UC Davis

UC Davis Electronic Theses and Dissertations

Title

Some Contributions to Uncertainty Quantification and Fairness in Statistical Machine Learning

Permalink

https://escholarship.org/uc/item/2zz8n7pp

Author

Wang, Jue

Publication Date

2024

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

Some Contributions to Uncertainty Quantification and Fairness in Statistical Machine Learning

By

Jue Wang

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

STATISTICS

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Thomas C. M. Lee

Christiana Drake

Jairo Fúquene-Patiño

Committee in Charge

2024

© Jue Wang, 2024. All rights reserved.

To the past two thousand days and nights.

Contents

Abstract v

Some Contributions to Uncertainty Quantification and Fairness in Statistical Machine Learning v

Acknowledgments vi

Chapter 1. Introduction 1

Chapter 2. Weighted Fair Regression under Selection Bias 4

2.1. Introduction 4

2.2. Preliminaries 6

2.3. Optimization with Weights W 9

2.4. Bias Correction through Optimization over Lagrangian Dual 19

2.5. Simulation Experiments 24

2.6. Benchmark Experiments on Selection Bias Correction 27

2.7. Benchmark Experiments without Selection Bias 31

Chapter 3. Advances on Seeded Region Growing on Graph 33

3.1. Introduction 33

3.2. `SRGonG`: Seeded Region Growing on Graphs 34

3.3. Improving the Greedy Merger Step in `SRGonG` 36

3.4. Adjustment for Exposure Map 39

Chapter 4. Uncertainty Quantification of Object Boundaries Extracted from Spatial Point Pattern Image 44

4.1. Introduction 44

4.2. Fourier Descriptors 45

4.3. Statistical Inference 48
4.4. Simulation Experiments 56
4.5. Applications 61
4.6. Discussion 63

Chapter 5. Joint Mediation Effects Estimation Under High Dimensional Linear Graphical Causal Model 67
5.1. Introduction 67
5.2. Definitions and Preliminaries 69
5.3. Estimating Joint Mediation Effect 75
5.4. Simulation Experiments 80
5.5. Application 84

Chapter 6. Conclusion 91

Appendix A. Appendix to Chapter 2 92
A.1. Derivation of Balanced Weights W in Section 2.2.6 92
A.2. Proof of Proposition 5 92
A.3. Newton's Method Implementation 94
A.4. Benchmark Experiments Setup 94
A.5. Figures 97
A.6. Tables 98

Appendix B. Appendix to Chapter 4 103
B.1. Model Selection Pitfull when *Enclosure Condition* is Violated 103

Appendix C. Appendix to Chapter 5 104
C.1. Proof of Lemma 5.2.3.1 104
C.2. Proof of Theorem 5.2.4 104
C.3. Proof of Theorem 5.3.6 104
C.4. Additional Real Data Analysis Results 108

Bibliography 109

Abstract

Some Contributions to Uncertainty Quantification and Fairness in Statistical Machine Learning

Statistical machine learning methods are widely used in various fields such as astrophysics, biomedical research, genetics, education, and recruitment. While much attention has been directed toward point estimation, several critical facets, including interval estimation, hypothesis testing, and algorithmic discrimination (where models perform differently for different groups), remain relatively underexplored. This dissertation includes three projects that address complex data analysis tasks across fields of selection bias correction, spatial point segmentation, and high-dimensional mediation analysis, and contributes to the problem of algorithmic fairness and uncertainty quantification in their respective domains. Chapter 2 addresses the fairness of regression problems, aiming at mitigating the disparity between two demographic groups. Our proposed method adapts to distribution shifts induced by selection bias through a weighting technique. Chapter 3 and Chapter 4 concentrate on refining methods for segmenting irregular regions of extended astronomical emission and quantifying associated uncertainty by constructing global confidence regions of a two-dimensional boundary. Lastly, Chapter 5 studies high-dimensional mediation analysis within biomedical research, proposing new techniques to fill methodological gaps and developing inference tools to identify mediators from observational data.

Acknowledgments

First of all, I want to extend my most sincere gratitude to my PhD advisor, Thomas C. M. Lee, for his continuous help and support throughout this journey. He has been guiding me through the ups and downs of research, and encouraged me to keep pushing forward. He is always open to discussing new ideas and willing to provide as much flexibility as he can for students to explore in different directions. Besides, Thomas is a great advisor who truly cares about students' well-being and future careers.

Next, I would like to express my great appreciation to Professor Raymond K. W. Wong, Professor David van Dyk, Doctor Vinay Kashyap, Professor Andreas Zezas, and Doctor Preetam Nandy for their invaluable support and assistance throughout our collaborations. Their expertise, guidance, and willingness to share their knowledge have been instrumental in shaping the outcomes of our joint endeavors.

I am also immensely grateful to the staff of our statistics department, including Nehad Ismail, Andi Carr, Amy C. McFarl, Olga I. Rodriguez, and Sarah Driver. They are incredibly kind-hearted and warm people who have always been willing to lend a helping hand.

I want to extend my thanks to Professor Christiana Drake, Professor Shizhe Chen, Professor Jairo Fúquene-Patiño, and Professor Andreas Zezas for their assistance in serving on both my Qualification and Defense committees. Their insightful suggestions have been invaluable to my research.

I'd like to thank my friends and PhD peers. It has been great meeting Zitong Zhang, Junwen Yao, Wei Du, and many other amazing peers at UC Davis.

Last but not least, I would like to express my deepest gratitude to my family and Xinwei Wang for their endless love and support. I could never imagine accomplishing all this work without them being by my side. Their trust and encouragement have given me the confidence to face every challenge bravely.

CHAPTER 1

Introduction

In the era of big data, massive amounts of information are generated every minute from people's daily routines and scientific activities. From healthcare monitoring to social media interactions, laboratory experiments, clinical trials, and astronomical observations, valuable data for exploration abounds. However, these data often present complex patterns, requiring advancements in statistical methods to tackle specific challenges within their respective fields. While some data are high-dimensional and challenging for traditional parametric approaches, others are influenced by historical decision-making policies, leading to inherent biases in the data collection process. This dissertation focuses on two critical issues in statistical machine learning amidst the complexity of data: algorithmic fairness and uncertainty quantification (UQ).

Algorithmic fairness addresses disparities in outcomes or performance across different demographic groups within automated decision-making systems. Common approaches to this problem include restricting the disparity metrics during optimization or constructing fair data representations. With applications spanning critical areas such as criminal justice, hiring, lending, healthcare, and education, algorithmic fairness aims to mitigate the potential harms of discrimination and promote equity and justice in decision-making processes. Despite its significance, algorithmic fairness remains relatively underdeveloped in artificial intelligence, particularly within the regression models explored in this dissertation. Uncertainty quantification (UQ), on the other hand, involves quantitatively characterizing and estimating uncertainties in computational and real-world applications, crucial for assessing the likelihood of outcomes when some aspects of the system are uncertain. It typically entails providing interval estimations, such as confidence intervals or confidence regions, and constructing hypothesis tests. Despite well-established methods for point estimation, many challenging problems still lack techniques to quantify uncertainty, as seen in the boundary estimation of 2D objects or mediation effects within high-dimensional graphical models explored in this dissertation.

We present three projects that demonstrate the development of statistical methods to tackle the problems of fairness and uncertainty quantification within their specific domains.

In Chapter 2, we study the fairness in machine learning algorithms, a critical concern in sectors such as employment, education, and criminal justice. Specifically, we aim at building a fair predictive model, whose accuracy discrepancy between two demographic groups is restricted below a predefined threshold. In this chapter, we address the fair regression problem and consider the mean square error (MSE) discrepancy as our fairness metric. We, in particular, consider the scenarios where bias selection exists, causing a covariate shift between training data distribution and population data distribution and, therefore, imposing the challenge of estimating the population risk from training data. Inspired from Wong and Chan (2018), our approach begins with the proposal of a weighting method aimed at estimating the population risk through a weighted empirical risk. Subsequently, we formulate our fair regression as a constrained optimization problem, aiming to minimize the overall (weighted) MSE while constraining the (weighted) MSE discrepancy under an arbitrary threshold. We obtain the global optimal solution of this constrained optimization problem by solving its Lagrangian dual problem. Furthermore, inspired by the algorithmic structure, we introduce a bias correction mechanism to address the bias arising from overfitting to further enhance the generalization of the fairness constraint. The effectiveness of our approach is demonstrated through both numerical simulations and experiments conducted on multiple benchmark datasets.

Chapter 3 and Chapter 4 focus on the segmentation of irregular regions of extended astronomical emission and uncertainty quantification of the algorithm. In high-energy astrophysics, detecting the boundaries of extended astronomical emission is challenging due to the complex features of the object and background noise. One strategy is to segment event lists by using algorithms that aggregate photons into distinct spatial clusters, typically based on their two-dimensional sky coordinates. Fan *et al.* (2023), for example, introduced "Seeded Region Growing on a Graph" (`SRGonG`), which combines a likelihood-based representation of an inhomogeneous Poisson process with the flexible non-parametric representation of the segments to derive maximum likelihood estimates of the segments, their number, and their Poisson rates. In Chapter 3, we present several advancements to `SRGonG` aimed at improving its performance and stability, particularly under low Signal-to-Noise

(S/N) conditions and variations in exposure time across the field of interest. In Charpter 4, we develop `UQ-BoSE`, a set of algorithms and statistical methods that quantify the uncertainty associated with a segment boundary. This is accomplished via a parameterization of the boundary derived using Fourier Descriptors along with spatial bootstrap techniques, including both parametric and non-parametric methods, to effectively assess the uncertainty of the parameter estimates. We also construct a small-size global confidence region for the true object shape and validate the overall `UQ-BoSE` procedure through comparisons between the empirical coverage rate of the global confidence region with its theoretical bounds under various numerical experiment scenarios. We apply the method to datasets of NGC 2300 galaxies.

In Charpter 5, we focus on high-dimensional mediation analysis and its application in biomedical research. In a previous study by Chakrabortty *et al.* (2018), the MIDA method was introduced to estimate individual mediation effects from high-dimensional observational data. However, an unresolved question in their work is that individual mediation effects do not sum up to the total mediation effect regulated by multiple mediators due to the (unknown) interrelations among mediators. To address this gap, we first generalize the definition of mediation effect to accommodate multiple mediators, achieving an exact decomposition of the total mediation effect of each individual mediator. We then propose the joint-MIDA method to estimate the identifiable versions of the mediation effects. We prove the consistency of our method under a high-dimensional linear setting and illustrate its effectiveness through multiple simulations. Lastly, we develop a non-parametric bootstrap inference tool and demonstrate its effectiveness in epigenetics on a DNA methylation dataset, which investigates the mediation pathways from childhood trauma to cortisol stress reactivity regulated by 385,882 DNA methylation loci. We identify two new loci that possess reasonable scientific interpretations in relating the effect of childhood trauma on cortisol stress reactivity.

CHAPTER 2

Weighted Fair Regression under Selection Bias

2.1. Introduction

The use of machine learning methods has become increasingly popular in critical sectors such as employment, education, healthcare, finance, and criminal justice. However, bias and unfairness arising from prediction algorithms have led to growing concerns. Notably, accuracy disparities across different demographic groups have been observed in various scenarios Barocas and Selbst (2016). For example, Seyyed-Kalantari *et al.* (2021) showed AI systems consistently underdiagnosed historically under-served patient populations, such as female patients and black patients. In another case, analysis conducted by Julia Angwin and Kirchner (2016) highlighted that the recidivism risk scale of COMPAS software exhibits a higher false positive rate for black defenders compared to their white counterparts.

While recent efforts have been devoted to mitigating algorithmic accuracy disparity (Agarwal *et al.*, 2019; Donini *et al.*, 2018; Chi *et al.*, 2021), most of them presume the training data unbiasedly represents the population, which is a usually violated assumption in the real world. In practice, datasets may suffer from *selection bias*, originating from historically biased decision-making, thereby distorting their representativeness of the population. Machine learning algorithms trained on such datasets induce bias on population-level estimands, posing a remaining violation of fairness, sometimes referred to as *residual unfairness* (Kallus and Zhou, 2018). Many sectors grappling with AI fairness concerns have encountered such biases in their data collection procedures. For instance, data on loan default is typically collected only from applicants who were historically approved. This dataset is then used to develop loan approval policies applicable to all applicants, introducing an inherent bias (Kallus and Zhou, 2018). In this chapter, we introduce reweighting techniques to adjust the model trained on a dataset subject to selection bias, thereby extending the model accuracy parity to the population level. Our primary focus is on regression

problems with continuous outcomes (Y) and a binary protected attribute (A). Some prior works (Chi *et al.*, 2021; Zhao, 2021; Oneto *et al.*, 2019) have made strides in mitigating regressor accuracy disparity without considering selection bias. Chi *et al.* (2021) reduced the accuracy disparity of a regressor through adversarially training a fair representation that minimizes the distance between two group distributions while maximizing the representation's predictive power. Similarly, Zhao (2021) showed for Lipschitiz continuous predictors, the accuracy disparity can be reduced by training the model on a fair representation that minimizes the Wasserstein distance between two group distributions. Oneto *et al.* (2019) converted the regression problem to binary classifications by discretizing the continuous label and adopted the same technique as in Donini *et al.* (2018) that reduces the equal opportunity between groups. However, research in mitigating accuracy disparity in regression under selection bias remains scarce. The most relevant work we could find is Du *et al.* (2022), which is developed under stronger assumptions than ours.

2.1.1. Related work. Many studies addressing the impact of selection bias on algorithmic fairness focus on scenarios where the training data consists of samples from the population of interest $p_{\mathcal{T}}$, with their outcomes (Y) selectively observed (Kallus and Zhou, 2018; Du *et al.*, 2022; Coston *et al.*, 2021a; Pessach and Shmueli, 2021). Kallus and Zhou (2018) used an estimated propensity score ratio between the target and training distribution to adjust the covariates shift caused by selection. Du *et al.* (2022) considered a special sample selection process that is linear in covariates and uses the Inverse Mills Ratio derived from the Heckman Model to correct the selection bias. Coston *et al.* (2021a) used extrapolation to infer the unobserved labels and de-biased the fairness estimation over a Rashomon Set of the model. Pessach and Shmueli (2021) addressed the scenario where the labels of the underprivileged group are not available in training time and proposed bias adjustment by penalizing the correlation between protected attribute and predictions.

In practice, data subject to selection may also exhibit incompleteness in their covariates (X). We provide a recruitment example to illustrate this scenario in Example 1 below. However, as far as we know, only Zhang and Long (2021) considered this more general selection bias setting characterized by incomplete covariates (X) and/or outcome (Y) in fairness-related literature. They utilize inverse probability weights (IPW) to assess the bias in fairness metric estimands.

2.2. Preliminaries

2.2.1. Problem Statement. Let (Z, A) be a vector of random variables, where A is a protected attribute, and $Z := \{X, Y\} \in \mathcal{X} \times \mathcal{Y}$ consists of predictors X and outcome Y. For simplicity, we assume A to be binary such that $A \in \{0, 1\}$. Let $g \in \mathcal{G} : \mathcal{X} \to \mathcal{Y}$ denote the predictive function. We define $R_{\mathcal{T}}(g) := \mathbb{E}_{\mathcal{T}}\{\ell(g; Z)\}$, where ℓ is a loss function, and the expectation $\mathbb{E}_{\mathcal{T}}$ is taken over a *population distribution* $p_{\mathcal{T}}$. In this chapter, we focus on regression problem, and therefore we let ℓ to be the $\mathcal{L}_2$ loss function, i.e.,

$$\ell(g; Z) = \big(g(X) - Y\big)^2.$$

We define the *population fair regressor* with fairness constraint $\mathcal{C}$:

$$g_\star := \operatorname*{argmin}_{g \in \mathcal{G} \cap \mathcal{C}} R_{\mathcal{T}}(g).$$

2.2.2. Fairness Notion. The fairness constraint is often in the form $\Delta(g) \leq \delta$ for disparity measure Δ and some $\delta \geq 0$. In this chapter, we consider the mean squared error (MSE) disparity as our fairness metric:

$$\Delta(g) = |\mathcal{E}_1(g) - \mathcal{E}_0(g)|,$$

where

$$\mathcal{E}_a(g) = \mathbb{E}_{\mathcal{T}}\{\ell(g; Z) | A = a\} \quad \text{and} \quad a = 0, 1,$$

denote the conditional MSEs. We, therefore, aim at building a function $g_\star$ that optimizes the following objective function,

$$\min_{g \in \mathcal{G}} \Big\{ \mathbb{E}_{\mathcal{T}}\{\ell(g(X); Y)\} : \Delta(g) \leq \delta \Big\}. \tag{2.1}$$

If we are given i.i.d. training samples $\{(Z_i, A_i)\}_{i=1}^n$ from $p_{\mathcal{T}}$, $\mathcal{E}_a(g)$ can be estimated by the data strata with $A = a$.

2.2.3. Training Data Distribution under Selection Bias. We assume the training samples are subject to selection bias. We use $S_i \in \{0, 1\}$ to indicate whether the i-th datum remains fully observed after selection. When $S_i = 1$, we obtain (Z_i, A_i). Conversely, when $S_i = 0$, part

of (Z_i, A_i) is not observed. The conditional distribution of (Z, A) given $S = 1$ is denoted by $p_{\mathcal{D}}$, which is also referred to as *training distribution*.

Due to the selection process, the training distribution $p_{\mathcal{D}}$ is not necessarily the same as the population distribution $p_{\mathcal{T}}$, and the empirical group loss can be biased for estimating the conditional loss. As such, adjustment will be needed to correct the selection bias and its effect on the estimated (conditional) loss. The issue has been discussed in Zhang and Long (2021); Kallus and Zhou (2018); Du *et al.* (2022); Coston *et al.* (2021b). We present an illustrative example in Example 1.

EXAMPLE 1 (Recruitment pre-screening). *We present an example in a recruitment system, where the success of a job interview (Y_i) is predicted based on a candidate's major (X_{i1}), undergraduate GPA (X_{i2}) and an assessment score (X_{i3}), with gender as the protected attribute A_i . A pre-screening process is first applied to filter candidates based on X_{i1} and X_{i2}. If a candidate failed this pre-screening ($S_i = 0$) process, their assessment score (X_{i3}) and interview outcome (Y_i) will become unrecorded (i.e., not selected). Therefore, if the pre-screening policy favors majors predominantly chosen by one gender group ($A_i = 1$), the other group ($A_i = 0$) may be underrepresented in the training data ($p_{\mathcal{D}}$). Consequently, models trained on such data can propagate biased estimates of risk $R_{\mathcal{T}}$ and fairness metrics in the population $p_{\mathcal{T}}$.*

2.2.4. Standard Selection Mechanism. We assume the following assumption holds for the sample selection mechanism.

ASSUMPTION 1. *There exists a sub-vector U of Z, which is not subject to selection, such that,*

$$\mathbb{E}_{\mathcal{T}}\{\ell(g; Z)|A = a, U\} = \mathbb{E}_{\mathcal{D}}\{\ell(g; Z)|A = a, U\}, \quad \textit{for } a = 0, 1 \textit{ and } g \in \mathcal{G}. \tag{2.2}$$

The condition (2.2) is clearly implied by the general condition that $p_{\mathcal{T}}(Z|U, A = a) = p_{\mathcal{D}}(Z|U, A = a)$, which is commonly assumed in many related works (Coston *et al.*, 2021b; Zhang and Long, 2021; Kallus and Zhou, 2018). Example 1 satisfies Assumption 1 as the pre-screening (selection process) depends on the candidate's major (X_{i1}) and undergraduate GPA (X_{i2}), which are always observable.

2.2.5. Weights Adjustment Approach. Note that the selection mechanism in (2.2) implies that

$$\mathbb{E}_{\mathcal{T}}\{\ell(g;Z)|A=a\} = \mathbb{E}_{\mathcal{T}}\{\mathbb{E}_{\mathcal{T}}\{\ell(g;Z)|U,A=a\}|A=a\} = \mathbb{E}_{\mathcal{T}}\{\mathbb{E}_{\mathcal{D}}\{\ell(g;Z)|U,A=a\}|A=a\}.$$

The outer (conditional) expectation in the RHS is taken over U, which can be estimated by the data as (U,A) is not subject to selection and is always observed.

Let $r(g,U,a) = \mathbb{E}_{\mathcal{T}}\{\ell(g;Z)|U,A=a\} = \mathbb{E}_{\mathcal{D}}\{\ell(g;Z)|U,A=a\}$. Suppose the weights control the error

$$\left|\frac{1}{n}\sum_{i=1}^{n} S_i W_i \mathbb{I}(A_i=a) r(g,U_i,a) - \frac{1}{N_a}\sum_{j=1}^{n}\mathbb{I}(A_j=a) r(g,U_j,a)\}\right|, \tag{2.3}$$

where $N_a = \sum_{i=1}^{n}\mathbb{I}(A_i=a)$. We will be able to estimate $\mathcal{E}_a(g)$ with the following:

$$\widetilde{\mathcal{E}}_a(g,W) := \frac{1}{n}\sum_{i=1}^{n} S_i W_i \mathbb{I}(A_i=a)\ell(g;Z_i) \tag{2.4}$$

$$= \frac{1}{n}\sum_{i=1}^{n} S_i W_i \mathbb{I}(A_i=a)\left\{\ell(g;Z_i) - r(g,U_i,a)\right\} \tag{2.5}$$

$$+ \frac{1}{n}\sum_{i=1}^{n} S_i W_i \mathbb{I}(A_i=a) r(g,U_i,a) - \frac{1}{N_a}\sum_{j=1}^{n}\mathbb{I}(A_j=a) r(g,U_j,a) \tag{2.6}$$

$$+ \frac{1}{N_a}\sum_{j=1}^{n}\mathbb{I}(A_j=a) r(g,U_j,a) - \mathbb{E}_{\mathcal{T}}\{r(g,U,a)|A=a\} \tag{2.7}$$

$$+ \mathbb{E}_{\mathcal{T}}\{r(g,U,a)|A=a\}. \tag{2.8}$$

Note that $\mathbb{E}_{\mathcal{T}}\{r(g,U,a)|A=a\} = \mathbb{E}_{\mathcal{T}}\{\ell(g;Z)|A=a\} = \mathcal{E}_a(g)$.

With the same sets of weights, we can then estimate

$$\mathbb{E}_{\mathcal{T}}\{\ell(g;Z)\} = \sum_{a=0,1}\mathbb{E}_{\mathcal{T}}\{\ell(g;Z)|A=a\}P(A=a)$$

by

$$\tilde{R}_{\mathcal{T}}(g,W) := \sum_{a=0,1}\widetilde{\mathcal{E}}_a(g,W)\frac{N_a}{n} = \frac{1}{n}\sum_{i=1}^{n} S_i \sum_{a=0,1}\mathbb{I}(A_i=a)\frac{W_{a,i}N_a}{n}\ell(g;Z_i). \tag{2.9}$$

2.2.6. Balanced Weights Estimation for General r. In this section, we aim to control the error term in (2.3) and denote the weights for controlling the error term for group a as W_a.

For each $a = 0, 1$, we would like to obtain the weights W_a such that

$$\sum_{i=1}^{n} \mathbb{I}(A_i = a) \left(\frac{S_i W_{a,i}}{n} - \frac{1}{N_a} \right) r(g, U_i, a) \approx 0. \tag{2.10}$$

Denote

$$H_n(W_a, r, a) = \left\{ \sum_{i=1}^{n} \mathbb{I}(A_i = a) \left(\frac{S_i W_{a,i}}{n} - \frac{1}{N_a} \right) r(g, U_i, a) \right\}^2,$$

we therefore need to control $H_n(W_a, r, a)$ to validate the balanced weights W_a. We have

$$H_n(W_a, r, a) \leq \frac{1}{n} \sum_{i=1}^{n} \mathbb{I}(A_i = a) r^2(g, U_i, a) \frac{1}{n} \sum_{i=1}^{n} \mathbb{I}(A_i = a) \left(S_i W_{a,i} - \frac{n}{N_a} \right)^2.$$

Following Wong and Chan (2018), we consider the following constrained minimizations for $a = 0, 1$ for some reproducing-kernel Hilbert space $\mathcal{H}$ with norm $|| \cdot ||_{\mathcal{H}}$:

$$\min_{W_{a,i} \geq \frac{n}{N_a}} \left[\sup_{r \in \tilde{\mathcal{H}}_n} \{ H_n(W_a, r, a) - \lambda_1 \|r\|_{\mathcal{H}} \} + \lambda_2 V_n(W_a) \right],$$

where

$$\begin{aligned} ||r||_n^2 &= n^{-1} \sum_{i=1}^{n} \mathbb{I}(A_i = a) r^2(g, U_i, a), \\ \tilde{\mathcal{H}}_n &= \{ r \in \mathcal{H} : ||r||_n = 1 \}, \\ V_n(W_a) &= \frac{1}{n} \sum_{i=1}^{n} \mathbb{I}(A_i = a) S_i W_{a,i}^2, \end{aligned}$$

and λ_1, λ_2 are tuning parameters. Since the counter part of $\frac{N_a}{n} W_{a,i}$ is $P(S_i = 1 | U_i, a)$, which is at least 1, we restrict $W_{a,i} \geq \frac{n}{N_a}$. The details of computing W_a can be found in Appendix A.1.

2.3. Optimization with Weights W

Under Assumption 2.2, we use weights estimated in Section 2.2.6 to adjust for the estimation bias caused by the shift between $p_{\mathcal{T}}$ and $p_{\mathcal{D}}$. In practice, some regularization may be needed to control the flexibility of $g(\cdot)$. We use the term $\boldsymbol{J}(g)$ to denote the roughness of $g(\cdot)$. We then

optimize the following empirical version of the problem in (2.1):

$$\min_{g\in\mathcal{G}} \left\{ \tilde{R}_{\mathcal{T}}(g,W) + \eta \boldsymbol{J}(g) : |\widetilde{\Delta}(g,W)| \leq \delta \right\}, \tag{2.11}$$

where $\widetilde{\Delta}(g,W)$ and $\tilde{R}_{\mathcal{T}}(g,W)$ are as defined in (2.4) and (2.9) respectively. Since the estimation and evaluation of g only depends on fully observed data, i.e., $\{X_i, Y_i | S_i = 1\}_{i=1}^n$, we reformulate the optimization with this set of samples. Without loss of generality, we assume that $S_i = 1$ for $i \leq m$ and $S_i = 0$ for $i > m$. Let $\boldsymbol{X} = (X_1, \ldots, X_m)^T$ denote the $m \times p$ design matrix, and $\boldsymbol{Y} = (Y_1, \ldots, Y_m)^T$ denote the vector of outcome. Let $\boldsymbol{W}$ denote a $m \times m$ diagonal matrix where $\boldsymbol{W}_{ii} = \sum_{a=0,1} \mathbb{I}(A_i = a) \frac{N_a W_{a,i}}{n^2}$. Let $\boldsymbol{D}$ denote a $m \times m$ diagonal matrix with $\boldsymbol{D}_{ii} = \mathbb{I}(A_i = 1)\frac{W_{1,i}}{n} - \mathbb{I}(A_i = 0)\frac{W_{0,i}}{n}$.

2.3.1. Parametric Estimation of g. We first consider a parametric form for $g(X) = X^T\boldsymbol{\beta}$, and set $\boldsymbol{J}(g) = ||g||_2^2$. Then, the optimization can be written as

$$\min_{\boldsymbol{\beta}} \left\{ (\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{\beta})^T \boldsymbol{W} (\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{\beta}) + \eta \boldsymbol{\beta}^T \boldsymbol{\beta} : \left| (\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{\beta})^T \boldsymbol{D} (\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{\beta}) \right| \leq \delta. \right\} \tag{2.12}$$

2.3.2. Nonparametric Estimation of g. This section introduces the nonparametric estimation of $g(\cdot)$. Consider $\boldsymbol{J}(g) = ||g||_{\mathcal{H}}^2$ for some reproducing kernel Hilbert space (RKHS) $\mathcal{H}$ with a reproducing kernel function $\boldsymbol{R}(\cdot,\cdot)$. It is proved by the representation theorem that the optimal solution of (2.11) lies in the space spanned by $\{\boldsymbol{R}(X_1,\cdot), \ldots, \boldsymbol{R}(X_m,\cdot)\}$. Therefore, let $\mathbf{K}$ denote the Gram matrix that $\mathbf{K}_{i,j} = \boldsymbol{R}(X_i, X_j)$ for $1 \leq i, j \leq m$. By plugging the form of

$$g(\cdot) = \sum_{i=1}^{m} \alpha_i \boldsymbol{R}(X_i, \cdot)$$

back into (2.11), we can reformulate the optimization problem in terms of $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_m)^T$,

$$\min_{\boldsymbol{\alpha}} \left\{ (\boldsymbol{Y} - \mathbf{K}\boldsymbol{\alpha})^T \boldsymbol{W} (\boldsymbol{Y} - \mathbf{K}\boldsymbol{\alpha}) + \eta \boldsymbol{\alpha}^T \mathbf{K} \boldsymbol{\alpha} \quad s.t. \quad |(\boldsymbol{Y} - \mathbf{K}\boldsymbol{\alpha})^T \boldsymbol{D} (\boldsymbol{Y} - \mathbf{K}\boldsymbol{\alpha})| < \delta \right\}. \tag{2.13}$$

The matrix $\mathbf{K}$ is nearly singular, which makes the following up Lagrangian optimization difficult. We will use a low-rank approximation of $\mathbf{K}$ to overcome this issue, i.e, we set $\mathbf{K} = \boldsymbol{L}\boldsymbol{L}^T$,

where $\boldsymbol{L}$ is a $m \times p$ matrix for some $p \ll m$. The above optimization becomes

$$\min_{\boldsymbol{\beta}} \Big\{ (\boldsymbol{Y} - \boldsymbol{L}\boldsymbol{\beta})^T \boldsymbol{W} (\boldsymbol{Y} - \boldsymbol{L}\boldsymbol{\beta}) + \eta \boldsymbol{\beta}^T \boldsymbol{\beta} \quad s.t. \quad |(\boldsymbol{Y} - \boldsymbol{L}\boldsymbol{\beta})^T \boldsymbol{D} (\boldsymbol{Y} - \boldsymbol{L}\boldsymbol{\beta})| < \delta \Big\}, \tag{2.14}$$

where $\boldsymbol{\beta}$ is equaivalent to $\boldsymbol{L}^T\boldsymbol{\alpha}$. The prediction of training samples are $\mathbf{K}\boldsymbol{\alpha} = \boldsymbol{L}\boldsymbol{L}^T\boldsymbol{\alpha} = \boldsymbol{L}\boldsymbol{\beta}$.

Then, any function $g(\cdot)$ in the span of $\{\boldsymbol{R}(X_1, \cdot), \ldots, \boldsymbol{R}(X_m, \cdot)\}$ can be written as

$$g(\cdot) = \sum_{i=1}^{m} \alpha_i \boldsymbol{R}(X_i, \cdot) = \boldsymbol{\alpha}^T \mathbf{K}\mathbf{K}^+ z(\cdot) = \boldsymbol{\beta}^T \mathbf{K}^+ z(\cdot), \tag{2.15}$$

where $z(\cdot) = (\{\boldsymbol{R}(X_1, \cdot), \ldots, \{\boldsymbol{R}(X_m, \cdot))^T$.

In the following sections, we focus on the optimization form in (2.14). Note that each of the two fairness constraints $(\boldsymbol{Y} - \boldsymbol{L}\boldsymbol{\beta})^T \boldsymbol{D} (\boldsymbol{Y} - \boldsymbol{L}\boldsymbol{\beta})$ and $-(\boldsymbol{Y} - \boldsymbol{L}\boldsymbol{\beta})^T \boldsymbol{D} (\boldsymbol{Y} - \boldsymbol{L}\boldsymbol{\beta})$ is a difference between two convex functions; therefore, we could use the disciplined convex-concave program (DCCP) to solve the optimization (Shen *et al.*, 2016). However, there is no guarantee that DCCP will have the global optimal solution. Therefore, we next introduce the Lagrangian duality and sufficient conditions to find the global optimal solution by solving the dual problem.

2.3.3. Lagrangian Duality and Global Optimal Solution. We formalize the constrained optimization problem as a convex problem through the Lagrangian duality. Specificaly, we want to solve the following problem

$$p^* := \inf_{\boldsymbol{\beta}} f_0(\boldsymbol{\beta}) : f_i(\boldsymbol{\beta}) \leq 0, i = 1, 2, \tag{2.16}$$

where

$$f_0(\boldsymbol{\beta}) = \boldsymbol{\beta}^T (\boldsymbol{L}^T \boldsymbol{W} \boldsymbol{L} + \eta \mathbf{I}) \boldsymbol{\beta} - 2\boldsymbol{Y}^T \boldsymbol{W} \boldsymbol{L} \boldsymbol{\beta}, \tag{2.17}$$

$$f_1(\boldsymbol{\beta}) = \boldsymbol{\beta}^T \boldsymbol{L}^T \boldsymbol{D} \boldsymbol{L} \boldsymbol{\beta} - 2\boldsymbol{Y}^T \boldsymbol{D} \boldsymbol{L} \boldsymbol{\beta} + \boldsymbol{Y}^T \boldsymbol{D} \boldsymbol{Y} - \delta, \tag{2.18}$$

$$f_2(\boldsymbol{\beta}) = -\boldsymbol{\beta}^T \boldsymbol{L}^T \boldsymbol{D} \boldsymbol{L} \boldsymbol{\beta} + 2\boldsymbol{Y}^T \boldsymbol{D} \boldsymbol{L} \boldsymbol{\beta} - \boldsymbol{Y}^T \boldsymbol{D} \boldsymbol{Y} - \delta. \tag{2.19}$$

Note that for simplicity, we dropped the constant term $\boldsymbol{Y}^T \boldsymbol{W} \boldsymbol{Y}$ from the original objective function to obtain (2.17).

2.3.3.1. *The Lagrangian Dual Form.* The Lagrangian form can be expressed with one free Lagrange multipler λ, or equivalently two nonnegative Lagrange multiplers

$$\lambda_+ := \max\{\lambda, 0\}, \lambda_- := -\min\{\lambda, 0\}, \lambda = (\lambda_+ - \lambda_-), \tag{2.20}$$

as follows:

$$\begin{aligned} L(\boldsymbol{\beta}, \lambda_+, \lambda_-) =& \boldsymbol{\beta}^T(\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{L} + \eta\mathbf{I})\boldsymbol{\beta} - 2\boldsymbol{Y}^T\boldsymbol{W}\boldsymbol{L}\boldsymbol{\beta} + \lambda_+(-\delta - \boldsymbol{\beta}^T\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}\boldsymbol{\beta} + 2\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{L}\boldsymbol{\beta} - \boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y}) \\ &+ \lambda_-(\boldsymbol{\beta}^T\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}\boldsymbol{\beta} - 2\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{L}\boldsymbol{\beta} + \boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta) \\ =& \boldsymbol{\beta}^T(\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I})\boldsymbol{\beta} - 2\boldsymbol{Y}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{L}\boldsymbol{\beta} \\ &+ \lambda_+(-\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta) - \lambda_-(-\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} + \delta) \end{aligned} \tag{2.21}$$

Then the dual of the constraint optimization in (2.16) becomes

$$\begin{aligned} d^* :=& \sup_{\lambda} -(\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{Y} + \eta\mathbf{I})^T(\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{L})^-(\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{Y}) \\ &+ \lambda_+(-\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta) - \lambda_-(-\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} + \delta) \\ &\text{subject to } \boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I} \succeq 0, \quad \boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{Y} \in \text{Range}(\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I}). \end{aligned} \tag{2.22}$$

2.3.3.2. *The Strong Duality.* We present the necessary and sufficient conditions for strong duality to hold between the primal form (2.16) and dual form (2.22) in the following Lemma.

ASSUMPTION 2.

(*i*) $\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L} \neq 0$.

(*ii*) $\exists\overline{\boldsymbol{\beta}} \in \mathbb{R}^d$, *such that* $-\delta < (\boldsymbol{Y} - \boldsymbol{L}\overline{\boldsymbol{\beta}})^T\boldsymbol{D}(\boldsymbol{Y} - \boldsymbol{L}\overline{\boldsymbol{\beta}}) < \delta$.

(*iii*) *The primal problem is bounded below.*

LEMMA 1 (Theorem 7 of Wang and Xia (2014)). *Under Assumption 2* (*i*) *and* (*ii*),

$$d^* = p^*. \tag{2.23}$$

Moreover, if Assumption 2 (iii) holds in addition to (i) and (ii), d^* is attained.

2.3.3.3. *The Conditions to Characterize a Global Optimal solution.* With Assumption 2, we have the following theorem to characterize the necessary and sufficient conditions for $\boldsymbol{\beta}^*$ being an optimal solution for problem (2.16).

THEOREM 1 (Theorem 2.3 of Pong and Wolkowicz (2014)). *If Assumption 2 holds, $\boldsymbol{\beta}^*$ is a solution to the primal problem(2.16) if and only if for some Lagrangian multipler $\lambda^* \in \mathbb{R}$, we have*

$$\left.\begin{aligned}(\boldsymbol{L}^T(\boldsymbol{W}-\lambda^*\boldsymbol{D})\boldsymbol{L}+\eta\mathbf{I})\boldsymbol{\beta}^* &= \boldsymbol{L}^T(\boldsymbol{W}-\lambda^*\boldsymbol{D})\boldsymbol{Y},\\ \boldsymbol{L}^T(\boldsymbol{W}-\lambda^*\boldsymbol{D})\boldsymbol{L}+\eta\mathbf{I} &\succeq 0,\end{aligned}\right\} \textit{dual fesibility}$$

$$-\delta \le (\boldsymbol{Y}-\boldsymbol{L}\boldsymbol{\beta}^*)^T\boldsymbol{D}(\boldsymbol{Y}-\boldsymbol{L}\boldsymbol{\beta}^*) \le \delta, \textit{ primal feasibility} \tag{2.24}$$

$$\left.\begin{aligned}\lambda_+^*(-\delta-(\boldsymbol{Y}-\boldsymbol{L}\boldsymbol{\beta}^*)^T\boldsymbol{D}(\boldsymbol{Y}-\boldsymbol{L}\boldsymbol{\beta}^*)) &= 0,\\ \lambda_-^*((\boldsymbol{Y}-\boldsymbol{L}\boldsymbol{\beta}^*)^T\boldsymbol{D}(\boldsymbol{Y}-\boldsymbol{L}\boldsymbol{\beta}^*)-\delta) &= 0,\end{aligned}\right\} \textit{complementary slackness}$$

Note that from the dual feasibility condition in (2.24), we have to ensure the semidefinite positivity of the matrix pencil $\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{L}+\eta\mathbf{I}-\lambda\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}$ when opitmizing for the Lagrangian multiplier λ. Without loss of geneality, we assume $\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{L}+\eta\mathbf{I} \succ 0$, then the matrix pencil $\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{L}+\eta\mathbf{I}-\lambda\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}$ is positively regular, i.e, there exist some $\hat{\lambda} \in \mathbb{R}$ such that

$$\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{L}+\eta\mathbf{I}-\hat{\lambda}\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L} \succ 0,$$

which guarantees the existence of an optimal solution for the primal problem (2.16) according to (Pong and Wolkowicz, 2014, Proposition 2.3).

Moreover, by (Pong and Wolkowicz, 2014, Theorem 2.4), there exists a nonempty open interval $(\underline{\lambda}, \overline{\lambda})$ such that $\boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{L}+\eta\mathbf{I} \succeq 0$ if and only if $\underline{\lambda} \le \lambda \le \overline{\lambda}$, and $\boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{L}+\eta\mathbf{I} \succ 0$ if and only if $\underline{\lambda} < \lambda < \overline{\lambda}$.

2.3.3.4. *Characterization of $\underline{\lambda}$ and $\overline{\lambda}$.* We denote $\lambda_{\max}(A,B)$ and $\lambda_{\min}(A,B)$ to be the maximum and minimum generalized eigenvalue of matrix pair (A,B). Then from Pong and Wolkowicz (2014) and (Nguyen and Ngan Nguyen, 2024, Theorem 1) we have the following proposition that can calculate $\underline{\lambda}$ and $\overline{\lambda}$ efficiently:

PROPOSITION 2. *1. If* $\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L} \succ 0$, $\underline{\lambda} = \lambda_{\min}(\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{L} + \eta\mathbf{I}, \boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}), \overline{\lambda} = \infty$.

2. If $\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L} \prec 0$, $\underline{\lambda} = -\infty, \overline{\lambda} = \lambda_{\max}(\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{L} + \eta\mathbf{I}, \boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L})$.

3. If $\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}$ *is indefinite, given* $\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{L}+\eta\mathbf{I} \succ 0$, $\underline{\lambda} = \frac{1}{\lambda_{\min}(\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}, \boldsymbol{L}^T\boldsymbol{W}\boldsymbol{L}+\eta\mathbf{I})}, \overline{\lambda} = -\frac{1}{\lambda_{\min}(-\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}, \boldsymbol{L}^T\boldsymbol{W}\boldsymbol{L}+\eta\mathbf{I})}$.

For each $\lambda \in$ *closure*$(\underline{\lambda}, \overline{\lambda})$, we define the first order stationary point,

$$\boldsymbol{\beta}(\lambda) := (\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I})^{-}\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{Y} \tag{2.25}$$

and the constraint evaluated at $\boldsymbol{\beta}(\lambda)$ as,

(2.26)

$$\begin{aligned}\psi(\lambda) :=& \boldsymbol{\beta}(\lambda)^T\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}\boldsymbol{\beta}(\lambda) - 2\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{L}\boldsymbol{\beta}(\lambda)\\ =&(\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{Y})^T(\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I})^{-}\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}^T(\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I})^{-}\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{Y}\\ &- 2\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{L}(\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I})^{-}\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{Y}\end{aligned}$$

Lemma 2.3, 2.4 and 2.5 in Pong and Wolkowicz (2014) guarantee that the following two propositions hold:

PROPOSITION 3. *Suppose that the Assumption 2 hold. Moreover, let* λ^* *denote the optimal solution of the dual problem (2.22), the following holds for primal problem (2.16):*

1. If $\lambda^* \in (\underline{\lambda}, \overline{\lambda})$, *then* $\boldsymbol{\beta}(\lambda^*)$ *is the unique global optimal solution of for (2.16).*

2. If $\lambda^* = \underline{\lambda}$, *then for any normalized* $\nu \in \boldsymbol{Null}(\boldsymbol{L}^T(\boldsymbol{W} - \underline{\lambda}\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I})$, *there exists an* $\alpha \in \mathbb{R}$ *such that* $\boldsymbol{\beta}(\lambda^*) + \alpha\nu$ *is the global optimal solution for (2.16).*

3. If $\lambda^* = \overline{\lambda}$, *then for any normalized* $\nu \in \boldsymbol{Null}(\boldsymbol{L}^T(\boldsymbol{W} - \overline{\lambda}\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I})$, *there exists an* $\alpha \in \mathbb{R}$ *such that* $\boldsymbol{\beta}(\lambda^*) + \alpha\nu$ *is the global optimal solution for (2.16).*

Also, since $0 \in (\underline{\lambda}, \overline{\lambda})$ and $\psi(\lambda)$ is strictly increasing on $(\underline{\lambda}, \overline{\lambda})$, we can reduce the two inequality constraints in the primal problem (2.16) to one equality constraint following the analysis in Section 2.2.2 Pong and Wolkowicz (2014):

PROPOSITION 4 (Case Check). *1. If* $-\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta \le \psi(0) \le -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} + \delta$, $\lambda^* = 0$.

2. If $\psi(0) < -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta$, *then* $\lambda^* > 0$. *By complementary slackness, if the equation* $\psi(\lambda) = -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta$ *is solvable within the* $(\underline{\lambda}, \overline{\lambda})$, λ^* *is the unique solution; otherwise,* $\lambda^* = \overline{\lambda}$ *and the*

α in Proposition 3 is chosen such that $(\boldsymbol{\beta}(\overline{\lambda})+\alpha\nu)^T\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}(\boldsymbol{\beta}(\overline{\lambda})+\alpha\nu) - 2\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{L}(\boldsymbol{\beta}(\overline{\lambda})+\alpha\nu) = -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y}-\delta$.

3. If $\psi(0) > -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y}+\delta$, *then* $\lambda^* < 0$. *By complementary slackness, if the equation* $\psi(\lambda) = -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y}+\delta$ *is solvable within the* $(\underline{\lambda},\overline{\lambda})$, λ^* *is the unique solution; otherwise,* $\lambda^* = \underline{\lambda}$ *and the* α *in Proposition 3 is chosen such that* $(\boldsymbol{\beta}(\underline{\lambda})+\alpha\nu)^T\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}(\boldsymbol{\beta}(\underline{\lambda})+\alpha\nu) - 2\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{L}(\boldsymbol{\beta}(\underline{\lambda})+\alpha\nu) = -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y}+\delta$.

