrt{T})$, where $d^* \leq d$ is the true sparsity of θ^* . This matches the regret obtained by oracle algorithms that know of the true sparsity d^* [68, 69]. We also apply our methodology to the case when there is a finite number of arms and obtain similar regret scaling as the oracle. `ALB-DIM` is the first algorithm to obtain such model selection guarantees. Prior state of art algorithm `ModCB` for model selection with dimension as a measure of complexity was proposed in [70], with a finite set of arms, where the regret guarantee was sub-optimal compared to the oracle. However, our regret bounds for dimension, though matches the oracle, depends on the minimum non-zero coordinate value and is thus not uniform over θ^* . Obtaining regret rates in this case that matches the oracle and is uniform over all θ^* is an interesting future work.

3. Empirical Validation - We conduct synthetic and real data experiments that demonstrate superior performance of `ALB` compared to state of art methods such as `OSOM` [67] in the mixture K armed MAB setting and `OFUL` [66] in the linear bandit setting. We further observe, that the performance of `ALB` is close to that of the oracle algorithms that know the true complexity. This indicates that the performance gains from `ALB` is fundamental, and not artifacts of mathematical bounds.

Motivating Example: Our model selection framework is applicable to personalized news recommendation platforms, that recommend one of K news outlets, to each of its users. The recommendation decisions to any fixed user, can be modeled as an instance of a MAB; the arms are the K different news outlets, and the platforms recommendation decision (to this user) on day t is the arm played at time t . On each day t , each news outlet i reports a story, that can be modeled by the vectors $\alpha_{i,t}$, which can be obtained by embedding the stories into a fixed dimension vector space by some common embedding schemes. The reward obtained by the platform in recommending news outlet i to this user on day t can be modeled as $\mu_i + \langle \alpha_{i,t}, \theta^* \rangle$, where μ_i captures the preference of this user to news outlet i and the vector θ^* captures the “interest” of the user. Thus, if a channel i on day t , publishes a news article $\alpha_{i,t}$, that this user “likes”, then most likely the content $\alpha_{i,t}$ is “aligned” to θ^* and have a large inner product $\langle \alpha_{i,t}, \theta^* \rangle$. Different users on the platform however may have different biases and θ^* . Some users have strong preference towards certain topics and will read content written by any outlet on this topic (these users will have a large value of $\|\theta^*\|$). Other users may be agnostic to topics, but may prefer a particular news outlet a lot (for ex. some users like fox news exclusively or CNN exclusively, regardless of the topic). These users will have low $\|\theta^*\|$.

In such a multi-user recommendation application, we show that our algorithm `ALB-NORM` that tailors the model class for each user separately is more effective (lesser regret), than to, employ

a (non-adaptive) linear bandit algorithm for each user. We further show that our algorithms are also more effective than state of art model selection algorithms such as OSOM [67], which posits a ‘binary’ model - users either assign a 0 weight to topic or assign a potentially large weight to topic. Furthermore the heterogeneous complexity in this application can also be captured by the cardinality of the support of θ^* ; different people are interested in different sub-vectors of θ^* which the recommendation platform is not aware of a priori. In this context, our adaptive algorithm ALB-Dim that tailors to the interest of the individual user achieves better performance compared to non-adaptive linear bandit algorithms.

8.2 Related Work

Model selection for MAB are only recently being studied [230, 231], with [67], [70] being the closest to our work. OSOM was proposed in [67] for model selection in the K armed mixture MAB from two hypothesis classes —a “simple model” where $\|\theta^*\| = 0$, or a “complex model”, where $0 < \|\theta^*\| \leq L$. OSOM was shown to obtain a regret guarantee of $O(\log(T))$ when the instance is simple and $\tilde{O}(L\sqrt{T})$ otherwise. We refine this to consider a *continuum* of hypothesis classes and propose ALB-NORM, which achieves regret $\tilde{O}(\|\theta^*\|\sqrt{T})$, a superior guarantee (which we also empirically verify) compared to OSOM. Model selection with dimension as a measure of complexity was recently initiated in [70], where an algorithm ModCB was proposed. The setup considered in [70] was that of contextual bandits [64] with a fixed and finite number of arms. ModCB in this setting was shown to achieve a regret scaling that is sub-optimal compared to the oracle. In contrast, we consider the linear bandit setting with a continuum of arms [66], and ALB-Dim achieves a regret scaling matching that of an oracle. The continuum of arms allows ALB-Dim a finer exploration of arms, that enables it to learn the support of θ^* reliably and thus obtain regret matching that of the oracle. However, our regret bounds depend on the magnitude of the minimum non-zero value of θ^* and is thus not uniform over all β^* . Obtaining regret rates matching the oracle that holds uniformly over all θ^* is an interesting future work.

Corral was proposed in [230], by casting the optimal algorithm for each hypothesis class as an expert, with the forecaster’s performance having low regret with respect to the best expert (best model class). However, Corral can only handle finitely many hypothesis classes and is not suited to our setting with continuum hypothesis classes.

Adaptive algorithms for linear bandits have also been studied in different contexts from ours. The papers of [232, 233] consider problems where the arms have an unknown structure, and propose algorithms adapting to this structure to yield low regret. The paper [234] proposes an algorithm in the adversarial bandit setup that adapt to an unknown structure in the adversary’s loss sequence, to obtain low regret. The paper of [235] consider adaptive algorithms, when the distribution changes over time. In the context of online learning with full feedback, there have been several works addressing model selection [236–239]. In the context of statistical learning, model selection has a long line of work (for eg. [240], [241], [242], [243], [244] [245]). However, the bandit feedback in our setups is much more challenging and a straightforward adaptation of algorithms developed for either statistical learning or full information to the setting with bandit feedback is not feasible.

8.3 Norm as a measure of Complexity

8.3.1 Problem Formulation

In this section, we formally define the problem. At each round $t \in [T]$, the player chooses one of the K available arms. Each arm has a context $\{\alpha_{i,t} \in \mathbb{R}^d\}_{i=1}^K$ that changes over time t . Similar to the standard stochastic contextual bandit framework, the context vectors for each arm is chosen independently of all other arms and of the past time instances.

We assume that there exists an underlying parameter $\theta^* \in \mathbb{R}^d$ and biases $\{\mu_1, \dots, \mu_K\}$ each taking value in $[-1, 1]$ such that the mean reward of an arm is a linear function of the context of the arm. The reward for playing arm i at time t is given by, $g_{i,t} = \mu_i + \langle \alpha_{i,t}, \theta^* \rangle + \eta_{i,t}$, where $\{\eta_{i,t}\}_{t=1}^T$ are i.i.d zero mean and σ sub-Gaussian noise. The context vector satisfies

$$\mathbb{E}[\alpha_{i,t} | \{\alpha_{j,s}, \eta_{j,s}\}_{j \in [K], s \in [t-1]}] = 0,$$

and

$$\mathbb{E}[\alpha_{i,t} \alpha_{i,t}^\top | \{\alpha_{j,s}, \eta_{j,s}\}_{j \in [K], s \in [t-1]}] \succcurlyeq \rho_{\min} I.$$

The above setting is popularly known as stochastic contextual bandit [67]. In the special case of $\theta^* = 0$, the above model reduces to $g_{i,t} = \mu_i + \eta_{i,t}$. Note that in this setting, the mean reward of arms are fixed, and not dependent on the context. Hence, this corresponds to a simple *multi-armed bandit* setup and standard algorithms (like UCB [63]) can be used as a learning rule. At round t , we define $i_t^* = \operatorname{argmax}_{i \in [K]} [\mu_i + \langle \theta^*, \alpha_{i,t} \rangle]$ as the best arm. Also let an algorithm play arm A_t at round t . The regret of the algorithm upto time T is given by,

$$R(T) = \sum_{s=1}^T [\mu_{i_s^*} + \langle \theta^*, \alpha_{i_s^*, s} \rangle - \mu_{A_s} - \langle \theta^*, \alpha_{A_s, s} \rangle].$$

Throughout the paper, we use $C, C_1, \dots, c, c_1, \dots$ to denote positive universal constants, the value of which may differ in different instances.

We define a new notion of complexity for stochastic linear bandits; and propose an algorithm that adapts to it. We define $\|\theta^*\|$ as the problem complexity for the linear bandit instance. Note that if $\|\theta^*\| = 0$, the linear bandit model reduces to the simple multi-armed bandit setting. Furthermore, the cumulative regret $R(T)$ of linear bandit algorithms (like OFUL [66] and OSOM [67]) scales linearly with $\|\theta^*\|$ ([67]). Hence, $\|\theta^*\|$ constitutes a natural notion of model complexity. In Algorithm 11, we propose an adaptive scheme which adapts to the true complexity of the problem, $\|\theta^*\|$. Instead of assuming an upper-bound on $\|\theta^*\|$, we use an initial exploration phase to obtain a rough estimate of $\|\theta^*\|$ and then successively refine it over multiple epochs. The cumulative regret of our proposed algorithm actually scales linearly with $\|\theta^*\|$.

8.3.2 Adaptive Linear Bandit (norm)—(ALB-Norm algorithm)

We present the adaptive scheme in Algorithm 11. Note that Algorithm 11 depends on the subroutine OFUL⁺. Observe that at each iteration, we estimate the bias $\{\mu_1, \dots, \mu_K\}$ and θ^* separately. The

Algorithm 11: Adaptive Linear Bandit (Norm)

- 1: **Input:** Initial exploration period τ , the phase length T_1 , $\delta_1 > 0$, $\delta_s > 0$.
 - 2: Select an arm at random, sample rewards 2τ times
 - 3: Obtain initial estimate (b_1) of $\|\theta^*\|$ according to Section 8.3.3
 - 4: **for** $t = 1, 2, \dots, K$ **do**
 - 5: Play arm t , receive reward $g_{t,t}$
 - 6: **end for**
 - 7: Define $\mathcal{S} = \{g_{i,i}\}_{i=1}^K$
 - 8: **for** epochs $i = 1, 2, \dots, N$ **do**
 - 9: Use \mathcal{S} as pure-exploration reward
 - 10: Play $\text{OFUL}_{\delta_i}^+(b_i)$ until the end of epoch i (denoted by \mathcal{E}_i)
 - 11: At $t = \mathcal{E}_i$, refine estimate of $\|\theta^*\|$ as, $b_{i+1} = \max_{\theta \in \mathcal{C}_{\mathcal{E}_i}} \|\theta\|$
 - 12: Set $T_{i+1} = 2T_i$, $\delta_{i+1} = \frac{\delta_i}{2}$.
 - 13: **end for**
 - 14: $\text{OFUL}_{\delta}^+(b)$:
 - 15: **Input:** Parameters $b, \delta > 0$, number of rounds \tilde{T}
 - 16: **for** $t = 1, 2, \dots, \tilde{T}$ **do**
 - 17: Select the best arm estimate as $j_t = \operatorname{argmax}_{i \in [K]} [\max_{\theta \in \mathcal{C}_{t-1}} \{\tilde{\mu}_{i,t-1} + \langle \alpha_{i,t}, \theta \rangle\}]$,
where $\tilde{\mu}_{i,t}$ and \mathcal{C}_t are given in Section 8.3.2.
 - 18: Play arm j_t , and update $\{\tilde{\mu}_{i,t}\}_{i=1}^K$ and \mathcal{C}_t
 - 19: **end for**
-

estimation of the bias involves a simple sample mean estimate with upper confidence level, and the estimation of θ^* involves building a confidence set that shrinks over time.

In order to estimate θ^* , we use a variant of the popular OFUL [66] algorithm with arm bias. We refer to the algorithm as OFUL^+ . Algorithm 11 is epoch based, and over multiple epochs, we successively refine the estimate of $\|\theta^*\|$. We start with a rough over-estimate of $\|\theta^*\|$ (obtained from a pure exploration phase), and based on the confidence set constructed at the end of the epoch, we update the estimate of $\|\theta^*\|$. We argue that this approach indeed correctly estimates $\|\theta^*\|$ with high probability over a sufficiently large time horizon T .

We now discuss the algorithm OFUL^+ . A variation of this was proposed in [67] in the context of model selection between linear and standard multi-armed bandits. We use $\tilde{\mu}_{i,t}$ to address the bias term, which we define shortly. The parameters b and δ are used in the construction of the confidence set \mathcal{C}_t . Suppose OFUL^+ is run for a total of \tilde{T} rounds and plays arm A_s at time s . Let $T_i(t)$ be the number of times OFUL^+ plays arm i until time t . Also, let b be the current estimate of $\|\theta^*\|$. We define,

$$\bar{g}_{i,t} = \frac{1}{T_i(t)} \sum_{s=1}^t g_{i,s} \mathbf{1}\{A_s = i\}.$$

With this, we have ¹

$$\tilde{\mu}_{i,t} = \bar{g}_{i,t} + c(\sigma + b) \sqrt{\frac{d}{T_i(t)} \log\left(\frac{1}{\delta}\right)}.$$

The confidence interval \mathcal{C}_t , is defined as

$$\mathcal{C}_t = \{\theta \in \mathbb{R}^d : \|\theta - \hat{\theta}_t\| \leq \mathcal{K}_\delta(b, t, \tilde{T})\},$$

where $\hat{\theta}_t$ is the least squares estimate defined as

$$\hat{\theta}_t = (\alpha_{K+1:t}^\top \alpha_{K+1:t} + I)^{-1} \alpha_{K+1:t}^\top G_{K+1:t}$$

with $\alpha_{K+1:t}$ as a matrix having rows $\alpha_{A_{K+1}, K+1}^\top, \dots, \alpha_{A_t, t}^\top$ and $G_{K+1:t} = [g_{A_{K+1}, K+1} - \tilde{\mu}_{A_{K+1}, K+1}, \dots, g_{A_t, t} - \tilde{\mu}_{A_t, t}]^\top$. The radius of \mathcal{C}_t is given by (see Section 8.8.1 for complete expression),

$$\mathcal{K}_\delta(b, t, \tilde{T}) = c \frac{(\sigma\sqrt{d} + b)}{\rho_{\min}\sqrt{t}} \sqrt{\log(K\tilde{T}/\delta)}.$$

Lemma 2 of [67] shows that $\theta^* \in \mathcal{C}_t$ with probability ² at least $1 - 4\delta$.

8.3.3 Construction of initial estimate b_1

We select an arm at random (without loss of generality, assume that this is arm 1), and sample rewards (in an i.i.d fashion) for 2τ times, where $\tau > 0$ is a parameter to be fed to the Algorithm 11. In order to kill the bias of arm 1, we take pairwise differences and form: $y(1) = g_{1,1} - g_{1,2}$, $y(2) = g_{1,3} - g_{1,4}$ and so on. Augmenting $y(\cdot)$, we obtain: $Y = \tilde{X}\theta^* + \tilde{\eta}$, where the i -th row of \tilde{X} is $(\alpha_{1,2i+1} - \alpha_{1,2i+2})^\top$, the i -th element of $\tilde{\eta}$ is $\eta_{1,2i+1} - \eta_{1,2i+2}$. Hence, the least squares estimate, $\hat{\theta}^{(\ell_s)}$ satisfies $\|\hat{\theta}^{(\ell_s)} - \theta^*\| \leq \sqrt{2}\sigma\sqrt{\frac{d}{\tau}\log(1/\delta_s)}$, with probability exceeding $1 - \delta_s$ ([13]). We set the initial estimate

$$b_1 = \max\{\|\hat{\theta}^{(\ell_s)}\| + \sqrt{2}\sigma\sqrt{\frac{d}{\tau}\log(1/\delta_s)}, 1\}$$

and this satisfies $b_1 \geq \|\theta^*\|$ and $b_1 \geq 1$ with probability at least $1 - \delta_s$.

8.3.4 Regret Guarantee of Algorithm 11

We now obtain an upper bound on the cumulative $R(T)$ with Algorithm 11 with high probability. For theoretical tractability, we assume that OFUL⁺ restarts at the start of each epoch. We have the following lemma regarding the sequence $\{b_i\}_{i=1}^\infty$ of estimates of $\|\theta^*\|$:

¹For complete expression, see Section 8.8.1

²There is a typo in the proof of regret in [67]. We correct the typo, and modify the definition of $\tilde{\mu}_{i,t}$ and $\mathcal{K}_\delta(b, t, \tilde{T})$. As a consequence, the high probability bounds change a little.

Lemma 54. *With probability exceeding $1 - 8\delta_1 - \delta_s$, the sequence $\{b_i\}_{i=1}^\infty$ converges to $\|\theta^*\|$ at a rate $\mathcal{O}(\frac{i}{2^i})$, and we obtain $b_i \leq (c_1\|\theta^*\| + c_2)$ for all i , provided $T_1 \geq C_1 (\max\{p, q\} b_1)^2 d$, where $C_1 > 9$, and $p = \lceil \frac{14 \log(\frac{2KT_1}{\delta_1})}{\sqrt{\rho_{\min}}} \rceil$, $q = \lceil \frac{2C\sigma \log(\frac{2KT_1}{\delta_1})}{\sqrt{\rho_{\min}}} \rceil$.*

Hence, the sequence converges to $\|\theta^*\|$ at an exponential rate. We have the following guarantee on the cumulative regret $R(T)$:

Theorem 20. *Suppose $T_1 > \max\{T_{\min}(\delta, T), C_1 (\max\{p, q\} b_1)^2 d\}$, where $C_1 > 9$ and $T_{\min}(\delta, T) = (\frac{16}{\rho_{\min}^2} + \frac{8}{3\rho_{\min}}) \log(\frac{2dT}{\delta})$. Then, with probability at least $1 - 18\delta_1 - \delta_s$, we have*

$$R(T) \leq C_1(2\tau + K)\|\theta^*\| + C(\|\theta^*\| + 1)(\sqrt{K} + \sqrt{d})\sqrt{T} \log(KT_1/\delta_1) \log(T/T_1).$$

Remark 32. Note that the regret bound depends on the problem complexity $\|\theta^*\|$, and we prove that Algorithm 11 adapts to this complexity. Ignoring the log factors, Algorithm 11 has a regret of $\tilde{\mathcal{O}}((1 + \|\theta^*\|)(\sqrt{K} + \sqrt{d})\sqrt{T})$ with high probability.

Remark 33. (Matches Linear Bandit algorithm) Note that the above bound matches the regret guarantee of the linear bandit algorithm with bias as presented in [67].

Remark 34. (Matches UCB when $\theta^* = 0$) When $\theta^* = 0$ (the simplest model, without any contextual information), Algorithm 1 recovers the minimax regret of UCB algorithm. Indeed, substituting $\|\theta^*\| = 0$ in the above regret bound yields $R(T) = \mathcal{O}(\sqrt{KT})$, with high probability, provided $K > d$. Hence, we obtain the “best of both worlds” results with simple model ($\theta^* = 0$) and contextual bandit model ($\theta^* \neq 0$).

8.4 Dimension as a Measure of Complexity - Continuum Armed Setting

In this section, we consider the standard stochastic linear bandit model in d dimensions [66], with the dimension as a measure of complexity. The setup in this section is almost identical to that in Section 8.3.1, with the 0 arm biases and a continuum collection of arms denoted by the set $\mathcal{A} := \{x \in \mathbb{R}^d : \|x\| \leq 1\}$ ³ Thus, the mean reward from any arm $x \in \mathcal{A}$ is $\langle x, \theta^* \rangle$, where $\|\theta^*\| \leq 1$. We assume that θ^* is $d^* \leq d$ sparse, where d^* is apriori unknown to the algorithm. Thus, unlike in Section 8.3, there is no i.i.d. context sampling in this section. We consider a sequence of d nested hypothesis classes, where each hypothesis class $i \leq d$, models θ^* as a i sparse vector. The goal of the forecaster is to minimize the regret, namely $R(T) := \sum_{t=1}^T [\langle x_t^* - x_t, \theta^* \rangle]$, where at any time t , x_t is the action recommended by an algorithm and $x_t^* = \operatorname{argmax}_{x \in \mathcal{A}} \langle x, \theta^* \rangle$. The regret $R(T)$ measures the loss in reward of the forecaster with that of an oracle that knows θ^* and thus can compute x_t^* at each time.

³Our algorithm can be applied to any compact set $\mathcal{A} \subset \mathbb{R}^d$, including the finite set as shown in Section 8.8.6.

Algorithm 12: Adaptive Linear Bandit (Dimension)

- 1: **Input:** Initial Phase length T_0 and slack $\delta > 0$.
 - 2: $\hat{\theta}_0 = \mathbf{1}, T_{-1} = 0$
 - 3: **for** Each epoch $i \in \{0, 1, 2, \dots\}$ **do**
 - 4: $T_i = 25^i T_0, \quad \varepsilon_i \leftarrow \frac{1}{2^i}, \quad \delta_i \leftarrow \frac{\delta}{2^i}$
 - 5: $\mathcal{D}_i := \{i : |\hat{\theta}_i| \geq \frac{\varepsilon_i}{2}\}$
 - 6: **for** Times $t \in \{T_{i-1} + 1, \dots, T_i\}$ **do**
 - 7: Play OFUL(1, δ_i) only restricted to coordinates in \mathcal{D}_i . Here δ_i is the probability slack parameter and 1 represents $\|\theta^*\| \leq 1$.
 - 8: **end for**
 - 9: **for** Times $t \in \{T_i + 1, \dots, T_i + 5^i \sqrt{T_0}\}$ **do**
 - 10: Play an arm from the action set \mathcal{A} chosen uniformly and independently at random.
 - 11: **end for**
 - 12: $\alpha_i \in \mathbb{R}^{S_i \times d}$ with each row being the arm played during all random explorations in the past.
 - 13: $\mathbf{y}_i \in \mathbb{R}^{S_i}$ with i -th entry being the observed reward at the i -th random exploration in the past
 - 14: $\hat{\theta}_{i+1} \leftarrow (\alpha_i^T \alpha_i)^{-1} \alpha_i \mathbf{y}_i$, is a d dimensional vector
 - 15: **end for**
-

8.4.1 ALB-Dim Algorithm

The algorithm is parametrized by $T_0 \in \mathbb{N}$, which is given in Equation (8.1) in the sequel and slack $\delta \in (0, 1)$. As in the previous case, ALB-Dim proceeds in phases numbered $0, 1, \dots$ which are non-decreasing with time. At the beginning of each phase, ALB-Dim makes an estimate of the set of non-zero coordinates of θ^* , which is kept fixed throughout the phase. Concretely, each phase i is divided into two blocks - (i) a regret minimization block lasting $25^i T_0$ time slots, (ii) followed by a random exploration phase lasting $5^i \lceil \sqrt{T_0} \rceil$ time slots. Thus, each phase i lasts for a total of $25^i T_0 + 5^i \lceil \sqrt{T_0} \rceil$ time slots. At the beginning of each phase $i \geq 0$, $\mathcal{D}_i \subseteq [d]$ denotes the set of ‘active coordinates’, namely the estimate of the non-zero coordinates of θ^* . Subsequently, in the regret minimization block of phase i , a fresh instance of OFUL [66] is spawned, with the dimensions restricted only to the set \mathcal{D}_i and probability parameter $\delta_i := \frac{\delta}{2^i}$. In the random exploration phase, at each time, one of the possible arms from the set \mathcal{A} is played chosen uniformly and independently at random. At the end of each phase $i \geq 0$, ALB-Dim forms an estimate $\hat{\theta}_{i+1}$ of θ^* , by solving a least squares problem using all the random exploration samples collected till the end of phase i . The active coordinate set \mathcal{D}_{i+1} , is then the coordinates of $\hat{\theta}_{i+1}$ with magnitude exceeding $2^{-(i+1)}$. The pseudo-code is provided in Algorithm 12, where, $\forall i \geq 0$, S_i in lines 15 and 16 is the total number of random-exploration samples in all phases upto and including i .

8.4.2 Main Result

We first specify, how to set the input parameter T_0 , as function of δ . For any $N \geq d$, denote by A_N to be the $N \times d$ random matrix with each row being a vector sampled uniformly and independently from the unit sphere in d dimensions. Denote by $M_N := \frac{1}{N} \mathbb{E}[A_N^T A_N]$, and by $\lambda_{max}^{(N)}, \lambda_{min}^{(N)}$, to be the largest and smallest eigenvalues of M_N . Observe that as M_N is positive semi-definite ($0 \leq \lambda_{min}^{(N)} \leq \lambda_{max}^{(N)}$) and almost-surely full rank, i.e., $\mathbb{P}[\lambda_{min}^{(N)} > 0] = 1$. The constant T_0 is the smallest integer such that

$$\sqrt{T_0} \geq \max \left(\frac{32\sigma^2}{(\lambda_{min}^{(\lceil \sqrt{T_0} \rceil)})^2} \ln(2d/\delta), \frac{4}{3} \frac{(6\lambda_{max}^{(\lceil \sqrt{T_0} \rceil)} + \lambda_{min}^{(\lceil \sqrt{T_0} \rceil)})(d + \lambda_{max}^{(\lceil \sqrt{T_0} \rceil)})}{(\lambda_{min}^{(\lceil \sqrt{T_0} \rceil)})^2} \ln(2d/\delta) \right) \quad (8.1)$$

Remark 35. T_0 in Equation (8.1) is chosen such that, at the end of phase 0, $\mathbb{P}[||\hat{\theta}_0 - \theta^*||_\infty \geq 1/2] \leq \delta$.

A formal statement of the Remark is provided in Lemma 55 in Appendix 8.8.2.

Theorem 21. *Suppose Algorithm 12 is run with input parameters $\delta \in (0, 1)$, and T_0 as given in Equation (8.1), then with probability at-least $1 - \delta$, the regret after a total of T arm-pulls satisfies*

$$R_T \leq \frac{50}{\gamma^{4.65}} T_0 + 25\sqrt{T} \left[1 + 4\sqrt{d^* \ln\left(1 + \frac{25T}{d^*}\right)} \left(1 + \sigma \sqrt{2 \ln\left(\frac{T}{T_0\delta}\right) + d^* \ln\left(1 + \frac{25T}{d^*}\right)} \right) \right].$$

The parameter $\gamma > 0$ is the minimum magnitude of the non-zero coordinate of θ^* , i.e., $\gamma = \min\{|\theta_i^*| : \theta_i^* \neq 0\}$ and d^* the sparsity of θ^* , i.e., $d^* = |\{i : \theta_i^* \neq 0\}|$.

In order to parse this result, we give the following corollary.

Corollary 4. *Suppose Algorithm 12 is run with input $\delta \in (0, 1)$, and $T_0 = \tilde{O}\left(d^2 \ln^2\left(\frac{1}{\delta}\right)\right)$ given in Equation (8.1), then with probability at-least $1 - \delta$, the regret after T times satisfies*

$$R_T \leq O\left(\frac{d^2}{\gamma^{4.65}} \ln^2(d/\delta)\right) + \tilde{O}(d^* \sqrt{T}).$$

Remark 36. The constants in the Theorem are not optimized. In particular, the exponent of γ can be made arbitrarily close to 4, by setting $\varepsilon_i = C^{-i}$ in Line 4 of Algorithm 12, for some appropriately large constant $C > 1$, and increasing $T_i = (C')^i T_0$, for appropriately large C' ($C' \approx C^4$).

Discussion - The regret of an oracle algorithm that knows the true complexity d^* scales as $\tilde{O}(d^* \sqrt{T})$ [68, 69], matching ALB-Dim's regret, upto an additive constant independent of time. ALB-Dim is the first algorithm to achieve such model selection guarantees. On the other hand, standard linear bandit algorithms such as OFUL achieve a regret scaling $\tilde{O}(d\sqrt{T})$, which is much larger compared to that of ALB-Dim, especially when $d^* \ll d$, and γ is a constant. Numerical simulations further confirms this deduction, thereby indicating that our improvements are fundamental and not from mathematical bounds. Corollary 4 also indicates that ALB-Dim has higher regret if γ is lower. A

small value of γ makes it harder to distinguish a non-zero coordinate from a zero coordinate, which is reflected in the regret scaling. Nevertheless, this only affects the *second order term as a constant*, and the dominant scaling term only depends on the true complexity d^* , and not on the underlying dimension d . However, the regret guarantee is not uniform over all θ^* as it depends on γ . Obtaining regret rates matching the oracles and that hold uniformly over all θ^* is an interesting avenue of future work.

