## Title

Neurosymbolic Learning and Reasoning for Trustworthy AI

## Permalink

https://escholarship.org/uc/item/0g79n9k0

## Author

Zeng, Zhe

## Publication Date

2024

Peer reviewed|Thesis/dissertation

# UNIVERSITY OF CALIFORNIA 

Los Angeles

Neurosymbolic Learning and Reasoning for Trustworthy AI

A dissertation submitted in partial satisfaction of the requirements for the degree<br>Doctor of Philosophy in Computer Science

by

Zhe Zeng

2024
© Copyright by
Zhe Zeng
2024

# ABSTRACT OF THE DISSERTATION 

Neurosymbolic Learning and Reasoning for Trustworthy AI
by

Zhe Zeng
Doctor of Philosophy in Computer Science
University of California, Los Angeles, 2024
Professor Guy Van den Broeck, Chair

Along with the ubiquitous applications of Artificial Intelligence (AI), the quest for developing trustworthy AI models intensifies. Deep neural networks, while powerful in learning, fall short in reasoning with domain knowledge and offering robustness guarantees. Neurosymbolic AI bridges this gap by melding the learning capabilities of neural networks and reasoning techniques from symbolic AI, thus building models that behave as intended. This dissertation demonstrates my work that addresses the two fundamental challenges in neurosymbolic AI: 1) enabling differentiable learning of deep neural networks under symbolic constraints and 2) performing scalable and reliable probabilistic reasoning over expressive symbolic constraints. It presents how these neurosymbolic approaches achieve trustworthiness through explainability, uncertainty quantification, and domain-knowledge incorporation. These contributions enable broader applications of neurosymbolic AI in various domains including scientific discoveries.

The dissertation of Zhe Zeng is approved.

Kai-Wei Chang<br>Cho-Jui Hsieh<br>Todd D. Millstein<br>Guy Van den Broeck, Committee Chair

University of California, Los Angeles

2024

To my family.

## TABLE OF CONTENTS

1 Introduction ..... 1
I Learning ..... 5
2 Differentiable Learning under Constraints ..... 6
2.1 Background ..... 6
2.2 Gradient Estimator for $k$-subset Sampling ..... 7
2.2.1 Problem Statement and Motivation ..... 7
2.2.2 Simple: Subset Implicit Likelihood Estimation ..... 10
2.2.3 Empirical Evaluation ..... 17
2.3 Simple for Graph Rewiring ..... 23
2.3.1 Background ..... 23
2.3.2 Probabilistically Rewired MPNNs ..... 25
2.3.3 Expressive Power of Probabilistically Rewired MPNNs ..... 28
2.3.4 Empirical Evaluation ..... 33
2.4 Discussion ..... 37
3 Linear-Equality Constrained Deep Generative Models ..... 41
3.1 Background ..... 41
3.2 Gradient Estimator for Linear Equalities ..... 42
3.2.1 Closed-Form Expected Loss ..... 44
3.2.2 Method ..... 45
3.2.3 Empirical Evaluation ..... 51
3.3 Discussion ..... 56
4 Learning from Weak Supervisions as Constraints ..... 58
4.1 Count-Based Weakly Supervised Learning ..... 58
4.1.1 Learning from Label Proportions ..... 60
4.1.2 Multiple Instance Learning ..... 60
4.1.3 Learning from Positive and Unlabeled Data ..... 61
4.2 A Unified Approach: Count Loss ..... 62
4.3 Empirical Evaluation ..... 69
4.3.1 Learning from Label Proportions ..... 70
4.3.2 Multiple Instance Learning ..... 70
4.3.3 Learning from Positive and Unlabeled Data ..... 73
4.4 Discussion ..... 74
II Reasoning ..... 78
5 Foundations: Weighted Model Integration ..... 79
5.1 Overview ..... 79
5.2 Formalization ..... 80
5.3 Related Work ..... 84
6 Exact Inference Over Constraints ..... 86
6.1 Search-based WMI Solver ..... 86
6.1.1 Structure in WMI Problems ..... 87
6.1.2 Method ..... 91
6.1.3 Empirical Evaluation ..... 99
6.2 Message-Passing WMI Solver ..... 102
6.2.1 Tractability Analysis ..... 102
6.2.2 Method ..... 106
6.2.3 Empirical Evaluation ..... 113
6.3 Discussion ..... 116
7 Approximate Inference Over Constraints ..... 117
7.1 On the hardness of WMI ..... 117
7.2 ReCoin: Approximate WMI Solver ..... 126
7.2.1 Relaxation: introducing and then "breaking" equivalence constraints ..... 126
7.3 Empirical Evaluation ..... 134
7.4 Discussion ..... 136
8 WMI for Uncertainty Quantification ..... 137
8.1 Bayesian Model Averaging as Weighted Volume Computation ..... 138
8.1.1 A Warm-Up Example ..... 140
8.1.2 General Reduction of BMA to WVC ..... 141
8.2 Approximating BMA by WMI ..... 141
8.3 Ciber: Collapsed Inference for Bayesian Deep Learning via WMI ..... 144
8.3.1 Approximation to Posteriors ..... 145
8.3.2 Encoding into WMI Problems ..... 146
8.3.3 Exact Integration in Collapsed BMA ..... 147
8.4 Empirical Evaluation ..... 151
8.4.1 Regression on Small and Large UCI Datasets ..... 151
8.4.2 Image Classification ..... 155
8.5 Discussion ..... 157
9 Conclusion ..... 158

## LIST OF FIGURES

2.1 A comparison of the bias and variance of the gradient estimators (left) and the average and standard deviation of the cosine distance of a single-sample gradient estimate to the exact gradient. We used the cosine distance, defined as ( $1-$ cosine similarity), in place of the euclidean distance as we only care about the direction of the gradient, not magnitude. The bias, variance and error were estimated using a sample of size

10,000 . The details of this experiment are provided in Section 2.2.3.1.
2.3 Bias and variance of Simple and Gumbel Softmax over 10k samples
2.4 ELBO against \# of epochs. (Left) Comparison of Simple against different flavors of IMLE on the 10 -subset DVAE, and (Right) against ST Gumbel Softmax on the 1 -subset DVAE.19
2.5 Overview of the probabilistically rewired MPNN framework. PR-MPNNs use an upstream model to learn priors $\boldsymbol{\theta}$ for candidate edges, parameterizing a probability mass function conditioned on exactly- $k$ constraints. Subsequently, we sample multiple $k$-edge adjacency matrices (here: $k=1$ ) from this distribution, aggregate these matrices (here: subtraction), and use the resulting adjacency matrix as input to a downstream model, typically an MPNN, for the final predictions task. On the backward pass, the gradients of the loss $\ell$ regarding the parameters $\boldsymbol{\theta}$ are approximated through the derivative of the exactly- $k$ marginals in the direction of the gradients of the point-wise loss $\ell$ regarding the sampled adjacency matrix. We use recent work to make the computation of these marginals exact and differentiable, reducing both
bias and variance.
2.7 Example graph from the Trees-LeafCount test dataset with radius 4 (left). PRMPNN rewires the graph, allowing the downstream MPNN to obtain the label information from the leaves in one massage-passing step (right).35
2.8 Test accuracy of our rewiring method on the Trees-NeighborsMatch [AY21] dataset, compared to the reported accuracies from [MGM23].
3.1 Comparisons of different gradient estimators for point-wise loss $\ell$ being L1 and L2 loss applied to Gaussian and Poisson variables are conducted. To evaluate the direction of the gradient, we utilized the cosine distance, defined as 1 - cosine similarity. The bias, variance, and error of the estimators are assessed using a sample size of 10,000. The bias and variance for Marginal Expectation applied to Poison random variable are extremely close to zero but not identically zero.49
3.2 Comparison of gradient estimators for VAE with constrained latent space. Negative log-likelihood (NLL), negative ELBO (NELBO), and reconstruction loss (RL) are averaged over 5 trials.
3.3 The first row displays the original MNIST images, while the second row shows the constrained versions that are input to the VAEs. We compare the reconstructed images produced by various VAE models and their contained counterparts integrated with linear equality.
4.1 An example of how to compute the count probability in a dynamic programming manner. Assume that an instance-level classifier predicts three instances to have $p\left(\boldsymbol{y}_{1}=1\right)=0.1, p\left(\boldsymbol{y}_{2}=1\right)=0.2$, and $p\left(\boldsymbol{y}_{3}=1\right)=0.3$ respectively. The algorithm starts from the top-left cell and propagates the results down right. A cell has its probability $p\left(\sum_{j=0}^{i} \boldsymbol{y}_{j}=s\right)$ computed by inputs from $p\left(\sum_{j=0}^{i-1} \boldsymbol{y}_{j}=s\right)$ weighted by $p\left(\boldsymbol{y}_{i}=0\right)$, and $p\left(\sum_{j=0}^{i-1} \boldsymbol{y}_{j}=s-1\right)$ weighted by $p\left(\boldsymbol{y}_{i}=1\right)$ respectively, as indicated by the arrows.
4.2 MIL MNIST dataset experiments with decreased numbers of training bags and lower bag size. Left: bag sizes sampled from $\mathcal{N}(10,2)$; Right: bag sizes sampled from $\mathcal{N}(5,1)$. We plot the mean test AUC (aggregated over 3 trials) with standard errors for 4 bag sizes. Best viewed in color.
4.3 A test bag from our MIL experiments, where we set only the digit 9 as a positive instance. Highlighted in red are digits identified to be positive with corresponding probability beneath.
4.4 MNIST17 setting for PU Learning: We compute the average discrete distribution for CL and CVIR, over 5 test bags, each of which contain 100 instances. A ground truth binomial distribution of counts is also shown.
5.1 Feasible region (left) of formula $\Gamma$ with one player and primal graph (right) of formula $\Gamma$ with $n$ players from Example 1. ..... 82
5.2 Feasible region of SMT theory $\Gamma_{i}$ from Example 2 ..... 83
6.1 WMI runtime on independent model in Example 6. ..... 87
6.2 Primal graph and search tree for $\left(y \vee x_{1}\right) \wedge\left(y \vee x_{2}\right)$. ..... 88
6.3 Primal graph and feasible region from Example 7. ..... 89
6.4 Continuous search trees for $\theta_{2}$ from Example 7. ..... 926.5 Piecewise polynomial $p(y)$ as defined in Proposition 17 for theory $\theta_{2}$ from Example 7,whose integration is $\operatorname{MI}\left(\theta_{2}\right)$. The two polynomials $p_{-}(y)$ and $p_{+}(y)$ are unknown, butwe can recover them from a finite number of points.93
6.6 (a)-(c) MI execution time on $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ with tree primal graphs. (d)-(f) Example tree primal graphs. ..... 98
6.7 Runtime and primal graph for house price model. ..... 101
6.8 The current landscape of classes of WMI problems. We enlarge the boundariesof tractable WMI inference from treeMI to treeWMI and prove the hardness of 2 MIand 2WMI.104
6.9 Factor graph (left) of formula $\Gamma$ with two players and piecewise polynomial mes-sages (right) sent from the three factor nodes to variable node $X_{\mathcal{T}}$ when solving theWMI in Example 3 by MP-WMI.1086.10 Results of the comparison between MP-WMI, WMI-PA and F-XSDD on WMI prob-lems with tree dependencies. In this setting, MP-WMI remarkably scales to problemshaving up to 60 variables on STAR, while solving SNOW and PATH problems hav-ing up to 90 variables, considerably "raising the bar" for the size of tractable WMIinference problems.114
6.11 Log-log plot of cumulative time (seconds, y -axis) for MP-WMI (orange, red, brown) and SMI (blue, green, purple) over STAR, SNOW and PATH primal graphs (see text) with 10,20 and 30 variables for increasing numbers of univariate and bivariate queries (x-axis). For every class, MP-WMI takes up to two order of magnitude less time when amortizing 100 queries, while being faster than SMI on a single query.
7.1 Primal graph $\mathcal{G}_{\Delta}$ used for the \#P-hardness reduction in Theorem 8. We construct the corresponding formula $\Delta$ such that $\mathcal{G}_{\Delta}$ has maximum diameter (it is a chain). We graphically augment graph $\mathcal{G}_{\Delta}$ by introducing blue nodes to indicate that integers $s_{i}$ in set $S$ are contained in clauses between two variables.
7.2 Primal graph used for \#P-hardness reduction in Theorem 7. We also put blue nodes to indicate that integer $s_{i}$ 's in set $S$ are contained in some clauses and that model integration over some cliques is the sum of some $s_{i}$ 's.
7.3 Average integrated absolute errors (left) and times in seconds (right) for 5 problems of increasing size ( $n$, x-axis) for ReCoIn and competitors. Number of compensating literals (2-4) or samples used are in parentheses. Mean values per problem size are connected by a line.
8.1 The integral surface of (a) the expected prediction in BMA, and (b) our proposed approximation. Both are highly non-convex and multi-modal. The z-axis is the weighted prediction $\boldsymbol{y} p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{w}) p(\boldsymbol{w} \mid \mathcal{D})$. Integration of (a) does not admit a closed-form solution, yet integration of $(\mathrm{b})$ is a close approximation that can be solved exactly and efficiently by WMI solvers.
8.2 Uncertainty estimates for regression. The red line is the ground truth. The dark blue line shows the predictive mean. The shaded region is the $90 \%$ confidence interval of the predictive distribution. For the same number of samples, (b) Ciber is closer than (a) small-sample HMC to (c) a highly accurate but slow HMC with a large number of samples.

### 8.3 Approximating the Gaussian distribution with a triangluar distribution.

8.4 Posterior predictive distributions in Bayesian linear regression. The $y$-axis shows the absolute difference between an estimated predictive distribution $p(y \mid \mathbf{x})$ and the ground-truth predictive distribution $q(y \mid \mathbf{x})$. Shaded regions are the $95 \%$ confidence interval.
8.5 KL divergence in Bayesian linear regression. The $x$-axis shows the number of samples the MC method uses for estimations, ranging from 50 to 150 . The blue curve shows the MC method, and the green dashed curve shows Ciber using 50 samples.

## LIST OF TABLES

2.1 Architectures of the three experiment settings. ..... 9
2.2 Results for three aspects with $k=10$ : test MSE and subset precision, both $\times 100$ ..... 22
2.3 Results for aspect Aroma: test MSE and subset precision, both $\times 100$, for $k \in\{5,10,15\}$. ..... 222.4 Comparison between PR-MPNN and baselines on three molecular property predic-tion datasets. We report results for PR-MPNN with different gradient estimators for$k$-subset sampling: Gumbel SoftSub-ST [MMT17a, JGP17a, XE19a], I-MLE [NMF21a],and Simple [AZN23a] and compare them with the base downstream model, and twograph transformer architectures. The variant using SIMPLE consistently outperformsthe base models and is competitive or better than the two graph transformers. Weuse green for the best model, blue for the second-best, and red for third. We notewith + Edge the instances where edge features are provided and with - Edge whenthey are not.38
2.5 Comparison between the base GIN model, PR-MPNN, and other more expressive models on the Exp dataset. ..... 39
2.6 Comparison between the base GIN model and probabilistic rewiring model on CSL dataset, w/o positional encodings. ..... 39
2.7 Comparison between PR-MPNN and other approaches as reported in [GRC23, KBM22,PMF21]. Our model outperforms existing approaches while keeping a lower variancein most of the cases, except for NCI1, where the WL Kernel is the best. We use greenfor the best model, blue for the second-best, and red for third.40
3.1 Summary of gradient estimators. The first block presents baseline estimators and the second presents our proposed ones. In the forward pass, we sample exactly from the constrained distribution by Proposition 7. In the backward pass, we use $\boldsymbol{m}(\boldsymbol{\theta})$ as a differntiable proxy.46
3.2 Comparison on VAE Generative Ability. The constrained VAE models achieve similar or better performance in terms of their generative ability while strictly satisfying the constraints. The unconstrained counterparts have a high constraint violation rate.
3.3 The prediction performances of different models for estimating partial charges on metal ions are presented. Comparing to the baseline MPNN (variance), both the closed-form loss function and Likelihood loss function yield superior Mean Absolute Deviation (MAD) results. Moreover, ensemble methods (second block) notably boost the predictive performance of all.57
4.1 A comparison of the tasks considered in the three weakly supervised settings, LLP (cf. Section 4.1.1), MIL (cf. Section 4.1.2) and PU learning (cf. Section 4.1.3), against the classical fully supervised setting for binary classification, using digits from the MNIST dataset.
4.2 A summary of the labels and objective functions for all the settings considered in this chapter.
4.3 LLP results across different bag sizes. We report the mean and standard deviation of the test AUC over 5 seeds for each setting. The highest metric for each setting is shown in boldface.
4.4 MIL experiment on the MNIST dataset. Each block represents a different distribution from which we draw bag sizes-First Block: $\mathcal{N}(10,2)$, Second Block: $\mathcal{N}(50,10)$, Third Block: $\mathcal{N}(100,20)$. We run each experiment for 3 runs and report mean test AUC with standard error. The highest metric for each setting is shown in boldface.
4.5 MIL: We report mean test accuracy, AUC, F1, precision, and recall averaged over 5 runs with std. error on the Colon Cancer dataset. The highest value for each metric is shown in boldface.
4.6 PU Learning: We report accuracy and standard deviation on a test set of unlabeled data, which is aggregated over 3 runs. The results from CVIR, nnPU, and uPU are aggregated over 10 epochs, as in [GWS21], while we choose the single best epoch based on validation for our approaches. The highest metric for each setting is shown in boldface.77
8.1 Average test log likelihood for the small UCI regression task. ..... 152
8.2 Average test RMSE for the small UCI regression task. ..... 153
8.3 Average test log likelihood for the large UCI regression task. ..... 154
8.4 Average test RMSE for the large UCI regression task. ..... 155
8.5 Average test performance for image classification tasks on CIFAR-10 and CIFAR-100. ..... 156
8.6 Average test performance for image transfer learning tasks. ..... 156

## ACKNOWLEDGMENTS

First and foremost, I would like to express my deepest gratitude to my advisor Guy Van den Broeck. I am incredibly fortunate and honored to have him as my advisor. Throughout my research journey, his wisdom and guidance have made the experience not only productive but also genuinely enjoyable. Guy's insightful suggestions and constructive feedback have been critical in shaping me into an independent researcher. His constant encouragement has helped clear my anxieties and given me the strength to carry on. He has instilled in me a passion for pursuing rigorous, elegant, and unexpected research. I will always miss our spontaneous conversations driven by intellectual curiosity and the inspiring journey we've shared. Thank you for everything!

I also want to extend my heartfelt thanks to my other committee members: Todd Millstein, Cho-Jui Hsieh, and Kai-Wei Chang, as well as Mathias Niepert who is in my oral qualification exam committee. Their feedback has played a crucial role in refining my research. Also, their shared experiences and insights have been instrumental in guiding my career planning.

I am extremely lucky to have wonderful collaborators, labmates, and friends: YooJung, Steven, Tal, Yitao, Chenyu, Antonio, Pasha, Kareem, Paolo, Honghua, Meihua, Anji, William, Nikil, Poorva, Daniel, Ellie, Paolo, Arthur, Yujia, Chi, Andy, Haiying, Tong, Chen, Jieyu (and her two adorable cats), Tao, Kuan-Hao, Liunian, Fan, Zijie, Da and so on. I will miss the lovely StarAI Lab.

My heartfelt appreciation goes out to Yuhang, my partner, best friend, and soulmate, by my side through all of life's ups and downs. With you, even the most challenging moments are less scary, and the joyful ones are filled with even more happiness.

Finally, I would like to express my heartfelt gratitude to my family. To my dad, Zhiqiang, and my mom, Yongyi, your unconditional love and unwavering support have been my greatest sources of strength and comfort. Thank you for always believing in me and standing by my side through every step of this journey. I also want to thank my cousin Lisi. She is the first PhD in the family and a true guiding star of my life who has inspired me to follow a similar path. I am constantly motivated by her wisdom, resilience and perseverance.

## PUBLICATIONS

Chendi Qian, Andrei Manolache, Kareem Ahmed, Zhe Zeng, Guy Van den Broeck, Mathias Niepert, and Christopher Morris. Probabilistically rewired message-passing neural networks. In Proceedings of the Eleventh International Conference on Learning Representations (ICLR), 2024

Zhe Zeng and Guy Van den Broeck. Collapsed inference for Bayesian deep learning. In Advances in Neural Information Processing Systems (NeurIPS), 2023

Vinay Shukla, Zhe Zeng, Kareem Ahmed, and Guy Van den Broeck. A unified approach to count-based weakly-supervised learning. In Advances in Neural Information Processing Systems (NeurIPS), 2023

Kareem Ahmed, Zhe Zeng, Mathias Niepert, and Guy Van den Broeck. SIMPLE: A gradient estimator for k-subset sampling. In Proceedings of the Eleventh International Conference on Learning Representations (ICLR), 2023

Wenzhe Li, Zhe Zeng, Antonio Vergari, and Guy Van den Broeck. Tractable computation of expected kernels. In Proceedings of the 37th Conference on Uncertainty in Artificial Intelligence
(UAI), 2021

Zhe Zeng, Paolo Morettin, Fanqi Yan, Antonio Vergari, and Guy Van den Broeck. Probabilistic inference with algebraic constraints: Theoretical limits and practical approximations. In Advances in Neural Information Processing Systems (NeurIPS), 2020.

Zhe Zeng, Paolo Morettin, Fanqi Yan, Antonio Vergari, and Guy Van den Broeck. Scaling up hybrid probabilistic inference with logical and arithmetic constraints via message passing. In Proceedings of the 37th International Conference on Machine Learning (ICML), 2020

Zhe Zeng and Guy Van den Broeck. Efficient search-based weighted model integration. In Proceedings of the 35th Conference on Uncertainty in Artificial Intelligence (UAI), 2019

## CHAPTER 1

## Introduction

Artificial Intelligence (AI), particularly deep learning, has become pervasive in our daily lives and safety-critical systems. Despite their profound impact, there are increasing concerns about the reliability and trustworthiness of AI models. Therefore, developing trustworthy AI is a pressing need and of great importance for the sustainable advancement of AI technologies.

Even though neural networks, the deep learning models, excel in pattern recognition and scalability, they inherently struggle to provide guarantees, explain their decisions, and incorporate domain knowledge - critical elements for ensuring trustworthiness [LQL23]. On the other hand, there is a branch of AI called symbolic AI that allows to define models in a transparent way with high explainability and to express the domain knowledge using constraints. Constraints can take various forms from graph structures to logical, arithmetic, and physical ones. Such constraints provide a high-level abstractions of the real-world knowledge, such as safety regulations and scientific phenomenons. Techniques for reasoning over constraints can provide guarantees that the model decisions are consistent with the knowledge.

Neurosymbolic AI emerges as a promising field that combine the good from both worlds for the next-generation AI models to enable and support decision-making in the presence of probabilistic uncertainty and symbolic knowledge [GL23, MDM24]. It has been shown to be powerful in solving challenging tasks that are previously out of the reach of pure deep learning models or pure symbolic models [TWL24]. The fusion of neural network-based learning with symbolic reasoning has the potential to revolutionize not only the capabilities of AI models but also their societal impact and roles in scientific discoveries. However, there are some unique challenges in
neurosymbolic AI when combining neural networks with constraints. One challenge is that most symbolic components are non-differentiable which prevents their integration into the gradientbased learning process of the deep learning models; The other challenge is that reasoning over the symbolic constraints is \#P-hard in general, meaning that the computation can take an unfeasibly long time during the inference process and hinder reliable and efficient reasoning over the neurosymbolic pipeline.

Trustworthy AI can be achieved by neurosymbolic methods in various ways. For example, when building neurosymbolic models, constraints can be enforced in the output space of a deep generative model which is a neural network model for data generation such that the generated data follows certain domain knowledge. Further, queries for trustworthiness such as "how confident the model is for its decision" can be answered by probabilistic reasoning over the neurosymbolic model. Besides, neurosymbolic methods allow for effective learning by restricting the training to the feasible space defined by the constraints, that is, for any assignments violating the constraints, they are not considered and always assigned low or zero probabilities.

My research aims to address the aforementioned two fundamental challenges in neurosymbolic AI and to build trustworthy AI models using neurosymbolic methods. The main contributions of this thesis are two-fold, which have been shown to achieve trustworthiness including explainability, uncertainty quantification and domain-knowledge incorporation:
i) Our proposed methods enable differentiable learning in the presence of symbolic constraints by building gradient estimators for the symbolic component, and it allows flexible integration of the constraints into any parts of the neural network architecture.
ii) We propose novel reasoning algorithms over complex constraints that can be applied to perform reliable reasoning over neural networks; this is achieved by translating part of the neural networks into symbolic representations, revealing a deep connection between neural and symbolic reasoning.

This thesis is organized as follows. It consists of two parts for neurosymbolic learning and reasoning respectively. Chapter 2, based on [AZN23b] provides how a $k$-subset constraint can be
enforced into neural network architectures and how the full pipeline can be effectively trained by our proposed gradient estimators such that the constraint can be differentiated. The proposed gradient estimator is further leveraged to derive a probabilistic graph rewiring algorithm for graph neural networks [QMA24].

Chapter 3 generalizes the same design philosophy of gradient estimators from the $k$-subset constraints over discrete domains to the linear equality defined over continuous domains, which allows introducing inductive bias into the data generation process of deep generative models, based on a work under review [LSB24].

While the previous two chapters consider hard constraints, that is, constraints with guaranteed satisfaction, Chapter 4 focuses on integrating the soft constraints into the training objectives in count-based weakly supervised learning tasks [SZA23].

The second part focuses on reasoning over complex constraints. Note that in this thesis, we use treat reasoning and inference as synonyms. Chapter 5 provides sufficient backgrounds in weighted model integration (WMI), a unified framework for probabilistic reasoning over algebraic constraints, including its formulations and motivations.

Chapter 6 introduce two exact WMI solvers that achieves state-of-art reasoning performance by exploiting problem structures. While the first solver is search-based and is the first exact WMI solver that comes with guarantees on tractability [ZB19], the second solver, built using a message-passing scheme, is able to solver a strictly larger WMI problem class in a tractable way and is the first WMI solver that can perform amortized inference [ZMY20a].