2.3.3.5. *Equality constraint.* Unless we have an interior solution, Proposition 4 allows us to reduce the two inequality constraints to a single equality constraint, i.e, we can instead optimize the following problem,

$$\begin{aligned}&\min_{\boldsymbol{\beta}} \boldsymbol{\beta}^T(\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{L}+\eta\mathbf{I})\boldsymbol{\beta} - 2\boldsymbol{Y}^T\boldsymbol{W}\boldsymbol{L}\boldsymbol{\beta}:\\ &\boldsymbol{\beta}^T\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}\boldsymbol{\beta} - 2\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{L}\boldsymbol{\beta} = C,\end{aligned} \tag{2.27}$$

for some suitable C chosen based on the value of $\psi(0)$. We can derive the optimal solution for reduced primal problem (2.27) through solving a λ^* that maximizes the following dual form of 2.27,

(2.28)

$$\begin{aligned}h_C(\lambda) =& -(\boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{Y})^T(\boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{L}+\eta\mathbf{I})^-(\boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{Y}) + \lambda C\\ &\text{subject to}\quad \boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{L}+\eta\mathbf{I}\succeq 0, \boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{Y}\in \text{Range}(\boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{L}+\eta\mathbf{I}).\end{aligned}$$

2.3.3.6. *Solution of α in Proposition 4.* Based on Lemma 2.5 in Pong and Wolkowicz (2014), when $\lambda^* = \underline{\lambda}$ or $\lambda^* = \overline{\lambda}$, for any normalized $v \in \mathbf{Null}(\boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{L})$, there exist a suitable α, such that

$$\boldsymbol{\beta}^* = (\boldsymbol{L}^T(\boldsymbol{W}-\lambda^*\boldsymbol{D})\boldsymbol{L}+\eta\mathbf{I})^-\boldsymbol{L}^T(\boldsymbol{W}-\lambda^*\boldsymbol{D})\boldsymbol{Y}+\alpha v$$

is primal feasible, i.e,

$$\boldsymbol{\beta}^{*T}\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}^T\boldsymbol{\beta}^* - 2\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{L}\boldsymbol{\beta}^* = C.$$

Algorithm 1 Optimization procedure to solve for problem 2.16

procedure MAIN ALGORITHM
  Calculate initial interval for $\lambda^*, (\underline{\lambda}, \overline{\lambda})$ as in Proposition 2.
  if $-\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta \leq \psi(0) \leq -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} + \delta$ ▷ Interior solution **then**
    $\lambda^* = 0, \boldsymbol{\beta}^* = \boldsymbol{\beta}(\lambda^*)$
  else if $\psi(0) < -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta$ ▷ Optimal solution reaches lower constraint **then**
    if $\overline{\lambda} = \infty$ or $\psi(\overline{\lambda}) < -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta$ **then**
      Use Newton's method 6 to solve $\lambda^* = \operatorname{argmax} h_C(\lambda)$ for $C = -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta$.
      $\boldsymbol{\beta}^* = \boldsymbol{\beta}(\lambda^*)$.
    else
      $\lambda^* = \overline{\lambda}$.
      Solve for suitable α as in Proposition 4; $\boldsymbol{\beta}^* = \boldsymbol{\beta}(\lambda^*) + \alpha\nu$.
  else ▷ Optimal solution reaches upper constraint
    if $\underline{\lambda} = -\infty$ or $\psi(\underline{\lambda}) > -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} + \delta$ **then**
      Use Newton's method 6 to solve $\lambda^* = \operatorname{argmax} h_C(\lambda)$ for $C = -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} + \delta$.
      $\boldsymbol{\beta}^* = \boldsymbol{\beta}(\lambda^*)$.
    else
      $\lambda^* = \underline{\lambda}$.
      Solve for suitable α as in Proposition 4; $\boldsymbol{\beta}^* = \boldsymbol{\beta}(\lambda^*) + \alpha\nu$.
  Return $(\lambda^*, \boldsymbol{\beta}^*)$

Indeed there are two such α exist, and here we choose the one with smaller magnititude:

(2.29)

$$a = \nu^T(\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L})\nu,$$

$$b = \nu^T(\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}(\boldsymbol{L}^T(\boldsymbol{W} - \lambda^*\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I})^-\boldsymbol{L}^T(\boldsymbol{W} - \lambda^*\boldsymbol{D})\boldsymbol{Y} - \boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{L}),$$

$$c = \boldsymbol{Y}^T(\boldsymbol{W} - \lambda^*\boldsymbol{D})\boldsymbol{L})(\boldsymbol{L}^T(\boldsymbol{W} - \lambda^*\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I})^-(\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L})(\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I})^-\boldsymbol{L}^T(\boldsymbol{W} - \lambda^*\boldsymbol{D})\boldsymbol{Y}$$

$$- 2\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{L}(\boldsymbol{L}^T(\boldsymbol{W} - \lambda^*\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I})^-\boldsymbol{L}^T(\boldsymbol{W} - \lambda^*\boldsymbol{D})\boldsymbol{Y} - S,$$

$$\alpha = \frac{c}{-b - \operatorname{sign}(b)\sqrt{b^2 - ac}}$$

2.3.3.7. *Algorithm.* We propose Algorithm 1 to solve the primal problem (2.16) by optimizing the Lagrangian multipler through Newton's method. Note that our method is different from the extend Rendl-Wolkowicz (ERW) algorithm proposed in Pong and Wolkowicz (2014), which solves the same primal problem as a parametrized eigenvalue problem. We, instead, focus on solving the dual problem (2.22) directly. In the experiments comparing ERW and Newton's method Pong and Wolkowicz (2014), Newton's method was shown to converge in a comparable number of iterations to ERW when $\lambda^* \in (\underline{\lambda}, \overline{\lambda})$. Following the case check in Proposition 4, our method provides the

closed-form global optimal solution when $\lambda^* = \underline{\lambda}$ or $\overline{\lambda}$. Moreover, by solving the dual problem, the optimal coefficients, $\boldsymbol{\beta}(\lambda)$, described in (2.25), will imposes a special structure related to $\boldsymbol{Y}$, which is not accessible with the ERW solution. That special structure enables us to perform model selection efficiently, as discussed in the following section.

2.3.4. Parameter tuning. In this section we use $\hat{g}_\eta$ to denote the predictive function whose complexity depends on η. A small value of η leads to a flexible model, and thereby small $\widetilde{\mathcal{E}}_0(\hat{g}_\eta, W)$ and $\widetilde{\mathcal{E}}_1(\hat{g}_\eta, W)$, as well as a small $\widetilde{\Delta}(\hat{g}_\eta, W)$. However, these models may not generalize well to the out of distribution sample as they overfit.

To mitigate the effect of overftting, a η has to be carefully selected to control the flexibilty of the fitted model. Cross-validation is commonly used strategy for model selection. We can split the datasets into k folds, training the models on $k-1$ folds and evaluate its performance in the kth fold. However, this approach requires reestimating the weights for each splitted datasets, imposing extra computation burden. In the following subsections, we introduce an approximate LOOCV to estimate the MSE on out of distribution samples.

2.3.4.1. *Approximate Group LOOCV.* One common way of estimating the generalization error $\mathbb{E}_{\mathcal{T}}\{(Y - \hat{g}_\eta(X))^2\}$ is through leave-one-out-cross-validation (LOOCV) Stone (1974). We would like to estimate the condition $\mathbb{E}_{\mathcal{T}}\{(Y - \hat{g}_\eta(X))^2 | A = a\}$ for each group a via a group LOOCV,

$$\widehat{\mathcal{E}}_a(\hat{g}_\eta, W) = \frac{1}{n}\sum_{i=1}^{n} S_i \mathbb{I}(A_i = a) W_{a,i}(Y_i - \hat{Y}_{i(-i)})^2, \tag{2.30}$$

where $\hat{Y}_{i(-i)}$ denote the estimator of $\mathbb{E}[Y_i]$ excluding the i-th observation.

Let $\boldsymbol{L}_{(-i)}$ and $\boldsymbol{Y}_{(-i)}$ denote the covariates matrix and outcome vector with i-th observation $(L_i, \boldsymbol{Y}_i)$ left out. Let $\boldsymbol{D}_{(-i)}$ and $\boldsymbol{W}_{(-i)}$ denote the corresponding $m-1 \times m-1$ weight matrices. Without loss of generality, we consider the following optimziations,

$$\hat{\boldsymbol{\beta}} = \underset{\boldsymbol{\beta}}{\operatorname{argmin}} (\boldsymbol{L}\boldsymbol{\beta} - \boldsymbol{Y})^T \boldsymbol{W}(\boldsymbol{L}\boldsymbol{\beta} - \boldsymbol{Y}) + \eta \boldsymbol{\beta}^T \boldsymbol{\beta} : (\boldsymbol{L}\boldsymbol{\beta} - \boldsymbol{Y})^T \boldsymbol{D}(\boldsymbol{L}\boldsymbol{\beta} - \boldsymbol{Y}) = -\delta, \tag{2.31}$$

(2.32)

$$\hat{\boldsymbol{\beta}}_{(-i)} = \underset{\boldsymbol{\beta}}{\operatorname{argmin}} (\boldsymbol{L}_{(-i)}\boldsymbol{\beta} - \boldsymbol{Y}_{(-i)})^T \boldsymbol{W}_{(-i)}(\boldsymbol{L}_{(-i)}\boldsymbol{\beta} - \boldsymbol{Y}_{(-i)}) + \eta \boldsymbol{\beta}^T \boldsymbol{\beta} : (\boldsymbol{L}_{(-i)}\boldsymbol{\beta} - \boldsymbol{Y}_{(-i)})^T \boldsymbol{D}_{(-i)}(\boldsymbol{L}_{(-i)}\boldsymbol{\beta} - \boldsymbol{Y}_{(-i)}) = -\delta.$$

PROPOSITION 5. *When the optimal solutions of the above constrained optimizations take the following forms,*

$$\hat{\boldsymbol{\beta}} = \boldsymbol{\beta}(\lambda) = (\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I})^{-1}\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{Y}, \tag{2.33}$$

for some $\lambda = \lambda^* \in (\underline{\lambda}, \overline{\lambda})$ *and*

$$\hat{\boldsymbol{\beta}}_{(-i)} = \boldsymbol{\beta}_{(-i)}(\lambda) = (\boldsymbol{L}_{(-i)}^T(\boldsymbol{W}_{(-i)} - \lambda\boldsymbol{D}_{(-i)})\boldsymbol{L}_{(-i)} + \eta\mathbf{I})^{-1}\boldsymbol{L}_{(-i)}^T(\boldsymbol{W}_{(-i)} - \lambda\boldsymbol{D}_{(-i)})\boldsymbol{Y}_{(-i)},$$

for some $\lambda = \lambda_{(-i)} \in (\underline{\lambda}, \overline{\lambda})$. *The following will hold,*

$$\hat{\boldsymbol{Y}}_{(i)} = \mathbf{H}(\lambda_{(-i)})\boldsymbol{Y}_{(i)}, \tag{2.34}$$

where

$$\mathbf{H}(\lambda) = \boldsymbol{L}(\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I})^{-1}\boldsymbol{L}^T(\boldsymbol{W} - \boldsymbol{L}\boldsymbol{D}), \tag{2.35}$$

and $\boldsymbol{Y}_{(i)} = (Y_1, \ldots, Y_{i-1}, \hat{Y}_{i(-i)}, Y_{i+1}, \ldots, Y_m)^T$, *with* $\hat{Y}_{i(-i)} = L_i^T\hat{\boldsymbol{\beta}}_{(-i)}$.

The proof of above proposition is included in Appendix A.2. We then use λ^* to approximate each $\lambda_{(-i)}$ and denote $\mathbf{H} = \mathbf{H}(\lambda^*)$. By Proposition 2.34 and the identity $\hat{\boldsymbol{Y}} = \mathbf{H}\boldsymbol{Y}$, we can approximate

$$Y_i - \hat{Y}_{i(-i)}$$

with

$$(Y_i - \hat{Y}_i)/(1 - \mathbf{H}_{i,i}).$$

Therefore, given a fixed η, the approximate group LOOCV can be simplified to the following:

$$\widehat{\mathcal{E}}_a(\hat{g}_\eta, W) \approx \frac{1}{n}\sum_{i=1}^{n} S_i\mathbb{I}(A_i = a)W_{a,i}\left(\frac{Y_i - \hat{Y}_i}{1 - \mathbf{H}_{i,i}}\right)^2.$$

When η becomes unexpectedly large, the assumption of Proposition 5 may not be satisfied, and we might encounter the edge case. In this situation, the approximate LOOCV may not serve as an accurate estimator of the out-of-sample MSE, but it still provides a manageable approach.

It's important to note that in both simulations and benchmark experiments, we haven't observed this case for a wide range of ηs.

2.3.4.2. *Tuning procedure.* For each fixed η, we estimate the out of sample unfairness via

$$\widehat{\Delta}(\hat{g}_\eta, W) = \left|\widehat{\mathcal{E}}_1(\hat{g}_\eta, W) - \widehat{\mathcal{E}}_0(\hat{g}_\eta, W)\right|,$$

and the MSE via

$$\widehat{R}_{\mathcal{T}}(\hat{g}_\eta, W) = \frac{N_1}{n}\widehat{\mathcal{E}}_1(\hat{g}_\eta, W) + \frac{N_0}{n}\widehat{\mathcal{E}}_0(\hat{g}_\eta, W).$$

For a given δ, we select the η that corresponds smallest $\widehat{R}_{\mathcal{T}}(\hat{g}_\eta)$ among all ηs such $\widehat{\Delta}(\hat{g}_\eta) \leq k\delta$, where k is a hyperparameter. If there is no η satisfies $\widehat{\Delta}(\hat{g}_\eta) \leq k\delta$, we select η coresponding to the smallest $\widehat{R}_{\mathcal{T}}(\hat{g}_\eta, W)$. We set $k = 1.2$ for all our experiments.

2.4. Bias Correction through Optimization over Lagrangian Dual

We use $\tilde{g}(\cdot)$ to denote an arbitary predictive model trained via constrained optimization (2.11). Although the difference between the empirical MSE and true risk of $\tilde{g}(\cdot)$, denoted by,

$$\text{bias}_a = \mathbb{E}_{\mathcal{T}}\{(Y - \tilde{g}(X))^2 | A = a\} - \widetilde{\mathcal{E}}_a(\tilde{g}, W), \tag{2.36}$$

will be small with certain regularizations on $\tilde{g}(\cdot)$, e.g., a suitable penalized parameter η, it still cannot be fully eliminated in general. As a result, even though the disparity of the empirical risk is constrained, i.e.,

$$\tilde{\Delta}(\tilde{g}, W) = |\widetilde{\mathcal{E}}_1(\tilde{g}, W) - \widetilde{\mathcal{E}}_0(\tilde{g}, W)| \leq \delta,$$

there is no gurantee that such empirical fairness can be generalized to the population so that the following can hold,

$$\Delta(\tilde{g}) = |\widetilde{\mathcal{E}}_1(\tilde{g}, W) + \text{bias}_1 - \widetilde{\mathcal{E}}_1(\tilde{g}, W) - \text{bias}_0| \leq \delta.$$

It important to note that the source of such bias is not from the selection bias, but is rather the consequence, or "optimism", of estimating the population risk by the empirical risk Efron (2004). Suppose ideally we have an estimator of such bias, $\widehat{\text{bias}}_a$, then to correct for such bias, we would like to optimize the objective function (2.11) subject to

$$|\widetilde{\mathcal{E}}_1(g, W) + \widehat{\text{bias}}_1 - \widetilde{\mathcal{E}}_0(g, W) - \widehat{\text{bias}}_0| \leq \delta.$$

However, note that $\widehat{\text{bias}}_a$ is an estimation of the bias of MSE of $\tilde{g}$, which is not easily accessible as the estimated bias depends on the estimator $\tilde{g}$ but at the same time is involved in the definition of the estimator $\tilde{g}$.

The importance of generalization fairness has been acknowledged in previous researches (e.g. Cotter *et al.*, 2019b), yet relatively few studies have directly addressed it. Indeed, dealing with the risk estimation bias during training stage without splitting the data is rather challenging because the bias of $\tilde{g}$ could not easily accessible. Woodworth *et al.* (2017) proposed a post-hoc correction approach to improve genreralization performance. Cotter *et al.* (2019a) framed the fairness problem as a constrained optimization while optimizing the model and enforcing the constraints on two independent datasets to estimate bias_a in training stage.

In the following subsections, we propose a bias correction procedure by modifying the constrained optimization (2.13) and its dual problem. Our method grounds on the structure of the first order stationary point solution of $\boldsymbol{\beta}(\lambda)$, and use the same LOOCV as in Section 2.3.4.1 to estimate the bias at each iteration of the Newton's method.

For simplicity, we use $\widehat{\text{bias}}_a(\tilde{g})$ to denote the estimated bias of empirical MSE of $\tilde{g}$ based on the training data, but it is important to note that $\widehat{\text{bias}}_a(\tilde{g})$ is not a function of $\tilde{g}$.

2.4.1. Optimizer with Bias Corrected. We first assume the following assumptions hold:

Assumption 3.

$$\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{Y} - \overline{\lambda}\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{Y} \notin \text{Range}(\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{L} + \eta\mathbf{I} - \overline{\lambda}\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}),$$

$$\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{Y} - \underline{\lambda}\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{Y} \notin \text{Range}(\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{L} + \eta\mathbf{I} - \underline{\lambda}\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}),$$

By Lemma 2.4 of Pong and Wolkowicz (2014), under Assumption 3, $\lim_{\lambda\to\underline{\lambda}} \psi(\lambda) = -\infty$ and $\lim_{\lambda\to\overline{\lambda}} \psi(\lambda) = \infty$.

Let $\tilde{\boldsymbol{\beta}} \in \Omega$ denote some arbitary fitted regression coefficients, and $\tilde{g}(\cdot)$ be the predictive function parametrized with $\tilde{\boldsymbol{\beta}}$. Suppose we have access to the following quantity,

$$\xi(\tilde{\boldsymbol{\beta}}) = \widehat{\text{bias}}_1(\tilde{g}) - \widehat{\text{bias}}_0(\tilde{g}), \tag{2.37}$$

which is essentially an estimation of the bias between empirical risk disparity and the population risk disparity that we aim to correct.

We could then define the a mapping, $L : \Omega \to \Omega$,

$$
\begin{aligned}
L(\tilde{\boldsymbol{\beta}}) := \underset{\boldsymbol{\beta}\in\Omega}{\operatorname{argmin}}\, & \boldsymbol{\beta}^T(\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{L} + \eta\mathbf{I})\boldsymbol{\beta} - 2\boldsymbol{Y}^T\boldsymbol{W}\boldsymbol{L}\boldsymbol{\beta} : \\
& \boldsymbol{\beta}^T\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}\boldsymbol{\beta} - 2\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{L}\boldsymbol{\beta} + \xi(\tilde{\boldsymbol{\beta}}) = C,
\end{aligned} \tag{2.38}
$$

for some suitably chosen C based on Proposition 4.

It is straightfoward to see that under Assumption 3,

$$L(\tilde{\boldsymbol{\beta}}) = \boldsymbol{\beta}(\tilde{\lambda}^*),$$

where $\boldsymbol{\beta}(\cdot)$ is as defined in (2.25), and $\tilde{\lambda}^* \in (\underline{\lambda}, \overline{\lambda})$ is the unique maximizer of the following objective function:

(2.39)

$$
\begin{aligned}
h_{\tilde{\boldsymbol{\beta}}}(\lambda) = & -(\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{Y})^T(\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I})^{-1}(\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{Y}) + \lambda(C - \xi(\tilde{\boldsymbol{\beta}})) \\
& \text{subject to} \quad \boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I} \succeq 0, \boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{Y} \in \text{Range}(\boldsymbol{L}^T(\boldsymbol{W} - \lambda\boldsymbol{D})\boldsymbol{L} + \eta\mathbf{I}).
\end{aligned}
$$

The mapping $L(\tilde{\boldsymbol{\beta}})$ is interpreted as the optimizer of a new constrained optimization where the bias of MSE estimation corresponding to $\tilde{\boldsymbol{\beta}}$ is intended to be corrected. If this mapping has a fixed point, i.e, there exists a $\boldsymbol{\beta}^{**}$ such that

$$\boldsymbol{\beta}^{**} = L(\boldsymbol{\beta}^{**}),$$

then the contraint in (2.38) gurantees that

$$\boldsymbol{\beta}^{**T}\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L}\boldsymbol{\beta}^{**} - 2\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{L}\boldsymbol{\beta}^{**} + \xi(\boldsymbol{\beta}^{**}) = C, \tag{2.40}$$

and hence $\boldsymbol{\beta}^{**}$ is our new optimizer whose bias is corrected.

2.4.2. Mapping on the Lagragian Dual Variable. Since every mapped $L(\tilde{\boldsymbol{\beta}})$ takes the form of $\boldsymbol{\beta}(\lambda)$ in (2.25), we can safely confine our attention to the subspace of Ω where all $\boldsymbol{\beta}$ takes the same form of $\boldsymbol{\beta}(\lambda)$. Let $\tilde{\lambda}$ to denote the optimal Lagrangian value corresponding to $\tilde{\boldsymbol{\beta}}$. We could use the same approximate LOOCV (Section 2.3.4.1) to estimate true conditional risk of $\tilde{g}$,

and denote the estimated bias as follows,

$$\widehat{\text{bias}}_a(\tilde{g}) = \frac{1}{n}\sum_{i=1}^{n} S_i \mathbb{I}(A_i = a) W_{a,i} \left(\frac{Y_i - \hat{Y}_i}{1 - \mathbf{H}_{i,i}}\right)^2 - \frac{1}{n}\sum_{i=1}^{n} S_i \mathbb{I}(A_i = a) W_{a,i} (Y_i - \hat{Y}_i)^2, \tag{2.41}$$

where $\mathbf{H} = \mathbf{H}(\tilde{\lambda})$ for the same $\mathbf{H}(\cdot)$ in (2.35). It follows that

$$\xi(\tilde{\boldsymbol{\beta}}) = \widehat{\text{bias}}_1(\tilde{g}) - \widehat{\text{bias}}_0(\tilde{g}) = \kappa(\tilde{\lambda}),$$

where

$$\kappa(\lambda) := \big(\boldsymbol{Y} - \mathbf{H}(\lambda)\boldsymbol{Y}\big)^T \Big(\boldsymbol{D}\big(\mathbf{I} - \text{diag}(\mathbf{H}(\lambda))\big)^{-2} - \boldsymbol{D}\Big)\big(\boldsymbol{Y} - \mathbf{H}(\lambda)\boldsymbol{Y}\big). \tag{2.42}$$

Note that the optimizer $\tilde{\lambda}^*$ of (2.39) solves the following equation uniquely:

$$\nabla h_{\tilde{\boldsymbol{\beta}}}(\lambda) = C - \xi(\tilde{\boldsymbol{\beta}}) + 2\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{L}\boldsymbol{\beta}(\lambda) - \boldsymbol{\beta}(\lambda)^T(\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L})\boldsymbol{\beta}(\lambda) = 0. \tag{2.43}$$

$\kappa(\tilde{\lambda})$ can be subsituded into (2.43) to replace $\xi(\tilde{\boldsymbol{\beta}})$, and $\tilde{\lambda}^*$ is then the unique solution of the following equation,

$$C - \kappa(\tilde{\lambda}) - \psi(\lambda) = 0, \tag{2.44}$$

with $\psi(\lambda)$ denoting the quantity $2\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{L}\boldsymbol{\beta}(\lambda) - \boldsymbol{\beta}(\lambda)^T(\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L})\boldsymbol{\beta}(\lambda)$.

We could define the solution of (2.44) as the following mapping of $\tilde{\lambda}$,

$$\tilde{\lambda}^* = T(\tilde{\lambda}) := \psi^{-1}(C - \kappa(\tilde{\lambda})).$$

It is then straightfoward to see that if λ^{**} is a fixed point of $T(\tilde{\lambda})$ such that

$$T(\lambda^{**}) = \lambda^{**}, \tag{2.45}$$

$\boldsymbol{\beta}^{**} = \boldsymbol{\beta}(\lambda^{**})$ is a fixed point of map L (2.38).

2.4.3. Modification to Algorithm 1 to Correct Bias. Based on the analysis, we introduce the following two modifications to our method to correct estimation bias: firstly, we refine the case check in Algorithm 1 to identify the interior solution and select the appropriate equality constraint

C; secondly, we modify the iterative step of updating λ_i in Algorithm 6 such that λ converges to a fixed point of (2.45).

We define

$$\phi(\lambda) := \psi(\lambda) + \kappa(\lambda) = \big(\boldsymbol{Y} - \mathbf{H}(\lambda)\boldsymbol{Y}\big)^T \boldsymbol{D}\big(\mathbf{I} - \operatorname{diag}(\mathbf{H}(\lambda))\big)^{-2}\big(\boldsymbol{Y} - \mathbf{H}(\lambda)\boldsymbol{Y}\big) - \boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y}.$$

For the first modification, we compare $\phi(0)$ with the two bounds $-\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta$ and $-\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} + \delta$. If $-\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta \le \phi(0) \le -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} + \delta$, we determine the interior solution and set optimal $\lambda^* = 0$. Otherwise, if $\phi(0) < -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta$, then $C = -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta$, and if $\phi(0) > -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} + \delta$, then $C = -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} + \delta$.

The second modification involves updating λ_i iteratively with a mapping step $\lambda_{i+1} = T(\lambda_i)$. However, this approach may not be very efficient because Newton's method must be repeatedly applied to each mapping step until it reaches convergence. Therefore, instead of using this approach, we opt for a linear approximation of $T(\lambda)$ to update λ. We introduce the following modification to the iterative step in Newton's method outlined in Algorithm 6.

$$\lambda_{i+1} = \lambda_i - \frac{\phi(\lambda_i) - C}{\phi'(\lambda_i)}, \tag{2.46}$$

where with $\phi'(\cdot)$ denoting the derivative of $\phi(\cdot)$.

Our algorithm with bias correction is summarized in Algorithm 2. It important to note that if such fixed point (2.45) does not exist, the Newton's method will not converge. In this case, we will simply use the result of Algorithm 1 without bias correction.

Algorithm 2 Optimization procedure with bias correction

procedure MAIN ALGORITHM
  Calculate initial interval for $\lambda^*, (\underline{\lambda}, \overline{\lambda})$ as in Proposition 2. Check condition in Assumption 3 is satisfied.
  if $-\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta \le \phi(0) \le -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} + \delta$ ▷ Interior solution **then**
    $\lambda^* = 0, \boldsymbol{\beta}^* = \boldsymbol{\beta}(\lambda^*)$
  else if $\phi(0) < -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta$ **then**
    Apply the Newton's method (2.46) to solve $\phi(\lambda) = C$ for $C = -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} - \delta$.
    $\boldsymbol{\beta}^* = \boldsymbol{\beta}(\lambda^*)$.
  else
    Apply modified Newton's method (2.46) to solve $\phi(\lambda) = C$ for $C = -\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{Y} + \delta$.
    $\boldsymbol{\beta}^* = \boldsymbol{\beta}(\lambda^*)$.
  Return $(\lambda^*, \boldsymbol{\beta}^*)$

2.5. Simulation Experiments

In this section, we use two simulations to illustrate the performance of the balanced weights proposed in Section 2.2.6 and our proposed optimization algorithms. In each simulation, we compare the balanced weights with inverse propensity weights (IPW), proposed by Zhang and Long (2021). In particular, we compare with two IPWs, where the propensity score is obtained by fitting selection indicator S to fully observed covariates set U with either a logistic regression (logistic) or a kernel support vector machine (kernel SVM) model with a Gaussian kernel.

2.5.1. Simulation: Parametric Regression. In this experiment, we aim at fitting a parametric model, where the true predictive covariates are available.

Let $X_i = (X_{i1}, X_{i2}, X_{i3}, X_{i4})^T \in \mathbb{R}^4$ follow the gaussian distribution such that $X_{ij} \overset{\text{i.i.d}}{\sim} N(0, \sigma^2_{A_i})$, with $\sigma_0 = 1.5$ and $\sigma_1 = 1$. We let $U_i = (X_{i1}, X_{i2})^T$ denoting the covariates not subject to selection. The selection mechanism is based on the following model $S_i|A_i \sim$ Bernoulli (sigmoid$(\pi_{A_i}(U_i))$), where

$$\pi_0(U_i) = -2 + 0.1\exp((U_{i1} - 1)/1.5) + 0.2((U_{i2} - 1)/0.8)^2,$$

$$\pi_1(U_i) = -1 + 2\sin(U_{i1} - 1)/1.2) - \log(|U_{i2}/5|).$$

Let $Y_i = X_i^T\boldsymbol{\beta} + (-4A_i + 1)X_{i2} + \epsilon$, for $\boldsymbol{\beta} = (1, -1, -1, 1)^T$, and $\epsilon_i \overset{\text{i.i.d}}{\sim} N(0, 1)$ indepdent of X_i and A_i. In each reptition, we generate 2000 observations from each of $A = 1, 0$ groups subject to selection as training data. We then estimate the weights and apply Algorithm 1 with $\delta \in \{0, 0.1, 0.3, 0.5, 0.7, 0.9\}$.

We generate two test datasets with the same data generating mechanism described above, except for one dataset we assume these exists no selection bias in order to assess the predictive model's performance under $p_{\mathcal{T}}$. We generate $100,000$ samples for each group from distribution $p_{\mathcal{T}}$ without selection. To evaluate the model's performance on the test dataset subject to selection bias, we generate another test dataset from $p_{\mathcal{D}}$. We first apply the same weight estimation techniques to this dataset, and then use the weighted MSE to assess the performance of the fitted models. We generate 2000 samples from each group. We summarize the average performance over 500 repetitions in Figure 2.1, Figure 2.2 and Figure 2.3.

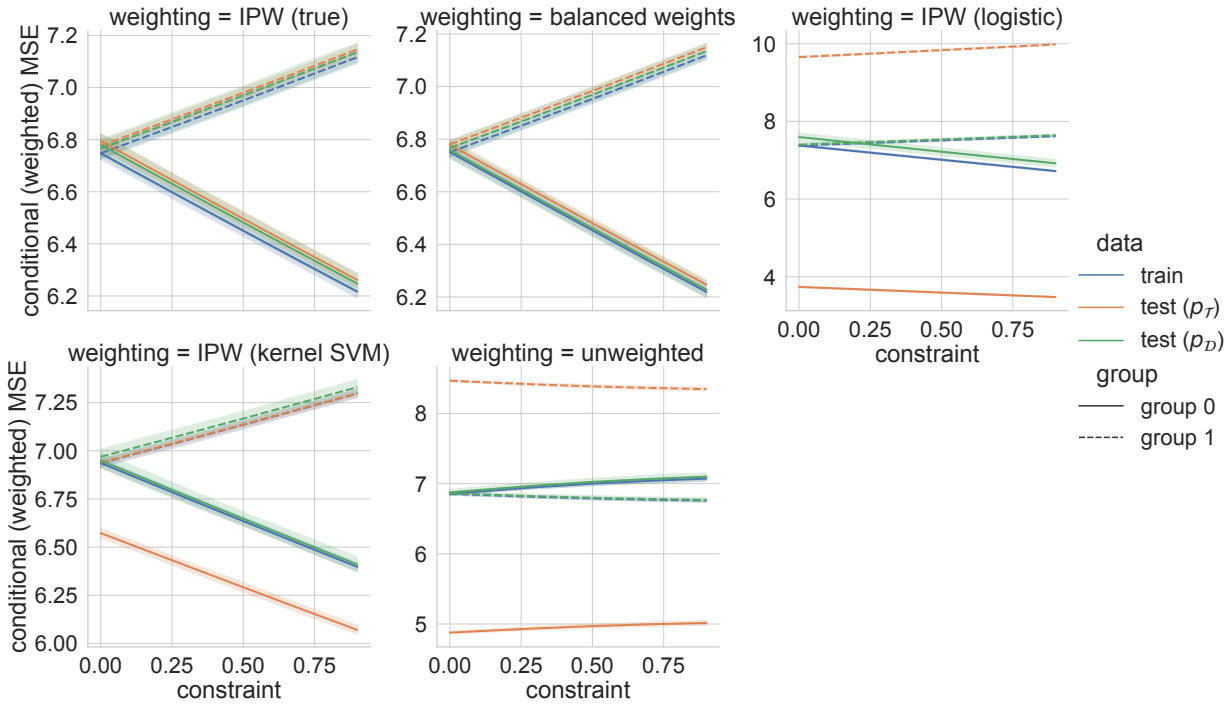


FIGURE 2.1. Conditional (weighted) MSE of parametric weighted fair regressor fitted by Algorithm 1.

In Figure 2.1, we plot the average conditional (weighted) MSE for different weights adjustments, with the band represented the 95% confidence interval. With balanced weights adjustments, the fitted model achieves better performance than adjustments with either logistic IPW or SVM IPW, both of which exhibit large bias in estimating $\mathcal{E}_a(g)$ with a weighted MSE, which is demonstrated by the large differences between the MSEs on the test ($p_{\mathcal{T}}$) and test($p_{\mathcal{D}}$) data. What is more, with balanced weights adjustment, the model achieves a smaller variance in prediction comparing to that with true IPW adjustment. This is expected because when the propensity score is unbounded from below, the true IPW has large variance and thereby imposes extra variation to the training process. As the result, model fitted with true IPW adjustment could not achieve the same trade-off between overall MSE and MSE disparity over $p_{\mathcal{T}}$ as balanced weights adjustments does, as presented in Figure 2.2 and Figure 2.3. The detailed experiments results can be found in Table A.1.

2.5.2. Simulation: Nonparametric Regression. In this experiment, we aim at fitting a nonparametric model, where only the transformed predictive covariates are available.

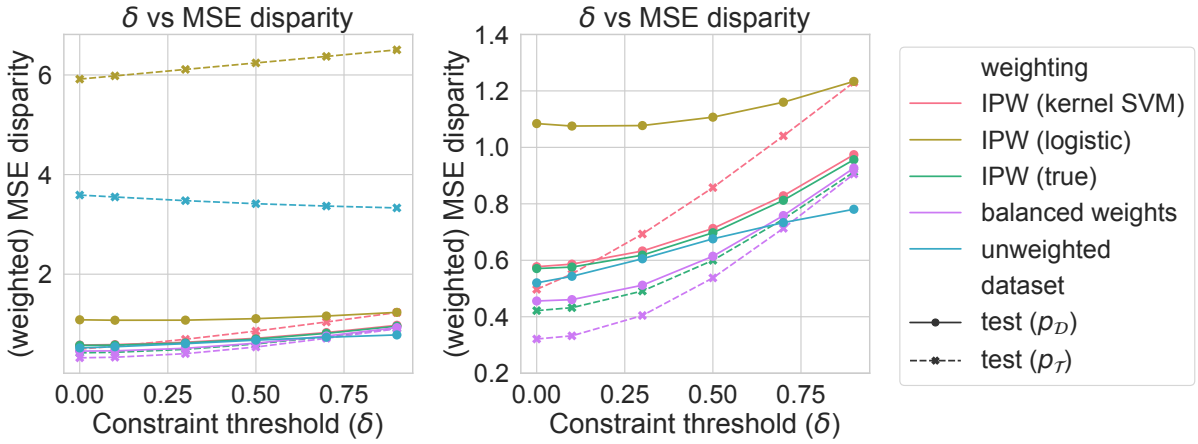


FIGURE 2.2. Performance of fitted parametric fair regressor with different weights adjustments. **Left:** (Weighted) MSE disparity evaluated at various δ; **Right:** Details of the left panel.

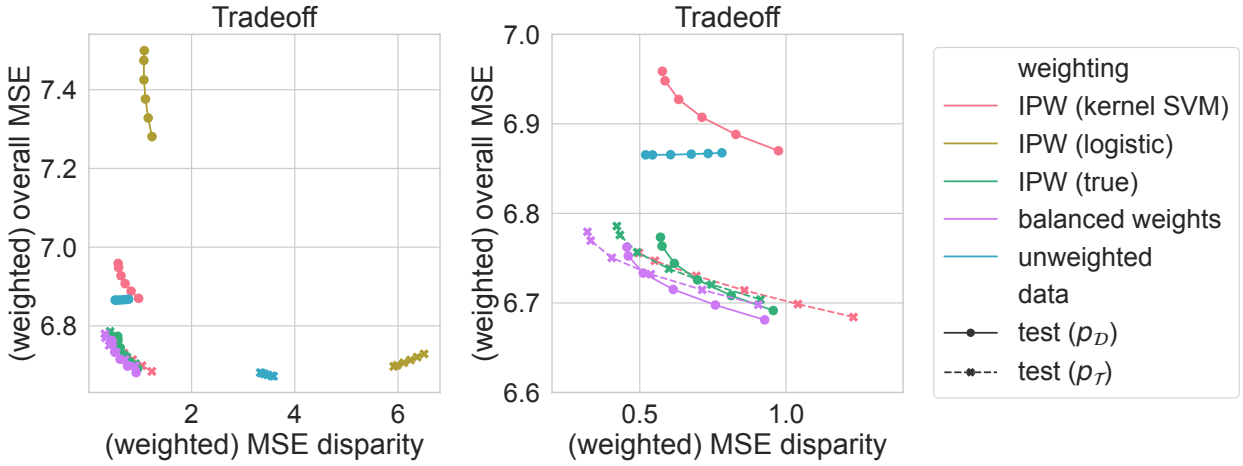


FIGURE 2.3. Performance of fitted parametric fair regressor with different weights adjustments. **Left: Left:** trade-off between overall (weighted) MSE and (weighted) MSE disparity. **Right:** Details of the middle panel.

Let $\widetilde{X}_i = (\widetilde{X}_{i1}, \widetilde{X}_{i2}, \widetilde{X}_{i3}, \widetilde{X}_{i4})^T \in \mathbb{R}^4$ follow the gaussian distribution such that $\widetilde{X}_{ij} \overset{\text{i.i.d}}{\sim} N(1 - 2A_i, 1)$.

The selection mechanism is based on the following model $S_i|A_i \sim \text{Bernoulli}(\text{sigmoid}(\pi_{A_i}(\widetilde{X}_i)))$, where

$$\pi_0(\widetilde{X}_i) = 1 - 0.8\widetilde{X}_{i1} - 0.2\widetilde{X}_{i2} - 0.25\widetilde{X}_{i3} - 0.1\widetilde{X}_{i4},$$

$$\pi_1(\widetilde{X}_i) = 0.5 + 0.5\widetilde{X}_{i1} - 2\widetilde{X}_{i2} - 0.2\widetilde{X}_{i3} - 0.1\widetilde{X}_{i4}.$$

The observed covariates $X_{i1} = \exp(\widetilde{X}_{i1}/2), X_{i2} = \widetilde{X}_{i2}/(1 + \exp(\widetilde{X}_{i1})) + 10, X_{i3} = (\widetilde{X}_{i3}/25 + 0.6)^3, X_{i4} = (\widetilde{X}_{i2} + \widetilde{X}_{i4} + 20)^2$. We let $U_i = (X_{i1}, X_{i2}, X_{i3}, X_{i4})^T$ denoting the covariates not subject to selection. Let $Y_i = 27.4\widetilde{X}_{i1} + 13.7\widetilde{X}_{i2} + 13.7\widetilde{X}_{i3} + 13.7\widetilde{X}_{i4}) + (-2A_i + 0.5)\widetilde{X}_{i1}\widetilde{X}_{i3} + \epsilon_i$ with $\epsilon_i \overset{\text{i.i.d}}{\sim} N(0, 1)$ indepdent of $\widetilde{X}_i$ and A_i.

For each reptition, we generate 2000 observations from each of $A = 1, 0$ groups subject to selection. We then estimate the weights and apply Algorithm 1 with $\delta \in \{0.1, 0.2, 0.4, 0.8, 1.4\}$.

Similar to the parametric model simulation, we generate two test datasets from both $p_{\mathcal{T}}$ and $p_{\mathcal{D}}$. For each repetition, we generate 1000 samples from each group for both test datasets. On the test dataset subject to selection bias we use the weighted MSE to assess the performance of the fitted model. We summarize the average performance over 500 repetitions in Figure 2.4 and Figure 2.6. With the balanced weights adjustments, the fitted model achieve the most stable MSE on both test datasets. It achieve the smallest MSE disparity under all levels of δ. Note that the MSE disparity starts to increase for some small values of δ partially due to the existence of bias discussed in Section 2.4. However, the weighted MSE on $p_{\mathcal{D}}$ dataset still serves as a good estimator of the risk over $p_{\mathcal{T}}$. Details of the experiment results can be found in Table A.2.

Next, we apply the Algorithm 2 to the same dataset. We present the performance with balanced weights adjustments as an illustration in Figure 2.7 and refer the readers to the Appendix A.5 for comprehensive results of other weights adjustments. Figure 2.7 shows that indeed the training MSE for group $A = 0$ in overrestricted in the Algorithm 1. With the bias correction modification, the restriction for group $A = 0$ is relaxed as needed, leading to better MSEs for both groups on both test datasets.

2.6. Benchmark Experiments on Selection Bias Correction

In this section, we compare various weighting techniques alongside the method proposed in Du *et al.* (2022) to adjust selection bias. As Du *et al.* (2022) specifically addresses scenario where only the outcome (Y) in training data is affected by selection machanism, we conduct our experiments under the same conditions.

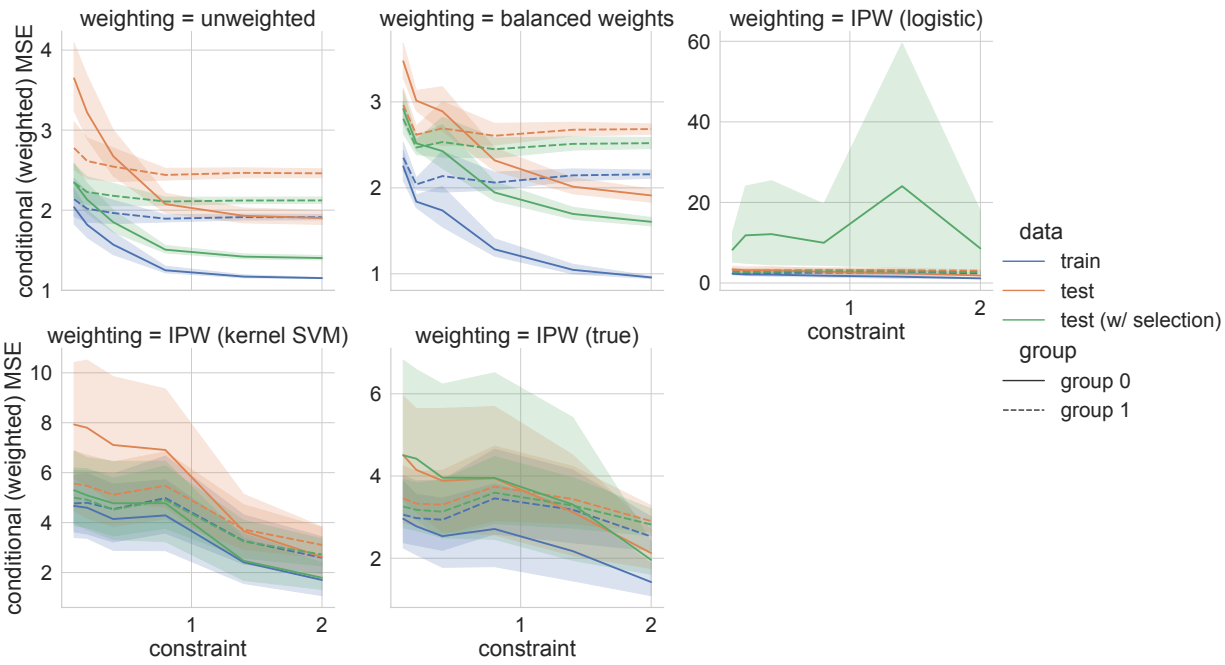


FIGURE 2.4. Conditional (weighted) MSE of nonparametric fair regressor fitted with Algorithm 1.

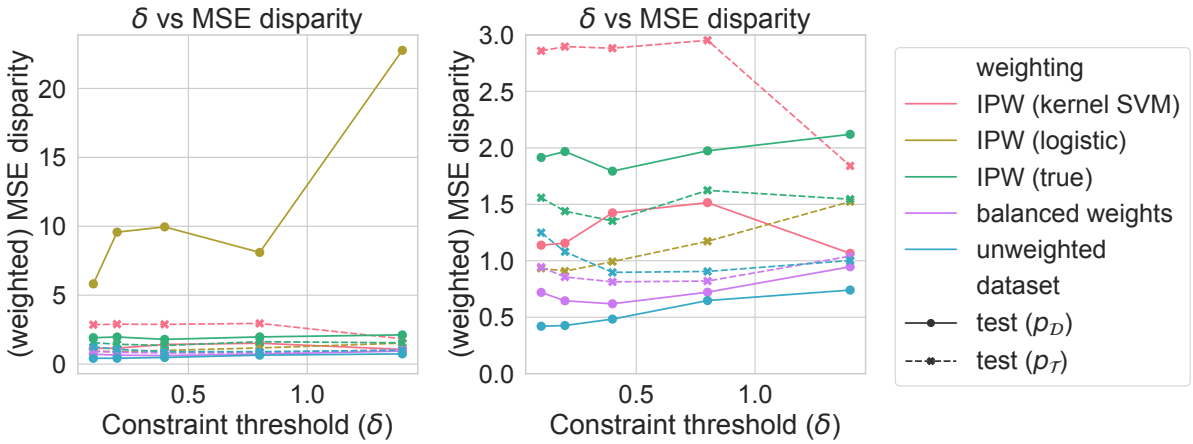


FIGURE 2.5. Performance of fitted nonparametric fair regressor with different weights adjustments. **Left:** (Weighted) MSE disparity evaluated on test dataset at various δ; **Right:** Details of the left panel.