8.5 Dimension as a Measure of Complexity - Finite Armed Setting

8.5.1 Problem Setup

In this section, we consider the model selection problem for the setting with finitely many arms in the framework studied in [70]. At each time $t \in [T]$, the forecaster is shown a context $X_t \in \mathcal{X}$, where \mathcal{X} is some arbitrary ‘feature space’. The set of contexts $(X_t)_{t=1}^T$ are i.i.d. with $X_t \sim \mathcal{D}$, a probability distribution over \mathcal{X} that is known to the forecaster. Subsequently, the forecaster chooses an action $A_t \in \mathcal{A}$, where the set $\mathcal{A} := \{1, \dots, K\}$ are the K possible actions chosen by the forecaster. The forecaster then receives a reward $Y_t := \langle \theta^*, \phi^M(X_t, A_t) \rangle + \eta_t$. Here $(\eta_t)_{t=1}^T$ is an i.i.d. sequence of 0 mean sub-gaussian random variables with sub-gaussian parameter σ^2 that is known to the forecaster. The function⁴ $\phi^M : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}^d$ is a known feature map, and $\theta^* \in \mathbb{R}^d$ is an unknown vector. The goal of the forecaster is to minimize its regret, namely $R(T) := \sum_{t=1}^T \mathbb{E}[\langle A_t^* - A_t, \theta^* \rangle]$, where at any time t , conditional on the context X_t , $A_t^* \in \arg\max_{a \in \mathcal{A}} \langle \theta^*, \phi^M(X_t, a) \rangle$. Thus, A_t^* is a random variable as X_t is random.

To describe the model selection, we consider a sequence of M dimensions $1 \leq d_1 < d_2, \dots < d_M := d$ and an associated set of feature maps $(\phi^m)_{m=1}^M$, where for any $m \in [M]$, $\phi^m(\cdot, \cdot) : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}^{d_m}$, is a feature map embedding into d_m dimensions. Moreover, these feature maps are nested, namely, for all $m \in [M - 1]$, for all $x \in \mathcal{X}$ and $a \in \mathcal{A}$, the first d_m coordinates of $\phi^{m+1}(x, a)$ equals $\phi^m(x, a)$. The forecaster is assumed to have knowledge of these feature maps. The unknown vector θ^* is such that its first d_{m^*} coordinates are non-zero, while the rest are 0. The forecaster does not know the true dimension d_{m^*} . If this were known, than standard contextual bandit algorithms such as LinUCB [64] can guarantee a regret scaling as $\tilde{O}(\sqrt{d_{m^*}T})$. In this section, we provide an algorithm in which, even when the forecaster is unaware of d_{m^*} , the regret scales as $O(\sqrt{d_{m^*}T})$. However, this result is non uniform over all θ^* as, we will show, depends on the minimum non-zero coordinate value in θ^* .

Model Assumptions We will require some assumptions identical to the ones stated in [70]. Let $\|\theta^*\|_2 \leq 1$, which is known to the forecaster. The distribution \mathcal{D} is assumed to be known to the forecaster. Associated with the distribution \mathcal{D} is a matrix $\Sigma_M := \frac{1}{K} \sum_{a \in \mathcal{A}} \mathbb{E}[\phi^M(x, a)\phi^M(x, a)^T]$ (where $x \sim \mathcal{D}$), where we assume its minimum eigen value $\lambda_{\min}(\Sigma_M) > 0$ is strictly positive.

⁴Superscript M will become clear shortly

Further, we assume that, for all $a \in \mathcal{A}$, the random variable $\phi^M(x, a)$ (where $x \sim \mathcal{D}$ is random) is a sub-gaussian random variable with (known) parameter τ^2 .

8.5.2 ALB-Dim Algorithm

The algorithm in this case is identical to that of Algorithm 12, except with the difference that in place of OFUL, we use `SupLinRel` of [246] as the black-box. The full details of the Algorithm are provided in Appendix 8.8.6.

8.5.3 Main Result

For brevity, we only state the Corollary of our main Theorem (Theorem 22) which is stated in Appendix 8.8.6.

Corollary 5. *Suppose Algorithm 13 is run with input parameters $\delta \in (0, 1)$, and $T_0 = \tilde{O}(d^2 \ln^2(\frac{1}{\delta}))$ given in Equation (8.15), then with probability at-least $1 - \delta$, the regret after T times satisfies*

$$R_T \leq O\left(\frac{d^2}{\gamma^{4.65}} \ln^2(d/\delta) \tau^2 \ln\left(\frac{TK}{\delta}\right)\right) + \tilde{O}(\sqrt{Td_m^*}),$$

where $\gamma = \min\{|\theta_i^*| : \theta_i^* \neq 0\}$ and d^* the sparsity of θ^* .

Discussion - Our regret scaling (in time) matches that of an oracle that knows the true problem complexity and thus obtains a regret scaling of $\tilde{O}(\sqrt{d_m^* T})$. This, thus improves on the rate compared to that obtained in [70], whose regret scaling is sub-optimal compared to the oracle. On the other hand however, our regret bound depends on γ and is thus not uniform over all θ^* , unlike the bound in [70] that is uniform over θ^* . Thus, in general, our results are not directly comparable to that of [70]. It is an interesting future work to close the gap and in particular, obtain the regret matching that of an oracle to hold uniformly over all θ^* .

8.6 Simulations

Synthetic Experiments We compare `ALB-Norm` with the (non-adaptive) `OFUL+` and an *oracle* that knows the problem complexity apriori. The oracle just runs `OFUL+` with the known problem complexity. We choose the bias $\sim \mathcal{U}[-1, 1]$, and the additive noise to be zero-mean Gaussian random variable with variance 0.5. At each round of the learning algorithm, we sample the context vectors from a d -dimensional standard Gaussian, $\mathcal{N}(0, I_d)$. We select $d = 50$, the number of arms, $K = 75$, and the initial epoch length as 100. In particular, we generate the true θ^* in 2 different ways: (i) $\|\theta^*\| = 0.1$, but the initial estimate $b_1 = 10$, and (ii) $\|\theta^*\| = 1$, with the initial estimate $b_1 = 10$.

In panel (a) and (b) of Figure 8.1, we observe that, in setting (i), `OFUL+` performs poorly owing to the gap between $\|\theta^*\|$ and b_1 . On the other hand, `ALB-Norm` is sandwiched between the `OFUL+`




Figure 8.1: Synthetic and real-data experiments, validating the effectiveness of Algorithm 1 and 12. All the results are averaged over 25 trials.

and the oracle. Similar things happen in setting (ii). In panel (c), we show that the norm estimates of ALB-NORM improves over epochs, and converges to the true norm very quickly.

In panel (d)-(f) of Figure 8.1, we compare the performance of ALB-DIM with the OFUL ([66]) algorithm and an *oracle* who knows the true support of θ^* apriori. For computational ease, we set $\varepsilon_i = 2^{-i}$ in simulations. We select θ^* to be $d^* = 20$ -sparse, with the smallest non-zero component, $\gamma = 0.12$. We have 2 settings: (i) $d = 500$ and (ii) $d = 200$. In panel (d) and (e), we observe a huge gap in cumulative regret between ALB-DIM and OFUL, thus showing the effectiveness of dimension adaptation. In panel (f), we plot the successive dimension refinement over epochs. We observe that within 4 – 5 epochs, ALB-DIM finds the sparsity of θ^* .

Real-data experiment Here, we evaluate the performance of ALB-NORM on Yahoo! ‘Learning to Rank Challenge’ dataset ([247]). In particular, we use the file `set2.test.txt`, which consists of 103174 rows and 702 columns. The first column denotes the rating, $\{0, 1, \dots, 4\}$ given by the user (which is taken as reward); the second column denotes the user id, and the rest 700 columns denote the context of the user. After selecting 20,000 rows and 50 columns at random (several other random selections yield similar results), we cluster the data by running k means algorithm with $k = 500$. We treat each cluster as a bandit arm with mean reward as the empirical mean of the individual rating in the cluster, and the context as the centroid of the cluster. This way, we obtain a bandit setting with $K = 500$ and $d = 50$.




Figure 8.2: Comparison with Stochastic Corral and ModCB

Assuming (reward, context) coming from a linear model (with bias, see Section 6.2), we use `ALB-NORM` to estimate the bias and θ^* simultaneously. In panel (g), we plot the cumulative reward accumulated over time. We observe that the reward is accumulated over time in an almost linear fashion. We also plot the norm estimate, $\|\theta^*\|$ over epochs in panel (h), starting with an initial estimate of 25. We observe that within 6 epochs the estimate stabilizes to a value of 11.1. This shows that `ALB-NORM` adapts to the actual $\|\theta^*\|$.

Comparison of ALB (dim): When θ^* is sparse, we compare ALB-Dim with 3 baselines: (i) the ModCB algorithm of [70] (ii) the Stochastic Corral algorithm of [248] and (iii) an oracle which knows the support of θ^* . We select θ^* to be $d^* = 20$ sparse, with dimension $d = 200$ and $d = 500$. The smallest non-zero component of θ^* is 0.12. For ModCB, we use `ILOVETOCONBANDITS` algorithm, similar to [70]. We select the cardinality of action set as 2 and select the sub-Gaussian parameter of the embedding as unity. In Figure 8.2, we observe that, the regret of ALB (dim) is better than ModCB and Stochastic Corral. The theoretical regret bound for ModCB scales as $\mathcal{O}(T^{2/3})$ (which is much larger than the ALB-Dim algorithm we propose), and Figure 1(c), validates this. The Stochastic Corral algorithm treats the base algorithms as bandit arms (with bandit feedback), as opposed to ALB-Dim which, at each arm-pull, updates the information about all the base algorithms. Thus, (Figs 1(d), 1(e)), ALB-Dim has a superior performance compared to Stochastic Corral.

8.7 Conclusion

In this paper, we considered refined model selection for linear bandits, by defining new notions of complexity. We gave two novel algorithms `ALB-NORM` and `ALB-DIM` that successively refines the hypothesis class and achieves model selection guarantees; regret scaling in the complexity of the smallest class containing the true model. This is the first such algorithm to achieve regret scaling similar to an oracle that knew the problem complexity. An interesting direction of future work is to derive regret bounds for the case when the dimension is a measure of complexity, that hold uniformly over all θ^* , i.e., have no explicit dependence on γ .

8.8 Proofs

8.8.1 Detailed Description of OFUL⁺

We now discuss the algorithm OFUL⁺. A variation of this was proposed in [67] in the context of model selection between linear and standard multi-armed bandits. As seen in the OFUL⁺ subroutine of Algorithm 11, we use $\tilde{\mu}_{i,t}$ to address the bias term in the observation, which we define shortly. The parameters b and δ appears in the construction of the confidence set and the regret guarantee. Furthermore, assume that the algorithm OFUL⁺ is run for \tilde{T} rounds.

Let A_s be the arm index played at time instant s and $T_i(t)$ be the number of times we play arm i until time t . Hence $T_i(t) = \sum_{s=1}^t \mathbf{1}\{A_s = i\}$. Also, let b be the current estimate of $\|\theta^*\|$. Also define,

$$\bar{g}_{i,t} = \frac{1}{T_i(t)} \sum_{s=1}^t g_{i,s} \mathbf{1}\{A_s = i\}.$$

With this, we have

$$\tilde{\mu}_{i,t} = \bar{g}_{i,t} + \sigma \left[\frac{1 + T_i(t)}{T_i^2(t)} \left(1 + 2 \log \left(\frac{K(1 + T_i(t))^{1/2}}{\delta} \right) \right) \right]^{1/2} + b \sqrt{\frac{2d}{T_i(t)} \log \left(\frac{1}{\delta} \right)} \quad (8.2)$$

In order to specify the confidence interval \mathcal{C}_t , we first talk about the least squares estimate $\hat{\theta}$ first. Using the notation of [67], we define

$$\hat{\theta}_t = (\alpha_{K+1:t}^\top \alpha_{K+1:t} + I)^{-1} \alpha_{K+1:t}^\top G_{K+1:t}$$

where $\alpha_{K+1:t}$ is a matrix with rows $\alpha_{A_{K+1},K+1}^\top, \dots, \alpha_{A_t,t}^\top$ and $G_{K+1:t} = [g_{A_{K+1},K+1} - \tilde{\mu}_{A_{K+1},K+1}, \dots, g_{A_t,t} - \tilde{\mu}_{A_t,t}]^\top$. With this, the confidence interval is defined as

$$\mathcal{C}_t = \left\{ \theta \in \mathbb{R}^d : \|\theta - \hat{\theta}_t\| \leq \mathcal{K}_\delta(b, t, \tilde{T}) \right\}, \quad (8.3)$$

and Lemma 2 of [67] shows that $\theta^* \in \mathcal{C}_t$ with probability at least $1 - 4\delta$.

We now define the quantity $\mathcal{K}_\delta(b, t, \tilde{T})$. Note that we track the dependence on the complexity

parameter $\|\theta^*\|$. We have

$$T_{\min}(\delta, \tilde{T}) = \left(\frac{16}{\rho_{\min}^2} + \frac{8}{3\rho_{\min}} \right) \log \left(\frac{2d\tilde{T}}{\delta} \right),$$

$$\mathcal{M}_\delta(b, t) = b + \sqrt{2\sigma^2 \left(\frac{d}{2} \log \left(1 + \frac{t}{d} \right) + \log \left(\frac{1}{\delta} \right) \right)}, \quad (8.4)$$

$$\Upsilon_\delta(b, t, \tilde{T}) = \frac{10}{3} \left(b + 2 + \sigma \sqrt{1 + 2 \log \left(\frac{2K\tilde{T}}{\delta} \right)} \right)$$

$$\times \left[\log \left(\frac{2K\tilde{T}}{\delta} \right) + \sqrt{t \log \left(\frac{2K\tilde{T}}{\delta} \right) + \log^2 \left(\frac{2K\tilde{T}}{\delta} \right)} \right], \quad (8.5)$$

$$\mathcal{K}_\delta(b, t, \tilde{T}) = \begin{cases} \mathcal{M}_\delta(b, t) + \Upsilon_\delta(b, t, \tilde{T}) & \text{if } 1 < t < T_{\min}, \\ \frac{\mathcal{M}_\delta(b, t)}{\sqrt{1 + \rho_{\min} t/2}} + \frac{\Upsilon_\delta(b, t, \tilde{T})}{1 + \rho_{\min} t/2} & \text{if } t > T_{\min}. \end{cases} \quad (8.6)$$

8.8.2 Proofs of the main results

In this section, we collect the proof of our main results. We start with the norm-based complexity measure.

8.8.3 Proof of Theorem 20

We first take Lemma 54 for granted and conclude the proof of Theorem 20 using the lemma. Suppose we play Algorithm 11 for N epochs. The cumulative regret is given by

$$R(T) \leq C_1(2\tau + K)\|\theta^*\| + \sum_{i=1}^N R(\delta_i, b_i)(T_i),$$

where $R(\delta_i, b_i)(T_i)$ is the cumulative regret of the $\text{OFUL}_{\delta_i}^+(b_i)$ in the i -th epoch. As seen (by tracking the dependence on $\|\theta^*\|$) in [67], the cumulative regret of $\text{OFUL}_{\delta_i}^+(b_i)$ scales linearly with b_i . Hence, we obtain

$$R(T) \leq \sum_{i=1}^N b_i R(\delta_i, 1)(T_i).$$

Using Lemma 1, we obtain, with probability at least $1 - 8\delta_1$,

$$R(T) \leq C_1(2\tau + K)\|\theta^*\| + (c_1\|\theta^*\| + c_2) \sum_{i=1}^N R(\delta_i, 1)(T_i)$$

Theorem 3 of [67] gives,

$$R(\delta_i, 1)(T_i) \leq C(\sqrt{K} + \sqrt{d})\sqrt{T_i} \log \left(\frac{KT_i}{\delta_i} \right) \quad (8.7)$$

with probability exceeding $1 - 5\delta_i$. With the doubling trick, we have

$$T_i = 2^{i-1}T_1, \quad \delta_i = \frac{\delta_1}{2^{i-1}}.$$

Substituting, we obtain

$$R(\delta_i, 1)(T_i) \leq C_1(\sqrt{K} + \sqrt{d})\sqrt{T_i} \log \left(\frac{KT_i}{\delta_i} \right) \left[(2i - 2) \log \left(\frac{KT_1}{\delta_1} \right) \right]$$

with probability at least $1 - 5\delta_i$.

Using the above expression, we obtain

$$R(T) \leq C_1(2\tau + K)\|\theta^*\| + (C_2\|\theta^*\| + C_3) \sum_{i=1}^N (\sqrt{K} + \sqrt{d})\sqrt{T_i} \left[(2i - 2) \log \left(\frac{KT_1}{\delta_1} \right) \right]$$

with probability

$$\begin{aligned} &\geq 1 - 8\delta_1 - 5\delta_1 \left(1 + \frac{1}{2} + \dots N\text{-th term} \right) \\ &\geq 1 - 8\delta_1 - 5\delta_1 \left(1 + \frac{1}{2} + \dots \right) \\ &= 1 - 8\delta_1 - 10\delta_1 \\ &= 1 - 18\delta_1, \end{aligned}$$

where the term $8\delta_1$ comes from Lemma 54. Also, from the doubling principle, we obtain

$$\sum_{i=1}^N 2^{i-1}T_1 = T \Rightarrow N = \log_2 \left(1 + \frac{T}{T_1} \right).$$

Using the above expression, we obtain

$$\begin{aligned}
 R(T) &\leq C_1(2\tau + K)\|\theta^*\| + (C_2\|\theta^*\| + C_3) \sum_{i=1}^N (\sqrt{K} + \sqrt{d})\sqrt{T_i} \left[(2i - 2) \log \left(\frac{KT_1}{\delta_1} \right) \right] \\
 &\leq C_1(2\tau + K)\|\theta^*\| + 2(C_2\|\theta^*\| + C_3)(\sqrt{K} + \sqrt{d}) \log \left(\frac{KT_1}{\delta_1} \right) \sum_{i=1}^N i\sqrt{T_i} \\
 &\leq C_1(2\tau + K)\|\theta^*\| + 2(C_2\|\theta^*\| + C_3)(\sqrt{K} + \sqrt{d}) \log \left(\frac{KT_1}{\delta_1} \right) N \sum_{i=1}^N \sqrt{T_i} \\
 &\leq C_1(2\tau + K)\|\theta^*\| + 2(C_2\|\theta^*\| + C_3)(\sqrt{K} + \sqrt{d}) \log \left(\frac{KT_1}{\delta_1} \right) \log \left(\frac{T}{T_1} \right) \sum_{i=1}^N \sqrt{T_i} \\
 &\leq C_1(2\tau + K)\|\theta^*\| + C(\|\theta^*\| + 1)(\sqrt{K} + \sqrt{d}) \log \left(\frac{KT_1}{\delta_1} \right) \log \left(\frac{T}{T_1} \right) \sqrt{T},
 \end{aligned}$$

where the last inequality follows from the fact that

$$\begin{aligned}
 \sum_{i=1}^N \sqrt{T_i} &= \sqrt{T_N} \left(1 + \frac{1}{\sqrt{2}} + \frac{1}{2} + \dots N\text{-th term} \right) \\
 &\leq \sqrt{T_N} \left(1 + \frac{1}{\sqrt{2}} + \frac{1}{2} + \dots \right) \\
 &= \frac{\sqrt{2}}{\sqrt{2} - 1} \sqrt{T_N} \\
 &\leq \frac{\sqrt{2}}{\sqrt{2} - 1} \sqrt{T}.
 \end{aligned}$$

The above regret bound holds with probability at least $1 - 18\delta_1$.

8.8.4 Proof of Lemma 54

Let us consider the i -th epoch, and let $\hat{\theta}_{\mathcal{E}_i}$ be the least square estimate of θ^* at the end of epoch i . From the above section, the confidence interval at the end of epoch i , is given by

$$\mathcal{C}_{\mathcal{E}_i} = \left\{ \theta \in \mathbb{R}^d : \|\theta - \hat{\theta}_{\mathcal{E}_i}\| \leq \mathcal{K}_{\delta_i}(b_i, T_i, T_i) \right\}$$

where we play OFUL $_{\delta_i}^+(b_i)$ during the i -th epoch, and T_i is the number of total rounds in the i -th epoch. By choosing $T_1 > T_{\min}(\delta, T)$, we ensure that $T_i \geq T_{\min}(\delta, T_i)$. From equation (8.6), and ignoring the non-dominant terms, we obtain

$$\mathcal{K}_{\delta_i}(b_i, T_i, T_i) = \frac{\mathcal{M}_{\delta_i}(b_i, T_i)}{\sqrt{1 + \rho_{\min} T_i/2}} + \frac{\Upsilon_{\delta_i}(b_i, T_i, T_i)}{1 + \rho_{\min} T_i/2},$$

with

$$\mathcal{M}_{\delta_i}(b_i, T_i) \leq b_i + c_1 \sigma \sqrt{d} \log \left(\frac{T_i}{d \delta_i} \right)$$

and

$$\Upsilon_{\delta_i}(b_i, T_i, T_i) = 4b_i \sqrt{T_i} \log \left(\frac{2KT_i}{\delta_i} \right) + c_2 \sigma \sqrt{T_i} \log \left(\frac{2KT_i}{\delta_i} \right)$$

Substituting the values, considering the dominating terms, and for a sufficiently large T_i , we obtain

$$\begin{aligned} \mathcal{K}_{\delta_i}(b_i, T_i, T_i) &\leq \frac{7b_i \log \left(\frac{2KT_i}{\delta_i} \right)}{\sqrt{\rho_{\min} T_i}} + C \frac{\sigma \sqrt{d}}{\sqrt{\rho_{\min} T_i}} \log \left(\frac{2KT_i}{\delta_i} \right) \\ &\leq \frac{7b_i \log \left(\frac{2KT}{\delta_i} \right)}{\sqrt{\rho_{\min} T_i}} + C \frac{\sigma \sqrt{d}}{\sqrt{\rho_{\min} T_i}} \log \left(\frac{2KT_i}{\delta_i} \right) \end{aligned}$$

where C is an universal constant. From Lemma 2 of [67], we know that $\theta^* \in \mathcal{C}_{\mathcal{E}_i}$ with probability at least $1 - 4\delta_i$. Hence, we obtain

$$\|\hat{\theta}_{\mathcal{E}_i}\| \leq \|\theta^*\| + 2\mathcal{K}_{\delta_i}(b_i, T_i, T_i) \leq \|\theta^*\| + \frac{14b_i \log \left(\frac{2KT_i}{\delta_i} \right)}{\sqrt{\rho_{\min} T_i}} + 2C \frac{\sigma \sqrt{d}}{\sqrt{\rho_{\min} T_i}} \log \left(\frac{2KT_i}{\delta_i} \right)$$

Recall from Algorithm 11 that at the end of the i -th epoch, we set the length $T_{i+1} = 2T_i$, and the estimate of $\|\theta^*\|$ is set to

$$b_{i+1} = \max_{\theta \in \mathcal{C}_{\mathcal{E}_i}} \|\theta\|.$$

From the definition of $\mathcal{C}_{\mathcal{E}_i}$, we obtain

$$b_{i+1} = \|\hat{\theta}_{\mathcal{E}_i}\| + \mathcal{K}_{\delta_i}(b_i, T_i, T_i) \leq \frac{7b_i \log \left(\frac{2KT_i}{\delta_i} \right)}{\sqrt{\rho_{\min} T_i}} + C \frac{\sigma \sqrt{d}}{\sqrt{\rho_{\min} T_i}} \log \left(\frac{2KT_i}{\delta_i} \right).$$

Re-writing the above expression, with probability at least $1 - 4\delta_i$, we obtain

$$\begin{aligned} b_{i+1} &\leq \|\theta^*\| + \left(\frac{7 \log \left(\frac{2KT_i}{\delta_i} \right)}{\sqrt{\rho_{\min}}} \right) \frac{b_i}{\sqrt{T_i}} + \left(\frac{C \sigma \log \left(\frac{2KT_i}{\delta_i} \right)}{\sqrt{\rho_{\min}}} \right) \frac{\sqrt{d}}{\sqrt{T_i}} \\ &\leq \|\theta^*\| + ip \frac{b_i}{\sqrt{T_i}} + iq \frac{\sqrt{d}}{\sqrt{T_i}} \\ &\leq \|\theta^*\| + ip \frac{b_i}{2^{\frac{i-1}{2}} \sqrt{T_1}} + iq \frac{\sqrt{d}}{2^{\frac{i-1}{2}} \sqrt{T_1}} \end{aligned} \tag{8.8}$$

where we use the fact that $\delta_i = \frac{\delta_1}{2^{i-1}}$ and $T_i = 2^{\frac{i-1}{2}} T_1$, and we have

$$p = \left(\frac{14 \log \left(\frac{2KT_1}{\delta_1} \right)}{\sqrt{\rho_{\min}}} \right)$$

and

$$q = \left(\frac{2C\sigma \log \left(\frac{2KT_1}{\delta_1} \right)}{\sqrt{\rho_{\min}}} \right).$$

Hence, we obtain

$$b_{i+1} - b_i \leq \|\theta^*\| + iq \frac{\sqrt{d}}{2^{\frac{i-1}{2}} \sqrt{T_1}} - \left(1 - ip \frac{1}{2^{\frac{i-1}{2}} \sqrt{T_1}} \right) b_i.$$

From the construction of b_i , we have $-b_i \leq -\|\theta^*\|$. Hence provided

$$T_1 \geq \frac{i^2 p^2}{2^{i-1}},$$

which is equivalent to the condition $T_1 \geq 3p^2$ (using the fact that $\frac{i^2}{2^{i-1}} \leq 3$ for $i \geq 1$), we obtain

$$b_{i+1} - b_i \leq \left(ip \frac{1}{2^{\frac{i-1}{2}} \sqrt{T_1}} \right) \|\theta^*\| + iq \frac{\sqrt{d}}{2^{\frac{i-1}{2}} \sqrt{T_1}}.$$

From the above expression, we obtain

$$\sup_i b_i < \infty.$$

with probability

$$\begin{aligned} &\geq 1 - 4\delta_1 \left(1 + \frac{1}{2} + \frac{1}{4} + \dots \right) \\ &= 1 - 8\delta_1. \end{aligned}$$

Invoking Equation (8.8) and using the above fact in conjunction yield (with probability at least $1 - 8\delta_1$)

$$\lim_{i \rightarrow \infty} b_i \leq \|\theta^*\|.$$

However, from construction $b_i \geq \|\theta^*\|$. Using this, along with the above equation, we obtain

$$\lim_{i \rightarrow \infty} b_i = \|\theta^*\|.$$

with probability exceeding $1 - 8\delta_1$. So, the sequence $\{b_0, b_1, \dots\}$ converges to $\|\theta^*\|$ with high probability, and hence our successive refinement algorithm is consistent.

Rate of Convergence: Since

$$b_i - b_{i-1} = \tilde{\mathcal{O}}\left(\frac{i}{2^i}\right), \quad (8.9)$$

with probability greater than $1 - 4\delta_i$, the rate of convergence of the sequence $\{b_i\}_{i=0}^{\infty}$ is exponential in the number of epochs.

Uniform upper bound on b_i for all i : We now compute a uniform upper bound on b_i for all i .

Consider the sequence $\left\{\frac{i}{2^{\frac{i-1}{2}}}\right\}_{i=1}^{\infty}$, and let t_j denote the j -th term of the sequence. It is easy to check that $\sup_i t_i = 1.5$, and that the sequence $\{t_i\}_{i=1}^{\infty}$ is convergent. With this new notation, we have

$$b_2 \leq \|\theta^*\| + t_1 \frac{pb_1}{\sqrt{T_1}} + t_1 \frac{q\sqrt{d}}{\sqrt{T_1}}.$$

with probability exceeding $1 - 4\delta_1$. Similarly, for b_3 , we have

$$\begin{aligned} b_3 &\leq \|\theta^*\| + t_2 \frac{pb_1}{\sqrt{T_1}} + t_2 \frac{q\sqrt{d}}{\sqrt{T_1}} \\ &\leq \left(1 + t_2 \frac{p}{\sqrt{T_1}}\right) \|\theta^*\| + \left(t_1 t_2 \frac{p}{\sqrt{T_1}} \frac{p}{\sqrt{T_1}} b_1\right) + \left(t_1 t_2 \frac{p}{\sqrt{T_1}} \frac{q\sqrt{d}}{\sqrt{T_1}} + t_2 \frac{q\sqrt{d}}{\sqrt{T_1}}\right). \end{aligned}$$

with probability at least $1 - 4\delta_1 - 4\delta_2 = 1 - 6\delta_1$. Similarly, we write expressions for b_4, b_5, \dots . Now, provided $T_1 \geq C_1 (\max\{p, q\} b_1)^2 d$, where $C_1 > 9$ is a sufficiently large constant, the expression for b_i can be upper-bounded as

$$b_i \leq (c_1 \|\theta^*\| + c_2), \quad (8.10)$$

with probability

$$\begin{aligned} &\geq 1 - 4\delta_1 \left(1 + \frac{1}{2} + \frac{1}{4} + \dots \text{upto } i\text{-th term}\right) \\ &\geq 1 - 4\delta_1 \left(1 + \frac{1}{2} + \frac{1}{4} + \dots\right) \\ &= 1 - 8\delta_1. \end{aligned}$$

Here c_1 and c_2 are constants, and are obtained from summing an infinite geometric series with decaying step size. We also use the fact that $b_1 \geq 1$, and the fact that $\delta_i = \frac{\delta_1}{2^{i-1}}$.