Chapter 7 is based on [ZMY20b] and starts with theoretical analysis of the inherent hardness of WMI problems. We derive the computational complexity analysis on WMI inference using graph structures to characterize its tractability boundary and discover the largest tractable WMI model class so far. Based on the theoretical understanding, we further propose an approximate WMI by the idea of performing exact inference on an approximate model.

Chapter 8, based on [ZB23b], shows that by combining exact WMI solvers with the collapsed
inference scheme from statistical machine learning, they scale to perform uncertain quantification for image classification tasks with large Bayesian neural networks, outperforming strong baselines in Bayesian deep learning including sampling and variational inference.

Chapter 9, the final one, summarizes the thesis and discuss about future research directions.

## Part I

## Learning

## CHAPTER 2

## Differentiable Learning under Constraints

This chapter tackles a fundamental challenge in Neurosymbolic AI: how to enforce constraints within the architecture of neural networks while still allowing for end-to-end training. Specifically, we consider a $k$-subset constraint which is ubiquitous in machine learning. The incorporation of constraints is problematic and requires relaxations or approximations in general which can harm the learning efficiency and model performance. We address this issue by minimizing such relaxations such that constraints can be incorporated into any part of the neural network with guaranteed satisfaction and meanwhile the constrained neural networks are effectively trained.

### 2.1 Background

$k$-subset sampling, sampling a subset of size $k$ of $n$ variables, is omnipresent in machine learning. It lies at the core of many fundamental problems that rely upon learning sparse features representations of input data, including stochastic high-dimensional data visualization [Maa09], parametric $k$-nearest neighbors [GWZ18], learning to explain [CSW18], discrete variational autoencoders [Rol17], and sparse regression, to name a few. All such tasks involve optimizing an expectation of an objective function with respect to a latent discrete distribution parameterized by a neural network, which are often assumed intractable. Score-function estimators offer a cloyingly simple solution: rewrite the gradient of the expectation as an expectation of the gradient, which can subsequently be estimated using a finite number of samples offering an unbiased estimate of the gradient. Simple as it is, score-function estimators suffer from very high variance which can interfere with training. This provided the impetus for other, low-variance, gradient
estimators, chief among them are those based on the reparameterization trick, which allows for biased, but low-variance gradient estimates. The reparameterization trick, however, does not allow for a direct application to discrete distributions thereby prompting continuous relaxations, e.g. Gumbel-softmax [JGP17b, MMT17b], that allow for reparameterized gradients w.r.t the parameters of a categorical distribution. Reparameterizable subset sampling [XE19b] generalizes the Gumbel-softmax trick to $k$-subsets which while rendering $k$-subset sampling amenable to backpropagation at the cost of introducing bias in the learning by using relaxed samples.

### 2.2 Gradient Estimator for $k$-subset Sampling

We set out with the goal of avoiding all such relaxations. Instead, we fall back to discrete sampling on the forward pass. On the backward pass, we reparameterize the gradient of the loss function with respect to the samples as a function of the exact marginals of the $k$-subset distribution. Computing the exact conditional marginals is, in general, intractable [Rot96a]. We give an efficient algorithm for computing the $k$-subset probability, and show that the conditional marginals correspond to partial derivatives, and are therefore tractable for the $k$-subset distribution. We show that our proposed gradient estimator for the $k$-subset distribution, coined SIMPLE, is reminiscent of the straight-through (ST) Gumbel estimator when $k=1$, with the gradients taken with respect to the unperturbed marginals. We empirically demonstrate that Simple exhibits lower bias and variance compared to other known gradient estimators, including the ST Gumbel estimator in the case $k=1$.

### 2.2.1 Problem Statement and Motivation

We consider models described by the equations

$$
\begin{equation*}
\boldsymbol{\theta}=h_{\boldsymbol{v}}(\boldsymbol{x}), \quad \boldsymbol{z} \sim p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right), \quad \hat{\boldsymbol{y}}=f_{\boldsymbol{u}}(\boldsymbol{z}, \boldsymbol{x}), \tag{2.1}
\end{equation*}
$$

where $\boldsymbol{x} \in \mathcal{X}$ and $\hat{\boldsymbol{y}} \in \mathcal{Y}$ denote feature inputs and target outputs, respectively, $h_{\boldsymbol{v}}: \mathcal{X} \rightarrow \Theta$ and $f_{u}: \mathcal{Z} \times \mathcal{X} \rightarrow \mathcal{Y}$ are smooth, parameterized maps and $\boldsymbol{\theta}$ are logits inducing a distribution over



Figure 2.1: A comparison of the bias and variance of the gradient estimators (left) and the average and standard deviation of the cosine distance of a single-sample gradient estimate to the exact gradient. We used the cosine distance, defined as (1- cosine similarity), in place of the euclidean distance as we only care about the direction of the gradient, not magnitude. The bias, variance and error were estimated using a sample of size 10,000 . The details of this experiment are provided in Section 2.2.3.1.
the latent binary vector $\boldsymbol{z}$. The induced distribution $p_{\boldsymbol{\theta}}(\boldsymbol{z})$ is defined as

$$
\begin{equation*}
p_{\boldsymbol{\theta}}(\boldsymbol{z})=\prod_{i=1}^{n} p_{\theta_{i}}\left(z_{i}\right), \text { with } p_{\theta_{i}}\left(z_{i}=1\right)=\operatorname{sigmoid}\left(\theta_{i}\right) \text { and } p_{\theta_{i}}\left(z_{i}=0\right)=1-\operatorname{sigmoid}\left(\theta_{i}\right) . \tag{2.2}
\end{equation*}
$$

The goal of our stochastic latent layer is not to simply sample from $p_{\boldsymbol{\theta}}(\boldsymbol{z})$, which would yield samples with a Hamming weight between 0 and $n$ (i.e., with an arbitrary number of ones). Instead, we are interested in sampling from the distribution restricted to samples with a Hamming weight of $k$, for any given $k$. That is, we are interested in sampling from the conditional distribution $p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)$.

Conditioning the distribution $p_{\boldsymbol{\theta}}(\boldsymbol{z})$ on this $k$-subset constraint introduces intricate dependencies between each of the $z_{i}$ 's. The probability of sampling any given $k$-subset vector $\boldsymbol{z}$, therefore, becomes

$$
p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)=p_{\boldsymbol{\theta}}(\boldsymbol{z}) / p_{\boldsymbol{\theta}}\left(\sum_{i} z_{i}=k\right) \cdot \llbracket \sum_{i} z_{i}=k \rrbracket
$$

where $\llbracket \rrbracket$ denotes the indicator function. In other words, the probability of sampling each $k$ subset is re-normalized by $p_{\boldsymbol{\theta}}\left(\sum_{i} z_{i}=k\right)$ - the probability of sampling exactly $k$ items from the unconstrained distribution induced by encoder $h_{\boldsymbol{v}}$. The quantity $p_{\boldsymbol{\theta}}\left(\sum_{i} z_{i}=k\right)=\sum_{\mathbf{z}} p_{\theta}(\mathbf{z})$.


Figure 2.2: The problem setting considered in this section. On the forward pass, a neural network $h_{v}$ outputs $\boldsymbol{\theta}$ parameterizing a discrete distribution over subsets of size $k$ of $n$ items, i.e., the $k$-subset distribution. We sample exactly, and efficiently, from this distribution, and feed the samples to a downstream neural network. On the backward pass, we approximate the true gradient by the product of the derivative of marginals and the gradient of the sample-wise loss.

| TASK | MAP $h_{\boldsymbol{v}}$ | MAP $f_{\boldsymbol{u}}$ | Loss $\ell$ |
| :---: | :---: | :---: | :---: |
| Discrete VAE (Sec. 2.2.3.2) | Encoder | Decoder | ELBO |
| Learn To Explain (Sec. 2.2.3.3) | Embedding | Regression | RMSE |
| Sparse Regression (Sec. 2.2.3.4) | Identity | Linear Regression | RMSE |

Table 2.1: Architectures of the three experiment settings.
$\llbracket \sum_{i} z_{i}=k \rrbracket$ appears to be intractable. We show that not to be the case, providing a tractable algorithm for computing it.

Given a set of samples $\mathcal{D}$, we are concerned with learning the parameters $\boldsymbol{\omega}=(\boldsymbol{v}, \boldsymbol{u})$ of the architecture in (3.1) through minimizing the training error $L$, which is the expected loss:

$$
\begin{equation*}
L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})=\mathbb{E}_{\boldsymbol{z} \sim p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)}\left[\ell\left(f_{\boldsymbol{u}}(\boldsymbol{z}, \boldsymbol{x}), \boldsymbol{y}\right)\right] \quad \text { with } \boldsymbol{\theta}=h_{\boldsymbol{v}}(\boldsymbol{x}), \tag{2.3}
\end{equation*}
$$

where $\ell: \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}^{+}$is a point-wise loss function. This formulation, illustrated in Figure 2.2, is general and subsumes many settings. Different choices of mappings $h_{v}$ and $f_{u}$, and sample-wise loss $\ell$ define various tasks. Table 2.1 presents some example settings used in our experimental evaluation.

Learning then requires computing the gradient of $L$ w.r.t. $\boldsymbol{\omega}=(\boldsymbol{v}, \boldsymbol{u})$. The gradient of $L$ w.r.t.
$\boldsymbol{u}$ is

$$
\begin{equation*}
\nabla_{\boldsymbol{u}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})=\mathbb{E}_{\boldsymbol{z} \sim p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)}\left[\partial_{\boldsymbol{u}} f_{\boldsymbol{u}}(\boldsymbol{z}, \boldsymbol{x})^{\top} \nabla_{\hat{\boldsymbol{y}}} \ell(\hat{\boldsymbol{y}}, \boldsymbol{y})\right], \tag{2.4}
\end{equation*}
$$

where $\hat{y}=f_{u}(\boldsymbol{z}, \boldsymbol{x})$ is the decoding of a latent sample $\boldsymbol{z}$. Furthermore, the gradient of $L$ w.r.t. $\boldsymbol{v}$ is

$$
\begin{equation*}
\nabla_{\boldsymbol{v}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})=\partial_{\boldsymbol{v}} h_{\boldsymbol{v}}(\boldsymbol{x})^{\top} \nabla_{\boldsymbol{\theta}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega}), \tag{2.5}
\end{equation*}
$$

where $\nabla_{\boldsymbol{\theta}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega}):=\nabla_{\boldsymbol{\theta}} \mathbb{E}_{\boldsymbol{z} \sim p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)}\left[\ell\left(f_{\boldsymbol{u}}(\boldsymbol{z}, \boldsymbol{x}), \hat{\boldsymbol{y}}\right)\right]$, the loss' gradient w.r.t. the encoder.
One challenge lies in computing the expectation in (3.2) and (3.3), which has no known closedform solution. This necessitates a Monte-Carlo estimate via sampling from $p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)$.

A second, and perhaps more substantial hurdle lies in computing $\nabla_{\boldsymbol{\theta}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})$ in (3.4) due to the non-differentiable nature of discrete sampling. One could rewrite $\nabla_{\boldsymbol{\theta}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})$ as

$$
\nabla_{\boldsymbol{\theta}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})=\mathbb{E}_{\boldsymbol{z} \sim p_{\theta}\left(\mathbf{z} \mid \sum_{i} z_{i}=k\right)}\left[\ell\left(f_{\boldsymbol{u}}(\boldsymbol{z}, \boldsymbol{x}), \boldsymbol{y}\right) \nabla_{\boldsymbol{\theta}} \log p_{\theta}\left(\mathbf{z} \mid \sum_{i} z_{i}=k\right)\right]
$$

which is known as the REINFORCE estimator [Wil92], or the score function estimator (SFE). It is typically avoided due to its notoriously high variance, despite its apparent simplicity. Instead, typical approaches [XE19b, PR18] reparameterize the samples as a deterministic transformation of the parameters, and some independent standard Gumbel noise, and relaxing the deterministic transformation, the top- $k$ function in this case, to allow for backpropagation.

### 2.2.2 Simple: Subset Implicit Likelihood Estimation

Our goal is to build a gradient estimator for $\nabla_{\boldsymbol{\theta}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})$. We start by envisioning a hypothetical sampling-free architecture, where the downstream neural network $f_{u}$ is a function of the marginals, $\boldsymbol{\mu}:=\mu(\boldsymbol{\theta}):=\left\{p_{\boldsymbol{\theta}}\left(z_{j} \mid \sum_{i} z_{i}=k\right)\right\}_{j=1}^{n}$, instead of a discrete sample $\boldsymbol{z}$, resulting in a loss $L_{m}$ s.t.

$$
\begin{equation*}
\nabla_{\boldsymbol{\theta}} L_{m}(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})=\partial_{\boldsymbol{\theta}} \mu(\boldsymbol{\theta})^{\top} \nabla_{\boldsymbol{\mu}} \ell_{m}\left(f_{\boldsymbol{u}}(\boldsymbol{\mu}, \boldsymbol{x}), \boldsymbol{y}\right) \tag{2.6}
\end{equation*}
$$

When the marginals $\mu(\boldsymbol{\theta})$ can be efficiently computed and differentiated, such a hypothetical pipeline can be trained end-to-end. Furthermore, [Dom10] observed that, for an arbitrary loss
function $\ell_{m}$ defined on the marginals, the Jacobian of the marginals w.r.t. the logits is symmetric:

$$
\begin{equation*}
\nabla_{\boldsymbol{\theta}} L_{m}(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})=\partial_{\boldsymbol{\theta}} \mu(\boldsymbol{\theta})^{\top} \nabla_{\boldsymbol{\mu}} \ell_{m}\left(f_{\boldsymbol{u}}(\boldsymbol{\mu}, \boldsymbol{x}), \boldsymbol{y}\right)=\partial_{\boldsymbol{\theta}} \mu(\boldsymbol{\theta}) \nabla_{\boldsymbol{\mu}} \ell_{m}\left(f_{\boldsymbol{u}}(\boldsymbol{\mu}, \boldsymbol{x}), \boldsymbol{y}\right) . \tag{2.7}
\end{equation*}
$$

Consequently, computing the gradient of the loss w.r.t. the logits, $\nabla_{\boldsymbol{\theta}} L_{m}(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})$, reduces to computing the directional derivative, or the Jacobian-vector product, of the marginals w.r.t. the logits in the direction of the gradient of the loss. This offers an alluring opportunity: the conditional marginals characterize the probability of each $z_{i}$ in the sample, and could be thought of as a differentiable proxy for the samples. Specifically, by reparameterizing $z$ as a function of the conditional marginal $\boldsymbol{\mu}$ under approximation $\partial_{\mu} \boldsymbol{z} \approx \mathrm{I}$ as proposed by [NMF21b], and using the straight-through estimator for the gradient of the sample w.r.t. the marginals on the backward pass, we approximate our true $\nabla_{\boldsymbol{\theta}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})$ as

$$
\begin{equation*}
\nabla_{\boldsymbol{\theta}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega}) \approx \partial_{\boldsymbol{\theta}} \mu(\boldsymbol{\theta}) \nabla_{\boldsymbol{z}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega}), \tag{2.8}
\end{equation*}
$$

where the directional derivative of the marginals can be taken along any downstream gradient, rendering the whole pipeline end-to-end learnable, even in the presence of non-differentiable sampling.

Now, estimating the gradient of the loss w.r.t. the parameters can be thought of as decomposing into two sub-problems: (P1) Computing the derivatives of conditional marginals $\partial_{\boldsymbol{\theta}} \mu(\boldsymbol{\theta})$, which requires the computation of the conditional marginals, and (P2) Computing the gradient of the loss w.r.t. the samples $\nabla_{z} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})$ using sample-wise loss, which requires drawing exact samples. These two problems are complicated by conditioning on the $k$-subset constraint, which introduces intricate dependencies to the distribution, and is infeasible to solve naively, e.g. by enumeration. We will show simple, efficient, and exact solutions to each problem, at the heart of which is the insight that we need not care about the variables' order, only their sum, introducing symmetries that simplify the problem.

### 2.2.2.1 Derivatives of Conditional Marginals

In many probabilistic models, marginal inference is \#P-hard [Rot96a, ZMY20b]. However, we observe that it is not the case for the $k$-subset distribution. We notice that the conditional marginals correspond to the partial derivatives of the log-probability of the $k$-subset constraint. To see this, note that the derivative of a multi-linear function with respect to a single variable retains all the terms referencing that variable, and drops all other terms; this corresponds exactly to the unnormalized conditional marginals. By taking the derivative of the log-probability, this introduces the $k$-subset probability in the denominator, leading to the conditional marginals. Intuitively, the rate of change of the $k$-subset probability w.r.t. a variable only depends on that variable through its length $-k$ subsets.

Theorem 1. Let $p_{\boldsymbol{\theta}}\left(\sum_{j} z_{j}=k\right)$ be the probability of exactly- $k$ of the unconstrained distribution parameterized by $\operatorname{logits} \boldsymbol{\theta}$. Let $\alpha_{i}:=\log p_{\boldsymbol{\theta}}\left(z_{i}\right)$ denote the log marginals. For every variable $Z_{i}$, its conditional marginal is

$$
\begin{equation*}
p_{\boldsymbol{\theta}}\left(z_{i} \mid \sum_{j} z_{j}=k\right)=\frac{\partial}{\partial \alpha_{i}} \log p_{\boldsymbol{\theta}}\left(\sum_{j} z_{j}=k\right) \tag{2.9}
\end{equation*}
$$

To establish the tractability of the above computation of the conditional marginals, we need to show that the probability of the exactly- $k$ constraint $p_{\boldsymbol{\theta}}\left(\sum_{i} z_{i}=k\right)$ can be obtained tractably, which we demonstrate next.

Proposition 1. The probability $p_{\boldsymbol{\theta}}\left(\sum_{i} z_{i}=k\right)$ of sampling exactly $k$ items from the unconstrained distribution $p_{\boldsymbol{\theta}}(\boldsymbol{z})$ over $n$ items as in Equation 2.2 can be computed exactly in time $\mathcal{O}(n k)$.

Proof. Our proof is constructive. As a base case, consider the probability of sampling $k=-1$ out of $n=0$ items. We can see that the probability of such an event is 0 . As a second base case, consider the probability of sampling $k=0$ out of $n=0$ items. We can see that the probably of such an event is 1 . Now assume that we are given the probability $p_{\boldsymbol{\theta}}\left(\sum_{i}^{n-1} z_{i}=k^{\prime}\right)$, for $k^{\prime}=0, \ldots, k$, and we are interested in computing $p_{\boldsymbol{\theta}}\left(\sum_{i}^{n} z_{i}=k\right)$. By the partition theorem, we

```
Algorithm 1 PrExactlyk \((\boldsymbol{\theta}, n, k)\)
Input: The logits \(\boldsymbol{\theta}\) of the distribution, the num-
    ber of variables \(n\), and the subset size \(k\)
Output: \(p_{\boldsymbol{\theta}}\left(\sum_{i} z_{i}=k\right)\)
    \(/ / a[i, j]=p_{\boldsymbol{\theta}}\left(\sum_{m=1}^{i} z_{m}=j\right)\) for all \(i, j\)
    initialize \(a\) to be 0 everywhere
    \(a[0,0]=1 / / p_{\theta}\left(\sum_{m=1}^{0} z_{m}=0\right)=1\)
    for \(i=1\) to \(n\) do
        for \(j=0\) to \(k\) do
            // cf. constructive proof of Prop. 1
        \(a[i, j]=a[i-1, j] \cdot p_{\theta_{i}}\left(z_{i}=0\right)\)
            \(+a[i-1, j-1] \cdot p_{\theta_{i}}\left(z_{i}=1\right)\)
    return \(a[n, k]\)
```

can see that

$$
p_{\boldsymbol{\theta}}\left(\sum_{i}^{n} z_{i}=k\right)=p_{\boldsymbol{\theta}}\left(\sum_{i}^{n-1} z_{i}=k\right) \cdot p_{\theta_{n}}\left(z_{n}=0\right)+p_{\boldsymbol{\theta}}\left(\sum_{i}^{n-1} z_{i}=k-1\right) \cdot p_{\theta_{n}}\left(z_{n}=1\right)
$$

as events $\sum_{i}^{n-1} z_{i}=k$ and $\sum_{i}^{n-1} z_{i}=k-1$ are disjoint and, for any $k$, partition the sample space. Intuitively, for any $k$ and $n$, we can sample $k$ out of $n$ items by choosing $k$ of $n-1$ items, and not the $n$-th item, or choosing $k-1$ of $n-1$ items, and the $n$-th item. The above process gives rise to Algorithm 1, which returns $p_{\boldsymbol{\theta}}\left(\sum_{i} z_{i}=k\right)$ in time $\mathcal{O}(n k)$.

By the construction described above, we obtain a closed-form $p_{\boldsymbol{\theta}}\left(\sum_{i}^{n} z_{i}=k\right)$, which allows us to compute conditional marginals $p_{\boldsymbol{\theta}}\left(z_{i} \mid \sum_{j} z_{j}=k\right)$ by Theorem 1 via auto-differentiation. This further allows the computation of the derivatives of conditional marginals

$$
\partial_{\boldsymbol{\theta}} \mu(\boldsymbol{\theta})_{i}=\partial_{\boldsymbol{\theta}} p_{\boldsymbol{\theta}}\left(z_{i} \mid \sum_{j} z_{j}=k\right)
$$

to be amenable to auto-differentiation, solving problem (P1) exactly and efficiently.

### 2.2.2.2 Gradients of Loss w.r.t. Samples

As alluded to in Section 2.2.2, we approximate $\nabla_{\boldsymbol{\theta}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})$ by the directional derivative of the marginals along the gradient of the loss w.r.t. discrete samples $\boldsymbol{z}, \nabla_{\boldsymbol{z}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})$, where $\boldsymbol{z}$ is drawn from the $k$-subset distribution $p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)$. What remains is to estimate the loss, necessitating faithful sampling from the $k$-subset distribution, which might appear daunting.

Exact $k$-subset Sampling Next we show how to sample exactly from the $k$-subset distribution $p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)$. We start by sampling the variables in reverse order, that is, we sample $z_{n}$ through $z_{1}$. The main intuition being that, having sampled $\left(z_{n}, z_{n-1}, \cdots, z_{i+1}\right)$ with a Hamming weight of $k-j$, we sample $Z_{i}$ with a probability of choosing $k-j$ of $n-1$ variables and the $n$-th variable given that we choose $k-j+1$ of $n$ variables. We formalize our intuition below.

Proposition 2. Let Sample be defined as in Algorithm 2. Given $n$ random variables $Z_{1}, \cdots, Z_{n}$, a subset size $k$, and a $k$-subset distribution $p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)$ parameterized by $\log$ probabilities $\boldsymbol{\theta}$, Algorithm 2 draws exact samples from $p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)$ in time $\mathcal{O}(n)$.

Proof. Assume that variables $Z_{n}, \cdots, Z_{i+1}$ are sampled and have their values to be $z_{n}, \cdots, z_{i+1}$ with $\sum_{m=i+1}^{n} z_{m}=k-j$. By Algorithm 2 we have that the probability for sampling $Z_{i}$ is

$$
\begin{aligned}
& p_{\text {Sample }}\left(z_{i}=1\right.\left.\mid z_{n}, \cdots, z_{i+1}\right)=\frac{p_{\boldsymbol{\theta}}\left(\sum_{m=i}^{n} z_{m}=k-j+1 \mid \sum_{m} z_{m}=k\right) p_{\theta_{i}}\left(z_{i}=1\right)}{p_{\boldsymbol{\theta}}\left(\sum_{m=i+1}^{n} z_{m}=k-j \mid \sum_{m} z_{m}=k\right)} \\
&= \frac{p_{\boldsymbol{\theta}}\left(\sum_{m=i+1}^{n} z_{m}=k-j \mid z_{i}=1, \sum_{m} z_{m}=k\right) p_{\theta_{i}}\left(z_{i}=1\right)}{p_{\boldsymbol{\theta}}\left(\sum_{m=i+1}^{n} z_{m}=k-j \mid \sum_{m} z_{m}=k\right)} \\
&=p_{\boldsymbol{\theta}}\left(z_{i}=1 \mid \sum_{m=i+1}^{n} z_{m}=k-j, \sum_{m} z_{m}=k\right) \text { (by Bayes' theorem) }
\end{aligned}
$$

It follows that samples drawn from Algorithm 2 distribute according to $p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)$.

### 2.2.2.3 Connection to Straight-Through Gumbel-Softmax

One might wonder if our gradient estimator reduces to the Straight-Through (ST) Gumbel-Softmax estimator, or relates to it in any way when $k=1$. On the forward pass, the ST Gumbel Softmax


Figure 2.3: Bias and variance of Simple and Gumbel Softmax over 10k samples
estimator makes use of the Gumbel-Max trick [MTM14], which states that we can efficiently sample from a categorical distribution by perturbing each of the logits with standard Gumbel noise, and taking the MAP, or more formally $\boldsymbol{z}=\operatorname{OneHot}\left(\arg \max _{i \in\{1, \ldots, k\}} \boldsymbol{\theta}_{i}+g_{i}\right) \sim p_{\boldsymbol{\theta}}$ where the $g_{i}$ 's are i.i.d $\operatorname{Gumbel}(0,1)$ samples, and OneHot encodes the sample as a binary vector.

Since arg max is non-differentiable, Gumbel-Softmax uses the perturbed relaxed samples, $\boldsymbol{y}=$ $\operatorname{Softmax}\left(\boldsymbol{\theta}+g_{i}\right)$ as a proxy for discrete samples $\boldsymbol{z}$ on the backward pass, using differentiable Softmax in place of the non-differentiable arg max, with the entire function returning $(\boldsymbol{z}-\boldsymbol{y})$. detach ()$+\boldsymbol{y}$ where detach ensures that the gradient flows only through the relaxed samples on the backward pass.

That is, just like Simple, ST Gumbel-Softmax returns exact, discrete samples. However, whereas SimpLe backpropagates through the exact marginals, ST Gumbel Softmax backpropagates through the perturbed marginals that result from applying the Gumbel-max trick. As can be seen in Figure 2.3, such a minor difference means that, empirically, Simple exhibits lower bias and variance compared to ST Gumbel Softmax while being exactly as efficient.