2.6.1. Dataset. We carry out the experiments on Law School Wightman (1998) and Community&Crime Redmond (2009) datasets. We adopt the same data generation and preprocess steps

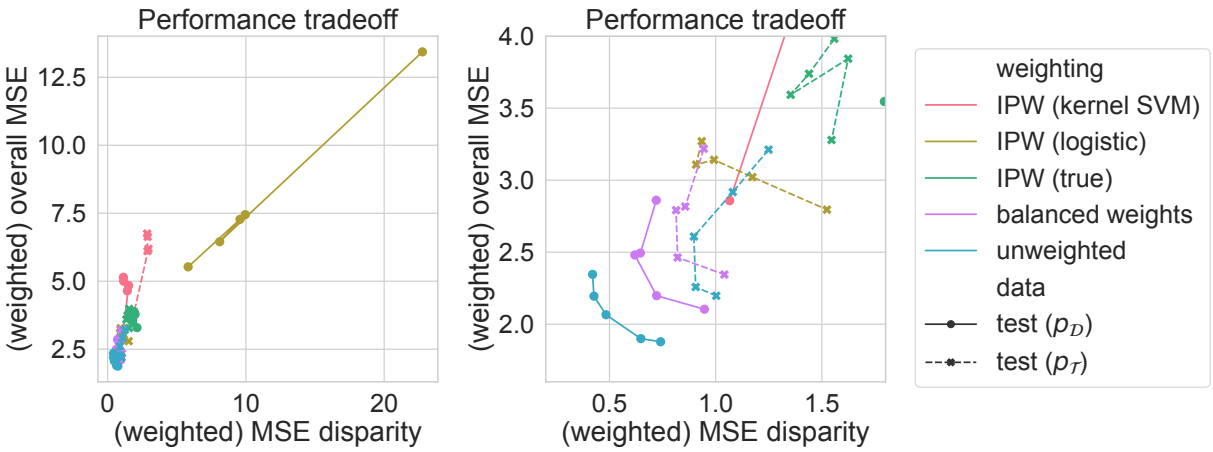


FIGURE 2.6. Performance of fitted nonparametric fair regressor with different weights adjustments. **Left:** trade-off between overall (weighted) MSE and (weighted) MSE disparity; **Right:** Details of the left panel.

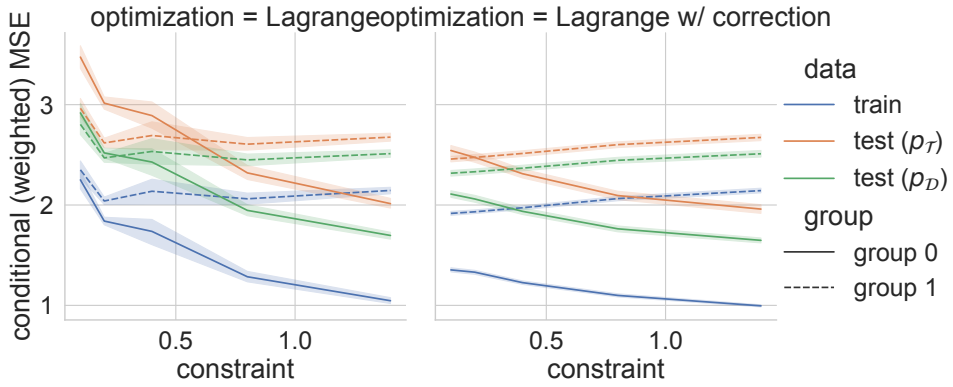


FIGURE 2.7. Conditional (weighted) MSE of nonparametric regressor fitted by weighted fair regression in Algorithm 1 (left) and weighted fair regression with bias correction in Algorithm 2 (right) .

as in Du *et al.* (2022).[1] In both datasets, a binary variable is selected as a protected attribute (A) and a continuous variable is selected as the outcome (Y). We refer readers to Appendix A.4.1 for more details about the datasets.

2.6.2. Method. We examine the performance of balanced weights adjustment. We further compare it with two inverse propensity weights (IPW), proposed in Zhang and Long (2021), where

[1]The source code of implementation as well as data preprocessing can be found at `https://drive.google.com/file/d/1cOdyTXUeaB1f4lxuyY-bCshUykHOqACA/view`.

the propensity score is estimated by fitting S to U using either logistic regression or a kernel support vector machine (kernel SVM) model. Du *et al.* (2022) employed the Heckman model for bias correction, assuming that the selection mechanism is linear in U. They derived the Inverse Mills Ratio (IMR) by modeling S as a probit classification model on U, treating it as an extra covariate alongside X. To prevent multicollinearity in predicting Y, they further assumed that X is a strict subset of U. Du *et al.* (2022) handled the constrained optimization by optimizing its Lagrangian dual form while controlling the constraint by manually tuning the dual variable. In our experiments, we use the default values provided in their code for each dataset.

In Crime Data, we use the selection mechanism based on the following propensity score:

$$S_i \sim \text{Bernoulli}(\Phi(-1.5\text{“NumUnderPov”} + 0.4)),$$

where $\Phi(\cdot)$ represents the cumulative distribution function of the standard normal distribution, and “NumUnderPov” attribute denotes the ratio of people under the poverty level in a community.

We carry out two experiments on this dataset, where the ”NumUnderPov” attribute is either included in the set of U or not, mirroring the setting used in Du *et al.* (2022).

In Law Data, we use the selection mechanism based on the following propensity score:

$$S_i \sim \text{Bernoulli}(\text{sigmoid}(-0.2\text{“age”} - 10.5)),$$

and “age” attribute is included in U as Du *et al.* (2022)’s setting.

2.6.3. Results and Analysis. In all experiemnts, we set constraint $\delta = 0.001$. For each repetition, we randomly split the dataset into 70% training and 30% testing sets. We apply Algorithm 1 with a model linear X on the training set and assess its performance on the testing set, where Y is observed for all samples. We summarize the average performance of 50 repetitions for each method in Figure 2.8. Our observations are as follows: (1) With Lagrangian optimization, we achieve a smaller MSE and MSE disparity level on both datasets compared to the optimization method proposed in Du *et al.* (2022). (2) Among all weighting techniques, balanced weights achieves the smallest MSE disparity for all three experiments.

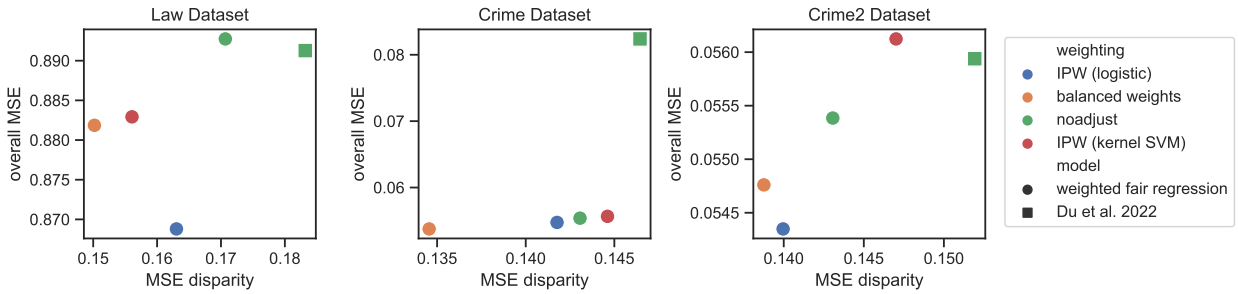


FIGURE 2.8. Overall performance: average test MSE and MSE disparity of different methods in two datasets. "NumUnderPov" attribute is not included in U in the Crime2 dataset.

2.7. Benchmark Experiments without Selection Bias

In this section, we assess the performance of our method proposed in Algorithm 1 and its modification with bias correction in Algorithm 2. Additionally, we compare our methods with the CENet method and WassersteinNet method [2] proposed in Chi *et al.* (2021). Both the CENet and WassersteinNet methods consider MSE disparity as the fairness metric and assume training dataset is representative of the test dataset. Therefore, to better compare with these two methods, we focus on the scenario without selection bias and do not apply weights adjustment.

2.7.1. Dataset. We conduct experiments on Law School Wightman (1998), Communities&Crime Redmond (2009) and Medical Insurance Cost Lantz (2013) datasets. We adopt the same data generation and preprocess steps as in Chi *et al.* (2021). For each dataset, a binary variable is selected as the protected attribute (A) and a continuous variable is selected as the outcome (Y). We refer readers to Appendix A.4.2 for more details on the datasets and preprocessing steps.

2.7.2. Methods. We train both ridge regression (RR) and kernel ridge regression (KRR) with different constraint levels δ with Algorithm 1. For KRR models, we further incorporate the bias correction with Algorithm 2. We test our methods across a wide range of δ values to understand its performance under various constraints. Both the CENet and WassersteinNet method aim to minimize the MSE disparity of neural network models by fitting them on a fair representation that minimizes the distance between two groups Chi *et al.* (2021). Algorithmic fairness is controlled

[2]The source code of implementation of both methods can be found at `https://github.com/JFChi/Understanding-and-Mitigating-Accuracy-Disparity-in-Regression`

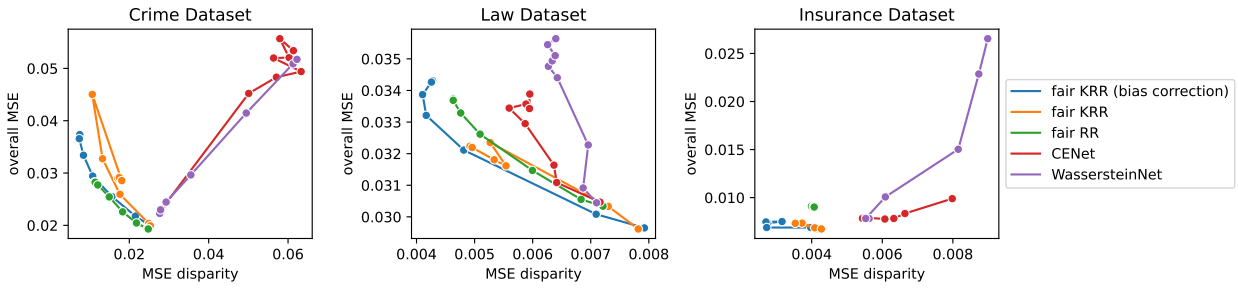


FIGURE 2.9. Overall performance: average test MSE and MSE disprepancy of different methods.

by a trade-off parameter λ, which penalizes the distance between conditional representations. In our experiments, we use the default neural network architectures from Chi *et al.* (2021) and vary λ across a wide range of values. Further details on parameter tuning for all methods can be found in Appendix A.4.2.

2.7.3. Results and Analysis. For each reptition, we randomly split the data into 80% training and 20% testing set. We train the models on training data and assess their performance on testing set. The average metrics at each value of δ or λ over 10 random splits are summarized in Figure 2.9. Here are our observations: (1) Our proposed methods achieve the best trade-offs in all datasets. (2) Our methods effectively reduce the MSE disparity. Among the three models, KRR with bias correction achieves the smallest MSE disparity. (3) Compared to our methods, the CENet method and WassersteinNet method can reduce the MSE disparity in some cases (such as in the Law Dataset), but overall, increasing λ does not consistently reduce the MSE disparity. We present other experiment details in Figure A.4.

CHAPTER 3

Advances on Seeded Region Growing on Graph

3.1. Introduction

In the realm of observational high-energy astrophysics, the advent of high-resolution imaging poses new challenges for data analyses. Unlike images captured at longer wavelengths, X-ray and γ-ray data exhibit complex features that are intricate to discern and characterize. A notable challenge in this domain is the analysis of the boundaries of extended X-ray sources, especially in scenarios characterized by low photon counts (e.g., van Dyk *et al.*, 2006).

To address this challenge, event-list segmentation algorithms have been developed (e.g., Scargle, 2002; Byers and Raftery, 2002). Recently, in Fan *et al.* (2023), we introduced an innovative non-parametric event list segmentation algorithm designed to partition a field of view into distinct emission components. The method constructs the segmentation using seeded region growing (Adams and Bischof, 1994) on a graph (hence its name, `SRGonG`, Seeded Region Growing on a Graph). Starting with a Voroni tessilation built on the location of the recorded photons, `SRGonG` sequntially grows an initial set of seeds along the dual graph of the Voroni tessilation until every photon is assigned to exactly one seed, thus forming a segmentation of the full event list. In practice, a large number of seeds are used to over segment the list. This allows us to use a combination of a greedy merging algorithm and the Bayesian Information Criterion (BIC) to selectively merge segments. In this way, we arrive at the penalized profile maximum likelihood estimate of the number of segments, the segmentation itself, and the segment-wise constant intensity function across the field of view. `SRGonG` is tailored to identify both point sources and diffuse spatial structures, even with sparse event list data.

Although the simulation studies conducted in Fan *et al.* (2023) demonstrate the ability of `SRGonG` to seperate event-list photons that are associated with an extended source from photons from a diffuse background, unsurprisingly the quality of the segmentation tends to degrade when

the S/N is low. This degredation is observed both in the general cases (as measured by the median of the Adjusted Rand Index or ARI) and the worst-case scenario (as measured by the lower bound of the ARI). To address this, we modify the greedy algorithm used by `SRGonG` to merge segments within the over-segmented event list, known as the Greedy Merger. Specifically, we devlop the Randomized Merger and the Beam Merger, which improve the performance of `SRGonG` under both a general case and the worst case, respectively. To enable the user to flexibily balance performance and computation time, we also devise a hybrid merging process that optimally switches between the Greedy Merger and Randomized/Beam Merger.

This chaprter is organized as following: Section 3.2 provides an short introduction on `SRGonG` . In Section we introduce our modidications to the Greedy Merger step and illustrate the effectiveness of our modidications through comprehensive simulation experiemnts.

In Section 3.4, we introduce an additional modifications to `SRGonG` : adjustment for the change of exposure map. This modificaiton aims to make `SRGonG` more flexible for real data applications.

3.2. `SRGonG`: Seeded Region Growing on Graphs

Instruments like the *Chandra X-Ray Observatory* record the energy, sky coordinates, and arrival time of X-ray photons. In this chapter, we confine our attention to the locations recorded by the sky coordinates, i.e., we ignore the instrumental point spread function. We model the recorded sky coordiantes of a collection of the n photons over a field of view $\mathcal{F}$,

$$\mathbf{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\},$$

as an inhomogeneous Poisson process. This means that the number of photons, $n \sim \text{Pois}(\mathbf{\Lambda})$, and that given n the sky coordinates are distributed $\mathbf{x}_i \overset{\text{iid}}{\sim} \lambda(\mathbf{x})/\mathbf{\Lambda}$, where $\lambda(\mathbf{x})$ is an intergrable intensity function over $\mathcal{F}$, and $\mathbf{\Lambda} = \int_{\mathcal{F}} \lambda(\mathbf{x}) d\mathbf{x}$. We further assume that $\mathcal{F}$ can be partitioned into K segments, denoted $\mathbf{S} = (\mathcal{S}_1, \ldots, \mathcal{S}_K)$, where $\lambda(\mathbf{x})$ is constant with intensity λ_k on each $\mathcal{S}_k$; we denote $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_K)$. Thus the intensity function can be expressed as

$$\lambda(\mathbf{x}) = \sum_{k=1}^{K} \lambda_k \mathbb{1}_{\mathcal{S}_k}(\mathbf{x}), \tag{3.1}$$

where $\mathbb{1}_{\mathcal{S}_k}(\mathbf{x})$ is and indicator function for $\mathbf{x} \in \mathcal{S}_k$.

The `SRGonG` algorithm Fan *et al.* (2023) is a novel non-parametric event-list segmentation algorithm for obtaining penalized maximum likelihood estimates of $\lambda(\mathbf{x})$, $\mathbf{S}$, and K under an BIC-type penality. The algorithm starts with the Voronoi diagram of the sky coordinates and uses a seeded region growing (SRG) step to combine voronoi cells into an oversegmented dataset. SRG requires that each cell have a well-defined set of neighbors and Fan *et al.* (2023) propose using the dual graph of the voronoi cell, known as the Delaunay triangulation for this purpose. The voronoi diagram and Delaunay triangulation are dual structures. The Voronoi cell for a photon is the spatial region that is closer to thisphoton than to any other photon in the dataset. The Delaunay triangulation is constructed from the Voronoi diagram by connecting the centers of adjacent Voronoi cells (Fan *et al.*, 2023, e.g., Section 2.2).

3.2.1. SRG Step. The penalized likelhood function favors segments that include Voronoi cells with similar estimated intensities. Following Fan *et al.* (2023), let $\widehat{\boldsymbol{\lambda}}_i^{\mathcal{V}_i}$ denote the estimated intensity for the cell corresponding to photon $\mathbf{x}_i$, i.e., $\widehat{\boldsymbol{\lambda}}_i^{\mathcal{V}_i}$ is the reciprical of the area of this cell. Similarly the estimated intensity of segment $\mathcal{S}$ is $\text{Num}(\mathcal{S})/\text{Area}(\mathcal{S})$, where $\text{Num}(\mathcal{S})$ is the number of photons with sky coordinates in $\mathcal{S}$ and $\text{Area}(\mathcal{S})$ is the area of $\mathcal{S}$. To obtain the oversegmentation, the SRG step starts with an overly large set of M_S seeds and at each iteration aggregates one unclustered photon, $\mathbf{x}_i$, into one segment, $\mathcal{S}_k$, where the (neighboring) photon and segment (i and k) are selected to maximize the change in penalized log-likelihood, i.e., by minimize the difference in their log estimated intensities,

$$\delta(i, \mathcal{S}_k) = \left| \log \widehat{\boldsymbol{\lambda}}_i^{\mathcal{V}_i} - \log\{\text{Num}(\mathcal{S}_k)/\text{Area}(\mathcal{S}_k)\} \right|. \tag{3.2}$$

This procedure continues until all photon in $\mathcal{F}$ are assigned to exactly one segment. Since we aim to have at least one seed set for each potential segment, but do not know the true number and location of the segments, we recommend using a large number of seeds, which leads to an oversegmentation of the data.

3.2.2. Greedy Merger Step. `SRGonG` uses a greedy algorithm to sequentially merge pairs of segments, reducing the estimate of K until the penalized profile likelihood is maximized. As noted, Fan *et al.* (2023) uses a BIC-type penality function that favors fewer segments, see Equation 3.4.

They also profile over $\boldsymbol{\lambda}$, replacing each λ_k with its maximum likelihood estimate given K and $\mathbf{S}$. Specifically, in a step with K segments, if two neighbouring segments, $\mathcal{S}_i$ and $\mathcal{S}_j$, are merged, the change in penalized log-likelihood is

$$\begin{aligned}\Delta\text{BIC}_{K,i,j} =& 2\ \text{Num}(\mathcal{S}_i)\log\frac{\text{Num}(\mathcal{S}_{i\cup j})\text{Area}(\mathcal{S}_i)}{\text{Area}(\mathcal{S}_{i\cup j})\text{Num}(\mathcal{S}_i)} \\ &+ 2\ \text{Num}(\mathcal{S}_j)\log\frac{\text{Num}(\mathcal{S}_{i\cup j})\text{Area}(\mathcal{S}_j)}{\text{Area}(\mathcal{S}_{i\cup j})\text{Num}(\mathcal{S}_j)} + m_{\text{seg}}\log n\,.\end{aligned} \tag{3.3}$$

Thus, the pair, $\mathcal{S}_i$ and $\mathcal{S}_j$, are selected to minimize $\Delta\text{BIC}_{K,i,j}, 1 \le i < j \le K$ resulting from their merger. By stepping through this procedure, we obtain a sequence of segmentations, with $k = 1, \ldots, M_S$ segments. The segmentation with the largest penalized profile likelihood is returned as the output of the greedy algorithm, where the penalized profile log-likehood value (BIC) is defined as

$$\text{BIC}(K, \mathbf{S}) = -2\mathcal{L}_{\text{profile}}(K, \mathbf{S}|\mathbf{X}, n) + Km_{\text{seg}}\log n, \tag{3.4}$$

with

$$\mathcal{L}_{\text{profile}}(K, \mathbf{S}|\mathbf{X}, n) = \sum_{k=1}^{K}\text{Num}(\mathcal{S}_k)\text{log}(\text{Num}(\mathcal{S}_k)/|\mathcal{S}_k|) - n - \text{log}n!. \tag{3.5}$$

The boundaries of a segmentation are extracted using the Voronoi diagram. If two photons located in different segments are adjacent in Delaunay triangulation, the edge shared by their Voronoi cells is taken as part of their common boundary. The boundary of each segment is the polygon constructed by connecting its set of edges.

3.3. Improving the Greedy Merger Step in `SRGonG`

Unfortunately, the merging procedure described in Section 3.2.2 is not robust to noise in the data. It is well known that step-wise greedy model selection algorithms, e.g., forward stepwise regression Efron *et al.* (2004) and the matching pursuit (MP) algorithm Mallat and Zhang (1993), in general, provide suboptimal solutions. From extensive simulations, we find this to be particularly problematic when the sample size, n, or SNR are low or when there is a large number of fitted segments.On the other hand, an exhaustive search leads to the optimal solution but is extremely

time-consuming when there is a large set of possible segmentations We, therefore, propose two alternative Merger Steps to improve the quality of the segmenation, one based on a randomized search and one based on a beam search. Both search strategies achieve a reasonable trade-off between time and segmentation quality.

3.3.1. Randomized Merger Step. A randomized search introduces randomization to each merging step by allowing sub-optimal steps and therefore expanding the search space. In each merge step, with say K associated segments, we label each segment $1 \leq i < j \leq K$ and compute $\Delta\text{BIC}_{K,i,j}$ for each neighbouring pair, (i, j). Keeping track of the M_C candidate pairs (i, j) with the smallest $\Delta\text{BIC}_{K,i,j}$, with equal probability, we randomly select one pair to merge. In this way, we obtain a sequence of nested segmentations, $\{\widehat{\mathbf{S}}(K), K = 1, \ldots, M_S\}$, each with a BIC value; we set $\widehat{K}$ and $\widehat{\mathbf{S}}(\widehat{K})$ to the values of the nested level with the smallest BIC as computed with Equation 3.4. This procedure is repeated M_R times, resulting M_R segmentations $\widehat{\mathbf{S}}^{(m_r)}(\widehat{K}^{(m_r)})$, for $m_r = 1, ..., M_R$. Finally, we select the fitted segmentation to be the $\widehat{\mathbf{S}}^{(m_r)}$ and $\widehat{K}^{(m_r)}$ that minimizes the BIC across all M_R segmentations.

3.3.2. Beam Merger Step. Beam search is another solution that provides more flexibility than the Greedy Merge Step. It is a heuristic search algorithm that explores the space of possible segmentations by expanding the most promising nodes in a limited set; it has been widely used in fields like speech recognition (Tillmann and Ney, 2003). We use one buffer to store M_B distinct possible segmentation for each $1 \leq K \leq M_S$. Initially, for $K = M_S$, the buffer contains one segmentation. For a fixed $K < M_S$, we compute the BIC of all possible mergers of all the segmentations, $\widehat{\mathbf{S}}^{(m_b)}(K + 1)$, $m_b = 1, ..., M_B$ in the Level $K + 1$ buffer, where the mergers we consider are segementations of size K formed by merging two neighboring segements in one of the segmentations of size $K + 1$ in the Level $K + 1$ buffer. The distinct mergers with the smallest M_B BIC values are stored in the Level K buffer. This is repeated until $K = 1$ and the segmentation with the smallest BIC value across all buffers is returned. When $M_B = 1$, the beam search is the same as the greedy search. As M_B increases, the result converges to that of an exhaustive search.

3.3.3. Hybrid Merging Process. The computational times for the Randomized Merger and the Beam Merger Steps scale with M_S times M_R and M_B, respectively. In practice, a large value of

M_S impacts algorithmic efficiency. We propose a hybrid merging process that combines the Greedy Merger Step and either the Randomized or Beam Merger Step and offers more flexibility between computation cost and model performance. Specifically, we excute the Greedy Merger Step until it reaches $K = M_s$, where M_s is a predetermined hyperparameter. Subsequently, we transition to either the Randomized or Beam Merger Step. The value of M_s should be set to be significantly larger than a reasonable guess as to the true underlying K; we typically set $M_s = 150$.

3.3.4. Simulations. We adapt the first simulation experiment in Section 4 of Fan *et al.* (2023) to compare the Greedy, Random, and Beam Merger Steps under several scenarios. In all simulation settings, four point-like sources of area 0.08 are embedded within a circular extended source of area 0.2. Each point-like source has $\beta\sigma$ counts, the diffuse extended source has $10\beta\sigma$ uniformly distributed counts, and there are 1000β background counts in expectatio spread uniformly across $\mathcal{F}$.

Our experiments are organized into a $3 \times 3 \times 3$ design, where the first factor, β, takes on values 0.5, 1, and 2, representing small, medium, and large sample size; the second factor, σ, takes on values 10, 20, and 30, representing small, medium, and large S/N; the third factor is the specification of grid size and seed size for the SRG Step.[1] We generate 500 replecate datasets for each of 9 (β, σ) pairs of simulation settings and apply the 3 settings of seeds specification to each dataset. We set the beam width to $M_B = 100$, the randomized candidate size to $M_C = 3$, and the iteration number to $M_R = 100$.

As in Fan *et al.* (2023), performance is evaluated via the *Adjusted Rand Index*. Figure 3.1 summarizes the median ARI of the 500 replicates for each of the 9 simulation settings and five seeds specification, with error bars representing the range of the first and third quartiles of the ARI. When the S/N is low ($\sigma = 10$), both the Beam and Randomized Merger Steps generally achieve a better ARI than the Greedy Merger Step. The ARI interquartile range for both the Beam and the Randomized Merger Step are smaller than that obtained with the Greedy Merger step. This illustrates the two alternative Merger Steps are more robust to noise and also more stable than the Greedy Merger Step when the S/N is low. Although the improvement is not significant with a

[1]We consider three combinations of grid size and seed size: (i) a 5×5 grid with and a seed size of 5, (ii) a 9×9 grid with a seed size of either 3,and (iii) a 17×17 grid with and a seed size of 1.We use a neighborhood size of 50 for the local maxima seed specification in addtion to the regular grid.

5×5 grid for the other two S/N settings, we observe both the Beam and the Randomized Merger Step still achieve higher lower bounds than that obtained with the Greedy Merger step. We also observe an increasingly larger boost in ARI with a larger number of seeds, and the improvement of Randomized or Beam Merger over the Greedy Merger is most significant for 17×17 grid. As discussed in Section 3.2, we recommand using a large number of initial seeds, and this is precisely the case when using the Randomized or the Beam Merger Step yields the largest improvement over the Greedy Merger Step.

Theoretically the computing time of the Randomized and Beam Merger Steps equal the computing time of the Greedy Merger Step times M_R and M_B, respectively. In practice, however, the Randomized Merger Step can proceed in a parallel, and is faster than the Beam Merger Step when M_R and M_B are equal. As an illustration, the average computing times of the Greedy, Randomized, and Beam Merger Steps with a 5×5 grid of seeds, $\sigma = 30$, and $\beta = 2$ are 0.117, 0.667 and 2.622 seconds repectively (using a 2021 Macbook Pro with M1 Pro Chip and 16GB memory).

3.4. Adjustment for Exposure Map

In the development of `SRGonG`, it was always assumed that the exposure map remained constant across the entire $\mathcal{F}$. The exposure map in a specific observation integrates data concerning the telescope's effective area and the detector, along with a record of the dwell time as a function of sky position. This dwell time information is compiled by tracking the telescope's pointing movements throughout the observation period Houck (-). In this section, we extend the `SRGonG` to adapt to a varying exposure across $\mathcal{F}$. Specifically, we consider a simplied scenario that the exposure map depends only on the dwell time and the sky position, i.e., exposure map is given as a measurable function $p(\mathbf{x})$ of position $\mathbf{x}$ over $\mathcal{F}$.

3.4.1. Profile loglikelihood. Given the function $p(\mathbf{x})$ defined on $\mathcal{F}$, we can rewrite the intensity function of the inhomogeneous Poisson process as

$$\tilde{\lambda}(\mathbf{x}) = \sum_{k=1}^{K} \lambda_k \mathbb{1}_{\mathcal{S}_k}(\mathbf{x}) p(\mathbf{x}), \tag{3.6}$$

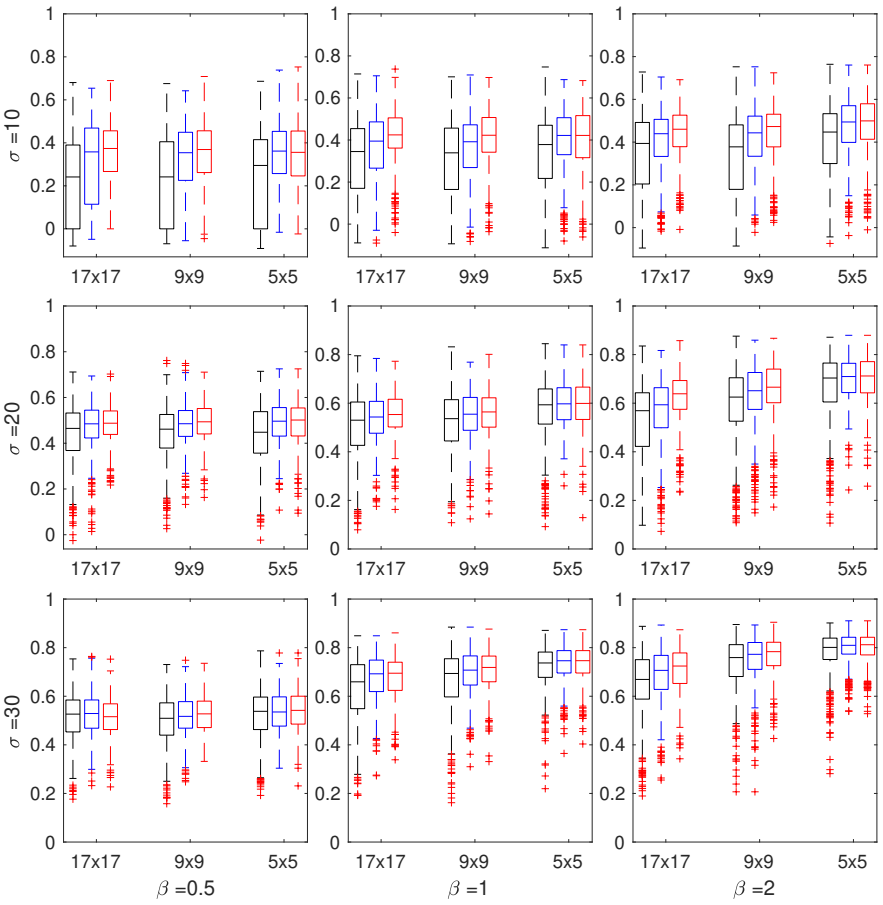


FIGURE 3.1. Simulation study comparing the Greedy (black), Randomized (red), and Beam (blue) Merger Steps. The boxplot of the Adjust Rand Index (ARI) obrtained on each of 500 replications is plotted.

Let $\tilde{\mathbf{\Lambda}} = \int_{\mathcal{F}} \tilde{\lambda}(\mathbf{x}) d\mathbf{x}$. Then $n \sim \text{Pois}(\tilde{\mathbf{\Lambda}})$ and given n the $\mathbf{x}_i \overset{\text{iid}}{\sim} \tilde{\lambda}(\mathbf{x})/\tilde{\mathbf{\Lambda}}$. Thus, their joint probability mass/density function under the inhomogeneous Poisson process is

$$\begin{aligned} f(\mathbf{x}_1, \ldots, \mathbf{x}_n, n) &= f(\mathbf{X} \mid n) \cdot f(n) \\ &= \frac{1}{\tilde{\mathbf{\Lambda}}^n} \prod_{i=1}^{n} \tilde{\lambda}(\mathbf{x}_i) \cdot \frac{\exp\left(-\tilde{\mathbf{\Lambda}}\right) \tilde{\mathbf{\Lambda}}^n}{n!} \\ &= \frac{\exp\left(-\tilde{\mathbf{\Lambda}}\right)}{n!} \prod_{i=1}^{n} \tilde{\lambda}(\mathbf{x}_i) \end{aligned} \tag{3.7}$$

and their log-likelihood function is given by

$$\mathcal{L}(K, \mathbf{S}, \boldsymbol{\lambda} \mid \mathbf{X}, n) = \log f(\mathbf{X}, n) = \sum_{i=1}^{n} \log \tilde{\lambda}(\mathbf{x}_i) - \int_{\mathcal{F}} \tilde{\lambda}(\mathbf{x}) d\mathbf{x} - \log n! \ , \tag{3.8}$$

Replacing $\tilde{\lambda}(\mathbf{x})$ with the intensity expression given in (3.6), we have

$$\sum_{i=1}^{n} \log \tilde{\lambda}(\mathbf{x}_i) = \sum_{k=1}^{K} \left(\mathrm{Num}(\mathcal{S}_k) \log \lambda_k + \sum_{\mathbf{x} \in S_k} \log p(\mathbf{x}) \right) = \sum_{k=1}^{K} \mathrm{Num}(\mathcal{S}_k) \log \lambda_k + \sum_{i=1}^{n} \log p(\mathbf{x}_i), \tag{3.9}$$

and

$$\begin{aligned} \int_{\mathcal{F}} \tilde{\lambda}(\mathbf{x}) d\mathbf{x} &= \int_{\mathcal{F}} \sum_{k=1}^{K} \lambda_k \mathbb{1}_{\mathcal{S}_k}(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} \\ &= \sum_{k=1}^{K} \lambda_k \int_{\mathcal{F}} \mathbb{1}_{\mathcal{S}_k}(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} \\ &= \sum_{k=1}^{K} \lambda_k \int_{\mathcal{S}_k} p(\mathbf{x}) d\mathbf{x}. \end{aligned} \tag{3.10}$$

Therefore,

$$\mathcal{L}(K, \mathbf{S}, \boldsymbol{\lambda} \mid \mathbf{X}, n) = \sum_{k=1}^{K} \left(\mathrm{Num}(\mathcal{S}_k) \log \lambda_k + \left(\lambda_k \int_{\mathcal{S}_k} p(\mathbf{x}) d\mathbf{x} \right) \right) + \sum_{i=1}^{n} \log p(\mathbf{x}_i) - \log n! \ . \tag{3.11}$$

We aim to maximize $\mathcal{L}$ as a function of K, $\mathbf{S}$, and $\boldsymbol{\lambda}$ to obtain their maximum likelihood estimates. For fixed K and $\mathbf{S}$, $\mathcal{L}$ is maximized as a function of $\boldsymbol{\lambda}$ by

$$\widehat{\lambda}_k = \mathrm{Num}(\mathcal{S}_k) / \left(\int_{\mathcal{S}_k} p(\mathbf{x}) d\mathbf{x} \right) \quad \text{for} \quad k = 1, \ldots, K \, . \tag{3.12}$$

Plugging the $\widehat{\lambda}_k$ into $\mathcal{L}$, we obtain the profile log-likelihood of K and $\mathbf{S}$, i.e.,

$$\mathcal{L}_{\text{profile}}(K, \mathbf{S} \mid \mathbf{X}, n) = \sum_{\substack{k=1 \\ \mathrm{Num}(\mathcal{S}_k) \neq 0}}^{K} \mathrm{Num}(\mathcal{S}_k) \log \left(\frac{\mathrm{Num}(\mathcal{S}_k)}{\int_{\mathcal{S}_k} p(\mathbf{x}) d\mathbf{x}} \right) + \sum_{i=1}^{n} \log p(\mathbf{x}_i) - n - \log n!. \tag{3.13}$$

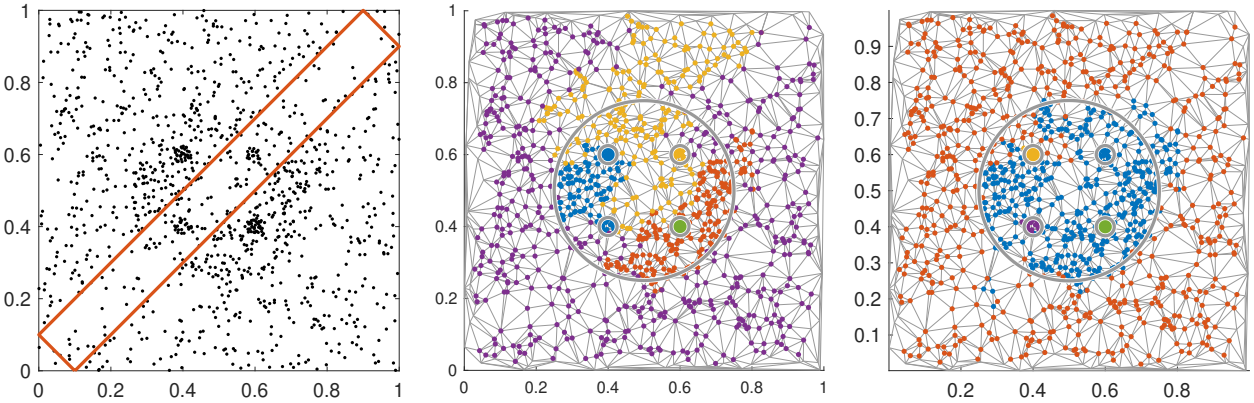


FIGURE 3.2. Simulation illustration of exposure adjustment. **Left**: Simulated photon locations for one simulation instance with $\beta = 1$ and signal-to-noise ratio $\sigma = 30$. The exposure map is constant function in and outside the rectangle with length ≈ 1.2 and width ≈ 0.14 across diagonal (solid red lines). The exposure rate in the effective area is 0.5. **Middle**: `SRGonG` Segmentation without adjustment. The ground truth boundaries are highlighted with grey lines. **Right**: `SRGonG` segmentation with adjustment.

Let $\tilde{\text{Area}}(\mathcal{S}_k) = \int_{\mathcal{S}_k} p(\mathbf{x})d\mathbf{x}$ denoted the normalized area for each segment. Let $\widetilde{\boldsymbol{\lambda}}_i^{\mathcal{V}_i} = 1/\left(\int_{\mathcal{V}_i} p(\mathbf{x})d\mathbf{x}\right)$ denote the normalized voronoi intensity of each photon. The operation in seeded region grow step and oversegmetation merge step follows by substituting the $\widehat{\boldsymbol{\lambda}}_i^{\mathcal{V}_i}$ and $\text{Area}(\mathcal{S}_k)$ by $\widetilde{\boldsymbol{\lambda}}_i^{\mathcal{V}_i}$ and $\tilde{\text{Area}}(\mathcal{S}_k)$.

3.4.2. Simulation. We adapt the same simulation setting in Section 3.3.4. In particular, we use sample size $\beta = 0.5, 1, 2$ and signal-to-noise ratio $\sigma = 10, 20, 30$ for the experiment. Let exposure map be described by a piece-wise constant function,

$$p(\mathbf{x}) = p\mathbb{1}_{E_1}(\mathbf{x}) + \mathbb{1}_{E_2}(\mathbf{x}), \tag{3.14}$$

where E_1 denote the rectangle across the diagonal of $\mathcal{F}$ visiualized in Figure 3.2 and $E_2 = \mathcal{F} \bigcap E_1^c$. We let $p = 0.5$.

We generate 500 data under the each setting of the sample size and signal-to-noise ratio. The first column of Figure 3.2 show one instance of data with $\beta = 1$ and signal-to-noise ratio $\sigma = 30$. The second and third column show the `SRGonG` segmentation result without and with exposure map adjustment. We can see that with adjustment, 3 out of 4 point sources are correctly detected and the extended source is not segmented to multiple segments. We summarize the values of *Adjusted Rand Index* of the 500 replaicates in Figure 3.3. The improvement in performance with exposure

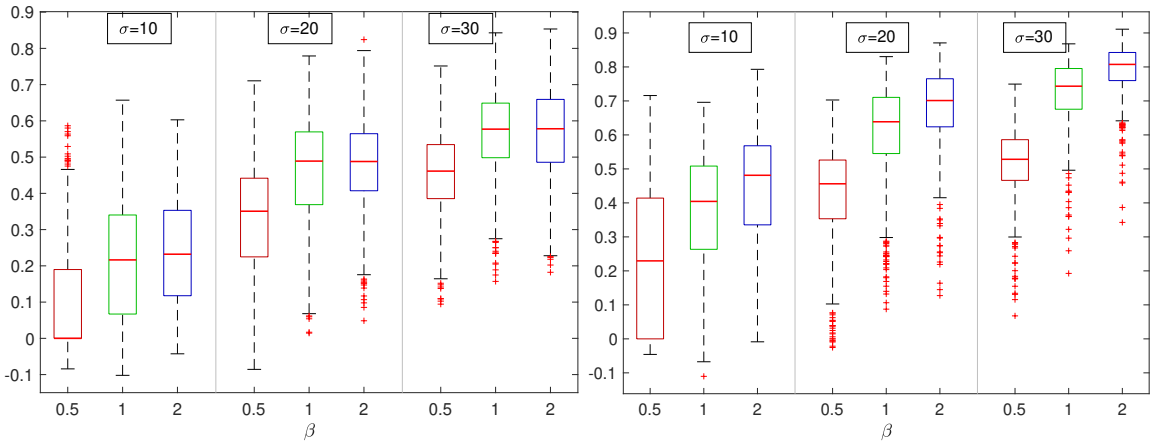


FIGURE 3.3. Simulation studies with different sample sizes and signal-to-noise raitos represented by parameters (σ, β). **Left Column**: Box plots of the *Adjusted Rand Index* for 500 replicates without exposure map adjustment. **Right Column**: Box plots of the *Adjusted Rand Index* for 500 replicates with exposure map adjustment.

map adjustment become more substantial as sample size increases for each signal-to-noise raito level.

CHAPTER 4

Uncertainty Quantification of Object Boundaries Extracted from Spatial Point Pattern Image

4.1. Introduction

Quantifying the uncertainty of a source boundary computed with `SRGonG` is a challenging task, owing to both the data-driven nature of `SRGonG` and the type of statistical noise (an inhomogeneous Poisson process). Indeed, we have not identified any existing uncertainty quantification method in the field of astrophysics for extended object boundary estimation when the input data consist of high-energy photon lists.

The quantification of uncertainty in medical image segmentation has garnered significant attention, especially at the pixel/voxel level (e.g., Hamu Goldberg *et al.*, 2018). These methods, however, typically rely on probabilistic atlases constructed from training sets of aligned images with expert annotations. Deep learning methods for probabilistic image segmentation have also been developed to encode prediction confidence and quantify ambiguity in binary segmentation problems (e.g., Gal and Ghahramani, 2015; Kohl *et al.*, 2018; Ding *et al.*, 2020). These methods model the pixel-wise class probability for a given image and visualize the uncertainty impact with probability maps or heat maps.

The problem addressed in this chapter is fundamentally different. `SRGonG` is not trained on a sample with expert annotations, and there is no ground-truth training set available. Thus, we must develop novel methods to quantify uncertainty in the `SRGonG` segmentation. Our methods focus on the case where a single extended source resides in the interior of the field of view. In practice, the source brightness and the intensity of the surrounding field vary continuously, leading `SRGonG` to produce a segmentation composed of more than two regions. Thus, the first step of our methods is to produce a single boundary estimate of the extended source. We propose `BoSE` that separates a

single boundary from `SRGonG` result and extract that boundary from segmentation. Our objective is to quantify the uncertainty of `BoSE` method.

To quantify the uncertainty of `BoSE`, the first step is to parameterize the boundary using Fourier Descriptors (Burger and Burge, 2016) and choose the number of parameters with BIC. With the parameterization in hand, we are able to develop inferential tools for computing both parameter estimates (endowed with geometric interpretations) and a global confidence region that effectively captures the uncertainty in the shape of the extended source. Using the Fourier Descriptors, we can further characterize the source in terms of its location, area, and eccentricity and quantify the uncertainty of these characterizations.

Our uncertainty quantification relies on resampling techniques (Efron, 1979). The nature of the photon list, however, prohibits the direct application of the standard bootstrap. Choosing events with replacement from the event list invariably produces multiple identical events, and thus any such resampled dataset is qualitatively different from the observed data. Thus, we develop a novel non-parametric spatial point bootstrap method that is applicable to event lists.

The chapter is organized as the following way. Section 4.2 reviews the Fourier Descriptor parametrization, and Section 4.3 introduces the main novel methodology, `UQ-BoSE`, of the chapter: a parametric estimate of the boundary of a spatially extended source and an asymptotically minimal global confidence region for this boundary. We validate our proposed methods with a simulation study using synthetic data and with experiments involving real data in Sections 4.4 and 4.5.1, respectively. Finally, in Section 4.6, we discuss practical issues concerning the application of our methods and discuss future extensions.

4.2. Fourier Descriptors

4.2.1. Review of Coordinate-Based Fourier Descriptors. In this section we review how a generic segment boundary, C, can be parameterized via Fourer descriptors (FD, Burger and Burge, 2009).