8.8.5 Proof of Theorem 21

We shall need the following lemma from [249], on the behaviour of linear regression estimates.

Lemma 55. *If $M \geq d$ and satisfies $M = O\left(\left(\frac{1}{\varepsilon^2} + d\right) \ln\left(\frac{1}{\delta}\right)\right)$, and $\widehat{\theta}^{(M)}$ is the least-squares estimate of θ^* , using the M random samples for feature, where each feature is chosen uniformly and independently on the unit sphere in d dimensions, then with probability 1, $\widehat{\theta}$ is well defined (the least squares regression has a unique solution). Furthermore,*

$$\mathbb{P}[\|\widehat{\theta}^{(M)} - \theta^*\|_\infty \geq \varepsilon] \leq \delta.$$

We shall now apply the theorem as follows. Denote by $\widehat{\theta}_i$ to be the estimate of θ^* at the beginning of any phase i , using all the samples from random explorations in all phases less than or equal to $i - 1$.

Remark 37. The choice $T_0 := O\left(d^2 \ln^2\left(\frac{1}{\delta}\right)\right)$ in Equation (8.1) is chosen such that from Lemma 56, we have that

$$\mathbb{P}\left[\|\widehat{\theta}^{(\lceil\sqrt{T_0}\rceil)} - \theta^*\|_\infty \geq \frac{1}{2}\right] \leq \delta$$

Lemma 56. *Suppose $T_0 = O\left(d^2 \ln^2\left(\frac{1}{\delta}\right)\right)$ is set according to Equation (8.1). Then, for all phases $i \geq 4$,*

$$\mathbb{P}\left[\|\widehat{\theta}_i - \theta^*\|_\infty \geq 2^{-i}\right] \leq \frac{\delta}{2^i}, \quad (8.11)$$

where $\widehat{\theta}_i$ is the estimate of θ^* obtained by solving the least squares estimate using all random exploration samples until the beginning of phase i .

Proof. The above lemma follows directly from Lemma 55. Lemma 55 gives that if $\widehat{\theta}_i$ is formed by solving the least squares estimate with at-least $M_i := O\left(\left(4^i + d\right) \ln\left(\frac{2^i}{\delta}\right)\right)$ samples, then the guarantee in Equation (8.11) holds. However, as $T_0 = O\left(\left(d + 1\right) \ln\left(\frac{2}{\delta}\right)\right)$, we have naturally that $M_i \leq 4^i i \sqrt{T_0}$. The proof is concluded if we show that at the beginning of phase $i \geq 4$, the total number of random explorations performed by the algorithm exceeds $i 4^i \lceil\sqrt{T_0}\rceil$. Notice that at the beginning of any phase $i \geq 4$, the total number of random explorations that have been performed is

$$\begin{aligned} \sum_{j=0}^{i-1} 5^j \lceil\sqrt{T_0}\rceil &= \lceil\sqrt{T_0}\rceil \frac{5^i - 1}{4}, \\ &\geq i 4^i \lceil\sqrt{T_0}\rceil, \end{aligned}$$

where the last inequality holds for all $i \geq 4$. □

The following corollary follows from a straightforward union bound.

Corollary 6.

$$\mathbb{P} \left[\bigcap_{i \geq 4} \{ \|\hat{\theta}_i - \theta^*\|_\infty \leq 2^{-i} \} \right] \geq 1 - \delta.$$

Proof. This follows from a simple union bound as follows.

$$\begin{aligned} \mathbb{P} \left[\bigcap_{i \geq 4} \{ \|\hat{\theta}_i - \theta^*\|_\infty \leq 2^{-i} \} \right] &= 1 - \mathbb{P} \left[\bigcup_{i \geq 4} \{ \|\hat{\theta}_i - \theta^*\|_\infty \geq 2^{-i} \} \right], \\ &\geq 1 - \sum_{i \geq 4} \mathbb{P} \left[\|\hat{\theta}_i - \theta^*\|_\infty \geq 2^{-i} \right], \\ &\geq 1 - \sum_{i \geq 4} \frac{\delta}{2^i}, \\ &\geq 1 - \sum_{i \geq 2} \frac{\delta}{2^i}, \\ &= 1 - \frac{\delta}{2}. \end{aligned}$$

□

We are now ready to conclude the proof of Theorem 21.

Proof of Theorem 21. We know from Corollary 6, that with probability at-least $1 - \delta$, for all phases $i \geq 4$, we have $\|\hat{\theta}_i - \theta^*\|_\infty \leq 2^{-i}$. Call this event \mathcal{E} . Now, consider the phase $i(\gamma) := \max \left(4, \log_2 \left(\frac{1}{\gamma} \right) \right)$. Now, when event \mathcal{E} holds, then for all phases $i \geq i(\gamma)$, \mathcal{D}_i is the correct set of d^* non-zero coordinates of θ^* . Thus, with probability at-least $1 - \delta$, the total regret upto time T can be upper bounded as follows

$$\begin{aligned} R_T \leq \sum_{j=0}^{i(\gamma)-1} \left(25^j T_0 + 5^j \lceil \sqrt{T_0} \rceil \right) &+ \sum_{j \geq i(\gamma)}^{\left\lceil \log_{25} \left(\frac{T}{T_0} \right) \right\rceil} \text{Regret}(\text{OFUL}(1, \delta_j; 25^j T_0)) \\ &+ \sum_{j=i(\gamma)}^{\left\lceil \log_{25} \left(\frac{T}{T_0} \right) \right\rceil} 5^j \lceil \sqrt{T_0} \rceil. \end{aligned} \tag{8.12}$$

The term $\text{Regret}(\text{OFUL}(L, \delta, T))$ denotes the regret of the OFUL algorithm [66], when run with parameters $L \in \mathbb{R}_+$, such that $\|\theta^*\| \leq L$, and $\delta \in (0, 1)$ denotes the probability slack and T is the time horizon. Equation (8.12) follows, since the total number of phases is at-most $\left\lceil \log_{25} \left(\frac{T}{T_0} \right) \right\rceil$. Standard result from [66] give us that, with probability at-least $1 - \delta$, we have

$$\text{Regret}(\text{OFUL}(1, \delta; T)) \leq 4\sqrt{Td^* \ln \left(1 + \frac{T}{d^*} \right)} \left(1 + \sigma \sqrt{2 \ln \left(\frac{1}{\delta} \right) + d^* \ln \left(1 + \frac{T}{d^*} \right)} \right).$$

Thus, we know that with probability at-least $1 - \sum_{i \geq 4} \delta_i \geq 1 - \frac{\delta}{2}$, for all phases $i \geq i(\gamma)$, the regret in the exploration phase satisfies

$$\begin{aligned} \text{Regret}(\text{OFUL}(1, \delta_i; 25^i T_0)) &\leq 4\sqrt{d^* 25^i T_0 \ln \left(1 + \frac{25^i T_0}{d^*} \right)} \\ &\quad \times \left(1 + \sigma \sqrt{2 \ln \left(\frac{2^i}{\delta} \right) + d^* \ln \left(1 + \frac{25^i T_0}{d^*} \right)} \right). \end{aligned} \quad (8.13)$$

In particular, for all phases $i \in [i(\gamma), \lceil \log_{25} \left(\frac{T}{T_0} \right) \rceil]$, with probability at-least $1 - \frac{\delta}{2}$, we have

$$\begin{aligned} \text{Regret}(\text{OFUL}(1, \delta_i; 25^i T_0)) &\leq 4\sqrt{d^* 25^i T_0 \ln \left(1 + \frac{25^i T_0}{d^*} \right)} \\ &\quad \times \left(1 + \sigma \sqrt{2 \ln \left(\frac{T}{T_0 \delta} \right) + d^* \ln \left(1 + \frac{25^i T_0}{d^*} \right)} \right), \\ &= \mathcal{C}(T, \delta, d^*) \sqrt{25^i T_0}, \end{aligned} \quad (8.14)$$

where the constant captures all the terms that only depend on T , δ and d^* . We can write that constant as

$$\mathcal{C}(T, \delta, d^*) = 4\sqrt{d^* \ln \left(1 + \frac{25T}{d^*} \right)} \left(1 + \sigma \sqrt{2 \ln \left(\frac{T}{T_0 \delta} \right) + d^* \ln \left(1 + \frac{25T}{d^*} \right)} \right).$$

Equation (8.14) follows, by substituting $i \leq \log_{25} \left(\frac{T}{T_0} \right)$ in all terms except the first 25^i term in Equation (8.13). As Equations (8.14) and (8.12) each hold with probability at-least $1 - \frac{\delta}{2}$, we can

combine them to get that with probability at-least $1 - \delta$,

$$\begin{aligned}
 R_T &\leq 2T_0 25^{i(\gamma)} + \sum_{j=0}^{\log_{25}\left(\frac{T}{T_0}\right)+1} \mathcal{C}(T, \delta, d^*) \sqrt{25^j T_0} + 25 \lceil \sqrt{T_0} \rceil 5^{\log_{25}\left(\frac{T}{T_0}\right)}, \\
 &\leq 2T_0 25^{i(\gamma)} + 25\sqrt{T} + \mathcal{C}(T, \delta, d^*) \sum_{j=0}^{\log_{25}\left(\frac{T}{T_0}\right)+1} \sqrt{25^j T_0}, \\
 &\stackrel{(a)}{\leq} 50T_0 \frac{2}{\gamma^{4.65}} + 25\sqrt{T} + 25\sqrt{T} \mathcal{C}(T, \delta, d^*), \\
 &= O\left(\frac{d^2}{\gamma^{4.65}} \ln^2\left(\frac{1}{\delta}\right)\right) + \tilde{O}\left(d^* \sqrt{T \ln\left(\frac{1}{\delta}\right)}\right).
 \end{aligned}$$

Step (a) follows from $25 \leq 2^{4.65}$. □

8.8.6 ALB-Dim for Stochastic Contextual Bandits with Finite Arms

8.8.7 ALB-Dim Algorithm for the Finite Armed Case

The algorithm given in Algorithm 13 is identical to the earlier Algorithm 12, except in Line 8, this algorithm uses `SupLinRel` [246] as opposed to `OFUL` used in the previous algorithm. In practice, one could also use `LinUCB` [64] in place of `SupLinRel`. However, we choose to present the theoretical argument using `SupLinRel`, as unlike `LinUCB`, has an explicit closed form regret bound [246]. The pseudocode is provided in Algorithm 13.

In phase $i \in \mathbb{N}$, the `SupLinRel` algorithm is instantiated with input parameter $25^i T_0$ denoting the time horizon, slack parameter $\delta_i \in (0, 1)$, dimension $d_{\mathcal{M}_i}$ and feature scaling $b(\delta)$. We explain the role of these input parameters. The dimension ensures that `SupLinRel` plays from the restricted dimension $d_{\mathcal{M}_i}$. The feature scaling implies that when a context $x \in \mathcal{X}$ is presented to the algorithm, the set of K feature vectors, each of which is $d_{\mathcal{M}_i}$ dimensional are $\frac{\phi^{d_{\mathcal{M}_i}}(x, 1)}{b(\delta)}, \dots, \frac{\phi^{d_{\mathcal{M}_i}}(x, K)}{b(\delta)}$. The constant $b(\delta) := O\left(\tau \sqrt{\log\left(\frac{TK}{\delta}\right)}\right)$ is chosen such that

$$\mathbb{P}\left[\sup_{t \in [0, T], a \in \mathcal{A}} \|\phi^M(x_t, a)\|_2 \geq b(\delta)\right] \leq \frac{\delta}{4}.$$

Such a constant exists since $(x_t)_{t \in [0, T]}$ are i.i.d. and $\phi^M(x, a)$ is a sub-gaussian random variable with parameter $4\tau^2$, for all $a \in \mathcal{A}$. Similar idea was used in [70].

Algorithm 13: Adaptive Linear Bandit (Dimension) with Finitely Many arms

- 1: **Input:** Initial Phase length T_0 and slack $\delta > 0$.
 - 2: $\widehat{\beta}_0 = \mathbf{1}, T_{-1} = 0$
 - 3: **for** Each epoch $i \in \{0, 1, 2, \dots\}$ **do**
 - 4: $T_i = 25^i T_0, \quad \varepsilon_i \leftarrow \frac{1}{2^i}, \quad \delta_i \leftarrow \frac{\delta}{2^i}$
 - 5: $\mathcal{D}_i := \{i : |\widehat{\beta}_i| \geq \frac{\varepsilon_i}{2}\}$
 - 6: $\mathcal{M}_i := \inf\{m : d_m \geq \max \mathcal{D}_i\}$.
 - 7: **for** Times $t \in \{T_{i-1} + 1, \dots, T_i\}$ **do**
 - 8: Play according to SupLinRel of [246] with time horizon of $25^i T_0$ with parameters $\delta_i \in (0, 1)$, dimension $d_{\mathcal{M}_i}$ and feature scaling $b(\delta) := O\left(\tau \sqrt{\log\left(\frac{TK}{\delta}\right)}\right)$.
 - 9: **end for**
 - 10: **for** Times $t \in \{T_i + 1, \dots, T_i + 5^i \sqrt{T_0}\}$ **do**
 - 11: Play an arm from the action set \mathcal{A} chosen uniformly and independently at random.
 - 12: **end for**
 - 13: $\alpha_i \in \mathbb{R}^{S_i \times d}$ with each row being the arm played during all random explorations in the past.
 - 14: $\mathbf{y}_i \in \mathbb{R}^{S_i}$ with i -th entry being the observed reward at the i -th random exploration in the past
 - 15: $\widehat{\beta}_{i+1} \leftarrow (\alpha_i^T \alpha_i)^{-1} \alpha_i \mathbf{y}_i$, is a d dimensional vector
 - 16: **end for**
-

8.8.8 Regret Guarantee for Algorithm 13

In order to specify a regret guarantee, we will need to specify the value of T_0 . We do so as before. For any N , denote by $\lambda_{max}^{(N)}$ and $\lambda_{min}^{(N)}$ to be the maximum and minimum eigen values of the following matrix: $\Sigma^N := \mathbb{E} \left[\frac{1}{K} \sum_{j=1}^K \sum_{t=1}^N \phi^M(x_t, j) \phi^M(x_t, j)^T \right]$, where the expectation is with respect to $(x_t)_{t \in [T]}$ which is an i.i.d. sequence with distribution \mathcal{D} . First, given the distribution of $x \sim \mathcal{D}$, one can (in principle) compute $\lambda_{max}^{(N)}$ and $\lambda_{min}^{(N)}$ for any $N \geq 1$. Furthermore, from the assumption on \mathcal{D} , $\lambda_{min}^{(N)} = \widetilde{O}\left(\frac{1}{\sqrt{d}}\right) > 0$ for all $N \geq 1$. Choose $T_0 \in \mathbb{N}$ to be the smallest integer such that

$$\sqrt{T_0} \geq b(\delta) \max \left(\frac{32\sigma^2}{(\lambda_{min}^{(\lceil \sqrt{T_0} \rceil)})^2} \ln(2d/\delta), \frac{4}{3} \frac{(6\lambda_{max}^{(\lceil \sqrt{T_0} \rceil)} + \lambda_{min}^{(\lceil \sqrt{T_0} \rceil)})(d + \lambda_{max}^{(\lceil \sqrt{T_0} \rceil)})}{(\lambda_{min}^{(\lceil \sqrt{T_0} \rceil)})^2} \ln(2d/\delta) \right). \quad (8.15)$$

As before, it is easy to see that

$$T_0 = O \left(d^2 \ln^2 \left(\frac{1}{\delta} \right) \tau^2 \ln \left(\frac{TK}{\delta} \right) \right).$$

Furthermore, following the same reasoning as in Lemmas 56 and 55, one can verify that for all $i \geq 4$, $\mathbb{P} \left[\|\widehat{\beta}_{i-1} - \beta^*\|_\infty \geq 2^{-i} \right] \leq \frac{\delta}{2^i}$.

Theorem 22. *Suppose Algorithm 13 is run with input parameters $\delta \in (0, 1)$, and T_0 as given in Equation (8.15), then with probability at-least $1 - \delta$, the regret after a total of T arm-pulls satisfies*

$$R_T \leq 2T_0 \max\left(25^4, \frac{2}{\gamma^{4.65}}\right) + 308(1 + \ln(2KT \ln T))^{3/2} \sqrt{Td_{m^*}} + 100\sqrt{T}.$$

The parameter $\gamma > 0$ is the minimum magnitude of the non-zero coordinate of β^* , i.e., $\gamma = \min\{|\beta_i^*| : \beta_i^* \neq 0\}$.

In order to parse the above theorem, the following corollary is presented.

Corollary 7. *Suppose Algorithm 13 is run with input parameters $\delta \in (0, 1)$, and $T_0 = \tilde{O}\left(d^2 \ln^2\left(\frac{1}{\delta}\right)\right)$ given in Equation (8.15), then with probability at-least $1 - \delta$, the regret after T times satisfies*

$$R_T \leq O\left(\frac{d^2}{\gamma^{4.65}} \ln^2(d/\delta) \tau^2 \ln\left(\frac{TK}{\delta}\right)\right) + \tilde{O}(\sqrt{Td_{m^*}}).$$

Proof of Theorem 22. The proof proceeds identical to that of Theorem 21. Observe from Lemmas 55 and 56, that the choice of T_0 is such that for all phases $i \geq 1$, the estimate

$\mathbb{P}\left[\|\widehat{\beta}_{i-1} - \beta^*\|_\infty \geq 2^{-i}\right] \leq \frac{\delta}{2^i}$. Thus, from an union bound, we can conclude that

$$\mathbb{P}\left[\bigcup_{i \geq 4} \|\widehat{\beta}_{i-1} - \beta^*\|_\infty \geq 2^{-i}\right] \leq \frac{\delta}{4}.$$

Thus at this stage, with probability at-least $1 - \frac{\delta}{2}$, the following events holds.

- $\sup_{t \in [0, T], a \in \mathcal{A}} \|\phi^M(x_t, a)\|_2 \leq b(\delta)$
- $\|\widehat{\beta}_{i-1} - \beta^*\|_\infty \leq 2^{-i}$, for all $i \geq 4$.

Call these events as \mathcal{E} . As before, let $\gamma > 0$ be the smallest value of the non-zero coordinate of β^* . Denote by the phase $i(\gamma) := \max\left(4, \log_2\left(\frac{2}{\gamma}\right)\right)$. Thus, under the event \mathcal{E} , for all phases $i \geq i(\gamma)$, the dimension $d_{\mathcal{M}_i} = d_{m^*}$, i.e., the SupLinRel is run with the correct set of dimensions.

It thus remains to bound the error by summing over the phases, which is done identical to that in Theorem 21. With probability, at-least $1 - \frac{\delta}{2} - \sum_{i \geq 4} \delta_i \geq 1 - \delta$,

$$\begin{aligned} R_T \leq & \sum_{j=0}^{i(\gamma)-1} \left(25^j T_0 + 5^j \sqrt{T_0}\right) + \sum_{j=i(\gamma)}^{\left\lceil \log_{25}\left(\frac{T}{T_0}\right) \right\rceil} \text{Regret}(\text{SupLinRel})(25^j T_0, \delta_j, d_{\mathcal{M}_j, b(\delta)}) \\ & + \sum_{j=i(\gamma)}^{\left\lceil \log_{25}\left(\frac{T}{T_0}\right) \right\rceil} 5^j \sqrt{T_0}, \end{aligned}$$

where $\text{Regret}(\text{SupLinRel})(25^i T_0, \delta_i, d_{\mathcal{M}_i, b(\delta)}) \leq 44(1 + \ln(2K25^i T_0 \ln 25^i T_0))^{3/2} \sqrt{25^i T_0 d_{\mathcal{M}_i}} + 2\sqrt{25^i T_0}$. This expression follows from Theorem 6 in [246]. We now use this to bound each of the three terms in the display above. Notice from straightforward calculations that the first term is bounded by $2T_0 25^{i(\gamma)}$ and the last term is bounded above by $25 \lceil \sqrt{T_0} \rceil 5^{\log_{25}(\frac{T}{T_0})}$ respectively. We now bound the middle term as

$$\begin{aligned} & \left[\log_{25} \left(\frac{T}{T_0} \right) \right] \\ & \sum_{j=i(\gamma)} \text{Reg}(\text{SupLinRel})(25^j T_0, \delta_i, d_m^*, b(\delta)) \\ & \leq b(\delta) \left(\sum_{j=i(\gamma)}^{\left[\log_{25} \left(\frac{T}{T_0} \right) \right]} 44(1 + \ln(2K25^j T_0 \ln 25^j T_0))^{3/2} \sqrt{25^j T_0 d_{\mathcal{M}_j}} + 2\sqrt{25^j T_0} \right). \end{aligned}$$

The first summation can be bounded as

$$\begin{aligned} & \left[\log_{25} \left(\frac{T}{T_0} \right) \right] \\ & \sum_{j=i(\gamma)} 44(1 + \ln(2K25^j T_0 \ln 25^j T_0))^{3/2} \sqrt{25^j T_0 d_{\mathcal{M}_j}} \\ & \leq \sum_{j=i(\gamma)}^{\left[\log_{25} \left(\frac{T}{T_0} \right) \right]} 44(1 + \ln(2KT \ln T))^{3/2} \sqrt{25^j T_0 d_m^*}, \\ & \leq 44(1 + \ln(2KT \ln T))^{3/2} 75^{\log_{25}(\frac{T}{T_0})} \sqrt{T_0 d_m^*}, \\ & = 308(1 + \ln(2KT \ln T))^{3/2} \sqrt{T d_m^*}, \end{aligned}$$

and the second by

$$\left[\log_{25} \left(\frac{T}{T_0} \right) \right] \sum_{j=i(\gamma)} 2\sqrt{25^j T_0} \leq 50\sqrt{T}.$$

Thus, with probability at-least $1 - \delta$, the regret of Algorithm 13 satisfies

$$R_T \leq 2T_0 25^{i(\gamma)} + 308(1 + \ln(2KT \ln T))^{3/2} \sqrt{T d_m^*} + 100\sqrt{T},$$

where $i(\gamma) := \max \left(4, \log_2 \left(\frac{2}{\gamma} \right) \right)$. Thus,

$$R_T \leq 2T_0 \max\left(25^4, \frac{2}{\gamma^{4.65}}\right) + 308(1 + \ln(2KT \ln T))^{3/2} \sqrt{Td_m^*} + 100\sqrt{T},$$

as $25 \leq 2^{4.65}$

□

Chapter 9

Model Selection for Contextual Bandits Beyond Linear Structure

In this chapter, we focus on a provable model selection guarantee for the (generic) stochastic contextual bandit problem. The results of Chapter 8 holds only for linear contextual bandits. The linear parameterization is quite restrictive. here, we remove the linear assumption and solve a non parametric problem. In particular we assume that we are given a set of M nested function classes and the reward function lies in one of them. Via careful exploration, we show that with high probability the correct model class identification is possible with minimal cost Θ (an additive constant in the regret term). We first discuss the contextual model selection problem.

9.1 Problem formulation

Let \mathcal{A} be the set of K actions, and let \mathcal{X} be the set of d dimensional contexts. At time t , Nature picks (x_t, r_t) in an i.i.d fashion, where $x_t \in \mathcal{X}$ and a context dependent $r_t \in [0, 1]^K$. Upon observing the context, the agent takes action $a_t \in \mathcal{A}$, and obtains the reward of $r_t(a_t)$. Note that, the reward $r_t(a_t)$ depends on the context x_t and the action a_t . Furthermore, we have the following realizability assumption:

Assumption 21. (Realizability:) *There exists a predictor $f^* \in \mathcal{F}$, such that $\mathbb{E}[r_t(a)|x] = f^*(x, a)$, for all x and a .*

Note that the realizability assumption is standard in the literature ([70, 71]. Also, in the contextual bandit literature ([71, 72]) it is assumed that the true regression function f^* is unknown, but we know the function class \mathcal{F} where it belongs. Hence, we pay some price (denoted by the regret) for this. To set up notation, for any $f \in \mathcal{F}$, we define a policy induced by the function, $\pi_f(x) = \operatorname{argmin}_{a \in \mathcal{A}} f(x, a)$. So, we need to compete with the policy induced by the true regressor

π_{f^*} . We define the regret over T rounds as the following:

$$R(T) = \sum_{t=1}^T [r_t(\pi_{f^*}(x_t)) - r_t(a_t)].$$

[71] proposed and analyzed a contextual bandit algorithm, namely FALCON (stands for FAsT Least-squares-regression-oracle CONtextual bandits), which gives provable guarantees for contextual bandits beyond linear structure. FALCON is an epoch based algorithm, and depends only on an *offline regression oracle*, which outputs an estimate \hat{f} of the regression function f^* at the beginning of each epoch. Using a randomization scheme, that depends on the inverse gap with respect to the estimate of the best action, it is showed that under the realizability assumption, with probability $1 - \delta$, FALCON yields a regret of

$$R(T) \leq \mathcal{O}(\sqrt{KT \log(|\mathcal{F}|T/\delta)}).$$

Although the above result makes sense only for the finite \mathcal{F} , an extension to the infinite \mathcal{F} is possible and was addressed in the same paper (see [71]).

We study the problem of model selection for contextual bandits. Hence, in contrast to the standard setting we do not know \mathcal{F} . Instead, we are given a nested class of M function classes, $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \dots \subset \mathcal{F}_M$. The smallest function class where the true regressor lies is denoted by \mathcal{F}_{d^*} , where $d^* \in [M]$. Hence, the regret of the contextual bandit algorithm should depend on \mathcal{F}_{d^*} . However, we do not know d^* , and our goal is to propose adaptive algorithms such that the regret depends on the *actual* problem complexity \mathcal{F}_{d^*} . First, let us rewrite the realizability assumption with the nested hypothesis class feature.

Assumption 22. *There exists a predictor $f_{d^*}^* \in \mathcal{F}_{d^*}$, such that $\mathbb{E}[r_t(a)|x] = f_{d^*}^*(x, a)$, for all x and a . Furthermore, the conditional variance of $r_t(\cdot)$ is bounded, i.e., given $x \in \mathcal{X}$,*

$$\mathbb{E}[r_t(a) - f_{d^*}^*(x, a)]^2 = \sigma^2 < \infty.$$

for all $a \in \mathcal{A}$.

Note that, the bounded second moment is also quite mild and is equivalent conditions appear in the stochastic optimization (SGD) literature.

We also have a separability condition as given in the following assumption. As we will see subsequently (in Section 9.4.1, this is equivalent to the condition on the minimum non-zero value of the optimal parameter in the linear contextual bandit setting.

Assumption 23. *For any $f \in \mathcal{F}_{d^*-1}$, we have*

$$\inf_f \mathbb{E}_{x,a} [f(x, a) - f_{d^*}^*(x, a)]^2 \geq \Delta,$$

for all pairs $(x, a) \in \mathcal{X} \times \mathcal{A}$, where $\Delta > 0$ is the minimum separation across the function classes.

Note that separability condition is quite standard in statistics, specially in the area of clustering and latent variable modelling (see [99, 250]). Model selection without separability condition is kept as an interesting future work.

Algorithm 14: Adaptive Contextual Bandits (ACB)

- 1: **Input:** epochs $0 = \tau_0 < \tau_1 < \tau_2 < \dots$, confidence parameter δ , Function classes $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \dots \subset \mathcal{F}_M$, threshold γ
 - 2: **for** epoch $m = 1, 2, \dots$, **do**
 - 3: $\delta_m = \delta/2^m$
 - 4: **for** function classes $j = 1, 2, \dots, M$ **do**
 - 5: Compute $\hat{f}_j^m = \operatorname{argmin}_{f \in \mathcal{F}_j} \sum_{t=\tau_{m-2}+1}^{\tau_{m-1}/2} (f(x_t, a_t) - r_t(a_t))^2$ via offline regression oracle
 - 6: Construct $T_j^m = \frac{1}{2^{m-2}} \sum_{t=\tau_{m-1}/2+1}^{\tau_{m-1}} (\hat{f}_j^m(x_t, a_t) - r_t(a_t))^2$
 - 7: **end for**
 - 8: **Model Selection:** Find the minimum index $\ell \in [M]$ such that $T_j^m \leq \gamma$. Let this class be denoted by \mathcal{F}_ℓ
 - 9: Set learning rate $\rho_m = (1/30)\sqrt{K\tau_{m-1} \log(|\mathcal{F}_\ell|(\tau_{m-1} - \tau_{m-2})/\delta_m)}$
 - 10: **for** round $t = \tau_{m-1} + 1, \dots, \tau_m$ **do**
 - 11: Observe context $x_t \in \mathcal{X}$
 - 12: Compute $\hat{f}_\ell^m(a)$ for all action $a \in \mathcal{A}$, set $\hat{a}_t = \operatorname{argmax}_{a \in \mathcal{A}} \hat{f}_\ell^m(a)$
 - 13: Define $p_t(a) = \frac{1}{K + \rho_m(\hat{f}_\ell^m(x_t, \hat{a}_t) - \hat{f}_\ell^m(x_t, a))} \forall a \neq \hat{a}_t, \quad p_t(\hat{a}_t) = 1 - \sum_{a \neq \hat{a}_t} p_t(a)$.
 - 14: Sample $a_t \sim p_t(\cdot)$ and observe reward $r_t(a_t)$.
 - 15: **end for**
 - 16: **end for**
-

9.2 Algorithm—Adaptive Contextual Bandits (ACB)

In this section, we provide a novel model selection algorithm that use successive refinements over epochs. We use the FALCON algorithm of [71], and add a model selection phase at the beginning of each epoch. In other words, over multiple epochs, we successively refine our estimates of the *proper* model class where the true regressor function f^* lies. The details are provided in Algorithm 14.