## Related Work

There is a large body of work on gradient estimation for categorical random variables. [MMT17b, JGP17b] propose the Gumbel-softmax distribution (named the concrete distribution by the former) to relax categorical random variables. For more complex distributions, such as the $k$-subset

```
Algorithm 3 The proposed algorithm for the \(k\)-subset distribution
    function ForwardPass \((\boldsymbol{\theta}) \quad\) function \(\operatorname{BackwardPass}\left(\nabla_{\boldsymbol{z}} \ell\left(f_{\boldsymbol{u}}(\boldsymbol{z}, \boldsymbol{x}), \boldsymbol{y}\right)\right)\)
    // \(p_{\boldsymbol{\theta}}\left(\sum_{m=1}^{i} z_{m}=j\right)\) for all \(i, j \quad\) load \(\boldsymbol{\theta}\) from the forward pass
    \(a=\operatorname{PrExact} \operatorname{lyk}(\boldsymbol{\theta}, n, k) \quad / /\) derivatives of \(p_{\boldsymbol{\theta}}\left(z \mid \sum_{i} z_{i}=k\right)\)
    // Sample from \(p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right) \quad \boldsymbol{\mu}=\nabla_{\boldsymbol{\theta}} \log a[n, k] / /\) by auto-diff
    \(\boldsymbol{z}=\operatorname{Sample}(\boldsymbol{\theta}, n, k) \quad\) // Return the directional derivative of the
    save \(a\) for the backward pass
    return \(z\)
    // marginals along the downstream gradients
    return \(\operatorname{JVP}\left(\boldsymbol{\mu}, \nabla_{\boldsymbol{z}} \ell\left(f_{\boldsymbol{u}}(\boldsymbol{z}, \boldsymbol{x})\right)\right.\)
```

distribution which we are concerned with in this section, existing approaches either use the straight-through and score function estimators or propose tailor-made relaxations (see for instance [KSE16, CSW18, GWZ18]). We directly compare to the score function and straight-through estimator as well as the tailored relaxations of [CSW18, GWZ18] and show that we are competitive and obtain a lower bias and/or variance than these other estimators. [TMM17, GCW18] develop parameterized control variates based on continuous relaxations for the score-function estimator. Lastly, [PCT20] offers a comprehensible work on relaxed gradient estimators, deriving several extensions of the softmax trick. All of the above works, ours included, assume the independence of the selected items, beyond there being $k$ of them. That is with the exception of [PCT20] which make use of a relaxation using pairwise embeddings, but do not make their code available. We leave that to future work.

A related line of work has developed and analyzed sparse variants of the softmax function, motivated by their potential computational and statistical advantages. Representative examples are [BMN20, PNM19, CNM19, MA16]. SparseMAP [NMB18] has been proposed in the context of structured prediction and latent variable models, also replacing the softmax with a sparser distribution. LP-SparseMAP [NM20] is an extension that uses a relaxation of the optimization problem rather than a MAP solver. Sparsity can also be exploited for efficient marginal inference in latent variable models [CNA20]. Contrary to our work, they cannot control the sparsity level
exactly through a $k$-subset constraint or guarantee a sparse output. Also, we aim at cases where samples in the forward pass are required.

Integrating specialized discrete algorithms into neural networks is growing in popularity. Examples are sorting algorithms [CTV19, BTB20, GWZ18], ranking [RMP20, KVW19], dynamic programming [MB18, CT19], and solvers for combinatorial optimization problems [BBT20, RSZ20, SBK20, NMF21b, MFN23, ZMY21a] or even probabilistic circuits over structured output spaces [ATC22a, Blo19]. There has also been work on making common programming language expression such as conditional statements, loops, and indexing differentiable through relaxations [PBK21]. [XDC20] propose optimal transport to obtain differentiable sorting methods for top- $k$ classification.

### 2.2.3 Empirical Evaluation

We conduct experiments on four different tasks: 1) A synthetic experiment designed to test the bias and variance, as well as the average deviation of Simple compared to a variety of wellestablished estimators in the literature. 2) A discrete $k$-subset Variational Auto-Encoder (DVAE) setting, where the latent space models a probability distribution over $k$-subsets. We will show that we can compute the evidence lower bound (ELBO) exactly, and that, coupled with exact sampling and our Simple gradient estimator, we attain a much lower loss compared to state of the art in sparse DVAEs. 3) The learning to explain (L2X) setting, where the aim is to select the $k$-subset of words that best describe the classifier's prediction, where we show an improved mean-squared error, as well as precision, across the board. 4) A novel, yet simple task, sparse linear regression, where, in a vein similar to L2X, we wish to select a $k$-subset of features that give rise to a linear regression model, avoiding overfitting the spurious features present in the data. Table 2.1 details the architecture with the objective functions. Our code is publicly available at github.com/UCLA-StarAI/SIMPLE.

### 2.2.3.1 Synthetic Experiments

We carried out a series of experiments with a 5 -subset distribution, and a latent space of dimension 10. We set the loss to $L(\boldsymbol{\theta})=\mathbb{E}_{z \sim p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)}\left[\|\boldsymbol{z}-\mathbf{b}\|^{2}\right]$, where $\mathbf{b}$ is the groundtruth logits sampled from $\mathcal{N}(0, \mathbf{I})$. Such a distribution is tractable: we only have $\binom{10}{5}=252 k$-subsets, which are easily enumerable and therefore, the exact gradient, the golden standard, can be computed in closed form.

In this experiment, we are interested in three metrics: bias, variance, and the average error of each gradient estimator, where the latter is measured by averaging the deviation of each singlesample gradient estimate from the exact gradient. We used the cosine distance, defined as $1-$ cosine similarity as the measure of deviation in our calculation of the metrics above, as we only care about direction.

We compare against four different baselines: exact, which denotes the exact gradient; Soft$S u b$ [XE19b], which uses an extension of the Gumbel-Softmax trick to sample relaxed $k$-subsets on the forward pass; I-MLE, which denotes the IMLE gradient estimator [NMF21b], where approximate samples are obtained using perturb-and-map (PAM) on the forward pass, approximating the marginals using PAM samples on the backward pass; and score function estimator, SFE.

We tease apart Simple's improvements by comparing three different flavors: SIMPLE-F, which only uses exact sampling, falling back to estimating the marginals using exact samples; SIMPLE-B, which uses exact marginals on the backward pass with approximate PAM samples on the forward pass; and Simple, coupling exact samples on the forward pass with exact marginals on the backward pass.

Our results are shown in Figure 2.1 As expected, we observe that SFE exhibits no bias, but high variance whereas SoftSub suffers from both bias and variance, due to the Gumbel noise injection into the samples to make them differentiable. We observe that I-MLE exhibits very high bias, as well as very low variance. This can be attributed to the PAM sampling, which in the case of $k$-subset distribution does not sample faithfully from the distribution, but is instead


Figure 2.4: ELBO against \# of epochs. (Left) Comparison of Simple against different flavors of IMLE on the 10 -subset DVAE, and (Right) against ST Gumbel Softmax on the 1 -subset DVAE.
biased to sampling only the mode of the distribution. This also means that, by approximating the marginals using PAM samples, there is a lot less variance to our gradients. On to our Simple gradient estimator, we see that it exhibits less bias as well as less variance compared to all the other gradient estimators. We also see that each estimated gradient is, on average, much more aligned with the exact gradient. To understand why that is, we compare Simple, Simple-f, and Simpleb. As hypothesized, we observe that exact sampling, Simple-f, reduces the bias, but increases the variance compared to I-MLE, this is since, unlike the PAM samples, our exact sample span the entire sample space.We also observe that, even compared to I-MLE, Simple-b, reduces the variance by marginalizing over all possible samples.

### 2.2.3.2 Discrete Variational Auto-Encoder

Next, we test our Simple gradient estimator in the $k$-subset discrete variational auto-encoder (DVAE) setting, where the latent variables model a probability distribution over $k$-subsets, and has a dimensionality of 20. Similar to prior work [JGP17b, NMF21b], the encoding and decoding functions of the VAE consist of three dense layers (encoding: 512-256-20x20; decoding: 256-512784). The DVAE is trained to minimize the sum of reconstruction loss and KL-divergence of the $k$-subset distribution and the constrained uniform distribution, known as the ELBO, on MNIST.

In prior work, the KL-divergence was approximated using the unconditional marginals, ob-

```
Algorithm 4 Entropy \((\boldsymbol{\theta}, n, k)\)
Output: \(H(\boldsymbol{z})=-\mathbb{E}_{\boldsymbol{z} \sim p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)}[\log p(\boldsymbol{z})]\)
    \(h=\operatorname{zeros}(n, k)\)
    for \(i=k\) to \(n\) do
        for \(j=0\) to \(k\) do
            // \(p\left(z_{i} \mid \sum_{m=1}^{i} z_{m}=j\right)\)
            \(p=a[i-1, j-1] * p_{\theta_{i}}\left(z_{i}=1\right) / a[i, j]\)
            \(h[i, j]=H_{b}(p)+p * h[i-1, j]+\)
                \((1-p) * h[i-1, j+1]\)
    return \(h\)
```

Input: The logits $\boldsymbol{\theta}$ of the distribution, the number of variables $n$, and the subset size $k$
tained simply through a Softmax layer. Instead we show that the KL-divergence between the $k$-subset distribution and the uniform distribution can be computed exactly. First note that, through simple algebraic manipulations, the KL-divergence between the $k$-subset distribution and the constrained uniform distribution can be rewritten as the sum of negative entropy, $-H(\boldsymbol{z})$, where $\boldsymbol{z} \sim p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)$ and $\log$ the number of $k$-subsets, $\log \binom{n}{k}$, reducing the hardness of computing the KL-divergence, to computing the entropy of a $k$-subset distribution, for which Algorithm 4 gives a tractable algorithm. Intuitively, the uncertainty in the distribution over a sequence of length $n, k$ of which are true, decomposes as the uncertainty over $Z_{n}$, and the average of the uncertainties over the remainder of the sequence.

Proposition 3. Let Entropy be defined as in Algorithm 4. Given variables, $Z_{1}, \cdots, Z_{n}$, and a $k$-subset distribution $p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)$, Algorithm 4 computes entropy of $p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)$.

We plot the loss ELBO against the number of epochs, as seen in Figure 2.4. We compared against I-MLE using sum-of-gamma noise as well as Gumbel noise for PAM sampling, on the 10subset DVAE, and against ST Gumbel Softmax on the 1-subset DVAE. We observe a significantly
lower loss on the test set on the 10 -subset DVAE, partly attributable to the exact ELBO computation, but also on the 1-subset DVAE compared to ST Gumbel Softmax, where the sole difference is the backward pass.

### 2.2.3.3 Learning to Explain

The BeerAdvocate dataset [MLJ12] consists of free-text reviews and ratings for 4 different aspects of beer: appearance, aroma, palate, and taste. The training set has 80 k reviews for the aspect Appearance and 70k reviews for all other aspects. In addition to the ratings for all reviews, each sentence in the test set contains annotations of the words that best describe the review score with respect to the various aspects. We address the problem introduced by the L2X paper [CSW18] of learning a $k$-subset distribution over words that best explain a given rating. We follow the architecture suggested in the L2X paper, consisting of four convolutional and one dense layer.

We compare to relaxation-based baselines L2X [CSW18] and SoftSub [XE19b] as well as to I-MLE which uses perturb-and-MAP to both compute an approximate sample in the forward pass and to estimate the marginals. Prior work has shown that the straight-through estimator (STE) did not work well and we omit it here. We used the standard hyperparameter settings of [CSW18] and choose the temperature parameter $t \in\{0.1,0.5,1.0,2.0\}$ for all methods. We used the standard Adam settings and trained separate models for each aspect using MSE as point-wise loss $\ell$. Table 2.3 lists results for $k \in\{5,10,15\}$ for the Aroma aspect. The mean-squared error (MSE) of Simple is almost always lower and its subset precision never significantly exceeded by those of the baselines. Table 2.2 shows results on the remaining aspects Appearance, Palate, and Taste for $k=10$.

### 2.2.3.4 Sparse Linear Regression

Given a library of feature functions, the task of sparse linear regression aims to learn from data which feature subset best describes the nonlinear partial differential equation (PDE) that the data are sampled from. We propose to tackle this task by learning a $k$-subset distribution over the feature functions. During learning, we first sample from the $k$-subset distribution to decide

| Method | Appearance |  | Palate |  | Taste |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Test MSE | Precision | Test MSE | Precision | Test MSE | Precision |
| Simple (Ours) | $2.35 \pm 0.28$ | $66.81 \pm 7.56$ | $2.68 \pm 0.06$ | $44.78 \pm 2.75$ | $2.11 \pm 0.02$ | $42.31 \pm \mathbf{0 . 6 1}$ |
| L2X ( $\mathrm{t}=0.1$ ) | $10.70 \pm 4.82$ | $30.02 \pm 15.82$ | $6.70 \pm 0.63$ | $50.39 \pm 13.58$ | $6.92 \pm 1.61$ | $32.23 \pm 4.92$ |
| SoftSub (t = 0.5) | $2.48 \pm 0.10$ | $52.86 \pm 7.08$ | $2.94 \pm 0.08$ | $39.17 \pm 3.17$ | $2.18 \pm 0.10$ | $41.98 \pm 1.42$ |
| I-MLE ( $\tau=30$ ) | $2.51 \pm 0.05$ | $65.47 \pm 4.95$ | $2.96 \pm 0.04$ | $40.73 \pm 3.15$ | $2.38 \pm 0.04$ | $41.38 \pm 1.55$ |

Table 2.2: Results for three aspects with $k=10$ : test MSE and subset precision, both $\times 100$

| Method | $k=5$ |  | $k=10$ |  | $k=15$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Test MSE | Precision | Test MSE | Precision | Test MSE | Precision |
| Simple (Ours) | $2.27 \pm 0.05$ | $57.30 \pm 3.04$ | $2.23 \pm 0.03$ | $47.17 \pm 2.11$ | $3.20 \pm 0.04$ | $53.18 \pm 1.09$ |
| L2X ( $\mathrm{t}=0.1$ ) | $5.75 \pm 0.30$ | $33.63 \pm 6.91$ | $6.68 \pm 1.08$ | $26.65 \pm 9.39$ | $7.71 \pm 0.64$ | $23.49 \pm 10.93$ |
| SoftSub (t = 0.5) | $2.57 \pm 0.12$ | $54.06 \pm 6.29$ | $2.67 \pm 0.14$ | $44.44 \pm 2.27$ | $2.52 \pm 0.07$ | $37.78 \pm 1.71$ |
| I-MLE ( $\tau=30$ ) | $2.62 \pm 0.05$ | $54.76 \pm 2.50$ | $2.71 \pm 0.10$ | $47.98 \pm 2.26$ | $2.91 \pm 0.18$ | $39.56 \pm 2.07$ |

Table 2.3: Results for aspect Aroma: test MSE and subset precision, both $\times 100$, for $k \in\{5,10,15\}$.
which feature function subset to choose. With $k$ chosen features, we perform linear regression to learn the coefficients of the features from data, and then update the $k$-subset distribution logit parameters by minimizing RMSE.

To test our proposed approach, we follow the experimental setting in PySINDy [SCQ20, KSF22] and use the dataset collected by PySINDy where the samples are collected from the Kuramoto-Sivashinsky (KS) equation, a fourth-order nonlinear PDE known for its chaotic behavior. This PDE takes the form $v_{t}=-v_{x x}-v_{x x x x}-v v_{x}$, which can be seen as a linear combination of feature functions $\mathcal{V}=\left\{v_{x x}, v_{x x x x}, v v_{x}\right\}$ with the coefficients all set to a value of -1 . At test time, we use the MAP estimation of the learned $k$-subset distribution to choose the $k$ feature
functions. For $k=3$, our proposed method achieves the same performance as the state-of-the-art solver on this task, PySINDy. It identifies the KS PDE from data by choosing exactly the ground truth feature function subset $\mathcal{V}$, obtaining an RMSE of 0.00622 after applying linear regression on $\mathcal{V}$.

### 2.3 Simple for Graph Rewiring

Besides the applications shown in the previous section, this section demonstrates how Simple excels at a complex tasks, probabilistically rewiring message-passing graph neural networks (MPNNs). MPNNs emerge as powerful tools for processing graph-structured input. However, they operate on a fixed input graph structure, ignoring potential noise and missing information. Furthermore, their local aggregation mechanism can lead to problems such as over-squashing and limited expressive power in capturing relevant graph structures. Existing solutions to these challenges have primarily relied on heuristic methods, often disregarding the underlying data distribution. Hence, devising principled approaches for learning to infer graph structures relevant to the given prediction task remains an open challenge. We devise probabilistically rewired MPNNs (PR-MPNNs), which learn to add relevant edges while omitting less beneficial ones using differentiable $k$-subset sampling and we show that when combined with Simple, PR-MPNNs show superior performance.

### 2.3.1 Background

Graph-structured data is prevalent across various application domains, including fields like chemoand bioinformatics [BO04, JEP21, RNE22], combinatorial optimization [CCK23], and social-network analysis [EK12], highlighting the need for machine learning techniques designed explicitly for graphs. In recent years, message-passing graph neural networks (MPNNs) [KW17, GSR17, SGT08a, VCC18] have become the dominant approach, showing promising performance in tasks such as predicting molecular properties [KGG20, JEP21] or enhancing combinatorial solvers [CCK23].

However, MPNNs have a limitation due to their local aggregation mechanism. They focus on encoding local structures, severely limiting their expressive power [MRF19, MLM21, XHL19]. In
addition, MPNNs struggle to capture global or long-range information, possibly leading to phenomena like under-reaching [BKM20] or over-squashing [AY21]. Over-squashing, as explained by [AY21], refers to excessive information compression from distant nodes due to a source node's extensive receptive field, occurring when too many layers are stacked. Existing strategies to mitigate over-squashing rely on heuristic rewiring methods that may not adapt well to a prediction task or employ computationally intensive global attention mechanisms. Furthermore, the impact of probabilistic rewiring on a model's expressive power remains unclear.

We formalize the necessary notations and concepts as below.
Notations Let $\mathbb{N}:=\{1,2,3, \ldots\}$. For $n \geq 1$, let $[n]:=\{1, \ldots, n\} \subset \mathbb{N}$. We use $\{\{\ldots\}$ to denote multisets, i.e., the generalization of sets allowing for multiple instances for each of its elements. A graph $G$ is a pair $(V(G), E(G))$ with finite sets of vertices or nodes $V(G)$ and edges $E(G) \subseteq\{\{u, v\} \subseteq V(G) \mid u \neq v\}$. If not otherwise stated, we set $n:=|V(G)|$, and the graph is of order $n$. We also call the graph $G$ an $n$-order graph. For ease of notation, we denote the edge $\{u, v\}$ in $E(G)$ by $(u, v)$ or $(v, u)$. Throughout this section, we use standard notations, e.g., we denote the neighborhood of a vertex $v$ by $N(v)$ and $\ell(v)$ denotes its discrete vertex label, and so on.

1-dimensional Weisfeiler-Leman algorithm The 1-WL or color refinement is a well-studied heuristic for the graph isomorphism problem, originally proposed by [WL68]. Formally, let $G=$ $(V(G), E(G), \ell)$ be a labeled graph. In each iteration, $t>0$, the 1-WL computes a node coloring $C_{t}^{1}: V(G) \rightarrow \mathbb{N}$, depending on the coloring of the neighbors. That is, in iteration $t>0$, we set

$$
C_{t}^{1}(v):=\operatorname{RELABEL}\left(\left(C_{t-1}^{1}(v),\left\{\left\{C_{t-1}^{1}(u) \mid u \in N(v)\right\}\right)\right),\right.
$$

for all nodes $v \in V(G)$, where RELABEL injectively maps the above pair to a unique natural number, which has not been used in previous iterations. In iteration 0 , the coloring $C_{0}^{1}:=\ell$. To test if two graphs $G$ and $H$ are non-isomorphic, we run the above algorithm in "parallel" on both graphs. If the two graphs have a different number of nodes colored $c \in \mathbb{N}$ at some iteration, the 1-WL distinguishes the graphs as non-isomorphic. Moreover, if the number of colors between
two iterations, $t$ and $(t+1)$, does not change, i.e., the cardinalities of the images of $C_{t}^{1}$ and $C_{i+t}^{1}$ are equal, or, equivalently,

$$
C_{t}^{1}(v)=C_{t}^{1}(w) \Longleftrightarrow C_{t+1}^{1}(v)=C_{t+1}^{1}(w)
$$

for all nodes $v$ and $w$ in $V(G)$, the algorithm terminates. For such $t$, we define the stable coloring $C_{\infty}^{1}(v)=C_{t}^{1}(v)$, for $v$ in $V(G)$. The stable coloring is reached after at most max $\{|V(G)|,|V(H)|\}$ iterations [Gro17]. It is easy to see that the algorithm cannot distinguish all non-isomorphic graphs [CFI92]. Nonetheless, it is a powerful heuristic that can successfully test isomorphism for a broad class of graphs [BK79]. A function $f: V(G) \rightarrow \mathbb{R}^{d}$, for $d>0$, is 1-WL-equivalent if $f \equiv C_{\infty}^{1}$.

Message-passing graph neural networks Intuitively, MPNNs learn a vectorial representation, i.e., a $d$-dimensional real-valued vector, representing each vertex in a graph by aggregating information from neighboring vertices. Let $\mathbf{G}=(G, \mathbf{L})$ be an attributed graph, following, [GSR17] and [SGT08b], in each layer, $t>0$, we compute vertex features

$$
\mathbf{h}_{v}^{(t)}:=\mathrm{UPD}^{(t)}\left(\mathbf{h}_{v}^{(t-1)}, \operatorname{AGG}^{(t)}\left(\left\{\mathbf{h}_{u}^{(t-1)} \mid u \in N(v)\right\}\right)\right) \in \mathbb{R}^{d},
$$

where $\mathrm{UPD}^{(t)}$ and $\mathrm{AGG}^{(t)}$ may be differentiable parameterized functions, e.g., neural networks, and $\mathbf{h}_{v}^{(t)}=\mathbf{L}_{v}$. In the case of graph-level tasks, e.g., graph classification, one uses

$$
\mathbf{h}_{G}:=\operatorname{READOUT}\left(\left\{\left\{\mathbf{h}_{v}^{(T)} \mid v \in V(G)\right\}\right) \in \mathbb{R}^{d},\right.
$$

to compute a single vectorial representation based on learned vertex features after iteration $T$. Again, READOUT may be a differentiable parameterized function, e.g., a neural network. To adapt the parameters of the above three functions, they are optimized end-to-end, usually through a variant of stochastic gradient descent, e.g., [KB15], together with the parameters of a neural network used for classification or regression.

### 2.3.2 Probabilistically Rewired MPNNs

By leveraging recent progress in differentiable $k$-subset sampling [AZN23a], we derive probabilistically rewired MPNNs (PR-MPNNs). Concretely, we utilize an upstream model to learn
prior weights for candidate edges. We then utilize the weights to parameterize a probability distribution constrained by so-called $k$-subset constraints. Subsequently, we sample multiple $k$-edge adjacency matrices from this distribution and process them using a downstream model, typically an MPNN, for the final predictions task. To make this pipeline trainable via gradient descent, we adapt recently proposed discrete gradient estimation and tractable sampling techniques [AZN23a, NMF21a, XE19a]; see Figure 2.5 for an overview of our architecture. Our theoretical analysis explores how PR-MPNNs overcome MPNNs' inherent limitations in expressive power and identifies precise conditions under which they outperform purely randomized approaches. Empirically, we demonstrate that our approach effectively mitigates issues like oversquashing and under-reaching. In addition, on established real-world datasets, our method exhibits competitive or superior predictive performance compared to traditional MPNN models and graph transformer architectures. Overall, PR-MPNNs pave the way for the principled design of more flexible MPNNs, making them less vulnerable to potential noise and missing information.

Here, we outline PR-MPNNs based on recent advancements in discrete gradient estimation and tractable sampling techniques [AZN23a]. Let $\mathfrak{A}_{n}$ denote the set of adjacency matrices of $n$ order graphs. Further, let $(G, \mathbf{X})$ be a $n$-order attributed graph with an adjacency matrix $\mathbf{A}(G) \in$ $\mathfrak{A}_{n}$ and node attribute matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$, for $d>0$. A PR-MPNN maintains a (parameterized) upstream model $h_{\boldsymbol{v}}: \mathfrak{A}_{n} \times \mathbb{R}^{n \times d} \rightarrow \Theta$, typically a neural network, parameterized by $\boldsymbol{v}$, mapping an adjacency matrix and corresponding node attributes to unnormalized edge priors $\boldsymbol{\theta} \in \Theta \subseteq \mathbb{R}^{n \times n}$.

In the following, we use the priors $\boldsymbol{\theta}$ as parameters of a (conditional) probability mass function,

$$
p_{\boldsymbol{\theta}}(\mathbf{A}(H)):=\prod_{i, j=1}^{n} p_{\theta_{i j}}\left(\mathbf{A}(H)_{i j}\right)
$$

assigning a probability to each adjacency matrix in $\mathfrak{A}_{n}$, where $p_{\theta_{i j}}\left(\mathbf{A}(H)_{i j}=1\right)=\operatorname{sigmoid}\left(\theta_{i j}\right)$ and $p_{\theta_{i j}}\left(\mathbf{A}(H)_{i j}=0\right)=1-\operatorname{sigmoid}\left(\theta_{i j}\right)$. Since the parameters $\boldsymbol{\theta}$ depend on the input graph $G$, we can view the above probability as a conditional probability mass function conditioned on the graph $G$.

Unlike previous probabilistic rewiring approaches, e.g., [FNP19], we introduce dependencies
between the graph's edges by conditioning the probability mass function $p_{\theta_{i j}}(\mathbf{A}(H))$ on a $k$ subset constraint. That is, the probability of sampling any given $k$-edge adjacency matrix $\mathbf{A}(H)$, becomes

$$
p_{(\boldsymbol{\theta}, k)}(\mathbf{A}(H)):=\left\{\begin{array}{ll}
p_{\boldsymbol{\theta}}(\mathbf{A}(H)) / Z & \text { if }\|\mathbf{A}(H)\|_{1}=k,  \tag{2.10}\\
0 & \text { otherwise },
\end{array} \quad \text { with } \quad Z:=\sum_{\mathbf{B} \in \mathfrak{A}_{n}:\|\mathbf{B}\|_{1}=k} p_{\boldsymbol{\theta}}(\mathbf{B}) .\right.
$$

The original graph $G$ is now rewired into a new adjacency matrix $\overline{\mathbf{A}}$ by combining $N$ samples $\mathbf{A}^{(i)} \sim p_{(\boldsymbol{\theta}, k)}(\mathbf{A}(G))$ for $i \in[N]$ together with the original adjacency matrix $\mathbf{A}(G)$ using a differentiable aggregation function $g: \mathfrak{A}_{n}^{(N+1)} \rightarrow \mathfrak{A}_{n}$, i.e., $\overline{\mathbf{A}}:=g\left(\mathbf{A}(G), \mathbf{A}^{(1)}, \ldots, \mathbf{A}^{(N)}\right) \in \mathfrak{A}_{n}$. Subsequently, we use the resulting adjacency matrix as input to a downstream model $f_{\mathbf{d}}$, parameterized by d, typically an MPNN, for the final predictions task.