Let $V_L(t)$ denote the following matrix,

$$V_L(t)^T = \begin{bmatrix} a(t) & c(t) \\ b(t) & d(t) \end{bmatrix} \tag{4.1}$$

where the elements of the matrix $V_J(t)$ are themselves row matrices of the form

$$a(t) = [\cos(-2\pi(\lfloor (L-1)/2 \rfloor t) \ \ldots \ \cos(-2\pi t) \ 1 \ \cos(2\pi t) \ \ldots \ \cos(2\pi \lfloor L/2 \rfloor t)] \ \in \mathbb{R}^L$$

$$b(t) = [\sin(-2\pi(\lfloor (L-1)/2 \rfloor t) \ \ldots \ \sin(-2\pi t) \ 1 \ \sin(2\pi t) \ \ldots \ \sin(2\pi \lfloor L/2 \rfloor t)] \ \in \mathbb{R}^L$$

$$c(t) = [-\sin(-2\pi(\lfloor (L-1)/2 \rfloor t) \ \ldots \ -\sin(-2\pi t) \ -1 \ -\sin(4\pi t) \ \ldots \ -\sin(2\pi \lfloor L/2 \rfloor t)] \ \in \mathbb{R}^{L-1}$$

$$d(t) = [\cos(-2\pi(\lfloor (L-1)/2 \rfloor t) \ \ldots \ \cos(-2\pi t) \ 1 \ \cos(4\pi t) \ \ldots \ \cos(2\pi \lfloor L/2 \rfloor t)] \ \in \mathbb{R}^{L-1}$$

Then C can be expressed as function $s : \mathbb{R} \to \mathbb{R}^2$,

$$s(t; \boldsymbol{\theta}, L) = \begin{bmatrix} s_x(t; \boldsymbol{\theta}, L) \\ s_y(t; \boldsymbol{\theta}, L) \end{bmatrix} = V_L(t)^T \boldsymbol{\theta}, \quad \text{for } t \in [0, 1], \tag{4.2}$$

where $\boldsymbol{\theta}$ are the parameters of the FD and $V_L(t)^T \boldsymbol{\theta}$ is a matrix representation of the parameterization. To derive $s(t; \boldsymbol{\theta}, L)$, we apply the Discrete Fourier Transformation (DFT) to C; specifically we select L evenly spaced points along C, resulting in a 2D sequence of coordinates $\{(x_l, y_l), l = 0 \ldots L-1\}$, where L can be arbitrarily large. The Fourier spectrum is the complex-valued vector, $\boldsymbol{G} = (G_0, \ldots, G_{L-1})$, of length L derived by applying the DFT to the set of complex signals $\boldsymbol{z} = (z_0, \ldots, z_{L-1})$, where $z_l = x_l + i \cdot y_l$, i.e.,

$$G_j = \frac{1}{L} \sum_{l=0}^{L-1} z_l \cdot \exp\left\{-i 2\pi j \frac{l}{L}\right\}, \text{ for } j = 0, \ldots, L-1. \tag{4.3}$$

Each G_j denotes a FD with frequency j. Because $\boldsymbol{G}$ is periodic, we have $G_j = G_{j \bmod L}$ for $j < 0$.

The FD parametrization of C is a smooth interpolating curve, $s(t; \boldsymbol{\theta}, L)$ given by[1],

$$\begin{aligned} s_x(t; \boldsymbol{\theta}, L) &= \sum_{j=-\lfloor (L-1)/2 \rfloor}^{\lfloor L/2 \rfloor} [\mathrm{Re}(G_j) \cdot \cos(2\pi j \cdot t) - \mathrm{Im}(G_j) \cdot \sin(2\pi j \cdot t)], \\ s_y(t; \boldsymbol{\theta}, L) &= \sum_{j=-(\lfloor L-1)/2 \rfloor}^{\lfloor L/2 \rfloor} [\mathrm{Im}(G_j) \cdot \cos(2\pi j \cdot t) + \mathrm{Re}(G_j) \cdot \sin(2\pi j \cdot t)], \end{aligned} \tag{4.4}$$

where $\boldsymbol{\theta} = (\mathrm{Re}(G_{-\lfloor (L-1)/2 \rfloor}), \ldots, \mathrm{Re}(G_{\lfloor L/2 \rfloor}), \mathrm{Im}(G_{-\lfloor (L-1)/2 \rfloor}), \ldots, \mathrm{Im}(G_{\lfloor L/2 \rfloor}))$.

[1] $\lfloor x/y \rfloor$ denotes the floor of x/y, i.e., the largest integer smaller than x/y.

The Inverse Fourier transformation (IFT) gurantees that $\big(s_x(l/L;\boldsymbol{\theta}), s_y(l/L;\boldsymbol{\theta})\big) = (x_l, y_l)$ for each l and thus gurantees that $s(t;\boldsymbol{\theta}, L)$ becomes arbitaryly close to C as L grows large. The high-frequency terms in Equation 4.4 capture local noise-driven flucuations and thus we typically truncate them and include only terms with $j = -\lfloor (J-1)/2 \rfloor, \ldots, \lfloor J/2 \rfloor$, for some J such that $2 \leq J < L$, i.e., we replace $s(t;\boldsymbol{\theta}, L)$ with $s(t;\boldsymbol{\theta}, J)$.

Geometrically, each pair of summands in Eq 4.4 with frequency $\pm j$ represents an ellipse with radii $|G_{-j}|+|G_j|$ and $\big||G_{-j}|-|G_j|\big|$, and orientation angle given by $\frac{1}{2}[\tan^{-1}(\frac{\mathrm{Im}(G_{-j})}{\mathrm{Re}(G_{-j})})+\tan^{-1}(\frac{\mathrm{Im}(G_j)}{\mathrm{Re}(G_j)})]$. The tuple of $\langle G_{-j}, G_j \rangle$ is referred as "Fourier descriptor pair" (FDP) with frequency j, with the low-frequency FDP describing the global information and the high-frequency FDP describing the local information of C. In particular, we use the FDP with frequency 1 to quantify shape's size, eccentricity and orientation.

4.2.2. Overparametrization. The FD generally depend on the choice of (x_0, y_0), as the choice of this value effects the other $L-1$ evenly space points used to compute the descriptors. The "shift property" of the DFT means that shifting the signal $\boldsymbol{g}$ by $-k_s$ position modifies the coefficients from G_j to $\tilde{G}_j = e^{ij\frac{2\pi k_s}{L}} G_j$. Hence, the value of $\boldsymbol{\theta}$ in Eq 4.4 is not unique and there is ambiguity in the parametrization of the curve defining the segment. Motivated by Ye *et al.* (2000), we constrain the FDs that normalize G_j to $\tilde{G}_j$ for each j. Specifically, we force $\tilde{G}_1 = |G_1|$, which is equivalent to shifting the phase of G_1 by $-\psi_1$, where

$$\psi_1 = \begin{cases} \tan^{-1}\frac{\mathrm{Im}(G_1)}{\mathrm{Re}(G_1)} & \text{if } \mathrm{Im}(G_1) \geq 0 \\ \left(\pi + \tan^{-1}\frac{\mathrm{Im}(G_1)}{\mathrm{Re}(G_1)}\right) & \text{if } \mathrm{Im}(G_1) < 0 \end{cases}. \tag{4.5}$$

It follows from the shift property of DFT that

$$\tilde{G}_j = \begin{cases} e^{-ij\cdot\psi_1} \cdot G_j, & \text{for } 0 \leq j \leq \lfloor L/2 \rfloor, \\ e^{ij\cdot\psi_1} \cdot G_j, & \text{for } \lfloor L/2 \rfloor < j \leq L-1. \end{cases} \tag{4.6}$$

In the following sections, we assume that $\boldsymbol{G}$ is normalized in this way.

4.3. Statistical Inference

We turn now to the primary contribution of this chapter: the extraction of Boundary of Segmented Events (`BoSE`) and the quantification of the uncertainty of the boundary (`UQ-BoSE`).

4.3.1. `BoSE` : The Merge-to-Two Step and Boundary Extraction. To isolate the boundary of a single extended source within a `SRGonG` segmentation, $\widehat{\mathbf{S}}(\widehat{K})$, we propose to merge segments to form a reduced segmentation of predetermined size, $\widetilde{K} < \widehat{K}$. To accomplish this, we use a "Forced-Merge" operation that applies a Beam Merger Step to $\widehat{\mathbf{S}}(\widehat{K})$ and arrives at the reduced segmentation, $\widetilde{\mathbf{S}}(\widetilde{K})$, by maximizing the profile log-likelihood over all potential subsequence segmentations with $\widetilde{K}$ segments. In this article, we are primarily focused on the case where $\widetilde{K} = 2$, and thus refer to the "Merge-to-Two" Step. We apply a Merge-to-Two Step to obtain $\widetilde{\mathbf{S}} = (\widetilde{\mathcal{S}}_1, \widetilde{\mathcal{S}}_2)$. Let $\widetilde{\mathcal{S}}_1$ be the central segment; we assume that $\widetilde{\mathcal{S}}_1$ has a single boundary.

The boundaries of a segmentation are extracted using the Voronoi diagram. If two photons located in different segments are adjacent in the Delaunay triangulation, the edge shared by their Voronoi cells is taken as part of their common boundary. The boundary of each segment is the polygon constructed by connecting its set of edges.

4.3.2. Estimate of the Object Boundary. In this section we use segmentation and parametrization tools to obtain a parametic estimate of an object's spatial boundary, which is essentially the first step of `UQ-BoSE` . We assume that the true boundary, $s^\star$, takes the same parameterized functional form as in Equation 4.4 and can be written $s^\star = s(t; \boldsymbol{\theta}^\star, J^\star)$, where $\boldsymbol{\theta}^\star$ and $J^\star$ are the "true" values of $\boldsymbol{\theta}$ and J. To be sure that we can estimate the complete boundary, we assume that $s(t; \boldsymbol{\theta}^\star, J^\star)$ is fully contained in $\mathcal{F}$. In order to compute the estimate, we first apply `SRGonG` to obtain the full segmentation $\widehat{\mathbf{S}} = (\widehat{\mathcal{S}}_1, \ldots, \widehat{\mathcal{S}}_K)$ and then apply `BoSE` to obtrain the two-region segmentation $\widetilde{\mathbf{S}} = (\widetilde{\mathcal{S}}_1, \widetilde{\mathcal{S}}_2)$, as well as the single boundary of $\widetilde{\mathcal{S}}_1$, denoted C_1. We further assume that $\widetilde{\mathcal{S}}_1$ shares this boundary with $\widetilde{\mathcal{S}}_2$ and that $\widetilde{\mathcal{S}}_2$ completely surrounds $\widetilde{\mathcal{S}}_1$.

Our estimate of the boundary is the following

$$\hat{s}(J) = \begin{bmatrix} s_x(t; \hat{\boldsymbol{\theta}}, J) \\ s_y(t; \hat{\boldsymbol{\theta}}, J) \end{bmatrix} = V_J(t)^T \hat{\boldsymbol{\theta}}, \quad \text{for} \quad t \in [0, 1], \tag{4.7}$$

where $\hat{\boldsymbol{\theta}} \in \mathbb{R}^{2J-1}$ represents the normalized FD parameters computed with a set of L points, $\{(x_l, y_l)\}$, evenly spaced along C_1, i.e., $\hat{\boldsymbol{\theta}}$ consists of the real and imaginary parts of the J lowest-frequency terms in $\boldsymbol{G}$,

(4.8)

$$\hat{\boldsymbol{\theta}} = (\mathrm{Re}(G_{-\lfloor (J-1)/2 \rfloor}), \ldots, \mathrm{Re}(G_{\lfloor J/2 \rfloor}), \mathrm{Im}(G_{-\lfloor (J-1)/2 \rfloor}), \ldots, \mathrm{Im}(G_0), \mathrm{Im}(G_2), \ldots, \mathrm{Im}(G_{\lfloor J/2 \rfloor}))^T \in \mathbb{R}^{2J-1},$$

where $\boldsymbol{G}$ is the Fourier spectrum of C_1 of length L. Because $\hat{\boldsymbol{\theta}}$ is determined by the FD parameters of C_1, once C_1 is determined, the estimate of the boundary depends only on the truncation cut off, J. Thus, we denote the estimate $\hat{s}(J)$ to emphasize its dependence on J. In the conventional Fourier parametrization, FDs are paired (Burger and Burge, 2016; Lin and Chellappa, 1987), i.e., the complete tuple $\langle G_{-j}, G_j \rangle$ is included in Equation 4.4. Following this convention, we assume $J^\star$ is an odd interger ≥ 3. If the user does not wish to pre-determine the value of J (e.g., based on the complexity of the boundary[2]), J can be estimated, as described in Section 4.3.3.

4.3.3. Estimating J via the Bayesian Information Criterion. We estimate J by maximizing its profile log-likelihood. Because the likelihoood depends on J through the segmentation it induces, we let $\mathcal{S}_{\hat{s}(J)}$ denote the interior region of $\hat{s}(J)$ and let $\mathbf{S}_{\hat{s}(J)} = (\mathcal{S}_{\hat{s}(J)}, \mathcal{F} \backslash \mathcal{S}_{\hat{s}(J)})$[3] be the partition of $\mathcal{F}$ into the interior and exterior of $\hat{s}(J)$. The profile log-likelihood of J can be expressed as $\mathcal{L}_{\text{profile}}(K = 2, \mathbf{S}_{\hat{s}(J)} | \mathbf{X}, n)$ as specified in Equation 3.5. We formlate the log-likelihood with $K = 2$ because we cannot substitute the parameterized boundary, $\hat{s}(J)$ for the `SRGonG` boundary, C_1 within a segmentaion with $K > 2$. This is becuase there would be discontinuities between an `SRGonG` segmentation and $\hat{s}(J)$, see Figure 4.1. Without a consistent segmentation, the log-likelihood is not defined. Thus, we evaluate the log-likelihood with $K = 2$. To accommodate the ensueing assumption that the Poission intensities over both $\mathcal{S}_{\hat{s}(J)}$ and $\mathcal{F} \backslash \mathcal{S}_{\hat{s}(J)}$ are constant, we compute $\mathcal{L}_{\text{profile}}$ using only photons recorded within a subset of $\mathcal{F}$ where the following holds.

Enclosure Assumption: There exists a subset $\mathcal{R} \subseteq \mathcal{F}$ that encloses $s^\star$, where photon arrivals can be appropriately modeled as an inhomogeneous Poisson process with an intensity function that is constant over both $\mathcal{S}^{\mathcal{R}}_{s^\star} = \mathcal{S}_{s^\star} \cap \mathcal{R}$ and $\mathcal{R} \backslash \mathcal{S}^{\mathcal{R}}_{s^\star} = (\mathcal{F} \backslash \mathcal{S}_{s^\star}) \cap \mathcal{R}$. By "encloses", we mean for all t, $s^\star = s(t; \boldsymbol{\theta}^\star, J^\star)$ is in the interior of $\mathcal{R}$.

[2]When $J = 3$, $s(t; \boldsymbol{\theta}, J)$ always describes an ellipse.

[3]$A \backslash B = \{x | x \in A \text{ and } x \notin B\}$.

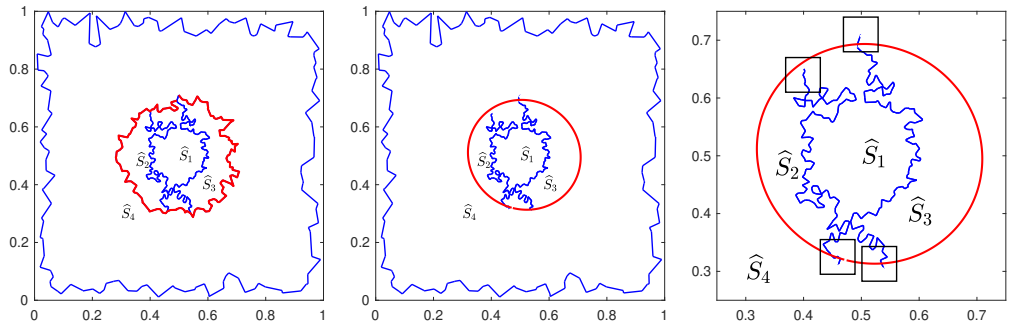


FIGURE 4.1. The potential issue of estimating J using a full segmentation, $\widehat{\mathbf{S}}$, with $K > 2$. This segmeationa is taken from a simulation study detailed in Section 4.4.1. **Left:** The blue and red lines represent the boundaries of segments in $\widehat{\mathbf{S}}$. The boundary of interest, C_1, was extracted from $\widetilde{\mathbf{S}}$ and is highlighted in red. **Middle:** Replacing the C_1 with $\hat{s}(J)$ within $\widehat{\mathbf{S}}$ (with $J = 3$); **Right:** Detail of the middle panel. The replacement of C_1 with $\hat{s}(J)$ causes multiple discontinuities in the boundaries of$\widehat{\mathcal{S}}_1$, $\widehat{\mathcal{S}}_2$, and $\widehat{\mathcal{S}}_3$. These are highlighted with black boxes. These discontinuties mean that the middle panel is not semgmentation and its log-likehood function is undefined.

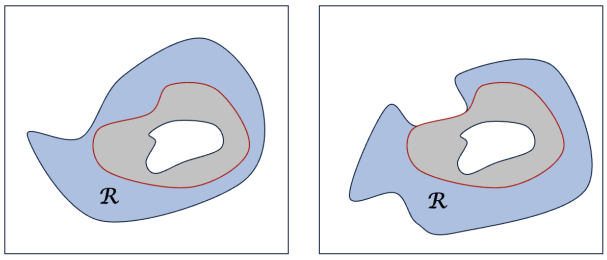


FIGURE 4.2. Illustration of *Enclosure Assumption*. The red curve is the object boundary $s(t; \boldsymbol{\theta}^\star, J^\star)$. The intensities over both blue and grey segments are assumed to be constant. Left Panel: A subset $\mathcal{R}$ that satisfies the Enclosure Assumption. Right Panel: A subset $\mathcal{R}$ that violates the Enclosure Assumption.

Figure 4.2 illustrates two segmentations, one following the Enclosure Assumption and one violating it. We discuss the pitfulls of estimating J when the Enclosure Assumption is violated in Appendix B.1. Because $s^\star$ is unknown, we approximate $\mathcal{R}$ with the minimal union of segments in $\widehat{\mathbf{S}}$, the full `SRGonG` segmentation, that fully encloses C_1, i.e.,

$$\mathcal{R} \approx \min \bigcup_k \widehat{\mathcal{S}}_k \text{ such that } C_1 \subset \text{interior}\left(\bigcup_k \widehat{\mathcal{S}}_k\right). \tag{4.9}$$

This is equivalent to the union of all $\widehat{\mathcal{S}}_k$ that share a boundary with C_1.

Under the Enclosure Asumption, we compute the profile log-likelihood of J using only the $n_{\mathcal{R}}$ photons observed in $\mathcal{R}$, i.e.,

$$\begin{aligned}\mathcal{L}^*_{\text{profile}}(J|\mathbf{X}_{\mathcal{R}}, n_{\mathcal{R}}) =& \text{Num}\left(\mathcal{S}^{\mathcal{R}}_{\hat{s}(J)}\right) \log\left(\frac{\text{Num}\left(\mathcal{S}^{\mathcal{R}}_{\hat{s}(J)}\right)}{\text{Area}\left(\mathcal{S}^{\mathcal{R}}_{\hat{s}(J)}\right)}\right) \\ &+ \text{Num}\left(\mathcal{R}\backslash\mathcal{S}^{\mathcal{R}}_{\hat{s}(J)}\right) \log\left(\frac{\text{Num}\left(\mathcal{R}\backslash\mathcal{S}^{\mathcal{R}}_{\hat{s}(J)}\right)}{\text{Area}\left(\mathcal{R}\backslash\mathcal{S}^{\mathcal{R}}_{\hat{s}(J)}\right)}\right) - n_{\mathcal{R}} - \log n_{\mathcal{R}}!,\end{aligned} \tag{4.10}$$

where $\mathbf{X}_{\mathcal{R}} = \{\mathbf{x}_i : \mathbf{x}_i \in \mathcal{R}\}$. Because $\mathcal{L}^*_{\text{profile}}(J|\mathbf{X}_{\mathcal{R}}, n_{\mathcal{R}})$ is an increasing function of J, simplying maximizing $\mathcal{L}_{\text{profile}}$ leads to overfitting of $s(t;\boldsymbol{\theta}, J)$. To avoid this, we introduce a BIC term to penalize model complexity. In particular, we use

$$\text{BIC}(J) = -2\mathcal{L}^*_{\text{profile}}(J|\mathbf{X}_{\mathcal{R}}, n_{\mathcal{R}}) + m_{\text{pseg}}(J)\log n_{\mathcal{R}}, \tag{4.11}$$

where $m_{\text{pseg}}(J)$ denotes the number of free parameters as a function of J, i.e., $m_{\text{pseg}}(J) = 2J + 1$, where $2J - 1$ corresponds to the free parameters in $\boldsymbol{\theta}$ given in Equation 4.8, and the remaining 2 account for the intensity parameters of $\mathcal{S}^{\mathcal{R}}_{\hat{s}(J)}$ and $\mathcal{R}\backslash\mathcal{S}^{\mathcal{R}}_{\hat{s}(J)}$. Consequently, the BIC estimate of J is

$$\hat{J}_{\text{BIC}} = \underset{J\in[3,L],\, \textit{with } J \textit{ odd}}{\arg\min} \text{BIC}(J). \tag{4.12}$$

BIC is known for its statistical consistency in model selection problems. We illustrate the performance of the estimates in Section 4.4.2.

The profile log-likelihood given in Equation 4.10 is closely related to that in Equation 3.5. Both encode the same Poisson model with its piecewise constant intensity function, with Equation 4.10 representing a 'local fit' in that it removes events outside $\mathcal{R}$. We use Equations 3.4 and 3.5 to estimate the number of segments K, while setting the value of m_{seg} to represent the number of parameters required to model each segment. In Equations 4.10 and 4.11, on the other hand, we fix $K = 2$ when fitting J to control the complexity of the boundary. The two BIC penalties in Equations 3.4 and 4.11 are the same up to how the number of parameters are counted in Km_{seg} and $m_{\text{pseg}}(J)$, respectively.

4.3.4. Construct the Global Confidence Region Through Bootstrapping. In this section we introduce the second step of `UQ-BoSE` : the construction of an asymptotic global confidence region for the true object boundary $s^\star$. In order to construct such global confidence region, we assume

$$\hat{\boldsymbol{\theta}} \overset{\text{asy}}{\sim} N(\boldsymbol{\theta}^\star, \Sigma_{\boldsymbol{\theta}}), \tag{4.13}$$

and thus from Equation (4.7) that

$$s(t; \hat{\boldsymbol{\theta}}, J) \overset{\text{asy}}{\sim} N\big(V_J(t)^T \boldsymbol{\theta}^\star, \Sigma_s(t)\big), \tag{4.14}$$

where $\Sigma_s(t) = V_J(t)^T \Sigma_{\boldsymbol{\theta}} V_J(t)$, and $\overset{\text{asy}}{\sim}$ indicates an asymptotic distribution. We follow the method proposed in Ye *et al.* (2000) to construct small-sized global confidence region for $s(t; \boldsymbol{\theta}^\star, J^\star)$. For a fixed $t \in [0,1]$, the smallest elliptical confidence region, $\hat{U}_\gamma(t)$, for $s(t; \boldsymbol{\theta}^\star, J^\star)$ with asymptotitc confidence level $\alpha \in [0,1]$, i.e,

$$\Pr\Big(s(t; \boldsymbol{\theta}^\star, J^\star) \in \hat{U}_\gamma(t)\Big) = \alpha \tag{4.15}$$

under the asymptotic distribution of $s(t; \hat{\boldsymbol{\theta}}, J)$ is

$$\hat{U}_\gamma(t) = \Big\{x \in \mathbb{R}^2 : \big(x - s(t; \hat{\boldsymbol{\theta}}, J)\big)^T \Sigma_s(t)^{-1} \big(x - s(t; \hat{\boldsymbol{\theta}}, J)\big) \ \leq \gamma^2\Big\}, \tag{4.16}$$

where γ^2 is the α quantile of a χ^2_2 distribution with two degrees of freedom[4] and $J = J^\star$. The global confidence region for the entire curve, $\hat{U}_\gamma$, is the union of the set of elliptical confidence regions for each t, i.e.,

$$\hat{U}_\gamma \overset{\Delta}{=} \bigcup_{t \in [0,1]} \hat{U}_\gamma(t). \tag{4.17}$$

We must address two differences between our scenario and that of Ye *et al.* (2000). First, the value of $J^\star$ is predetermined in Ye *et al.* (2000) but not here. Second, Ye *et al.* (2000) is able to estimate $\Sigma_{\boldsymbol{\theta}}$ from multiple observations while we have only one observed dataset. We address the

[4]I.e., $\gamma^2 = \chi^2_2(\alpha)$, where $\chi^2_2(\cdot)$ is the inverse cumulative distribution function of a chi-square random variable with 2 degrees of freedom.

Algorithm 3 Construct asymptotic minimum global confidence region for $s^\star$

Require: Coordinates of observed photons $\mathbf{X}$ of size n; γ that corresponds to confidence level of coverage.

Ensure: $\hat{U}_\gamma$, global confidence region of extended object boundary $s(t;\boldsymbol{\theta}^\star, J^\star), t = [0,1]$.

1: Derive $\widehat{\mathbf{S}}$ (full segments) and $\widetilde{\mathbf{S}}$ (two segments) by apply segmentation algorithm to input data $\mathbf{X}$. Extract curve C_1 from $\widetilde{\mathbf{S}}$.
2: Construct estimator $\hat{\boldsymbol{\theta}}$ with length $2\hat{J}-1$ by applying DFT to C_1 and a model selection method that estimates J.
3: **for** $N_b = 1$ to N_B **do**
4: Genetate bootstrap data $\mathbf{X}^{(N_b)}$ of size n with suitable bootstrapping procedure.
5: Derive $\widetilde{\mathbf{S}}^{(N_b)}$ by segmenting $\mathbf{X}^{(N_b)}$ with the Merge-to-Two Step. Extract ${C_1}^{(N_b)}$ from $\mathbf{S}^{(N_b)}$. Derive $\boldsymbol{\theta}^{(N_b)}$ with length $2\hat{J}-1$ by applying DFT to $C^{(N_b)}$.
6: Estimate $\widetilde{\Sigma}_{\boldsymbol{\theta}}$ from $\boldsymbol{\theta}^{(N_b)}, N_b = 1,\ldots,B$. Construct $\hat{U}_\gamma(t)$ over $t \in [0,1]$.
7: **Return**$\hat{U}_\gamma \triangleq \bigcup_{t\in[0,1]} \hat{U}_\gamma(t)$.

first issue by estimating $\hat{J}$ using the model selection procedure described in Section 4.3.3 and the second via a bootstrap method.

To esimate $\Sigma_{\boldsymbol{\theta}}$ via the bootstrap, we first generate N_B replicate datasets, $\mathbf{X}^{(N_b)}$, for $N_b = 1,\ldots,N_B$. Our procedure for doing this is described in Section 4.3.6. With the replicate datasets in hand, we apply `BoSE` to each dataset and derive $\widetilde{\mathbf{S}}^{(N_b)}$ with two segments. Let $\boldsymbol{G}^{(N_b)}$ denote the Fourier spectrum of segment boundary ${C_1}^{(N_b)}$ extracted from $\widetilde{\mathbf{S}}^{(N_b)}$. Fixing J at $\hat{J}$, $\boldsymbol{\theta}^{(N_b)}$ consists of the $2\hat{J}-1$ lowest-frequency components in $\boldsymbol{G}^{(N_b)}$.We estimate $\Sigma_{\boldsymbol{\theta}}$ with

$$\widetilde{\Sigma}_{\boldsymbol{\theta}} = \frac{1}{N_B - 1}\sum_{N_b=1}^{N_B}(\boldsymbol{\theta}^{(N_b)} - \overline{\boldsymbol{\theta}})(\boldsymbol{\theta}^{(N_b)} - \overline{\boldsymbol{\theta}})^T \quad \text{with} \quad \overline{\boldsymbol{\theta}} = \frac{1}{N_B}\sum_{N_b=1}^{N_B}\boldsymbol{\theta}^{(N_b)} \in \mathbb{R}^{2\hat{J}-1}. \tag{4.18}$$

The confidence region for $s^\star$ is constructed by replacing Σ_θ in Equation 4.16 with $\widetilde{\Sigma}_\theta$. The procedure is summarized in Algorithm 3 and we evaluate its performance through numerical experiments with extended sources of various shapes in Section 4.4.2. Furthermore, we use Monte-Carlo data to estimate the true crossing-level probability and compare it with the theoretical bounds proposed in Ye *et al.* (2000).

4.3.5. Lower Bounds for the Global Confidence Level of FD-based Confidence Region. In this section, we delve into the examination of the global confidence level of FD-based Confidence Region constructed using Algorithm 3, i.e., $\Pr(s(t;\boldsymbol{\theta}^\star, J^\star) \in \hat{U}_\gamma, \forall t)$. We introduce two lower bounds for $\Pr(s(t;\boldsymbol{\theta}, J^\star) \in \hat{U}_\gamma, \forall t)$ and their empirical estimations when the exact value is

not directly accessible due to the absence of the true covariance $\Sigma_{\boldsymbol{\theta}}$. Additionally, we conduct a comparative analysis of the empirical global confidence level of $\hat{U}_\gamma$ across various selections of γ within the framework of three simulation scenarios, as outlined in Section 4.4.2. To evaluate the performance of the procedure in Section 4.3.4, we simulate Monte-Carlo data to estimate the true crossing level probability and compare it with the empirical theoretical bounds. We generate synthetic data with various shapes of extended source $s(t;\boldsymbol{\theta}^\star, J^\star)$, whose complexity is measured by $J^\star$. For a fixed γ, we construct a global confidence region $\hat{U}_\gamma^{(N_m)}$ for each of the N_M simulated data and estimate the true level-crossing probability, $1-\Pr(s(t;\boldsymbol{\theta}^\star, J^\star) \in \hat{U}_\gamma,\ \forall t)$ by empirical level-crossing probability,

$$P_e = \frac{1}{N_M}\sum_{N_m=1}^{N_M}(1-\mathbb{1}\{s(t;\boldsymbol{\theta}^\star, J^\star) \in \hat{U}_\gamma^{(N_m)},\ \forall t\}). \tag{4.19}$$

In the construction of $\hat{U}$, we consider two senarios, where J is true $J^\star$ and J is esitmated with the BIC. We denote the empirical level-crossing probability as P_e and P_e^{BIC} when $J = J^\star$ and $J = \hat{J}_{\text{BIC}}$, respectively. We compare the P_e and P_e^{BIC} with the gamma bound and the estimated level-crossing bound.

4.3.5.1. *Empirical Estimation of the Bounds.* The exact value of $\Pr(s(t;\boldsymbol{\theta}^\star, J^\star) \in \hat{U}_\gamma, \forall t)$ is not accessible, but following the analysis in Ye *et al.* (2000), we can derive two lower bounds, incomplete gamma bound and level-crossing bound, for the probability when $J = J^\star$ in the construction of $\hat{U}$. The incomplete gamma bound is written as [5]:

$$\Pr(s(t;\boldsymbol{\theta}^\star, J^\star) \in \hat{U}_\gamma, \forall t \in [0,1]) \geq \Gamma(\frac{2J^\star-1}{2}, \frac{\gamma^2}{2}). \tag{4.20}$$

The level-crossing bound is written as:

$$\Pr(s(t;\boldsymbol{\theta}^\star, J^\star) \in \hat{U}_\gamma, \forall t \in [0,1]) \geq 1-(1+\kappa_o\gamma)\exp(-\frac{\gamma^2}{2}), \tag{4.21}$$

where κ_o is given by

$$\kappa_o \triangleq \int_0^1 (\frac{\zeta(t)}{2\pi})^{1/2} dt, \tag{4.22}$$

[5] $\Gamma(\cdot,\cdot)$ denotes the regularized gamma function.

and suppose $\Sigma_{\boldsymbol{\theta}}^{1/2} V_{J^\star}(t)$ has full column rank and let $A = \Sigma_{\boldsymbol{\theta}}^{1/2} V_{J^\star}(t) V_{J^\star}(t)^T \Sigma_{\boldsymbol{\theta}}^{1/2}$,

$$\zeta(t) = \text{trace}(P_A^{\perp} \Sigma_{\boldsymbol{\theta}}^{1/2} \dot{V}_{J^\star} \Sigma_s^{-1} \dot{V}_{J^\star} \Sigma_{\boldsymbol{\theta}}^{1/2}), \tag{4.23}$$

where $\dot{V}_{J^\star}$ denotes the derivative of $V_{J^\star}(t)$ with respect to t, and $P_A^{\perp}$ denotes the projection on the orthogonal complement of column span of A. To approximate the level-crossing bound, we estimate the $\Sigma_{\boldsymbol{\theta}}$ in eqn. 4.23 with N_M Monte-Carlo simulated data, i.e.,

$$\hat{\Sigma}_{\boldsymbol{\theta}} = \frac{1}{N_M} \sum_{m=1}^{N_M} (\hat{\boldsymbol{\theta}}^{(N_m)} - \overline{\boldsymbol{\theta}})(\hat{\boldsymbol{\theta}}^{(N_m)} - \overline{\boldsymbol{\theta}})^T, \tag{4.24}$$

where $\overline{\boldsymbol{\theta}} = \frac{1}{N_M} \sum_{N_m=1}^{N_M} \hat{\boldsymbol{\theta}}^{(N_m)}$, and $\hat{\boldsymbol{\theta}}^{(N_m)} \in \mathbb{R}^{2J^\star - 1}$ denote the truncated FD parameters of segment's boundary of one set of N_M simulated observations. Let $\hat{\kappa}_0$ denote the estimated κ_0 in 4.22, we have the following two upper bounds for the level-crossing probability P_e:

$$P_{LC} = (1 + \hat{\kappa}_o \gamma) \exp(-\frac{\gamma^2}{2}), \tag{4.25}$$

$$P_G = 1 - \Gamma(\frac{2J^\star - 1}{2}, \frac{\gamma^2}{2}) \tag{4.26}$$

Furthermore, a useful comparison with P_{LC} and P_G is given by:

$$P_l = \exp(-\frac{\gamma^2}{2}). \tag{4.27}$$

According to Ye *et al.* (2000), P_l could be an approximation to P_e if the pointwise probability that $s(t; \boldsymbol{\theta}^\star, J^\star)$ lies outside of $\hat{U}_\gamma(t)$ was approximately equal to P_e.

4.3.6. Bootstrapping. There are numerous non-parametric bootstrapping techniques to reconstruct spatial point patterns. Some widely used procedures include tiling (e.g., Hall, 1985) as well as its many variants and extension (e.g., Loh and Stein, 2004). An adaptive reconstruction method was proposed by Tscheschel and Stoyan (2006) that iteratively updates the points based on certain summary characteristics of the data. All of these methods, however, assume the spatial data are isotropic or aim at estimating some statisitcs depending on local patterns in a homogeneous way, which does not apply in our case. Thus, we propose two bootstrap procedures, one parametric and one nonparametric, that work in different scenarios.

4.3.6.1. *Fourier-descriptors-based Parametric Bootstrap.* The parametric bootstrap utilizes the FDs and the piece-wise inhomogeneous Possion assumption. More specifically, we assume that the intensity over both $\mathcal{S}_{s^\star}$, the segment enclosed by $s^\star$, and $\mathcal{F}\backslash\mathcal{S}_{s^\star}$ are constant and estimate the underlying intensity function as

$$\hat{\lambda}(\mathbf{x}) = \frac{\text{Num}(\mathcal{S}_{\hat{s}(J)})}{\text{Area}(\mathcal{S}_{\hat{s}(J)})}\mathbb{1}\{\mathbf{x} \in \mathcal{S}_{\hat{s}(J)}\} + \frac{\text{Num}(\mathcal{F}\backslash\mathcal{S}_{\hat{s}(J)})}{\text{Area}(\mathcal{F}\backslash\mathcal{S}_{\hat{s}(J)})}\mathbb{1}\{\mathbf{x} \in \mathcal{F}\backslash\mathcal{S}_{\hat{s}(J)}\}. \tag{4.28}$$

Setting $\hat{\Lambda}$ equal to the integral of $\hat{\lambda}(\mathbf{x})$ over $\mathcal{F}$, i.e., $\hat{\Lambda} = \text{Num}(\mathcal{F})$, we obtain a bootstrap replicate, $\mathbf{X}^{(N_b)}$, by sampling from the Poisson process. I.e., we obtain an independent sample of size n from the density $\hat{\lambda}(\mathbf{x})/\hat{\Lambda}$.

4.3.6.2. *Nonparametric Bootstrap.* The parametric bootstrapping procedure generate piece-wise constant Possion data, which may not fully reflect the variance of the original Poisson process. This becomes particular problematic in real data that are observations of diffuse sources. Therefore, we propose an alternative nonparametric solution. We first sample n photons with replacement from the observed event list, $\mathbf{X}$, and then create the Voronoi diagram from the unique photons that were selected. If an observed photon was selected $k > 1$ times, we independently place k uniformly distriubted photons in its Voronoi cell.

4.4. Simulation Experiments

4.4.1. Simulation Study: The Effects of Contrast and S/N on `BoSE`. In this section, Monte-Carlo simulated data with various contrasts and signal-to-noise ratios (S/N) are generated to evaluate the performance of `BoSE`: We consider the scenario of a circular extended source surrounded by a halo. The expected photon count from a uniformly distributed background is denoted by μ_B and that of the circular source and halo are denoted by μ_{SC} and μ_{SH}, respectively. The field of view,$\mathcal{F}$, is assumed to be $[0,1] \times [0,1]$ and the areas of the circular extended source and the halo are fixed at 0.03 and 0.09. Let $\mu_S = \mu_{SC} + \mu_{SH}$ be the total expected source count. The S/N is defined as $\text{S/N} = \mu_S/\sqrt{\mu_S + 0.12\mu_B}$ and the contrast is defined by $c = 4\mu_{SC}/\mu_S$. Our experiments are organized into a $3 \times 3 \times 3$ design, where the first factor, μ_B, takes values 100, 500, and 2500, representing small, medium, and high background noise; the second factor, S/N, takes values 3, 10, and 30, representing a small, medium, and high S/N; and the third factor, c, takes

S/N	μ_B \ c	1.4	2	2.8
3	100	(6,11)	(8,8)	(11,5)
	500	(10,18)	(14,14)	(20,8)
	2500	(20,37)	(29,28)	(40,17)
10	100	(39,72)	(56,55)	(78,33)
	500	(50,92)	(71,71)	(99,43)
	2500	(81,149)	(115,115)	(161,69)
30	100	(319,593)	(456,456)	(638,274)
	500	(335,621)	(478,478)	(669,287)
	2500	(398,739)	(569,568)	(796,341)

TABLE 4.1. Expected source counts (μ_{SC}, μ_{SH}) of each simulation setting, specified in terms of S/N, μ_B, and c.

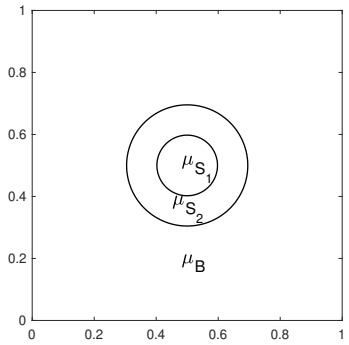


FIGURE 4.3. Illustrative figure of the simulations in Section 4.4.1. An extended circular shape of area 0.03 is surrounded by a halo of area 0.09.

values 1.4, 2 and 2.8, representing small, medium and high contrast. For each of the background and the two source objects, the photons are simulated from a Poisson process constructed with the given means, that is μ_B, μ_{SC} and μ_{SH} respectively. The total number of photons is fixed to be $\mu_B + \mu_{SC} + \mu_{SH}$, and the photon locations within each region are uniformly distributed. The values of (μ_{SC}, μ_{SH}) of each simulation setting appear in Table 4.1. An illustrative figure of the extended source can be found in Figure 4.3.

We generate $N_M = 500$ Monte-Carlo replicate datasets for each of the 27 simulation setting. We specify the initial seeds on regular grids. For $\mu_B = 100, 500$, and 2500, we use 7×7, 10×10 and 10×10 grids with seed size=1, 3, and 5, respectively. We apply `SRGonG` with a Random Merger Step to each simulated dataset, specifying $M_C = 3$ and $M_R = 1000$. In the Merge-to-Step step, we apply a Beam Merger Step with $M_B = 20$.

We summarize the performance of `BoSE` in estimating the source radius. When $J = 2$, the radius of the fitted boundary is $|G_1|$. Thus, Figure 4.4 examines the distribution of $|G_1|$ in each of the 27 simiulation settings. For some of these settings, such as low background brightness and low S/N, the BIC criteria frequently results in all subsegments being merged into one, i.e., $\widehat{K} = 1$ and there is no source in the fitted model. Even if $\widehat{K} > 2$, there may not be an identifiable source in the Merge-to-Two result, $\widetilde{\mathbf{S}}$, if there is no such $\widetilde{\mathcal{S}}_1$ whose boundary does not intersect with the

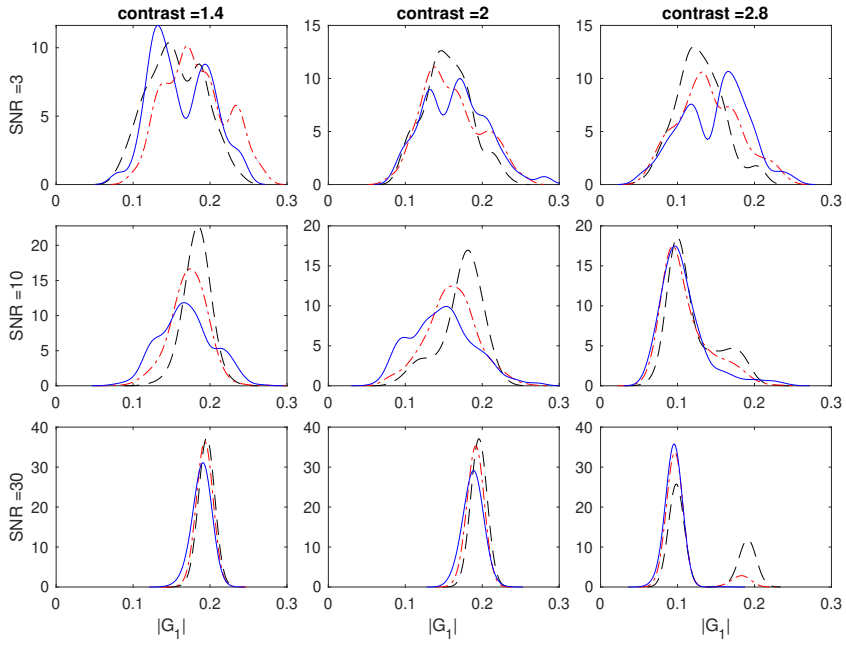


FIGURE 4.4. Simulation study for `BoSE` . The distribution of G_1 computed from the segmented object is plotted. The black dashed line, red dot-dash line and blue solid lines represent the frequency for $\mu_B = 100, 500$ and 2500, respectively.

boundary of $\mathcal{F}$ frame. When the source is identified (i.e., the boundary of $\widetilde{\mathcal{S}}_1$ does not intersect with $\mathcal{F}$ frame), we can extract its boundary, C_1. We then apply DFT to each extracted C_1 and summarize the distribution of $|G_1|$ in Figure 4.4. Under the low and medium contrast levels, we expect `SRGonG` to identify the inner circular extend source along with its halo. (These combine to form a circular shape with radius ≈ 0.2.) Under the high contrast setting, however, we expect `SRGonG` to only identify the circular extended source. (This is a a circular shape with radius $\approx$ 0.1.) Figure 4.4 illustrates that although $|G_1|$ fluctuates under small S/N, most of the time `SRGonG` identrifies a segement with $0.1 < |G_1| < 0.2$. The segmentation results are less ambiguous under the middle and high S/Ns levels.

4.4.2. Simulaiton Study: Performance of `UQ-BoSE`. We introduce three extended sources of varying complexity to illustrate the performance of `UQ-BoSE` including its FD boundary estimates, model selection techniques, and parametric bootstrap method for constucting a global confidence region. The shapes of the extended sources that we consider appear in the first column of Figure 4.5, specically, an ellipse, a bridge, and a spiral; these shapes can be constructed with 3, 7, and 11 Fourier

descriptors, respectively. Each extended source has an area of 0.2. For each source we simulate $N_M = 300$ replicate data sets each with 1300 photons. The photons are randomly split between source and background with probability 3/13 of being assinged to the extended source. Finally, the source and background photons are randomly distributed (uniformly) across their respective regions. This results in a surface brightness contrast between the source and the background of 1.5×.