We assume that epochs have doubling epoch length. Let τ_0, τ_1, \dots be epoch instances, with $\tau_0 = 0$, and $\tau_m = 2^m$. Before the beginning of the m -th epoch, using all the data of the $m - 1$ -th epoch, we add a model selection module, as shown in Algorithm 14.

Note that we feed the samples of the $m - 1$ -th epoch to the offline regression oracle. Moreover, we split the samples in 2 equal halves. We use the first half to compute the regression estimate $\{\hat{f}_j^m\}_{j=1}^M$, and then use the rest to construct T_j^m . We do not use the same set of samples to remove any dependence issues with \hat{f}_j^m and the samples $\{x_t, a_t, r_t(a_t)\}_{t=\tau_{m-1}/2+1}^{\tau_{m-1}}$.

We compare T_j^m to a pre-calculated threshold, and pick the model class where T_j^m falls below such threshold (with smallest j , see Algorithm 23). After selecting the regressor function, the randomization follows exactly that of FALCON, i.e., we follow the inverse gap probability distribution $p_t(\cdot)$, and sample action $a_t \sim p_t(\cdot)$ and henceforth observe reward $r_t(a_t)$.

9.2.1 Regret Guarantee

We now analyze the performance of the model selection procedure of Algorithm 14. We have the doubling epochs, i.e., $\tau_m = 2^m$. Without loss of generality, we simply assume $\tau_1 = 2$. Also, assume that we are at the beginning of epoch m , and hence we have the samples from epoch $m - 1$. So, we have total of 2^{m-1} samples, out of which, we use 2^{m-2} to construct the regression functions and the rest 2^{m-2} to obtain the testing function T_j^m . Furthermore, we want the model selection procedure to succeed with probability at least $1 - \delta/2^m$, since we want a guarantee that holds for all m , and a simple application of the union bound yields that. We have the following Lemma.

Lemma 57. *Suppose we run Algorithm 14 with the threshold $\gamma = \sigma^2 + c_0$, where c_0 is a universal (small) constant. Then, provided*

$$2^m \gtrsim \max\left\{\frac{\log(1/\delta)}{(\Delta - c_0)^2}, \frac{1}{c_0^2} \max\{\log(|\mathcal{F}_M|/\delta), \log(1/\delta)\}\right\},$$

and $\Delta > c_0$, Algorithm 14 identifies the correct model class \mathcal{F}_{d^*} , with probability exceeding

$$1 - M \sum_{m=1}^{\infty} 3\delta/2^m \geq 1 - 6M\delta.$$

With the above lemma, we obtain the following regret bound for Algorithm 14.

Theorem 23. *Suppose Assumptions 2 and 3 hold and $\Delta > c_0$. Then with probability at least $1 - 6M\delta - \delta$, running Algorithm 14 for T iterations yield*

$$R(T) \leq C \max\left\{\frac{\log(1/\delta)}{(\Delta - c_0)^2}, \frac{1}{c_0^2} \max\{\log(|\mathcal{F}_M|/\delta), \log(1/\delta)\}\right\} + \mathcal{O}(\sqrt{KT \log(|\mathcal{F}_{d^*}|T/\delta)}).$$

Remark 38. The first term of the regret does not scale with T . Hence, the regret scaling (with respect to T) is $\mathcal{O}(\sqrt{KT \log(|\mathcal{F}_{d^*}|T/\delta)})$, with high probability, which is exactly the model selection guarantee.

Remark 39. The first term can be interpreted as the cost of model selection. Hence, the model selection procedure only adds an additive constant to the regret performance of FALCON.

Remark 40. Note that the first term is a problem dependent quantity and implies the complexity of the problem. If Δ is small, the model selection problem is hard, and the regret will be worse. Note that this is reflected via $\sim 1/\Delta^2$ dependence.

Remark 41. Note that the constant c_0 in the threshold can be chosen as a small enough constant, and so the condition on the gap, given by $\Delta > c_0$ is a mild condition. Since we do not know the gap Δ , from an implementation point of view, the value of c_0 should be kept sufficiently small. Of course, this results in a cost of $\mathcal{O}(1/c_0^2)$ in the constant (independent of T) term in the regret.

Remark 42. Note that if an oracle provides an estimate of the gap Δ , one can avoid the condition $\Delta > c_0$ by choosing the threshold γ as a function of Δ .

Algorithm 15: ETC for Model Selection for Contextual Bandits

- 1: **Input:** Action set \mathcal{A} , Function classes $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \dots \subset \mathcal{F}_M$, time horizon T , confidence parameter δ , threshold γ
 - 2: **Explore:**
 - 3: **for** $t = 1, 2, \dots, C\sqrt{T}$ **do**
 - 4: Observe context reward pair (x_t, r_t)
 - 5: Select action a_t uniformly at random from \mathcal{A} , independent of x_t
 - 6: Observe reward $r_t(a_t)$
 - 7: **end for**
 - 8: Compute regression estimator $\hat{f}_j = \operatorname{argmin}_{f \in \mathcal{F}_j} \frac{1}{C\sqrt{T}} \sum_{t=1}^{C\sqrt{T}} [f(x_t, a_t) - r_t(a_t)]^2$ (via offline regression oracle) for all $j \in [M]$
 - 9: **Model Selection test:**
 - 10: Obtain another set of $C\sqrt{T}$ fresh samples of (x_t, r_t, a_t) via pure exploration (similar to line 4-6)
 - 11: Construct the test statistic $T_j = \frac{1}{C\sqrt{T}} \sum_{t=1}^{C\sqrt{T}} (\hat{f}_j(x_t, a_t) - r_t(a_t))^2$ for all $j \in [M]$
 - 12: **Thresholding:** Find the minimum index $\ell \in [M]$ such that $T_j \leq \gamma$. We obtain the regressor $\hat{f}_\ell \in \mathcal{F}_\ell$
 - 13: **Commit:**
 - 14: **for** $t = 2C\sqrt{T} + 1, \dots, T$ **do**
 - 15: Observe context reward pair (x_t, r_t)
 - 16: Select action a_t according to FALCON algorithm of [71] with the function class \mathcal{F}_ℓ
 - 17: Observe reward $r_t(a_t)$
 - 18: **end for**
-

9.3 Explore-then-Commit algorithm for model selection

In the previous section, we analyze FALCON with model selection, and obtain high probability regret bounds. Here, instead, we use a simple explore-then-commit (ETC) algorithm for selecting the correct model, and then commit to it during the exploitation phase. After a round of exploration, we do a (one-time) threshold based testing to estimate the function class, and after that, exploit the estimated function class for the rest of the iterations. We show that this simple strategy finds the optimal function class \mathcal{F}_{d^*} with high probability. The details are given in Algorithm 15. We now explain the exploration and exploitation phases of this algorithm.

Exploration: Let the time horizon be T . For now, we assume that we know the value of T . We keep $2C\sqrt{T}$ time for exploration, where C is a universal constant. The exploration works in 2 phases.

Collect samples For $2C\sqrt{T}$ time epochs, we sample from the Nature randomly. Precisely, the context-reward pair (x_t, r_t) is being sampled by Nature in an i.i.d fashion, and the action the agent

takes is chosen uniformly at random from the action set \mathcal{A} . In particular, the action is chosen independent of the context x_t . Hence, this is a pure exploration strategy.

Obtain regression estimate Based on the samples of the first $C\sqrt{T}$ rounds, we estimate the regression function \hat{f} for all the (hypothesis) function classes $\mathcal{F}_1, \dots, \mathcal{F}_M$. More precisely, we call the offline regression oracle (as mentioned in [71]) and compute the following function estimate:

$$\hat{f}_j = \operatorname{argmin}_{f \in \mathcal{F}_j} \left(\sum_{t=1}^{C\sqrt{T}} f(x_t, a_t) - r_t(a_t) \right)^2$$

for all $j \in [M]$.

Testing on fresh samples To remove dependence issues, we use the remaining $C\sqrt{T}$ samples obtained from the sampling phase. Here we actually compute the following test statistic for all hypothesis classes:

$$T_j = \frac{1}{C\sqrt{T}} \sum_{t=1}^{C\sqrt{T}} (\hat{f}_j(x_t, a_t) - r_t(a_t))^2$$

for all $j \in [M]$. The idea is to perform a thresholding on $\{T_j\}_{j=1}^M$. Intuitively, since the (conditional) expectation of $r(a)$ is $f_{d^*}^*(x, a)$, we can expect T_j to be small for hypothesis classes that contain $f_{d^*}^*$. Otherwise, thanks to the separation condition, we expect T_j to be large. So, a simple thresholding on T_j is sufficient to identify the smallest hypothesis class where $f_{d^*}^*$ lies and we show that this procedure succeeds with high probability.

As mentioned earlier, we provide guarantees with ETC under 2 settings—(a) gap constrained and (b) long horizon.

ETC with (large) Gap constraint

Recall from Assumption 23 that, Δ is the minimum separation across the function classes. Here, we analyze Algorithm 15 with the following assumption on the gap Δ .

Assumption 24. *We assume that gap satisfies*

$$\Delta \geq \sigma^2 + C \left(T^{-1/2} \log(|\mathcal{F}_M|/\delta) + T^{-1/4} \sqrt{\log(1/\delta)} \right)$$

where T is the time horizon, $\delta \in (0, 1)$ and C is a universal constant.

Note that the above assumption suggests that for large T , the gap satisfies $\Delta \geq \sigma^2 + \text{small-term}$. Armed with the above assumption, we have the following result:

Lemma 58. *Suppose we set the threshold in Algorithm 15 as*

$$\gamma = \sigma^2 + C' \left(T^{-1/2} \log(|\mathcal{F}_M|/\delta) + T^{-1/4} \sqrt{\log(1/\delta)} \right),$$

for a universal constant C' . Then, with probability at least $1 - 2M\delta$, Algorithm 15 identifies the correct model class \mathcal{F}_{d^} .*

Regret Guarantee We now analyze the regret performance of Algorithm 15. The regret $R(T)$ is comprised of 2 stages; (a) exploration and (b) commit (exploitation). We have the following result.

Theorem 24. *Suppose Assumptions 2 and 3 hold. Then with probability at least $1 - (2M + 1)\delta$, running Algorithm 15 for T iterations yield*

$$R(T) \leq \mathcal{O}(\sqrt{KT \log(|\mathcal{F}_{d^*}|T/\delta)}).$$

ETC with long horizon

One of the potential issues of the previous section is the assumption on the gap. In this section, we remove this restriction. In other words, we trade off the requirement on the gap with the time horizon T . We show that for any gap Δ (as long as $\Delta > c_0$ for some small constant c_0), our proposed model selection procedure succeeds provided $T \gtrsim 1/\Delta^4$.

Guarantee We have the following result.

Theorem 25. *Suppose we choose the threshold $\gamma = \sigma^2 + c_0$ in Algorithm 15. The model selection procedure succeeds with probability exceeding $1 - 2(M + 1)\delta$ provided,*

$$T \geq C \max\left\{ \frac{\log^2(1/\delta)}{(\Delta - c_0)^4}, \frac{1}{c_0^4} \log^2(|\mathcal{F}_M|/\delta), \log^2(1/\delta) \right\},$$

and running Algorithm 15 for T iterations yield

$$R(T) \leq \mathcal{O}(\sqrt{KT \log(|\mathcal{F}_{d^*}|T/\delta)}).$$

Remark 43. Note that if the separation $\Delta \gtrsim \sigma^2$ (high SNR regime), the first condition is equivalent to $T \gtrsim \tilde{\mathcal{O}}(1)$, and we get the setting with large gap.

Remark 44. Note that we still require a gap constraint, $\Delta > c_0$. But this is a mild condition compared to Assumption 24, as one can choose c_0 sufficiently small constant.

9.4 Model Selection with infinite function classes

The results in Section 9 holds for finite function classes, since the regret bound depends on the cardinality of the function class. However, as shown in [71], it can be easily extended to the infinite function class setting. Exploiting the notion of the complexity of infinite function classes, the reduction is done.

Like before, we consider a nested sequence of M function classes $\mathcal{F}_1 \subset \dots \subset \mathcal{F}_M$. The reward is sampled from an unknown function f_{d^*} lying in the (smallest) function class indexed by $d^* \in [M]$, which is unknown. Given the function classes, our job is to find the function class \mathcal{F}_{d^*} , and subsequently exploit the class to obtain sub-linear regret. Let us first rewrite the separability assumption.

We assume that the function classes $\mathcal{F}_1 \subset \dots \subset \mathcal{F}_M$ are compact. This, in conjunction with the extreme value theorem ensures the existence of the following minimizers: for $j < d^*$, we define

$$\bar{f}_j = \operatorname{arginf}_{f \in \mathcal{F}_j} \mathbb{E}_{x,a} [f(x, a) - f_{d^*}(x, a)]^2 \quad (9.1)$$

for all pairs (x, a) . For $j \geq d^*$, we know that this minimizer is indeed f_{d^*} . This comes directly from the realizability assumption. Note that we require the existence of the minimizer (regression function) in order to use it for selecting actions in the contextual bandit framework (see [71])

Having defined the minimizers, we have the following separability assumption.

Assumption 25. For any \bar{f}_j , where $j < d^*$, we have

$$\mathbb{E}_{x,a} [\bar{f}_j(x, a) - f_{d^*}(x, a)]^2 \geq \Delta,$$

for all pairs $(x, a) \in \mathcal{X} \times \mathcal{A}$. Here we interpret Δ as the minimum separation across function classes.

Similar to [71], here, we are not worried about the explicit form of the regression functions \bar{f}_j . Rather, we assume the following performance guarantee of the offline regressor. For $j \geq d^*$ (meaning, the class containing the true regressor f_{d^*}), we have the following assumption.

Assumption 26. Given n i.i.d data samples $(x_1, a_1, r_1(a_1)), (x_2, a_2, r_2(a_2)), \dots, (x_n, a_n, r_n(a_n))$, the offline regression oracle returns a function \hat{f}_j , such that for $\delta > 0$, with probability at least $1 - \delta$,

$$\mathbb{E}_{x,a} [\hat{f}_j(x, a) - f_{d^*}(x, a)] \leq \xi_{\mathcal{F}_j, \delta}(n)$$

This assumption is taken from [71], and it is used to obtain the regret guarantee of the FALCON algorithm. As discussed in the above-mentioned paper, the quantity $\xi_{(\cdot)}(n)$ is a decreasing function of n , e.g., $\xi_{(\cdot)}(n) = \tilde{O}(1/n)$. As an instance, consider the class of all linear regressors in \mathbb{R}^d . In that case, $\xi_{(\cdot)}(n) \sim \tilde{O}(d/n)$. For function classes with finite VC dimension (or related quantities like VC-sub graph or fat-shattering dimension; pseudo dimension in general, denoted by \tilde{d}), we have $\xi_{(\cdot)}(n) \sim \tilde{O}(\tilde{d}/n)$.

Here, to avoid repetition, we do not present all our previous results in the infinite function class setting. We consider two instances:

1. The ETC algorithm with large gap condition
2. The adaptive contextual bandit (ACB) algorithm

The model-selection algorithm remains more-or-less the same overall. For Option I, we explore for the first $2C\sqrt{T}$ rounds. The first $C\sqrt{T}$ rounds are used to collect samples (x_t, r_t, a_t) via pure exploration. Feeding this samples to the offline regression oracle, and focusing on the individual function classes $\{\mathcal{F}_j\}_{j=1}^M$ separately, we obtain $(\hat{f}_j, \xi_{\mathcal{F}_j, \delta}(C\sqrt{T}))$ for all $j \in [M]$. Thereafter, we perform another round of pure exploration, and obtain $C\sqrt{T}$ fresh samples. Like in the finite case, we construct statistic T_j for all $j \in [M]$.

For Option II, we collect all the samples from the previous epoch of the FALCON algorithm, split the samples, to obtain the regression estimate \hat{f}_j^m and similarly construct test statistic T_j^m for all $j \in [M]$.

Similar to Algorithm 15 and 14, we choose the correct model based on a threshold on the test statistic T_j (for Option II, it is T_j^m). Let us assume that the threshold is γ . We show that for all $j \geq d^*$, $T_j \leq \gamma$, and for all $j < d^*$, $T_j > \gamma$ with high probability. We also obtain an explicit form of γ . Once this is shown, the model selection procedure follows exactly as Algorithm 15, i.e., we find the smallest index $\ell \in [M]$, for which $T_\ell \leq \gamma$. With high probability, we show that $\ell = d^*$.

Regret Guarantee We first show the guarantee for Option I, and then Option II.

Theorem 26. (ETC) *Suppose Assumptions 2, 6 and 5 hold and the gap satisfies*

$$\Delta \geq \sigma^2 + C_1 \left(\xi_{\mathcal{F}_M, \delta}(C\sqrt{T}) + T^{-1/4} \sqrt{\log(1/\delta)} \right).$$

Then running Algorithm 15 for T iterations with threshold

$$\gamma = \sigma^2 + C' \left(\xi_{\mathcal{F}_M, \delta}(C\sqrt{T}) + T^{-1/4} \sqrt{\log(1/\delta)} \right),$$

yields, with probability at least $1 - (2M + 1)\delta$, the regret

$$R(T) \leq \mathcal{O} \left(\sqrt{K \xi_{\mathcal{F}_{d^*}, \delta/2T}(T) T} \right).$$

Theorem 27. (ACB) *Suppose Assumptions 2 and 3 hold, and $\Delta > c_0$. Then with probability at least $1 - 6M\delta - \delta$, running Algorithm 14 for T iterations yield*

$$R(T) \leq C \max \left\{ \frac{\log(1/\delta)}{(\Delta - c_0)^2}, \frac{1}{c_0^2} \max \{ \log(|\mathcal{F}_M|/\delta), \log(1/\delta) \} \right\} + \mathcal{O} \left(\sqrt{K \xi_{\mathcal{F}_{d^*}, \delta/2T}(T) T} \right).$$

Remark 45. The proof of these theorems parallels exactly similar to the finite function class setting. The only difference is that instead of upper-bounding the prediction error using [72], we use the the definition of $\xi(\cdot)$ to accomplish this.

9.4.1 Special Case—Stochastic Linear Bandits

Here we consider the specific instance of the above-discussed generic contextual bandits—the linear setup. We see that order-wise, the generic Algorithm 14 recovers the regret guarantees of the linear bandit setup (with finite number of arms).

Recall the problem setup of Chapter 8 for finite set of actions. Furthermore, assume that the context embeddings $\phi^j(x, a)$ is a unit d_j dimensional Gaussian random variable. Note that this is stronger than the sub-Gaussian assumption.

Rewriting the nested hypothesis class, this corresponds to setting the function classes \mathcal{F}_j to be the linear class as the following:

$$\mathcal{F}_j = \{(x, a) \mapsto \langle \theta_j, \phi^j(x, a) \rangle \mid \theta_j \in \mathbb{R}^{d_j}, \|\theta\| \leq 1\}.$$

Also, m^* is the smallest index such that the optimal regressor is realized, i.e.,

$$f_{m^*}^*(x, a) = \langle \theta^*, \phi^{m^*}(x, a) \rangle.$$

So, this can be cast-ed as a model selection problem in our framework.

Let us look at the separability assumption of 23. For $j < m^*$, we first compute

$$\bar{f}_j = \operatorname{arginf}_{f \in \mathcal{F}_j} \mathbb{E}_{x,a} [f(x, a) - f_{m^*}^*(x, a)]^2,$$

and then compute the quantity

$$\mathbb{E}_{x,a} [\bar{f}_j(x, a) - f_{m^*}^*(x, a)]^2.$$

Substituting $f_{m^*}^*(x, a) = \langle \theta^*, \phi^{m^*}(x, a) \rangle$ and optimizing over θ_j , we obtain

$$\mathbb{E}_{x,a} [\bar{f}_j(x, a) - f_{m^*}^*(x, a)]^2 \geq \Upsilon^2 (d_{m^*} - j) \geq \Upsilon^2,$$

where Υ is the minimum magnitude of the non-zero coordinate of θ^* , i.e., $\Upsilon = \min\{|\theta_i^*| : \theta_i^* \neq 0\}$. Note that we exploit the Gaussian distribution of $\phi(\cdot)$ to obtain the above calculation.

Hence for linear stochastic bandits, one may take

$$\Delta = \Upsilon^2,$$

where Υ is the minimum magnitude of the non-zero coordinate of θ^* , i.e., $\Upsilon = \min\{|\theta_i^*| : \theta_i^* \neq 0\}$. Substituting in the regret expression of Theorem 27, with $\xi_{\mathcal{F}_{d_{m^*}}}(\cdot)(T) = \tilde{\mathcal{O}}(d_{m^*}/T)$ (linear class of d_{m^*} dimensional functions, VC-dimension is $d_{m^*} + 1$), we obtain

$$R_{lin}(T) \leq \tilde{\mathcal{O}}\left(\frac{d}{\Upsilon^4} + \sqrt{KTd_{m^*}}\right).$$

with high probability. Several remarks are in order:

Remark 46. If the number of actions $K = \mathcal{O}(1)$, the above regret scaling matches to that of Theorem 22. Hence, in this setting we recover the performance of Algorithm 13.

Remark 47. Note that in general, the performance of Algorithm 14 specialized to the linear setting is worse than Theorem 22. In particular, there is no K dependence in Theorem 22, but we have a \sqrt{K} term in the leading factor here. This matches our intuition. Algorithm 14 is applicable for any generic contextual bandit problem, whereas Algorithm 13 is specialized to the linear case only. The price in regret can be viewed as the cost of generalization.

Remark 48. Let us now focus on the additive (minor) term with no T dependence. It scales as $1/\Upsilon^4$ here, whereas in Theorem 22 it scales as $1/\Upsilon^{4.65}$. Note that, we remarked (after Theorem 13) that via carefully choosing the problem constants, the dependence in Theorem 22 can be made arbitrarily close to $1/\Upsilon^4$. So, we have the same dependence on Υ in both the settings.

Remark 49. Finally, note that the additive term is linearly dependent (d) here, whereas it has a quadratic dependence (d^2) in Theorem 22. We believe this stems from the analysis of Algorithm 13. In Algorithm 13, we are successively estimating the support of the underlying true parameter, and it is not clear whether support recovery is indeed required to ensure low regret.

9.5 Proofs

9.5.1 Proof of Lemma 57

Let us first show that T_j concentrates around its expectation. We show it via a simple application of the Hoeffdings inequality.

Fix a particular m and $j \in [M]$. Note that \hat{f}_j^m is computed based on 2^{m-2} samples. Also, in the testing phase, we use a fresh set of 2^{m-2} samples, and so \hat{f}_j^m is independent of the second set of samples, used in constructing T_j . Since we have $r(\cdot) \in [0, 1]$. Furthermore, we have $\hat{f}_j^m(\cdot) \in [0, 1]$. Note that this assumption is justified since our goal is obtain an estimate of the reward function via regression function, and this assumption also features in [71]. So the random variable $(\hat{f}_j^m(x_t, a_t) - r_t(a_t))^2$ is upper-bounded by 4, and hence sub-Gaussian with a constant parameter.

Note that we are using only the samples from the previous epoch. In the, FALCON algorithm, the regression estimate actually remains fixed over an entire epoch. Hence, conditioning on the filtration consisting of (context, action, reward) triplet upto the end of the $m - 2$ -th epoch, the random variables $\{(\hat{f}_j^m(x_t, a_t) - r_t(a_t))^2\}_{t=\tau_{m-1}/2+1}^{\tau_{m-1}}$ (a total of 2^{m-2} samples) are independent. Note that similar argument is given in [71, Section 4.1] (the FALCON+ algorithm) to argue the independence of the (context, action, reward) triplet, accumulated over just the previous epoch.

Hence using Hoeffdings inequality for sub-Gaussian random variables, we have

$$\mathbb{P}(T_j - \mathbb{E}T_j \geq \ell) \leq \exp(-n\ell^2/32).$$

Let us look at the expression $\mathbb{E}T_j$.

Realizable classes: Fix m and consider $j \in [M]$ such that $j \geq d^*$. So, for this realizable setting, we obtain the excess risk as:

$$\begin{aligned} & \mathbb{E}_{x,r,a}[\hat{f}_j^m(x,a) - r(a)]^2 - \inf_{f \in \mathcal{F}_j} \mathbb{E}_{x,r,a}[f(x,a) - r(a)]^2 \\ &= \mathbb{E}_{x,r,a}[\hat{f}_j^m(x,a) - r(a)]^2 - \mathbb{E}_{x,r,a}[f_{d^*}^*(x,a) - r(a)]^2 \\ &= \mathbb{E}_{x,a}[\hat{f}_j^m(x,a) - f_{d^*}^*(x,a)]^2. \end{aligned}$$

So, we have, for the realizable function class,

$$\begin{aligned} \mathbb{E}T_j^m &= \frac{1}{2^{m-2}} \mathbb{E}_{x_t, r_t, a_t} \sum_{t=1}^{2^{m-2}} [\hat{f}_j^m(x_t, a_t) - r_t(a_t)]^2 \\ &= \frac{1}{2^{m-2}} \sum_{t=1}^{2^{m-2}} \mathbb{E}_{x_t, r_t, a_t} [f_{d^*}^*(x_t, a_t) - r_t(a_t)]^2 + \frac{1}{2^{m-2}} \sum_{t=1}^{2^{m-2}} \mathbb{E}_{x_t, a_t} [\hat{f}_j^m(x, a) - f_{d^*}^*(x, a)]^2 \\ &\leq \sigma^2 + C_1 \log(2^m |\mathcal{F}_j| / \delta) / (2^{m-2}), \end{aligned}$$

with probability at least $1 - \delta/2^m$. Here, the first term comes from the second moment bound of σ^2 , and the second term comes from [72, Lemma 4.1]. So, by applying Hoeffding's inequality, we finally have.

$$T_j^m \leq \sigma^2 + C \frac{\log(|\mathcal{F}_j|/\delta)}{2^m} + C_1 \frac{m}{2^m} + C_3 \frac{\sqrt{\log(1/\delta)}}{2^{m/2}} + C_4 \frac{\sqrt{m}}{2^{m/2}}$$

with probability at least $1 - 2\delta/2^m$. Note that, provided

$$2^m \gtrsim \frac{1}{c_0^2} \max\{\log(|\mathcal{F}_M|/\delta), \log(1/\delta)\}, \quad (9.2)$$

with carefully chosen constants, we have

$$T_j^m \leq \sigma^2 + c_0$$

with probability at least $1 - 2\delta/2^m$, where c_0 is a small universal constant.

Non-Realizable classes: For the non realizable classes, we have the following calculation. For any $f \in \mathcal{F}_j$, where $j < d^*$, we have

$$\begin{aligned} & \mathbb{E}_{x,r,a}[f(x,a) - r(a)]^2 - \mathbb{E}_{x,r,a}[r(a) - f_{d^*}^*(x,a)]^2 \\ &= \mathbb{E}_{x,a,r}[(f(x,a) - f_{d^*}^*(x,a))(f(x,a) + f_{d^*}^*(x,a) - 2r(a))] \\ &= \mathbb{E}_{x,a} \mathbb{E}_{r|x}[(f(x,a) - f_{d^*}^*(x,a))(f(x,a) + f_{d^*}^*(x,a) - 2r(a))] \\ &= \mathbb{E}_{x,a}[(f(x,a) - f_{d^*}^*(x,a))(f(x,a) + f_{d^*}^*(x,a) - 2\mathbb{E}_{r|x}r(a))] \\ &= \mathbb{E}_{x,a}[f(x,a) - f_{d^*}^*(x,a)]^2, \end{aligned}$$

where the third inequality follows from the fact that given context x , the distribution of r is independent of a (see [72, Lemma 4.1]).