We have so far assumed that the upstream MPNN computes one set of priors $h_{v}: \mathfrak{A}_{n} \times \mathbb{R}^{n \times d} \rightarrow$ $\mathbb{R}^{n \times n}$ which we use to generate a new adjacency matrix $\overline{\mathbf{A}}$ through sampling and then aggregating the adjacency matrices $\mathbf{A}^{(1)}, \ldots, \mathbf{A}^{(N)}$. In Section 7.3, we show empirically that having multiple sets of priors from which we sample is beneficial. Multiple sets of priors mean that we learn an upstream model $h_{v}: \mathfrak{A}_{n} \times \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{n \times n \times M}$ where $M$ is the number of priors. We can then sample and aggregate the adjacency matrices from these multiple sets of priors.

Learning to sample To learn the parameters of the up- and downstream model $\boldsymbol{\omega}=(\boldsymbol{v}, \boldsymbol{u})$ of the PR-MPNN architecture, we minimize the expected loss

$$
L(\mathbf{A}(G), \mathbf{X}, y ; \boldsymbol{\omega}):=\mathbb{E}_{\mathbf{A}^{(i)} \sim p_{(\boldsymbol{\theta}, k)}(\mathbf{A}(G))}\left[\ell\left(f_{\boldsymbol{u}}\left(g\left(\mathbf{A}(G), \mathbf{A}^{(1)}, \ldots, \mathbf{A}^{(N)}\right), \mathbf{X}\right), y\right)\right]
$$

with $y \in \mathcal{Y}$, the targets, $\ell$ a point-wise loss such as the cross-entropy or MSE, and $\boldsymbol{\theta}=h_{\boldsymbol{v}}(\mathbf{A}(G), \mathbf{X})$. To minimize the above expectation using gradient descent and backpropagation, we need to efficiently draw Monte-Carlo samples from $p_{(\boldsymbol{\theta}, k)}(\mathbf{A}(G))$ and estimate $\nabla_{\boldsymbol{\theta}} L$ the gradients of an expectation regarding the parameters $\boldsymbol{\theta}$ of the distribution $p_{(\boldsymbol{\theta}, k)}$.

Sampling To sample an adjacency matrix $\mathbf{A}^{(i)}$ from $p_{(\boldsymbol{\theta}, k)}(\mathbf{A}(G))$ conditioned on $k$-edge constraints, and to allow PR-MPNNs to be trained end-to-end, we use Simple [AZN23a], a recently
proposed gradient estimator. Concretely, we can use Simple to sample exactly from the $k$-edge adjacency matrix distribution $p_{(\boldsymbol{\theta}, k)}(\mathbf{A}(G))$ on the forward pass. On the backward pass, we compute the approximate gradients of the loss (which is an expectation over a discrete probability mass function) regarding the prior weights $\boldsymbol{\theta}$ using

$$
\nabla_{\boldsymbol{\theta}} L \approx \partial_{\boldsymbol{\theta}} \mu(\boldsymbol{\theta}) \nabla_{\mathbf{A}} \ell \text { with } \mu(\boldsymbol{\theta}):=\left\{p_{(\boldsymbol{\theta}, k)}\left(\mathbf{A}(G)_{i j}\right)\right\}_{i, j=1}^{n} \in \mathbb{R}^{n \times n}
$$

with an exact and efficient computation of the marginals $\mu(\boldsymbol{\theta})$ that is differentiable on the backward pass, achieving lower bias and variance. We show empirically that Simple [AZN23a] outperforms other sampling and gradient approximation methods such as Gumbel SoftSub-ST [XE19a] and I-MLE [NMF21a], improving accuracy without incurring a computational overhead.

Computational complexity The vectorized complexity of the exact sampling and marginal inference step is $\mathcal{O}(\log k \log l)$, where $k$ is from our $k$-subset constraint, and $l$ is the maximum number of edges that we can sample. Assuming a constant number of layers, PR-MPNN's worstcase training complexity is $\mathcal{O}(l)$ for both the upstream and downstream models. Let $n$ be the number of nodes in the initial graph, and $l=\max \left(\left\{l_{\mathrm{add}}, l_{\mathrm{rm}}\right\}\right)$, with $l_{\mathrm{add}}$ and $l_{\mathrm{rm}}$ being the maximum number of added and deleted edges. If we consider all of the possible edges for $l_{\text {add }}$, the worst-case complexity becomes $\mathcal{O}\left(n^{2}\right)$. Therefore, to reduce the complexity in practice, we select a subset of the possible edges using simple heuristics, such as considering the top $l_{\text {add }}$ edges of the most distant nodes. During inference, since we do not need gradients for edges not sampled in the forward pass, the complexity is $\mathcal{O}(l)$ for the upstream model and $\mathcal{O}(L)$ for the downstream model, with $L$ being the number of edges in the rewired graph.

### 2.3.3 Expressive Power of Probabilistically Rewired MPNNs

We now, for the first time, explore the extent to which probabilistic MPNNs overcome the inherent limitations of MPNNs in expressive power caused by the equivalence to $1-W L$ in distinguishing non-isomorphic graphs [XLT18, MRF19]. Moreover, we identify formal conditions under which PR-MPNNs outperform popular randomized approaches such as those dropping nodes and edges uniformly at random. We first make precise what we mean by probabilistically separating graphs

Forward pass


Figure 2.5: Overview of the probabilistically rewired MPNN framework. PR-MPNNs use an upstream model to learn priors $\boldsymbol{\theta}$ for candidate edges, parameterizing a probability mass function conditioned on exactly- $k$ constraints. Subsequently, we sample multiple $k$-edge adjacency matrices (here: $k=1$ ) from this distribution, aggregate these matrices (here: subtraction), and use the resulting adjacency matrix as input to a downstream model, typically an MPNN, for the final predictions task. On the backward pass, the gradients of the loss $\ell$ regarding the parameters $\boldsymbol{\theta}$ are approximated through the derivative of the exactly- $k$ marginals in the direction of the gradients of the point-wise loss $\ell$ regarding the sampled adjacency matrix. We use recent work to make the computation of these marginals exact and differentiable, reducing both bias and variance.
by introducing a probabilistic and generally applicable notion of graph separation.
Let us assume a conditional probability mass function $p: \mathfrak{A}_{n} \rightarrow[0,1]$ conditioned on a given $n$-order graph, defined over the set of adjacency matrices of $n$-order graphs. In the context of PRMPNNs, $p$ is the probability mass function defined in Section 2.3.2 but it could also be any other conditional probability mass function over graphs. Moreover, let $f: \mathfrak{A}_{n} \rightarrow \mathbb{R}^{d}$, for $d>0$, be a permutation-invariant, parameterized function mapping a sampled graph's adjacency matrix to a vector in $\mathbb{R}^{d}$. The function $f$ could be the composition of an aggregation function $g$ that removes the sampled edges from the input graph $G$ and of a downstream MPNN. Now, the conditional
probability mass function $p$ separates two graphs $G$ and $H$ with probability $\rho$ with respect to $f$ if

$$
\mathbb{E}_{\bar{G} \sim p(\cdot \mid G), \bar{H} \sim p(\cdot \mid H)}[f(\mathbf{A}(\bar{G})) \neq f(\mathbf{A}(\bar{H}))]=\rho,
$$

that is, if in expectation over the conditional probability distribution, the vectors $f(\mathbf{A}(\bar{G}))$ and $f(\mathbf{A}(\bar{H}))$ are distinct with probability $\rho$.

In what follows, we analyze the case of $p$ being the exactly- $k$ probability distribution defined in Equation 2.10 and $f$ being the aggregation function removing edges and a downstream MPNN. However, our framework readily generalizes to the case of node removal. Following Section 2.3.2, we sample adjacency matrices with exactly $k$ edges and use them to remove edges from the original graph. We aim to understand the separation properties of the probability mass function $p_{(k, \boldsymbol{\theta})}$ in this setting and for various types of graph structures. Most obviously, we do not want to separate isomorphic graphs and, therefore, remain isomorphism invariant, a desirable property of MPNNs.

Theorem 2. For sufficiently large $n$, for every $\varepsilon \in(0,1)$ and $k>0$, we have that for almost all pairs, in the sense of [BES80], of isomorphic n-order graphs $G$ and $H$ and all permutation-invariant, 1-WL-equivalent functions $f: \mathfrak{A}_{n} \rightarrow \mathbb{R}^{d}$, $d>0$, there exists a probability mass function $p_{(\boldsymbol{\theta}, k)}$ that separates the graph $G$ and $H$ with probability at most $\varepsilon$ with respect to $f$.

Theorem 2 relies on the fact that most graphs have a discrete 1-WL coloring. For graphs where the $1-W L$ stable coloring consists of a discrete and non-discrete part, the following result shows that there exist distributions $p_{(\boldsymbol{\theta}, k)}$ not separating the graphs based on the partial isomorphism corresponding to the discrete coloring.

Proposition 4. Let $\varepsilon \in(0,1), k>0$, and let $G$ and $H$ be graphs with identical 1-WL stable colorings. Let $V_{G}$ and $V_{H}$ be the subset of nodes of $G$ and $H$ that are in color classes of cardinality 1. Then, for all choices of 1-WL-equivalent functions $f$, there exists a conditional probability distribution $p_{(\boldsymbol{\theta}, k)}$ that separates the graphs $G\left[V_{G}\right]$ and $H\left[V_{H}\right]$ with probability at most $\varepsilon$ with respect to $f$.

Existing methods such as DropGNN [PMF21] or DropEdge [RHX20] are more likely to separate two (partially) isomorphic graphs by removing different nodes or edges between discrete
color classes, i.e., on their (partially) isomorphic subgraphs. For instance, we prove that pairs of graphs with $m$ edges exist where the probability of non-separation under uniform edge sampling is at most $1 / m$. This is undesirable as it breaks the MPNNs' permutation-invariance in these parts.

Now that we have established that distributions with priors from upstream MPNNs are more likely to preserve (partial) isomorphism between graphs, we turn to analyze their behavior in separating the non-discrete parts of the coloring. The following theorem shows that PR-MPNNs are more likely to separate non-isomorphic graphs than probability mass functions that remove edges or nodes uniformly at random.

Theorem 3. For every $\varepsilon \in(0,1)$ and every $k>0$, there exists a pair of non-isomorphic graphs $G$ and $H$ with identical and non-discrete 1-WL stable colorings such that for every 1-WL-equivalent function $f$,
(1) there exists a probability mass function $p_{(k, \boldsymbol{\theta})}$ that separates $G$ and $H$ with probability at least $(1-\varepsilon)$ with respect to $f$;
(2) removing edges uniformly at random separates $G$ and $H$ with probability at most $\varepsilon$ with respect to $f$.

Finally, we can also show a negative result, namely that there exist classes of graphs for which PR-MPNNs cannot do better than random sampling.

Proposition 5. For every $k>0$, there exist non-isomorphic graphs $G$ and $H$ with identical 1-WL colorings such that every probability mass function $p_{(\boldsymbol{\theta}, k)}$ separates the two graphs with the same probability as the distribution that samples edges uniformly at random.

## Related Work

MPNNs are inherently biased towards encoding local structures, which limits their expressive power [MRF19, MLM21, XHL19]. Specifically, they are at most as powerful as distinguishing


Figure 2.6: Comparison between PR-MPNN and DropGNN on the 4-Cycles dataset. PR-MPNN rewiring is almost always better than randomly dropping nodes, and is always better with 10 priors.
non-isomorphic graphs or nodes with different structural roles as the 1-dimensional WeisfeilerLeman algorithm [WL68], a simple heuristic for the graph isomorphism problem; see Section 2.3.1. Additionally, they cannot capture global or long-range information, often linked to phenomena such as under-reaching [BKM20] or over-squashing [AY21], with the latter being heavily investigated in recent works.

Graph rewiring Several recent works aim to circumvent over-squashing via graph rewiring. Perhaps the most straightforward way of graph rewiring is incorporating multi-hop neighbors. For example, [BYS22] rewires the graphs with $k$-hop neighbors and virtual nodes and also augments them with positional encodings. MixHop [APK19], SIGN [FRE20], DIGL [GWG19], and SP-MPNN [ADC22] can also be considered as graph rewiring as they can reach further-away neighbors in a single layer. Particularly, [GDB23] rewires the graph similarly to [ADC22] but with a novel delay mechanism, showcasing promising empirical results. Several rewiring methods depend on particular metrics, e.g., Ricci or Forman curvature [BMS22] and balanced Forman curvature [TDC21]. In addition, [DLV22, SVV23] utilize expander graphs to enhance message passing and connectivity, while [KBM22] resort to spectral techniques, and [BKW22] propose a greedy random edge flip approach to overcome over-squashing. Refining [TDC21], [DGB23] analyzed how the architectures' width and graph structure contribute to the over-squashing problem,
showing that over-squashing happens among nodes with high commute time, stressing the importance of rewiring techniques. Contrary to our proposed method, these strategies to mitigate over-squashing either rely on heuristic rewiring methods or use purely randomized approaches that may not adapt well to a given prediction task. Furthermore, the impact of existing rewiring methods on a model's expressive power remains unclear and we close this gap with our work.

Graph transformers Different from the above, graph transformers [DRG22, HHL23, MGM23, RGD22, COB22] and similar global attention mechanisms [LWJ21, WJW21] marked a shift from local to global message passing, aggregating over all nodes. While not understood in a principled way, empirical studies indicate that graph transformers possibly alleviate over-squashing; see, e.g., [MGM23]. However, all transformers suffer from their quadratic space and memory requirements due to computing an attention matrix.

### 2.3.4 Empirical Evaluation

We explore to what extent our probabilistic graph rewiring leads to improved predictive performance on synthetic and real-world datasets. Concretely, we answer the following questions.

Q1 Can probabilistic graph rewiring mitigate the problems of over-squashing and under-reaching in synthetic datasets?

Q2 Is the expressive power of standard MPNNs enhanced through probabilistic graph rewiring? That is, can we verify empirically that the separating probability mass function of Section 2.3.3 can be learned with PR-MPNNs and that multiple priors help?

Q3 Does the increase in predictive performance due to probabilistic rewiring apply to (a) graphlevel molecular prediction tasks and (b) node-level prediction tasks involving heterophilic data?

Datasets To answer Q1, we utilized the Trees-NeighborsMatch dataset [AY21]. Additionally, we created the Trees-LeafCount dataset to investigate whether our method could mitigate under-reaching issues. To tackle Q2, we performed experiments with the Exp [ACG20] and CSL
datasets [MSR19] to assess how much probabilistic graph rewiring can enhance the models' expressivity. In addition, we utilized the 4-Cycles dataset from [Lou20, PMF21] and set it against a standard DropGNN model [PMF21] for comparison while also ablating the performance difference concerning the number of priors and samples per prior. To answer Q3 (a), we used the established molecular graph-level regression datasets Alchemy [CCH19], Zinc [JCB17, DJL20], ogbg-molhiv [HFZ20], QM9 [HYL17], LRGB [DRG22] and five datasets from the TUDataset repository [MKB20]. To answer Q3 (b), we used the Cornell, Wisconsin, Texas node-level classification datasets [PWC20].

Baseline and model configurations For our upstream model $h_{\boldsymbol{v}}$, we use an MPNN, specifically the GIN layer [XHL19]. For an edge $(v, w) \in E(G)$, we compute $\boldsymbol{\theta}_{v w}=\phi\left(\left[\mathbf{h}_{v}^{T}| | \mathbf{h}_{w}^{T}\right]\right) \in \mathbb{R}$, where $[\cdot \| \cdot]$ is the concatenation operator and $\phi$ is an MLP. After obtaining the prior $\boldsymbol{\theta}$, we rewire our graphs by sampling two adjacency matrices for deleting edges and adding new edges, i.e., $g\left(\mathbf{A}(G), \mathbf{A}^{(1)}, \mathbf{A}^{(2)}\right):=\left(\mathbf{A}(G)-\mathbf{A}^{(1)}\right)+\mathbf{A}^{(2)}$ where $\mathbf{A}^{(1)}$ and $\mathbf{A}^{(2)}$ are two sampled adjacency matrices with a possibly different number of edges, respectively. Finally, the rewired adjacency matrix (or multiple adjacency matrices) is used in a downstream model $f_{u}: \mathfrak{A}_{n} \times \mathbb{R}^{n \times d} \rightarrow \mathcal{Y}$, typically an MPNN, with parameters $\boldsymbol{u}$ and $\mathcal{Y}$ the prediction target set. For the instance where we have multiple priors, as described in Section 2.3.2, we can either aggregate the sampled adjacency matrices $\mathbf{A}^{(1)}, \ldots, \mathbf{A}^{(N)}$ into a single adjacency matrix $\mathbb{X}$ that we send to a downstream model as described in Figure 2.5, or construct a downstream ensemble with multiple aggregated matrices $\mathbf{A}_{1}, \ldots, \mathbf{X}_{M}$. In practice, we always use a downstream ensemble before the final projection layer when we rewire with more than one adjacency matrix, and we do rewiring by both adding and deleting edges.

All of our downstream models $f_{\mathbf{d}}$ and base models are MPNNs with GIN layers. When we have access to edge features, we use the GINE variant [HLG20] for edge feature processing. For graph-level tasks, we use mean pooling, while for node-level tasks, we take the node embedding $\mathbf{h}_{v}^{T}$ for a node $v$. The final embeddings are then processed and projected to the target space by an MLP.


Figure 2.7: Example graph from the Trees-LeafCount test dataset with radius 4 (left). PRMPNN rewires the graph, allowing the downstream MPNN to obtain the label information from the leaves in one massage-passing step (right).

For Zinc, Alchemy, and ogbg-molhiv, we compare our rewiring approaches with the base downstream model, both with and without positional embeddings. Further, we compare to GPS [RGD22] and SAT [COB22], two state-of-the-art graph transformers. For the TUDATASEt, we compare with the reported scores from [GRC23] and use the same evaluation strategy as in [XHL19, GRC23], i.e., running 10 -fold cross-validation and reporting the maximum average validation accuracy. For different tasks, we search for the best hyperparameters for sampling and our upstream and downstream models. For Zinc, Alchemy, and ogbg-molhiv, we evaluate multiple gradient estimators in terms of predictive power and computation time. Specifically, we compare Gumbel SoftSub-ST [MMT17a, JGP17a, XE19a], I-MLE [NMF21a], and Simple [AZN23a]. The results in terms of predictive power are detailed in Table 2.4.

Experimental results and discussion Concerning Q1, our rewiring method achieves perfect test accuracy up to a problem radius of 6 on both Trees-NeighborsMatch and TreesLeafCount, demonstrating that it can successfully alleviate over-squashing and under-reaching, see Figure 2.8. For Trees-LeafCount, our model can create connections directly from the leaves to the root, achieving perfect accuracy with a downstream model containing a single MPNN layer. We provide a qualitative result in Figure 2.7. Concerning Q2, on the 4-Cycles dataset, our probabilistic rewiring method matches or outperforms DropGNN. This advantage is most pronounced with 5 and 10 priors, where we achieve $100 \%$ task accuracy using 20 samples, as


Figure 2.8: Test accuracy of our rewiring method on the Trees-NeighborsMatch [AY21] dataset, compared to the reported accuracies from [MGM23].
detailed in Figure 2.6. On the Exp dataset, we showcase the expressive power of probabilistic rewiring by achieving perfect accuracy, see Table 2.5. Besides, our rewiring approach can distinguish the regular graphs from the CSL dataset without any positional encodings, whereas the 1-WL-equivalent GIN obtains only random accuracy. Concerning Q3 (a), the results in Table 2.4 show that our rewiring methods consistently outperform the base models on Zinc, Alchemy, and ogbg-molhiv and are competitive or better than the state-of-the-art GPS and SAT graph transformer methods. On TUDATASET, see Table 2.7, our probabilistic rewiring method outperforms existing approaches and obtains lower variance on most of the datasets, with the exception being NCI1, where our method ranks second, after the WL kernel. Hence, our results indicate that probabilistic graph rewiring can improve performance for molecular prediction tasks. Concerning Q3 (b), we obtain performance gains over the base model and other existing MPNNs, indicating that data-driven rewiring has the potential of alleviating the effects of over-smoothing by removing undesirable edges and making new ones between nodes with similar features. The graph transformer methods outperform the rewiring approach and the base models, except on the Texas dataset, where our method gets the best result. We speculate that GIN's aggregation mechanism for the downstream models is a limiting factor on heterophilic data. We leave the analysis of
combining probabilistic graph rewiring with downstream models that address over-smoothing for future investigations.

### 2.4 Discussion

The key to the success of Simple is the capability of exact reasoning over the $k$-subset constraint as the constraint probabilities of the exact samples from the constrained distribution serve as an informative proxy for gradient updates. This paradigm can be potentially generalized to any constraints; that is, the problem of differentiable learning under constraints can be reduced to the problem of tractable computation of the constraint probability. An interesting direction for future work is to come up with efficient approximations of Simple for complex constraints.

Table 2.4: Comparison between PR-MPNN and baselines on three molecular property prediction datasets. We report results for PR-MPNN with different gradient estimators for $k$-subset sampling: Gumbel SoftSub-ST [MMT17a, JGP17a, XE19a], I-MLE [NMF21a], and Simple [AZN23a] and compare them with the base downstream model, and two graph transformer architectures. The variant using SIMPLE consistently outperforms the base models and is competitive or better than the two graph transformers. We use green for the best model, blue for the second-best, and red for third. We note with + Edge the instances where edge features are provided and with Edge when they are not.

|  |  | Zinc |  | $\begin{aligned} & \text { OGBG-molhiv } \\ & + \text { Edge } \uparrow \end{aligned}$ | Alchemy <br> + Edge $\downarrow$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | - Edge $\downarrow$ | + Edge $\downarrow$ |  |  |
|  | K-ST SAT | $0.166 \pm 0.007$ | $0.115 \pm 0.005$ | $0.625 \pm 0.039$ | N/A |
|  | K-SG SAT | $0.162 \pm 0.013$ | $0.095 \pm 0.002$ | $0.613 \pm 0.010$ | N/A |
|  | Base | $0.258 \pm 0.006$ | $0.207 \pm 0.006$ | $0.775 \pm 0.011$ | $11.12 \pm 0.690$ |
|  | Base w. PE | $0.162 \pm 0.001$ | $0.101 \pm 0.004$ | $0.764 \pm 0.018$ | $7.197 \pm 0.094$ |
|  | PR-MPNN ${ }_{\text {Gmb }}$ (ours) | $0.153 \pm 0.003$ | $0.103 \pm 0.008$ | $0.760 \pm 0.025$ | $6.858 \pm 0.090$ |
|  | PR-MPNN ${ }_{\text {IMle }}$ (ours) | $\mathbf{0 . 1 5 1} \pm 0.001$ | $0.104 \pm 0.008$ | $\mathbf{0 . 7 7 4} \pm 0.015$ | $\mathbf{6 . 6 9 2} \pm 0.061$ |
|  | PR-MPNN ${ }_{\text {SIM }}$ (ours) | $\mathbf{0 . 1 3 9} \pm 0.001$ | $\mathbf{0 . 0 8 5} \pm 0.002$ | $\mathbf{0 . 7 9 5} \pm 0.009$ | $6.447 \pm 0.057$ |
| $\underset{i}{k}$ | GPS | N/A | $\mathbf{0 . 0 7 0} \pm 0.004$ | $\mathbf{0 . 7 8 8} \pm 0.010$ | N/A |
|  | K-ST SAT | $0.164 \pm 0.007$ | $0.102 \pm 0.005$ | $0.625 \pm 0.039$ | N/A |
|  | K-SG SAT | $\mathbf{0 . 1 3 1} \pm 0.002$ | $\mathbf{0 . 0 9 4} \pm 0.008$ | $0.613 \pm 0.010$ | N/A |

Table 2.5: Comparison between the base GIN model, PR-MPNN, and other more expressive models on the Exp dataset.

| Model | Accuracy $\uparrow$ |
| :--- | :---: |
| GIN | $0.511 \pm 0.021$ |
| GIN + ID-GNN | $1.000 \pm 0.000$ |
| OSAN | $1.000 \pm 0.000$ |
| PR-MPNN (ours) | $1.000 \pm 0.000$ |

Table 2.6: Comparison between the base GIN model and probabilistic rewiring model on CSL dataset, w/o positional encodings.

| Model | Accuracy $\uparrow$ |
| :--- | :--- |
| GIN | $0.100 \pm 0.000$ |
| GIN + PosEnc | $1.000 \pm 0.000$ |
| PR-MPNN (ours) | $0.998 \pm 0.008$ |
| PR-MPNN + PosEnc (ours) | $1.000 \pm 0.000$ |

Table 2.7: Comparison between PR-MPNN and other approaches as reported in [GRC23, KBM22, PMF21]. Our model outperforms existing approaches while keeping a lower variance in most of the cases, except for NCI1, where the WL Kernel is the best. We use green for the best model, blue for the second-best, and red for third.

| ModeL | MutaG | PTC_MR | Proteins | NCI1 | NCI109 |
| :--- | :--- | :--- | :--- | :--- | :--- |
| GK $(k=3)$ [SVP09] | $81.4 \pm 1.7$ | $55.7 \pm 0.5$ | $71.4 \pm 0.3$ | $62.5 \pm 0.3$ | $62.4 \pm 0.3$ |
| PK [NGB16] | $76.0 \pm 2.7$ | $59.5 \pm 2.4$ | $73.7 \pm 0.7$ | $82.5 \pm 0.5$ | N/A |
| WL KERNEL [SSV11] | $90.4 \pm 5.7$ | $59.9 \pm 4.3$ | $75.0 \pm 3.1$ | $86.0 \pm 1.8$ | N/A |
| DGCNN [ZCN18] | $85.8 \pm 1.8$ | $58.6 \pm 2.5$ | $75.5 \pm 0.9$ | $74.4 \pm 0.5$ | N/A |
| IGN [MBS19B] | $83.9 \pm 13.0$ | $58.5 \pm 6.9$ | $76.6 \pm 5.5$ | $74.3 \pm 2.7$ | $72.8 \pm 1.5$ |
| GIN [XHL19] | $89.4 \pm 5.6$ | $64.6 \pm 7.0$ | $76.2 \pm 2.8$ | $82.7 \pm 1.7$ | N/A |
| PPGNs [MBS19A] | $90.6 \pm 8.7$ | $66.2 \pm 6.6$ | $77.2 \pm 4.7$ | $83.2 \pm 1.1$ | $82.2 \pm 1.4$ |
| NATURAL GN [HCW20] | $89.4 \pm 1.6$ | $66.8 \pm 1.7$ | $71.7 \pm 1.0$ | $82.4 \pm 1.3$ | $83.0 \pm 1.9$ |
| GSN [BFZ22] | $92.2 \pm 7.5$ | $68.2 \pm 7.2$ | $76.6 \pm 5.0$ | $83.5 \pm 2.0$ | $83.5 \pm 2.3$ |
| CIN [BFW21] | $92.7 \pm 6.1$ | $68.2 \pm 5.6$ | $77.0 \pm 4.3$ | $83.6 \pm 1.4$ | $84.0 \pm 1.6$ |
| CAN [GBT23] | $94.1 \pm 4.8$ | $72.8 \pm 8.3$ | $78.2 \pm 2.0$ | $84.5 \pm 1.6$ | $83.6 \pm 1.2$ |
| CIN++ [GRC23] | $94.4 \pm 3.7$ | $73.2 \pm 6.4$ | $80.5 \pm 3.9$ | $85.3 \pm 1.2$ | $84.5 \pm 2.4$ |
| FoSR [KBM22] | $86.2 \pm 1.5$ | $58.5 \pm 1.7$ | $75.1 \pm 0.8$ | $72.9 \pm 0.6$ | $71.1 \pm 0.6$ |
| DROPGNN [PMF21] | $90.4 \pm 7.0$ | $66.3 \pm 8.6$ | $76.3 \pm 6.1$ | $81.6 \pm 1.8$ | $80.8 \pm 2.6$ |
| PR-MPNN (10-FOLD CV) | $98.4 \pm 2.4$ | $74.3 \pm 3.9$ | $80.7 \pm 3.9$ | $85.6 \pm 0.8$ | $84.6 \pm 1.2$ |

## CHAPTER 3

## Linear-Equality Constrained Deep Generative Models

This chapter focuses on the same challenge of differentiable learning under constraints as the previous chapter but extends from the discrete domains to the continuous ones, from $k$-subset constraint to linear-equality constraints. Enforcing linear-equality constraints poses unique challenges to the design of gradient estimators. Our aim is to tackle these complexities, demonstrating that our methods, while broadly applicable to any neural networks, enable robust and efficient enforcement of constraints, particularly in deep generative models.