We apply `BoSE` to each simulated dataset with $M_R = 1000$ and using a 9×9 regular seed grid, initialized with 5 seeds. To visualize the results, we randomly select one simulated data set from each of the three extended-source shape settings. The first column of Figure 4.5 shows these simulated datasets along with the corresponding ground truth. The second and third columns represent the `BoSE` segmentation results and the extracted segment boundary obtained using a Voronoi diagram as described in Section 3.2. `UQ-BoSE` results including the FD esitmates of the boundary obtained with the same three simulations appear in Figure 4.6. The first column plots the FD boundaries, $s(t; \hat{\boldsymbol{\theta}}, \hat{J}_{\text{BIC}})$, along with the simulated data (with $\hat{J}_{\text{BIC}}$ computed as in Equation 4.12), and the second and third columns illustrate the 39% and 68% and global confidence regions, respectively, all constructed with $N_B = 200$ resampled parametric bootstrap datasets. We choose these two confidence level a 39% corresponds $\gamma = 1$, which straightforwardly illustrate the width of the band, and 68% is standard confidence level that is commonly used in statistics tests.

Figure 4.7 summarizes the sampling distribution of $\hat{J}_{\text{BIC}}$. BIC excels in modeling simple shapes but in some cases tends to favor overly simplistic models. This can result in underfitting, especially in cases involving intricate shapes like the spiral example shown in the third row of Figure 4.5. We suggest using BIC when the underlying extended object is not expected to exhibit sharp corners. As object complexity increases, characterized by additional bulbs or sharp corners, the FD parametrization requires more parameters.

Lastly, we use the same simulation setting as in 4.4.2 to study the coverage probably of our global confidence region. For each scenario (shape corresponding to different true $J^\star$), we generate 300 replicates and estimate the empirical level-crossing bound and gamma bound for each γ value. For each replicate, we construct the global confidence region described in Algorithm 3.

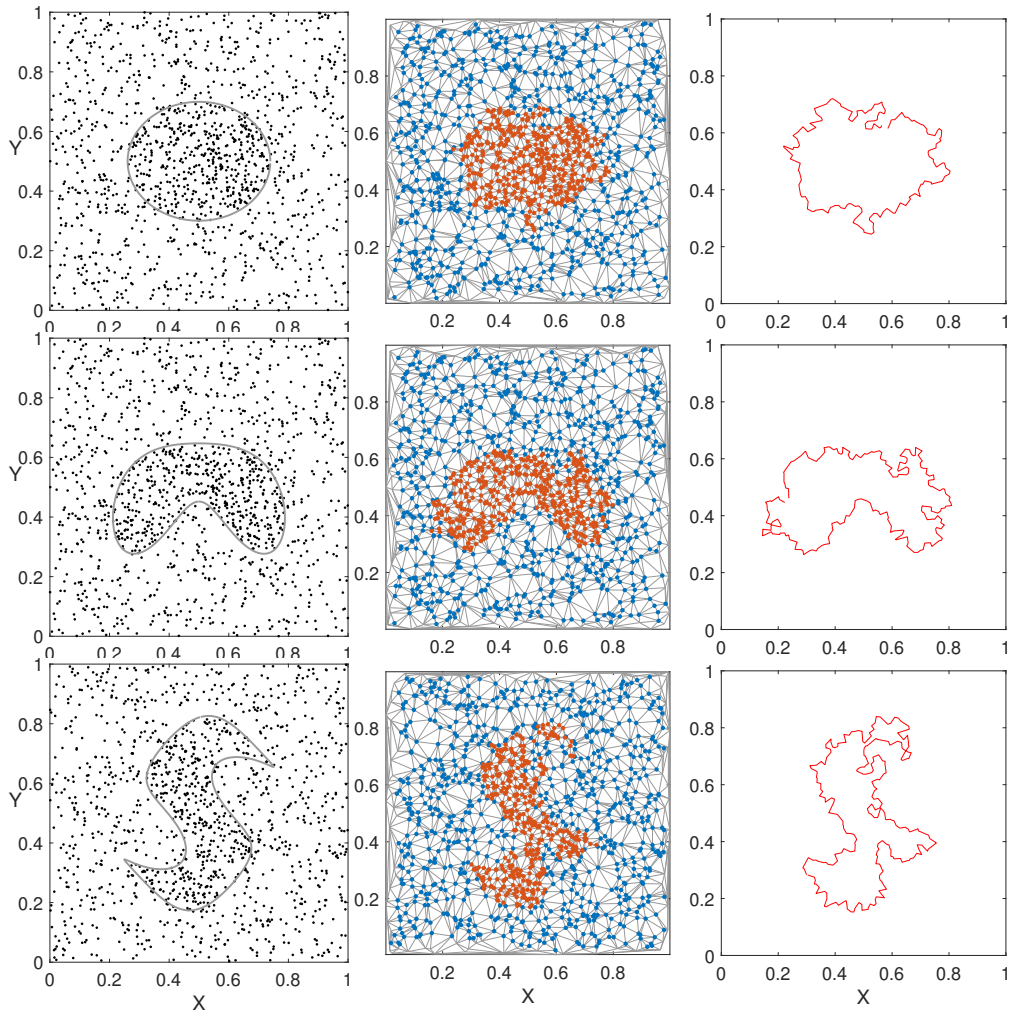


FIGURE 4.5. Simulation experiments with various extended sources. Each extended source has a size of 0.2 and the expected number of photons from the background and the source is 1000 and 300 respectively. **Left**: Simulated observations whose groundtruth of source boundaries is modeled with FD parametrization with $J = 3, 7, 11$ from top to bottom; **Middle**: `SRGonG` with Merge-to-Two Step results; **Right**: Segment's boundaries.

The level-crossing probability and their empirical bounds are displayed in Figure 4.8 for each scenario. We denote the empirical level-crossing probability using BIC model selection by P_e^{BIC} and that using true number of FD as P_e. The simulation results show that under this particular simulation setting, both gamma bound and estimated level-crossing bound the empirical level-crossing probability from above when $\hat{U}_\gamma$ is constructed with true $J^\star$. For both ellipse and bridge cases, $\hat{U}_\gamma$ exhibits high global confidence level. The P_e^{BIC} value is slightly higher than P_e when

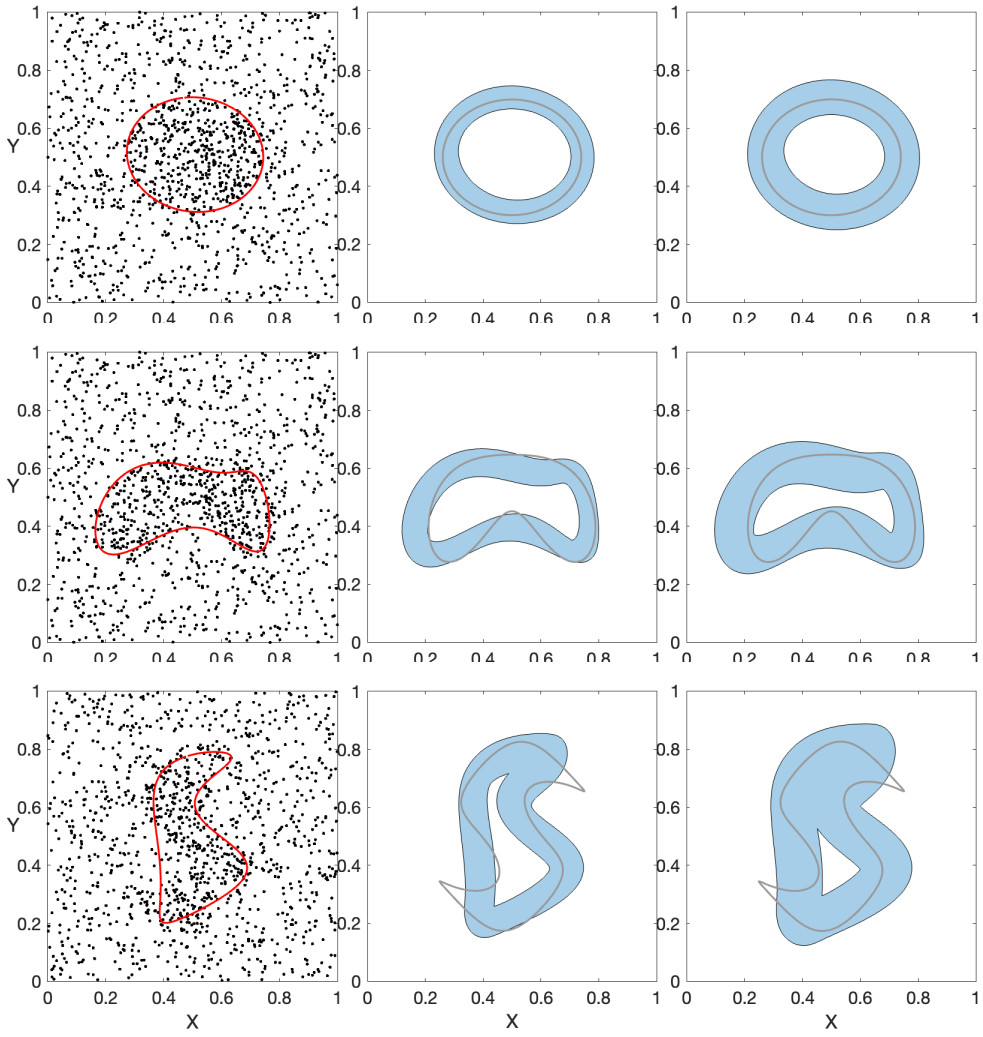


FIGURE 4.6. Simulation experiments on ellipse, bridge and spiral shape. **Left**: Estimated extended source boundary $f(t;\hat{\boldsymbol{\theta}};\hat{J})$ with $\hat{J}_{\mathrm{BIC}} = 3, 7, 7$ for the three shapes, respectively; **Middle**: global confidence region $\hat{U}_\gamma$ with $\gamma = 1$, corresponding to $\alpha \approx 0.39$; **Right**: global confidence region $\hat{U}_\gamma$ with $\gamma = 1.51$, corresponding to $\alpha = 0.68$.

$J^\star = 7$ due to the underfitting of BIC. The performance of $\hat{U}_\gamma$ degenerates under spiral case mainly due to the challenge of fitting the hook-like shape at both ends of the spiral.

4.5. Applications

In this section, we illustrate the Fourier descriptor parametrization and related analysis on real data. NGC 2300 is an early-type galaxy residing in the NGC 2300 group. The surface brightness

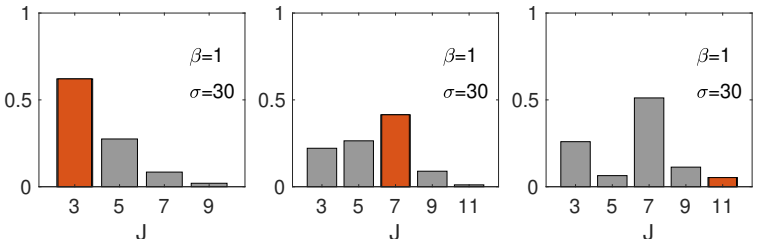


FIGURE 4.7. Simulation studies with various shape of extended sources. The histogram of BIC fitted value $\hat{J}$ for 300 replicates with the true number of J highlighted.

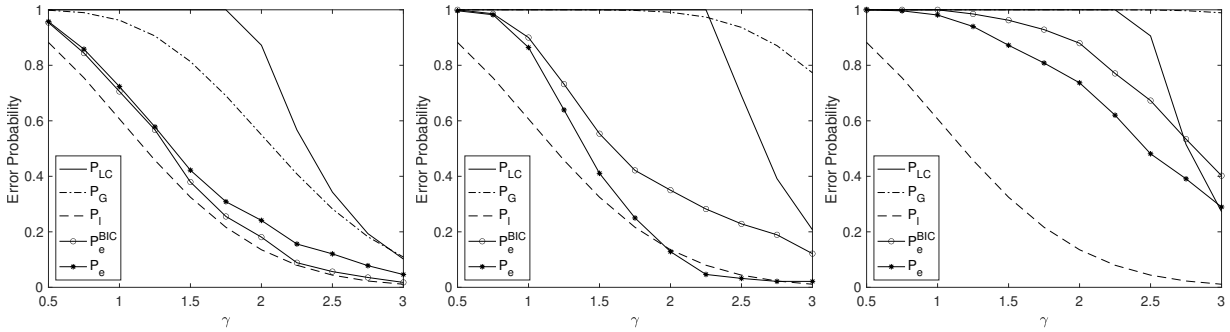


FIGURE 4.8. Comparisons of the Monte Carlo simulated empirical level-crossing probability P_e^{BIC} and P_e with the gamma bound P_G, estimated level-crossing bound P_{LC} and P_l for the shape in 4.5; **Left**: $J = 3$, **Middle**: $J = 7$, **Right**: $J = 11$).

of its X-ray emission follows a β profile and it is fairly smooth (Davis et al. 1996). This makes it an excellent test case for identifying the boundary of extended sources such as early type galaxies, galaxy groups, or higher redshift galaxy clusters. Since it has been observed with both *Chandra* and XMM-*Newton* it allows us to compare `UQ-BoSE` on two datasets with very different spatial resolution.

4.5.1. Implementing `BoSE`. To apply `BoSE` we must convert each dataset to a $[0,1] \times [0,1]$ field of view by linearly transforming their sky coordinates. The conversion parameter are given in Table 4.2; the center coordinate is mapped to $(0.5, 0.5)$, and each 0.1×0.1 area corresponds to an area of $dx \times dy$ in sky coordinates, with dx and dy given in Table 4.2.

We set the `SRGonG` seeds via voroni-area stratified sampling (Fan *et al.*, 2023), construct the over-segmented event list, and apply the Hybrid Merging Process with $M_s = 150$ Randomized Merger Steps. We repeat the procedure with 24 combinations of the voroni-area stratified sampling seeding parameters, specifically, six values for the number strata, i.e, m_{R} = [5,8,11,14,17,20]

Object	Insturment	Source Center	RA scale	Dec scale
NGC 2300	Chandra	(3415.01, 3830.98)	34.7"	33.5"
NGC 2300	XMM	(26107.5, 28046.5)	24.0"	24.0"

TABLE 4.2. Rescaling parameters for the three datasets. The source center corresponds to the physical pixel valus (x,y) in the original image. The scale corresponds to the angular size of a a length=0.1 in therescaled image.

crossed with four values for the minimum number of photons required for a Voronoi-area seed to be accepted, i.e., $m_{\text{vorthr}} = [5,10,15,20]$. We apply `BoSE` to each of the 24 resulting segmentations and choose the $(m_{\text{R}}, m_{\text{vorthr}})$ pair with the maximum penalized profile log-likelihood for its two-region segmentation. We then apply `UQ-BoSE` including the estimation of J via BIC and the generation of N_B non-parametric bootstrap datasets as described in Section 4.3.3 and 4.3.6.2. Finally, by applying `BoSE` on each of the N_B resampled datasets, we obtain the global confidence region for the source boundary as described in Algorithm 3.

4.6. Discussion

4.6.1. Using AIC for Model Selection Procedure. BIC is known for its consistency in model selection problems , but may suffer from underfitting Burnham and Anderson (2002). Here we consider using AIC as an alternative to BIC for the penilty term when estimating J. Both the BIC and the AIC are available in implementation. The estimate of J under the AIC penality is

$$\hat{J}_{\text{AIC}} = \underset{J\in[3,L],\, \textit{with } J \textit{ odd}}{\arg\min} -2\mathcal{L}^*_{\text{profile}}(J|\mathbf{X}_{\mathcal{R}}, n_{\mathcal{R}}) + 2m_{\text{pseg}}(J), \tag{4.29}$$

where $m_{\text{pseg}}(J)$ and $\mathcal{L}^*_{\text{profile}}$ are as in Equation 4.11.AIC allows for more flexible fitting of the boundary, but may favor overly complex fits, i.e., AIC is prone to overfitting. As the complexity of the boundary increases, perhaps characterized by additional bulbs or sharp corners, the FD requires more parameters, making the AIC a preferable choice. It is important to note, however, that complex shapes often entail wider global confidence regions as more parameters lead to larger model variance. emphasizing the need for an upper limit constraint on AIC, possibly guided by domain-specific information.

4.6.2. Extending `UQ-BoSE` to Multiple Segments. Our current work is designed with a specific focus on conducting statistical inference for external boundary of a single extended sources

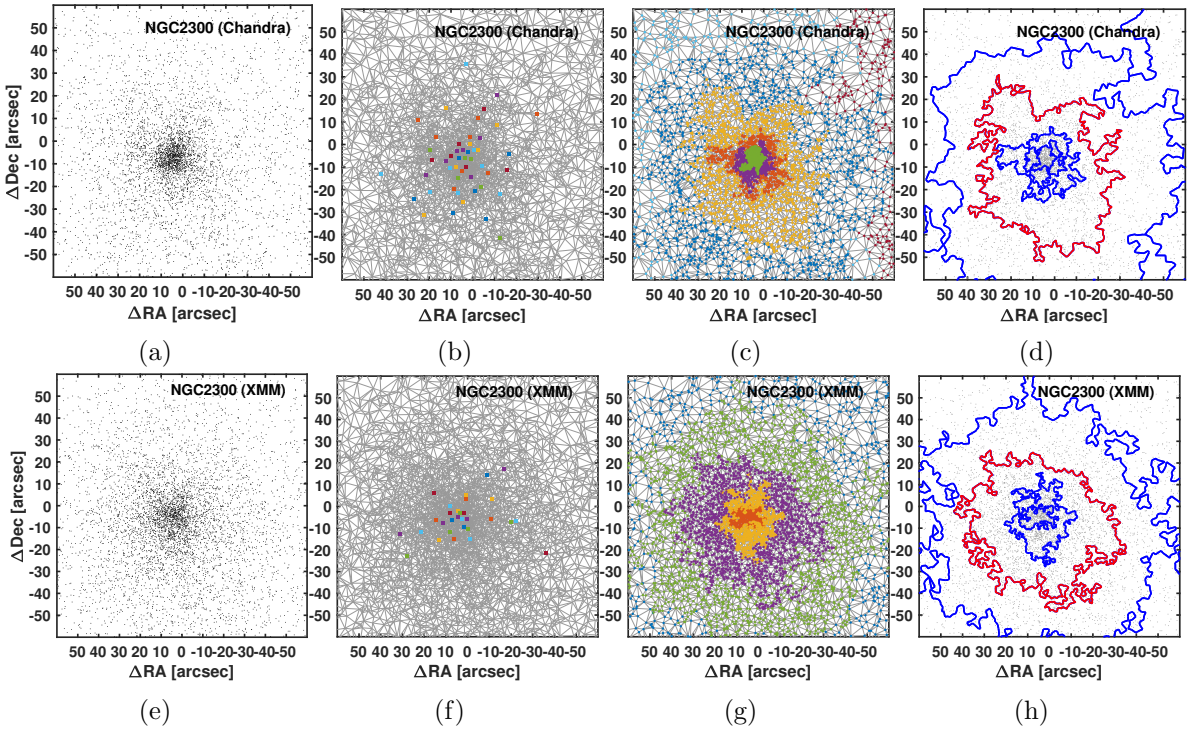


FIGURE 4.9. Representation of `SRGonG` on NGC2300 chandra (top) and NGC2300 XMM-Newton (bottom) data. **Leftmost:** scatter plots of the photons Chandra and XMM-Newton data. **Center-left:** seed specification for `SRGonG` via a Nearest-neighbor Labeling methods proposed in Fan *et al.* (2023). The hyperparameter pair $(m_{\mathrm{R}}, m_{\mathrm{vorthr}})$ are tuned via a grid search. chandra: $M_S = 80$; XMM: $M_S = 53$. **Center-right:** raw `SRGonG` segmentation $\widehat{\mathbf{S}}$. $\widehat{K} = 7$ for chandra and $\widehat{K} = 5$ for XMM data. **Rightmost:** boundaries of all $\widehat{\mathcal{S}}_k$ in $\widehat{\mathbf{S}}$. The C_1 extracted from two-segment $\widetilde{\mathbf{S}}$ is highlighted in red.

within the entire field of view $\mathcal{F}$. We presume that the external boundary of this extended source is fully contained within the field of view. While it is recognized that FD-based parametrization is inherently well-suited for addressing closed contours Burger and Burge (2009), it is pertinent to note that this parametrization framework can also be adapted to accommodate a partial two-dimensional shape parametrization, as discussed in Lin and Chellappa (1987). One avenue of future research expansion entails the extension of our uncertainty quantification methodology to scenarios where the underlying extended source is only partially contained within the field of view. In such cases, inference methodology can be applied to complete the boundary first, as proposed by Lin and Chellappa (1987), and then quantify the uncertainty of the inferred shape.

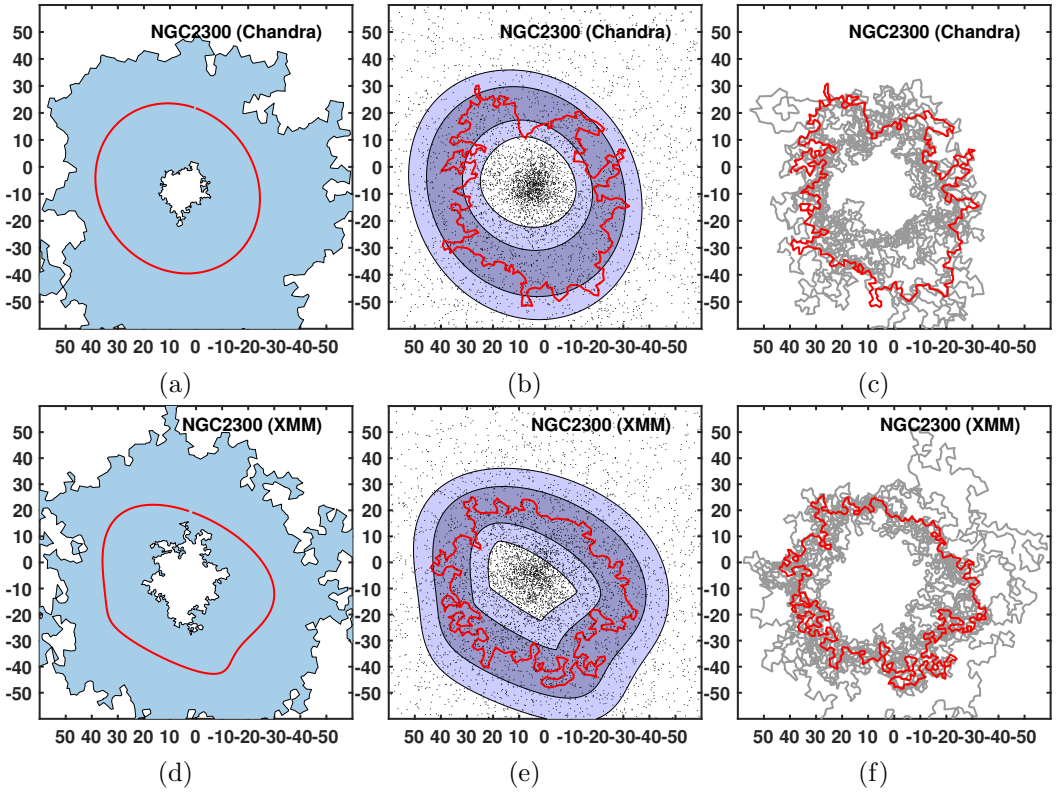


FIGURE 4.10. FD parameterization and uncertainty quantification on NGC2300chandra (top) and NGC2300XMM (bottom) dataset. **Left:** Red curve: FD parameterization of C_1 with $\hat{J}$ estimated by BIC (NCG2300chandra: $\hat{J}_{\text{BIC}} = 3$, NCG2300XMM: $\hat{J}_{\text{BIC}} = 7$); Blue region: $\mathcal{R}$, smallest region fully encloses C_1 identified within $\widehat{\mathbf{S}}$; **Middle:** $\hat{U}_\gamma$, the FD-based global confidence region constructed using nonparametric bootstrap data. Purple band: $\hat{U}_1$, light purple band: $\hat{U}_2$, red curve: C_1; **Right:** Illustrations of `BoSE` applied 10 non parametric bootstrap data. Red curve: C_1, grey curves: $C_1^{(N_b)}$ for $N_b = 1, \ldots, 10$.

While we acknowledge that real-world scenarios often involve the presence of multiple diffusive objects, our approach offers a foundational framework for statistical inference on individual objects, laying the groundwork for future developments in handling multiple segments within a unified statistical framework. One of the principal challenges impeding the development of an inference tool capable of handling multiple segments lies in the difficulty of aligning segmentations across various bootstrap datasets. Our observations indicate that, particularly in scenarios characterized by low signal-to-noise ratios (S/N), the `SRGonG` algorithm may result in the segmentation of a single diffusive region into multiple segments or merge it with other segments. This introduces ambiguity

in identifying segments that estimate the same region for further analysis. Addressing this challenge necessitates enhancing the robustness of the `SRGonG` algorithm to effectively manage such scenarios.

CHAPTER 5

Joint Mediation Effects Estimation Under High Dimensional Linear Graphical Causal Model

5.1. Introduction

Mediation analysis aims to investigate the role of intermediate variables played in the pathway through which a treatment affects a response VanderWeele (2015). Originating from Baron and Kenny's work in 1986 Baron and Kenny (1986), mediation analysis has gained significant interest in biological research over the past decades to understand complex biological, genetic, and behavioral phenomena Huang and Yang (2017); Zhang *et al.* (2016); Nguyen *et al.* (2016). The rise of high-dimensional data, such as high-throughput genomic data, has further emphasized the need for suitable mediation effect estimation and inference tools for datasets with a large number of potential mediators. Despite mediation methods being well-established for analyzing one or multiple mediators in low-dimensional settings, they are often not well-suited for high-dimensional datasets and can yield conservative p-values and limited power Zeng *et al.* (2021).

Our work is motivated by the MIDA method, which was developed to estimate the individual mediation effect from high-dimensional observational data Chakrabortty *et al.* (2018). The method assumed the causal relationship among all variables could be summarized as a DAG and relied on the estimation of an underlying CPDAG to produce a multi-set of possible mediation effects, which would be further summarized through averaging. The authors proved the high-dimensional consistency of MIDA with sub-Gaussian errors and derived distributional convergence results for MIDA in high-dimensional settings. The effectiveness of the algorithm was demonstrated through extensive numerical simulations as well as real data application. However, the open question left over is how MIDA could be generalized to estimate the joint mediation effect of multiple mediators.

This generalization is particularly important because intervention experiments are considered to be the gold standard for making causal inferences Kim and Steiner (2016), and multiple interventions are concurrently applied in many areas where the mediators of interest are potentially interrelated. For instance, double or triple gene knockout is popular in genomics as simultaneous modification of multiple homeologous genes, i.e., genes with similar but not identical DNA sequences, is often needed to obtain a desired phenotype Ramanagoudr-Bhojappa *et al.* (2018); Endo *et al.* (2015). However, as the number of possible pairs or tuples of mediators groups can impede comprehensive experiments, estimating the joint mediation effect through a potentially interrelated group of mediators could be extremely helpful in planning and prioritizing the experiments.

Consider the following case presented in Figure 5.1, where the causal effect $M_1 \to M_2$ exists. The MIDA method estimates the individual mediation effect as the marginalized path effects going through the mediator in the causal graph. More specifically, the mediation effect of M_1 is $T \to M_1 \to Y + T \to M_1 \to M_2 \to Y$, and the mediation effect of M_2 is $T \to M_2 \to Y + T \to M_1 \to M_2 \to Y$. In this case, the total mediation effect $T \to (M_1, M_2) \to Y$ cannot be decomposed into the sum of individual mediation effects due to the interrelation between potential mediators. In the direct summation, the effect through $T \to M_1 \to M_2 \to Y$ is double counted.

In this chapter, we first extend the definition of the individual mediation effect in Chakrabortty *et al.* (2018) to multiple mediators. We propose a new definition of the marginal mediation effect that can be decomposed into the sum of each potential mediator's component-wise mediation effect. In the above illustrative example, the marginal mediation effect through M_1, M_2 are decomposed to the component-wise mediation effect through M_1,

$$T \to M_1 \to Y,$$

and the component-wise mediation effect through M_2,

$$T \to M_2 \to Y + T \to M_1 \to M_2 \to Y,$$

respectively. The interpretation of the component-wise mediation effect of M_i is the marginalized mediation effects that go through M_i but none of its descendants. Our proposed definition not only allows the calculation of the total mediation effect regulated by a set of mediators but also provides

a reasonable way to interpret the portion of the mediation effect regulated by each mediator. The definition can be generalized to an arbitrary number of potential mediators with unknown causal ordering. We then develop an IDA-type method, named joint-MIDA, that estimates these effects from high-dimensional observational data. We prove its consistency under a high-dimensional setting and illustrate its effectiveness through both simulations and a real data application.

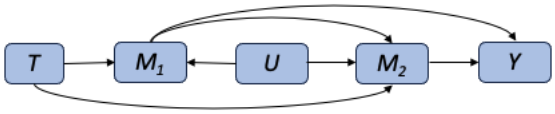


FIGURE 5.1. Causal diagram for an example with treatment T, two mediators M_1, M_2, one confounder U and response Y. M_1 is the post-treatment confounder of M_2.

Our work is summarized as follows: Section 5.2 provides background knowledge about graphs and defines the joint mediation effect. Section 5.3 introduces a scalable method, joint-MIDA, to estimate the identifiable version of the joint mediation effect and prove the consistency of the estimator. Then, in section 5.4, we present numerical experiments to illustrate the consistency of our estimator empirically and compare its performance in identifying pairs of mediator effects under different settings. Lastly, in Section 5.5, we apply the joint-MIDA method together with a non-parametric bootstrap inference tool to a DNA methylation dataset consisting of 385,882 DNA methylation loci. Our analysis identifies two new loci that possess reasonable scientific interpretations in relating the effect of childhood trauma on cortisol stress reactivity.

5.2. Definitions and Preliminaries

5.2.1. Graph Terminology. We consider graphs $\mathcal{H} = (\mathbf{X}, E)$ with vertex set $\mathbf{X} = \{X_1, \ldots, X_p\}$ and edge set E. For any set $\{X_i, \ldots, X_k\}$, we can write $\{X_{i,\ldots,k}\}$ equivalently. There is at most one edge between any pair of vertices $\{X_i, X_j\}$, which can either be directed ($X_i \to X_j$) or undirected ($X_i - X_j$). If $\mathcal{H}$ contains only (un)directed edges, it is called *(un)directed graph*; otherwise, it is called *partially directed*. The *skeleton* of a partially directed graph is the undirected graph that results from replacing all directed edges with undirected ones.

A pair of vertices $\{X_i, X_j\}$ is called *adjacent* if they are connected by an edge. Vertex X_i is a *parent* of X_j and X_j is a *child* of X_i if there exists $(X_i \to X_j)$ in $\mathcal{H}$. We denote the set of all parents and all adjacent vertices of X_j in $\mathcal{H}$ by $\mathbf{PA}_{\mathcal{H}}(X_j)$ and $\mathbf{Adj}_{\mathcal{H}}(X_j)$. A *path* between X_i and X_j is a sequence of distinct vertices such that all pairs of successive vertices are adjacent. A path is called *directed* if all edges in the path are directed. X_i is the *ancestor* of X_j, and X_j is the *descendant* of X_i if there is a directed path from X_i to X_j.

5.2.2. Linear Structural Equation Models and Causal Effects.

DEFINITION 5.2.1. *Let $\mathcal{G}_0 = (\mathbf{X}, E)$ be a Directed Acyclic Graph (DAG) with $|\mathbf{X}| = p$ vertices. Let $B_{\mathcal{G}_0}$ be a $p \times p$ weight matrix such that $(B_{\mathcal{G}_0})_{ij} \neq 0$ if and only if $X_i \in \mathbf{PA}_{\mathcal{G}_0}(X_j)$. $(B_{\mathcal{G}_0})_{ij}$ represents the weight of edge $(X_i \to X_j)$. Let $\boldsymbol{\epsilon} = (\epsilon_1, \ldots, \epsilon_p)$ be a zeros mean random vector of jointly independent error variables. Then $\mathbf{X} = (X_1, \ldots, X_p)^T$ is said to be generated from a linear structural equation model (LSEM) characterized by the pair $(B_{\mathcal{G}_0}, \boldsymbol{\epsilon})$ if*

$$(\mathbf{X} - \boldsymbol{\mu}) \leftarrow B_{\mathcal{G}_0}^T(\mathbf{X} - \boldsymbol{\mu}) + \boldsymbol{\epsilon}, \text{ where } \boldsymbol{\mu} := \mathbb{E}[\mathbf{X}]. \tag{5.1}$$

We call $\mathcal{G}_0$ the *causal DAG*. We assume there are no hidden confounders other than those specified in $\mathbf{X}$. Thus, the joint independence of errors, $\boldsymbol{\epsilon}$, is ensured.

5.2.3. Markov Equivalence and Completed Partially Directed Acyclic Graph. A *Markov equivalence class* refers to a set of DAGs that encode the same set of conditional independence relationships among a given set of random variables. In other words, two DAGs are considered *Markov equivalent* if they have the same skeleton (underlying undirected graph) and the same set of v-structures, where a v-structure is a vertex triplet $X_i \to X_k \leftarrow X_j$ and X_i, X_j are non-adjacent Pearl (2000). Therefore, the *Markov equivalence class* of a causal DAG $\mathcal{G}_0$ can be uniquely presented by a *Completed Partially Directed Acyclic Graph* (CPDAG) $\mathcal{C}_0$ Spirtes *et al.* (2000a) if the followings are satisfied: (1) $\mathcal{C}_0$ and $\mathcal{G}_0$ share the same skeleton; (2) $X_j \to X_i$ in $\mathcal{C}_0$ if $X_j \to X_i$ in every DAG that is in the same *Markov equivalence class* of $\mathcal{G}_0$. It has been shown by Chickering (2002) that two DAGs have the same CPDAG if and only if they belong to the same Markov equivalence class.

5.2.4. Intervention Calculus. In this section, we give a brief introduction to the intervention calculus based on Pearl (2009). We consider the above $p+1$ vertices $\mathbf{X} = \{X_1, \ldots, X_p\}$. We let $\mathbf{PA}_j = \mathbf{PA}(X_j)$ denote the parent set of vertex X_j in the DAG. Theorem 3.1 in Pearl (2009) states that the distribution of $\{X_1, \ldots, X_p\}$ that is generate from the DAG can be factorized to

$$f(x_1, \ldots, x_{p+1}) = \prod_{j=1}^{p+1} f(x_j | \mathbf{PA}_j)). \tag{5.2}$$

In a "do-operation", the distribution of $f(x_1, \ldots, x_{p+1} | do(X_i = x_i'))$ is then given by

$$f(x_1, \ldots, x_{p+1} | do(X_i = x_i')) = \begin{cases} \prod_{j=1, j\neq i}^{p+1} f(x_j | \mathbf{PA}_j), & \text{if } x_i = x_i' \\ 0, & \text{otherwise,} \end{cases} \tag{5.3}$$

where $\prod_{j=1, j\neq i}^{p+1} f(x_j | \mathbf{PA}(x_j))$ are the pre-intervention conditional probability Pearl (2009). Therefore, the distribution of X_{p+1} after a $\text{do}(X_i = x_i)$ can be written as following,

$$f(x_{p+1} | do(X_i = x_i')) = \begin{cases} f(x_{p+1}) & \text{if } X_{p+1} \in \mathbf{PA}(X_i) \\ \int f(x_p | x_i', \mathbf{PA}_i) f(\mathbf{PA}_i) \partial \mathbf{PA}_i, & \text{otherwise,} \end{cases} \tag{5.4}$$

where $f(\cdot)$ and $f(\cdot | x_i', \mathbf{PA}(x_i))$ denote the pre-intervention distribution. Under LSEM, $do(X_i = x_i')$ can be interpreted as holding X_i constant at x_i while removing all edges coming into X_i in $\mathcal{G}_0$ Pearl (2000).

5.2.5. Problem Setup. We assume that $\mathbf{X} = \{X_1, \ldots, X_p\}$ is generated from LSEM in Equation 5.1, and the causal effect is described with a DAG $\mathcal{G}_0$. We let $\{X_1, \ldots, X_{t-1}\}$ be the set of confounders; X_t be the treatment variable; $\mathbf{X}' = \{X_{t+1}, \ldots, X_{p-1}\}$ be the potential mediator set, and X_p be the response variable. We assume that no potential mediators are a direct cause of $\{X_1, \ldots, X_t\}$ and X_p is not the direct cause of any X_j for all $j < p$. We assume no other causal ordering in $\mathbf{X}$ besides those encoded by $B_{\mathcal{G}_0}$.

Maathuis *et al.* (2009) defined the causal effect of $X_i \to X_p$ based on *intervention* on X_i with *Pearl's do-operator*, "$do(X_i = x_i)$" or equivalently, "$do(x_i)$" introduced in Section 5.2.4. Let $\mathbb{E}[X_p | do(X_i = x_i)]$ denote the post-interventional expectation of X_p conditioning on intervening

X_i, the *total causal effect* of $X_i \to X_p$ is given by Maathuis *et al.* (2009),

$$\theta_{ip} := \frac{\partial}{\partial x_i} \mathbb{E}\left[X_p | do(X_i = x_i)\right],$$

Nandy *et al.* (2017) extended the causal effect of a single treatment to multiple ones and defined *total joint effect* of $(X_i, \ldots, X_k)$ on X_p,

$$\boldsymbol{\theta}_p^{(i,\ldots,k)} := (\theta_{ip}^{(i,\ldots,k)}, \ldots, \theta_{kp}^{(i,\ldots,k)}),$$

where

$$\theta_{ip}^{(i,\ldots,k)} := \frac{\partial}{\partial x_i} \mathbb{E}\left[X_p | do(X_i = x_i, \ldots, X_k = x_k)\right],$$

denote the total effect of X_i on X_p in a *joint intervention* of $do(X_i = x_i, \ldots, X_k = x_k)$.

It is proved that under the LSEM assumption, $\mathbb{E}\left[X_p | do(X_i = x_i)\right]$ is linear in x_i Maathuis *et al.* (2009); Nandy *et al.* (2017), and therefore, θ_{ip} and $\boldsymbol{\theta}_p^{(i,\ldots,k)}$ are independent of the value of x_i and $\{x_i, \ldots, x_k\}$, respectively. Note that in a nonlinear SEM, θ_{ip} and $\theta_{ip}^{(i,\ldots,k)}$ are nontrivial functions of x_i and $\{x_i, \ldots, x_k\}$ respectively.

Under LSEM, graphical interpretations of θ_{ip} and $\theta_{ip}^{(i,\ldots,k)}$ are given by the *path method* Wright (1921). Let $\{X_{i_0} \to X_{i_1} \to \ldots \to X_{i_{k+1}}\}$ denote a direct path from X_{i_0} to $X_{i_{k+1}}$ in $\mathcal{G}_0$. Then we can define the effect of X_{i_0} to $X_{i_{k+1}}$ through this path as $\prod_{r=0}^{k} (B_{\mathcal{G}_0})_{i_r i_{r+1}}$. The total causal effect θ_{ip} can be interpreted as the sum of effects through all such paths from X_i to X_p. The joint causal effect $\theta_{ip}^{(i,\ldots,k)}$ can be interpreted as the sum of effects through all such paths from X_i to X_p that does not go through any descendent of X_i, $X_j, i < j \leq k$. We then use $\theta_{ip}(\mathcal{G}_0)$ and $\theta_{ip}^{(i,\ldots,k)}(\mathcal{G}_0)$ to indicate the dependence of the causal effect on the weighted DAG $\mathcal{G}_0$.

We define the *marginal (joint) mediation effect* of an arbitrary potential mediator set $\{X_i, \ldots X_k\}$ on X_p as the portion of *total causal effect* of treatment variable X_t on response variable X_p that go through one of $\{X_i, \ldots, X_k\}$ for any $t < i \leq k < p$.

Definition 5.2.2 (Marginal (Joint) Mediation Effect). *The marginal mediation effect, $\eta_{i,\ldots,k}$, w.r.t. an arbitrary subset of all potential mediators, $\{X_{i,\ldots,k}\}$, is defined as*

$$\eta_{i,\ldots,k} := \frac{\partial}{\partial x_t} \mathbb{E}[X_p | do(X_t = x_t)] - \frac{\partial}{\partial x_t} \mathbb{E}[X_p | do(X_t = x_t, X_i = x_i, \ldots, X_k = x_k)]$$

REMARK 5.2.1. The interpretation of the marginal mediation effect $\eta_{i,\dots,k}$ is the change in marginal causal effect of treatment X_t on X_p when the potential mediators $\{X_i, \dots, X_k\}$ are all knocked out from the causal graph $\mathcal{G}_0$ by the intervention $do(X_i = x_i, \dots, X_k = x_k)$. This interpretation generally holds for any (non-)linear structural equation model, while it should be noted that under nonlinearity, $\eta_{i,\dots,k}$ is a nontrivial function of $\{x_t, x_i, \dots, x_k\}$. Note that the definition of *individual mediation effect* proposed in Chakrabortty *et al.* (2018) is essentially a special case of our definition when $|\{i, \dots, k\}| = 1$.

We further define the *component-wise (joint) mediation effect* that attributes the marginal mediation effects to each individual mediator.

DEFINITION 5.2.3 (Component-wise (Joint) Mediation Effect). *Without loss of generality, assume $\{X_i, \dots, X_k\}$ is causally ordered potential mediators. The component-wise mediation effect through X_j, $\eta_j^{(i,\dots,k)}$, w.r.t. an arbitrary subset of all potential mediators, $\{X_{i,\dots,k}\}$, is defined as*

$$\eta_j^{(i,\dots,k)} := \eta_{j,\dots,k} - \eta_{j+1,\dots,k}.$$

Next, we show that under LSEM, $\eta_{i,\dots,k}$ equals the dot product of the vector of total causal effects of X_t to $\{X_{i,\dots,k}\}$ and the total joint effect of $\{X_{i,\dots,k}\}$ on X_p.

Let $[i,k]$ and $[i,k] \setminus j$ denote $(i, \dots, k)$ and $(i, \dots j-1, j+1, \dots, k)$, respectively. First, we would like to prove the following lemma. The proof can be found in Appendix C.1:

LEMMA 5.2.3.1. *Let $\mathbf{X}$ be generated from an LSEM. If $\{X_{i,\dots,k}\}$ follows the causal order characterized by $\mathcal{G}_0$, i.e., there is no direct path in $\mathcal{G}_0$ from X_j to X_i if $i < j$. Then $\theta_{jp}^{[i,k]} = \theta_{jp}^{[j,k]}$ and $\theta_{tj}^{[i,k]} = \theta_{tj}^{[i,j]}$ for arbitrary $i \le j \le k$.*

THEOREM 5.2.4. *Let $\mathbf{X}$ be generated from an LSEM. Let $\boldsymbol{\beta}_t^{(i,\dots,k)} = (\theta_{ti}, \dots, \theta_{tk})$ denote the vector of total causal effects of X_t on each individual X_j, $j = i, \dots, k$. Let $\boldsymbol{\theta}_p^{(i,\dots,k)} = (\theta_{ip}^{(i,\dots,k)}, \dots, \theta_{kp}^{(i,\dots,k)})$ denote the total joint effect of $\{X_{i,\dots,k}\}$ on X_p, then*

$$\eta_{i,\dots,k} = {\boldsymbol{\beta}_t^{(i,\dots,k)}}^T \boldsymbol{\theta}_p^{(i,\dots,k)} \tag{5.5}$$

The proof can be found in Appendix C.2. As a direct result of proof of Theorem 5.2.4, the component-wise mediation effect $\eta_j^{(i,\ldots,k)}$ can be uniquely defined without assuming knowing the causal ordering of $\{X_i, \ldots, X_k\}$. We summarize the result in the following proposition.

PROPOSITION 5.2.1 (Uniqueness). *Let* $\mathbf{X}$ *be generated from an LSEM. The definition of* $\eta_j^{(i,\ldots,k)}$ *does not depend on the chosen causal ordering of the potential mediators and*

$$\eta_j^{(i,\ldots,k)} = \theta_{tj} \times \theta_{jp}^{(i,\ldots,k)}.$$

By path method, $\eta_j^{(i,\ldots,k)}$ is interpreted as the portion of marginal mediation effect through X_j but not any descendent of X_j in $\{X_{i,\ldots,k}\}$.