So, we have

$$\begin{aligned}\mathbb{E}_{x,r,a}[f(x,a) - r(a)]^2 &\geq \mathbb{E}_{x,r,a}[r(a) - f_{d^*}^*(x,a)]^2 + \mathbb{E}_{x,a}[f(x,a) - f_{d^*}^*(x,a)]^2 \\ &\geq \Delta + \sigma^2,\end{aligned}$$

where the last inequality comes from the separability assumption along with the assumption on the second moment. Since the regressor $\hat{f}_j^m \in \mathcal{F}_j$, we have

$$\begin{aligned}\mathbb{E}_{x,r,a}[\hat{f}_j^m(x,a) - r(a)]^2 &\geq \mathbb{E}_{x,r,a}[r(a) - f_{d^*}^*(x,a)]^2 + \mathbb{E}_{x,a}[f(x,a) - f_{d^*}^*(x,a)]^2 \\ &\geq \Delta + \sigma^2.\end{aligned}$$

Now, using 2^{m-2} samples, we obtain

$$T_j^m \geq \Delta + \sigma^2 - C_3 \frac{\sqrt{\log(1/\delta)}}{2^{m/2}} - C_4 \frac{\sqrt{m}}{2^{m/2}}$$

with probability at least $1 - \delta/2^m$. Now, suppose m is such that

$$\frac{\sqrt{\log(1/\delta)}}{2^{m/2}} + \frac{\sqrt{m}}{2^{m/2}} \lesssim \Delta - c_0.$$

where c_0 is a small universal constant (see equation (9.2)) such that $\Delta > c_0$. Observe that, provided

$$2^m \gtrsim \frac{\log(1/\delta)}{(\Delta - c_0)^2}.$$

and with appropriately chosen constants, we have

$$T_j^m \geq \sigma^2 + c_0.$$

with probability at least $1 - \delta/2^m$.

Based on the above calculation, we choose the threshold $\gamma = \sigma^2 + c_0$. Finally, we say that provided

$$2^m \gtrsim \max\left\{\frac{\log(1/\delta)}{(\Delta - c_0)^2}, \frac{1}{c_0^2} \max\{\log(|\mathcal{F}_M|/\delta), \log(1/\delta)\}\right\},$$

the model selection procedure succeeds with probability exceeding

$$1 - \sum_{m=1}^{\infty} 3\delta/2^m \geq 1 - 6\delta.$$

9.5.2 Proof of Theorem 23

The above calculation shows that as soon as

$$2^m \gtrsim \max\left\{\frac{\log(1/\delta)}{(\Delta - c_0)^2}, \frac{1}{c_0^2} \max\{\log(|\mathcal{F}_M|/\delta), \log(1/\delta)\}\right\},$$

the model selection procedure will succeed with high probability. Until the above condition is satisfied, we do not have any handle on the regret and hence the regret in that phase will be linear. This corresponds the first term in the regret expression. After the condition is satisfied, the regret is given by (see [71])

$$\sum_m \mathcal{O}\left(\sqrt{K(\tau_m - \tau_{m-1}) \log(|\mathcal{F}_{d^*}|(\tau_m - \tau_{m-1})/\delta)}\right) \leq \mathcal{O}\left(\sqrt{KT \log(|\mathcal{F}_{d^*}|T/\delta)}\right),$$

with probability exceeding $1 - \delta$. Finally, a union bound on the error probability yields the theorem.

9.5.3 Proof of Lemma 58

Since, we have samples from pure exploration, let us first show that T_j concentrates around its expectation. We show it via a simple application of the Hoeffdings inequality.

Fix a particular $j \in [M]$. Note that \hat{f}_j is computed based on the first set of $C\sqrt{T}$ samples. Also, in the testing phase, we again sample $C\sqrt{T}$ samples, and so \hat{f}_j is independent of the second set of $C\sqrt{T}$ samples, used in constructing T_j . Note that we have $r(\cdot) \in [0, 1]$. Furthermore, we have $\hat{f}_j(\cdot) \in [0, 1]$. Note that this assumption is justified since our goal is obtain an estimate of the reward function via regression function. This assumption also features in [71]. So the random variable $(\hat{f}_j(x_t, a_t) - r_t(a_t))^2$ is upper-bounded by 4, and hence sub-Gaussian with a constant parameter. Also, note that since we are choosing an action independent of the context, the random variables $\{(\hat{f}_j(x_t, a_t) - r_t(a_t))^2\}_{t=1}^{C\sqrt{T}}$ are independent. Hence using Hoeffdings inequality for sub-Gaussian random variables, we have

$$\mathbb{P}(T_j - \mathbb{E}T_j \geq \ell) \leq \exp(-n\ell^2/32).$$

Re-writing the above, we obtain

$$T_j - \mathbb{E}T_j \leq \sqrt{\frac{32 \log(1/\delta)}{C\sqrt{T}}}$$

with probability at least $1 - \delta$. Let us look at the expression $\mathbb{E}T_j$.

$$\mathbb{E}T_j = \mathbb{E}\left(\frac{1}{C\sqrt{T}} \sum_{t=1}^{C\sqrt{T}} (\hat{f}_j(x_t, a_t) - r_t(a_t))^2\right).$$

Case I: Realizable Class First consider the case that $j \geq d^*$, meaning that $f_{d^*}^* \in \mathcal{F}_j$. So, for this realizable setting, we obtain the excess risk as (using [72])

$$\begin{aligned} & \mathbb{E}_{x,r,a}[\hat{f}_j(x,a) - r(a)]^2 - \inf_{f \in \mathcal{F}_j} \mathbb{E}_{x,r,a}[f(x,a) - r(a)]^2 \\ &= \mathbb{E}_{x,r,a}[\hat{f}_j(x,a) - r(a)]^2 - \mathbb{E}_{x,r,a}[f_{d^*}^*(x,a) - r(a)]^2 \\ &= \mathbb{E}_{x,a}[\hat{f}_j(x,a) - f_{d^*}^*(x,a)]^2. \end{aligned}$$

So, we have, for the realizable function class,

$$\begin{aligned} \mathbb{E}T_j &= \frac{1}{C\sqrt{T}} \mathbb{E}_{x_t, r_t, a_t} \sum_{t=1}^{C\sqrt{T}} [\hat{f}_j(x_t, a_t) - r_t(a_t)]^2 \\ &= \frac{1}{C\sqrt{T}} \sum_{t=1}^{C\sqrt{T}} \mathbb{E}_{x_t, r_t, a_t} [f_{d^*}^*(x_t, a_t) - r_t(a_t)]^2 + \frac{1}{C\sqrt{T}} \sum_{t=1}^{C\sqrt{T}} \mathbb{E}_{x_t, a_t} [\hat{f}(x, a) - f_{d^*}^*(x, a)]^2 \\ &\leq \sigma^2 + C_1 \log(|\mathcal{F}_j|/\delta)/(C\sqrt{T}), \end{aligned}$$

with probability at least $1 - \delta$, where C_1 is a known universal constant. Here, we also use the fact that the second moment bound is σ^2 . So, we finally have.

$$T_j \leq \sigma^2 + C_2 T^{-1/2} \log(|\mathcal{F}_j|/\delta) + C_3 T^{-1/4} \sqrt{\log(1/\delta)}$$

with probability at least $1 - 2\delta$.

Case II: Non-realizable class We now consider the case when $j < d^*$, meaning that $f_{d^*}^*$ does not lie in \mathcal{F}_j . We have

$$\mathbb{E}_{x,r,a}[\hat{f}_j(x,a) - r(a)]^2 - \inf_{f \in \mathcal{F}_j} \mathbb{E}_{x,r,a}[f(x,a) - r(a)]^2 \geq 0.$$

Continuing, we obtain

$$\begin{aligned} \mathbb{E}T_j &= \frac{1}{C\sqrt{T}} \mathbb{E}_{x_t, r_t, a_t} \sum_{t=1}^{C\sqrt{T}} [\hat{f}_j(x_t, a_t) - r_t(a_t)]^2 \\ &\geq \inf_{f \in \mathcal{F}_j} \mathbb{E}_{x,a} [f(x,a) - f_{d^*}^*(x,a)]^2, \\ &\geq \Delta. \end{aligned}$$

So, in this setting,

$$\begin{aligned} T_j &\geq \mathbb{E}T_j - \sqrt{\frac{32 \log(1/\delta)}{C\sqrt{T}}} \\ &\geq \Delta - \sqrt{\frac{32 \log(1/\delta)}{C\sqrt{T}}}. \end{aligned}$$

with probability at least $1 - \delta$.

Choice of threshold If we can ensure the lower bound (on T_j) in the non-realizable setting is larger than the upper bound in the realizable case, we can ensure identifiability. Based on the above calculation, we fix the following threshold on the test statistic T_j ,

$$\begin{aligned} & \sigma^2 + C' \left(T^{-1/2} \log(|\mathcal{F}_j|/\delta) + T^{-1/4} \sqrt{\log(1/\delta)} \right) \\ & \leq \sigma^2 + C' \left(T^{-1/2} \log(|\mathcal{F}_M|/\delta) + T^{-1/4} \sqrt{\log(1/\delta)} \right) = \gamma \end{aligned}$$

For all $j \in [M]$, we compare T_j with this threshold γ . From the above calculation, it is clear that

$$T_j \leq \gamma$$

with high probability for all $j \leq d^*$. Furthermore, recall that from Assumption 24). For sufficiently large T , it is ensured that

$$T_j > \gamma$$

for all $j < d^*$ (with high probability). This implies that we can identify whether a function class is realizable or not simply based on the threshold γ . Finally, a simple union bound yields the result.

9.5.4 Proof of Theorem 24

The regret $R(T)$ can be decomposed in 2 stages, namely exploration and exploitation.

$$R(T) = R_{\text{explore}} + R_{\text{exploit}}$$

Since we spend $2C\sqrt{T}$ time steps in exploration, and $r_t(\cdot) \in [0, 1]$, the regret incurred in this stage

$$R_{\text{explore}} \leq C_1 \sqrt{T}.$$

Now, at the end of the explore stage, provided Assumptions 2 and 3, we know, with probability at least $1 - 2\delta$, we obtain the true function class \mathcal{F}_{d^*} . The threshold is set in such a way that we obtain the above result. Now, we would just commit to the function class and use the contextual bandit algorithm, namely FALCON. From [71], the regret guarantee of FALCON is

$$\begin{aligned} R_{\text{exploit}} & \leq \mathcal{O} \left(\sqrt{K(T - 2C\sqrt{T}) \log(|\mathcal{F}_{d^*}|(T - 2C\sqrt{T})/\delta)} \right) \\ & \leq \mathcal{O} \left(\sqrt{KT \log(|\mathcal{F}_{d^*}|T/\delta)} \right), \end{aligned}$$

with probability exceeding $1 - \delta$. Combining the above expressions yield the result.

9.5.5 Proof of Lemma 25

As usual, we consider 2 cases:

Relizable class: The calculation here remains the same as before,

$$T_j \leq \sigma^2 + C_2 T^{-1/2} \log(|\mathcal{F}_j|/\delta) + C_3 T^{-1/4} \sqrt{\log(1/\delta)}$$

with probability at least $1 - 2\delta$. We now choose T such that

$$\begin{aligned} C_2 T^{-1/2} \log(|\mathcal{F}_j|/\delta) + C_3 T^{-1/4} \sqrt{\log(1/\delta)} &\leq c_0 \\ \Rightarrow T &\gtrsim \frac{1}{c_0^4} \max\{\log^2(|\mathcal{F}_M|/\delta), \log^2(1/\delta)\}, \end{aligned}$$

for appropriate choice of constants. Using this, we get

$$T_j \leq \sigma^2 + c_0$$

with probability at least $1 - 2\delta$.

Non-realizable class For any $f \in \mathcal{F}_j$, where $j < d^*$, we have

$$\begin{aligned} &\mathbb{E}_{x,r,a}[f(x,a) - r(a)]^2 - \mathbb{E}_{x,r,a}[r(a) - f_{d^*}^*(x,a)]^2 \\ &= \mathbb{E}_{x,a,r}[(f(x,a) - f_{d^*}^*(x,a))(f(x,a) + f_{d^*}^*(x,a) - 2r(a))] \\ &= \mathbb{E}_{x,a} \mathbb{E}_{r|x}[(f(x,a) - f_{d^*}^*(x,a))(f(x,a) + f_{d^*}^*(x,a) - 2r(a))] \\ &= \mathbb{E}_{x,a}[(f(x,a) - f_{d^*}^*(x,a))(f(x,a) + f_{d^*}^*(x,a) - 2\mathbb{E}_{r|x}r(a))] \\ &= \mathbb{E}_{x,a}[f(x,a) - f_{d^*}^*(x,a)]^2, \end{aligned}$$

where the third inequality follows from the fact that given context x , the distribution of r is independent of a . This comes from the basic model of contextual bandits.

So, we have

$$\begin{aligned} \mathbb{E}_{x,r,a}[f(x,a) - r(a)]^2 &\geq \mathbb{E}_{x,r,a}[r(a) - f_{d^*}^*(x,a)]^2 + \mathbb{E}_{x,a}[f(x,a) - f_{d^*}^*(x,a)]^2 \\ &\geq \Delta + \sigma^2, \end{aligned}$$

where the last inequality comes from the separability assumption along with the assumption on the variance. Since the regressor $\hat{f}_j \in \mathcal{F}_j$, we have

$$\begin{aligned} \mathbb{E}_{x,r,a}[\hat{f}_j(x,a) - r(a)]^2 &\geq \mathbb{E}_{x,r,a}[r(a) - f_{d^*}^*(x,a)]^2 + \mathbb{E}_{x,a}[f(x,a) - f_{d^*}^*(x,a)]^2 \\ &\geq \Delta + \sigma^2. \end{aligned}$$

Now, using $C\sqrt{T}$ samples, we obtain

$$T_j \geq \Delta + \sigma^2 - C_3 T^{-1/4} \sqrt{\log(1/\delta)}$$

with high probability. Note that, provided $\Delta > c_0$ and

$$T > C \frac{\log^2(1/\delta)}{(\Delta - c_0)^4},$$

for a sufficiently large constant C , we have

$$T_j \geq \sigma^2 + c_0,$$

So we have

$$T_j > \sigma^2 + c_0$$

with high probability.

Choice of threshold: Based on the above calculation, we choose the threshold $\gamma = \sigma^2 + c_0$. Finally we can say that our model selection procedure succeeds with high probability provided,

$$T \geq C \max\left\{\frac{\log^2(1/\delta)}{(\Delta - c_0)^4}, \frac{1}{c_0^4} \log^2(|\mathcal{F}_M|/\delta), \log^2(1/\delta)\right\}.$$

9.5.6 Proof of Theorem 26

Case I: Realizable Class Consider $j \geq d^*$. Using calculations similar to the finite cardinality setting, we obtain

$$\mathbb{E}T_j \leq \sigma^2 + \xi_{\mathcal{F}_j, \delta}(C\sqrt{T}).$$

Hence, invoking Hoeffding's inequality, we obtain

$$T_j \leq \sigma^2 + \xi_{\mathcal{F}_j, \delta}(C\sqrt{T}) + C_1 T^{-1/4} \sqrt{\log(1/\delta)}$$

with probability at least $1 - 2\delta$.

Case II: Non-realizable Class We now consider the setting where $j < d^*$, meaning that $f_{d^*}^*$ does not lie in \mathcal{F}_j . In this case, we have

$$\mathbb{E}_{x,r,a}[\hat{f}_j(x,a) - r(a)]^2 - \inf_{f \in \mathcal{F}_j} \mathbb{E}_{x,r,a}[f(x,a) - r(a)]^2 \geq 0.$$

Continuing, we obtain

$$\begin{aligned} \mathbb{E}T_j &= \frac{1}{C\sqrt{T}} \mathbb{E}_{x_t, r_t, a_t} \sum_{t=1}^{C\sqrt{T}} [\hat{f}_j(x_t, a_t) - r_t(a_t)]^2 \\ &\geq \inf_{f \in \mathcal{F}_j} \mathbb{E}_{x,a} [f(x,a) - f_{d^*}^*(x,a)]^2, \\ &\geq \Delta, \end{aligned}$$

and hence

$$\begin{aligned} T_j &\geq \mathbb{E}T_j - \sqrt{\frac{32 \log(1/\delta)}{C\sqrt{T}}} \\ &\geq \Delta - \sqrt{\frac{32 \log(1/\delta)}{C\sqrt{T}}}. \end{aligned}$$

Fixing a Threshold: Based on the above calculation, if we select the threshold

$$\gamma = \sigma^2 + \xi_{\mathcal{F}_M, \delta}(C\sqrt{T}) + 2C_1 T^{-1/4} \sqrt{\log(1/\delta)}.$$

With the gap condition

$$\Delta \geq \sigma^2 + C_1 \left(\xi_{\mathcal{F}_M, \delta}(C\sqrt{T}) + T^{-1/4} \sqrt{\log(1/\delta)} \right),$$

we are guaranteed that $T_j > \gamma$ for the non realizable classes and $T_j \leq \gamma$ for the realizable classes with high probability. Hence exploiting the nested hypothesis class structure, the smallest index where $T_j \leq \gamma$ corresponds to selecting the correct model with high probability.

After obtaining the correct model class, the regret expression comes directly from [71] in the infinite function class setting.

9.5.7 Proof of Theorem 27

The proof directly follows by combining the proof of Theorem 23 and 26, and hence omitted.

Chapter 10

Model Selection for Reinforcement Learning with Function Approximation

We study the problem of model selection in Reinforcement Learning with function approximation. We restrict our attention to the framework of model based episodic *linear* MDP. In particular, similar to Chapter 8, we study both the *norm* and *sparsity/dimension* of the problem parameter as problem complexity and obtain model selection guarantees, meaning that the performance (regret) of our algorithm depends on the complexity of the problem instance, as compared to some predefined upper-bound on the complexity parameters. We first set the problem up, in the next section, and in the subsequent sections, take norm and sparsity as complexity measures and propose algorithms with regret guarantee.

10.1 Formulation and Linear Model

In this chapter, we consider an heterogeneous, episodic Markov decision process denoted by a tuple $M(\mathcal{S}, \mathcal{A}, H, \{\mathbb{P}_h\}_{h=1}^H, \{r_h\}_{h=1}^H)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, H is the length of each episode, \mathbb{P}_h is the transition probability at step h such that $\mathbb{P}_h(\cdot | s, a)$ is the distribution over states provided action a is taken for state s at step h , and $r_h : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the deterministic reward function at step h . A policy π is a collection of H functions $\{\pi_h : \mathcal{S} \rightarrow \mathcal{A}\}_{h \in [H]}$, where each of them maps a state s to an action a .

In each episode, an initial state s_1 is first picked by the environment. Then, at each step $h \in [H]$, the agent observes the state s_h , picks an action a_h according a certain policy π , receives a reward $r_h(s_h, a_h)$, and then transitions to the next state s_{h+1} , which is drawn from the distribution $P_h(\cdot | s_h, a_h)$. The episode ends when x_{H+1} is reached. For each $(s, a) \in \mathcal{S} \times \mathcal{A}$, we define action values $Q_h^\pi(s, a)$ and state values $V_h^\pi(s)$ as

$$Q_h^\pi(s, a) := r_h(s, a) + \mathbb{E} \left[\sum_{h'=h+1}^H r_{h'}(s_{h'}, \pi_{h'}(s_{h'})) \mid s_h = s, a_h = a \right]$$

$$V_h^\pi(s) := Q_h^\pi(s, \pi_h(s))$$

with $V_{H+1}^\pi(s) = 0$. The optimal value function $V_h^*(\cdot)$ and the optimal action-value function $Q_h^*(\cdot, \cdot)$ are given by $V_h^*(s) = \sup_\pi V_h^\pi(s)$ and $Q_h^*(s, a) = \sup_h Q_h^\pi(s, a)$. For brevity, we denote $[\mathbb{P}_h V_{h+1}](s, a) := \mathbb{E}_{s' \sim \mathbb{P}(\cdot|s, a)} V_{h+1}(s')$. With this notation, Bellman equations for policy π can be written as

$$Q_h^\pi(s, a) = (r_h + \mathbb{P}_h V_{h+1}^\pi)(s, a)$$

and the Bellman optimality equation is given by

$$Q_h^*(s, a) = (r_h + \mathbb{P}_h V_{h+1}^*)(s, a)$$

In the online learning setting, the agent interacts with the environment for K episodes to learn the unknown transition kernels $\{\mathbb{P}_h\}$ and hence to improve its policy. For each $k \geq 1$, the environment picks a starting state s_1^k , and the agent chooses a policy π^k that will be followed through the whole k -th episode. The objective is to design a learning algorithm that constructs a sequence $\{\pi^k\}$ that minimizes the total regret

$$R(K) = \sum_{k=1}^K \left[V_1^*(s_1^k) - V_1^{\pi^k}(s_1^k) \right].$$

In this chapter, we consider a special class of MDPs called *linear kernel MDPs* (a.k.a., linear mixture MDPs) [73]. Roughly speaking, it means that the transition kernel $\{\mathbb{P}_h\}$ can be represented as a linear function of a given feature map $\phi : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^d$. Formally, we have the following definition.

Definition 13 (Linear kernel MDP). $M(\mathcal{S}, \mathcal{A}, H, \{\mathbb{P}_h\}_{h=1}^H, \{r_h\}_{h=1}^H)$ is called an *heterogeneous, episodic B -bounded linear kernel MDP* if there exists a known feature mapping $\phi : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^d$ and an unknown vector $\theta_h \in \mathbb{R}^d$ with $\|\theta_h\|_2 \leq B$ such that

(i) For any state-action-state triplet $(s, a, s') \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$ and step $h \in [H]$, $\mathbb{P}(s'|s, a) = \langle \phi(s' | s, a), \theta_h \rangle$.

(ii) For any bounded function $V : \mathcal{S} \rightarrow [0, 1]$ and any tuple $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have $\|\phi_V(s, a)\|_2 \leq 1$, where $\phi_V(s, a) := \sum_{s' \in \mathcal{S}} \phi(s' | s, a) V(s')$.

In the learning problem, the vectors $\{\theta_h^*\}_{h=1}^H$ are unknown to the learner. The episodic MDP is parameterized by $\Theta^* = \{\theta_h^*\}_{h=1}^H$ and we denote the MDP by M_{Θ^*} . In what follows, we define two natural complexity measure of M_{Θ^*} : (a) $\|\theta_h^*\|$ for all $h \in [H]$, and (b) $\|\theta_h^*\|_0$ for all $h \in [H]$.

Notation We keep the notation consistent with the previous chapters. For positive integer r , by $[r]$, we denote the set of integers $\{1, 2, \dots, r\}$. Also $\|\cdot\|$ denotes ℓ_2 norm unless otherwise specified. Moreover, $\gamma_{\min}(A)$ denotes the minimum eigenvalue of the matrix A .

Algorithm 16: Adaptive Reinforcement Learning (norm)-ARL-NORM

- 1: **Input:** Over estimate of norm B , the initial phase length K_1 , Slack $\delta_1 = \delta > 0$
 - 2: Initial estimates: $b_h^{(1)} = \max\{B, 1\}$, for all $h \in [H]$.
 - 3: **for** epochs $i = 1, 2, \dots, N$ **do**
 - 4: Play UCRL-VTR⁺ with norm estimates $\{b_h^{(i)}\}_{h=1}^H$,
 until the end of epoch i (denoted by \mathcal{E}_i)
 - 5: At $k = \mathcal{E}_i$, refine estimate of $\|\theta_h^*\|$ as, $b_h^{(i+1)} = \max_{\theta \in \hat{\mathcal{C}}_{\mathcal{E}_i, h}} \|\theta\|$
 - 6: Set $K_{i+1} = 2K_i$, $\delta_{i+1} = \frac{\delta_i}{2}$.
 - 7: **end for**
-

10.2 Norm as Complexity measure

In this section, we define $\{\|\theta_h^*\|\}_{h=1}^H$ as a natural measure of complexity of the problem. Recall that the transition kernel is parameterized by θ_h^* , i.e., $\mathbb{P}(s' | s, a) = \langle \phi(s' | s, a), \theta_h^* \rangle$. Hence, if θ_h^* is close to 0 (hence $\|\theta_h^*\|$ is very small), the set of states s' for the next step will have a small cardinality. Similarly, when θ_h^* is high (with $\|\theta_h^*\|$ is large), the above-mentioned cardinality will be quite large. We mention here that, via the same set of arguments the sparsity (number of non-zero coordinates) of θ_h^* also serves a natural measure of complexity. We consider this in the subsequent section.

Here, we propose and analyze an algorithm, that adapts to the problem complexity $\|\theta_h^*\|$, and as a result, the regret obtained will dependent on $\|\theta_h^*\|$. In prior works, usually it is assumed that $\{\theta_h^*\}_{h=1}^H$ lies in a norm ball with known radius, i.e., $\|\theta_h^*\| \leq B$ (see [73]). This is a non-adaptive algorithm and the algorithm uses B as a proxy for the problem complexity, which can be a huge over-estimate. In sharp contrast, we start with this over estimate of $\|\theta_h^*\|$, and successively refine this estimate over multiple epochs. We show that this refinement strategy yields a consistent sequence of estimates of $\|\theta_h^*\|$, and as a consequence, our regret bound depends on $\|\theta_h^*\|$, not B .

Our idea: We take the UCRL-VTR⁺ of [74] as a blackbox, and successively refine the estimates of $\{\theta_h^*\}_{h=1}^H$ over epochs. Note that we take UCRL-VTR⁺ as our baseline algorithm since the regret of the same is near optimal. In particular UCRL-VTR⁺ tracks the variance of the value function (via updating a parameter $\bar{\sigma}_{k,h}$) along with a careful estimation of covariance matrix corresponding to the embeddings $\phi(\cdot)$.

We now proceed for the successive model refinement based model selection algorithm. The details are given in Algorithm 16. In particular we choose the doubling epoch. As mentioned, we use the model based RL algorithm of [74], namely UCRL-VTR⁺. We begin with some over estimate of $\{\theta_h^*\}_{h=1}^H$ (in particular, we assume $\|\theta_h^*\| \leq B$ for all $h \in [H]$), and over epochs, we refine this estimate. In particular, as seen in Algorithm 16, at the end of the epoch, based on the confidence set built, we choose the new estimate of $\{\theta_h^*\}_{h=1}^H$. We argue that this sequence of estimates is indeed consistent, and as a result, the regret depends on $\{\theta_h^*\}_{h=1}^H$.

Let us have a look how UCRL-VTR⁺ constructs the confidence ellipsoid $\hat{\mathcal{C}}_{k,h}$ for episode k .

From [74], we obtain the following

$$\hat{\mathcal{C}}_{k,h} = \left\{ \theta : \|\theta - \hat{\theta}_{k,h}\|_{\hat{\Sigma}_{k,h}} \leq \hat{\beta}_{k,h} \right\},$$

where $\hat{\theta}_{k,h}$ is an estimate of θ^* (a regularized least square estimator with adjusted variance) computed by UCRL-VTR⁺, and

$$\hat{\beta}_{k,h} = 8\sqrt{d \log(1 + k/\lambda) \log(4k^2 H/\delta)} + 4\sqrt{d} \log(4k^2 H/\delta) + \sqrt{\lambda} b_h,$$

where b_h is the current estimate of $\|\theta_h^*\|$ and λ is the regularization in the least squares problem. Also, we have

$$\hat{\Sigma}_{k,h} = \lambda I + \sum_{t=1}^k \bar{\sigma}_{t,h}^{-2} \phi_{V_{t,h+1}}(s_h^t, a_h^t) \phi_{V_{t,h+1}}^\top(s_h^t, a_h^t)$$

Then, [74] shows that

$$\mathbb{P} \left(\theta_h^* \in \hat{\mathcal{C}}_{k,h} \right) \geq 1 - 3\delta.$$

Now, let us present the main result of this section. We show that the sequence of estimators build via Algorithm 16 actually converges to the true norm estimates. We first make the following assumption on the embeddings:

Assumption 27. *We have, for all $h \in [H]$,*

$$\mathbb{E}_{k-1} \left[\bar{\sigma}_{k,h}^{-2} \phi_{V_{k,h+1}}(s_h^k, a_h^k) \phi_{V_{k,h+1}}^\top(s_h^k, a_h^k) \right] = \Sigma_c \succcurlyeq \rho_{\min} I,$$

where, $\rho_{\min} > 0$, and $\mathbb{E}_{k-1}[\cdot]$ is a conditional expectation, conditioned on the observation upto episode $k - 1$.

We emphasize that similar assumptions have featured in existing literature, e.g., see [67, 70]. Note that we had similar assumptions in Chapter 8. In fact removing such spectral assumption is an interesting open problem (see COLT open problem, [251]). Model selection without the above assumption is an interesting future direction.

Armed with the above assumption, we obtain the following result, which is crucial in understanding the rate at which the confidence ellipsoid of $\{\theta_h^*\}_{h=1}^H$ shrinks.