### 3.1 Background

The linear equality constraint is ubiquitous not only in machine learning such as learning sparse features [CSW18] but also in scientific applications such as modeling charge-neutral molecules in chemistry [RSS20] and count-aware cell type deconvolution in biology [LWZ23]. Even though deep generative models (DGMs) have made great progress in synthesizing realistic data in various domains, they struggle to learn such constraints from data while these constraints can encode necessary domain knowledge of the task at hand. For example, a DGM trained on a charge-neutral molecule dataset often generates molecules where the charges of each atom do not sum to zero. That is, DGMs are excellent at approximating complex distributions but can fail to capture simple constraints [SDC24]. However, such constraints should be integrated into the DGMs to generate realistic data guaranteed to follow the domain knowledge.

While its discrete counterpart, known as $k$-subset constraints, has been thoroughly studied for being enforced into deep learning models [XE19c, NMF21b, AZN23a], how to integrate linear
equality into DGMs with continuous or discrete domains is underexplored. [AAB19, FNC20] study the problem when the linear equalities are constraints to the optimization problems instead of data distributions. While one can ignore the constraints in the training and use post-processing of the generated data, it harms the generation performance. [SDC24] considers linear inequalities enforced on DGMs by minimally moving the sample into the feasible space defined by the linear inequalities but it does not apply to linear equalities.

### 3.2 Gradient Estimator for Linear Equalities

This section shows how to integrate linear constraints that encode background knowledge into different DGMs. We start by characterizing the settings when the training objective admits a closed-form expression even though the data distribution is constrained, and thus the DGM allows end-to-end training as in standard settings. We show in the experiment section that such scenarios are common in real-world applications. For the more general settings when no closedfrom objectives are available, we propose several gradient estimators by the principle that constrained marginal probabilities can serve as an informative proxy for gradient updates as observed in [NMF21b, AZN23a] and provide a comparison of their bias and variance. The proposed estimators enable linear equality to be enforced not just on the data distribution but also on the distribution of latent variables. With our method, DGMs are transformed into the corresponding constrained DGMs that generate data guaranteed to satisfy the linear constraints.

To evaluate our approach, we modify the MNIST dataset to be constrained by a constant sum of pixel values and conduct an extensive experimental analysis to compare DGMs and their constrained counterparts. Results show that DGMs often fail to satisfy the constraints, that is, they are incapable of learning the domain knowledge. Their constrained counterparts, on the contrary, always generate compliant samples and have better generation quality due to the inductive bias from the constraints. By further comparing the training time, we demonstrate that the injection of the constraints brings little overhead. We also include a real-world experiment on a scientific application where the task is to predict the atomic partial charges on metal-organic frameworks
under a charge-neutral constraint where our method achieves state-of-the-art predictive performance.

Problem Statement We consider a general constrained model described by the equations

$$
\begin{equation*}
\boldsymbol{\theta}=h_{\boldsymbol{v}}(\boldsymbol{x}), \quad \boldsymbol{z} \sim p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right), \quad \hat{\boldsymbol{y}}=f_{\boldsymbol{u}}(\boldsymbol{z}) \tag{3.1}
\end{equation*}
$$

where $\boldsymbol{x} \in \mathcal{X}$ and $\hat{\boldsymbol{y}} \in \mathcal{Y}$ denote feature inputs and target outputs, respectively, $h_{\boldsymbol{v}}: \mathcal{X} \rightarrow \Theta$ and $f_{u}: \mathcal{Z} \rightarrow \mathcal{Y}$ are smooth, parameterized mappings. $\boldsymbol{\theta}$ are parameters inducing a distribution over the latent variables $\boldsymbol{z}$ conform to a Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ where parameters $\boldsymbol{\theta}=(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ consist of the mean vector $\boldsymbol{\mu} \in \mathbb{R}^{n}$ and the covariance matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{n \times n}$. That is, $\boldsymbol{z}$ has its probability density function (p.d.f.) defined as $p_{\boldsymbol{\theta}}(\boldsymbol{z})=\frac{1}{(2 \pi)^{n / 2}\left|\boldsymbol{\Sigma}_{\boldsymbol{\theta}}\right|^{1 / 2}} \exp \left(-\frac{1}{2}\left(\boldsymbol{z}-\boldsymbol{\mu}_{\boldsymbol{\theta}}\right)^{\top} \boldsymbol{\Sigma}_{\boldsymbol{\theta}}^{-1}\left(\boldsymbol{z}-\boldsymbol{\mu}_{\boldsymbol{\theta}}\right)\right)$. A linear equality $\sum_{i} z_{i}=k$ with $k \in \mathbb{R}$ is enforced over the distribution $p_{\boldsymbol{\theta}}(\boldsymbol{z})$ inducing a conditional distribution $p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)$. This formulation is general and it subsumes various neural network models that integrate the linear equality constraint in 1 ) output (when the mapping $f_{u}$ is the identity function), where the goal of constrained generative modeling is to learn the parameters $\boldsymbol{\theta}$ such that the constrained model distribution $p_{\boldsymbol{\theta}}$ approximates the underlying data distribution; or 2) latent space, where the constrained generative modeling learns a constrained posterior distribution $p_{\boldsymbol{\theta}}$ over latent variables. The training of such models is performed by optimizing an expected loss as below,

$$
\begin{equation*}
L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})=\mathbb{E}_{\boldsymbol{z} \sim p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)}\left[\ell\left(f_{\boldsymbol{u}}(\boldsymbol{z}), \boldsymbol{y}\right)\right] \quad \text { with } \boldsymbol{\omega}=(\boldsymbol{v}, \boldsymbol{u}) \text { and } \boldsymbol{\theta}=h_{\boldsymbol{v}}(\boldsymbol{x}), \tag{3.2}
\end{equation*}
$$

where $\ell: \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}^{+}$is a point-wise loss function. Figure 2.2 shows a visualization of the model. Note that in this work we formulate the linear equality constraint as an unweighted sum of variables for ease of notation while all of our theoretical results can be extended to a weighted sum of variables of the form $\sum_{i} \alpha_{i} z_{i}=k$ with $\alpha_{i} \in \mathbb{R}$ for any $i$. We explore other the distributional assumptions and present results on Poisson distributions for discrete variables with infinite domain in Section 3.2.2.3.

### 3.2.1 Closed-Form Expected Loss

Before we tackle the challenging task of deriving gradient estimators for training DGMs, it is crucial to first address a fundamental question: when is a gradient estimator necessary, and when is it not? A short answer is that when the expected loss in Equation 3.2 admits a closed-form expression, standard training can be applied to the constrained DGMs and thus no gradient estimator is needed. This section presents our theoretical results on characterizing when such closed forms are available.

The first assumption we make is that the mapping $f_{u}$ is an identity function, that is, $\hat{\boldsymbol{y}}=\boldsymbol{z}$, as it can introduce high non-linearity that easily makes a closed-form impossible. Then we show that when the element-wise loss $\ell$ is $L 1$ or $L 2$ loss, the expected loss admits a closed-form expression as below.

Proposition 6 (Gaussian Closed-form Expected Loss). Let $\boldsymbol{z} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Let $\boldsymbol{y}=\left(y_{1}, \ldots, y_{n}\right)^{T}$ be the ground truth vector subject to the equality constraint $\sum_{j=1}^{n} y_{j}=k$. Then it holds that
i) when $\ell$ is L1 loss, $L(\boldsymbol{\theta})=\sum_{i=1}^{n} \bar{\sigma}_{i} \sqrt{\frac{2}{\pi}} \exp \left(\frac{-\left(\bar{\mu}_{i}-y_{i}\right)^{2}}{2 \bar{\sigma}_{i}^{2}}\right)+\left(\bar{\mu}_{i}-y_{i}\right) \operatorname{erf}\left(\frac{\bar{\mu}_{i}-y_{i}}{\sqrt{2 \bar{\sigma}_{i}^{2}}}\right)$;
ii) when $\ell$ is $L 2$ loss, $L(\boldsymbol{\theta})=\sum_{i=1}^{n} \bar{\mu}_{i}^{2}+\bar{\sigma}_{i}^{2}-2 y_{i} \bar{\mu}_{i}+y_{i}^{2}$,
with $\bar{\mu}_{i}:=\boldsymbol{\mu}_{i}+\frac{\boldsymbol{\Sigma}_{i, i}+\sum_{j \neq i}^{n} \boldsymbol{\Sigma}_{i, j}}{\sum_{t=1}^{n} \boldsymbol{\Sigma}_{t, t}+\sum_{j \neq t}^{n} \boldsymbol{\Sigma}_{t, j}}\left(k-\sum_{j=1}^{n} \boldsymbol{\mu}_{j}\right)$ and $\bar{\sigma}_{i}^{2}:=\boldsymbol{\Sigma}_{i, i}-\frac{\left(\boldsymbol{\Sigma}_{i, i}+\sum_{j \neq i}^{n} \boldsymbol{\Sigma}_{i, j}\right)^{2}}{\sum_{t=1}^{n} \boldsymbol{\Sigma}_{t, t}+\sum_{j \neq t}^{n} \boldsymbol{\Sigma}_{t, j}}$.
We further generalize our results to a system of linear equality constraints.

Corollary 1. The closed-form expressions for $L(\boldsymbol{\theta})$ defined under L1 loss and L2 loss under the constraint $\boldsymbol{A} \boldsymbol{z}=\boldsymbol{k}$, where $\boldsymbol{A} \in \mathbb{R}^{a \times n}$ and $\operatorname{rank}(\boldsymbol{A})=a<n$, are identical to those presented in Proposition 6, except that the parameters are defined as follows,

$$
\bar{\mu}_{i}=\mu_{i}+\boldsymbol{e}_{i}^{T} \boldsymbol{\Sigma} \boldsymbol{A}\left(\boldsymbol{A} \boldsymbol{\Sigma} \boldsymbol{A}^{T}\right)^{-1}(\boldsymbol{k}-\boldsymbol{A} \boldsymbol{\mu}) \quad \text { and } \quad \bar{\sigma}_{i}^{2}=\boldsymbol{e}_{i}^{T} \boldsymbol{\Sigma} \boldsymbol{e}_{i}-\boldsymbol{e}_{i}^{T} \boldsymbol{\Sigma} \boldsymbol{A}^{T}\left(\boldsymbol{A} \boldsymbol{\Sigma} \boldsymbol{A}^{T}\right)^{-1} \boldsymbol{A} \boldsymbol{\Sigma} \boldsymbol{e}_{i}
$$

where $\boldsymbol{e}_{i}$ is a standard basis vector in $\mathbb{R}^{n}$ with only $i$-th entry being 1 and otherwise 0.

Later we will empirically show that the closed-form expected loss derived above leads to state-of-the-art predictive performance in a scientific application as in Section 3.2.3.4.

### 3.2.2 Method

As we move on to the general setting where the expected loss in Equation 3.2 does not admit a closed form, standard auto-differentiation can not be directly applied to the expected loss due to two main obstacles. First, for the gradient of $L$ w.r.t. parameters $\boldsymbol{u}$ in the decoder mapping $f_{\boldsymbol{u}}$ defined as

$$
\begin{equation*}
\nabla_{\boldsymbol{u}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})=\mathbb{E}_{\boldsymbol{z} \sim p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)}\left[\partial_{\boldsymbol{u}} f_{\boldsymbol{u}}(\boldsymbol{z}, \boldsymbol{x})^{\top} \nabla_{\hat{\boldsymbol{y}}} \ell(\hat{\boldsymbol{y}}, \boldsymbol{y})\right] \tag{3.3}
\end{equation*}
$$

with $\hat{y}=f_{\boldsymbol{u}}(\boldsymbol{z})$ being the decoding of a latent sample $\boldsymbol{z}$, the expectation does not allow closedform solution in general and requires Monte-Carlo estimations by sampling $\boldsymbol{z}$ from the constrained distribution $p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)$. Another issue arises in the gradient of $L$ w.r.t. parameters $\boldsymbol{v}$ in the encoder mapping which is defined as

$$
\begin{equation*}
\nabla_{\boldsymbol{v}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})=\partial_{\boldsymbol{v}} h_{\boldsymbol{v}}(\boldsymbol{x})^{\top} \nabla_{\boldsymbol{\theta}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega}) \tag{3.4}
\end{equation*}
$$

where the obstacle lies in the computation of the gradient of $L$ w.r.t. $\boldsymbol{\theta}$ defined as $\nabla_{\boldsymbol{\theta}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega}):=$ $\nabla_{\boldsymbol{\theta}} \mathbb{E}_{\boldsymbol{z} \sim p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)}\left[\ell\left(f_{\boldsymbol{u}}(\boldsymbol{z}, \boldsymbol{x}), \hat{\boldsymbol{y}}\right)\right]$ that requires gradient estimators. In this section, we tackle the gradient estimation for the linear equality constraint by solving the aforementioned two subproblems: ( $\mathbf{P} 1$ ) how to sample exactly from the constrained distribution $p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)$ and ( $\mathbf{P} 2$ ) how to estimate $\nabla_{\boldsymbol{\theta}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega})$. By combining solutions to these two subproblems, we manage to train the constrained DGMs in an end-to-end manner. We summarize our proposed estimators in Table 3.1.

### 3.2.2.1 Exact Sampling

We find that the constrained distribution $p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k\right)$ is still a Gaussian distribution and thus it is straightforward to sample exactly from this distribution as long as we obtain the mean vector and the covariance matrix of this Gaussian. We formally state our findings as below.

Proposition 7 (Gaussian Constrained Distribution). Given $\boldsymbol{z}=\left(z_{1}, \ldots, z_{n}\right)^{T} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, the constrained distribution $p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{j=1}^{n} z_{j}=k\right)$ is equivalent to an $n-1$ dimensional multivariate

Table 3.1: Summary of gradient estimators. The first block presents baseline estimators and the second presents our proposed ones. In the forward pass, we sample exactly from the constrained distribution by Proposition 7. In the backward pass, we use $\boldsymbol{m}(\boldsymbol{\theta})$ as a differntiable proxy.

| Gradient Estimator | Proxy $\boldsymbol{m}(\boldsymbol{\theta})$ | Description |
| :--- | :--- | :--- |
| Random | - | Sample a random gradient from $\mathcal{N}(\mathbf{0}, \boldsymbol{I})$ |
| Unconstrained Marginal | $p_{\boldsymbol{\theta}}\left(z_{i}\right)$ | P.d.f. of unconstrained $\boldsymbol{z}$ as a proxy for $\boldsymbol{z}$ |
| Constrained Marginal | $p_{\boldsymbol{\theta}}\left(z_{i} \mid \sum_{j=1}^{n} z_{j}=k\right)$ | P.d.f. of conditional marginals as a proxy for $\boldsymbol{z}$ |
| Marginal Expectation | $\mathbb{E}_{z_{i} \sim p_{\boldsymbol{\theta}}\left(z_{i} \mid \sum_{j=1}^{n} z_{j}=k\right)}\left[z_{i}\right]$ | Expectation of conditional marginals as a proxy for $\boldsymbol{z}$ |

Gaussian distribution with mean $\overline{\boldsymbol{\mu}} \in \mathbb{R}^{n-1}$ and covariance matrix $\bar{\Sigma} \in \mathbb{R}^{n-1 \times n-1}$ defined as below,

$$
\begin{gathered}
\overline{\boldsymbol{\mu}}_{i}=\boldsymbol{\mu}_{i}+\frac{\boldsymbol{\Sigma}_{i, i}+\sum_{j \neq i}^{n} \boldsymbol{\Sigma}_{i, j}}{\sum_{t=1}^{n} \boldsymbol{\Sigma}_{t, t}+\sum_{j \neq t}^{n} \boldsymbol{\Sigma}_{t, j}}\left(k-\sum_{j=1}^{n} \boldsymbol{\mu}_{j}\right) \\
\overline{\boldsymbol{\Sigma}}_{i, j}= \begin{cases}\boldsymbol{\Sigma}_{i, i}-\frac{\left(\boldsymbol{\Sigma}_{i, i}+\sum_{j \neq i}^{n} \boldsymbol{\Sigma}_{i, j}\right)^{2}}{\sum_{t=1}^{n} \boldsymbol{\Sigma}_{t, t}+\sum_{j \neq t}^{n} \boldsymbol{\Sigma}_{t, j}} & i=j \\
\boldsymbol{\Sigma}_{i, j}-\frac{\left(\boldsymbol{\Sigma}_{i, i}+\sum_{j \neq i}^{n} \boldsymbol{\Sigma}_{i, j}\right)\left(\boldsymbol{\Sigma}_{j, j}+\sum_{k \neq j}^{n} \boldsymbol{\Sigma}_{j, k}\right)}{\sum_{t=1}^{n} \boldsymbol{\Sigma}_{t, t}+\sum_{j \neq t}^{n} \boldsymbol{\Sigma}_{t, j}} & i \neq j\end{cases}
\end{gathered}
$$

### 3.2.2.2 Gradient Estimators

[NMF21b, AZN23a] observe that in the discrete case when $\boldsymbol{z}$ are Bernoulli variables, the problematic term in Equation 3.4 can be effectively approximated as

$$
\begin{equation*}
\nabla_{\boldsymbol{\theta}} L(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega}) \approx \partial_{\boldsymbol{\theta}} \boldsymbol{m}(\boldsymbol{\theta}) \nabla_{\boldsymbol{z}} \ell(\boldsymbol{x}, \boldsymbol{y} ; \boldsymbol{\omega}) \tag{3.5}
\end{equation*}
$$

where $\boldsymbol{m}(\boldsymbol{\theta}):=\left\{p_{\boldsymbol{\theta}}\left(z_{i} \mid \sum_{j=1}^{n} z_{i}=k\right)\right\}_{j=1}^{n}$ denotes the conditional marginal probabilities. The intuition behind this is that the conditional marginals characterize the probability of each $z_{i}$ in the generated sample and can serve as a differentiable proxy allowing for end-to-end training. We generalize this intuition to the linear-equality constrained distribution in the continuous domains and perform an analysis of its effectiveness. The continuous setting mostly differs from
the discrete one in that there are two choices of $\boldsymbol{m}(\boldsymbol{\theta}): 1)$ the conditional marginal probability density $p_{\boldsymbol{\theta}}\left(z_{i} \mid \sum_{j=1}^{n} z_{j}=k\right)$; and 2) the expectation of $z_{i}$ under the conditional marginal, that is, $\mathbb{E}_{z_{i} \sim p_{\boldsymbol{\theta}}\left(z_{i} \mid \sum_{j=1}^{n} z_{j}=k\right)}\left[z_{i}\right]$.

While these two quantities are the same in discrete cases, we make an interesting observation in our empirical study that in continuous domains, the use of expectation is consistently more effective than the conditional marginals. We further provide a baseline estimator that chooses $\boldsymbol{m}(\boldsymbol{\theta})$ to be the unconstrained marginals $p_{\boldsymbol{\theta}}\left(z_{i}\right)$, meaning that the constraint is ignored during the training process; empirical results show that such ignorance can harm model performance even though the constraint is enforced at inference time. We refer the readers to the experimental section for the empirical results.

The remaining question is how to obtain the closed-form $\boldsymbol{m}(\boldsymbol{\theta})$ for different estimators. We present below the theoretical results for the constrained marginals and their expectations.

Proposition 8 (Gaussian Conditional Marginal and Expectations). Given $\boldsymbol{z}=\left(z_{1}, \ldots, z_{n}\right)^{T} \sim$ $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, the conditional marginal $p_{\boldsymbol{\theta}}\left(z_{i} \mid \sum_{j=1}^{n} z_{j}=k\right)$ follows a univariate Gaussian distribution with mean $\bar{\mu}_{i}=\boldsymbol{\mu}_{i}+\frac{\boldsymbol{\Sigma}_{i, i}+\sum_{j \neq i}^{n} \boldsymbol{\Sigma}_{i, j}}{\sum_{t=1}^{n} \boldsymbol{\Sigma}_{t, t}+\sum_{j \neq t}^{n} \boldsymbol{\Sigma}_{t, j}}\left(k-\sum_{j=1}^{n} \boldsymbol{\mu}_{j}\right)$ and variance $\bar{\sigma}_{i}^{2}=\boldsymbol{\Sigma}_{i, i}-$ $\frac{\left(\boldsymbol{\Sigma}_{i, i}+\sum_{j \neq i}^{n} \boldsymbol{\Sigma}_{i, j}\right)^{2}}{\sum_{t=1}^{n} \boldsymbol{\Sigma}_{t, t}+\sum_{j \neq t}^{n} \boldsymbol{\Sigma}_{t, j}}$. Further, the expectation of the marginal distribution is $\bar{\mu}_{i}$.

We also present theoretical results on a system of linear equality constraints.
Proposition 9. Let $\boldsymbol{z} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. The constrained distribution $p_{\boldsymbol{\theta}}(\boldsymbol{z} \mid \boldsymbol{A} \boldsymbol{z}=\boldsymbol{k})$ is equivalent to an $n-a$ dimensional multivariate Gaussian distribution with mean $\overline{\boldsymbol{\mu}} \in \mathbb{R}^{n-a}$ and covariance matrix $\bar{\Sigma} \in \mathbb{R}^{(n-a) \times(n-a)}$ defined as below,

$$
\bar{\mu}=\boldsymbol{E} \boldsymbol{\mu}+\boldsymbol{E} \boldsymbol{\Sigma} \boldsymbol{A}^{T}\left(\boldsymbol{A} \boldsymbol{\Sigma} \boldsymbol{A}^{T}\right)^{-1}(\boldsymbol{k}-\boldsymbol{A} \boldsymbol{\mu}) \quad \text { and } \quad \overline{\boldsymbol{\Sigma}}=\boldsymbol{E} \boldsymbol{\Sigma} \boldsymbol{E}^{T}-\boldsymbol{E} \boldsymbol{\Sigma} \boldsymbol{A}^{T}\left(\boldsymbol{A} \boldsymbol{\Sigma} \boldsymbol{A}^{T}\right)^{-1} \boldsymbol{A} \boldsymbol{\Sigma} \boldsymbol{E}^{T}
$$

where $\boldsymbol{E} \in \mathbb{R}^{n-a \times n}$ contains the first $n-a$ rows of an identity matrix $\boldsymbol{I} \in \mathbb{R}^{n \times n}$. The marginal distribution subject to the constraint $p_{\boldsymbol{\theta}}\left(z_{i} \mid \boldsymbol{A} \boldsymbol{z}=\boldsymbol{k}\right)$ can also be computed in closed form. Specifically, with $\boldsymbol{e}_{i}$ being a standard basis vector, it holds that $p_{\boldsymbol{\theta}}\left(z_{i} \mid \boldsymbol{A} \boldsymbol{z}=\boldsymbol{k}\right) \sim \mathcal{N}\left(\bar{\mu}_{i}, \bar{\sigma}_{i}^{2}\right)$, with $\bar{\mu}_{i}=\mu_{i}+\boldsymbol{e}_{i}^{T} \boldsymbol{\Sigma} \boldsymbol{A}\left(\boldsymbol{A} \boldsymbol{\Sigma} \boldsymbol{A}^{T}\right)^{-1}(\boldsymbol{k}-\boldsymbol{A} \boldsymbol{\mu})$ and $\bar{\sigma}_{i}^{2}=\boldsymbol{e}_{i}^{T} \boldsymbol{\Sigma} \boldsymbol{e}_{i}-\boldsymbol{e}_{i}^{T} \boldsymbol{\Sigma} \boldsymbol{A}^{T}\left(\boldsymbol{A} \boldsymbol{\Sigma} \boldsymbol{A}^{T}\right)^{-1} \boldsymbol{A} \boldsymbol{\Sigma} \boldsymbol{e}_{i}$.