EXAMPLE 5.2.5. *We consider an example with* $p = 8, t = 2$. *The causal structure includes one pre-treatment variable* X_1, *the treatment variable* X_2, *the potential mediators* $\{X_3, \ldots, X_7\}$, *and the response variable* X_8. *The variables are generated from the following LSEM:* $X_1 \leftarrow \epsilon_1, X_1 \leftarrow \epsilon_1, X_2 \leftarrow X_1 + \epsilon_2, X_3 \leftarrow 1.3X_1 + \epsilon_3, X_4 \leftarrow 0.3X_2 + 1.4X_7 + \epsilon_4, X_5 \leftarrow 1.2X_3 + 0.6X_2 + \epsilon_5, X_6 \leftarrow 1.5X_5 + \epsilon_6, X_7 \leftarrow 0.5X_2 + 0.7X_1 + \epsilon_7, X_8 \leftarrow 1.2X_4 + 0.8X_6 + 0.2X_7 + \epsilon_8$. *The DAG in Figure 5.2 represents the causal structure. We present the total causal effect of the treatment variable on each of the potential mediators, the total joint effect of each of the potential mediators on the response variable, and the component-wise mediation effect of each potential mediator in Table 5.1. Note that the component-wise mediation effect of* X_3 *and* X_5 *are both 0 since when* X_6 *is knocked out, the total causal effect of* X_2 *on* X_8 *will not change if either* X_3 *or* X_5 *is further knocked out.*

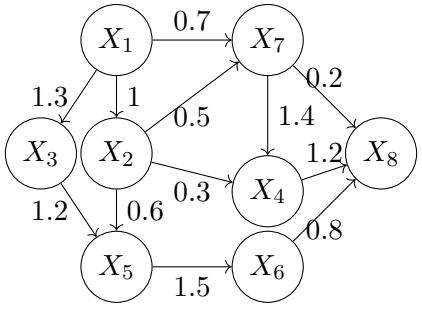


FIGURE 5.2. Example of weight causal DAG representing the schema of data generation in Example 5.2.5

TABLE 5.1. component-wise mediation effects: Illustration for the DAG in Example 5.2.5.

j	θ_{2j}	$\theta_{2j}^{(3,\ldots,7)}$	$\eta_j^{(3,\ldots,7)}$
3	0	$1.2 \times 1.5 \times 0.8$	0
4	$0.3 + 0.5 \times 1.4$	1.2	1.2
5	0.6	0	0
6	0.6×0.5	0.8	0.24
7	0.5	0.2	0.1

5.3. Estimating Joint Mediation Effect

Maathuis *et al.* (2009) developed the IDA (Intervention when the DAG is Absent) method to estimate the total causal effect via covariates adjustment, i.e., regressing the outcome to the mediator and the parent set of the mediator. Nandy *et al.* (2017) further extended the IDA algorithm to joint-IDA that estimates the total joint effect of multiple variables that may have a correlation among themselves. Applying the joint-IDA method, Chakrabortty *et al.* (2018) developed the MIDA method to estimate the marginal mediation effect of a single mediator . In this section, we introduce joint-MIDA, an IDA-based algorithm extending from MIDA, that estimates $\eta_{i,\ldots,k}$ as well as each $\eta_j^{(i,\ldots,k)}$ through covariates adjustments. We first introduce notations and necessary conditions and present the algorithm in section 5.3.2.3.

5.3.1. Notations and Assumptions. Let $\mathcal{G}_0 = (\mathbf{X}, E)$ denote the underlying causal structure of $\mathbf{X}$, and $\mathcal{C}_0$ be the CPDAG of $\mathcal{G}_0$. Let $\mathcal{G}_0'$ denote the subgraph of $\mathcal{G}_0$ by deleting vertices representing confounders $\{X_1, \ldots, X_{t-1}\}$, treatment $\{X_p\}$ and response $\{X_p\}$ as well as their corresponding edges. Let $\mathcal{C}_0'$ denote the CPDAG of $\mathcal{G}_0'$. Let $B_{\mathcal{G}_0'}$ be the submatrix of $B_{\mathcal{G}_0}$ that corresponds to $\mathbf{X}' = \{X_{t+1}, \ldots, X_{p-1}\}$ and $\boldsymbol{\epsilon}' = (\epsilon_{t+1}, \ldots, \epsilon_{p-1})^T$.

We denote the covariance matrix of $\mathbf{X}$ by Σ_0. For any set $S \subset \{1, \ldots, p\}$, we denote the corresponding random vector $\{X_r : r \in S\}$, i.e., the restriction of $\mathbf{X}$ on S, by $\mathbf{X}_S$. We denote the $i-$th column of the $k \times k$ identity matrix by $e_{i,k}$. For $i \neq k$ and any set $S \subset \{1, \ldots, p\} \setminus \{i, k\}$, we denote the coefficient of X_i in the linear regression of X_k on $\{\mathbf{X}_{\{i\}\cup S}\}$ by $\beta_{ik|S}$ or by $\beta_{ik|\mathbf{X}_S}$. For simplicity, we denote $\beta_{ik|\emptyset}$ as β_{ik}. Further more, we allow $\beta_{ik|S}$ to be *model free*, i.e., $\beta_{ik|S} = e_{1,|S|+1}^T((\Sigma_0)_{(i,S)(i,S)})^{-1}(\Sigma_0)_{(i,S)k}$ regardless of the distribution of conditional expectation $\mathbb{E}[X_k|\mathbf{X}_{\{i\}\cup S}]$. The sample version of $\beta_{ik|S}$, $\hat{\beta}_{ik|S}$, is equal to $e_{1,|S|+1}^T((\hat{\Sigma}_0)_{(i,S)(i,S)})^{-1}(\hat{\Sigma}_0)_{(i,S)k}$.

We assume that the conditional distribution of $\mathbf{X}'$ given $\{X_{i,\ldots,t}\}$ is *faithful* to $\mathcal{G}_0$, which states that any independence in the distribution is encoded in the d-separation properties of $\mathcal{G}'_0$ (Koller and Friedman, 2009, Definition 3.8). The faithfulness assumption is necessary for learning causal structures from observational data Chickering (2002) but is not needed when the underlying causal structure is known.

5.3.2. Mediation Effect Estimator: joint-MIDA. Assuming we have n i.i.d. observational sample $\mathbf{X}$ from an LSEM. We would like to estimate component-wise mediation effect $\eta_j^{(i,\ldots,k)} = \theta_{tj} \times \theta_{jp}^{(i,\ldots,k)}$ for each $j = i, \ldots, k$ and marginal mediation effect $\eta_{i,\ldots,k} = \boldsymbol{\beta}_t^{(i,\ldots,k)^T} \boldsymbol{\theta}_p^{(i,\ldots,k)}$ from the observational samples.

5.3.2.1. *Calculating Causal Effect Via Covariate Adjustment.* It has been proved that under LSEM, the total causal effect of X_j on X_p, θ_{jp}, is equivalent to $\beta_{jp|S}$ for any appropriately chosen set $\mathbf{X}_S$ that satisfies Pearl's back-door criterion if $X_p \notin \mathbf{PA}_{\mathcal{G}_0}(X_j)$ and 0 otherwise Maathuis *et al.* (2009); Nandy *et al.* (2017). $\mathbf{X}_S$ satisfies the back-door criterion if $\mathbf{X}_S$ block all the path from X_j to X_p and none of $\mathbf{X}_S$ is the descendent of X_p *Pearlet al.* (2000). One choice of such covariates set, which was used in Maathuis *et al.* (2009); Nandy *et al.* (2017), is the parent set of X_j, denoted as $\mathbf{PA}_{\mathcal{G}_0}(X_j)$. An alternative choice proposed in Chakrabortty *et al.* (2018) is $\mathbf{PA}_{\mathcal{G}'_0}(X_j) \cup \mathbf{X}_{\{1,\ldots t\}}$, i.e,

$$\theta_{jp} = \beta_{jp|\mathbf{PA}_{\mathcal{G}'_0}(X_j) \cup \mathbf{X}_{\{1,\ldots t\}}}, \tag{5.6}$$

leveraging on the facts that no potential mediator is a direct cause of a variable in $\{X_{1,\ldots,t}\}$. Similarly, we have

$$\theta_{tj} = \beta_{tj|\mathbf{X}_{\{1,\ldots,t-1\}}}. \tag{5.7}$$

5.3.2.2. *Recursive Regressions for Causal Effect Method.* Let $\mathbf{PA}'_{\mathcal{G}'_0}(X_i)$ denote $\mathbf{PA}_{\mathcal{G}'_0}(X_i) \cup X_{\{1,\ldots,t\}}$. Applying the causal effect calculation through covariates adjustment in Equation 5.6, we calculate the total joint effect $\theta_{ip}^{(i,\ldots,k)}$ based on the oracle version of the *Recursive regressions for*

causal effect (RRC) method proposed in Nandy *et al.* (2017), i.e,

$$\theta_{ip}^{[i,k]} = \theta_{ip}^{[i,k]\setminus\{j\}} - \theta_{ik}^{[i,k]\setminus\{j\}}\theta_{kp}^{[i,k]\setminus\{j\}}$$

$$\theta_{ip} = \begin{cases} 0, & \text{if } X_p \in \mathbf{PA}(X_i), \\ \beta_{ip|\mathbf{PA}'_{\mathcal{G}'_0}(X_i)} & \text{otherwise.} \end{cases} \tag{5.8}$$

This oracle calculates the joint effect via performing covariates adjustment recursively, assuming the true distribution of $\mathbf{X}$ is fully known. The sample version of the RRC method (RRC estimator) is defined as follows,

$$\hat{\theta}_{ip}^{[i,k]} = \begin{cases} 0, & \text{if } X_p \in \mathbf{PA}(X_i), \\ \hat{\theta}_{ip}^{[i,k]\setminus\{j\}} - \hat{\theta}_{ik}^{[i,k]\setminus\{j\}}\hat{\theta}_{kp}^{[i,k]\setminus\{j\}} & \text{otherwise,} \end{cases} \tag{5.9}$$

and

$$\hat{\theta}_{ip} = \begin{cases} 0, & \text{if } X_p \in \mathbf{PA}(X_i), \\ \hat{\beta}_{ip|\mathbf{PA}'_{\mathcal{G}'_0}(X_i)} & \text{otherwise,} \end{cases} \tag{5.10}$$

where $\hat{\beta}_{ip|\mathbf{PA}'_{\mathcal{G}'_0}(X_i)}$ is the sample regression coefficient of X_i in the linear regression $X_p \sim X_i + \mathbf{PA}'_{\mathcal{G}'_0}(X_i)$.

Let $\mathcal{PA}'_{\mathcal{G}'_0}(X_i, \ldots, X_k) = \left(\mathbf{PA}'_{\mathcal{G}'_0}(X_i), \ldots, \mathbf{PA}'_{\mathcal{G}'_0}(X_k)\right)$. We denote $\boldsymbol{\theta}_p^{(i,\ldots,k)}(\mathcal{G}'_0) = \boldsymbol{\theta}_p^{(i,\ldots,k)}\left(\mathcal{PA}'_{\mathcal{G}'_0}(X_i, \ldots, X_k)\right)$ to indicate $\boldsymbol{\theta}_p^{(i,\ldots,k)}$ is calculated with the RRC method via adjustment of the particular set of covariates $\mathcal{PA}'_{\mathcal{G}'_0}(X_i, \ldots, X_k)$ depending on $\mathcal{G}'_0$. Similarly, we use $\hat{\boldsymbol{\theta}}_p^{(i,\ldots,k)}(\mathcal{G}'_0)$ to denote the RRC estimator of $\boldsymbol{\theta}_p^{(i,\ldots,k)}(\mathcal{G}'_0)$.

5.3.2.3. *Identifiable joint-MIDA Estimator.* When $\mathbf{X}$ is given and the true partial DAG $\mathcal{G}'_0$ is known, $\eta_j^{(i,\ldots,k)} = \theta_{tj} \times \theta_{jp}^{(i,\ldots,k)}$ can be estimated by plugging in the RRC estimators of θ_{tj} and $\theta_{jp}^{(i,\ldots,k)}$, i.e.,

$$\hat{\eta}_j^{(i,\ldots,k)} = \hat{\beta}_{tj|\mathbf{X}_{\{1,\ldots,t-1\}}} \times \hat{\theta}_{jp}^{(i,\ldots,k)}(\mathcal{PA}'_{\mathcal{G}'_0}(X_i, \ldots, X_k)).$$

$\eta_{i,\ldots,k}$ could then be estimated by $\sum_{j=i}^{k} \hat{\eta}_j^{(i,\ldots,k)}$.

When $\mathcal{G}'_0$ is not known, we would like to estimate $\mathcal{G}'_0$, which is usually not identifiable from the observational data, so is the set $\mathcal{PA}'_{\mathcal{G}'_0}(X_i, \ldots, X_k)$. Therefore, we follow the approach in Maathuis

et al. (2009); Nandy *et al.* (2017) and aim to estimate the identifiable version of the joint mediation effects.

Let $\mathcal{C}_0'$ denote the CPDAG of $\mathcal{G}_0'$ and MEC($\mathcal{C}_0'$) denote the Markov equivalent class uniquely represented by the $\mathcal{C}_0'$. Let aver$(A) = |A|^{-1}\sum_{a\in A} a$ for an arbitrary set A. We aim to estimate the following identifiable version $\eta_j^{(i,\ldots,k)}(\mathcal{C}_0')$ and $\eta_{i,\ldots,k}(\mathcal{C}_0')$ of $\eta_j^{(i,\ldots,k)}$ and $\eta_{i,\ldots,k}$,

$$\eta_j^{(i,\ldots,k)}(\mathcal{C}_0') = \theta_{tj} \times \text{aver}\left(\Theta_{jp}^{(i,\ldots,k)}(\mathcal{C}_0')\right),$$

$$\eta_{i,\ldots,k}(\mathcal{C}_0') = \sum_{j=i}^{k} \eta_j^{(i,\ldots,k)}(\mathcal{C}_0')$$

where

$$\Theta_{jp}^{(i,\ldots,k)}(\mathcal{C}_0') = \left\{\theta_{jp}^{(i,\ldots,k)}\left(\mathcal{G}'\right) : \mathcal{G}' \in \text{MEC}(\mathcal{C}_0')\right\},$$

is a multi-set of possible total causal effect X_j on X_p in intervention of $\{X_i, \ldots, X_k\}$.

Lastly, we apply the method proposed in Section 3.2 of Chakrabortty *et al.* (2018) to estimate the CPDAG $\mathcal{C}_0'$. To summarize, we first replace the data corresponding to each X_j with the residuals of the linear regression $X_j \sim X_1 + \ldots + X_t$ in order to remove the effect of $\mathbf{X}_{1,\ldots,t}$ on each of $X_j, j = t+1, \ldots, p-1$. We then apply any suitable structure learning algorithm, e.g. (AR)GES Nandy *et al.* (2018a) or PC algorithm Spirtes *et al.* (2000b), on the residuals to estimate $\hat{\mathcal{C}}_0'$.

According to Lemma 3.2 of Chakrabortty *et al.* (2018), the distribution of $\mathbf{X}' - \mathbb{E}[\mathbf{X}'|\mathbf{X}_{1,\ldots,t}]$ is faithful to $\mathcal{G}_0'$, which justify the estimation of $\mathcal{C}_0'$ from the residuals.

The RRC method is then applied to $\hat{\mathcal{C}}_0'$ estimate $\hat{\Theta}_{jp}^{(i,\ldots,k)}(\hat{\mathcal{C}}_0')$ for each $j \in \{i, \ldots, k\}$. We summarize the steps of the joint-MIDA approach in the Algorithm 4:

5.3.3. Uniform Consistency in High-Dimensional Sparse Settings. We now consider an asymptotic scenario where the sample size n as well as the number of potential mediators $(p-t-1)$ grows to infinity, while $t \geq 1$ is fixed. We prove the high-dimensional consistency of the joint-MIDA estimators $\hat{\eta}_j^{(i,\ldots,k)}(\hat{\mathcal{C}}_0')$ and $\hat{\eta}_{i,\ldots,k}(\hat{\mathcal{C}}_0')$ when the following assumptions hold.

ASSUMPTION 5.3.1. $\mathbf{X}$ *is generated from a linear SEM with sub-Gaussian error variables satisfying* $\max_{1\leq i\leq p} ||\epsilon_i||_{\psi_2} \leq C_1$ *for some absolute constant* $C_1 > 0$.

ASSUMPTION 5.3.2. $p = O(n^a)$ *for some* $a \geq 0$.

Algorithm 4 joint-MIDA

Require: n i.i.d observations of $\mathbf{X}$ (data)

Ensure: Estimated $\eta_{i,\dots,k}(\mathcal{C}_0')$ and $\eta_j^{(i,\dots,k)}(\mathcal{C}_0')$ for $j = i, \dots, k$.

1: for $j = t+1, \dots, p-1$, obtain the vector of residuals $\mathbf{r}_j = (r_j^{(1)}, \dots, r_j^{(n)})$ and coefficient $\hat{\beta}_{tj|\mathbf{X}_{1,\dots,t-1}}$ from regression of X_j on $X_{1,\dots,t}$;

2: apply any suitable structure learning algorithm on the data $\mathbf{r}_{t+1}, \dots, \mathbf{r}_{p-1}$ to obtain an estimate $\hat{\mathcal{C}}_0'$ of the CPDAG $\mathcal{C}_0'$;

3: obtain a multi-set of possible joint causal effects $\hat{\mathbf{\Theta}}_p^{(i,\dots,k)}(\hat{\mathcal{C}}_0') := \{\hat{\boldsymbol{\theta}}_p^{(i,\dots,k)}(\mathcal{G}') : \mathcal{G}' \in \mathrm{MEC}(\hat{\mathcal{C}}_0')\}$ with the RRC method;

4: **return** $\hat{\eta}_j^{(i,\dots,k)}(\hat{\mathcal{C}}_0') := \hat{\beta}_{tj|\mathbf{X}_{1,\dots,t-1}} \times \mathrm{aver}(\hat{\Theta}_{jp}^{(i,\dots,k)}(\hat{\mathcal{C}}_0')), \forall j \in \{i, \dots, k\}$, and $\hat{\eta}_{i,\dots,k}(\hat{\mathcal{C}}_0') := \sum_{j=i}^{k} \hat{\eta}_j^{(i,\dots,k)}(\hat{\mathcal{C}}_0')$.

ASSUMPTION 5.3.3. *Let* $q := \max_{t<j<p} |\mathbf{Adj}_{\mathcal{C}_0'}(X_j)|$ *denote the maximum degree in* $\mathcal{C}_0'$. *Then* $q = O(n^{1-b_1})$ *for some* $0 < b_1 \leq 1$. *This condition enforces the sparsity of* $\mathcal{C}_0'$.

ASSUMPTION 5.3.4. $\mathbb{P}$ *The estimated CPDAG* $\hat{\mathcal{C}}_0'$ *in Algorithm* 5.3.2.3 *is a consistent estimator of* $\mathcal{C}_0'$, *i.e.* $\mathbb{P}(\hat{\mathcal{C}}_0' \neq \mathcal{C}_0') \to 0$.

ASSUMPTION 5.3.5. *For any* $(q+t+2) \times (q+t+2)$ *principle matrix* Σ *of* $\Sigma_0 = \mathrm{Cov}(\mathbf{X})$,

$$C_2 \leq 1/||\Sigma^{-1}||_2 \leq ||\Sigma||_2 \leq C_3,$$

for some absolute constant $C_2, C_3 > 0$, *where* $||\cdot||_2$ *denotes the spectral norm.*

For simplicity, we only consider double mediators. The proof can be found in Appendix C.3.

THEOREM 5.3.6. *(Uniform consistency of Joint-MIDA estimator under sub-Gaussian errors). Let* $\hat{\Theta}_{ip}^{(i,j)}(\hat{\mathcal{C}}_0')$ *be the multi-set of estimated joint effect* X_i *on* X_p *in double intervention of* X_i, X_j. *Let* $\hat{\eta}_{i,j}(\hat{\mathcal{C}}_0')$ *and* $\hat{\eta}_i^{(i,j)}(\hat{\mathcal{C}}_0')$ *joint-MIDA estimators of the marginal mediation effect and component-wise mediation effect. Then under assumptions 5.3.1-5.3.5, we have* $\mathbb{P}$

$$\max_{t<i,j<p, i\neq j} \left| aver(\hat{\Theta}_{ip}^{(i,j)}(\hat{\mathcal{C}}_0')) - aver(\Theta_{ip}^{(i,j)}(\mathcal{C}_0') \right| \xrightarrow{\mathbb{P}} 0, \tag{5.11}$$

$$\max_{t<i,j<p, i\neq j} \left| \hat{\eta}_i^{(i,j)}(\hat{\mathcal{C}}_0') - \eta_i^{(i,j)}(\mathcal{C}_0') \right| \xrightarrow{\mathbb{P}} 0, \tag{5.12}$$

$$\max_{t<i,j<p, i\neq j} \left|\hat{\eta}_{i,j}(\hat{\mathcal{C}}'_0) - \eta_{i,j}(\mathcal{C}'_0)\right| \xrightarrow{\mathbb{P}} 0. \tag{5.13}$$

5.4. Simulation Experiments

5.4.1. Simulations Experiment 1. In this simulation experiment, we provide empirical support for Theorem 5.3.6. We also compare the results of joint-MIDA estimators with that of covariates adjustment using only the treatment and confounders (referred to as "empty graph").

We use the R-package *pcalg* Kalisch *et al.* (2024) to simulate samples from random linear LSEMs by constructing a block structure specified as follows: we generate n_{graph} random weighted DAGs, $\{\mathcal{G}'_m\}_{m=1}^{n_{graph}}$; each $\mathcal{G}'_m$ contains p vertices, and each vertex has an expected degree equal to 2. We then obtain the DAG $\mathcal{G}_0$ from $\mathcal{G}'_0 = \bigcup_{m=1}^{n_{graph}} \mathcal{G}'_m$ by randomly adding direct edges from treatment vertex, X_1, to each mediator in $\mathcal{G}'_0$ with probability 0.2 for, and from each mediator vertex within a randomly chosen DAG $\mathcal{G}'_q$ to the response variable, $X_{n_{graph}p+2}$, with probability 0.2. All edge weights of $\mathcal{G}_0$ are independently sampled from a Uniform($[-1,-0.5] \cup [0.5,1]$) distribution. We let $p = 10$ and $n_{graph} = 50$, resulting in a graph $\mathcal{G}_0$ with 502 vertices.

For each repetition $r = 1, \ldots, n_{itr}$, we generate a random graph $\mathcal{G}_0^{(r)}$ specified by the above mechanism. The DAG that contains true mediators is denoted as $\mathcal{G}_{q_r}^{\prime(r)}$. Let $B_{\mathcal{G}_0^{(r)}}$ denote the weight matrix of DAG $\mathcal{G}_0^{(r)}$. Let $\boldsymbol{\epsilon}^{(r)}$ denote a vector of errors, each of which is an independent Gaussian variable with variance independently drawn from a Uniform($[0.5,1]$) distribution. The weight matrix $B_{\mathcal{G}_0^{(r)}}$ and the random error vector $\boldsymbol{\epsilon}^{(r)}$ define the distribution of $\mathbf{X}^{(r)}$ via

$$\mathbf{X}^{(r)} = B^T_{\mathcal{G}_0^{(r)}} \mathbf{X}^{(r)} + \boldsymbol{\epsilon}^{(r)}. \tag{5.14}$$

For each r, we generate a sample $D_{n;r}$ of size n via Eq (5.14) and standardize all the variables. We then randomly select a pair of mediators, $\{X_{i_r}, X_{j_r}\}$, while restricting the pair to belong to the $\mathcal{G}_{q_r}^{\prime(r)}$ disconnected block, which contains true mediators, to avoid getting too many 0 values. We then apply the joint-MIDA method to each $D_{n,r}$ to estimate both $\eta_{i_r,j_r}(\mathcal{C}_0^{\prime\,(r)})$ and $\eta_{i_r}^{i_r,j_r}(\mathcal{C}_0^{\prime\,(r)})$. The estimated CPDAG $\hat{\mathcal{C}}_0^{\prime(r)}$ is learned via the (AR)GES algorithm Nandy *et al.* (2018b) with tuning parameter = 0.01.

We let $r = 5000$ and let $n = 50, 100, 200, 400, 800, 1600, 3200$.

We summarize the simulation results in Figure 5.3. The left column contains the log-log plots of the mean absolute error (MAE) and sample size n when the magnitudes of the true statistics are small, i.e., when $|\eta_{i_r}^{i_r,j_r}(\mathcal{C}_0'^{(r)})| < 0.1$ and $|\eta_{i_r,j_r}(\mathcal{C}_0'^{(r)})| < 0.1$ for the top and the bottom row, respectively. The right column contains the log-log plots of mean absolute relative error (MARE) and sample size n for the large magnitudes case. We also include the results when assuming that the mediator variables have no parents other than the known treatment and confounders, which is denoted by an "empty graph". The estimation errors of both individual and marginal mediation effect decay with n exponentially for both small and large effects when we apply joint-MIDA with estimated $\hat{\mathcal{C}}_0$ and show the consistency of the estimators.

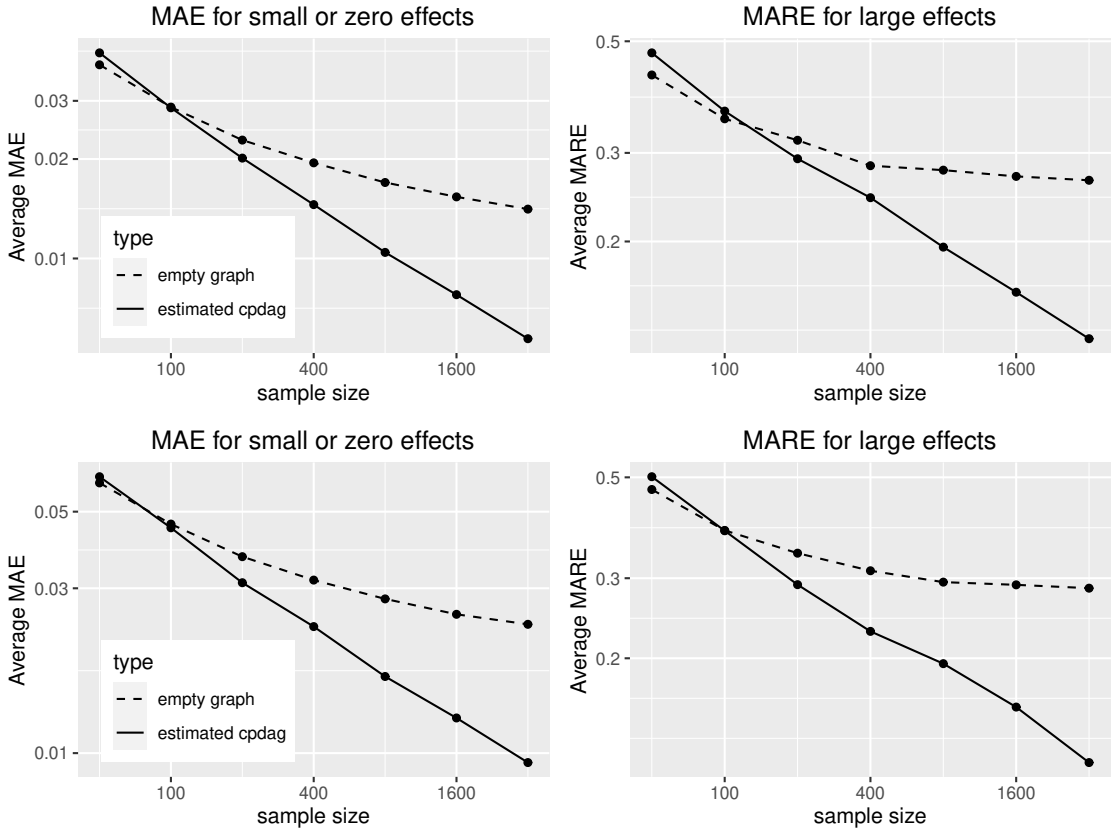


FIGURE 5.3. Log-log plot of the mean absolute error (MAE) and sample size n for zero or small joint mediation effect (left), and related mean absolute error (RMAE) and sample size n for large joint mediation effects (right) for estimating the $\eta_{i_r}^{(i_r,j_r)}(\mathcal{C}_0'^{(r)})$ (upper) and $\eta_{i_r,j_r}(\mathcal{C}_0'^{(r)})$ (lower).

5.4.2. Simulations Experiment 2. In this section, we perform simulations to illustrate the empirical performance of joint-MIDA in identifying mediator pairs with true large marginal mediation effects and true large component-wise mediation effects in a double intervention mechanism. We adopt the same setting as in Nandy *et al.* (2016), and include it here for completeness.

We will start with an empty DAG $\mathcal{G}_0'$ with presumably $p-2$ vertices, each with an expected neighbor size ens. Then we choose a number of p^* $\sim$Binomial($p-2$,q_1) vertices randomly from vertices set $\{1,\ldots,p-2\}$, denoted by $X_{i_1},\ldots,X_{i_{p^*}}$. We assume $i_1 < \cdots < i_{p^*}$. For each $z \in 1,\ldots,p^*$, we randomly select Binomial($p-2-i_z, q_2$) vertices from the set $X_{i_z+1},\ldots,X_{p-2}$, and add an directed edge from X_{i_z} to each selected vertices. For a given q_1, we choose $q_2 = ens/((p-2-1)q_1)$ to ensure the expected neighborhood size is ens. Small values of q_1 lead to hub-structure, which commonly appears in biological networks Barabasi and Oltvai (2004). We then obtain a graph $\mathcal{G}_0$ from each $\mathcal{G}_0'$ by randomly adding direct edges from treatment vertex X_1 to each vertex in $\mathcal{G}_0'$ with probability 0.15 and direct edges from each vertex in $\mathcal{G}_0'$ to the response vertex X_p with probability 0.1. All edge weights are independently sampled from a Uniform($[-1,-0.5]\cup[0.5,1]$) distribution.

We simulate n_{graph} DAGs, $\{{\mathcal{G}_0}^{(m)}\}_{m=1}^{n_{graph}}$, each consisting of p vertices, from the setting described above.

For each $m = 1,\ldots,n_{graph}$, we generate $\mathbf{X}^{(m)}$ from

$$\mathbf{X}^{(m)} = B^T_{\mathcal{G}_0^{(m)}}\mathbf{X}^{(m)} + \boldsymbol{\epsilon}^{(m)}, \tag{5.15}$$

where $B_{\mathcal{G}_0^{(m)}}$ denotes the weight matrix of $\mathcal{G}_0^{(m)}$ and $\boldsymbol{\epsilon}^{(m)}$ denotes independent Gaussian error vector where each error variance are independently drawn from Uniform($[0.5,1]$). We standardize all p variables to have variance = 1. In this simulation setting, we let $q_1 = 0.2$, $n_{graph} = 5$, $p = 402$, and $ens = 3.5$.

For each of $\mathbf{X}^{(m)}$, we randomly draw p pairs of vertices, $\{(X_{i_{pr}^m}, X_{j_{pr}^m})\}_{pr=1}^{p}$, from $\{X_2,\ldots,X_{p-1}\}$. We then generate n_{itr} random replicates $D_{n;r}^{(m)}$ of size n from the joint distributions of $\mathbf{X}^{(m)}$ for $r = 1,\ldots,n_{itr}$. We aim to estimate the component-wise mediation effect ,$\eta^{(m)}_{i_{pr}^m,j_{pr}^m}$, and the marginal mediation effect, $\eta^{(i_{pr}^m,j_{pr}^m)^{(m)}}_{i_{pr}^m}$, for each of the p mediator pairs by applying the joint-MIDA method described in Algorithm 4 to each simulated data $D_{n;r}^{(m)}$ for $r = 1,\ldots,n_{itr}$. The graph used in the Algorithm 4 step 3 is chosen as one of the following: an estimated CPDAG $\hat{\mathcal{C}}_0'$ from simulated data

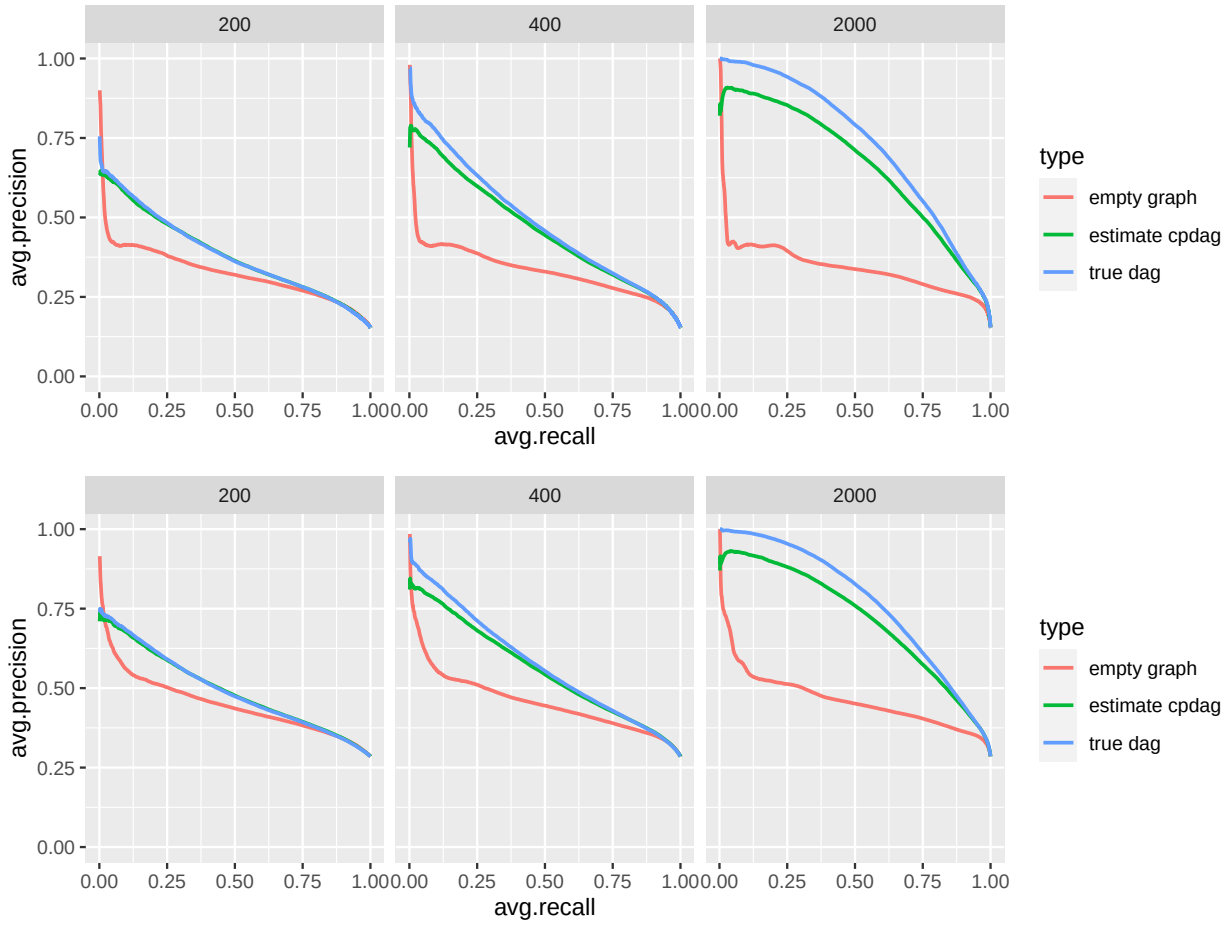


FIGURE 5.4. Precision-Recall curves for the identifying target set $\{\bigcup_{m=1}^{n_{graph}} \eta^{(m)}_{i^m_{pr},j^m_{pr}} : |\eta^{(m)}_{i^m_{pr},j^m_{pr}}| > 0;\ pr = 1\ldots,p\}$ (top) and $\{\bigcup_{m=1}^{n_{graph}} \eta^{(i^m_{pr},j^m_{pr})^{(m)}}_{i^m_{pr}} : |\eta^{(i^m_{pr},j^m_{pr})^{(m)}}_{i^m_{pr}}| > 0;\ pr = 1\ldots,p\}$ (bottom) for all $n_{graph}p$ pairs of vertices in $\bigcup_{m=1}^{n_{graph}}\{(X_{i^m_{pr}}, X_{j^m_{pr}})\}_{pr=1}^{p}$ with different sample sizes.

$D^{(m)}_{n;r}$ by applying (AR)GES algorithm Nandy *et al.* (2018b) with tuning parameter = 0.05 to the residuals, the true DAG $\mathcal{G}_0'^{(m)}$, or an empty graph.

Lastly, we calculate the precision and recall over the estimations of $\{\bigcup_{m=1}^{n_{graph}} \eta^{(m)}_{i^m_{pr},j^m_{pr}} : pr = 1\ldots,p\}$ as well as $\{\bigcup_{m=1}^{n_{graph}} \eta^{(i^m_{pr},j^m_{pr})^{(m)}}_{i^m_{pr}} : pr = 1\ldots,p\}$. The target sets are $\{\bigcup_{m=1}^{n_{graph}} \eta^{(m)}_{i^m_{pr},j^m_{pr}} : |\eta^{(m)}_{i^m_{pr},j^m_{pr}}| > 0;\ pr = 1\ldots,p\}$ and $\{\bigcup_{m=1}^{n_{graph}} \eta^{(i^m_{pr},j^m_{pr})^{(m)}}_{i^m_{pr}} : |\eta^{(i^m_{pr},j^m_{pr})^{(m)}}_{i^m_{pr}}| > 0;\ pr = 1\ldots,p\}$, respectively. We let $n = \{200, 400, 2000\}$ and $n_{itr} = 200$. We plot the averaged Precision-Recall curves over 200 replicates in Figure 5.4. As we expect, using the true DAG $\mathcal{G}_0'$ in joint-MIDA results in the best

performance. Applying the estimated CPDAG outperforms using the empty graph, and the gaining of using an estimated CPDAG increases with a larger sample size n.

5.5. Application

5.5.1. Dataset and Related Studies. In this section, we demonstrate the performance of joint-MIDA on identifying DNA methylation loci that mediate childhood trauma's causal effect on cortisol stress reactivity. We work on a dataset that consists of 385,882 DNA methylation loci information, the increased area under the curve (iAUC) in cortisol level after a stress test, a one-dimensional score on the childhood trauma questionnaire, and several covariates such as gender and age for 85 individuals. [1]

A number of mediation analyses have been previously devoted to this dataset to study how DNA methylation plays a role in the regulation of human stress reactivity Houtepen *et al.* (2016); van Kesteren and Oberski (2019); Guo *et al.* (2022). Among those works, Houtepen *et al.* (2016) conducted a separate linear mediation analysis on each individual locus and identified three mediating loci that may regulate the causal effect of childhood trauma on cortisol stress reactivity. van Kesteren and Oberski (2019) addressed the dependency among loci with coordinate-wise mediation filter . They identified three mediating loci that have no overlap with the three identified by Houtepen et al. Guo *et al.* (2022) identified loci that have a large causal effect on iAUC in cortisol stress level via partial penalized least square regression approach proposed in Guo *et al.* (2023). They further compared their results with that of Houtepen *et al.* (2016) and van Kesteren and Oberski (2019) by conducting hypothesis tests on the causal effects of loci with different sets of loci included in the partial penalized least squares regression Guo *et al.* (2022). They aimed to assess how the significance levels vary with the inclusion of different loci sets. They observed that the causal effect of cg27512205 (KITLG gene), cg12500973(HNRNPF gene), and cg16657538 (ZSCAN30 gene) could change from significant to insignificant after the inclusion of more mediators. They explained this change in significance by the association among mediators, such as pairwise and partial correlation between mediators.

[1]The full dataset can be obtained from ArrayExpress: `https://www.ebi.ac.uk/biostudies/arrayexpress/studies/E-GEOD-77445`.

Motivated by those different results from previous studies, we apply the joint-MIDA method to the same dataset to estimate the component-wise mediation effects of each of the candidate loci. We also perform a non-parametric hypothesis test on the candidate set to identify those loci that play significant mediating roles when multiple loci in a joint intervention mechanism.

5.5.2. Analysis. The setup of the experiment is as follows. We consider age, sex, B cell proportion, CD4 T cell proportion, CD8 T cell proportion, Monocytes cell proportion, Granulocytes cell proportion and Natural Killercell proportion as confounding variables, denoted as $X_1, \ldots, X_8$. The treatment $(X_t, t = 9)$ is a one-dimensional score on a childhood trauma questionnaire test. The 385,882 DNA methylation loci in the blood are considered potential mediators. The iAUC in cortisol level after a stress test is the response variable. Our analysis is organized into 3 steps described as follows.

Step 1: Pre-screening. We first apply a pre-screening procedure similar to the one conducted by van Kesteren and Oberski (2019) and Guo *et al.* (2022). The pre-screening retains the top 1000 potential mediators based on the rank of their absolute correlation with the treatment variable (X_9). We include the additional 11 methylation loci identified by the previous work van Kesteren and Oberski (2019); Guo *et al.* (2022); Houtepen *et al.* (2016), resulting in a total of 1011 potential mediators, denoted by $X_{10}, \ldots, X_{1020}$. The response variable, the increased area under the curve (iAUC) in cortisol level after a stress test, is then denoted as X_{1021}.

Step 2: Identify active mediators. In this step, we aim to identify a candidate set that contains active mediators that have a nonzero marginal mediation effect when no other mediators are considered jointly. We apply the MIDA hypothesis test Chakrabortty *et al.* (2018) on the mediation effect η_j of each of 1,011 DNA methylation loci individually. More specifically, we conduct the following hypothesis tests:

$$H_{0j} : \eta_j = 0 \quad \text{vs} \quad H_{aj} : \eta_j \neq 0, \quad j = 10, \ldots, 1020. \tag{5.16}$$

We set the significant level at 0.05 as used in Chakrabortty *et al.* (2018). In this experiment, 19 mediators are considered having significant mediation effects: cg21875401(X_{49}), cg04602284(X_{57}), cg00097357(X_{122}), cg02032299(X_{207}), cg25829425(X_{247}), cg16247183(X_{323}), cg06613263(X_{425}), cg14693391(X_{481}), cg07635747(X_{499}), cg17710779(X_{512}), cg05781698(X_{533}), cg10631546(X_{572}),

cg04485065(X_{649}), cg04295094(X_{662}), cg08158408(X_{699}), cg18703581(X_{782}), cg18449062(X_{846}), g22178613(X_{868}) and cg06688323(X_{992}).

Step 3: Multiple non-parametric hypothesis tests. We then perform hypothesis tests on the component-wise mediation effect of all active mediators. For a better comparison, we consider the 19 mediators identified by MIDA together with the 11 mediators identified in the previous works, resulting in a set of 30 mediators. We denote the set of indices of those mediators as $\mathbf{M}$.

We perform the Bias-Corrected and Accelerated (BCa) non-parametric bootstrap hypothesis tests Efron and Tibshirani (1993), on the component-wise mediation effect of each of the mediators in the candidate set, i.e.,

$$H_{0j} : \eta_j^{(\mathbf{M})} = 0 \quad \text{vs} \quad H_{aj} : \eta_j^{(\mathbf{M})} \neq 0, \quad j \in \mathbf{M}. \tag{5.17}$$

Let $\hat{\eta}_j^{(\mathbf{M})^{(b)}}$ and $\hat{\eta}_{j_{(i)}}^{(\mathbf{M})}$ denote the bootstrap estimator and leave observation i out estimator of $\eta_j^{(\mathbf{M})}$, respectively. The two-sided p-value of BCa bootstrap test H_{0j} vs H_{aj} in Equation 5.17 is

$$\text{p-value} = \begin{cases} 2\Phi(z) & \text{if } z > 0, \\ 2(1 - \Phi(z)) & \text{otherwise,} \end{cases} \tag{5.18}$$

where

$$z = \frac{w_0 - \hat{z}_0}{1 + \hat{a}(w_0 - \hat{z}_0)} - \hat{z}_0,$$

$$w_0 = \Phi^{-1}\left(\frac{\#(\hat{\eta}_j^{(\mathbf{M})^{(b)}} < 0)}{B}\right),$$

$$\hat{z}_0 = \Phi^{-1}\left(\frac{\#(\hat{\eta}_j^{(\mathbf{M})^{(b)}} < \hat{\eta}_j^{(\mathbf{M})})}{B}\right),$$

$$\hat{a} = \frac{\sum_{i=1}^{n}(\hat{\eta}_{j_{(\cdot)}}^{(\mathbf{M})} - \hat{\eta}_{j_{(i)}}^{(\mathbf{M})})^3}{6\{\sum_{i=1}^{n}(\hat{\eta}_{j_{(\cdot)}}^{(\mathbf{M})} - \hat{\eta}_{j_{(i)}}^{(\mathbf{M})})^2\}^{3/2}},$$

$$\hat{\eta}_{j_{(\cdot)}}^{(\mathbf{M})} = (\sum_{i=1}^{n} \hat{\eta}_{j_{(i)}}^{(\mathbf{M})})/n.$$

Lastly, we apply Bonferroni correction on the p-values to select $\eta_j^{(\mathbf{M})}$'s that are significant.

We summarize the non-parametric inference procedure in the following Algorithm 5.

Algorithm 5 Multiple hypothesis tests on component-wise mediation effect

Require: n i.i.d observations of $\mathbf{X}$ (data); candidate mediator index $\mathbf{M}$; significant level α.
Ensure: X_j's with significant component-wise mediation effects $\eta_j^{(\mathbf{M})}$.

1: Perform step 1 and 2 from Algorithm 4 to estimate $\hat{\mathcal{C}}_0$. Calculate $\hat{\eta}_j^{(\mathbf{M})}(\hat{\mathcal{C}}_0')$, for each $j \in \mathbf{M}$, by applying the RRC method to $\mathbf{X}$.
3: For each b $\in 1, \ldots, B$, independently sample n observations from $\mathbf{X}$ with replacement, denoted as $\mathbf{X}^{(b)}$. Calculate $\hat{\eta}_j{}^{(\mathbf{M})^{(b)}}(\hat{\mathcal{C}}_0')$ by applying the RRC method to $\mathbf{X}^{(b)}$.
4: For each $1 \in 1, \ldots, n$, calculate $\hat{\eta}_{j_{(i)}}^{(\mathbf{M})}(\hat{\mathcal{C}}_0')$ by applying the RRC method to $\mathbf{X}_{(i)}$, the sample with observation i left out.
5: Calculate the BCa two-sided p-value for each H_{0j} as specified in Equation 5.18.
6: **return** X_j's with p-value $\leq \frac{\alpha}{|\mathbf{M}|}$ as specified in Equation 5.17.