Lemma 59. *Consider the matrix $M_{k,h} = \lambda I + \sum_{t=1}^k \bar{\sigma}_{t,h}^{-2} \phi_{V_{t,h+1}}(s_h^t, a_h^t) \phi_{V_{t,h+1}}^\top(s_h^t, a_h^t)^\top$. For all $k \in \{\tau_{\min}(\delta), \dots, K\}$ and all $h \in [H]$, we have*

$$\gamma_{\min}(M_{k,h}) = \lambda I + \left(\sum_{t=1}^k \bar{\sigma}_{t,h}^{-2} \phi_{V_{t,h+1}}(s_h^t, a_h^t) \phi_{V_{t,h+1}}^\top(s_h^t, a_h^t)^\top \right) \geq \lambda + \rho_{\min} k/2,$$

with probability at least $1 - \delta$, where

$$\tau_{\min}(\delta) = \left(\frac{16d}{\rho_{\min}^2 H^2} + \frac{8\sqrt{d}}{3H\rho_{\min}} \right) \log(2dK/\delta).$$

We now show that the norm estimates produced by `ARL-norm` indeed converges to the true norm $\{\theta_h^*\}_{h=1}^H$ at an exponential rate with high probability.

Lemma 60. *With probability exceeding $1 - 8H\delta$, the sequence $\{b_h^{(i)}\}_{i=1}^\infty$ converges to $\|\theta_h^*\|$ at a rate $\mathcal{O}(\frac{i}{2^i})$, and we obtain $b_h^{(i)} \leq (c_1\|\theta_h^*\| + c_2)$ for all i , provided the length of the initial episode K_1 satisfies the following*

$$K_1 \geq \max\{\tau_{\min}(\delta) \log_2(1 + \frac{K}{K_1}), C_1 \left(\max\{p, q\} b_h^{(1)}\right)^2 d\}$$

where C_1 , and $p = 3\sqrt{\frac{\lambda}{\rho_{\min}}}$ and $q = \frac{C}{\sqrt{\rho_{\min}}} \left(\sqrt{\log(K_1/\lambda) \log(K_1 H/\delta_1)} + \log(K_1 H/\delta_1)\right)$

Hence, the sequence converges to $\|\theta_h^*\|$ at an exponential rate. Armed with the above lemma, we finally focus on the regret of `ARL (norm)`. We have the following:

Theorem 28. *Suppose the initial epoch length K_1 satisfies,*

$$K_1 > \max\{\tau_{\min}(\delta) \log_2(1 + \frac{K}{K_1}), C_1 \left(\max\{p, q\} b_h^{(1)}\right)^2 d\}$$

Then, with probability exceeding $1 - 10H\delta$, we have

$$R(K) \leq (c_1\sqrt{\|\bar{\theta}^*\|} + c_2)\tilde{\mathcal{O}} \left(\sqrt{(d^2 H^3 + dH^4)\lambda^{1/2}} \text{polylog}(K_1/\delta) \right) \text{polylog}\left(\frac{K}{K_1}\right)\sqrt{K},$$

where $\|\bar{\theta}^*\| = \max_{h \in [H]} \|\theta_h^*\|$.

Remark 50. If the algorithm is run for T rounds, with $T = KH$, we have

$$R(T) \leq \tilde{\mathcal{O}} \left(\sqrt{\|\bar{\theta}^*\|} \sqrt{(d^2 H^2 + dH^3)\lambda^{1/2}} \text{polylog}(K_1/\delta) \text{polylog}\left(\frac{K}{K_1}\right) \sqrt{T} \right)$$

with probability at least $1 - 8H\delta - 2\delta$.

Remark 51. The regret depends on $\{\|\theta_h^*\|\}_{h=1}^H$, instead of an upper-bound on the norm. Hence, our algorithm adapts to the problem complexity.

Remark 52. (Cost of model selection) Note that the regret bound here matches to that of the near-optimal regret of `UCRL-VTR+`, and so the model selection is (order-wise) free, in terms of regret.

10.3 Dimension as complexity measure

In this section, we consider the linear MDP, with dimension as a measure of complexity. We propose and analyze an adaptive algorithm that tailors to the sparsity of $\{\theta_h^*\}_{h=1}^H$. We denote the sparsity of the parameters as $\{d_h^*\}_{h=1}^H$, and define $\bar{d}^* = \max_h d_h^*$. In what follows, we consider estimating the

Algorithm 17: Adaptive Reinforcement Learning (dim)-ARL-Dim

- 1: **Input:** Initial Phase length k_0 , slack $\delta > 0$.
 - 2: $\hat{\theta}_0 = \mathbf{1}, T_{-1} = 0$
 - 3: **for** epochs $i = 0, 1, 2 \dots$ **do**
 - 4: $k_i = 4^i k_0, \quad \delta_i \leftarrow \frac{\delta}{2^i}$
 - 5: $\mathcal{D}_i := \{i : |\hat{\theta}_i| \geq (0.9)^{i+1}\}$
 - 6: **for** times $t \in \{k_{i-1} + 1, \dots, k_i\}$ **do**
 - 7: Play UCRL-VTR⁺(1, $\delta/2^i$) only restricted to $(\mathcal{D}_h^i)_{h \in [H]}$ co-ordinates.
 Here, $\delta/2^i$ is the slack parameter and 1 represents $\|\theta_h^*\| \leq 1$ for all $h \in [H]$.
 - 8: **end for**
 - 9: **for** times $t \in \{k_i + 1, \dots, k_i + 2^i \sqrt{k_0}\}$ **do**
 - 10: Play UCRL-VTR⁺(1, $\delta/2^i$) starting from where we left of in epoch $i - 1$.
 - 11: **end for**
 - 12: **end for**
-

sparsity parameter, d_h^* of the h -th parameter θ_h^* . For notational convenience, we drop the subscript h from d_h^* and θ_h^* .

As seen in Algorithm 17, our proposed scheme works over multiple epochs, and we use diminishing thresholds to estimate the support of θ^* . The algorithm is parameterized by k_0 , and the slack $\delta \in (0, 1)$. ARL-Dim proceeds in phases numbered $0, 1, \dots$, increasing with time. Each phase starts with a support estimation routine of θ^* , and is kept fixed throughout the phase.

Each phase i is divided into two blocks - (i) a regret minimization block lasting $4^i k_0$ time slots, (ii) followed by a support estimation phase lasting $2^i \lceil \sqrt{k_0} \rceil$ time slots. Thus, each phase i lasts for a total of $4^i k_0 + 2^i \lceil \sqrt{k_0} \rceil$ time slots.

At the beginning of phase $i \geq 0$, $\mathcal{D}_i \subseteq [d]$ denotes the set of ‘active coordinates’, namely the estimate of the non-zero coordinates of θ^* , and by notation, $\mathcal{D}_0 = [d]$. In the regret minimization block of phase i , a fresh instance of UCRL-VTR⁺ is spawned, with the dimensions restricted only to the set \mathcal{D}_i and probability parameter $\delta_i := \frac{\delta}{2^i}$.

On the other hand, in the support estimation phase, we continue running the UCRL-VTR⁺ algorithm in full d dimension, from the point where we left of in epoch $i - 1$. Concretely, one should think of the support estimation phases over epochs as a single run of UCRL-VTR⁺ in the full d dimension. We choose the support estimation phases length in such a way that this does not affect the final regret of ARL-Dim. Furthermore, choosing the error (slack) probability halving over epochs, we ensure that final regret bound holds with high probability. At the end of each phase $i \geq 0$, ARL-Dim forms an estimate $\hat{\theta}_{i+1}$ of θ^* . The active coordinate set \mathcal{D}_{i+1} , is then the coordinates of $\hat{\theta}_{i+1}$ with magnitude exceeding $(0.9)^{(i+1)}$. The pseudo-code is provided in Algorithm 17. By this careful choice of exploration periods and thresholds, we show that the estimated support of θ^* is equal to the true support, for all but finitely many phases. Thus after a finite number of epochs, the true support of θ^* *locks-in*, and thereafter the agent incurs the optimal regret that an oracle knowing the true support would incur. Hence, the extra regret we incur with respect to an oracle (which knows the support of θ^*) is the additive constant independent of time (depending on

the smallest non-zero value of θ^*).

10.3.1 Regret Guarantee

Suppose we choose the parameter

$$k_0 = 2\tau_{\min}^2(\delta) + O\left(\frac{d^2 \log^4(K^2 H/\delta)}{\rho_{\min}^2(0.9)^4}\right).$$

This ensures that the estimate of θ^* after the 0-th epoch, $\hat{\theta}_1$ satisfy

$$\mathbb{P}\left[\|\hat{\theta}_1 - \theta^*\|_{\infty} \geq (0.9)\right] \leq 2\delta$$

Furthermore, for all epochs $i \geq 2$, we have

$$\mathbb{P}\left[\|\hat{\theta}_i - \theta^*\|_{\infty} \geq (0.9)^i\right] \leq \frac{2\delta}{2^i},$$

where $\hat{\theta}_i$ is the estimate of θ^* obtained by collecting all samples of the support estimation phase of UCRL-VTR⁺ algorithm until the beginning of phase i . A formal statement of the above is deferred to the proof section.

We are now ready to present the main theorem of this section. Recall that $\bar{d}^* = \max_h d_h^*$.

Theorem 29. *Suppose Algorithm 17 is run with parameter k_0 chosen above and with slack $\delta > 0$ for K episodes. Then, with probability at least $1 - 3H\delta$, the regret upper bound is given by*

$$\begin{aligned} R(K) \leq & \frac{1}{\bar{\gamma}^{16}} \left(\tau_{\min}^2(\delta) + \frac{d^2 \log^4(K^2 H/\delta)}{\rho_{\min}^2(0.9)^4} \right) \\ & + \left(\sqrt{(\bar{d}^*)^2 H^3 + \bar{d}^* H^4 \text{polylog}(\lceil \frac{K}{k_0} \rceil)} \text{polylog}(K/\delta) \right) \sqrt{K} \end{aligned}$$

where $\bar{\gamma} = \min_{h \in [H]} \left\{ \min\{|\theta_h^*(j)| : \theta_h^*(j) \neq 0\} \right\}$, where $\theta_h^*(j)$ denotes the j -th coordinate of θ_h^* .

Remark 53. Note that we maintain a \sqrt{K} regret dependence, which is near optimal (see [74]). However, the regret depends on $\bar{\gamma}$, and hence it is a problem instance dependent bound.

Remark 54. (Dependence on $\bar{\gamma}$;) In the above regret bound, we haven't optimized over the dependence on $\bar{\gamma}$. Note that the $1/\bar{\gamma}^{16}$ dependence can be improved significantly by choosing an aggressive growing epoch length. For example, if we choose the support estimation period as $6^i \lceil \sqrt{k_0} \rceil$, the regret minimization period as $36^i k_0$ and the threshold as $(0.5)^i$, the dependence on $\bar{\gamma}$ can be significantly reduced to $1/\bar{\gamma}^{5.12}$. The support estimation period, regret minimization period and threshold selection may be kept as tuning parameters.

Remark 55. (Cost of model selection:) Note that the cost of model selection is the first term in the regret expression, which is independent of K , and hence a constant term.

10.4 Conclusion

We address the model selection problem for episodic MDP with linear function approximation. In particular, similar to the contextual bandits, we take norm and dimension as complexity parameter and propose algorithms that adapts to these. An immediate future work is to relax the linearity structure for RL, and obtain model selection for classes beyond linearity.

10.5 Proofs

10.5.1 Proof of Lemma 59

We use matrix Freedman inequality, and we critically use the bounded ness of $\phi_{V_{t,h+1}}(s_h^t, a_h^t)$. Note that a similar proof technique is used in [67] in the problem of contextual linear bandits.

For notational simplicity, we denote $\phi_{V_{t,h+1}}(s_h^t, a_h^t)$ by ϕ_{V_t} . In particular we have the bound $\|\phi_{V_t}\| \leq 1$ for all t , and from construction $\bar{\sigma}_{t,h} \geq \frac{H}{\sqrt{d}}$ for all t and h . With this, similar to [67], we construct the following matrix martingale,

$$Z_k = \sum_{t=1}^k \bar{\sigma}_{t,h}^{-2} \phi_{V_t} \phi_{V_t}^\top - k \Sigma_c,$$

with $Z_0 = 0$, and consider the martingale difference sequence, $Y_k = Z_k - Z_{k-1}$. Since, $\|\bar{\sigma}_{t,h}^{-1} \phi_{V_t}\| \leq \frac{\sqrt{d}}{H}$, we have

$$\|\Sigma_c\|_{op} = \|\mathbb{E}_{k-1} \bar{\sigma}_{k,h}^{-2} \phi_{V_k} \phi_{V_k}^\top\| \leq \frac{\sqrt{d}}{H},$$

and as a result

$$\|Y_k\|_{op} \leq \frac{\sqrt{d}}{H} + \frac{\sqrt{d}}{H} \leq 2 \frac{\sqrt{d}}{H}.$$

Furthermore, a simple calculation yields

$$\|\mathbb{E}_{k-1}[Y_k Y_k^\top]\|_{op} = \|\mathbb{E}_{k-1}[Y_k^\top Y_k]\|_{op} \leq 2 \frac{d}{H^2}.$$

Now, applying matrix Freedman inequality (Theorem 13 of [67]) with $R = 2\sqrt{d}/H$, $\omega^2 = 2td/H^2$, $u = \rho_{\min}t/2$, we obtain

$$\mathbb{P}\left(\|Z_k\| \geq \frac{\rho_{\min}k}{2}\right) \leq \frac{\delta}{k},$$

for all k satisfying $k \geq \left(\frac{16d}{\rho_{\min}^2 H^2} + \frac{8\sqrt{d}}{3H\rho_{\min}}\right) \log(2dK/\delta)$. Finally, using Weyl's inequality, we obtain the result.

10.5.2 Proof of Lemma 60

We consider doubling epochs, with initial epoch length K_1 and $K_i = 2^{i-1}K_1$ for $i \in \{1, \dots, N\}$. The number of epochs is N .

Let us consider the i -th epoch, and let $\hat{\theta}_{\mathcal{E}_i, h}$ be the least square estimate of θ_h^* at the end of epoch i . The confidence interval at the end of epoch i , is given by

$$\mathcal{C}_{\mathcal{E}_i, h} = \left\{ \theta \in \mathbb{R}^d : \|\theta - \hat{\theta}_{\mathcal{E}_i, h}\|_{\hat{\Sigma}_{\mathcal{E}_i, h}} \leq \hat{\beta}_{\mathcal{E}_i, h} \right\}.$$

Using Lemma 59, one can rewrite $\mathcal{C}_{\mathcal{E}_i, h}$ as

$$\mathcal{C}_{\mathcal{E}_i, h} = \left\{ \theta \in \mathbb{R}^d : \|\theta - \hat{\theta}_{\mathcal{E}_i, h}\| \leq \frac{\hat{\beta}_{\mathcal{E}_i, h}}{\sqrt{1 + \rho_{\min} K_i/2}} \right\},$$

with probability at least $1 - \delta_i = 1 - \delta/2^{i-1}$. Here we use the fact that $\gamma_{\min}(M_{K_i}, h) \geq 1 + \rho_{\min} K_i/2$, provided $K_i \geq \tau_{\min}(\delta/2^i)$. To ensure this condition, we take (the sufficient condition) $K_1 \geq \tau_{\min}(\delta)N$, where N is the number of epochs. Also, from the doubling principle, we obtain

$$\sum_{i=1}^N 2^{i-1} K_1 = K \Rightarrow N = \log_2 \left(1 + \frac{K}{K_1} \right).$$

Hence, with K_1 satisfying $K_1 \geq \tau_{\min}(\delta) \log_2 \left(1 + \frac{K}{K_1} \right)$, we ensure that $\gamma_{\min}(M_{K_i}, h) \geq 1 + \rho_{\min} K_i/2$.

Also, we know that $\theta_h^* \in \mathcal{C}_{\mathcal{E}_i, h}$ with probability at least $1 - 3\delta_i$. Hence, we obtain

$$\|\hat{\theta}_{\mathcal{E}_i, h}\| \leq \|\theta_h^*\| + 2 \frac{\hat{\beta}_{\mathcal{E}_i, h}}{\sqrt{1 + \rho_{\min} K_i/2}}$$

with probability at least $1 - 4\delta_i$. Recall from Algorithm 16 that at the end of the i -th epoch, we set the length $K_{i+1} = 2K_i$, and the estimate of $\|\theta^*\|$ is set to

$$b_h^{(i+1)} = \max_{\theta \in \mathcal{C}_{\mathcal{E}_i, h}} \|\theta\|.$$

From the definition of $\mathcal{C}_{\mathcal{E}_i, h}$, we obtain

$$\begin{aligned} b_h^{(i+1)} &= \|\hat{\theta}_{\mathcal{E}_i, h}\| + \frac{\hat{\beta}_{\mathcal{E}_i, h}}{\sqrt{1 + \rho_{\min} K_i/2}} \\ &\leq \|\theta_h^*\| + 3 \frac{\hat{\beta}_{\mathcal{E}_i, h}}{\sqrt{1 + \rho_{\min} K_i/2}} \end{aligned}$$

with probability exceeding $1 - 4\delta_i$. Let us look at

$$\hat{\beta}_{\mathcal{E}_{i,h}} = 8\sqrt{d \log(1 + K_i/\lambda) \log(4K_i^2 H/\delta_i)} + 4\sqrt{d} \log(4K_i^2 H/\delta_i) + \sqrt{\lambda} b_h^{(i)},$$

We now substitute $K_i = 2^{i-1} K_1$ and $\delta_i = \frac{\delta}{2^{i-1}}$. We obtain

$$\sqrt{1 + \rho_{\min} K_i/2} \geq 2^{\frac{i-2}{2}} \sqrt{\rho_{\min} K_1},$$

and

$$\begin{aligned} \frac{\hat{\beta}_{\mathcal{E}_{i,h}}}{\sqrt{1 + \rho_{\min} K_i/2}} &\leq \frac{\sqrt{\lambda} b_h^{(i)}}{2^{\frac{i-2}{2}} \sqrt{\rho_{\min} K_1}} \\ &\quad + C \frac{i}{2^{\frac{i-2}{2}} \sqrt{\rho_{\min} K_1}} \left(\sqrt{d \log(K_1/\lambda) \log(K_1 H/\delta_1)} + \sqrt{d} \log(K_1 H/\delta_1) \right). \end{aligned}$$

Using this, we obtain

$$\begin{aligned} b_h^{(i+1)} &\leq \|\theta_h^*\| + 3 \frac{\sqrt{\lambda} b_h^{(i)}}{2^{\frac{i-2}{2}} \sqrt{\rho_{\min} K_1}} \\ &\quad + C \frac{i}{2^{\frac{i-2}{2}} \sqrt{\rho_{\min} K_1}} \left(\sqrt{d \log(K_1/\lambda) \log(K_1 H/\delta_1)} + \sqrt{d} \log(K_1 H/\delta_1) \right) \\ &= \|\theta_h^*\| + b_h^{(i)} \frac{p}{2^{\frac{i-2}{2}}} \sqrt{\frac{1}{K_1}} + \frac{iq}{2^{\frac{i-2}{2}}} \sqrt{\frac{d}{K_1}}, \end{aligned}$$

with probability at least $1 - 4\delta_i$, where

$$p = 3\sqrt{\frac{\lambda}{\rho_{\min}}} \quad \text{and} \quad q = \frac{C}{\sqrt{\rho_{\min}}} \left(\sqrt{\log(K_1/\lambda) \log(K_1 H/\delta_1)} + \log(K_1 H/\delta_1) \right).$$

Hence, we obtain,

$$b_h^{(i+1)} - b_h^{(i)} \leq \|\theta_h^*\| + \left(1 - \frac{p}{2^{\frac{i-2}{2}}} \sqrt{\frac{1}{K_1}}\right) b_h^{(i)} + \frac{iq}{2^{\frac{i-2}{2}}} \sqrt{\frac{d}{K_1}}$$

By construction, $b_h^{(i)} \geq \|\theta_h^*\|$. Hence, provided $K_1 > \frac{4p^2}{2^i} > 2p^2$, we have

$$b_h^{(i+1)} - b_h^{(i)} \leq \frac{p}{2^{\frac{i-2}{2}}} \sqrt{\frac{1}{K_1}} \|\theta_h^*\| + \frac{iq}{2^{\frac{i-2}{2}}} \sqrt{\frac{d}{K_1}}$$

From the above expression,

$$\sup_i b_h^{(i)} < \infty$$

with probability greater than or equal to

$$1 - \sum_i 4\delta_i = 1 - \sum_i 4\delta/2^{i-1} \geq 1 - 8\delta.$$

. From the expression of $b_h^{(i)}$ and using the above fact in conjunction yield

$$\lim_{i \rightarrow \infty} b_h^{(i)} \leq \|\theta_h^*\|.$$

However, from construction $b_h^{(i)} \geq \|\theta_h^*\|$. Using this, along with the above equation, we obtain

$$\lim_{i \rightarrow \infty} b_h^{(i)} = \|\theta_h^*\|.$$

with probability exceeding $1 - 8\delta$. So, the sequence $\{b_h^{(1)}, b_h^{(2)}, \dots\}$ converges to $\|\theta_h^*\|$ with probability at least $1 - 8H\delta$ for all $h \in [H]$, and hence our successive refinement algorithm is consistent.

Rate of Convergence: Since

$$b_h^{(i)} - b_h^{(i-1)} = \tilde{O}\left(\frac{i}{2^i}\right), \quad (10.1)$$

with high probability, the rate of convergence of the sequence $\{b_h^{(i)}\}_{i=1}^{\infty}$ is exponential in the number of epochs.

Uniform upper bound on $b_h^{(i)}$ for all i : We now compute a uniform upper bound on $b_h^{(i)}$ for all i .

Consider the sequences $\left\{\frac{i}{2^{\frac{i-2}{2}}}\right\}_{i=1}^{\infty}$ and $\left\{\frac{1}{2^{\frac{i-2}{2}}}\right\}_{i=1}^{\infty}$, and let t_j and u_j denote the j -th term of the sequences respectively. It is easy to check that $\sup_i t_i = \sqrt{2}$ and $\sup_i u_i = 2$, and that the sequences $\{t_i\}_{i=1}^{\infty}$ and $\{u_i\}_{i=1}^{\infty}$ are convergent. With this new notation, we have

$$b_h^{(2)} \leq \|\theta^*\| + u_1 \frac{pb_h^{(1)}}{\sqrt{K_1}} + t_1 \frac{q\sqrt{d}}{\sqrt{K_1}}.$$

with probability exceeding $1 - 4\delta$. Similarly, for $b_h^{(3)}$, we have

$$\begin{aligned} b_h^{(3)} &\leq \|\theta^*\| + u_2 \frac{pb_h^{(2)}}{\sqrt{K_1}} + t_2 \frac{q\sqrt{d}}{\sqrt{K_1}} \\ &\leq \left(1 + u_2 \frac{p}{\sqrt{K_1}}\right) \|\theta^*\| + \left(u_1 u_2 \frac{p}{\sqrt{K_1}} \frac{p}{\sqrt{K_1}} b_h^{(1)}\right) + \left(t_1 u_2 \frac{p}{\sqrt{K_1}} \frac{q\sqrt{d}}{\sqrt{K_1}} + t_2 \frac{q\sqrt{d}}{\sqrt{K_1}}\right). \end{aligned}$$

with probability at least $1 - 4\delta - 2\delta = 1 - 6\delta$. Similarly, we write expressions for $b_h^{(4)}, b_h^{(5)}, \dots$. Now, provided $K_1 \geq C_1 \left(\max\{p, q\} b_h^{(1)} \right)^2 d$, where C_1 is a sufficiently large constant, the expression for $b_h^{(i)}$ can be upper-bounded as

$$b_i \leq (c_1 \|\theta_h^*\| + c_2), \quad (10.2)$$

with probability

$$\begin{aligned} &\geq 1 - 4\delta \left(1 + \frac{1}{2} + \frac{1}{4} + \dots \text{upto } i\text{-th term} \right) \\ &\geq 1 - 4\delta \left(1 + \frac{1}{2} + \frac{1}{4} + \dots \right) \\ &= 1 - 8\delta. \end{aligned}$$

Here c_1 and c_2 are constants, and are obtained from summing an infinite geometric series with decaying step size. We also use the fact that $b_1 \geq 1$, and the fact that $\delta_i = \frac{\delta_1}{2^{i-1}}$.

10.5.3 Proof of Theorem 28

Suppose we play Algorithm 16 for N epochs. The cumulative regret is given by

$$R(T) \leq \sum_{i=1}^N R(\delta_i, \bar{b}^{(i)})(K_i),$$

where $R(\delta_i, \bar{b}^{(i)})(K_i)$ is the cumulative regret of the UCRL-VTR⁺ in the i -th epoch. As seen (by tracking the dependence on $\{\|\theta_h^*\|\}_{h=1}^H$) in UCRL-VTR⁺, the cumulative regret of UCRL-VTR⁺ is given by

$$R(\delta_i, \bar{b}^{(i)})(K_i) = \tilde{O} \left(\sqrt{(d^2 H^3 + d H^4) \lambda^{1/2} \bar{b}^{(i)}} \sqrt{K_i} \text{polylog}(K_i/\delta_i) \right)$$

with probability at least $1 - \delta_i$, where $\bar{b}^{(i)} = \max_{h \in [H]} b_h^{(i)}$. Hence, we have

$$R(T) \leq \sum_{i=1}^N \sqrt{\bar{b}^{(i)}} R(\delta_i, 1)(K_i).$$

Using Lemma 60, we obtain, with probability at least $1 - 8H\delta$,

$$R(T) \leq (c_1 \sqrt{\|\bar{\theta}^*\|} + c_2) \sum_{i=1}^N R(\delta_i, 1)(K_i),$$

where $\|\bar{\theta}^*\| = \max_{h \in [H]} \|\theta_h^*\|$. With the doubling trick, we have

$$K_i = 2^{i-1} K_1, \quad \delta_i = \frac{\delta}{2^{i-1}}.$$

Substituting, we obtain

$$R(\delta_i, 1)(K_i) \leq \tilde{\mathcal{O}} \left(\sqrt{(d^2 H^2 + dH^3) \lambda^{1/2}} \sqrt{K_i} \text{poly}(i) \text{polylog}(K_1/\delta) \right)$$

with probability greater than $1 - \delta_i$.

Using the above expression, we obtain

$$R(K) \leq (c_1 \sqrt{\|\bar{\theta}^*\|} + c_2) \sum_{i=1}^N \tilde{\mathcal{O}} \left(\sqrt{(d^2 H^2 + dH^3) \lambda^{1/2}} \sqrt{K_i} \text{poly}(i) \text{polylog}(K_1/\delta) \right)$$

with probability at least $1 - 8H\delta - 2\delta$.

Also, from the doubling principle, we obtain

$$\sum_{i=1}^N 2^{i-1} K_1 = K \Rightarrow N = \log_2 \left(1 + \frac{K}{K_1} \right).$$

Using the above expression, we obtain

$$\begin{aligned} R(K) &\leq (c_1 \sqrt{\|\bar{\theta}^*\|} + c_2) \sum_{i=1}^N \tilde{\mathcal{O}} \left(\sqrt{(d^2 H^3 + dH^4) \lambda^{1/2}} \sqrt{K_i} \text{poly}(i) \text{polylog}(K_1/\delta) \right) \\ &\leq (c_1 \sqrt{\|\bar{\theta}^*\|} + c_2) \tilde{\mathcal{O}} \left(\sqrt{(d^2 H^3 + dH^4) \lambda^{1/2}} \text{polylog}(K_1/\delta) \right) \text{poly}(N) \sqrt{K} \\ &\leq (c_1 \sqrt{\|\bar{\theta}^*\|} + c_2) \tilde{\mathcal{O}} \left(\sqrt{(d^2 H^3 + dH^4) \lambda^{1/2}} \text{polylog}(K_1/\delta) \right) \text{polylog}\left(\frac{K}{K_1}\right) \sum_{i=1}^N \sqrt{K_i} \\ &\leq (c_1 \sqrt{\|\bar{\theta}^*\|} + c_2) \tilde{\mathcal{O}} \left(\sqrt{(d^2 H^3 + dH^4) \lambda^{1/2}} \text{polylog}(K_1/\delta) \right) \text{polylog}\left(\frac{K}{K_1}\right) \sum_{i=1}^N \sqrt{K_i} \\ &\leq (c_1 \sqrt{\|\bar{\theta}^*\|} + c_2) \tilde{\mathcal{O}} \left(\sqrt{(d^2 H^3 + dH^4) \lambda^{1/2}} \text{polylog}(K_1/\delta) \right) \text{polylog}\left(\frac{K}{K_1}\right) \sqrt{K}, \end{aligned}$$

where the last inequality follows from the fact that

$$\begin{aligned}
 \sum_{i=1}^N \sqrt{K_i} &= \sqrt{K_N} \left(1 + \frac{1}{\sqrt{2}} + \frac{1}{2} + \dots N\text{-th term} \right) \\
 &\leq \sqrt{K_N} \left(1 + \frac{1}{\sqrt{2}} + \frac{1}{2} + \dots \right) \\
 &= \frac{\sqrt{2}}{\sqrt{2}-1} \sqrt{K_N} \\
 &\leq \frac{\sqrt{2}}{\sqrt{2}-1} \sqrt{K}.
 \end{aligned}$$

The above regret bound holds with probability greater than or equal to $1 - 8H\delta - 2\delta$.