## Related Work

A substantial amount of research has been devoted to estimating gradients for categorical random variables. [MMT16] and [JGP16] proposed to refactor the non-differentiable sample from a categorical distribution with a differentiable sample from a novel Gumbel-Softmax distribution, which enables automatic differentiation. Gradient estimation under linear equality constraints for categorical random variables has been widely studied. Existing methods either employ the score function and straight-through estimator or suggest custom relaxation [KSE16, CSW18, GWZ18, XE19b]. [XE19b] extends the Gumbel-softmax technique to k-subsets, enabling backpropagation for k-subset sampling. However, this comes at the trade-off of introducing some bias in the learning process due to the use of relaxed samples. While score function estimators offer a seemingly simple solution, it is widely acknowledged that they are prone to exhibiting exceedingly high variance. A recently introduced gradient estimator known as SIMPLE [AZN23a] surpasses its predecessors but is limited to Bernoulli random variables. Our work deviates from previous research on gradient estimation under linear equality constraints by investigating exact sampling as well as closed-form loss function for Gaussian and Poisson random variables, which are commonly used in scientific applications. We also extended the SIMPLE estimator [AZN23a] to Gaussian and Poisson random variables.

Extensive research has been conducted on numerical sampling from multivariate normal distributions while adhering to various constraints. [AMD14] reviewed classical Gibbs Sampling on a standard simplex (samples are positive and sum to one) and proposed using Hamiltonian Monte Carlo (HMC) methods. Efficient sampling methods for multivariate normal distributions truncated by hyperplanes $(\boldsymbol{A x}=\boldsymbol{b}$, where $\operatorname{dim}(\boldsymbol{x})=N$ and $\operatorname{rank}(\boldsymbol{A})=n<N)$ were investigated by [MBR22] and [CCZ17]. While these studies focus on numerical simulations, our study aims to derive closed-form solutions for not only linear equality constraint but also a system of linear equality constraint. Our theoretical findings, combined with closed-form loss functions, enable researchers to seamlessly incorporate constraints into machine learning models while


Figure 3.1: Comparisons of different gradient estimators for point-wise loss $\ell$ being L1 and L2 loss applied to Gaussian and Poisson variables are conducted. To evaluate the direction of the gradient, we utilized the cosine distance, defined as 1 - cosine similarity. The bias, variance, and error of the estimators are assessed using a sample size of 10,000 . The bias and variance for Marginal Expectation applied to Poison random variable are extremely close to zero but not identically zero.
maintaining the probabilistic nature of random variables.
Integrating constraints as background knowledge into deep learning models has been widely studied. Traditional methods propose to add penalty terms to the loss function for constraint violations [DGM12, XZF18, FBD19, BGS22, SGL24]. While these methods are relatively straightforward to implement, they do not guarantee constraint satisfaction. Alternative approaches, such as [ATC22b], [GL21], and [SDC24], incorporate background knowledge into the topology of the network itself. In the context of generative tasks, [DAG20] incorporate propositional logic constraints in Generative Adversarial Networks for structured object generation, while [MMS22] demonstrate the integration of [DKT07] with variational autoencoders.

### 3.2.2.3 Beyond Gaussian

In this section, we present the theoretical results when $\boldsymbol{z}$ are Poisson variables defined over discrete domains. Similar to the Gaussian setting, we find that when the element-wise loss $\ell$ is $L 1$ or $L 2$ loss and the mapping $f_{\boldsymbol{u}}$ is an identity function, the expected loss admits closed-form expressions.

Proposition 10 (Poisson Closed-form Expected Loss). Let $\boldsymbol{z}=\left(z_{1}, \ldots, z_{n}\right)^{T}$, where $z_{i} \sim \operatorname{Poisson}\left(\theta_{i}\right)$. Let $\boldsymbol{y}=\left(y_{1}, y_{2}, \ldots, y_{n}\right)^{T}$ be the ground truth vector subject to the equality constraint $\sum_{i=1}^{n} y_{i}=k$ with $k \in \mathbb{N}^{+}$. Then it holds that
i) when $\ell$ is L1 loss,

$$
\begin{aligned}
L(\boldsymbol{\theta}) & =\sum_{i=1}^{n}\left(k-\left\lfloor k p_{i}-d_{i}\right\rfloor\right) p_{i} \operatorname{Bin}\left(\left\lfloor k p_{i}-d_{i}\right\rfloor ; k, p_{i}\right) \\
& +\left\lfloor k p_{i}-d_{i}\right\rfloor\left(1-p_{i}\right) \operatorname{Bin}\left(\left\lfloor k p_{i}-d_{i}\right\rfloor ; k, p_{i}\right)-2 d_{i} F\left(\left\lfloor k p_{i}-d_{i}\right\rfloor ; k, p_{i}\right)+d_{i}
\end{aligned}
$$

ii) when $\ell$ is L2 loss, $L(\boldsymbol{\theta})=\sum_{i=1}^{n} k p_{i}\left(1-p_{i}\right)+k^{2} p_{i}^{2}-2 y_{i} k p_{i}+y_{i}^{2}$,
where Bin denotes the probability mass function (p.m.f.) of a binomial distribution and $F$ denotes a regularized incomplete beta function. $d_{i}=k p_{i}-y_{i}$ and $p_{i}=\frac{\theta_{i}}{\sum_{j=1}^{n} \theta_{j}}$.

In the general setting when there is no closed-form solution for the expected loss, we first show that exact sampling from the constrained distribution can be achieved as it takes its form as a multinomial distribution.

Proposition 11 (Poisson Constrained Distribution). Given $\boldsymbol{z}=\left(z_{1}, \ldots, z_{n}\right)^{T}$ with $z_{i} \sim \operatorname{Poisson}\left(\theta_{i}\right)$, the constrained distribution $p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{j=1}^{n} z_{n}=k\right)$ is equivalent to a multinomial distribution with parameter $k$ and probabilities $\frac{\theta_{1}}{\sum_{j=1}^{n} \theta_{j}}, \ldots, \frac{\theta_{n}}{\sum_{j=1}^{n} \theta_{j}}$.

In the backward pass, the results below allow us to derive gradient estimators with either the conditional marginals or the expectation of the conditional marginal as a differentiable proxy.

Proposition 12 (Poisson Conditional Marginal and Expectations). Given $\boldsymbol{z}=\left(z_{1}, \ldots, z_{n}\right)^{T}$ with $z_{i} \sim \operatorname{Poisson}\left(\theta_{i}\right)$, the conditional marginal $p_{\boldsymbol{\theta}}\left(z_{i} \mid \sum_{j=1}^{n} z_{j}=k\right)$ follows a binomial distribution with parameter $k$ and probability $\frac{\theta_{i}}{\sum_{j=1}^{\theta_{j}} \theta_{j}}$. Further, its expectation is $\frac{k \theta_{i}}{\sum_{j=1}^{k} \theta_{j}}$.

### 3.2.3 Empirical Evaluation

We conduct a comprehensive empirical analysis of our proposed method across four different tasks: 1) Synthetic experiments where the ground truth gradients are available and thus we can compare the bias and variance of different estimators. 2) A variational autoencoder (VAE) setting where the posterior of the latent variables is constrained by the linear equality and gradient estimators are necessary for end-to-end training. It allows us to compare the effectiveness of different estimators in terms of generation performance. 3) Another VAE setting where the underlying data distribution is constrained by linear equality. We compare various VAE models and their constrained counterparts, where the success of linear equality integration consistently boosts model performance. 4) A scientific application where a charge-neutral constraint is enforced into the charge predictions of the message-passing neural networks (MPNNs) and the expected loss admits a closed-form solution in this setting.

### 3.2.3.1 Synthetic Experiments

We consider synthetic settings where the ground truth gradients can be obtained by taking derivatives of the closed-form expected loss as stated in Section 3.2.1 such that we can compare how good the gradient approximations are for each gradient estimator. The distance between the estimated and the ground truth gradient vectors is measured by the cosine distance, defined as 1 cosine similarity. We evaluate the performance of gradient estimators on three metrics: bias, variance, and average error. We compare the two proposed gradient estimators, Constrained Marginal and Marginal Expectation, with two baseline estimators, Random and Unconstrained Marginal as defined in Table 3.1.

Results are shown in Figure 3.1, where Marginal Expectation significantly outperforms the


Figure 3.2: Comparison of gradient estimators for VAE with constrained latent space. Negative log-likelihood (NLL), negative ELBO (NELBO), and reconstruction loss (RL) are averaged over 5 trials.
others in all cases. Estimator Unconstrained Marginal has similar performances to Random which is expected since it discards the constraint information. What is interesting is that estimator Constrained Marginal, even though it is informed by the constraint, it also performs as bad as Random. Even though in the Bernoulli setting, Marginal Expectation and Constrained Marginal are the same estimator as shown in [AZN23a], we show that in the Gaussian and Poisson setting, the former is capable of providing decent gradient approximations while the latter is not.

### 3.2.3.2 Constrained Latent Space

We further test the four gradient estimators in Table 3.1 under the setting where the VAE model has its latent space constrained by linear equality, that is, the latent variables should sum to a constant. The VAE is trained on the MNIST dataset using the evidence lower bound (ELBO) as objective, which consists of a reconstruction loss (RL) and the KL divergence between a constrained approximate posterior $p_{\boldsymbol{\theta}}\left(\boldsymbol{z} \mid \sum_{i} z_{i}=k, \boldsymbol{x}\right)$ and a prior of the latent space.

Experiment results are presented in Figure 3.2 where the generative performance is evaluated using test negative likelihood, estimated using importance sampling [BGS16], ELBO, and reconstruction loss. Results show that the estimator Marginal Expectation outperforms all other estimators, which is consistent with the approximation results that we observe in the synthetic
experiments in Figure 3.1. Unconstrained Marginal and Constrained Margianl have similar performance, both better than Random.

### 3.2.3.3 Constrained Data Generation

We consider a setting where the underlying data generation distribution is constrained by linear equality as domain knowledge. We modify the MNIST dataset such that for each image, all the pixel values sum to be 100. The first row in Figure 3.3 shows the original MNIST images and our processed ones in the second row. The processed images are then the inputs to various VAE models and their constrained counterparts.

We consider three VAE models: Vanilla VAE [KW22], Ladder VAE [SRM16], and Graph VAE [HGM18]. We compare the performance of these models and their contained counterparts integrated with linear equality using estimator Marginal Expectation as it shows the best performance so far. All the VAE models are trained using ELBO as objectives.

The performance of the VAEs is evaluated from three aspects. (1) Generative Ability: We use test $\log$-likelihood (LL) $\log p_{\boldsymbol{\theta}}(\boldsymbol{x})$ estimated using 5,000 importance-weighted samples, ELBO, and Reconstruction Loss (RL) to measure the models' performance, accuracy, and generalization capabilities. (2) Constraint Violation Rate: We measure the proportion of reconstructed testing data that violates the equality constraint. (3) Training Time: We record the average training time of all the models for one epoch. All results are averaged over 5 independent runs. The results are summarized in Table 3.2.

We find out that the vanilla VAE and other VAEs have high constraint violation rates. On the contrary, our method can constrain the model such that their generated data satisfy the constraint while also achieving better generative performance due to the inductive bias. In order to show that the addition of Marginal Expectation has basically no impact on speed, for each model we measure the training time of 1 epoch. We also test the average training time where the results show that the constrained versions are less than $10 \%$ slower than the unconstrained version. Marginal Expectation adds less than 1 second of additional training time for VAE and Ladder

Table 3.2: Comparison on VAE Generative Ability. The constrained VAE models achieve similar or better performance in terms of their generative ability while strictly satisfying the constraints. The unconstrained counterparts have a high constraint violation rate.

| Model | $\mathbf{L L} \uparrow$ | $\mathbf{E L B O} \uparrow$ | $\mathbf{R L} \downarrow$ | Violation $\downarrow$ |
| :--- | ---: | ---: | ---: | ---: |
| VAE | $-22.42 \pm 0.29$ | $-23.41 \pm 0.22$ | $15.00 \pm 0.46$ | $0.30 \pm 0.06$ |
| Constrained VAE | $-\mathbf{2 1 . 4 8} \pm \mathbf{0 . 1 8}$ | $-\mathbf{2 2 . 6 2} \pm \mathbf{0 . 0 7}$ | $\mathbf{1 2 . 7 9} \pm \mathbf{0 . 1 1}$ | $\mathbf{0} \pm \mathbf{0}$ |
| Ladder VAE | $-24.25 \pm 0.07$ | $-30.84 \pm 0.51$ | $\mathbf{2 3 . 0 6} \pm \mathbf{0 . 5 4}$ | $0.38 \pm 0.02$ |
| Constrained Ladder VAE | $-\mathbf{2 3 . 8 6} \pm \mathbf{0 . 0 6}$ | $-\mathbf{3 0 . 7 8} \pm \mathbf{0 . 0 8}$ | $\mathbf{2 3 . 4 0} \pm \mathbf{0 . 1 6}$ | $\mathbf{0} \pm \mathbf{0}$ |
| Graph VAE | $-22.74 \pm 0.11$ | $-23.54 \pm 0.18$ | $15.45 \pm 0.41$ | $0.29 \pm 0.09$ |
| Constrained Graph VAE | $-\mathbf{2 1 . 6 1} \pm \mathbf{0 . 2 0}$ | $-\mathbf{2 2 . 5 3} \pm \mathbf{0 . 0 6}$ | $\mathbf{1 2 . 7 3} \pm \mathbf{0 . 2 1}$ | $\mathbf{0} \pm \mathbf{0}$ |

VAE.

### 3.2.3.4 Partial Charge Predictions for Metal-Organic Frameworks

Metal-organic frameworks (MOFs) represent a class of materials with a wide range of applications in chemistry and materials science. Predicting properties of MOFs, such as partial charges on metal ions, is essential for understanding their reactivity and performance in chemical processes. However, it is challenging due to the complex interactions between metal ions and ligands and the requirement that the predictions need to satisfy the charge neutral constraint, that is, an exactly-zero constraint.

We adopt the same model as [RSS20] where the charges are assumed to be Gaussian and the element loss is L1. Our model architecture extends the Message Passing Neural Network (MPNN) framework and incorporates linear-equality constraint for Gaussian variables, ensuring strict adherence to the critical constraint. The core innovation involves replacing the conventional L1 loss with the closed-form Gaussian loss function 6 as well as the negative log likelihood of the con-


Figure 3.3: The first row displays the original MNIST images, while the second row shows the constrained versions that are input to the VAEs. We compare the reconstructed images produced by various VAE models and their contained counterparts integrated with linear equality.
strained multivariate Gaussian. The L1 loss function penalizes deviations from the linear-equality constraint while considering the probabilistic nature of Gaussian variables, and the negative log likelihood loss models the observed data with higher probability. This comprehensive approach not only enables our model to capture complex structural relationships in MOFs but also ensures accurate predictions of partial charges while respecting the crucial linear-equality constraint, enhancing its applicability in a wide range of graph-based applications, including those pertaining to metal-organic frameworks.

Additionally, we also devise an ensemble methodology to enhance the predictive performance and robustness of our linear-equality constrained MPNN model. To achieve this, we adopt a systematic approach encompassing model variability, aggregation strategies and cross-validation. Two instances of the linear-equality constrained MPNN model are trained with variations in initialization. We apply the averaging aggregation technique to combine the predictions from
these models. Performance assessment is conducted through cross-validation techniques. The ensemble's performance is evaluated on a separate test dataset to ascertain its generalization ability. This ensemble approach not only elevates predictive accuracy but also fortifies the model's resilience, rendering it highly effective for complex tasks, including those pertaining to metalorganic frameworks.

The prediction performance of our four proposed approaches is presented in Table 3.3, compared with baseline approaches reported by [RSS20]. Results show that training using negative log likelihood loss and closed-form expected loss achieves better performance as MPNN (variance) which is considered to be the strongest baseline approach. When further combined with the ensemble method, our approach achieves significantly better predictions.

### 3.3 Discussion

In this chapter, we present the continuous counterpart of Simple and the subtlety of how to define the constrained marginal probabilities as a proxy for the gradient updates. For discrete samples, both constrained marginal probabilities and the constrained marginal expectation of a sample are the same but for continuous variables, they are different, and interestingly only the latter works well as shown in the empirical evaluations, showing that it is consistently more informative than the previous one.

Table 3.3: The prediction performances of different models for estimating partial charges on metal ions are presented. Comparing to the baseline MPNN (variance), both the closed-form loss function and Likelihood loss function yield superior Mean Absolute Deviation (MAD) results. Moreover, ensemble methods (second block) notably boost the predictive performance of all.

| METHOD | MAD $\downarrow$ | NLL $\downarrow$ |
| :--- | ---: | ---: |
| (charge neutrality enforcement) | mean $\pm$ std | mean $\pm$ std |
| Constant Prediction | $0.324 \pm 0.007$ | - |
| Element-mean (uniform) | $0.154 \pm 0.002$ | - |
| Element-mean (variance) | $0.153 \pm 0.002$ | - |
| MPNN (uniform) | $0.0260 \pm 0.0008$ | $109.8 \pm 6.9$ |
| MPNN (variance) | $0.0251 \pm 0.0010$ | $-19.9 \pm 71.1$ |
| Closed-form (ours) | $\mathbf{0 . 0 2 4 5} \pm \mathbf{0 . 0 0 0 9}$ | $1.14 \mathrm{e}+8 \pm 3.15 \mathrm{e}+8$ |
| Likelihood (ours) | $0.0248 \pm 0.0008$ | $\mathbf{- 2 5 2} \pm \mathbf{2 4 . 7}$ |
| MPNN (variance) + Ensemble | $0.0238 \pm 0.0007$ | $-45.2 \pm 55.8$ |
| Closed-form + Ensemble (ours) | $\mathbf{0 . 0 2 3 0} \pm \mathbf{0 . 0 0 0 8}$ | $1.51 \mathrm{e}+7 \pm 1.29 \mathrm{e}+7$ |
| Likelihood + Ensemble (ours) | $\mathbf{0 . 0 2 3 1} \pm \mathbf{0 . 0 0 0 7}$ | $\mathbf{- 1 8 0} \pm \mathbf{3 8 . 3}$ |

## CHAPTER 4

## Learning from Weak Supervisions as Constraints

In the previous two chapters, we consider "hard" constraints, that is, such constraints are guaranteed to be satisfied in the deep learning pipeline. In this chapter, we consider "soft" constraints instead, where the models are encouraged to satisfy such constraints in the learning process. Further, instead of enforcing the constraints into the model architectures, this chapter is dedicated to the integration of soft constraints by modifying the loss functions. Specifically, we choose the count-based weakly supervised learning tasks to show that by translating the count supervision into soft constraints and using our methods to incorporate such constraints, model performance can be greatly boosted.

### 4.1 Count-Based Weakly Supervised Learning

Weakly supervised learning [Zho18] enables a model to learn from data with restricted, partial or inaccurate labels, often known as weakly-labeled data. Weakly supervised learning fulfills a need arising in many real-world settings that are subject to privacy or budget constraints, such as privacy sensitive data [WIB11], medical image analysis [BDO18], clinical practice [QCC], personalized advertisement [BD20] and knowledge base completion [GTH15, ZD18], to name a few. In some settings, instance-level labels are unavailable. Instead, instances are grouped into bags with corresponding bag-level labels that are a function of the instance labels, e.g., the proportion of positive labels in a bag. A key insight that we bring forth is that such weak supervision can very often be construed as enforcing constraints on label counts of data.

More concretely, we consider three prominent weakly supervised learning paradigms. The

(a) Classical

(b) LLP

(c) MIL

(d) PU Learning

Table 4.1: A comparison of the tasks considered in the three weakly supervised settings, LLP (cf. Section 4.1.1), MIL (cf. Section 4.1.2) and PU learning (cf. Section 4.1.3), against the classical fully supervised setting for binary classification, using digits from the MNIST dataset.
first paradigm is known as learning from label proportions [QSC08]. Here the weak supervision consists in the proportion of positive labels in a given bag, which can be interpreted as the count of positive instances in such a bag. The second paradigm, whose supervision is strictly weaker than the former, is multiple instance learning [ML97, DLL01]. Here the bag labels only indicate the existence of at least one positive instance in a bag, which can be recast as to whether the count of positive instances is greater than zero. The third paradigm, learning from positive and unlabeled data [DDG99, LDG00], grants access to the ground truth labels for a subset of only the positive instances, providing only a class prior for what remains. We can recast the class prior as a distribution of the count of positive labels.

We formally introduce the aforementioned weakly supervised learning paradigms below. For notation, let $\mathcal{X} \in \mathbb{R}^{d}$ be the input feature space over $d$ features and $\mathcal{Y}=\{0,1\}$ be a binary label space. We write $\boldsymbol{x} \in \mathcal{X}$ and $\boldsymbol{y} \in \mathcal{Y}$ for the input and output random variables respectively. Recall that in fully-supervised binary classification, it is assumed that each feature and label pair $(\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{X} \times \mathcal{Y}$ is sampled independently from a joint distribution $p(\boldsymbol{x}, \boldsymbol{y})$. A classifier
$f$ is learned to minimize the risk $R(f)=\mathbb{E}_{(\boldsymbol{x}, \boldsymbol{y}) \sim p}[\ell(f(\boldsymbol{x}), \boldsymbol{y})]$ where $\ell:[0,1] \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$ is the cross entropy loss function. Typically, the true distribution $p(\boldsymbol{x}, \boldsymbol{y})$ is implicit and cannot be observed. Therefore, a set of $n$ training samples, $\mathcal{D}=\left\{\left(\boldsymbol{x}_{i}, \boldsymbol{y}_{i}\right)\right\}_{i=1}^{n}$, is used and the empirical risk, $\hat{R}(f)=\frac{1}{n} \sum_{i=1}^{n} \ell\left(f\left(\boldsymbol{x}_{i}\right), \boldsymbol{y}_{i}\right)$, is minimized in practice. In the count-based weakly supervised learning settings, the supervision is given at a bag level instead of an instance level. We formally introduce these settings as below.

### 4.1.1 Learning from Label Proportions

Learning from label proportions (LLP) [QSC08] assumes that each instance in the training set is assigned to bags and only the proportion of positive instances in each bag is known. One example is in light of the coronavirus pandemic, where infection rates were typically reported based on geographical boundaries such as states and counties. Each boundary can be treated as a bag with the infection rate as the proportion annotation.

The goal of LLP is to learn an instance-level classifier $f: \mathcal{X} \rightarrow[0,1]$ even though it is trained on bag-level labeled data. Formally, the training dataset consists of $m$ bags, denoted by $\mathcal{D}=\left\{\left(B_{i}, \tilde{y}_{i}\right)\right\}_{i=1}^{m}$ where each bag $B_{i}=\left\{\boldsymbol{x}_{j}\right\}_{j=1}^{k}$ consist of $k$ instances and this $k$ could vary among different bags. The bag proportions are defined as $\tilde{y}_{i}=\sum_{j=1}^{k} \boldsymbol{y}_{j} / k$ with $\boldsymbol{y}_{j}$ being the instance label that cannot be accessed and only $\tilde{y}_{i}$ is available during training. An example is shown in Figure 4.1b. We do not assume that the bags are non-overlapping while some existing work suffers from this limitation including [SZ20].

### 4.1.2 Multiple Instance Learning

Multiple instance learning (MIL) [ML97, DLL01] refers to the scenario where the training dataset consists of bags of instances, and labels are provided at bag level. However, in MIL, the bag label is a single binary label indicating whether there is a positive instance in the bag or not as opposed to a bag proportion defined in LLP. A real-world application of MIL lies in the field of drug activity [DLL01]. We can observe the effects of a group of conformations but not for any specific molecule, motivating a MIL setting. Formally, in MIL, the training dataset consists of $m$

Table 4.2: A summary of the labels and objective functions for all the settings considered in this chapter.

| Task | Label | Label Level | ObJECtive |
| :--- | :--- | :--- | :--- |
| Classical Fully Supervised | Binary $\boldsymbol{y}$ | Instance Level | $-\boldsymbol{y} \log p(\boldsymbol{y})-(1-\boldsymbol{y}) \log (1-p(\boldsymbol{y}))$ |
| Learning from Label Proportion | Continuous $\tilde{y}=\sum_{i} \boldsymbol{y}_{i} / k$ | Bag Level | $-\log p\left(\sum \hat{y}_{i}=k \tilde{y}\right)$ |
| Multiple Instance Learning | Binary $\tilde{y}=\max \left\{y_{i}\right\}$ | Bag Level | $-\tilde{y} \log p\left(\sum \hat{y}_{i} \geq 1\right)-(1-\tilde{y}) \log p\left(\sum_{i} \hat{y}_{i}=0\right)$ |
| Learning from Positive | Binary $\tilde{y}$ | Instance Level | $1) \mathbb{D}_{K L}\left(\operatorname{Bin}(k, \beta) \\| p\left(\sum_{i} \hat{y}_{i}\right)\right)$ <br> and Unlabeled Data |

bags, denoted by $\mathcal{D}=\left\{\left(B_{i}, \tilde{y}_{i}\right)\right\}_{i=1}^{m}$, with a bag consisting of $k$ instances, i.e., $B_{i}=\left\{\boldsymbol{x}_{j}\right\}_{j=1}^{k}$. The size $k$ can vary among different bags. For each instance $\boldsymbol{x}_{j}$, there exists an instance-level label $\boldsymbol{y}_{j}$ which is not accessible. The bag-level label is defined as $\tilde{y}_{i}=\max _{j}\left\{\boldsymbol{y}_{j}\right\}$. An example is shown in Figure 4.1c.

The main goal of MIL is to learn a model that predicts a bag label while a more challenging goal is to learn an instance-level predictor that is able to discover positive instances in a bag. In this work, we aim to tackle both by training an instance-level classifier whose predictions can be combined into a bag-level prediction as the last step.