5.5.3. Results. At a significant level $\alpha = 0.05$ with Bonferroni correction, six loci, cg06613263(X_{425}), cg07635747(X_{499}), cg05608730(X_{1011}), cg02309301(X_{1013}), cg25626453(X_{1016}), cg13136721(X_{1017}) are considered to have significant component-wise mediation effects. Table 5.2 presents the estimated causal effects from treatment X_t to mediators X_j, $\hat{\theta}_{tj}$, the estimated joint causal effects from mediators X_j to response variable X_p, $\hat{\theta}_{jp}^{(\mathbf{M})}$, the estimated mediation effects $\hat{\eta}_j^{(\mathbf{M})}$ and p-values for those significant loci. We annotate the locations of the selected loci with the *BioConductor: FDb.InfiniumMethylation.hg18* R package Tim Triche (2014) to obtain the nearest protein-coding gene. Table 5.3 lists the 6 DNA methylation loci with their gene location as well as some field knowledge of these genes from existing research. Among the six loci, the two new loci, cg06613263 (X_{425}) and cg07635747 (X_{499}), are located in NR3C1 gene and CDH2 gene, respectively.

The NR3C1 gene has been widely researched in many preclinical studies. Perroud *et al.* (2011) showed that childhood maltreatment, through epigenetic modification of the glucocorticoid receptor gene (NR3C1), influences the hypothalamic-pituitary-adrenal axis (HPA axis). Multiple researchers have shown that the NR3C1 gene is associated with blunted cortisol reactivity level Alexander *et al.* (2018); Bakusic *et al.* (2020). Therefore, gene NR3C1 possesses strong neurological evidence to be the true mediator between childhood trauma and cortisol change. Experiments conducted on male pups and adult mice revealed that CDH2 could be affected by early life stress through gene ontology analysis Reemst *et al.* (2022). Also, Ferreira *et al.* (2023) had shown that CDH2 is a candidate gene for hypopituitarism, which could cause insufficient ACTH production or secretion and lead to

TABLE 5.2. Estimated total causal effects, estimated joint mediation effects, and p-values of significant mediators.

Locus and variable	$\hat{\theta}_{tj}$	$\hat{\theta}_{jp}^{(\mathbf{M})}$	$\hat{\eta}_{j}^{(\mathbf{M})}$	p-value
cg06613263(X_{425})	-0.009	-219.347	2.057	0.001026
cg07635747(X_{499})	0.004	-478.282	-2.12	6e-04
cg05608730(X_{1011})	0.007	-430.764	-3.08	7.6e-05
cg02309301(X_{1013})	-0.012	177.274	-2.064	0.000408
cg25626453(X_{1016})	-0.012	320.87	-3.917	0.000255
cg13136721(X_{1017})	0.012	162.28	1.954	0.000296

a subsequent abnormal cortisol level Bradley and Wass (2004). Hence, it is reasonable to suggest that the identified loci, or the genes they harbor, play a role in regulating neurological pathways and partially explain how cortisol stress reactivity could be affected by childhood trauma.

We include the comprehensive hypothesis test results in Table C.1 in the Appendix.

Locus and variable	Gene	Field knowledge from literature
cg06613263(X_{425})	NR3C1	Encoded the glucocorticoid receptor and regulated the cellular response to stress hormones (Liu and Nusslock, 2018).
cg07635747(X_{499})	CDH2	Involved in the development of the nervous system by mediating cell-cell adhesion(László and Lele, 2022).
cg05608730(X_{1011})	C1QTNF2	Involved in regulation of lipid metabolism and adipocyte differentiation (Park *et al.*, 2009).
cg02309301 (X_{1013})	ARGLU1	Involved in modulating glucocorticoid signaling in the central nervous system (Magomedova *et al.*, 2019).
cg25626453(X_{1016})	PRRC2A	Associated with the age-at-onset of diabetes.
cg13136721(X_{1017})	RPTOR	Serve as a scaffold protein that integrates upstream signals to control cell growth, metabolism, and autophagy (Saxton and Sabatini, 2017).

TABLE 5.3. Annotation and summary of selected mediator loci

5.5.4. Comparison with Guo *et al.* (2022). In this section, we further compare our method with the partial penalized least square regression approach used in Guo *et al.* (2022). We apply the above analysis in Section 5.5.2 on the same set of potential mediators used by Guo *et al.* (2022). More specifically, we use $\mathbf{M}_{(1)}$, $\mathbf{M}_{(2)}$, $\mathbf{M}_{(3)}$, $\mathbf{M}_{(4)}$ to denote if a locus is identified by in Houtepen *et al.* (2016), van Kesteren and Oberski (2019), Guo *et al.* (2022) or through MIDA hypothesis test.

We let $\mathbf{M'} = \cup_{i=1}^{3} \mathbf{M}_{(i)}$, and $\mathbf{M'}$ has a size equals to 11. We then apply the non-parametric bootstrap hypothesis test introduced in Section 5.5.2 Step 3 to each of $\eta_j^{(\mathbf{M'})}, j \in \mathbf{M'}$, and present our results in Table 5.4. Note that the component-wise mediation effects of variable X_{1010} (cg27512205), X_{1014} (cg12500973) and X_{1015} (cg16657538) have p-values greater than 0.05.

On the other hand, Guo et al. Guo *et al.* (2022) considered the following linear model,

$$\begin{aligned} X_p &= \boldsymbol{\alpha_0}^T \mathbf{X}_{\mathbf{M'}} + \alpha_1 X_t + \boldsymbol{\alpha}_2^T \mathbf{X}_{1,\ldots,t-1} + \epsilon_1, \\ \mathbf{X}_{\mathbf{M'}} &= \boldsymbol{\gamma_1} X_t + \boldsymbol{\Gamma}_2^T \mathbf{X}_{1,\ldots,t-1} + \boldsymbol{\epsilon}_2, \end{aligned} \tag{5.19}$$

where X_t, X_p, $\mathbf{X}_{1,\ldots,t-1}$ and $\mathbf{X}_{\mathbf{M'}}$ denote the treatment variable, response variable, the vector of confounders and the vector of potential mediators, respectively. Under this setting, $\boldsymbol{\alpha_0}$ is a vector of size 1×11; $\boldsymbol{\gamma_1}$ is a vector of size 1×11; $\boldsymbol{\Gamma}_2$ is a matrix of size $t - 1 \times 11$.

The authors are interested in a locus that both have a significant causal effect on response variable X_p in $\boldsymbol{\alpha_0}$ and corresponds to a significant causal effect in $\boldsymbol{\gamma_1}$, which is the causal effect of X_t on that mediator locus.

The authors, therefore, conduct hypothesis tests on each component in $\boldsymbol{\alpha_0}$ based on the partial penalized least square regression model Guo *et al.* (2023). Their results show that the $\boldsymbol{\alpha}_{0i}$'s corresponding to X_{1010}, X_{1014} and X_{1015} have p-value greater than 0.05. They also conduct a hypothesis test on each component in $\boldsymbol{\gamma_1}$ in linear regression and conclude that all components in $\boldsymbol{\gamma_1}$ are significant effects at a significance level 0.05 except the one corresponding to X_{1017} Guo *et al.* (2022).

In summary, our method reaches a similar conclusion as Guo *et al.* (2022) when the same mediators are considered.

TABLE 5.4. Estimated causal effect $\hat{\theta}_{tj}$, total joint effect $\hat{\theta}_{jp}^{(\mathbf{M}')}$, component-wise joint mediation effect $\hat{\eta}_{j}^{(\mathbf{M}')}$, with their bootstrap standard deviation in parenthesis. At a significance level of 0.05, a locus with a significant $\hat{\eta}^{(\mathbf{M}')}j$ after Bonferroni correction is annotated with "**"; a locus with a significant $\hat{\eta}^{(\mathbf{M}')}j$ before Bonferroni correction is annotated with "*".

Locus and variable	$\hat{\theta}_{tj}$	$\hat{\theta}_{jp}^{(\mathbf{M}')}$	$\hat{\eta}_{j}^{(\mathbf{M}')}$	p-value	sig	source
cg27512205(X_{1010})	0.005(0.003)	-106.734(91.81)	-0.517(1.229)	0.310232		
cg05608730(X_{1011})	0.007(0.003)	-348.963(132.804)	-2.495(1.406)	0.007657	*	$\mathbf{M}_{(1)}$
cg26179948(X_{1012})	0.006(0.003)	-443.553(133.565)	-2.802(1.318)	0.005514	*	
cg02309301(X_{1013})	-0.012(0.004)	243.533(105.05)	-2.835(2.051)	0.001071	**	
cg12500973(X_{1014})	-0.011(0.004)	34.424(108.424)	-0.386(1.234)	0.704202		
cg16657538(X_{1015})	0.023(0.006)	65.18(61.006)	1.484(1.466)	0.072158		$\mathbf{M}_{(2)}$
cg25626453(X_{1016})	-0.012(0.009)	417.014(47.109)	-5.091(1.231)	0.0025	**	
cg13136721(X_{1017})	0.012(0.003)	214.857(114.752)	2.586(0.985)	0.001504	**	
cg19230917(X_{1018})	-0.006(0.005)	216.022(97.37)	-1.373(2.371)	0.017501	*	
cg06422529(X_{1019})	-0.008(0.003)	465.654(136.805)	-3.725(1.394)	0.016357	*	$\mathbf{M}_{(3)}$
cg03199124(X_{1020})	-0.006(0.003)	387.076(148.27)	-2.363(0.844)	0.011278	*	

CHAPTER 6

Conclusion

This dissertation studied two important aspects of statistical machine learning: uncertainty quantification and algorithmic fairness. We developed novel solutions to mitigate the accuracy disparity under selection bias and to perform inferenfce on 2D object boundaries and high-dimensional mediation effects. All of our proposed solutions demonstrated their effectiveness and practical value through simulation experiments and real-data examples.

It is worth noting that some of our methods can be further generalized to new scenarios. Firstly, the weighting techniques proposed in Chapter 2 to correct selection bias could be applied to classification problems, further expanding its application scenarios. We hope our contributions could stimulate more interest in uncertainty quantification and algorithmic fairness in the statistical machine learning field. Moreover, we use Fourier Descriptor to parameterize the object boundary in Chapter 4, which focuses more on capturing the global information (centroid, size, eccentricity) of the shape. Our framework can be incorporated with wavelet descriptor or B-spline to better capture the local variation of the curve. On the other hand, other UQ methodologies, such as Fiducial inference (Han and Lee, 2022), can be adopted to quantify the uncertainty of the FD parameters.

APPENDIX A

Appendix to Chapter 2

A.1. Derivation of Balanced Weights W in Section 2.2.6

Let K be the reproducing kernel of $\mathcal{H}$, and $r(g, U, a)$ lies in a finite-dimensional subspace span $\{K(U_j, \cdot) : j = 1, ..., n; A_j = a\}$. We can then write $r(g, U, a) = \sum_{j=1, A_j=a}^{n} \alpha_{a,j} K(U_j, \cdot)$. Let M denote the $N_a \times N_a$ Gram matrix, i.e., $M_{ij} = K(U_{a_i}, U_{a_j}), A_{a_i} = A_{a_j} = a$. We have

$$H_n(W_a, \sum_{j=1, A_j=a}^{n} \alpha_{a,j} K(U_j, \cdot), a) = \frac{1}{n^2} \alpha_a^T M Q(W_a, a) M \alpha_a,$$

where

$$Q(W_a, a) = q(W_a, a) q(W_a, a)^T \quad \text{with} \quad q(W_a, a) = \left((S_{a_1} W_{a_1} - \frac{n}{N_a}), ..., (S_{a_{N_a}} W_{a_{N_a}} - \frac{n}{N_a}) \right)^T.$$

Let the eigenvalue decomposition of M be

$$M = \begin{pmatrix} P_1 & P_2 \end{pmatrix} \begin{pmatrix} Q_1 & 0 \\ 0 & Q_2 \end{pmatrix} \begin{pmatrix} P_1^T \\ P_2^T \end{pmatrix}.$$

The weights $\boldsymbol{W}_a$ for each $a = \{0, 1\}$ is the solution of the following problem:

$$\min_{W \geq \frac{n}{N_a}} \left[\sigma_{\max} \left\{ \frac{1}{n} P_1 Q(W_a, a) P_1 - n\lambda_1 Q_1^{-1} \right\} + \lambda_2 V_n(W) \right].$$

A.2. Proof of Proposition 5

PROOF. Let $\boldsymbol{L}_{(-i)}$ and $\boldsymbol{Y}_{(-i)}$ denote the covariate matrix and outcome vector with the i-th observation $(L_i, \boldsymbol{Y}_i)$ left out. Let $\boldsymbol{D}_{(-i)}$ and $\boldsymbol{W}_{(-i)}$ denote the corresponding $(m-1) \times (m-1)$ weight matrices. Note that primal feasibility condition of the solution of Equaiton 2.32 ensures

$$(\boldsymbol{L}_{(-i)} \hat{\boldsymbol{\beta}}_{(-i)} - \boldsymbol{Y}_{(-i)})^T \boldsymbol{D}_{(i)} (\boldsymbol{L}_{(-i)} \hat{\boldsymbol{\beta}}_{(-i)} - \boldsymbol{Y}_{(-i)}) + \delta = 0 \tag{A.1}$$

To prove the position, it suffices to show that $\hat{\boldsymbol{\beta}}_{(-i)}$ is also the optimial solution of the following constraint optimization,

$$\min_{\boldsymbol{\beta}}(\boldsymbol{L}\boldsymbol{\beta}-\boldsymbol{Y}_{(i)})^T\boldsymbol{W}(\boldsymbol{L}\boldsymbol{\beta}-\boldsymbol{Y}_{(i)})+\eta\boldsymbol{\beta}^T\boldsymbol{\beta}:(\boldsymbol{L}\boldsymbol{\beta}-\boldsymbol{Y}_{(i)})^T\boldsymbol{D}(\boldsymbol{L}\boldsymbol{\beta}-\boldsymbol{Y}_{(i)})=-\delta \tag{A.2}$$

Firstly, note the lagrangian form of above optimization is

$$\tilde{L}(\boldsymbol{\beta},\lambda)=\boldsymbol{\beta}^T\boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{L}\boldsymbol{\beta}-2\boldsymbol{Y}_{(i)}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{L}\boldsymbol{\beta}+\eta\boldsymbol{\beta}^T\boldsymbol{\beta}+\boldsymbol{Y}_{(i)}^T\boldsymbol{W}\boldsymbol{Y}_{(i)}+\lambda(-\boldsymbol{Y}_{(i)}^T\boldsymbol{D}\boldsymbol{Y}_{(i)}-\delta). \tag{A.3}$$

We can write the first order stationary point of (A.3) as

$$\tilde{\boldsymbol{\beta}}(\lambda)=(\boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{L}+\eta\mathbf{I})^{-1}\boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{Y}_{(i)}.$$

As $\hat{Y}_{i(-i)}=L_i^T\hat{\boldsymbol{\beta}}_{(-i)}$, by definition,

$$\hat{\boldsymbol{\beta}}_{(-i)}=\operatorname*{argmin}_{\boldsymbol{\beta}}\sum_{j\neq i}(\boldsymbol{W}_{jj}-\lambda_{(-i)}\boldsymbol{D}_{jj})(L_j^T\boldsymbol{\beta}-Y_j)^2+(\boldsymbol{W}_{ii}-\lambda_{(-i)}\boldsymbol{D}_{ii})(L_i^T\boldsymbol{\beta}-\hat{Y}_{i(-i)})^2+\eta\boldsymbol{\beta}^T\boldsymbol{\beta}+\lambda_{(-i)}(-\delta),$$

it follows that $\hat{\boldsymbol{\beta}}_{(-i)}$ is the minimizer of $\tilde{L}(\boldsymbol{\beta},\lambda)$ in (A.3) at $\lambda=\lambda_{(-i)}$, and we have $\hat{\boldsymbol{\beta}}_{(-i)}=\tilde{\boldsymbol{\beta}}(\lambda_{(-i)})$. All left is to show $\max_\lambda\min_{\boldsymbol{\beta}}\tilde{L}(\boldsymbol{\beta},\lambda)$ is attained at $\lambda=\lambda_{(-i)}$.

The dual problem of (A.2) is

$$\begin{aligned}\tilde{h}(\lambda)=&-(\boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{Y}_{(i)})^T(\boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{L}+\eta\mathbf{I})^-(\boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{Y}_{(i)})+\lambda(-\boldsymbol{Y}_{(i)}^T\boldsymbol{D}\boldsymbol{Y}_{(i)}-\delta).\\ &\boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{L}+\eta\mathbf{I}\succeq 0,\\ &\boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{Y}_{(i)}\in\text{Range}(\boldsymbol{L}^T(\boldsymbol{W}-\lambda\boldsymbol{D})\boldsymbol{L}+\eta\mathbf{I}).\end{aligned} \tag{A.4}$$

According to Proposition 3, $\tilde{\boldsymbol{\beta}}(\tilde{\lambda})$ is the global optimal solution of (A.2) if $\tilde{\lambda}$ is the optimal solution of $\tilde{h}(\lambda)$ for some $\tilde{\lambda}\in(\underline{\lambda},\overline{\lambda})$.

The unique optimizer of (A.4) solves the following equation

$$\nabla\tilde{h}(\lambda)=-\delta-\big(\boldsymbol{L}\tilde{\boldsymbol{\beta}}(\lambda)-\boldsymbol{Y}_{(i)}\big)^T\boldsymbol{D}\big(\boldsymbol{L}\tilde{\boldsymbol{\beta}}(\lambda)-\boldsymbol{Y}_{(i)}\big)=0. \tag{A.5}$$

Since $\lambda_{(-i)} \in (\underline{\lambda}, \overline{\lambda})$ by assumption, function in (A.4) takes its optimal values at $\lambda = \lambda_{(-i)}$ as

$$\big(\boldsymbol{L}\tilde{\boldsymbol{\beta}}(\lambda_{(-i)}) - \boldsymbol{Y}_{(i)}\big)^T \boldsymbol{D}\big(\boldsymbol{L}\tilde{\boldsymbol{\beta}}(\lambda_{(-i)}) - \boldsymbol{Y}_{(i)}\big) \tag{A.6}$$

$$=\big(\boldsymbol{L}\hat{\boldsymbol{\beta}}_{(-i)} - \boldsymbol{Y}_{(i)}\big)^T \boldsymbol{D}\big(\boldsymbol{L}\hat{\boldsymbol{\beta}}_{(-i)} - \boldsymbol{Y}_{(i)}\big) \tag{A.7}$$

$$=(\boldsymbol{L}_{(-i)}\hat{\boldsymbol{\beta}}_{(-i)} - \boldsymbol{Y}_{(-i)})^T \boldsymbol{D}_{(i)}(\boldsymbol{L}_{(-i)}\hat{\boldsymbol{\beta}}_{(-i)} - \boldsymbol{Y}_{(-i)}) \tag{A.8}$$

$$= -\delta, \tag{A.9}$$

where the first quality follows from $\hat{\boldsymbol{\beta}}_{(-i)} = \tilde{\boldsymbol{\beta}}(\lambda_{(-i)})$, and second equality follows from the fact $\hat{Y}_{i(-i)} = L_i^T \hat{\boldsymbol{\beta}}_{(-i)}$, and the third equality is the primal feaibility condition in (A.1). □

A.3. Newton's Method Implementation

Algorithm 6 Optimization procedure to solve $\lambda^* = \operatorname{argmax} h_C(\lambda)$.

procedure NEWTON'S METHOD
  $\lambda_0 = 0$
  while epoch ¡ maxepoch; epoch ++ **do**
    $\Delta\lambda = \frac{\nabla h_C(\lambda_i)}{\nabla^2 h_C(\lambda_i)}$
    $\nabla h_C(\lambda_i) = C + 2\boldsymbol{Y}^T\boldsymbol{D}\boldsymbol{L}\boldsymbol{\beta}(\lambda_i) - \boldsymbol{\beta}(\lambda_i)^T(\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L})\boldsymbol{\beta}(\lambda_i)$,
$\nabla^2 h_C(\lambda_i) = -2(\boldsymbol{\beta}(\lambda_i)\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L} - \boldsymbol{L}^T\boldsymbol{D}\boldsymbol{Y})^T(\boldsymbol{L}^T\boldsymbol{W}\boldsymbol{L} - \lambda_i\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L} + \eta\mathbf{I})^{-1}(\boldsymbol{\beta}(\lambda_i)\boldsymbol{L}^T\boldsymbol{D}\boldsymbol{L} - \boldsymbol{L}^T\boldsymbol{D}\boldsymbol{Y})$
    if $\Delta\lambda > 0$ **then**
      $\Delta\lambda = \min(\Delta\lambda, 0.95(\overline{\lambda} - \lambda_i))$
    else
      $\Delta\lambda = \max(\Delta\lambda, 0.95(\underline{\lambda} - \lambda_i))$
    $\lambda_{i+1} = \lambda_i + \Delta\lambda$; $\boldsymbol{\beta}_{i+1} = \boldsymbol{\beta}(\lambda_{i+1})$
    if
$$\max\left\{\frac{|\nabla h_C(\lambda_i)|}{|S|+1}, \frac{|f_0(\boldsymbol{\beta}_{i+1}) - h(\lambda_{i+1})|}{|f_0(\boldsymbol{\beta}_{i+1})|+1}, \left(\frac{||(\boldsymbol{L}^T(\boldsymbol{W} - \lambda_{i+1}\boldsymbol{D})\boldsymbol{L})\boldsymbol{\beta}_{i+1} - (\boldsymbol{L}^T(\boldsymbol{W} - \lambda_{i+1}\boldsymbol{D})\boldsymbol{Y})||}{||\boldsymbol{L}^T(\boldsymbol{W} - \lambda_{i+1}\boldsymbol{D})\boldsymbol{L}|| + ||\boldsymbol{L}^T(\boldsymbol{W} - \lambda_{i+1}\boldsymbol{D})\boldsymbol{Y}|| + 1}\right)^2\right\} \leq tol$$
    then
      Break.
  Return λ_{i+1}

A.4. Benchmark Experiments Setup

A.4.1. Experiment Details of Section 2.6.

A.4.1.1. *Crime Data.* Crime dataset Redmond (2009) was collected from the 1990 US Census and contains socio-economic information from 1994 communities. The crime rate of a given community is the target variable (Y), and the African American Population Ratio (AAPR) is selected as the sensitive attribute. Communities with an AAPR exceeding 50% are labeled as protected ($A = 1$), resulting in 219 protected communities and 1775 non-protected ones. As specified in Du *et al.* (2022), 15 attributes are selected as the prediction covariates set X: "racePctHisp", "agePct12t21", "agePct12t29", "agePct16t24", "agePct65up", "numbUrban", "pctUrban", "medIncome", "pctWWage", "pctWFarmSelf", "pctWInvInc","pctWSocSec", "pctWPubAsst", "pctWRetire", "medFamInc". Additionally, the following 5 attributes are included in U to avoid the multicollinearity issue of Heckman model: "population", "householdsize", "racepctblack", racePctWhite", racePctAsian". We further test the peformance of the model when "NumUnderPov" is also included in U. After removing attributes with missing values and standardizing all attributes to have zero mean and unit variance, the IMR covariate and weights are learned using the full set of U. Subsequently, a fair linear regression model is trained on X (with or without the additional IMR covariate).

A.4.1.2. *Law Data.* Law dataset Wightman (1998) was collected from the Law School Admissions Council's National Longitudinal Bar Passage Study, containing personal records of law students who eventually took the bar exam, including their LSAT scores, age, race, and more. The objective is to predict the GPA of a student based on other attributes. We follow the same preprocessing steps for sensitive attributes and covariates as outlined in Du *et al.* (2022). The race attribute (black/non-black) is chosen as the sensitive attribute, with black being considered the protected ($A = 1$) group. With a total of 20,649 records, we randomly select 3,000 records for the final dataset, including 1,000 protected and 2,000 non-protected samples. As specified in Du *et al.* (2022), 5 atrributes are included in the prediction covariates set X: "fulltime", "fam_inc", "age", "gender", "pass". In addition to X, the following 6 attributes are included in U but not X to avoid the multicollinearity issue of Heckman model: "cluster", "lsat", "ugpa", "zgpa".

A.4.2. Experiment Details of Section 2.7. All experiments in this section follow the data preprocessing and covariate sets outlined in Chi *et al.* (2021). We'll provide a brief overview of the

general setup, including the sensitive and target variables. For more detailed information and an introduction to the datasets, readers can refer to the Appendix of Chi *et al.* (2021).

A.4.2.1. *Crime data.* In this experiment, the race atrribute is used as the sensitive attribute with $A = 1$ represents when the population percentage of the white is greater or equal to 80% and 0 otherwise. Each instance has an input dimension of 96. For ridge regression, we set the tuning hyper-parameter range to be 20 numbers spaced evenly on a log scale between 10^{-9} and 10^{-3}. In Kernel ridge regression, we implement the rbf kernel using random Fourier features and choose the bandwidth through median heuristic. We select the penalizing hyper-parameter η from 20 evenly spaced numbers on a log scale between 10^{-8} and 1. The threshold δ for our methods is set to be $\{10^{-6}, 10^{-3}, 10^{-4}, 10^{-3}, 0.005, 0.01, 0.015, 0.02, 0.03\}$; the penalizing hyper-parameter λ for CENet and WassersteinNet in Chi *et al.* (2021) is set as $\{0.01, 0.1, 1.0, 5.0, 10.0, 15.0, 20.0, 35.0, 50.0, 100.0\}$.

A.4.2.2. *Law Data.* In this experiment, we use gender as the sensitive attribute and undergraduate GPA as the target variable. For ridge regression, we set the tuning hyper-parameter range to be 20 numbers spaced evenly on a log scale between 10^{-9} and 10^{-3}. In kernel ridge regression, we use the Sobolev kernel. We select the penalizing hyper-parameter η from 20 evenly spaced numbers on a log scale between 10^{-8} and 1. The MSE disparity threshold δ for our methods is set to be $\{10^{-6}, 10^{-3}, 10^{-4}, 10^{-3}, 0.0025, 0.005, 0.075, 0.01\}$; the penalizing hyper-parameter λ for CENet and WassersteinNet in Chi *et al.* (2021) is set as $\{0.01, 0.1, 1.0, 5.0, 10.0, 15.0, 20.0, 50.0, 100.0\}$.

A.4.2.3. *Medical Insurance Cost Data.* The medical insurance cost dataset Lantz (2013) is a simulated dataset created using real-world demographic statistics from the U.S. Census Bureau. It contains 1,338 medical expense examples for patients in the United States. In this experiment, gender is selected as the sensitive attribute and the charged medical expenses as the target variable. In this experiment, gender is chosen as the sensitive attribute, and the charged medical expenses are the target variable. We apply the same sub-sample process as in Chi *et al.* (2021), by randomly selecting 5% of examples with male gender and 50% of examples with female gender, resulting in a total of 364 examples. For ridge regression, we set the tuning hyper-parameter range to be 20 numbers spaced evenly on a log scale between 10^{-9} and 10^{-3}. We use the Sobolev kernel in kernel ridge regression. We select the penalizing hyper-parameter η from 20 evenly spaced numbers on a log scale between 10^{-8} and 1. The MSE disparity threshold δ for our methods is set to

be $\{10^{-6}, 10^{-3}, 10^{-4}, 10^{-3}, 0.0025, 0.005, 0.01\}$; the penalizing hyper-parameter λ for CENet and WassersteinNet in Chi *et al.* (2021) is set as $\{0.01, 0.1, 0.5, 1.0, 2.0, 5.0\}$.

A.5. Figures

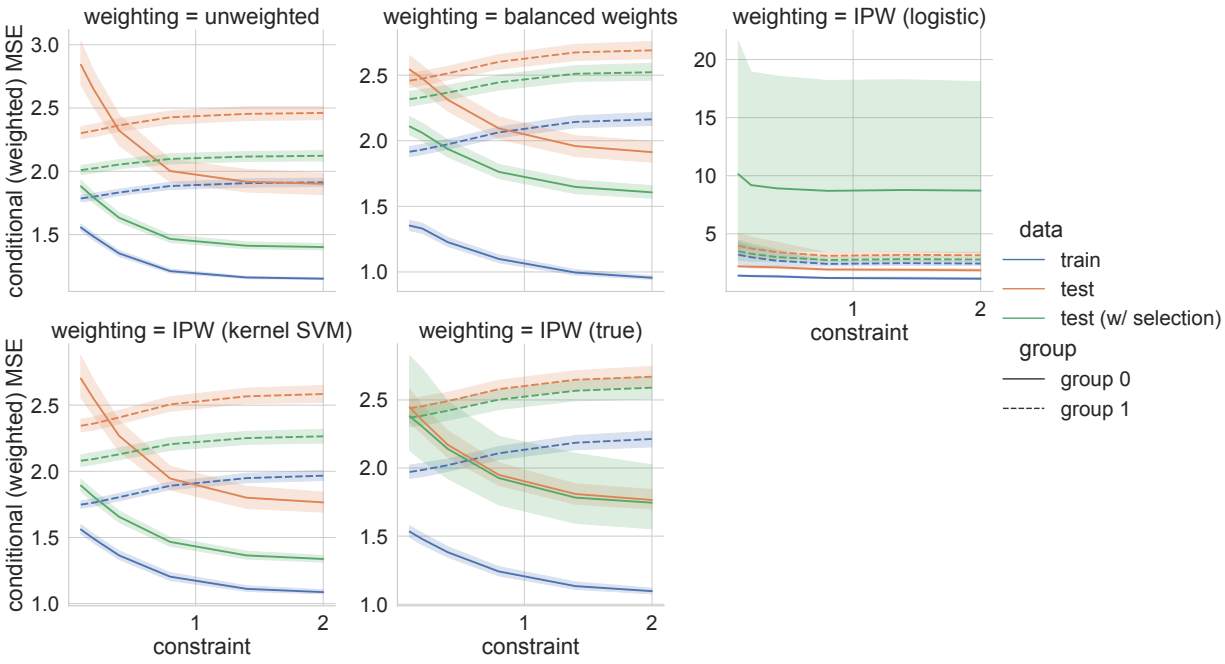


FIGURE A.1. Conditional (weighted) MSE of nonparametric fair regressor fifted with Algorithm 2.

We provide additional results and analysis for the experiments conducted in Section 2.7. We present the average performance as well as the standard error of the MSE disparity under various δ and λ for all three datasets.

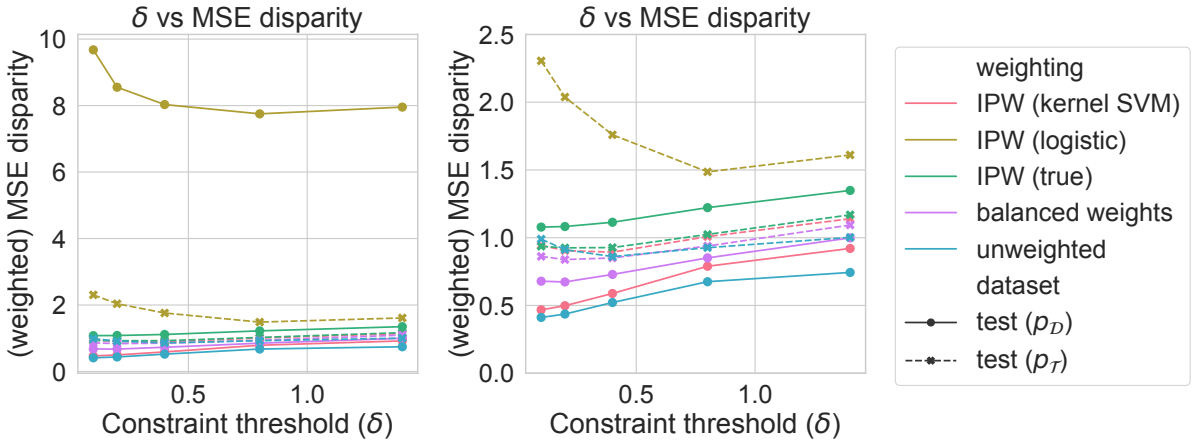


FIGURE A.2. Performance of fitted nonparametric models fitted with Algorithm 2. **Left:** (Weighted) MSE disparity evaluated on test dataset at various δ; **Right:** Details for the left panel.

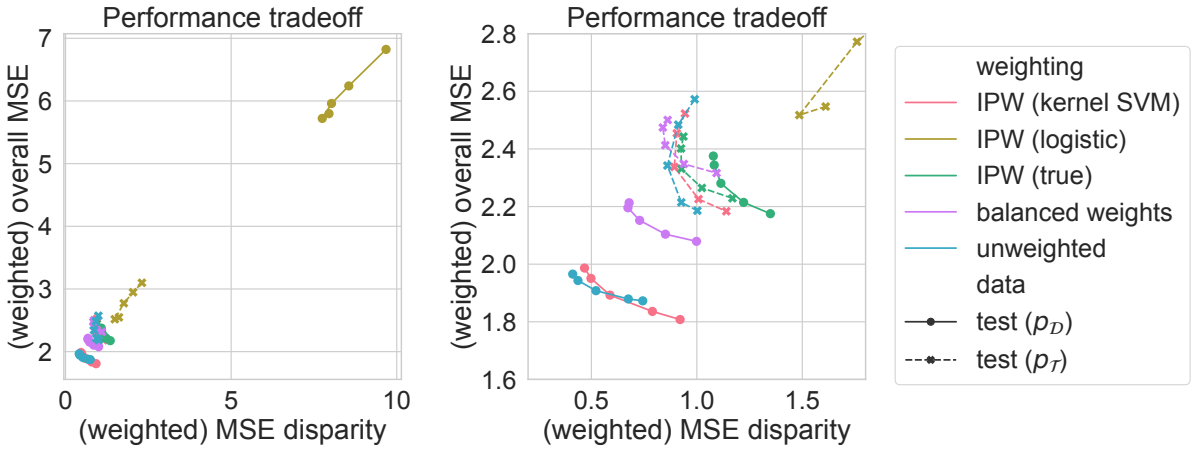


FIGURE A.3. Performance of fitted nonparametric models fitted with Algorithm 2. **Left:** Tradeoff between overall (weighted) MSE and (weighted) MSE disparity; **Right:** Details for the left panel.

A.6. Tables

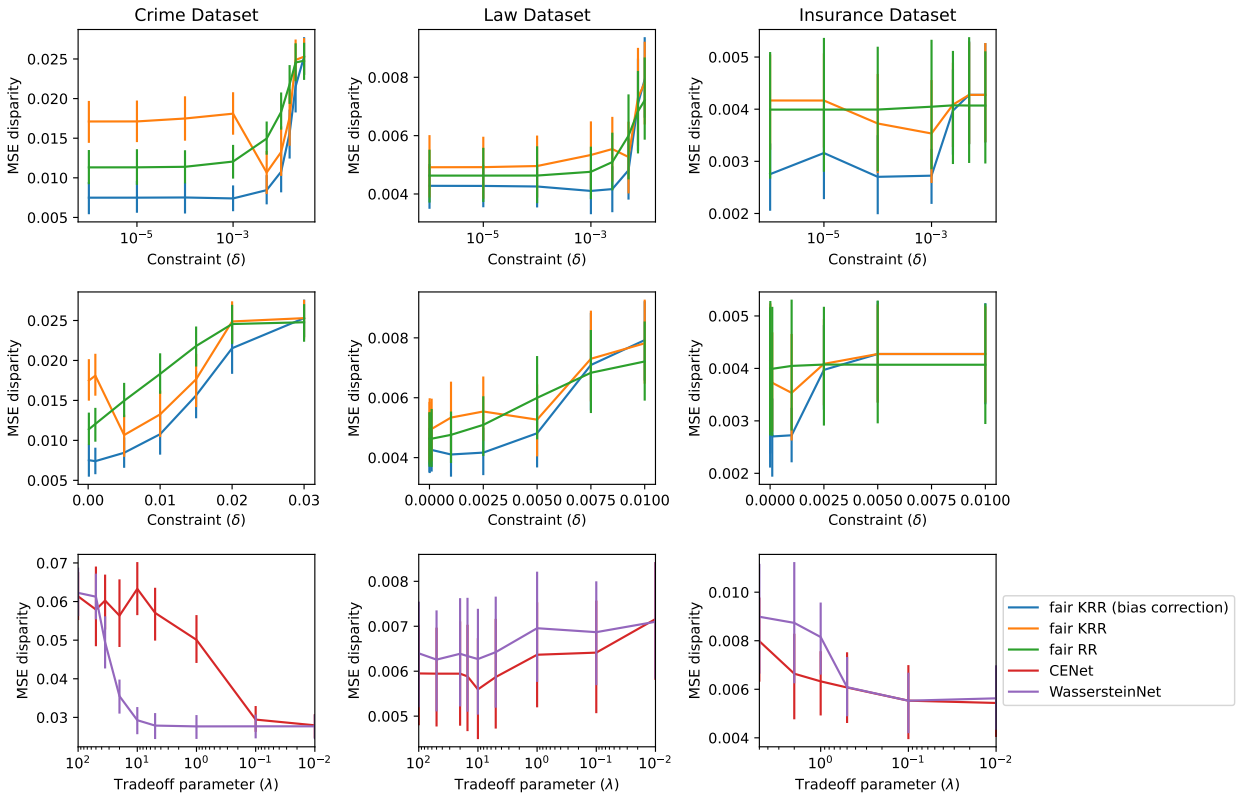


FIGURE A.4. MSE disparity at various levels of fairness controlled by δ or λ. The model is supposed to be more "fair" with a smaller δ or larger λ. The standard error of 10 repetitions is represented by the error bars.

TABLE A.1. Mean and standard error of MSE of parametric model fitted with Algrotihm 1 evaluated on testing dataset from $p_{\mathcal{T}}$; the numbers reported are obtained from 500 simulated datasets; all the values are multiplied by 100.

δ	weights	MSE	MSE$\|A=1$	MSE$\|A=0$	MSE disparity
	IPW (kernel SVM)	675.587 (0.249)	693.889 (0.886)	657.285 (1.3)	49.704 (1.559)
	IPW (logistic)	669.682 (0.284)	965.535 (1.503)	373.829 (1.207)	591.706 (2.666)
0.0	IPW (true)	678.58 (0.273)	677.703 (0.951)	679.457 (1.43)	42.168 (1.427)
	balanced weights	677.949 (0.205)	678.063 (0.731)	677.834 (1.06)	32.112 (1.039)
	unweighted	667.191 (0.234)	846.634 (0.836)	487.748 (1.061)	358.885 (1.852)
	IPW (kernel SVM)	674.712 (0.245)	697.8 (0.889)	651.625 (1.293)	55.165 (1.69)
	IPW (logistic)	669.983 (0.286)	969.064 (1.499)	370.901 (1.195)	598.164 (2.651)
10.0	IPW (true)	677.585 (0.269)	681.721 (0.955)	673.448 (1.424)	43.207 (1.407)
	balanced weights	676.959 (0.202)	682.078 (0.734)	671.84 (1.055)	33.237 (1.066)
	unweighted	667.296 (0.236)	844.799 (0.903)	489.792 (1.139)	355.007 (2.0)
	IPW (kernel SVM)	673.022 (0.237)	705.682 (0.895)	640.362 (1.28)	69.312 (1.892)
	IPW (logistic)	670.622 (0.29)	976.186 (1.492)	365.057 (1.172)	611.13 (2.621)
30.0	IPW (true)	675.657 (0.26)	689.82 (0.963)	661.494 (1.413)	49.113 (1.531)
	balanced weights	675.042 (0.195)	690.17 (0.741)	659.915 (1.045)	40.445 (1.297)
	unweighted	667.532 (0.241)	841.345 (1.07)	493.719 (1.318)	347.627 (2.352)
	IPW (kernel SVM)	671.41 (0.229)	713.646 (0.902)	629.173 (1.268)	85.749 (2.049)
	IPW (logistic)	671.312 (0.294)	983.395 (1.486)	359.23 (1.15)	624.165 (2.591)
50.0	IPW (true)	673.814 (0.251)	698.003 (0.973)	649.625 (1.401)	59.983 (1.744)
	balanced weights	673.211 (0.189)	698.347 (0.747)	648.074 (1.035)	53.755 (1.546)
	unweighted	667.763 (0.244)	838.51 (1.231)	497.015 (1.487)	341.495 (2.686)
	IPW (kernel SVM)	669.876 (0.221)	721.692 (0.909)	618.061 (1.256)	104.054 (2.106)
	IPW (logistic)	672.056 (0.299)	990.691 (1.479)	353.421 (1.128)	637.27 (2.562)
70.0	IPW (true)	672.057 (0.242)	706.273 (0.982)	637.841 (1.39)	74.573 (1.949)
	balanced weights	671.464 (0.183)	706.61 (0.753)	636.318 (1.026)	71.329 (1.677)
	unweighted	667.957 (0.248)	836.353 (1.356)	499.56 (1.621)	336.793 (2.948)
	IPW (kernel SVM)	668.423 (0.213)	729.821 (0.917)	607.025 (1.244)	123.038 (2.115)
	IPW (logistic)	672.855 (0.303)	998.079 (1.473)	347.631 (1.106)	650.448 (2.533)
90.0	IPW (true)	670.386 (0.234)	714.629 (0.992)	626.144 (1.38)	91.26 (2.134)
	balanced weights	669.804 (0.177)	714.959 (0.76)	624.649 (1.016)	90.536 (1.736)
	unweighted	668.121 (0.252)	834.662 (1.457)	501.58 (1.734)	333.082 (3.163)

TABLE A.2. Mean and standard error of MSE of nonparametric model fitted with Algrotihm 1 evaluated on testing dataset from $p_{\mathcal{T}}$; the numbers reported are obtained from 500 simulated datasets; all the values are multiplied by 100.

δ	weights	MSE	MSE$\|A=1$	MSE$\|A=0$	MSE disparity
10.0	IPW (kernel SVM)	706.168 (93.993)	585.711 (70.11)	826.625 (119.788)	291.014 (56.012)
	IPW (logistic)	327.416 (27.312)	315.067 (22.054)	339.765 (33.285)	92.997 (13.732)
	IPW (true)	414.159 (53.952)	355.667 (37.004)	472.651 (71.812)	168.384 (37.147)
	balanced weights	323.547 (9.878)	297.526 (9.855)	349.567 (10.837)	94.313 (5.137)
	unweighted	376.614 (26.339)	324.82 (21.91)	428.408 (31.88)	141.469 (14.109)
20.0	IPW (kernel SVM)	778.634 (102.295)	640.915 (75.172)	916.353 (131.47)	334.894 (62.775)
	IPW (logistic)	314.124 (27.107)	307.758 (21.906)	320.489 (33.032)	90.538 (13.662)
	IPW (true)	432.662 (55.791)	368.39 (37.576)	496.934 (74.916)	187.939 (39.502)
	balanced weights	291.0 (6.578)	269.798 (6.283)	312.202 (8.062)	87.449 (4.914)
	unweighted	337.52 (23.328)	297.929 (18.626)	377.111 (29.094)	126.203 (13.809)
40.0	IPW (kernel SVM)	975.42 (117.774)	827.021 (90.146)	1123.819 (148.083)	405.192 (66.888)
	IPW (logistic)	330.014 (35.1)	330.059 (28.37)	329.969 (42.35)	101.073 (15.759)
	IPW (true)	442.202 (65.352)	395.226 (48.27)	489.178 (83.694)	179.968 (39.229)
	balanced weights	292.657 (14.981)	282.801 (15.592)	302.512 (15.045)	85.628 (5.225)
	unweighted	336.462 (26.344)	314.533 (21.522)	358.392 (32.049)	118.912 (13.421)
80.0	IPW (kernel SVM)	950.936 (117.743)	828.358 (90.984)	1073.515 (146.259)	409.115 (60.586)
	IPW (logistic)	323.637 (34.549)	345.062 (28.687)	302.212 (40.957)	120.449 (14.174)
	IPW (true)	420.18 (63.603)	404.722 (47.915)	435.638 (79.955)	167.559 (33.783)
	balanced weights	265.565 (14.13)	279.298 (14.707)	251.832 (14.065)	84.976 (4.077)
	unweighted	295.792 (21.383)	305.785 (18.832)	285.798 (24.38)	105.262 (6.906)
140.0	IPW (kernel SVM)	492.887 (79.608)	470.876 (61.133)	514.898 (99.161)	239.665 (40.981)
	IPW (logistic)	287.297 (35.212)	324.604 (23.27)	249.99 (47.694)	153.412 (25.239)
	IPW (true)	375.241 (58.012)	378.524 (42.276)	371.957 (74.229)	178.737 (32.697)
	balanced weights	257.865 (14.103)	291.166 (14.73)	224.565 (14.001)	106.388 (4.07)
	unweighted	221.001 (3.69)	247.991 (4.14)	194.011 (5.071)	100.585 (4.102)
200.0	IPW (kernel SVM)	385.445 (60.62)	390.708 (47.579)	380.182 (74.527)	195.035 (29.337)
	IPW (logistic)	246.469 (5.696)	306.499 (6.944)	186.438 (5.979)	146.81 (4.882)
	IPW (true)	257.976 (24.079)	295.928 (19.578)	220.023 (28.935)	131.343 (9.945)
	balanced weights	231.887 (3.482)	270.036 (3.873)	193.737 (4.879)	112.0 (3.956)
	unweighted	218.074 (2.445)	245.879 (2.766)	190.27 (4.486)	102.213 (4.111)

TABLE A.3. Mean and standard error of MSE of nonparametric model fitted with Algorithm 2 evaluated on testing dataset from $p_{\mathcal{T}}$; the numbers reported are obtained from 500 simulated datasets; all the values are multiplied by 100.