10.5.4 Proof of Theorem 29

First, we need the following set of results. Throughout the calculation, we take $\lambda = \Theta(1)$.

Lemma 61. *Suppose UCRL-VTR⁺ is run for M episodes, where $\tau_{\min}(\delta) \leq M < K$. If $M = O\left(\frac{d \log^2(K^2 H/\delta)}{\rho_{\min} \varepsilon^2}\right)$, and $\hat{\theta}^{(M)}$ is the estimate of θ^* , then we have*

$$\mathbb{P}[\|\hat{\theta}^{(M)} - \theta^*\|_{\infty} \geq \varepsilon] \leq 2\delta.$$

Proof. The proof of this comes from the analysis of UCRL-VTR⁺. Note that with Assumption 27 (and effectively Lemma 1), if we let UCRL-VTR⁺ run for M episodes, we obtain

$$\left\| \hat{\theta}^{(M)} - \theta^* \right\|_{\infty} \leq \left\| \hat{\theta}^{(M)} - \theta^* \right\| \leq \frac{8\sqrt{d \log(1 + M/\lambda) \log(4K^2 H/\delta)} + 4\sqrt{d} \log(4K^2 H/\delta) + \sqrt{\lambda}}{\sqrt{\lambda + \rho_{\min} M/2}} \quad (10.3)$$

with probability at least $1 - 2\delta$. Here $\tau_{\min}(\delta) = \left(\frac{16d}{\rho_{\min}^2 H^2} + \frac{8\sqrt{d}}{3H\rho_{\min}}\right) \log(2dK/\delta)$. Hence, with the choice of M , and using the fact that $M < K$, we get the lemma. \square

We shall now apply the lemma as follows. Denote by $\hat{\theta}_i$ to be the estimate of θ^* at the beginning of any phase i , using all the samples from random explorations in all phases less than or equal to $i - 1$.

Remark 56. The choice $\sqrt{k_0} = \tau_{\min}(\delta) + O\left(\frac{d \log^2(K^2 H/\delta)}{\rho_{\min} (0.9)^2}\right)$ (we have added the $\tau_{\min}(\delta)$ to make the calculations easier), ensures that

$$\mathbb{P}\left[\|\hat{\theta}^{(\lceil \sqrt{k_0} \rceil)} - \theta^*\|_{\infty} \geq 0.9\right] \leq 2\delta$$

Lemma 62. Suppose $k_0 = 2\tau_{\min}^2(\delta) + O\left(\frac{d^2 \log^4(K^2 H/\delta)}{\rho_{\min}^2(0.9)^4}\right)$. Then, for all phases $i \geq 2$,

$$\mathbb{P}\left[\|\widehat{\theta}_i - \theta^*\|_\infty \geq (0.9)^i\right] \leq \frac{2\delta}{2^i}, \quad (10.4)$$

where $\widehat{\theta}_i$ is the estimate of θ^* obtained by collecting all samples of the support estimation phase of UCRL-VTR⁺ algorithm until the beginning of phase i .

Proof. The above lemma follows directly from Lemma 61. Lemma 61 gives that if $\widehat{\theta}_i$ is computed by considering the samples of UCRL-VTR⁺ over at least M_i episodes, where $\tau_{\min}(\delta) \leq M_i \leq O\left(\frac{d \log^2(K^2 H 2^i/\delta)}{\rho_{\min}(0.9)^{2i}}\right)$ episodes, then we have equation 10.4. However, note that $\sqrt{k_0} = \tau_{\min}(\delta) + O\left(\frac{d \log^2(K^2 H/\delta)}{\rho_{\min}(0.9)^2}\right)$, and that we require $M_i \leq i(1.2)^i \sqrt{k_0}$. The proof is concluded if we show that at the beginning of phase $i \geq 2$, the total number of support estimation episodes performed by the algorithm exceeds $i(1.2)^i \lceil \sqrt{k_0} \rceil$. Notice that at the beginning of any phase $i \geq 2$, the total number of support estimation episodes

$$\begin{aligned} \sum_{j=0}^{i-1} 2^j \lceil \sqrt{k_0} \rceil &= \lceil \sqrt{k_0} \rceil (2^i - 1), \\ &\geq i(1.2)^i \lceil \sqrt{k_0} \rceil, \end{aligned}$$

where the last inequality holds for all $i \geq 2$. □

We now calculate the probability of the following event, which follows from a straightforward union bound.

Corollary 8.

$$\mathbb{P}\left[\bigcap_{i \geq 2} \left\{ \|\widehat{\theta}_i - \theta^*\|_\infty \leq (0.9)^i \right\}\right] \geq 1 - \delta.$$

Proof. This follows from a simple union bound as follows.

$$\begin{aligned} \mathbb{P}\left[\bigcap_{i \geq 2} \left\{ \|\widehat{\theta}_i - \theta^*\|_\infty \leq (0.9)^i \right\}\right] &= 1 - \mathbb{P}\left[\bigcup_{i \geq 2} \left\{ \|\widehat{\theta}_i - \theta^*\|_\infty \geq (0.9)^i \right\}\right], \\ &\geq 1 - \sum_{i \geq 2} \mathbb{P}\left[\|\widehat{\theta}_i - \theta^*\|_\infty \geq (0.9)^i\right], \\ &\geq 1 - \sum_{i \geq 2} \frac{2\delta}{2^i}, \\ &\geq 1 - \sum_{i \geq 2} \frac{\delta}{2^i}, \\ &= 1 - \delta. \end{aligned}$$

□

Proof of Theorem 29. We know from Corollary 8, that with probability at-least $1 - \delta$, for all phases $i \geq 2$, we have $\|\widehat{\theta}_i - \theta^*\|_\infty \leq (0.9)^i$. Call this event \mathcal{E} . Now, consider the phase $i(\gamma) := \max\left(2, \log_{(1/0.9)}\left(\frac{1}{\gamma}\right)\right)$. Now, when event \mathcal{E} holds, then for all phases $i \geq i(\gamma)$, \mathcal{D}_i is the correct set of d^* non-zero coordinates of θ^* . Thus, with probability at-least $1 - \delta$, the total regret upto time K episodes can be upper bounded as follows

$$\begin{aligned}
 R(K) \leq & \sum_{j=0}^{i(\gamma)-1} \left(4^j k_0 + 2^j \lceil \sqrt{k_0} \rceil\right) + \sum_{j=i(\gamma)}^{\lceil \log_4 \left(\frac{K}{k_0}\right) \rceil} \text{Regret}(\text{UCRL-VTR}^+(1, \delta_i; 4^j k_0)) \\
 & + \sum_{j=i(\gamma)}^{\lceil \log_4 \left(\frac{K}{k_0}\right) \rceil} \text{Regret}(\text{UCRL-VTR}^+(1, \delta_i; 2^j \lceil \sqrt{k_0} \rceil)). \quad (10.5)
 \end{aligned}$$

The term $\text{Regret}(\text{UCRL-VTR}^+(1, \delta, T))$ denotes the regret of the UCRL-VTR^+ algorithm, when run with $\|\theta^*\| \leq L$, and $\delta \in (0, 1)$ denotes the probability slack and K is the time horizon. Equation (10.5) follows, since the total number of phases is at-most $\mathcal{O}\left(\lceil \log_4 \left(\frac{K}{k_0}\right) \rceil\right)$. Also, note that the third term in Equation (10.5) is dominated by the second term, and hence handling the first two terms is sufficient.

We now use the near optimal regret bound of [74] to upper bound the second term of Equation (10.5). In particular, for all phases $i \in [i(\gamma), \lceil \log_4 \left(\frac{K}{k_0}\right) \rceil]$, we have, with probability at-least $1 - \delta$,

$$\text{Regret}(\text{UCRL-VTR}^+(1, \delta, K)) \leq \tilde{O}\left(\sqrt{(d^*)^2 H^3 + d^* H^4} \text{polylog}(K/\delta) \sqrt{K}\right).$$

Thus, we know that with probability at-least $1 - \sum_{i \geq 2} \delta_i \geq 1 - \delta$, for all phases $i \geq i(\gamma)$, the regret in the exploration phase satisfies

$$\text{Regret}(\text{UCRL-VTR}^+(1, \delta_i, 4^i k_0)) \quad (10.6)$$

$$\begin{aligned}
 & \leq \tilde{O}\left(\sqrt{(d^*)^2 H^3 + d^* H^4} \sqrt{4^i k_0} \text{poly}(i) \text{polylog}(K/\delta)\right) \\
 & \leq \tilde{O}\left(\sqrt{(d^*)^2 H^3 + d^* H^4} \sqrt{4^i k_0} \text{poly}(\lceil \log_4 \left(\frac{K}{k_0}\right) \rceil) \text{polylog}(K/\delta)\right) \\
 & \leq \mathcal{C}(H, K, \delta, d^*) \sqrt{4^i k_0} \quad (10.7)
 \end{aligned}$$

$$\begin{aligned}
 \mathcal{C}(H, K, \delta, d^*) & = \tilde{O}\left(\sqrt{(d^*)^2 H^3 + d^* H^4} \text{poly}(\lceil \log_4 \left(\frac{K}{k_0}\right) \rceil) \text{polylog}(K/\delta)\right) \\
 & = \tilde{O}\left(\sqrt{(d^*)^2 H^3 + d^* H^4} \text{polylog}(\lceil \left(\frac{K}{k_0}\right) \rceil) \text{polylog}(K/\delta)\right)
 \end{aligned}$$

Equation the above follows, by substituting $i \leq \mathcal{O}(\log_4 \left(\frac{K}{k_0}\right))$. We have with high probability, the (order-wise) regret is

$$\begin{aligned}
 R(K) &\leq 2k_0 4^{i(\gamma)} + 2 \sum_{j=0}^{\log_4 \left(\frac{K}{k_0}\right) + 1} \mathcal{C}(H, K, \delta, d^*) \sqrt{4^j k_0}, \\
 &\leq 2k_0 4^{i(\gamma)} + \mathcal{C}(H, K, \delta, d^*) \sum_{j=0}^{\log_4 \left(\frac{K}{k_0}\right) + 1} \sqrt{4^j k_0}, \\
 &\stackrel{(a)}{\leq} 2k_0 \frac{1}{\gamma^{16}} + 4\sqrt{K} + 4\sqrt{K} \mathcal{C}(H, K, \delta, d^*), \\
 &\leq \frac{1}{\gamma^{16}} \left(\tau_{\min}^2(\delta) + \frac{d^2 \log^4(K^2 H / \delta)}{\rho_{\min}^2 (0.9)^4} \right) + \mathcal{C}(H, K, \delta, d^*) \sqrt{K} \\
 &\leq \frac{1}{\gamma^{16}} \left(\tau_{\min}^2(\delta) + \frac{d^2 \log^4(K^2 H / \delta)}{\rho_{\min}^2 (0.9)^4} \right) \\
 &\quad + \left(\sqrt{(d^*)^2 H^3 + d^* H^4} \text{polylog}(\lceil \frac{K}{k_0} \rceil) \text{polylog}(K/\delta) \right) \sqrt{K}
 \end{aligned}$$

Step (a) follows from $4 \leq 1.1^{16}$.

Let us now compute the probability with which the above holds. Note that the support estimation procedure succeeds with probability at least $1 - H\delta$. Furthermore, we have 2 parallel runs of UCRL-VTR⁺, each of which succeeds with probability at least $1 - \sum_{i \geq 2} \delta_i \geq 1 - \delta$. Hence, with probability at least $1 - 3H\delta$, the above regret bound holds. □

Bibliography

- [1] H Brendan McMahan et al. “Communication-efficient learning of deep networks from decentralized data”. In: *arXiv preprint arXiv:1602.05629* (2016).
- [2] Jakub Konečný et al. “Federated optimization: distributed machine learning for on-device intelligence”. In: *arXiv preprint arXiv:1610.02527* (2016).
- [3] Brendan McMahan and Daniel Ramage. *Federated Learning: Collaborative Machine Learning without Centralized Training Data*. <https://research.googleblog.com/2017/04/federated-learning-collaborative.html>. 2017.
- [4] Jacob Devlin et al. “BERT: Pre-training of deep bidirectional transformers for language understanding”. In: *arXiv preprint arXiv:1810.04805* (2018).
- [5] Leslie Lamport, Robert Shostak, and Marshall Pease. “The Byzantine Generals Problem”. In: *ACM Trans. Program. Lang. Syst.* 4.3 (July 1982), pp. 382–401.
- [6] Virginia Smith et al. “Federated multi-task learning”. In: *Advances in Neural Information Processing Systems*. 2017, pp. 4424–4434.
- [7] Felix Sattler, Klaus-Robert Müller, and Wojciech Samek. “Clustered Federated Learning: Model-Agnostic Distributed Multi-Task Optimization under Privacy Constraints”. In: *arXiv preprint arXiv:1910.01991* (2019).
- [8] Yihan Jiang et al. “Improving federated learning personalization via model agnostic meta learning”. In: *arXiv preprint arXiv:1909.12488* (2019).
- [9] Yishay Mansour et al. “Three approaches for personalization with applications to federated learning”. In: *arXiv preprint arXiv:2002.10619* (2020).
- [10] Badrul M Sarwar et al. “Recommender systems for large-scale e-commerce: Scalable neighborhood formation using clustering”. In: *Proceedings of the fifth international conference on computer and information technology*. Vol. 1. 2002, pp. 291–324.
- [11] Qing Li and Byeong Man Kim. “Clustering approach for hybrid recommender system”. In: *Proceedings IEEE/WIC International Conference on Web Intelligence (WI 2003)*. IEEE. 2003, pp. 33–38.
- [12] Rich Caruana. “Multitask learning”. In: *Machine learning* 28.1 (1997), pp. 41–75.
- [13] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*. Vol. 48. Cambridge University Press, 2019.

- [14] Dan Alistarh et al. “The convergence of sparsified gradient methods”. In: *Advances in Neural Information Processing Systems*. 2018, pp. 5973–5983.
- [15] Sebastian U Stich, Jean-Baptiste Cordonnier, and Martin Jaggi. “Sparsified sgd with memory”. In: *Advances in Neural Information Processing Systems*. 2018, pp. 4447–4458.
- [16] Nikita Ivkin et al. “Communication-efficient distributed SGD with Sketching”. In: *arXiv preprint arXiv:1903.04488* (2019).
- [17] Dan Alistarh et al. “Communication-efficient stochastic gradient descent, with applications to neural networks”. In: (2017).
- [18] Lili Su and Nitin H Vaidya. “Fault-tolerant multi-agent optimization: optimal iterative distributed algorithms”. In: *Proceedings of the 2016 ACM symposium on principles of distributed computing*. ACM. 2016, pp. 425–434.
- [19] Jiashi Feng, Huan Xu, and Shie Mannor. “Distributed robust learning”. In: *arXiv preprint arXiv:1409.5937* (2014).
- [20] Sai Praneeth Karimireddy et al. “Error feedback fixes signsgd and other gradient compression schemes”. In: *International Conference on Machine Learning*. PMLR. 2019, pp. 3252–3261.
- [21] Dan Alistarh et al. “QSGD: Communication-efficient SGD via gradient quantization and encoding”. In: *Advances in Neural Information Processing Systems*. 2017, pp. 1709–1720.
- [22] Jeremy Bernstein et al. “signsgd with majority vote is communication efficient and Byzantine fault tolerant”. In: *arXiv preprint arXiv:1810.05291* (2018).
- [23] Dong Yin et al. “Byzantine-Robust Distributed Learning: Towards Optimal Statistical Rates”. In: *Proceedings of the 35th International Conference on Machine Learning*. Ed. by Jennifer Dy and Andreas Krause. Vol. 80. Proceedings of Machine Learning Research. Stockholmsmässan, Stockholm Sweden: PMLR, Oct. 2018, pp. 5650–5659.
- [24] Ohad Shamir, Nati Srebro, and Tong Zhang. “Communication-efficient distributed optimization using an approximate newton-type method”. In: *International conference on machine learning*. 2014, pp. 1000–1008.
- [25] Yuchen Zhang and Xiao Lin. “Disco: Distributed optimization for self-concordant empirical loss”. In: *International conference on machine learning*. 2015, pp. 362–370.
- [26] Shusen Wang et al. “GIANT: Globally improved approximate Newton method for distributed optimization”. In: *Advances in Neural Information Processing Systems*. 2018, pp. 2332–2342.
- [27] Rixon Crane and Fred Roosta. “DINGO: Distributed Newton-Type Method for Gradient-Norm Optimization”. In: *Advances in Neural Information Processing Systems*. 2019, pp. 9494–9504.
- [28] Fred Roosta et al. “Newton-MR: Newton’s Method Without Smoothness or Convexity”. In: *arXiv preprint arXiv:1810.00303* (2018).

- [29] Sashank J Reddi et al. “Aide: Fast and communication efficient distributed optimization”. In: *arXiv preprint arXiv:1608.06879* (2016).
- [30] Yudong Chen, Lili Su, and Jiaming Xu. “Distributed statistical machine learning in adversarial settings: Byzantine gradient descent”. In: *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 1.2 (2017), p. 44.
- [31] Dong Yin et al. “Defending Against Saddle Point Attack in Byzantine-Robust Distributed Learning”. In: *Proceedings of the 36th International Conference on Machine Learning*. Ed. by Kamalika Chaudhuri and Ruslan Salakhutdinov. Vol. 97. Proceedings of Machine Learning Research. Long Beach, California, USA: PMLR, Sept. 2019, pp. 7074–7084.
- [32] Avishek Ghosh et al. “Robust federated learning in a heterogeneous environment”. In: *arXiv preprint arXiv:1906.06629* (2019).
- [33] Peva Blanchard et al. “Byzantine-tolerant machine learning”. In: *arXiv:1703.02757* (2017).
- [34] Michal Derezhinski and Michael W Mahoney. “Distributed estimation of the inverse Hessian by determinantal averaging”. In: *Advances in Neural Information Processing Systems*. 2019, pp. 11401–11411.
- [35] Daniel Soudry and Yair Carmon. *No bad local minima: Data independent training error guarantees for multilayer neural networks*. 2016. arXiv: 1605.08361 [stat.ML].
- [36] Rong Ge, Chi Jin, and Yi Zheng. “No Spurious Local Minima in Nonconvex Low Rank Problems: A Unified Geometric Analysis”. In: *Proceedings of the 34th International Conference on Machine Learning*. Vol. 70. Proceedings of Machine Learning Research. International Convention Centre, Sydney, Australia: PMLR, June 2017, pp. 1233–1242.
- [37] Kenji Kawaguchi. “Deep Learning without Poor Local Minima”. In: *Advances in Neural Information Processing Systems*. Ed. by D. Lee et al. Vol. 29. Curran Associates, Inc., 2016, pp. 586–594.
- [38] Prateek Jain et al. *Global Convergence of Non-Convex Gradient Descent for Computing Matrix Squareroot*. 2017. arXiv: 1507.05854 [math.NA].
- [39] Ju Sun, Qing Qu, and John Wright. “A Geometric Analysis of Phase Retrieval”. In: *CoRR* abs/1602.06664 (2016). arXiv: 1602.06664.
- [40] Yann N Dauphin et al. “Identifying and attacking the saddle point problem in high-dimensional non-convex optimization”. In: *Advances in Neural Information Processing Systems*. Vol. 27. 2014, pp. 2933–2941.
- [41] Srinadh Bhojanapalli, Behnam Neyshabur, and Nathan Srebro. *Global Optimality of Local Search for Low Rank Matrix Recovery*. 2016. arXiv: 1605.07221 [stat.ML].
- [42] Ju Sun, Qing Qu, and John Wright. “Complete Dictionary Recovery Over the Sphere I: Overview and the Geometric Picture”. In: *IEEE Transactions on Information Theory* 63.2 (Feb. 2017), pp. 853–884.
- [43] Yurii Nesterov and Boris T Polyak. “Cubic regularization of Newton method and its global performance”. In: *Mathematical Programming* 108.1 (2006), pp. 177–205.

- [44] Jonas Moritz Kohler and Aurelien Lucchi. “Sub-sampled cubic regularization for non-convex optimization”. In: *International Conference on Machine Learning*. PMLR. 2017, pp. 1895–1904.
- [45] Zhe Wang et al. “Cubic regularization with momentum for nonconvex optimization”. In: *Uncertainty in Artificial Intelligence*. PMLR. 2020, pp. 313–322.
- [46] Chi Jin et al. “How to escape saddle points efficiently”. In: *International Conference on Machine Learning*. PMLR. 2017, pp. 1724–1732.
- [47] Chih-Chung Chang and Chih-Jen Lin. “LIBSVM: A library for support vector machines”. In: *ACM transactions on intelligent systems and technology (TIST)* 2.3 (2011), p. 27.
- [48] Robert W Harrison. “Phase problem in crystallography”. In: *JOSA a* 10.5 (1993), pp. 1046–1055.
- [49] C Fienup and J Dainty. “Phase retrieval and image reconstruction for astronomy”. In: *Image Recovery: Theory and Application* 231 (1987), p. 275.
- [50] Anwei Chai, Miguel Moscoso, and George Papanicolaou. “Array imaging using intensity-only measurements”. In: *Inverse Problems* 27.1 (2010), p. 015005.
- [51] Fajwel Fogel, Irène Waldspurger, and Alexandre d’Aspremont. “Phase retrieval for imaging problems”. In: *Mathematical programming computation* 8.3 (2016), pp. 311–335.
- [52] EM Bronshtein. “ ε -entropy of convex sets and functions”. In: *Siberian Mathematical Journal* 17.3 (1976), pp. 393–398.
- [53] Adityanand Guntuboyina and Bodhisattva Sen. “Covering numbers for convex functions”. In: *IEEE Transactions on Information Theory* 59.4 (2013), pp. 1957–1965.
- [54] Fuchang Gao and Jon A Wellner. “Entropy of Convex Functions on \mathbb{R}^d ”. In: *Constructive approximation* 46.3 (2017), pp. 565–592.
- [55] Alessandro Magnani and Stephen P Boyd. “Convex piecewise-linear fitting”. In: *Optimization and Engineering* 10.1 (2009), pp. 1–17.
- [56] Yuxin Chen and Emmanuel Candes. “Solving random quadratic systems of equations is nearly as easy as solving linear systems”. In: *Advances in Neural Information Processing Systems*. 2015, pp. 739–747.
- [57] Grigoris Paouris. “Small ball probability estimates for log-concave measures”. In: *Transactions of the American Mathematical Society* 364.1 (2012), pp. 287–308.
- [58] Mark Rudelson and Roman Vershynin. “Small ball probabilities for linear images of high-dimensional distributions”. In: *International Mathematics Research Notices* 2015.19 (2014), pp. 9594–9617.
- [59] Irene Waldspurger. “Phase Retrieval With Random Gaussian Sensing Vectors by Alternating Projections”. In: *IEEE Transactions on Information Theory* 64.5 (May 2018), pp. 3301–3312.

- [60] Teng Zhang. “Phase retrieval using alternating minimization in a batch setting”. In: *Applied and Computational Harmonic Analysis* (2019).
- [61] Praneeth Netrapalli, Prateek Jain, and Sujay Sanghavi. “Phase retrieval using alternating minimization”. In: *Advances in Neural Information Processing Systems*. 2013, pp. 2796–2804.
- [62] Tze Leung Lai and Herbert Robbins. “Asymptotically efficient adaptive allocation rules”. In: *Advances in applied mathematics* 6.1 (1985), pp. 4–22.
- [63] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. “Finite-time analysis of the multiarmed bandit problem”. In: *Machine learning* 47.2-3 (2002), pp. 235–256.
- [64] Wei Chu et al. “Contextual bandits with linear payoff functions”. In: *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*. 2011, pp. 208–214.
- [65] Varsha Dani, Thomas P Hayes, and Sham M Kakade. “Stochastic linear optimization under bandit feedback”. In: (2008).
- [66] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. “Improved algorithms for linear stochastic bandits”. In: *Advances in Neural Information Processing Systems*. 2011, pp. 2312–2320.
- [67] Niladri S Chatterji, Vidya Muthukumar, and Peter L Bartlett. “Osom: A simultaneously optimal algorithm for multi-armed and linear contextual bandits”. In: *arXiv preprint arXiv:1905.10040* (2019).
- [68] Alexandra Carpentier and Rémi Munos. “Bandit theory meets compressed sensing for high dimensional stochastic linear bandit”. In: *Artificial Intelligence and Statistics*. 2012, pp. 190–198.
- [69] Hamsa Bastani and Mohsen Bayati. “Online decision making with high-dimensional covariates”. In: *Operations Research* 68.1 (2020), pp. 276–294.
- [70] Dylan J Foster, Akshay Krishnamurthy, and Haipeng Luo. “Model selection for contextual bandits”. In: *Advances in Neural Information Processing Systems*. 2019, pp. 14714–14725.
- [71] David Simchi-Levi and Yunzong Xu. *Bypassing the Monster: A Faster and Simpler Optimal Algorithm for Contextual Bandits under Realizability*. 2020. arXiv: 2003.12699 [cs.LG].
- [72] Alekh Agarwal et al. “Contextual bandit learning with predictable rewards”. In: *Artificial Intelligence and Statistics*. PMLR. 2012, pp. 19–26.
- [73] Zeyu Jia et al. “Model-based reinforcement learning with value-targeted regression”. In: *Learning for Dynamics and Control*. PMLR. 2020, pp. 666–686.
- [74] Dongruo Zhou, Quanquan Gu, and Csaba Szepesvari. “Nearly Minimax Optimal Reinforcement Learning for Linear Mixture Markov Decision Processes”. In: *arXiv preprint arXiv:2012.08507* (2020).

- [75] Jonathan Lee et al. “Online model selection for reinforcement learning with function approximation”. In: *International Conference on Artificial Intelligence and Statistics*. PMLR. 2021, pp. 3340–3348.
- [76] Sebastian Caldas et al. “Leaf: A benchmark for federated settings”. In: *arXiv preprint arXiv:1812.01097* (2018).
- [77] Wayne S DeSarbo and William L Cron. “A maximum likelihood methodology for cluster-wise linear regression”. In: *Journal of classification* 5.2 (1988), pp. 249–282.
- [78] Yuekai Sun, Stratis Ioannidis, and Andrea Montanari. “Learning Mixtures of Linear Classifiers.” In: *ICML*. 2014, pp. 721–729.
- [79] Martin Zinkevich et al. “Parallelized stochastic gradient descent”. In: *Advances in neural information processing systems*. 2010, pp. 2595–2603.
- [80] Benjamin Recht et al. “Hogwild: A lock-free approach to parallelizing stochastic gradient descent”. In: *Advances in neural information processing systems*. 2011, pp. 693–701.
- [81] Mu Li et al. “Scaling distributed machine learning with the parameter server”. In: *11th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 14)*. 2014, pp. 583–598.
- [82] Andrew Hard et al. “Federated learning for mobile keyboard prediction”. In: *arXiv preprint arXiv:1811.03604* (2018).
- [83] Yue Zhao et al. “Federated Learning with Non-IID Data”. In: *arXiv:1806.00582* (2018).
- [84] Anit Kumar Sahu et al. “On the Convergence of Federated Optimization in Heterogeneous Networks”. In: *CoRR abs/1812.06127* (2018). arXiv: 1812.06127.
- [85] Liping Li et al. “RSA: Byzantine-Robust Stochastic Aggregation Methods for Distributed Learning from Heterogeneous Datasets”. In: *CoRR abs/1811.03761* (2018). arXiv: 1811.03761.
- [86] Felix Sattler et al. “Robust and Communication-Efficient Federated Learning from Non-IID Data”. In: *CoRR abs/1903.02891* (2019). arXiv: 1903.02891.
- [87] Xiang Li et al. “On the convergence of fedavg on non-iid data”. In: *arXiv:1907.02189* (2019).
- [88] Mehryar Mohri, Gary Sivek, and Ananda Theertha Suresh. “Agnostic Federated Learning”. In: *CoRR abs/1902.00146* (2019). arXiv: 1902.00146.
- [89] Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. “Personalized federated learning: A meta-learning approach”. In: *arXiv preprint arXiv:2002.07948* (2020).
- [90] Fei Chen et al. “Federated Meta-Learning with Fast Convergence and Efficient Communication”. In: *arXiv preprint arXiv:1802.07876* (2018).
- [91] Lei Xu and Michael I Jordan. “On convergence properties of the EM algorithm for Gaussian mixtures”. In: *Neural computation* 8.1 (1996), pp. 129–151.