### 4.1.3 Learning from Positive and Unlabeled Data

Learning from positive and unlabeled data or PU learning [DDG99, LDG00] refers to the setting where the training dataset consists of only positive instances and unlabeled data, and the unlabeled data can contain both positive and negative instances. A motivation of PU learning is persistence in the case of shifts to the negative-class distribution [PNS15], for example, a spam filter. An attacker may alter the properties of a spam email, making a traditional classifier require a new negative dataset [PNS15]. We note that taking a new unlabeled sample would be more efficient, motivating PU learning. Formally, in PU learning, the training dataset $\mathcal{D}=\mathcal{D}_{p} \cup \mathcal{D}_{u}$ where $\mathcal{D}_{p}=\left\{\left(\boldsymbol{x}_{i}, \tilde{y}_{i}=1\right)\right\}_{i=1}^{n_{p}}$ is the set of positive instances with $\boldsymbol{x}_{i}$ from $p(\boldsymbol{x} \mid \boldsymbol{y}=1)$ and $\tilde{y}$ denoting
whether the instance is labeled, and $\mathcal{D}_{u}=\left\{\left(\boldsymbol{x}_{i}, \tilde{y}_{i}=0\right)\right\}_{i=1}^{n_{u}}$ the unlabeled set with $\boldsymbol{x}_{i}$ from

$$
\begin{equation*}
p_{u}(\boldsymbol{x})=\beta p(\boldsymbol{x} \mid \boldsymbol{y}=1)+(1-\beta) p(\boldsymbol{x} \mid \boldsymbol{y}=0) \tag{4.1}
\end{equation*}
$$

where the mixture proportion $\beta:=p(\boldsymbol{y}=1 \mid \tilde{y}=0)$ is the fraction of positive instances among the unlabeled population. Although the instance label $\boldsymbol{y}$ is not accessible, its information can be inferred from the binary selection label $\tilde{y}$ : if the selection label $\tilde{y}=1$, it belongs to the positively labeled set, i.e., $p(\boldsymbol{y}=1 \mid \tilde{y}=1)=1$; otherwise, the instance $\boldsymbol{x}$ can be either positive or negative. An example of such a dataset is shown in Figure 4.1d.

The goal of PU learning is to train an instance-level classifier. However, it is not straightforward to learn from PU data and it is necessary to make assumptions to enable learning with positive and unlabeled data [BD20]. In this work, we make a commonly-used assumption for PU learning, selected completely at random (SCAR), which lies at the basis of many PU learning methods.

Definition 1 (SCAR). Labeled instances are selected completely at random, independent from input features and the positive distribution $p(\boldsymbol{x} \mid \boldsymbol{y}=1)$, that is, $p(\tilde{y}=1 \mid \boldsymbol{x}, \boldsymbol{y}=1)=p(\tilde{y}=1 \mid \boldsymbol{y}=1)$.

### 4.2 A Unified Approach: Count Loss

In this section, we derive objectives for the three weakly supervised settings, LLP, MIL, and PU learning, from first principles. Our proposed objectives bridge between neural outputs, which can be observed as counts, and arithmetic constraints derived from the weakly supervised labels. The idea is to capture how close the classifier is to satisfying the arithmetic constraints on its outputs. They can be easily integrated with deep learning models, and allow them to be trained end-to-end. For the three objectives, we show that they share the same computational building block: given $k$ instances $\left\{\boldsymbol{x}_{i}\right\}_{i=1}^{k}$ and an instance-level classifier $f$ that predicts $p\left(\hat{y}_{i} \mid \boldsymbol{x}_{i}\right)$ with $\hat{y}$ denoting the prediction variable, the problem of inferring the probability of the constraint on counts $\sum_{i=1}^{k} \hat{y}_{i}=s$ is to compute the count probability defined below:

$$
p\left(\sum_{i=1}^{k} \hat{y}_{i}=s \mid\left\{\boldsymbol{x}_{i}\right\}_{i=1}^{k}\right):=\sum_{\hat{\boldsymbol{y}} \in \mathcal{Y}^{k}} \llbracket \sum_{i=1}^{k} \hat{y}_{i}=s \rrbracket \prod_{i=1}^{k} p\left(\hat{y}_{i} \mid \boldsymbol{x}_{i}\right)
$$

Algorithm 5 Count Probability $p\left(\sum_{i=1}^{k} \hat{y}_{i}=s\right)$
Input: A set of $k \log$ probabilities $\left\{t_{i}\right\}_{i=1}^{k}$ with $t_{i}:=\log p\left(\hat{y}_{i}=1\right)$, the number of instances $k$, and a label sum $s$

Output: $\log$ probabilities $\log p\left(\sum_{i=1}^{k} \hat{y}_{i}=s\right)$ or a set of $\log$ probability $\left\{\log p\left(\sum_{i=1}^{k} \hat{y}_{i}=s\right)\right\}_{s=0}^{k}$ $/ / A[i, m]=\log p\left(\sum_{j=1}^{i} \boldsymbol{y}_{j}=m\right) \forall i, m$
Initialize an array $A$ to be -Inf everywhere
$A[0,0]=0 \quad / / p\left(\sum_{j=1}^{0} \boldsymbol{y}_{j}=0\right)=1$
Compute $t_{i}^{\prime} \leftarrow \log 1 \operatorname{mexp}\left(t_{i}\right) / / \log p\left(\boldsymbol{y}_{i}=0\right)$
for $i=1$ to $k$ do
for $m=0$ to $s$ do

$$
\begin{aligned}
& a_{+}=A[i-1, m-1]+t_{i} \\
& a_{-}=A[i-1, m]+t_{i}^{\prime}
\end{aligned}
$$

$$
A[i, m]=\operatorname{logsumexp}\left(a_{+}, a_{-}\right)
$$

return $A[k, s]$ or $A[k,:]$
where $\llbracket \rrbracket$ denotes the indicator function and $\hat{\boldsymbol{y}}$ denotes the vector $\left(\hat{y}_{1}, \cdots, \hat{y}_{k}\right)$. For succinctness, we omit the dependency on the input and simply write the count probability as $p\left(\sum_{i=1}^{k} \hat{y}_{i}=s\right)$. Next, we show how the objectives derived from first principles can be solved by using the count probability as an oracle. We summarize all proposed objectives in Table 4.2. Later, we will show how this seemingly intractable count probability can be efficiently computed by our proposed algorithm.

LLP setting. Given a bag $B=\left\{\boldsymbol{x}_{i}\right\}_{i=1}^{k}$ of size $k$ and its weakly supervised label $\tilde{y}$, by definition, it can be inferred that the number of positive instances (count) in the bag is $k \tilde{y}$. Our objective is to minimize the negative $\log$ probability $-\log p\left(\sum_{i} \hat{y}_{i}=k \tilde{y}\right)$. Notice that when each bag consists of only one instance, that is, when the bag-level supervisions are reduced to instance-level ones, this objective is exactly cross-entropy loss. We further show that our method is risk-consistent, that is, the optimal classifier under our proposed loss provides predictions consistent with the underlying risk as in the supervised learning setting.

MIL setting. Given a bag $B=\left\{\boldsymbol{x}_{i}\right\}_{i=1}^{k}$ of size $k$ and a single binary label $\tilde{y}$ as its weakly supervised label, we propose a cross-entropy loss as below

$$
\ell(B, \tilde{y})=-\tilde{y} \log p\left(\sum \hat{y}_{i} \geq 1\right)-(1-\tilde{y}) \log p\left(\sum \hat{y}_{i}=0\right) .
$$

Notice that in the above loss, the probability term $p\left(\sum \hat{y}_{i}=0\right)$ is accessible to the oracle for computing count probability, and the other probability term $p\left(\sum \hat{y}_{i} \geq 1\right)$ can simply be obtained from $1-p\left(\sum \hat{y}_{i}=0\right)$, i.e., the same call to the oracle since all prediction variables $\hat{y}_{i}$ are binary.

PU Learning setting. Recall that for the unlabeled data $\mathcal{D}_{u}$ in the training dataset, an unlabeled instance $\boldsymbol{x}_{i}$ is drawn from a mixture distribution as shown in Equation 4.1 parameterized by a mixture proportion $\beta=p(\boldsymbol{y}=1 \mid \tilde{y}=0)$. Under the SCAR assumption, even though only a class prior is given, we show that the mixture proportion can be estimated from the dataset.

Proposition 13. With SCAR assumption and a class prior $\alpha:=p(\boldsymbol{y}=1)$, the mixture proportion $\beta:=p(\boldsymbol{y}=1 \mid \tilde{y}=0)$ can be estimated from dataset $\mathcal{D}$.

Proof. First, the label frequency $p(\tilde{y}=1 \mid \boldsymbol{y}=1)$ denoted by $c$ can be obtained by

$$
c=\frac{p(\tilde{y}=1, \boldsymbol{y}=1)}{p(\boldsymbol{y}=1)}=\frac{p(\tilde{y}=1)}{p(\boldsymbol{y}=1)} \text { (by the definition of PU learning). }
$$

that is, $c=p(\tilde{y}=1) / \alpha$. Notice that $p(\tilde{y}=1)$ can be estimated from the dataset $\mathcal{D}$ by counting the proportion of the labeled instances. Thus, we can estimate the mixture proportion as below,

$$
\beta=\frac{p(\tilde{y}=0 \mid \boldsymbol{y}=1) p(\boldsymbol{y}=1)}{p(\tilde{y}=0)}=\frac{(1-p(\tilde{y}=1 \mid \boldsymbol{y}=1)) p(\boldsymbol{y}=1)}{1-p(\tilde{y}=1)}=\frac{(1-c) \alpha}{1-\alpha c} .
$$

The probabilistic semantic of the mixture proportion is that if we randomly draw an instance $\boldsymbol{x}_{i}$ from the unlabeled population, the probability that the true label $\boldsymbol{y}_{i}$ is positive would be $\beta$. Further, if we randomly draw $k$ instances, the distribution of the summation of the true labels $\sum_{i=1}^{k} \boldsymbol{y}_{i}$ conforms to a binomial distribution $\operatorname{Bin}(k, \beta)$ parameterized by the mixture proportion $\beta$, i.e.,

$$
\begin{equation*}
p\left(\sum_{i=1}^{k} \boldsymbol{y}_{i}=s\right)=\binom{k}{s} \beta^{s}(1-\beta)^{k-s} . \tag{4.2}
\end{equation*}
$$

Based on this observation, we propose an objective to minimize the KL divergence between the distribution of predicted label sum and the binomial distribution parameterized by the mixture proportion for a random subset drawn from the unlabeled population, that is,

$$
\mathbb{D}_{K L}\left(\operatorname{Bin}(k, \beta) \| p\left(\sum_{i=1}^{k} \hat{y}_{i}\right)\right)=\sum_{s=0}^{k} \operatorname{Bin}(s ; k, \beta) \log \frac{\operatorname{Bin}(s ; k, \beta)}{p\left(\sum_{i=1}^{k} \hat{y}_{i}=s\right)}
$$

where $\operatorname{Bin}(s ; k, \beta)$ denotes the probability mass function of the binomial distribution $\operatorname{Bin}(k, \beta)$. Again, the KL divergence can be obtained by $k+1$ calls to the oracle for computing count probability $p\left(\sum_{i=1}^{k} \hat{y}_{i}=s\right)$. The KL divergence is further combined with a cross entropy defined over labeled data $\mathcal{D}_{p}$ as in the classical binary classification training as the overall objective.

As an alternative, we propose an objective for the unlabeled data that requires fewer calls to the oracle: instead of matching the distribution of the predicted label sum with the binomial distribution, this objective matches only the expectations of the two distributions, that is, to $\operatorname{maximize} p\left(\sum_{i=1}^{k} \hat{y}_{i}=k \beta\right)$ where $k \beta$ is the expectation of the binomial distribution $\operatorname{Bin}(k, \beta)$. We present empirical evaluations of both proposed objectives in the experimental section.


Figure 4.1: An example of how to compute the count probability in a dynamic programming manner. Assume that an instance-level classifier predicts three instances to have $p\left(\boldsymbol{y}_{1}=1\right)=0.1$, $p\left(\boldsymbol{y}_{2}=1\right)=0.2$, and $p\left(\boldsymbol{y}_{3}=1\right)=0.3$ respectively. The algorithm starts from the top-left cell and propagates the results down right. A cell has its probability $p\left(\sum_{j=0}^{i} \boldsymbol{y}_{j}=s\right)$ computed by inputs from $p\left(\sum_{j=0}^{i-1} \boldsymbol{y}_{j}=s\right)$ weighted by $p\left(\boldsymbol{y}_{i}=0\right)$, and $p\left(\sum_{j=0}^{i-1} \boldsymbol{y}_{j}=s-1\right)$ weighted by $p\left(\boldsymbol{y}_{i}=1\right)$ respectively, as indicated by the arrows.

Tractable Computation of Count Probability In the previous section, we show how the count probability $p\left(\sum_{i=1}^{k} \hat{y}_{i}=s\right)$ serves as a computational building block for the objectives derived from first principles for the three weakly supervised learning settings. With a closer look at the count probability, we can see that given a set of instances, the classifier predicts an instancelevel probability for each and it requires further manipulation to obtain count information; actually, the number of joint labelings for the set can be exponential in the number of instances. Intractable as it seems, we show that it is indeed possible to derive a tractable computation for the count probability based on a result from [AZN23b].

Proposition 14. The count probability $p\left(\sum_{i=1}^{k} \hat{y}_{i}=s\right)$ of sampling $k$ prediction variables that sums to s from an unconstrained distribution $p(\boldsymbol{y})=\prod_{i=1}^{k} p\left(\hat{y}_{i}\right)$ can be computed exactly in time $\mathcal{O}(k s)$. Moreover, the set $\left\{p\left(\sum_{i=1}^{k} \hat{y}_{i}=s\right)\right\}_{s=0}^{k}$ can also be computed in time $\mathcal{O}\left(k^{2}\right)$.

The above proposition can be proved in a constructive way where we show that the count probability $p\left(\sum_{i=1}^{k} \hat{y}_{i}=s\right)$ can be computed in a dynamic programming manner. We provide an illustrative example of this computation in Figure 4.1. In practice, we implement this
computation in $\log$ space for numeric stability which we summarized as Algorithm 5, where function $\log 1$ mexp provides a numerically stable way to compute $\log 1 \operatorname{mexp}(x)=\log (1-$ $\exp (x))$ and function logsumexp a numerically stable way to compute $\log \operatorname{sumexp}(x, y)=$ $\log (\exp (x)+\exp (y))$. Notice that since we show it is tractable to compute the set $\left\{p\left(\sum_{i=1}^{k} \hat{y}_{i}=\right.\right.$ $s)\}_{s=0}^{k}$, for any two given label sum $s_{1}$ and $s_{2}$, a count probability $p\left(s_{1} \leq \sum_{i} \hat{y}_{i} \leq s_{2}\right)$ where the count lies in an interval, can also be exactly and tractably computed. This implies that our tractable computation of count probabilities can potentially be leveraged by other count-based applications besides the three weakly supervised learning settings in the last section.

## Related Work

Weakly Supervised Learning. Besides settings explored in our work there are many other weakly-supervised settings. One of which is semi-supervised learning, a close relative to PU Learning with the difference being that labeled samples can be both positive and negative [ZG22, Zhu05]. Another is label noise learning, which occurs when our instances are mislabeled. Two common variations involve whether noise is independent or dependent on the instance [FV13, SKP22]. A third setting is partial label learning, where each instance is provided a set of labels of which exactly one is true [CST11a]. An extension of this is partial multi-label learning, where among a set of labels, a subset is true [XH18].

Unified Approaches. There exists some literature in regards to "general" approaches for weakly supervised learning. One example being the method proposed in [Hul14], which provides a procedure that minimizes the empirical risk on "fuzzy" sets of data. The paper also establishes guarantees for model identification and instance-level recognition. Co-Training and Self-Training are also examples of similar techniques that are applicable to a wide variety of weakly supervised settings [BM98, Yar95]. Self-training involves progressively incorporating more unlabeled data via our model's prediction (with pseudo-label) and then training a model on more data as an iterative algorithm [KMZ21]. Co-Training leverages two models that have different "views" of the data and iteratively augment each other's training set with samples they deem as "well-classified".

They are traditionally applied to semi-supervised learning but can extend to multiple instance learning settings [LZW11, XTX13, LZS23].

LLP. [QSC08] first introduced an exponential family based approach that used an estimation of mean for each class. Others seek to minimize "empirical proportion risk" or EPR as in [YCK14], which is centered around creating an instance-level classifier that is able to reproduce the label proportions of each bag. As mentioned previously, more recent methods use bag posterior approximation and neural-based approaches [AC17, TL20]. One such method is Proportion Loss (PL) [TL20], which we contrast to our approach. This is computed by binary cross entropy between the averaged instance-level probabilities and ground-truth bag proportion.

MIL. MIL finds its earlier approaches with SVMs, which have been used quite prolifically and still remain one of the most common baselines. We start with MI-SVM/mi-SVM [ATH02] which are examples of transductive SVMs [CCG18] that seek a stable instance classification through repeated retraining iterations. MI-SVM is an example of an instance space method [CCG18], which identifies methods that classify instances as a preliminary step in the problem. This is in contrast to bag-space or embedded-space methods that omit the instance classification step. Furthermore, [WYT18] remains one of the hallmarks of the use of neural networks for MultiInstance Learning. [ITW18], utilize a similar approach but with attention-based mechanisms.

PU learning. [BD20] groups PU Learning paradigms into three main classes: two step, biased, and class prior incorporation. Biased learning techniques train a classifier on the entire dataset with the understanding that negative samples are subject to noise [BD20]. We will focus on a subset of biased learning techniques (Risk Estimators) as they are considered state-of-theart and relevant to us as baselines. The Unbiased Risk Estimator (uPU) provides an alternative to the inefficiencies in manually biasing unlabeled data [PNS14, PNS15]. Later, Non-negative Risk Estimator (nnPU) [KND17] accounted for weaknesses in the unbiased risk estimator such as overfitting.

Count Loss. To our knowledge, viewing the computation of the "bag posterior" as probabilistic is new. However, the prior approaches do this implicitly. Many approaches have tried to approximate the "bag posterior" by averaging the instance-level probabilities in a bag [AC17, TL20]. In MIL settings, among instance-level approaches, the MIL-pooling is an implicit "bag posterior" computation. These include mean, max, and log-sum-exp pooling to approximate the likelihood that a bag has at least one positive instance [WYT18]. But again, these are all approximations of what our computation does exactly. In PU Learning, to our best knowledge, the view of unlabeled data as a bag annotated with the mixture proportion is new.

Neuro-Symbolic Losses. In this section, we have dealt with a specific form of distributional constraint. Conversely, there has been a plethora of work exploring the integration of hard symbolic constraints into the learning of neural networks. This can take the form of enforcing a hard constraint [ATC22a], whereby the network's predictions are guaranteed to satisfy the prespecified constraints. Or it can take the form of a soft constraint [XZF18, MDK18, AWC21, AWC22, ALT22, ACB23] whereby the network is trained with an additional loss term that penalizes the network for placing any probability mass on predictions that violate the constraint. While in this work we focus on discrete linear inequality constraints defined over binary variables, there is existing work focusing on hybrid linear inequality constraints defined over both discrete and continuous variables and their tractability [BPB15, ZMY21b, ZMY20b]. The development of inference algorithms for such constraints and their applications such as Bayesian deep learning remain an active topic [ZB19, KMZ19, ZMY20a, ZB23a].

### 4.3 Empirical Evaluation

In this section, we present a thorough empirical evaluation of our proposed count loss on the three weakly supervised learning problems, $L L P, M I L$, and $P U$ learning. ${ }^{1}$

[^0]
### 4.3.1 Learning from Label Proportions

We experiment on two datasets: 1) Adult with 8192 training samples where the task is to predict whether a person makes over 50 k a year or not given personal information as input; 2) Magic Gamma Ray Telescope with 6144 training samples where the task is to predict whether the electromagnetic shower is caused by primary gammas or not given information from the atmospheric Cherenkov gamma telescope [DG17]. ${ }^{2}$

We follow [SZ20] where two settings are considered: one with label proportions uniformly on $\left[0, \frac{1}{2}\right]$ and the other uniformly on $\left[\frac{1}{2}, 1\right]$. Additionally, we experiment on a third setting with label proportions distributing uniformly on $[0,1]$ which is not considered in [SZ20] but is the most natural setting since the label proportion is not biased toward either 0 or 1 . We experiment on four bag sizes $n \in\{8,32,128,512\}$.

Count loss (CL) denotes our proposed approach using the loss objective defined in Table 4.2 for LLP. We compare our approach with a mutual contamination framework for LLP (LMMCM) [SZ20] and against Proportion Loss (PL) [TL20].

Results and Discussions We show our results in Table 4.3. Our method showcases superior results against the baselines on both datasets and variations in bag sizes. Especially in cases with lower bag sizes, i.e. 8,32 , CL greatly outperforms all other methodologies. Among our baselines are methods that approximate the bag posterior (PL), which we show to be less effective than optimizing the exact bag posterior with CL.

### 4.3.2 Multiple Instance Learning

We first experiment on the MNIST dataset [LeC98] and follow the MIL experimental setting in [ITW18]: the training and test set bags are randomly sampled from the MNIST training and test set respectively; each bag can have images of digits from 0 to 9 , and bags with the digit 9 are labeled positive. Moreover, the dataset is constructed in a balanced way such that there is an equal

[^1]


Figure 4.2: MIL MNIST dataset experiments with decreased numbers of training bags and lower bag size. Left: bag sizes sampled from $\mathcal{N}(10,2)$; Right: bag sizes sampled from $\mathcal{N}(5,1)$. We plot the mean test AUC (aggregated over 3 trials) with standard errors for 4 bag sizes. Best viewed in color.
amount of positively and negatively labeled bags as in [ITW18]. The task is to train a classifier that is able to predict bag labels; the more challenging task is to discover key instances, that is, to train a classifier that identifies images of digit 9. Following [ITW18], we consider three settings that vary in the bag generation process: in each setting, bags have their sizes generated from a normal distribution being $\mathcal{N}(10,2), \mathcal{N}(50,10), \mathcal{N}(100,20)$ respectively. The number of bags in training set $n$ is in $\{50,100,150,200,300,400,500\}$. Thus, we have $3 \times 7=21$ settings in total. Additionally, we introduce experimental analysis on how the performance of the learning methods would degrade as the number of bags and total samples in training set decreases, by modulating the number of training bags $n$ to be $\{10,20,30,40\}$ and selecting bag sizes from $\mathcal{N}(5,1)$ and $\mathcal{N}(10,2)$.

We also experiment on the Colon Cancer dataset [SRT16] to simulate a setting where bag instances are not independent. The dataset consists of 100 total hematoxylin-eosin (H\&E) stained images, each of which contains images of cell nuclei that are classified as one of: epithelial, inflammatory, fibroblast, and miscellaneous. Each image represents a bag and instances are $27 \times 27$ patches extracted from the original image. A positively labeled bag or image is one that contains the epithelial nuclei. For both datasets, we include the Attention and Gated Attention mecha-


Figure 4.3: A test bag from our MIL experiments, where we set only the digit 9 as a positive instance. Highlighted in red are digits identified to be positive with corresponding probability beneath.
nism [ITW18] as baselines. We also use the MIL objective defined in Table 4.2.
Results and Discussions For the MNIST experiments, CL is able to outperform all other baselines or exhibit highly comparable performance for bag-level predictions as shown in Ta ble 4.4. A more interesting setting is to compare how robust the learning methods are if the number of training bags decreases. [WYT18] claim that instance-level classifiers tend to lose against embedding-based methods. However, we show in our experiment that this is not true in all cases as seen in Figure 4.2. While Attention and Gated Attention are based on embedding, they suffer from a more severe drop in predictive performance than CL when the number of training bags drops from 40 to 10; our method shows great robustness and consistently outperforms all baselines. The rationale we provide is that with a lower number of training instances, we need more supervision over the limited samples we have. Our constraint provides this additional supervision, which accounts for the difference in performance.

We provide an additional investigation in Figure 4.3 to show that our approach learns effectively and delivers accurate instance-level predictions under bag-level supervision. In Figure 4.3, we can see that even though the classifier is trained on feedback about whether a bag contains the digit 9 or not, it accurately discovers all images of digit 9 .

Our experimental results on the Colon Cancer dataset are shown in Table 4.5. We show that both our proposed objectives are able to consistently outperform baseline methods on all met-


Figure 4.4: MNIST17 setting for PU Learning: We compute the average discrete distribution for CL and CVIR, over 5 test bags, each of which contain 100 instances. A ground truth binomial distribution of counts is also shown.
rics. Interestingly, we do not expect CL to perform well when instances in a bag are dependent; however, the results indicate that our count loss is robust to these settings.

### 4.3.3 Learning from Positive and Unlabeled Data

We experiment on dataset MNIST and CIFAR-10 [KH09a], following the four simulated settings from [GWS21]: 1) Binarized MNIST: the training set consist of images of digits $0-9$ and images with digits in range [0, 4] are positive instances while others as negative; 2) MNIST17: the training set consist of images of digits 1 and 7 and images with digit 1 are defined as positive while 7 as negative; 3) Binarized CIFAR: the training set consists of images from ten classes and images from the first five classes is defined as positive instances while others as negative; 4) CIFAR Cat vs. Dog: the training set consist of images of cats and dogs and images of cats are defined as positive while dogs as negative. The mixture proportion is 0.5 in all experiments. The performance is evaluated using the accuracy on a test set of unlabeled data.

As shown in Table 4.2, we propose two objectives for PU learning. Our first objective is denoted by CL whereas the second approach is denoted by CL-expect. We compare against the

Conditional Value Ignoring Risk approach (CVIR) [GWS21], nnPU [KND17], and uPU [PNS15].
Results and Discussions Accuracy results are presented in Table 4.6 where we can see that our proposed methods perform better than baselines on 3 out of the 4 simulated PU learning settings. CL-expect builds off a similar "exactly-k" count approach, which we have shown to work well in the label proportion setting. The more interesting results are from CL where we fully leverage the information from a distribution as supervision instead of simply using the expectation. We think of this as applying a loss on each count weighted by their probabilities from the binomial distribution. We provide further evidence that our proposed count loss effectively guides the classifier towards predicting a binomial distribution as shown in Figure 4.4: we plot the count distributions predicted by CL and CVIR as well as the ground-truth binomial distribution. We can see that CL is able to generate the expected distribution, proving the efficacy of our approach.