δ	weights	MSE	MSE$\|A=1$	MSE$\|A=0$	MSE disparity
10.0	IPW (kernel SVM)	252.295 (4.09)	234.437 (2.504)	270.154 (7.959)	94.381 (7.551)
	IPW (logistic)	309.853 (27.098)	398.529 (52.652)	221.177 (6.05)	230.49 (51.352)
	IPW (true)	244.307 (3.439)	243.893 (3.29)	244.722 (6.31)	93.645 (6.034)
	balanced weights	250.05 (2.953)	245.794 (2.678)	254.306 (5.453)	86.171 (4.919)
	unweighted	257.243 (4.418)	230.194 (2.387)	284.291 (8.778)	98.975 (8.581)
20.0	IPW (kernel SVM)	245.441 (3.824)	236.065 (2.535)	254.817 (7.363)	90.507 (6.861)
	IPW (logistic)	294.717 (24.144)	372.876 (46.614)	216.559 (5.813)	203.777 (45.246)
	IPW (true)	240.151 (3.207)	245.372 (3.314)	234.931 (5.722)	92.511 (5.418)
	balanced weights	247.441 (2.843)	247.438 (2.724)	247.443 (5.15)	83.883 (4.631)
	unweighted	248.419 (4.039)	232.027 (2.419)	264.81 (7.953)	91.186 (7.646)
40.0	IPW (kernel SVM)	233.739 (3.151)	240.69 (2.582)	226.788 (5.859)	89.327 (5.163)
	IPW (logistic)	277.206 (19.415)	342.43 (36.37)	211.983 (6.517)	175.938 (34.565)
	IPW (true)	233.084 (2.847)	249.14 (3.317)	217.028 (4.818)	92.716 (4.565)
	balanced weights	241.279 (2.669)	251.361 (2.751)	231.196 (4.714)	85.057 (4.17)
	unweighted	234.236 (3.369)	236.323 (2.463)	232.148 (6.481)	86.042 (5.996)
80.0	IPW (kernel SVM)	222.572 (2.628)	250.571 (2.756)	194.573 (4.55)	100.804 (3.858)
	IPW (logistic)	251.746 (7.278)	310.381 (11.595)	193.111 (5.341)	148.592 (9.868)
	IPW (true)	226.462 (2.654)	257.952 (3.482)	194.971 (4.074)	102.413 (4.025)
	balanced weights	234.775 (2.56)	260.151 (2.896)	209.4 (4.241)	93.853 (3.746)
	unweighted	221.45 (2.547)	242.666 (2.582)	200.233 (4.622)	92.632 (4.064)
140.0	IPW (kernel SVM)	218.369 (2.458)	256.676 (3.031)	180.063 (4.173)	113.98 (3.841)
	IPW (logistic)	254.72 (8.554)	319.025 (13.273)	190.414 (6.232)	161.068 (10.887)
	IPW (true)	222.878 (2.528)	264.818 (3.609)	180.937 (3.843)	116.876 (4.093)
	balanced weights	231.64 (2.45)	267.34 (3.135)	195.939 (4.081)	109.358 (3.899)
	unweighted	218.569 (2.472)	245.379 (2.728)	191.758 (4.498)	100.237 (4.064)
200.0	IPW (kernel SVM)	217.472 (2.385)	258.42 (3.158)	176.524 (4.041)	117.871 (3.932)
	IPW (logistic)	250.817 (7.238)	315.359 (11.108)	186.276 (6.063)	163.082 (9.53)
	IPW (true)	221.788 (2.472)	267.056 (3.795)	176.52 (3.713)	122.201 (4.294)
	balanced weights	230.129 (2.402)	268.95 (3.264)	191.309 (3.982)	113.337 (4.035)
	unweighted	218.155 (2.45)	246.15 (2.789)	190.159 (4.486)	102.595 (4.121)

APPENDIX B

Appendix to Chapter 4

B.1. Model Selection Pitfull when *Enclosure Condition* is Violated

In this section, we would construct a simple example to illustrate the potential issue of fitting J when the Enclosure Assumption is voilated. Consider the example in the left panel of Figure B.1. We parametrize the object boundary C, highlighted by red curve, with FD and aim the fit the number of FD with the BIC. The Enclosure Assumption is voilated because partial C is not interior point of $\mathcal{R}$, the blue segment. In the right panel of Figure B.1, we displays the paramerized function $s(t; \hat{\theta}, J)$ with $J = 2$. Notably, the paramerized curve now intersects with the boundary of $\mathcal{R}$ and creates an extra region. The presence of extra regions introduces ambiguity during the derivation of the profile log-likelihood function for the corresponding segmentation due to the undefined intensity within these additional segments.

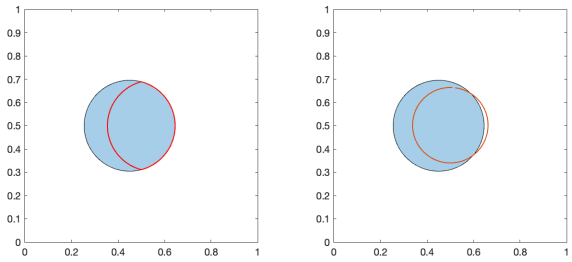


FIGURE B.1. Illustration of pitfull caused when Enclosure Assumption is violated. **Left**: Red curve: object boundary C; Blue segment: A construction of $\mathcal{R}$ that violates the Enclosure Assumption as partial C is not interior point of $\mathcal{R}$; **Right**: Red curve: Paramerized function $s(t; \hat{\theta}, J)$ with $J = 2$. Its intersection with the boundary of $\mathcal{R}$ creates an extra region.

APPENDIX C

Appendix to Chapter 5

C.1. Proof of Lemma 5.2.3.1

PROOF. It suffices to prove $\theta_{jp}^{[j-1,k]} = \theta_{jp}^{[j,k]}$ for the first equation. From theorem 3.2 Nandy *et al.* (2017), $\theta_{jp}^{[j-1,k]} = \theta_{jp}^{[j,k]} - \theta_{jj-1}^{[j,k]}\theta_{jp}^{[j-1,k]\setminus\{j\}}$. Since there is no path in $\mathcal{G}_0$ from X_j to X_{j-1}, $\theta_{jj-1}^{[j,k]} = 0$. The second equation follows from a similar argument. □

C.2. Proof of Theorem 5.2.4

PROOF. Without loss of generality, we can assume that $X_i, \ldots, X_k$ are in causal order.

Base case: We start with proving $\eta_k = \theta_{tk}\theta_{kp}^{[i,k]}$. From Theorem 3.1 in Nandy *et al.* (2017) and Lemma 5.2.3.1, $\eta_k = \theta_{tk}\theta_{kp} = \theta_{tk}\theta_{kp}^{[i,k]}$.

Induction step: We will then prove for any $j \in [i,k)$, $\eta_{[j,k]} - \eta_{[j+1,k]} = \theta_{tj}\theta_{jp}^{[j,k]}$. Based on Definition 5.2.2, we have

$$\begin{aligned}\eta_{[j,k]} - \eta_{[j+1,k]} =& (\theta_{tp} - \theta_{tp}^{\{t\}\cup[j+1,k]}) - (\theta_{tp} - \theta_{tp}^{\{t\}\cup[j,k]}) \\ =& \theta_{tp}^{\{t\}\cup[j,k]} - \theta_{tp}^{\{t\}\cup[j+1,k]} \\ =& \theta_{tj}^{\{t\}\cup[j+1,k]}\theta_{jp}^{[j,k]},\end{aligned} \tag{C.1}$$

where the last equality follows from Theorem 3.2 in Nandy *et al.* (2017). Lemma 5.2.3.1 guarantees $\theta_{tj}^{\{t\}\cup[j+1,k]} = \theta_{tj}$, and the result follows.

If $\eta_{[j+1,k]} = \boldsymbol{\beta}_t^{(j+1,\ldots,k)^T}\boldsymbol{\theta}_p^{(j+1,\ldots,k)}$ holds, from Equation C.1, we have

$$\eta_{[j,k]} = \boldsymbol{\beta}_t^{(j+1,\ldots,k)^T}\boldsymbol{\theta}_p^{(j+1,\ldots,k)} + \theta_{tj}\theta_{jp}^{[j,k]} = \boldsymbol{\beta}_t^{(j,\ldots,k)^T}\boldsymbol{\theta}_p^{(j,\ldots,k)},$$

where the last equality holds due to Lemma 5.2.3.1. We then complete the proof by induction. □

C.3. Proof of Theorem 5.3.6

To prove Theorem 5.3.6, we first prove Lemma C.3.0.2 and Lemma C.3.1.1.

LEMMA C.3.0.1. *((A.8) Chakrabortty* et al. *(2018)). For any* $i \neq k \in \{t+1, \ldots, p\}$, $\mathbb{P}(|\hat{\theta}_{ik}(\mathbf{PA}'_{\mathcal{G}'_r}(X_i)) - \theta_{ik}(\mathbf{PA}'_{\mathcal{G}'_r}(X_i))| > \delta) \leq 2\exp(-C_7 n\delta^2)$ *for some absolute constant* $C_7 > 0$ *depending on* C_2, C_3, C_4 *in assumption 5.3.5.*

Note that although the original argument in ($(A.8) Chakrabortty$ *et al.* (2018)) proves for the case $k = p$ only, the conclusion can be generalized to any other pair of X_i, X_k such that $i \neq k \in \{t+1, \ldots, p\}$ with the same argument. Therefore, we omit the proof here.

LEMMA C.3.0.2. $\mathbb{P}(|\hat{\theta}^{(i,j)}_{ip}(\mathcal{PA}'_{\mathcal{G}'_r}(X_i, X_j)) - \theta^{(i,j)}_{ip}(\mathcal{PA}'_{\mathcal{G}'_r}(X_i, X_j))| > \delta) \leq 2\exp(-C_7^* n\delta^2)$ *for some absolute constant* $C_7^* > 0$ *depending on* C_2, C_3, C_4 *in assumptions 5.3.5.*

PROOF. If $X_p \in \mathbf{PA}'_{\mathcal{G}'_r}(X_i)$, then $\hat{\theta}^{(i,j)}_{ip}(\mathcal{G}'_r) = 0$. Otherwise, we can have the following decomposition based on the Oracle in Theorem 3.1 and the definition of the RRC estimator in Nandy *et al.* (2017):

$$
\begin{aligned}
&|\hat{\theta}^{(i,j)}_{ip}(\mathcal{G}'_r) - \theta^{(i,j)}_{ip}(\mathcal{G}'_r)| \\
\leq &|\hat{\theta}_{ip}(\mathcal{G}'_r) - \theta_{ip}(\mathcal{G}'_r)| + |\hat{\theta}_{ij}(\mathcal{G}'_r)\hat{\theta}_{jp}(\mathcal{G}'_r) - \theta_{ij}(\mathcal{G}'_r)\theta_{jp}(\mathcal{G}'_r)| \\
\leq &|\hat{\theta}_{ip}(\mathcal{G}'_r) - \theta_{ip}(\mathcal{G}'_r)| + |\hat{\theta}_{ij}(\mathcal{G}'_r) - \theta_{ij}(\mathcal{G}'_r)||\hat{\theta}_{jp}(\mathcal{G}'_r) - \theta_{jp}(\mathcal{G}'_r)| \\
&+ |\theta_{ij}(\mathcal{G}'_r)||\hat{\theta}_{jp}(\mathcal{G}'_r) - \theta_{jp}(\mathcal{G}'_r)| + |\theta_{jp}(\mathcal{G}'_r)||\hat{\theta}_{ij}(\mathcal{G}'_r) - \theta_{ij}(\mathcal{G}'_r)|
\end{aligned}
$$

We can complete the proof by showing that $|\theta_{ij}(\mathcal{G}'_r)|$ and $|\theta_{jp}(\mathcal{G}'_r)|$ are bounded, and the result follows from lemma C.3.0.1. It is shown in (A.12) Chakrabortty *et al.* (2018) that under assumption 5.3.3, $|\theta_{ik}(\mathcal{G}'_r)| \leq \frac{C_2}{C_3}$ for any $i, k \in t+1, \ldots, p$. □

Denote the multiset of all possible total joint effect of X_i and X_j on X_p in $\mathcal{C}'_0$ as follows:

$$
\mathbf{\Theta}^{(i,j)}_p(\mathcal{C}'_0) = \left\{ \left(\theta^{(i,j)}_{ip}(\mathcal{G}'), \theta^{(i,j)}_{jp}(\mathcal{G}')\right) : \mathcal{G}' \in \mathrm{MEC}(\mathcal{C}'_0) \right\}
$$

for $i \neq j \in \{t+1, \ldots, p-1\}$.

We first introduce a distance metric $d(A, B)$ of two multisets A and B. We then prove the uniform consistency of the RRC estimator $\hat{\mathbf{\Theta}}^{(i,j)}_p(\hat{\mathcal{C}}'_0)$ in Lemma C.3.1.1 .

DEFINITION C.3.1. *For any two multisets of scalars* $A = \{a_1, \ldots, a_{m_1}\}$ *and* $B = \{b_1, \ldots, b_{m_2}\}$ *with ordered statistics* $a_{(1)}, ..., a_{(m_1)}$ *and* $b_{(1)}, ..., b_{(m_2)}$*, we define*

$$d(A, B) = \begin{cases} \sup\limits_{i=1,\ldots,m_1} |a_{(i)} - b_{(i)}|, & \text{if } m_1 = m_2, \\ \infty, & \text{otherwise.} \end{cases}$$

For multisets of 2-dimensional vectors $A = \{(a_{11}, a_{21})^T, \ldots, (a_{1m_1}, a_{2m_1})^T\}$*, and* $B = \{(b_{11}, b_{21})^T, \ldots, (b_{1m_1}, b_{2m_1})^T\}$*, we define*

$$d(A, B) = \max\big(d(A_1, B_1), d(A_2, B_2)\big),$$

where $A_i = \{a_{i1}, \ldots, a_{im_1}\}, B_i = \{b_{i1}, \ldots, b_{im_2}\}$ *for* $i = 1, 2$.

LEMMA C.3.1.1. *(Uniform consistency of the RRC estimator under sub-Gaussian errors). Under the assumptions 5.3.1-5.3.5, there exists a sequence of* $\alpha_n \to 0$ *such that*

$$\max_{t<i,j<p, i\neq j} d\left(\hat{\mathbf{\Theta}}_p^{(i,j)}(\hat{\mathcal{C}}'(\alpha_n)), \mathbf{\Theta}_p^{(i,j)}(\mathcal{C}_0')\right) \xrightarrow{\mathbb{P}} 0$$

Let $\{\mathcal{G}_1', \ldots, \mathcal{G}_m'\}$ be the Markov equivalence class of $\mathcal{C}_0'$, such that each $\mathcal{G}_r'$ in the set is a high-dimensional sparse DAG satisfying assumptions 5.3.2 and 5.3.3.

PROOF. Define $A_n := \{\hat{C}'(\alpha_n) = C_0'\}$ where $\hat{C}'(\alpha_n)$ denote the output of a suitable structural learning algorithm and α_n is the tuning parameter. Since $\mathbb{P}(A_n^C) \to 0$, it suffices to prove for any δ,

$$\mathbb{P}\left(\max_{t<i,j<p, i\neq j} d\big(\hat{\mathbf{\Theta}}_p^{(i,j)}(\mathcal{C}_0'), \mathbf{\Theta}_p^{(i,j)}(\mathcal{C}_0')\big) > \delta, A_n\right) \to 0.$$

On A_n, $\hat{\theta}_{ip}^{(i,j)}(\mathcal{C}_0')$ and $\theta_{ip}^{(i,j)}(\mathcal{C}_0')$ have the same cardinality. We have

$$d\big(\hat{\theta}_{ip}^{(i,j)}(\mathcal{C}_0'), \theta_{ip}^{(i,j)}(\mathcal{C}_0')\big) \leq \max_{r=1,\ldots,m} \big|\hat{\theta}_{ip}^{(i,j)}(\mathcal{PA}_{\mathcal{G}_r'}(X_i, X_j)) - \theta_{ip}^{(i,j)}(\mathcal{PA}_{\mathcal{G}_r'}(X_i, X_j))\big|$$

since the pairing of $(\hat{\theta}_{ip}^{(i,j)}, \theta_{ip}^{(i,j)})$ with $d(\cdot, \cdot)$ is the optimal pairing for the supremum distance. Also, note that the distinct element in $\{\mathcal{PA}_{\mathcal{G}_r'}'(X_i, X_j) : \mathcal{G}_r' \in \text{MEC}(\mathcal{C}_0')\}$ is bounded above by

$2^{|\mathbf{Adj}_{\mathcal{C}_0'}(X_i)|} \times 2^{|\mathbf{Adj}_{\mathcal{C}_0'}(X_j)|}$, which is bounded by 4^q by assumption 5.3.3. Hence,

$$\begin{aligned}
&\mathbb{P}\left(\max_{t<i,j<p,i\neq j} d(\hat{\mathbf{\Theta}}_p^{(i,j)}(\mathcal{C}_0'), \mathbf{\Theta}_p^{(i,j)}(\mathcal{C}_0')) > \delta, A_n\right) \\
&\leq (p-2)^2 \max_{t<i,j<p,i\neq j} \mathbb{P}\left(\max_{r=1,\ldots,m} |\hat{\theta}_{ip}^{(i,j)}(\mathcal{G}_r') - \theta_{ip}^{(i,j)}(\mathcal{G}_r')| > \delta\right) \\
&\leq (p-2)^2 4^q \max_{t<i,j<p,i\neq j} \max_{r=1,\ldots,m} \mathbb{P}\left(|\hat{\theta}_{ip}^{(i,j)}(\mathcal{G}_r') - \theta_{ip}^{(i,j)}(\mathcal{G}_r')| > \delta\right) \\
&\leq 2(p-2)^2 4^q \exp(-C_7^* n\delta^2)
\end{aligned}$$

where the third inequality is ensured by Lemma C.3.0.2, and the last inequality is ensured by assumption 5.3.2 and 5.3.3. □

Proof of Theorem 5.3.6. Here we present the proof of Theorem 5.3.6.

Proof. The first part of the theorem follows as a direct corollary of Lemma C.3.1.1. On $A_n := \{\hat{C}'(\alpha_n) = C_0'\}$, we have

$$\begin{aligned}
&\left|\text{aver}(\hat{\Theta}_{ip}^{(i,j)}(\mathcal{C}_0')) - \text{aver}(\Theta_{ip}^{(i,j)}(\mathcal{C}_0')\right| \\
&= \frac{1}{m}\Big|\sum_{r=1}^{m}(\hat{\theta}_{ip}^{(i,j)}(\mathcal{G}_r') - \theta_{ip}^{(i,j)}(\mathcal{G}_r'))\Big| \\
&\leq \frac{1}{m} m \sup_{\text{i}=1,\ldots,\text{m}} \left|\hat{\theta}_{ip_{(\text{i})}}^{(i,j)} - \theta_{ip_{(\text{i})}}^{(i,j)}\right| \\
&= d\left(\hat{\Theta}_{ip}^{(i,j)}(\mathcal{C}_0'), \Theta_{ip}^{(i,j)}(\mathcal{C}_0')\right) \\
&\leq d\left(\hat{\mathbf{\Theta}}_p^{(i,j)}(\mathcal{C}_0'), \mathbf{\Theta}_p^{(i,j)}(\mathcal{C}_0')\right),
\end{aligned}$$

where $\theta^{(i,j)}{}_{ip_{(\text{i})}}$ and $\hat{\theta}_{ip_{(\text{i})}}^{(i,j)}$ are the i^{th} ordered element in $\Theta_{ip}^{(i,j)}(\mathcal{C}_0')$ and $\hat{\Theta}_{ip}^{(i,j)}(\mathcal{C}_0')$, respectively.

Next, we prove the second part of the theorem, i.e., $|\hat{\beta}_{ti|\mathbf{X}_{1,\ldots,t-1}} \times \text{aver}(\hat{\Theta}_{ip}^{(i,j)}(\hat{\mathcal{C}}_0')) - \theta_{ti} \times \text{aver}(\Theta_{ip}^{(i,j)}(\mathcal{C}_0'))|$ converges to 0 in probability.

Following $|\hat{a}\hat{b} - ab| \leq |a||\hat{b} - b| + |b||\hat{a} - a| + |\hat{a} - a||\hat{b} - b|$, we let $\hat{a} = \hat{\beta}_{ti|\mathbf{X}_{1,\ldots,t-1}}, \hat{b} = \text{aver}(\hat{\Theta}_{ip}^{(i,j)}(\hat{\mathcal{C}}_0')), a = \theta_{ti}, b = \text{aver}(\Theta_{ip}^{(i,j)}(\mathcal{C}_0'))$. By showing that $|\text{aver}(\Theta_{ip}^{(i,j)}(\mathcal{C}_0'))|$ is bounded, the conclusion will follow from $(A.11)$ Chakrabortty *et al.* (2018) and the first part of theorem 5.3.6. Note that it has been shown in $(A.12)$ Chakrabortty *et al.* (2018) that $|\theta_{ti}| \leq \frac{C_3}{C_2}$. Therefore, $|\text{aver}(\Theta_{ip}^{(i,j)}(\mathcal{C}_0'))| \leq \max_{\mathcal{G}_0'} |\theta_{ip}(\mathcal{G}_0') - \theta_{ij}(\mathcal{G}_0')\theta_{jp}(\mathcal{G}_0')| \leq \frac{C_3}{C_2}(1 + \frac{C_3}{C_2})$, and we complete the proof.

The third part will follow from the second part directly as the result of triangle inequality.

□

C.4. Additional Real Data Analysis Results

TABLE C.1. Estimated causal effect $\hat{\theta}_{tj}$, total joint effect $\hat{\theta}_{jp}^{(\mathbf{M})}$, component-wise joint mediation effect $\hat{\eta}_j^{(\mathbf{M})}$, with their bootstrap standard deviation in parenthesis. At a significance level of 0.05, a locus with a significant $\hat{\eta}^{(\mathbf{M})}j$ after Bonferroni correction is annotated with "**"; a locus with a significant $\hat{\eta}^{(\mathbf{M})}j$ before Bonferroni correction is annotated with "*".

Locus and variable	Gene	$\hat{\theta}_{tj}$	$\hat{\theta}_{jp}^{(\mathbf{M})}$	$\hat{\eta}_j^{(\mathbf{M})}$	p-value	sig	source
cg27512205(X_{1010})	KITLG	0.005(0.002)	-152.777(157.847)	-0.74(1.003)	0.089284		
cg05608730(X_{1011})	C1QTNF2	0.007(0.002)	-430.764(146.097)	-3.08(1.027)	0.000116	**	$\mathbf{M}_{(1)}$
cg26179948(X_{1012})	JAZF1	0.006(0.003)	-176.26(87.54)	-1.113(0.744)	0.033675	*	
cg02309301(X_{1013})	ARGLU1	-0.012(0.002)	177.274(119.989)	-2.064(0.868)	0.000508	**	
cg12500973(X_{1014})	HNRNPF	-0.011(0.004)	-35.356(69.787)	0.396(0.779)	0.24614		
cg16657538(X_{1015})	ZSCAN30	0.023(0.006)	33.071(44.584)	0.753(0.943)	0.117237		$\mathbf{M}_{(2)}$
cg25626453(X_{1016})	PRRC2A	-0.012(0.002)	320.87(157.581)	-3.917(1.101)	0.000237	**	
cg13136721(X_{1017})	RPTOR	0.012(0.003)	162.28(88.494)	1.954(0.621)	0.000179	**	
cg19230917(X_{1018})	RAB5IF	-0.006(0.009)	146.289(31.146)	-0.929(0.797)	0.022734	*	
cg06422529(X_{1019})	CPQ	-0.008(0.003)	325.408(90.889)	-2.603(0.871)	0.004645	*	$\mathbf{M}_{(3)}$
cg03199124(X_{1020})	AGPAT1	-0.006(0.003)	307.091(151.185)	-1.874(0.871)	0.004146	*	
cg21875401(X_{49})	MYOCD	-0.009(0.003)	-84.228(63.985)	0.799(0.633)	0.212951		
cg04602284(X_{57})	CFAP91	0.009(0.003)	253.645(76.731)	2.2(0.654)	0.00214	*	
cg00097357(X_{122})	SYT10	0.008(0.004)	2.418(99.007)	0.019(1.042)	0.851406		
cg02032299(X_{207})	ZCCHC14	0.01(0.003)	-51.199(66.417)	-0.52(0.567)	0.095291		
cg25829425(X_{247})	NRN1	0.007(0.005)	315.384(76.117)	2.195(1.607)	0.007009	*	
cg16247183(X_{323})	ENAH	0.006(0.002)	60.002(143.582)	0.36(1.218)	0.271834		
cg06613263(X_{425})	NR3C1	-0.009(0.003)	-219.347(99.186)	2.057(0.801)	0.00115	**	
cg14693391(X_{481})	FRZB	0.006(0.003)	376.024(115.463)	2.099(0.71)	0.002937	*	
cg07635747(X_{499})	CDH2	0.004(0.003)	-478.282(77.46)	-2.12(0.631)	0.000632	**	
cg17710779(X_{512})	FGFR2	0.009(0.004)	-53.568(54.824)	-0.493(0.595)	0.425986		$\mathbf{M}_{(4)}$
cg05781698(X_{533})	IGF1R	-0.007(0.003)	-60.037(69.781)	0.409(0.874)	0.642688		
cg10631546(X_{572})	ARSJ	-0.006(0.003)	-218.464(97.067)	1.397(1.002)	0.016751	*	
cg04485065(X_{649})	FBLN5	0.004(0.003)	-179.017(76.021)	-0.766(0.681)	0.08148		
cg04295094(X_{662})	KLF17	-0.008(0.003)	-44.006(147.919)	0.363(0.79)	0.312587		
cg08158408(X_{699})	G0S2	0.008(0.003)	-49.257(103.698)	-0.398(1.051)	0.198749		
cg18703581(X_{782})	RGL2	-0.009(0.003)	-106.131(111.463)	0.945(1.19)	0.014987	*	
cg18449062(X_{846})	CNNM2	0.005(0.002)	-31.335(94.575)	-0.168(0.691)	0.571051		
cg22178613(X_{868})	SFRP2	0.007(0.004)	-54.793(77.552)	-0.399(1.284)	0.193354		
cg06688323(X_{992})	MSX1	0.006(0.004)	4.751(80.576)	0.028(0.97)	0.707256		

Bibliography

Adams, R. and Bischof, L. (1994) Seeded region growing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **16**, 641–647.

Agarwal, A., Dudík, M. and Wu, Z. S. (2019) Fair regression: Quantitative definitions and reduction-based algorithms. In *International Conference on Machine Learning*, 120–129. PMLR.

Alexander, N., Kirschbaum, C., Wankerl, M., Stauch, B. J., Stalder, T., Steudte-Schmiedgen, S., Muehlhan, M. and Miller, R. (2018) Glucocorticoid receptor gene methylation moderates the association of childhood trauma and cortisol stress reactivity. *Psychoneuroendocrinology*, **90**, 68–75. URL `https://www.sciencedirect.com/science/article/pii/S0306453017313550`.

Bakusic, J., Vrieze, E., Ghosh, M., Bekaert, B., Claes, S. and Godderis, L. (2020) Increased methylation of nr3c1 and slc6a4 is associated with blunted cortisol reactivity to stress in major depression. *Neurobiology of Stress*, **13**, 100272.

Barabasi, A.-L. and Oltvai, Z. N. (2004) Network biology: understanding the cell's functional organization. *Nature reviews genetics*, **5**, 101–113.

Barocas, S. and Selbst, A. D. (2016) Big data's disparate impact. *California law review*, 671–732.

Baron, R. M. and Kenny, D. A. (1986) The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of personality and social psychology*, **51**, 1173.

Bradley, K. and Wass, J. (2004) Hypopituitarism. In *Encyclopedia of Endocrine Diseases* (ed. L. Martini), 652–661. New York: Elsevier. URL `https://www.sciencedirect.com/science/article/pii/B012475570400706X`.

Burger, W. and Burge, M. J. (2009) *Principles of Digital Image Processing*, vol. 111. Springer.

Burger, W. and Burge, M. J. (2016) *The Discrete Fourier Transform in 2D*, 481–501. London: Springer London. URL `https://doi.org/10.1007/978-1-4471-6684-9_19`.

Burnham, K. P. and Anderson, D. R., eds. (2002) *Advanced Issues and Deeper Insights*, 267–351. New York, NY: Springer New York. URL https://doi.org/10.1007/978-0-387-22456-5_6.

Byers, S. D. and Raftery, A. E. (2002) *Bayesian Estimation and Segmentation of Spatial Point Processes Using Voronoi Tilings*, 109–119. Chapman and Hall/CRC.

Chakrabortty, A., Nandy, P. and Li, H. (2018) Inference for individual mediation effects and interventional effects in sparse high-dimensional causal graphical models. *arXiv preprint arXiv:1809.10652*.

Chi, J., Tian, Y., Gordon, G. J. and Zhao, H. (2021) Understanding and mitigating accuracy disparity in regression. In *International Conference on Machine Learning*, 1866–1876. PMLR.

Chickering, D. M. (2002) Learning equivalence classes of bayesian-network structures. *The Journal of Machine Learning Research*, **2**, 445–498.

Coston, A., Rambachan, A. and Chouldechova, A. (2021a) Characterizing fairness over the set of good models under selective labels. In *Proceedings of the 38th International Conference on Machine Learning* (eds. M. Meila and T. Zhang), vol. 139 of *Proceedings of Machine Learning Research*, 2144–2155. PMLR. URL https://proceedings.mlr.press/v139/coston21a.html.

Coston, A., Rambachan, A. and Chouldechova, A. (2021b) Characterizing fairness over the set of good models under selective labels.

Cotter, A., Gupta, M., Jiang, H., Srebro, N., Sridharan, K., Wang, S., Woodworth, B. and You, S. (2019a) Training well-generalizing classifiers for fairness metrics and other data-dependent constraints. In *International Conference on Machine Learning*, 1397–1405. PMLR.

Cotter, A., Jiang, H. and Sridharan, K. (2019b) Two-player games for efficient non-convex constrained optimization. In *Algorithmic Learning Theory*, 300–332. PMLR.

Ding, Y., Liu, J., Xu, X., Huang, M., Zhuang, J., Xiong, J. and Shi, Y. (2020) Uncertainty-aware training of neural networks for selective medical image segmentation. In *Medical Imaging with Deep Learning*, 156–173. PMLR.

Donini, M., Oneto, L., Ben-David, S., Shawe-Taylor, J. S. and Pontil, M. (2018) Empirical risk minimization under fairness constraints. *Advances in Neural Information Processing Systems*, **31**.

Du, W., Wu, X. and Tong, H. (2022) Fair regression under sample selection bias. In *2022 IEEE International Conference on Big Data (Big Data)*, 1435–1444. IEEE.

Efron, B. (1979) Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics*, **7**, 1 – 26. URL `https://doi.org/10.1214/aos/1176344552`.

Efron, B. (2004) The estimation of prediction error: covariance penalties and cross-validation. *Journal of the American Statistical Association*, **99**, 619–632.

Efron, B., Hastie, T., Johnstone, I. and Tibshirani, R. (2004) Least angle regression.

Efron, B. and Tibshirani, R. J. (1993) *An Introduction to the Bootstrap*. No. 57 in Monographs on Statistics and Applied Probability. Boca Raton, Florida, USA: Chapman & Hall/CRC.

Endo, M., Mikami, M. and Toki, S. (2015) Multigene knockout utilizing off-target mutations of the crispr/cas9 system in rice. *Plant and Cell Physiology*, **56**, 41–47.

Fan, M., Wang, J., Kashyap, V. L., Lee, T. C. M., van Dyk, D. A. and Zezas, A. (2023) Identifying diffuse spatial structures in high-energy photon lists. *The Astronomical Journal*, **165**, 66.

Ferreira, N. G., Madeira, J. L., Gergics, P., Kertsz, R., Marques, J. M., Trigueiro, N. S., Benedetti, A. F. F., Azevedo, B. V., Fernandes, B. H., Bissegatto, D. D. *et al.* (2023) Homozygous cdh2 variant may be associated with hypopituitarism without neurological disorders. *Endocrine Connections*, **12**.

Gal, Y. and Ghahramani, Z. (2015) Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. URL `https://arxiv.org/abs/1506.02142`.

Guo, X., Li, R., Liu, J. and Zeng, M. (2022) High-dimensional mediation analysis for selecting dna methylation loci mediating childhood trauma and cortisol stress reactivity. *Journal of the American Statistical Association*, **117**, 1110–1121.

Guo, X., Li, R., Liu, J. and Zeng, M. (2023) Statistical inference for linear mediation models with high-dimensional mediators and application to studying stock reaction to covid-19 pandemic. *Journal of Econometrics*, **235**, 166–179.

Hall, P. (1985) Resampling a coverage pattern. *Stochastic processes and their applications*, **20**, 231–246.

Hamu Goldberg, H. B., Mushkin, J., Raviv, T. R. and Sochen, N. (2018) Sampling technique for defining segmentation error margins with application to structural brain MRI. In *2018 25th*

IEEE International Conference on Image Processing (ICIP), 734–738.

Han, Y. and Lee, T. C. (2022) Uncertainty quantification for sparse estimation of spectral lines. *IEEE Transactions on Signal Processing*, **70**, 6243–6256.

Houck, J. (-) An introduction to exposure maps. URL `https://space.mit.edu/ASC/docs/expmap_intro.pdf`.

Houtepen, L. C., Vinkers, C. H., Carrillo-Roa, T., Hiemstra, M., Van Lier, P. A., Meeus, W., Branje, S., Heim, C. M., Nemeroff, C. B., Mill, J. *et al.* (2016) Genome-wide dna methylation levels and altered cortisol stress reactivity following childhood trauma in humans. *Nature communications*, **7**, 10967.

Huang, Y.-T. and Yang, H.-I. (2017) Causal mediation analysis of survival outcome with multiple mediators. *Epidemiology (Cambridge, Mass.)*, **28**, 370.

Julia Angwin, Jeff Larson, S. M. and Kirchner, L. (2016) Machine bias: There's software used across the country to predict future criminals. and it's biased against blacks. `https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing`. Accessed: 2016-05-23.

Kalisch, M., Hauser, A., Maechler, M., Colombo, D., Entner, D., Hoyer, P., Hyttinen, A., Peters, J., Andri, N., Perkovic, E. *et al.* (2024) Package 'pcalg'.

Kallus, N. and Zhou, A. (2018) Residual unfairness in fair machine learning from prejudiced data. In *International Conference on Machine Learning*, 2439–2448. PMLR.

Kim, Y. and Steiner, P. (2016) Quasi-experimental designs for causal inference. *Educational psychologist*, **51**, 395–405.

Kohl, S. A. A., Romera-Paredes, B., Meyer, C., De Fauw, J., Ledsam, J. R., Maier-Hein, K. H., Eslami, S. M. A., Rezende, D. J. and Ronneberger, O. (2018) A probabilistic u-net for segmentation of ambiguous images. URL `https://arxiv.org/abs/1806.05034`.

Koller, D. and Friedman, N. (2009) *Probabilistic Graphical Models: Principles and Techniques*. Adaptive computation and machine learning. MIT Press. URL `https://books.google.co.in/books?id=7dzpHCHzNQ4C`.

Lantz, B. (2013) Machine learning with r: learn how to use r to apply powerful machine learning methods and gain an insight into real-world applications. *(No Title)*.

László, Z. I. and Lele, Z. (2022) Flying under the radar: Cdh2 (n-cadherin), an important hub molecule in neurodevelopmental and neurodegenerative diseases. *Frontiers in Neuroscience*, **16**, 972059.

Lin, C. C. and Chellappa, R. (1987) Classification of partial 2-d shapes using fourier descriptors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **PAMI-9**, 686–690.

Liu, P. Z. and Nusslock, R. (2018) How stress gets under the skin: early life adversity and glucocorticoid receptor epigenetic regulation. *Current genomics*, **19**, 653–664.

Loh, J. M. and Stein, M. L. (2004) Bootstrapping a spatial point process. *Statistica Sinica*, **14**, 69–101. URL `http://www.jstor.org/stable/24307180`.

Maathuis, M. H., Kalisch, M. and Bühlmann, P. (2009) Estimating high-dimensional intervention effects from observational data. *The Annals of Statistics*, **37**, 3133–3164.

Magomedova, L., Tiefenbach, J., Zilberman, E., Le Billan, F., Voisin, V., Saikali, M., Boivin, V., Robitaille, M., Gueroussov, S., Irimia, M. *et al.* (2019) Arglu1 is a transcriptional coactivator and splicing regulator important for stress hormone signaling and development. *Nucleic Acids Research*, **47**, 2856–2870.

Mallat, S. G. and Zhang, Z. (1993) Matching pursuits with time-frequency dictionaries. *IEEE Transactions on signal processing*, **41**, 3397–3415.

Nandy, P., Hauser, A. and Maathuis, M. H. (2018a) High-dimensional consistency in score-based and hybrid structure learning. *The Annals of Statistics*, **46**, 3151–3183. URL `https://www.jstor.org/stable/26542900`.

Nandy, P., Hauser, A. and Maathuis, M. H. (2018b) High-dimensional consistency in score-based and hybrid structure learning. *The Annals of Statistics*, **46**, 3151 – 3183.

Nandy, P., Maathuis, M. H. and Richardson, T. (2016) Supplement to "estimating the effect of joint interventions from observational data in sparse high-dimensional settings.".

Nandy, P., Maathuis, M. H. and Richardson, T. S. (2017) Estimating the effect of joint interventions from observational data in sparse high-dimensional settings. *The Annals of Statistics*, **45**, 647–674.

Nguyen, T. Q., Webb-Vargas, Y., Koning, I. M. and Stuart, E. A. (2016) Causal mediation analysis with a binary outcome and multiple continuous or ordinal mediators: Simulations and application

to an alcohol intervention. *Structural equation modeling: a multidisciplinary journal*, **23**, 368–383.

Nguyen, V.-B. and Ngan Nguyen, T. (2024) Positive semidefinite interval of matrix pencil and its applications to the generalized trust region subproblems. *Linear Algebra and its Applications*, **680**, 371–390.

Oneto, L., Donini, M. and Pontil, M. (2019) General fair empirical risk minimization. URL `https://arxiv.org/abs/1901.10080`.

Park, S.-Y., Choi, J. H., Ryu, H. S., Pak, Y. K., Park, K. S., Lee, H. K. and Lee, W. (2009) C1q tumor necrosis factor α-related protein isoform 5 is increased in mitochondrial dna-depleted myocytes and activates amp-activated protein kinase. *Journal of biological chemistry*, **284**, 27780–27789.

Pearl, J. (2000) *Causality: Models, Reasoning, and Inference*. Cambridge University Press.

Pearl, J. (2009) Causal inference in statistics: An overview. *Statistics surveys*, **3**, 96–146.

Pearl, J. *et al.* (2000) Models, reasoning and inference. *Cambridge, UK: CambridgeUniversityPress*, **19**.

Perroud, N., Paoloni-Giacobino, A., Prada, P., Olié, E., Salzmann, A., Nicastro, R., Guillaume, S., Mouthon, D., Stouder, C., Dieben, K., Huguelet, P., Courtet, P. and Malafosse, A. (2011) Increased methylation of glucocorticoid receptor gene (nr3c1) in adults with a history of childhood maltreatment: a link with the severity and type of trauma. *Translational Psychiatry*, **1**, e59–e59. URL `https://doi.org/10.1038/tp.2011.60`.

Pessach, D. and Shmueli, E. (2021) Improving fairness of artificial intelligence algorithms in privileged-group selection bias data settings. *Expert Systems with Applications*, **185**, 115667.

Pong, T. K. and Wolkowicz, H. (2014) The generalized trust region subproblem. *Computational Optimization and Applications*, **58**, 273–322. URL `https://doi.org/10.1007/s10589-013-9635-7`.

Ramanagoudr-Bhojappa, R., Carrington, B., Ramaswami, M., Bishop, K., Robbins, G. M., Jones, M., Harper, U., Frederickson, S. C., Kimble, D. C., Sood, R. *et al.* (2018) Multiplexed crispr/cas9-mediated knockout of 19 fanconi anemia pathway genes in zebrafish revealed their roles in growth, sexual development and fertility. *PLoS genetics*, **14**, e1007821.

Redmond, M. (2009) Communities and Crime. UCI Machine Learning Repository. DOI: https://doi.org/10.24432/C53W3X.

Reemst, K., Kracht, L., Kotah, J. M., Rahimian, R., van Irsen, A. A., Congrains Sotomayor, G., Verboon, L. N., Brouwer, N., Simard, S., Turecki, G. *et al.* (2022) Early-life stress lastingly impacts microglial transcriptome and function under basal and immune-challenged conditions. *Translational Psychiatry*, **12**, 507.

Saxton, R. A. and Sabatini, D. M. (2017) mtor signaling in growth, metabolism, and disease. *Cell*, **168**, 960–976.

Scargle, J. D. (2002) Bayesian blocks in two or more dimensions: Image segmentation and cluster analysis. In *AIP Conference Proceedings*, vol. 617, 163–173. AIP.

Seyyed-Kalantari, L., Zhang, H., McDermott, M. B., Chen, I. Y. and Ghassemi, M. (2021) Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. *Nature medicine*, **27**, 2176–2182.

Shen, X., Diamond, S., Gu, Y. and Boyd, S. (2016) Disciplined convex-concave programming. URL `https://arxiv.org/abs/1604.02639`.

Spirtes, P., Glymour, C. and Scheines, R. (2000a) *Causation, Prediction, and Search*. MIT press, 2nd edn.

Spirtes, P., Glymour, C. N. and Scheines, R. (2000b) *Causation, prediction, and search*. MIT press.

Stone, M. (1974) Cross-validatory choice and assessment of statistical predictions. *Journal of the royal statistical society: Series B (Methodological)*, **36**, 111–133.

Tillmann, C. and Ney, H. (2003) Word reordering and a dynamic programming beam search algorithm for statistical machine translation. *Computational linguistics*, **29**, 97–133.

Tim Triche, J. (2014) *FDb.InfiniumMethylation.hg19: Annotation package for Illumina Infinium DNA methylation probes*. URL `https://www.bioconductor.org/packages/release/data/annotation/html/FDb.InfiniumMethylation.hg19.html`. R package version 2.2.0.

Tscheschel, A. and Stoyan, D. (2006) Statistical reconstruction of random point patterns. *Computational Statistics & Data Analysis*, **51**, 859–871. URL `https://EconPapers.repec.org/RePEc:eee:csdana:v:51:y:2006:i:2:p:859-871`.

van Dyk, D., Connors, A., Esch, D. N., Freeman, P., Kang, H., Karovska, M., Kashyap, V., Siemiginowska, A., Zezas, A. *et al.* (2006) Deconvolution in high-energy astrophysics: Science, instrumentation, and methods. *Bayesian Analysis*, **1**, 189–235.

van Kesteren, E.-J. and Oberski, D. L. (2019) Exploratory mediation analysis with many potential mediators. *Structural Equation Modeling: A Multidisciplinary Journal*, **26**, 710–723.

VanderWeele, T. (2015) *Explanation in causal inference: methods for mediation and interaction*. Oxford University Press.

Wang, S. and Xia, Y. (2014) Strong duality for generalized trust region subproblem: S-lemma with interval bounds. *Optimization Letters*, **9**, 1063–1073. URL `http://dx.doi.org/10.1007/s11590-014-0812-0`.

Wightman, L. F. (1998) Lsac national longitudinal bar passage study. lsac research report series.

Wong, R. K. and Chan, K. C. G. (2018) Kernel-based covariate functional balancing for observational studies. *Biometrika*, **105**, 199–213.

Woodworth, B., Gunasekar, S., Ohannessian, M. I. and Srebro, N. (2017) Learning non-discriminatory predictors. In *Conference on Learning Theory*, 1920–1953. PMLR.

Wright, S. (1921) Correlation and causation. *National Agricultural Library Digital Collections*, 557–585.

Ye, J. C., Bresler, Y. and Moulin, P. (2000) Asymptotic global confidence regions in parametric shape estimation problems. *IEEE Transactions on Information Theory*, **46**, 1881–1895.

Zeng, P., Shao, Z. and Zhou, X. (2021) Statistical methods for mediation analysis in the era of high-throughput genomics: current successes and future challenges. *Computational and structural biotechnology journal*, **19**, 3209–3224.

Zhang, H., Zheng, Y., Zhang, Z., Gao, T., Joyce, B., Yoon, G., Zhang, W., Schwartz, J., Just, A., Colicino, E. *et al.* (2016) Estimating and testing high-dimensional mediation effects in epigenetic studies. *Bioinformatics*, **32**, 3150–3154.

Zhang, Y. and Long, Q. (2021) Assessing fairness in the presence of missing data.

Zhao, H. (2021) Costs and benefits of fair regression. URL `https://arxiv.org/abs/2106.08812`.