- [92] Dar-Shyang Lee. “Effective Gaussian mixture learning for video background subtraction”. In: *IEEE transactions on pattern analysis and machine intelligence* 27.5 (2005), pp. 827–832.
- [93] Michel Wedel and Wagner A Kamakura. “Mixture regression models”. In: *Market segmentation*. Springer, 2000, pp. 101–124.
- [94] Dong Yin et al. “Learning mixtures of sparse linear regressions using sparse graph codes”. In: *IEEE Transactions on Information Theory* 65.3 (2018), pp. 1430–1451.
- [95] James R Fienup. “Phase retrieval algorithms: a comparison”. In: *Applied optics* 21.15 (1982), pp. 2758–2769.
- [96] Rick P Millane. “Phase retrieval in crystallography and optics”. In: *JOSA A* 7.3 (1990), pp. 394–411.
- [97] Constantinos Daskalakis, Christos Tzamos, and Manolis Zampetakis. “Ten steps of EM suffice for mixtures of two Gaussians”. In: *arXiv preprint arXiv:1609.00368* (2016).
- [98] Xinyang Yi, Constantine Caramanis, and Sujay Sanghavi. “Solving a mixture of many random linear equations by tensor decomposition and alternating minimization”. In: *arXiv preprint arXiv:1608.05749* (2016).
- [99] Sivaraman Balakrishnan, Martin J Wainwright, and Bin Yu. “Statistical guarantees for the EM algorithm: From population to sample-based analysis”. In: *The Annals of Statistics* 45.1 (2017), pp. 77–120.
- [100] Irene Waldspurger. “Phase retrieval with random gaussian sensing vectors by alternating projections”. In: *IEEE Transactions on Information Theory* 64.5 (2018), pp. 3301–3312.
- [101] Avishek Ghosh and Kannan Ramchandran. “Alternating Minimization Converges Super-Linearly for Mixed Linear Regression”. In: *arXiv preprint arXiv:2004.10914* (2020).
- [102] Prateek Jain, Praneeth Netrapalli, and Sujay Sanghavi. “Low-rank matrix completion using alternating minimization”. In: *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. ACM, 2013, pp. 665–674.
- [103] Xinyang Yi, Constantine Caramanis, and Sujay Sanghavi. “Alternating minimization for mixed linear regression”. In: *International Conference on Machine Learning*. 2014, pp. 613–621.
- [104] Bowei Yan, Mingzhang Yin, and Purnamrita Sarkar. “Convergence of gradient EM on multi-component mixture of Gaussians”. In: *Advances in Neural Information Processing Systems*. 2017, pp. 6956–6966.
- [105] Ofer Dekel et al. “Optimal distributed online prediction using mini-batches”. In: *Journal of Machine Learning Research* 13.Jan (2012), pp. 165–202.
- [106] Yann LeCun et al. “Gradient-based learning applied to document recognition”. In: *Proceedings of the IEEE* 86.11 (1998), pp. 2278–2324.

- [107] Alex Krizhevsky and Geoffrey Hinton. “Learning multiple layers of features from tiny images”. In: *Technical Report* (2009).
- [108] David Lopez-Paz and Marc’Aurelio Ranzato. “Gradient episodic memory for continual learning”. In: *Advances in Neural Information Processing Systems*. 2017, pp. 6467–6476.
- [109] Ian J Goodfellow et al. “An empirical investigation of catastrophic forgetting in gradient-based neural networks”. In: *arXiv preprint arXiv:1312.6211* (2013).
- [110] James Kirkpatrick et al. “Overcoming catastrophic forgetting in neural networks”. In: *Proceedings of the national academy of sciences* 114.13 (2017), pp. 3521–3526.
- [111] Roman Vershynin. “Introduction to the non-asymptotic analysis of random matrices”. In: *arXiv preprint arXiv:1011.3027* (2010).
- [112] Mark Rudelson and Roman Vershynin. “Hanson-Wright inequality and sub-Gaussian concentration”. In: *Electronic Communications in Probability* 18 (2013).
- [113] Ananda Theertha Suresh et al. “Distributed mean estimation with limited communication”. In: *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org. 2017, pp. 3329–3337.
- [114] Hongyi Wang et al. “Atomo: Communication-efficient learning via atomic sparsification”. In: *Advances in Neural Information Processing Systems*. 2018, pp. 9850–9861.
- [115] Wei Wen et al. “Terngrad: Ternary gradients to reduce communication in distributed deep learning”. In: *Advances in neural information processing systems*. 2017, pp. 1509–1519.
- [116] Venkata Gandikota, Raj Kumar Maity, and Arya Mazumdar. “vqSGD: Vector Quantized Stochastic Gradient Descent”. In: *arXiv preprint arXiv:1911.07971* (2019).
- [117] Jeremy Bernstein et al. “signSGD: Compressed optimisation for non-convex problems”. In: *arXiv preprint arXiv:1802.04434* (2018).
- [118] Nikko Strom. “Scalable distributed DNN training using commodity GPU cloud computing”. In: *Sixteenth Annual Conference of the International Speech Communication Association*. 2015.
- [119] Frank Seide et al. “1-bit stochastic gradient descent and its application to data-parallel distributed training of speech dnns”. In: *Fifteenth Annual Conference of the International Speech Communication Association*. 2014.
- [120] Yuchen Zhang et al. “Information-theoretic lower bounds for distributed statistical estimation with communication constraints”. In: *Advances in Neural Information Processing Systems*. 2013, pp. 2328–2336.
- [121] Samuel Horvath et al. “Stochastic Distributed Learning with Gradient Quantization and Variance Reduction”. In: *arXiv e-prints*, arXiv:1904.05115 (Apr. 2019), arXiv:1904.05115. arXiv: 1904.05115.
- [122] El Mahdi El Mhamdi, Rachid Guerraoui, and Sébastien Rouault. “The hidden vulnerability of distributed learning in byzantium”. In: *arXiv preprint arXiv:1802.07927* (2018).

- [123] Georgios Damaskinos et al. “AGGREGATHOR: Byzantine Machine Learning via Robust Gradient Aggregation”. In: (2019). Published in the Conference on Systems and Machine Learning (SysML) 2019, Stanford, CA, USA., p. 19.
- [124] *AWS News Blog*. <https://tinyurl.com/yxe4hu4w>. Accessed: 2019-10-08.
- [125] K. Lee et al. “Coded computation for multicore setups”. In: *2017 IEEE International Symposium on Information Theory (ISIT)*. June 2017, pp. 2413–2417.
- [126] Paolo Costa et al. “R2C2: A Network Stack for Rack-scale Computers”. In: *SIGCOMM 2015*. ACM - Association for Computing Machinery, Aug. 2015.
- [127] Simon Haykin. *An introduction to analog and digital communication*. John Wiley, 1994.
- [128] Sai Praneeth Karimireddy et al. “SCAFFOLD: Stochastic controlled averaging for on-device federated learning”. In: *arXiv preprint arXiv:1910.06378* (2019).
- [129] Prathamesh Mayekar and Himanshu Tyagi. “Limits on Gradient Compression for Stochastic Optimization”. In: *arXiv preprint arXiv:2001.09032* (2020).
- [130] Shusen Wang, Alex Gittens, and Michael W Mahoney. “Sketched ridge regression: Optimization perspective, statistical perspective, and model averaging”. In: *The Journal of Machine Learning Research* 18.1 (2017), pp. 8039–8088.
- [131] Vipul Gupta et al. “Oversketching newton: Fast convex optimization for serverless systems”. In: *arXiv preprint arXiv:1903.08857* (2019).
- [132] Mert Pilanci and Martin J Wainwright. “Newton sketch: A near linear-time optimization algorithm with linear-quadratic convergence”. In: *SIAM Journal on Optimization* 27.1 (2017), pp. 205–245.
- [133] Petros Drineas et al. “Fast approximation of matrix coherence and statistical leverage”. In: *Journal of Machine Learning Research* 13.Dec (2012), pp. 3475–3506.
- [134] Petros Drineas and Michael W Mahoney. “RandNLA: randomized numerical linear algebra”. In: *Communications of the ACM* 59.6 (2016), pp. 80–90.
- [135] Michael W Mahoney. “Randomized algorithms for matrices and data”. In: *Foundations and Trends® in Machine Learning* 3.2 (2011), pp. 123–224.
- [136] Swarm2. *Swarm User Documentation*. <https://people.cs.umass.edu/~swarm/index.php?n=Main.NewSwarmDoc>. Accessed: 2018-01-05. 2018.
- [137] Anna Choromanska et al. *The Loss Surfaces of Multilayer Networks*. 2015. arXiv: 1412.0233 [cs.LG].
- [138] Yi Xu, Rong Jin, and Tianbao Yang. “First-order stochastic algorithms for escaping from saddle points in almost linear time”. In: *arXiv preprint arXiv:1711.01944* (2017).
- [139] Zeyuan Allen-Zhu and Yuanzhi Li. “Neon2: Finding local minima via first-order oracles”. In: *arXiv preprint arXiv:1711.06673* (2017).
- [140] Avishek Ghosh et al. “Communication-efficient and byzantine-robust distributed learning”. In: *2020 Information Theory and Applications Workshop (ITA)*. IEEE, 2020, pp. 1–28.

- [141] Avishek Ghosh, Raj Kumar Maity, and Arya Mazumdar. “Distributed Newton Can Communicate Less and Resist Byzantine Workers”. In: *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*. 2020.
- [142] Jason D Lee et al. “Gradient descent converges to minimizers”. In: *arXiv:1602.04915* (2016).
- [143] Jason D Lee et al. “First-order methods almost always avoid saddle points”. In: *arXiv preprint arXiv:1710.07406* (2017).
- [144] Simon S Du et al. “Gradient descent can take exponential time to escape saddle points”. In: *arXiv preprint arXiv:1705.10412* (2017).
- [145] Coralia Cartis, Nicholas IM Gould, and Philippe L Toint. “Adaptive cubic regularisation methods for unconstrained optimization. Part I: motivation, convergence and numerical results”. In: *Mathematical Programming* 127.2 (2011), pp. 245–295.
- [146] Coralia Cartis, Nicholas IM Gould, and Philippe L Toint. “Adaptive cubic regularisation methods for unconstrained optimization. Part II: worst-case function-and derivative-evaluation complexity”. In: *Mathematical programming* 130.2 (2011), pp. 295–319.
- [147] N Tripuraneni et al. “Stochastic cubic regularization for fast nonconvex optimization”. In: *Advances in Neural Information Processing Systems*. 2018, pp. 2899–2908.
- [148] Dongruo Zhou, Pan Xu, and Quanquan Gu. “Stochastic variance-reduced cubic regularized Newton methods”. In: *International Conference on Machine Learning*. PMLR. 2018, pp. 5990–5999.
- [149] Zhe Wang et al. “Stochastic variance-reduced cubic regularization for nonconvex optimization”. In: *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR. 2019, pp. 2731–2740.
- [150] Yair Carmon and John C Duchi. “Gradient descent efficiently finds the cubic-regularized non-convex newton step”. In: *arXiv preprint arXiv:1612.00547* (2016).
- [151] Naman Agarwal et al. “Finding approximate local minima faster than gradient descent”. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 2017, pp. 1195–1199.
- [152] Qiyang Han and Jon A Wellner. “Multivariate convex regression: global risk bounds and adaptation”. In: *arXiv preprint arXiv:1601.06844* (2016).
- [153] Lauren A Hannah and David B Dunson. “Multivariate convex regression with adaptive partitioning”. In: *The Journal of Machine Learning Research* 14.1 (2013), pp. 3261–3294.
- [154] Gábor Balázs. “Convex regression: theory, practice, and applications”. PhD thesis. University of Alberta, 2016.
- [155] Jerry Ladd Prince and Alan S Willsky. “Reconstructing convex sets from support line measurements”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12.4 (1990), pp. 377–389.

- [156] Jens Gregor and Fernando R Rannou. “Three-dimensional support function estimation and application for projection magnetic resonance imaging”. In: *International journal of imaging systems and technology* 12.1 (2002), pp. 43–50.
- [157] Richard J. Gardner. *Geometric Tomography*. 2nd ed. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2006.
- [158] Adityanand Guntuboyina. “Optimal rates of convergence for convex set estimation from support functions”. In: *The Annals of Statistics* 40.1 (2012), pp. 385–411.
- [159] Yong Sheng Soh and Venkat Chandrasekaran. “Fitting Tractable Convex Sets to Support Function Evaluations”. In: *arXiv preprint arXiv:1903.04194* (2019).
- [160] Sara A van de Geer. *Regression analysis and empirical processes*. Vol. 45. CWI Tract. Stichting Mathematisch Centrum, Centrum voor Wiskunde en Informatica, Amsterdam, 1988, pp. vi+161.
- [161] Matthew Fickus et al. “Phase retrieval from very few measurements”. In: *Linear Algebra and its Applications* 449 (2014), pp. 475–499.
- [162] Emilio Seijo and Bodhisattva Sen. “Nonparametric least squares estimation of a multivariate convex regression function”. In: *The Annals of Statistics* 39.3 (2011), pp. 1633–1657.
- [163] Eunji Lim and Peter W Glynn. “Consistency of multidimensional convex regression”. In: *Operations Research* 60.1 (2012), pp. 196–208.
- [164] Rahul Mazumder et al. “A computational framework for multivariate convex regression and its variants”. In: *Journal of the American Statistical Association* 114.525 (2019), pp. 318–331.
- [165] R Gerchberg and W Saxton. “A practical algorithm for the determination of phase from image and diffraction plane pictures”. In: *Optik* 35 (1972), pp. 237–246.
- [166] Evelyn ML Beale and Roderick JA Little. “Missing values in multivariate analysis”. In: *Journal of the Royal Statistical Society: Series B (Methodological)* 37.1 (1975), pp. 129–145.
- [167] Herman O Hartley. “Maximum likelihood estimation from incomplete data”. In: *Biometrics* 14.2 (1958), pp. 174–194.
- [168] Michael I Jordan and Robert A Jacobs. “Hierarchical mixtures of experts and the EM algorithm”. In: *Neural computation* 6.2 (1994), pp. 181–214.
- [169] Arun Tejasvi Chaganty and Percy Liang. “Spectral experts for estimating mixtures of linear regressions”. In: *International Conference on Machine Learning*. 2013, pp. 1040–1048.
- [170] CF Jeff Wu. “On the convergence properties of the EM algorithm”. In: *The Annals of statistics* 11.1 (1983), pp. 95–103.
- [171] Paul Tseng. “An analysis of the EM algorithm and entropy-like proximal point methods”. In: *Mathematics of Operations Research* 29.1 (2004), pp. 27–44.

- [172] Stéphane Chrétien and Alfred O Hero. “On EM algorithms and their proximal generalizations”. In: *ESAIM: Probability and Statistics* 12 (2008), pp. 308–326.
- [173] Yuchen Zhang et al. “Spectral methods meet EM: A provably optimal algorithm for crowd-sourcing”. In: *Advances in neural information processing systems*. 2014, pp. 1260–1268.
- [174] Kai Zhong, Prateek Jain, and Inderjit S Dhillon. “Mixed linear regression with multiple components”. In: *Advances in neural information processing systems*. 2016, pp. 2190–2198.
- [175] Chi Jin et al. “Local maxima in the likelihood of Gaussian mixture models: Structural results and algorithmic consequences”. In: *Advances in neural information processing systems*. 2016, pp. 4116–4124.
- [176] Constantinos Daskalakis, Christos Tzamos, and Manolis Zampetakis. “Ten Steps of EM Suffice for Mixtures of Two Gaussians”. In: *30th Annual Conference on Learning Theory*. 2017.
- [177] Ji Xu, Daniel J Hsu, and Arian Maleki. “Global analysis of expectation maximization for mixtures of two Gaussians”. In: *Advances in Neural Information Processing Systems*. 2016, pp. 2676–2684.
- [178] Jeongyeol Kwon et al. “Global Convergence of EM Algorithm for Mixtures of Two Component Linear Regression”. In: *arXiv preprint arXiv:1810.05752v3* (2018).
- [179] Daniel Hsu and Sham M Kakade. “Learning mixtures of spherical Gaussians: moment methods and spectral decompositions”. In: *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*. ACM. 2013, pp. 11–20.
- [180] Constantine Caramanis Xinyang Yi and Sujay Sanghavi. “Solving a Mixture of Many Random Linear Equations by Tensor Decomposition and Alternating Minimization”. In: (2016).
- [181] Yan Shuo Tan and Roman Vershynin. “Phase Retrieval via Randomized Kaczmarz: Theoretical Guarantees”. In: *arXiv e-prints*, arXiv:1706.09993 (June 2017), arXiv:1706.09993. arXiv: 1706.09993 [math.NA].
- [182] Georges M Tallis. “The moment generating function of the truncated multi-normal distribution”. In: *Journal of the Royal Statistical Society: Series B (Methodological)* 23.1 (1961), pp. 223–229.
- [183] Yanyao Shen and Sujay Sanghavi. “Iterative Least Trimmed Squares for Mixed Linear Regression”. In: *arXiv preprint arXiv:1902.03653* (2019).
- [184] Guillaume Lecué and Shahar Mendelson. “Minimax rate of convergence and the performance of ERM in phase recovery”. In: *arXiv preprint arXiv:1311.5024* (2013).
- [185] Yong Sheng Soh. “Fitting Convex Sets to Data: Algorithms and Applications”. PhD thesis. California Institute of Technology, 2019.
- [186] Raghunandan H Keshavan, Andrea Montanari, and Sewoong Oh. “Matrix completion from a few entries”. In: *IEEE transactions on information theory* 56.6 (2010), pp. 2980–2998.

- [187] Hanie Sedghi, Majid Janzamin, and Anima Anandkumar. “Provable tensor methods for learning mixtures of generalized linear models”. In: *Proceedings of Machine Learning Research* 51 (2014), pp. 1223–1231.
- [188] Kai Zhong et al. “Recovery guarantees for one-hidden-layer neural networks”. In: *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org. 2017, pp. 4140–4149.
- [189] Amit Daniely, Sivan Sabato, and Shai S Shwartz. “Multiclass learning approaches: A theoretical comparison with implications”. In: *Advances in Neural Information Processing Systems*. 2012, pp. 485–493.
- [190] Dmitry Babichev and Francis Bach. “Slice inverse regression with score functions”. In: *Electronic Journal of Statistics* 12.1 (2018), pp. 1507–1543.
- [191] Surbhi Goel et al. “Reliably learning the ReLU in polynomial time”. In: *Conference on Learning Theory*. PMLR. 2017, pp. 1004–1042.
- [192] Adityanand Guntuboyina and Bodhisattva Sen. “Nonparametric shape-restricted regression”. In: *Statistical Science* 33.4 (2018), pp. 568–594.
- [193] Adityanand Guntuboyina and Bodhisattva Sen. “Global risk bounds and adaptation in univariate convex regression”. In: *Probability Theory and Related Fields* 163.1-2 (2015), pp. 379–411.
- [194] Pierre C. Bellec. “Sharp oracle inequalities for Least Squares estimators in shape restricted regression”. In: *Ann. Statist.* 46.2 (Apr. 2018), pp. 745–780.
- [195] Andres M Medina and Mehryar Mohri. “Learning theory and algorithms for revenue optimization in second price auctions with reserve”. In: *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*. 2014, pp. 262–270.
- [196] Jamie Morgenstern and Tim Roughgarden. “Learning simple auctions”. In: *Conference on Learning Theory*. 2016, pp. 1298–1318.
- [197] Ker-Chau Li. “Sliced inverse regression for dimension reduction”. In: *Journal of the American Statistical Association* 86.414 (1991), pp. 316–327.
- [198] Joel L Horowitz. *Semiparametric and nonparametric methods in econometrics*. Vol. 12. Springer, 2009.
- [199] Robert A Jacobs et al. “Adaptive mixtures of local experts.” In: *Neural computation* 3.1 (1991), pp. 79–87.
- [200] Seniha Esen Yuksel, Joseph N Wilson, and Paul D Gader. “Twenty years of mixture of experts”. In: *IEEE transactions on neural networks and learning systems* 23.8 (2012), pp. 1177–1193.
- [201] Teng Zhang. “Phase Retrieval by Alternating Minimization with Random Initialization”. In: *arXiv preprint arXiv:1812.01255* (2018).

- [202] Avishek Ghosh et al. “Max-Affine Regression II: Universal parameter estimation for small-ball designs [Online] <https://tinyurl.com/tae7z5r>”. In: 2019.
- [203] T Tony Cai, Xiaodong Li, and Zongming Ma. “Optimal rates of convergence for noisy sparse phase retrieval via thresholded Wirtinger flow”. In: *The Annals of Statistics* 44.5 (2016), pp. 2221–2251.
- [204] Marek Kanter and Harold Propp. “Reduction of variance for Gaussian densities via restriction to convex sets”. In: *Journal of Multivariate Analysis* 7.1 (1977), pp. 74–81.
- [205] Santosh S Vempala. “Learning convex concepts from Gaussian distributions with PCA”. In: *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. IEEE. 2010, pp. 124–130.
- [206] Michel Ledoux. *The concentration of measure phenomenon*. 89. American Mathematical Soc., 2001.
- [207] Barry James, Kang James, and Yongcheng Qi. “Limit distribution of the sum and maximum from multivariate Gaussian sequences”. In: *Journal of multivariate analysis* 98.3 (2007), pp. 517–532.
- [208] Yi Yu, Tengyao Wang, and Richard J Samworth. “A useful variant of the Davis–Kahan theorem for statisticians”. In: *Biometrika* 102.2 (2014), pp. 315–323.
- [209] Andrew F Siegel. “A surprising covariance involving the minimum of multivariate normal variables”. In: *Journal of the American Statistical Association* 88.421 (1993), pp. 77–80.
- [210] Jun S Liu. “Siegel’s formula via Stein’s identities”. In: *Statistics & Probability Letters* 21.3 (1994), pp. 247–251.
- [211] BG Manjunath and Stefan Wilhelm. *Moments calculation for the double truncated multivariate normal density*. 2009.
- [212] Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. 1st. Springer Publishing Company, Incorporated, 2008.
- [213] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.
- [214] Vladimir N Vapnik and Aleksei Yakovlevich Chervonenkis. “The uniform convergence of frequencies of the appearance of events to their probabilities”. In: *Doklady Akademii Nauk*. Vol. 181. 4. Russian Academy of Sciences. 1968, pp. 781–783.
- [215] Monika Csikos, Andrey Kupavskii, and Nabil H Mustafa. “Optimal bounds on the VC-dimension”. In: *arXiv preprint arXiv:1807.07924* (2018).
- [216] Gilles Hargé. “A convex/log-concave correlation inequality for Gaussian measure and an application to abstract Wiener spaces”. In: *Probability theory and related fields* 130.3 (2004), pp. 415–440.
- [217] Yaozhong Hu. “Itô-Wiener chaos expansion with exact residual and correlation, variance inequalities”. In: *Journal of Theoretical Probability* 10.4 (1997), pp. 835–848.

- [218] John Burkardt. “The truncated normal distribution”. In: (2014).
- [219] Aad W van der Vaart and Jon A. Wellner. *Weak convergence and empirical processes*. Springer Series in Statistics. With applications to statistics. Springer-Verlag, New York, 1996, pp. xvi+508.
- [220] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*. Vol. 47. Cambridge University Press, 2018.
- [221] Avishek Ghosh et al. “Max-Affine Regression I: Parameter Estimation for Gaussian Designs”. In: <https://tinyurl.com/spyjmgv> (2019).
- [222] Avishek Ghosh et al. “Max-Affine Regression: Provable, Tractable, and Near-Optimal Statistical Estimation”. In: *arXiv preprint arxiv:1906.09255* (2019).
- [223] Larry Wasserman. *All of nonparametric statistics*. Springer Science & Business Media, 2006.
- [224] Emmanuel J Candes, Xiaodong Li, and Mahdi Soltanolkotabi. “Phase retrieval via Wirtinger flow: Theory and algorithms”. In: *IEEE Transactions on Information Theory* 61.4 (2015), pp. 1985–2007.
- [225] Yonina C Eldar and Shahar Mendelson. “Phase retrieval: Stability and recovery guarantees”. In: *Applied and Computational Harmonic Analysis* 36.3 (2014), pp. 473–494.
- [226] John C Duchi and Feng Ruan. “Solving (most) of a set of quadratic equalities: Composite optimization for robust phase retrieval”. In: *arXiv preprint arXiv:1705.02356* (2017).
- [227] Anthony Carbery and James Wright. “Distributional and L^q norm inequalities for polynomials over convex bodies in \mathbb{R}^n ”. In: *Mathematical research letters* 8.3 (2001), pp. 233–248.
- [228] Mark Rudelson and Roman Vershynin. “The Littlewood–Offord problem and invertibility of random matrices”. In: *Advances in Mathematics* 218.2 (2008), pp. 600–633.
- [229] Jean-Yves Audibert and Sébastien Bubeck. “Best arm identification in multi-armed bandits”. In: 2010.
- [230] Alekh Agarwal et al. “Corralling a band of bandit algorithms”. In: *arXiv:1612.06246* (2016).
- [231] Avishek Ghosh, Sayak Ray Chowdhury, and Aditya Gopalan. “Misspecified linear bandits”. In: *Thirty-First AAAI Conference on Artificial Intelligence*. 2017.
- [232] Andrea Locatelli and Alexandra Carpentier. “Adaptivity to Smoothness in X-armed bandits”. In: *Conference on Learning Theory*. 2018, pp. 1463–1492.
- [233] Akshay Krishnamurthy, Zhiwei Steven Wu, and Vasilis Syrgkanis. “Semiparametric contextual bandits”. In: *arXiv preprint arXiv:1803.04204* (2018).
- [234] Thodoris Lykouris, Karthik Sridharan, and Éva Tardos. “Small-loss bounds for online learning with partial information”. In: *arXiv preprint arXiv:1711.03639* (2017).

- [235] Peter Auer, Pratik Gajane, and Ronald Ortner. “Adaptively tracking the best arm with an unknown number of distribution changes”. In: *European Workshop on Reinforcement Learning*. Vol. 14. 2018, p. 375.
- [236] Haipeng Luo and Robert E Schapire. “Achieving all with no parameters: Adanormalhedge”. In: *Conference on Learning Theory*. 2015, pp. 1286–1304.
- [237] Brendan McMahan and Jacob Abernethy. “Minimax optimal algorithms for unconstrained linear optimization”. In: *Advances in Neural Information Processing Systems*. 2013, pp. 2724–2732.
- [238] Francesco Orabona. “Simultaneous model selection and optimization through parameter-free stochastic learning”. In: *Advances in Neural Information Processing Systems*. 2014, pp. 1116–1124.
- [239] Ashok Cutkosky and Kwabena Boahen. “Online learning without prior information”. In: *arXiv preprint arXiv:1703.02629* (2017).
- [240] Vladimir Vapnik. *Estimation of dependences based on empirical data*. Springer Science & Business Media, 2006.
- [241] Lucien Birgé, Pascal Massart, et al. “Minimum contrast estimators on sieves: exponential bounds and rates of convergence”. In: *Bernoulli* 4.3 (1998), pp. 329–375.
- [242] Gábor Lugosi, Andrew B Nobel, et al. “Adaptive model selection using empirical complexities”. In: *The Annals of Statistics* 27.6 (1999), pp. 1830–1864.
- [243] Sylvain Arlot, Peter L Bartlett, et al. “Margin-adaptive model selection in statistical learning”. In: *Bernoulli* 17.2 (2011), pp. 687–713.
- [244] Vladimir Cherkassky. “Model complexity control and statistical learning theory”. In: *Natural computing* 1.1 (2002), pp. 109–133.
- [245] Luc Devroye, László Györfi, and Gábor Lugosi. *A probabilistic theory of pattern recognition*. Vol. 31. Springer Science & Business Media, 2013.
- [246] Peter Auer. “Using confidence bounds for exploitation-exploration trade-offs”. In: *Journal of Machine Learning Research* 3.Nov (2002), pp. 397–422.
- [247] Olivier Chapelle and Yi Chang. “Yahoo! Learning to Rank Challenge Overview”. In: *Proceedings of the 2010 International Conference on Yahoo! Learning to Rank Challenge - Volume 14*. YLRC’10. Haifa, Israel: JMLR.org, 2010, pp. 1–24.
- [248] Aldo Pacchiano et al. “Model selection in contextual stochastic bandit problems”. In: *arXiv preprint arXiv:2003.01704* (2020).
- [249] Michael Krikheli and Amir Leshem. “Finite sample performance of linear least squares estimators under sub-Gaussian martingale difference noise”. In: *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE. 2018, pp. 4444–4448.

- [250] Jeongyeol Kwon and Constantine Caramanis. “The EM Algorithm gives Sample-Optimality for Learning Mixtures of Well-Separated Gaussians”. In: *Conference on Learning Theory*. PMLR. 2020, pp. 2425–2487.
- [251] Dylan J Foster, Akshay Krishnamurthy, and Haipeng Luo. “Open Problem: Model Selection for Contextual Bandits”. In: *Conference on Learning Theory*. PMLR. 2020, pp. 3842–3846.