### 4.4 Discussion

In this chapter, we present a unified approach to several weakly-supervised tasks, i.e., LLP, MIL, PU. We construct our approach based on the idea of using weak labels to constrain count-based probabilities computed from model outputs. A future direction for our work can be to extend to multi-class classification as well as explore the applicability to other weakly-supervised settings, e.g. label noise learning, semi-supervised learning, and partial label learning [CST11b, NDR13, Zhu05].

Table 4.3: LLP results across different bag sizes. We report the mean and standard deviation of the test AUC over 5 seeds for each setting. The highest metric for each setting is shown in boldface.

| Dataset | Dist | Method | 8 | 32 | 128 | 512 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| Adult | $\left[0, \frac{1}{2}\right]$ | PL | $0.8889 \pm 0.0024$ | $0.8782 \pm 0.0036$ | $\mathbf{0 . 8 7 4 3} \pm \mathbf{0 . 0 0 3 9}$ | $0.8678 \pm 0.0085$ |
| Adult | $\left[0, \frac{1}{2}\right]$ | LMMCM | $0.8728 \pm 0.0019$ | $0.8693 \pm 0.0047$ | $0.8669 \pm 0.0041$ | $0.8674 \pm 0.0040$ |
| Adult | $\left[0, \frac{1}{2}\right]$ | CL (Ours) | $\mathbf{0 . 8 9 8 4} \pm \mathbf{0 . 0 0 1 3}$ | $\mathbf{0 . 8 8 4 8} \pm \mathbf{0 . 0 0 4 1}$ | $\mathbf{0 . 8 7 4 3} \pm \mathbf{0 . 0 0 5 2}$ | $\mathbf{0 . 8 7 0 3} \pm \mathbf{0 . 0 0 7 0}$ |
| Adult | $\left[\frac{1}{2}, 1\right]$ | PL | $0.8781 \pm 0.0038$ | $0.8731 \pm 0.0035$ | $\mathbf{0 . 8 6 9 9} \pm \mathbf{0 . 0 0 5 7}$ | $0.8556 \pm 0.0180$ |
| Adult | $\left[\frac{1}{2}, 1\right]$ | LMMCM | $0.8584 \pm 0.0164$ | $0.8644 \pm 0.0052$ | $0.8601 \pm 0.0045$ | $0.8500 \pm 0.0186$ |
| Adult | $\left[\frac{1}{2}, 1\right]$ | CL (Ours) | $\mathbf{0 . 8 8 5 4} \pm \mathbf{0 . 0 0 2 2}$ | $\mathbf{0 . 8 7 3 8} \pm \mathbf{0 . 0 0 3 9}$ | $0.8675 \pm 0.0043$ | $\mathbf{0 . 8 6 0 7} \pm \mathbf{0 . 0 0 5 6}$ |
| Adult | $[0,1]$ | PL | $0.8884 \pm 0.0030$ | $0.8884 \pm 0.0008$ | $\mathbf{0 . 8 8 7 9} \pm \mathbf{0 . 0 0 2 5}$ | $\mathbf{0 . 8 8 2 8} \pm \mathbf{0 . 0 0 5 1}$ |
| Adult | $[0,1]$ | LMMCM | $0.8831 \pm 0.0026$ | $0.8819 \pm 0.0006$ | $0.8821 \pm 0.0017$ | $0.8786 \pm 0.0052$ |
| Adult | $[0,1]$ | CL (Ours) | $\mathbf{0 . 8 9 8 5} \pm \mathbf{0 . 0 0 1 0}$ | $\mathbf{0 . 8 8 9 1} \pm \mathbf{0 . 0 0 1 3}$ | $0.8871 \pm 0.0021$ | $0.8790 \pm 0.0056$ |
| Magic | $\left[0, \frac{1}{2}\right]$ | PL | $0.8900 \pm 0.0095$ | $0.8510 \pm 0.0032$ | $0.8405 \pm 0.0110$ | $0.8332 \pm 0.0149$ |
| Magic | $\left[0, \frac{1}{2}\right]$ | LMMCM | $0.8918 \pm 0.0077$ | $0.8799 \pm 0.0113$ | $0.8753 \pm 0.0157$ | $0.8734 \pm 0.0092$ |
| Magic | $\left[0, \frac{1}{2}\right]$ | CL (Ours) | $\mathbf{0 . 9 0 8 8} \pm \mathbf{0 . 0 0 5 6}$ | $\mathbf{0 . 8 8 3 0} \pm \mathbf{0 . 0 0 9 7}$ | $\mathbf{0 . 8 9 2 6} \pm \mathbf{0 . 0 0 4 9}$ | $\mathbf{0 . 8 8 6 4} \pm \mathbf{0 . 0 1 0 7}$ |
| Magic | $\left[\frac{1}{2}, 1\right]$ | PL | $0.9066 \pm 0.0016$ | $0.8818 \pm 0.0108$ | $0.8769 \pm 0.0101$ | $0.8429 \pm 0.0443$ |
| Magic | $\left[\frac{1}{2}, 1\right]$ | LMMCM | $0.8911 \pm 0.0083$ | $0.8790 \pm 0.0091$ | $0.8684 \pm 0.0046$ | $0.8567 \pm 0.0292$ |
| Magic | $\left[\frac{1}{2}, 1\right]$ | CL (Ours) | $\mathbf{0 . 9 1 0 5} \pm \mathbf{0 . 0 0 2 0}$ | $\mathbf{0 . 8 9 8 0} \pm \mathbf{0 . 0 0 5 9}$ | $\mathbf{0 . 8 8 5 1} \pm \mathbf{0 . 0 2 5 5}$ | $\mathbf{0 . 8 8 1 6} \pm \mathbf{0 . 0 0 8 3}$ |
| Magic | $[0,1]$ | PL | $0.9039 \pm 0.0029$ | $0.8870 \pm 0.0037$ | $0.9002 \pm 0.0092$ | $0.8807 \pm 0.0200$ |
| Magic | $[0,1]$ | LMMCM | $0.9070 \pm 0.0026$ | $0.9048 \pm 0.0058$ | $0.9113 \pm 0.0058$ | $0.8934 \pm 0.0097$ |
| Magic | $[0,1]$ | CL (Ours) | $\mathbf{0 . 9 1 7 3} \pm \mathbf{0 . 0 0 1 8}$ | $\mathbf{0 . 9 1 0 2} \pm \mathbf{0 . 0 0 5 7}$ | $\mathbf{0 . 9 1 4 6} \pm \mathbf{0 . 0 0 5 1}$ | $\mathbf{0 . 9 0 8 8} \pm \mathbf{0 . 0 0 3 9}$ |

Table 4.4: MIL experiment on the MNIST dataset. Each block represents a different distribution from which we draw bag sizes-First Block: $\mathcal{N}(10,2)$, Second Block: $\mathcal{N}(50,10)$, Third Block: $\mathcal{N}(100,20)$. We run each experiment for 3 runs and report mean test AUC with standard error. The highest metric for each setting is shown in boldface.

| Training Bags | 50 | 100 | 150 | 200 | 300 | 400 | 500 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| Gated Attention | $0.775 \pm 0.034$ | $0.894 \pm 0.012$ | $0.935 \pm 0.005$ | $0.939 \pm 0.006$ | $\mathbf{0 . 9 6 3} \pm \mathbf{0 . 0 0 2}$ | $0.959 \pm 0.002$ | $\mathbf{0 . 9 6 6} \pm \mathbf{0 . 0 0 3}$ |
| Attention | $0.807 \pm 0.026$ | $\mathbf{0 . 9 1 3} \pm \mathbf{0 . 0 0 6}$ | $\mathbf{0 . 9 4 0} \pm \mathbf{0 . 0 0 4}$ | $0.942 \pm 0.007$ | $0.957 \pm 0.002$ | $0.961 \pm 0.005$ | $0.965 \pm 0.004$ |
| CL (Ours) | $\mathbf{0 . 8 1 8} \pm \mathbf{0 . 0 2 4}$ | $0.906 \pm 0.009$ | $0.929 \pm 0.005$ | $\mathbf{0 . 9 4 6} \pm \mathbf{0 . 0 0 1}$ | $0.952 \pm 0.004$ | $\mathbf{0 . 9 6 2} \pm \mathbf{0 . 0 0 2}$ | $0.963 \pm 0.002$ |
| Gated Attention | $\mathbf{0 . 9 4 3} \pm \mathbf{0 . 0 0 5}$ | $0.949 \pm 0.009$ | $\mathbf{0 . 9 7 0} \pm \mathbf{0 . 0 0 5}$ | $\mathbf{0 . 9 7 7} \pm \mathbf{0 . 0 0 1}$ | $0.983 \pm 0.002$ | $0.986 \pm 0.004$ | $\mathbf{0 . 9 8 7} \pm \mathbf{0 . 0 0 2}$ |
| Attention | $0.936 \pm 0.010$ | $\mathbf{0 . 9 6 2} \pm \mathbf{0 . 0 0 6}$ | $\mathbf{0 . 9 7 0} \pm \mathbf{0 . 0 0 1}$ | $\mathbf{0 . 9 7 7} \pm \mathbf{0 . 0 0 2}$ | $0.981 \pm 0.002$ | $\mathbf{0 . 9 8 7} \pm \mathbf{0 . 0 0 1}$ | $\mathbf{0 . 9 8 7} \pm \mathbf{0 . 0 0 2}$ |
| CL (Ours) | $0.939 \pm 0.010$ | $0.960 \pm 0.002$ | $0.964 \pm 0.007$ | $0.972 \pm 0.002$ | $\mathbf{0 . 9 8 2} \pm \mathbf{0 . 0 0 3}$ | $0.982 \pm 0.001$ | $\mathbf{0 . 9 8 7} \pm \mathbf{0 . 0 0 2}$ |
| Gated Attention | $0.975 \pm 0.003$ | $0.981 \pm 0.004$ | $0.992 \pm 0.002$ | $0.987 \pm 0.004$ | $\mathbf{0 . 9 9 6} \pm \mathbf{0 . 0 0 1}$ | $\mathbf{0 . 9 9 8} \pm \mathbf{0 . 0 0 1}$ | $0.990 \pm 0.004$ |
| Attention | $\mathbf{0 . 9 8 4} \pm \mathbf{0 . 0 0 1}$ | $0.982 \pm 0.001$ | $\mathbf{0 . 9 9 6} \pm \mathbf{0 . 0 0 0}$ | $0.987 \pm 0.007$ | $0.992 \pm 0.004$ | $0.994 \pm 0.002$ | $0.998 \pm 0.000$ |
| CL (Ours) | $0.981 \pm 0.007$ | $\mathbf{0 . 9 8 9} \pm \mathbf{0 . 0 0 0}$ | $\mathbf{0 . 9 9 6} \pm \mathbf{0 . 0 0 2}$ | $\mathbf{0 . 9 9 5} \pm \mathbf{0 . 0 0 1}$ | $\mathbf{0 . 9 9 6} \pm \mathbf{0 . 0 0 2}$ | $0.993 \pm 0.003$ | $\mathbf{0 . 9 9 9} \pm \mathbf{0 . 0 0 1}$ |

Table 4.5: MIL: We report mean test accuracy, AUC, F1, precision, and recall averaged over 5 runs with std. error on the Colon Cancer dataset. The highest value for each metric is shown in boldface.

| Method | Accuracy | AUC | F1 | Precision | Recall |
| :--- | :--- | :--- | :--- | :--- | :--- |
| Gated Attention | $0.909 \pm 0.014$ | $0.908 \pm 0.013$ | $0.886 \pm 0.021$ | $0.916 \pm 0.020$ | $0.879 \pm 0.020$ |
| Attention | $0.893 \pm 0.015$ | $0.890 \pm 0.008$ | $0.876 \pm 0.017$ | $0.908 \pm 0.016$ | $0.879 \pm 0.018$ |
| CL (Ours) | $\mathbf{0 . 9 1 5} \pm \mathbf{0 . 0 0 8}$ | $\mathbf{0 . 9 1 2} \pm \mathbf{0 . 0 1 0}$ | $\mathbf{0 . 9 0 3} \pm \mathbf{0 . 0 1 0}$ | $\mathbf{0 . 9 3 6} \pm \mathbf{0 . 0 1 4}$ | $\mathbf{0 . 8 9 8} \pm \mathbf{0 . 0 0 7}$ |

Table 4.6: PU Learning: We report accuracy and standard deviation on a test set of unlabeled data, which is aggregated over 3 runs. The results from CVIR, nnPU, and uPU are aggregated over 10 epochs, as in [GWS21], while we choose the single best epoch based on validation for our approaches. The highest metric for each setting is shown in boldface.

| Dataset | Network | CL-expect (Ours) | CL (Ours) | CVIR | $\mathbf{n n P U}$ | $\mathbf{n P U}$ |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| Binarized MNIST | MLP | $95.9 \pm 0.15$ | $\mathbf{9 6 . 4} \pm \mathbf{0 . 0 1}$ | $96.3 \pm 0.07$ | $96.1 \pm 0.14$ | $95.2 \pm 0.19$ |
| MNIST17 | MLP | $98.7 \pm 0.17$ | $\mathbf{9 9 . 0} \pm \mathbf{0 . 1 9}$ | $98.7 \pm 0.09$ | $98.4 \pm 0.20$ | $98.4 \pm 0.09$ |
| Binarized CIFAR | ResNet | $79.2 \pm 0.27$ | $80.1 \pm 0.34$ | $\mathbf{8 2 . 3} \pm \mathbf{0 . 1 8}$ | $77.2 \pm 1.03$ | $76.7 \pm 0.74$ |
| CIFAR Cat vs. Dog | ResNet | $\mathbf{7 6 . 5} \pm \mathbf{1 . 8 6}$ | $74.8 \pm 1.64$ | $73.3 \pm 0.94$ | $71.8 \pm 0.33$ | $68.8 \pm 0.53$ |

Part II

## Reasoning

## CHAPTER 5

## Foundations: Weighted Model Integration

In the previous chapters, we demonstrate that the learning performance under constraints highly depends on the capability of probabilistic reasoning over constraints. This chapter addresses another fundamental challenge in Neurosymbolic AI: how to perform scalable and reliable probabilistic reasoning over expressive symbolic constraints. Instead of considering specific constraints, we aim for a general framework that handles complex constraints and performs probabilistic reasoning in a principled way.

Weighted model integration (WMI) is a such a framework for performing advanced probabilistic inference in hybrid domains, i.e., on distributions over mixed continuous-discrete random variables and in the presence of complex logical and arithmetic constraints. This chapter is dedicated to provide the background on WMI.

### 5.1 Overview

In many real-world scenarios, performing probabilistic inference requires reasoning over domains with complex logical and arithmetic constraints while dealing with variables that are heterogeneous in nature, i.e., both continuous and discrete.

Consider for example the task of matching players in a game by their skills. Performing probabilistic inference for this task has been popularized by [MCZ18] and is at the core of several online gaming services. A probabilistic model for this task has to deal with continuous variables, such as the player and team performance, and reason over discrete attributes such as membership in a squad and the achieved scores. Moreover, such a model would need to take into account
constraints such as the team performance being bounded by that of the players in it, and that forming a squad boosts performance. Ultimately, this translates into performing probabilistic inference in the presence of logical and arithmetic constraints and dependencies.

These hybrid scenarios are beyond the reach of probabilistic models including variational autoencoders [KW13] and generative adversarial networks [GPM14], whose inference capabilities, despite their recent success, are limited. Classical probabilistic graphical models [KF09], while providing more flexible inference routines, are generally incapacitated when dealing with continuous and discrete variables at once [SW11], or they tend to make simplistic [HG95, LW89] or overly strong assumptions about their parametric forms [YBR14]. Even recent efforts in modeling these hybrid scenarios while delivering tractable inference [MVD18, VMP19] can not perform inference in the presence of complex constraints.

Weighted Model Integration (WMI) [BPB15, MPS17] is a recent framework for probabilistic inference that offers all the aforementioned "ingredients" needed for hybrid probabilistic reasoning with logical constraints, by design. WMI leverages the expressive language of Satisfiability Modulo Theories (SMT) [BMR10] for describing problems over continuous and discrete variables. Moreover, WMI provides a principled way to perform hybrid probabilistic inference: asking for the probability of a complex query with logical and arithmetic constraints can be done by integrating weight functions over the regions that satisfy the constraints and query at hand.

### 5.2 Formalization

Notation. We use uppercase letters for random variables (e.g., $X, B$ ) and lowercase letters for their assignments (e.g., $x, b$ ). Bold uppercase letters denote sets of variables (e.g., $\mathcal{X}, \mathbf{B}$ ) and their lowercase denote their assignments (e.g., $\boldsymbol{x}, \boldsymbol{b}$ ). We represent logical formulas by capital Greek letters, (e.g., $\Lambda, \Phi, \Delta$ ), and literals (i.e., atomic formulas or their negation) by lowercase ones (e.g., $\phi, \delta$ ) or $\ell$. We denote satisfaction of a formula $\Phi$ by one assignment $\boldsymbol{x}$ by $\boldsymbol{x} \mid=\Phi$ and we denote its corresponding indicator function as $\llbracket \boldsymbol{x} \mid=\Phi \rrbracket$. For undirected graphs, neigh denotes the set of neighboring nodes; for directed ones, pa and ch denote the parent node and the set of
child nodes respectively.

Satisfiability Modulo Theories (SMT). SMT [BT18] generalizes the well-known SAT problem [BHM09] to determining the satisfiability of a logical formula w.r.t. a decidable theory. Rich mixed logical/arithmetic constraints can be expressed in SMT for hybrid domains. In particular, we consider quantifier-free SMT formulas in the theory of linear arithmetic over the reals, or $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$. Here, formulas are propositional combinations of atomic Boolean literals and of atomic $\mathcal{L R} \mathcal{A}$ literals over real variables, for which satisfaction is defined in a natural way. W.l.o.g. we assume SMT formulas to be in conjunctive normal form (CNF). In the following, we will use the shorthand SMT to denote $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$.

Example 1 (SMT representation of a skill matching system). In a skill rating system for online games, the team performance $X_{\mathcal{T}}$ of each team $\mathcal{T}$ is defined by the performance $X_{i}$ of each player $i$ in team $\mathcal{T}$, both of which are real variables. The team performance $X_{\mathcal{T}}$ is also related to a Boolean variable B indicating whether players in the team form a squad, i.e., a group of friends, which offsets (boosts) the team performance. We can build an SMT formula $\Gamma$ of the relationship among these variables as follows. For brevity, we omit the domains for real variables in the formula.

$$
\Gamma:=\bigwedge_{i \in \mathcal{T}}\left|X_{\mathcal{T}}-X_{i}\right|<1 \bigwedge\left(B \Rightarrow X_{\mathcal{T}}>2\right)
$$

We show in Figure 5.1 the feasible regions offormula $\Gamma$ i.e., the volumes for which the constraints are satisfied.

Example 2 (SMT representation of a house price model). For a house $i$, let price ${ }_{i}$ be its price and sqft its square footage. We can build a simple $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ formula of the relationship between these real variables, with the corresponding solution space depicted in Figure 5.2. That is, $\operatorname{SMT}(\mathcal{L} \mathcal{R} \mathcal{A})$ formula $\Gamma_{i}$ is

$$
\begin{gathered}
\Gamma_{i}:=\left(\text { price }_{i}<10 \cdot \text { sqft }_{i}+1000\right) \vee\left(\text { price }_{i}<20 \cdot \text { sqft }_{i}+100\right) \\
\wedge\left(0<\text { price }_{i}<3000\right) \wedge\left(0<\text { sqft }_{i}<200\right) .
\end{gathered}
$$



Figure 5.1: Feasible region (left) of formula $\Gamma$ with one player and primal graph (right) of formula $\Gamma$ with $n$ players from Example 1.

Weighted Model Integration (WMI). Weighted Model Integration (WMI) [BPB15, MPS17] provides a framework for probabilistic inference with models defined over the logical constraints given by SMT formulas.

Definition 2. (WMI) Let $\mathcal{X}$ be a set of continuous random variables defined over $\mathbb{R}$, and $\mathbf{B}$ a set of Boolean random variables defined over $\mathbb{B}=\{$ true, false $\}$. Given an SMT formula $\Delta$ over $\mathcal{X}$ and $\mathbf{B}$, and a weight function $w:(\boldsymbol{x}, \boldsymbol{b}) \mapsto \mathbb{R}^{+}$belonging to some parametric weight function family $\Omega$, the weighted model integration (WMI) task computes

$$
\begin{equation*}
\operatorname{WMI}(\Delta, w ; \mathcal{X}, \mathbf{B}) \triangleq \sum_{\boldsymbol{b} \in \mathbb{B}^{|\mathbb{B}|}} \int_{(\boldsymbol{x}, \boldsymbol{b}) \neq \Delta} w(\boldsymbol{x}, \boldsymbol{b}) d \boldsymbol{x} . \tag{5.1}
\end{equation*}
$$

That is, summing over all possible Boolean assignments $\boldsymbol{b} \in \mathbb{B}^{|\mathbf{B}|}$ while integrating over the weighted assignments of $\mathcal{X}$ making the evaluation of the formula SAT: $(\boldsymbol{x}, \boldsymbol{b})=\Delta$.

Weight functions $w$ are usually defined as products of literal weights [BPB15, CD08, ZB19]. That is, for a set of literals $\mathcal{L}$, a set of per-literal weight functions $\mathcal{W}=\left\{w_{\ell}(\boldsymbol{x})\right\}_{\ell \in \mathcal{L}}$ is given, with weight functions $w_{\ell}$ defined over variables in literal $\ell$. Then, the weight of assignment $(\boldsymbol{x}, \boldsymbol{b})$ is:

$$
w(\boldsymbol{x}, \boldsymbol{b})=\prod_{\ell \in \mathcal{L}} w_{\ell}(\boldsymbol{x})^{\llbracket \boldsymbol{x}, \boldsymbol{b} \mid=\ell]} .
$$

When all variables are Boolean (i.e., $\mathcal{X}=\emptyset$ ), the per-literal weights $w_{\ell}(\boldsymbol{x})$ are constants and we retrieve the original definition of the well-known weighted model counting (WMC) task [CD08] as a special case of WMI. In this section, we assume that all per-literal weights are from some certain weight function family, and for literals not in the set $\mathcal{L}$, their weights are the constant


Figure 5.2: Feasible region of SMT theory $\Gamma_{i}$ from Example 2
function one. This setting is expressive enough to approximate many continuous distributions [BPB15].

Example 3 (WMI formulation of a skill matching system). Consider the team performance SMT model $\Gamma$ in Example 1. Assume that a set of per-literal weights $w_{\ell_{i}}\left(X_{\mathcal{T}}, X_{i}\right)=0.1 \cdot\left(X_{\mathcal{T}}+X_{i}-\right.$ $6)^{2}$ is associated to literals $\ell_{i}=X_{\mathcal{T}}-X_{i}<1$, quantifying how likely the team performance is upper bounded by player performances. Then the WMI of the formula $\Gamma$ with two players is $\operatorname{WMI}(\Gamma, w ; \mathcal{X}, \mathbf{B}) \approx 170.69$.

Example 4 (WMI formulation of a house price model). Consider a formula $(b \vee \neg b) \wedge \Gamma_{i}$ where $b$ is a Boolean variable and $\Gamma_{i}$ is as defined in Example 2. Consider the set of literals $\mathcal{L}=\{b,(0<$ price $\left.\left._{i}<3000\right)\right\}$ and per-literal weight functions $\mathcal{P}=\left\{p_{b}, p_{\left(0<\text { price }_{i}<3000\right)}\right\}$, with $p_{b}(\boldsymbol{x})=1.5$ and $p_{\left(0<\text { price }_{i}<3000\right)}(\boldsymbol{x})=$ price $_{i}^{2}$. Then, in worlds with both literals in $\mathcal{L}$ satisfied, our weight function is

$$
p_{b}\left(\text { price }_{i}, \text { sqft }_{i}\right) \cdot p_{\left(0<\text { price }_{i}<3000\right)}\left(\text { price }_{i}, \text { sqft }_{i}\right)=1.5 \cdot \text { price }_{i}^{2} .
$$

In worlds where $b$ is false and only $\left(0<\right.$ price $\left._{i}<3000\right)$ is satisfied, the weight function is price ${ }_{i}^{2}$.

Intuitively, $\operatorname{WMI}(\Delta, w ; \mathcal{X}, \mathbf{B})$ equals the partition function of the unnormalized probability distribution induced by weights $w$ on formula $\Delta$. In the following, we will adopt the shorthand $\mathrm{WMI}(\Delta, w)$ for computing the WMI with all the variables in $\Delta$ in scope. The set of weight functions $w$ together act as an unormalized probability density while the formula $\Delta$ represents logical constraints defining its structure. Therefore, it is possible to compute the (now normalized) probability of any logical query $\Phi$ expressible as an SMT formula involving complex constraints:

$$
\operatorname{Pr}_{\Delta}(\Phi)=\operatorname{WMI}(\Delta \wedge \Phi, w) / \operatorname{WMI}(\Delta, w)
$$

Example 5 (WMI inference for skill rating). Suppose we want to quantify the squad effect in a $2 v 2$ game. Specifically, given two teams $\mathcal{T}_{1}$ and $\mathcal{T}_{2}$ whose players have the same performance, but team $\mathcal{T}_{1}$ is a squad while $\mathcal{T}_{2}$ is not, that is, $\Phi_{c}=\left(B_{1}=\right.$ true $\wedge B_{2}=$ false $)$. We wonder what is the probability of query $\Phi=X_{\mathcal{T}_{1}}>X_{\mathcal{T}_{2}}$, that is team $\mathcal{T}_{1}$ wins and $\mathcal{T}_{2}$ loses. The probability of query $\Phi$ can be computed by two WMI tasks as follows.

$$
\operatorname{Pr}_{\Delta}\left(\Phi \mid \Phi_{c}\right)=\frac{\mathrm{WMI}\left(\Delta \wedge \Phi_{c} \wedge \Phi, w\right)}{\operatorname{WMI}\left(\Delta \wedge \Phi_{c}, w\right)}=\frac{4,206}{7,225} \approx 58.22 \%
$$

with the SMT formula $\Delta:=\Gamma_{1} \wedge \Gamma_{2}$ where the two sub-formulas $\Gamma_{1}$ and $\Gamma_{2}$ model the two teams as in Example 1.
W.l.o.g, from here on we will focus on WMI problems on continuous variables only. We can safely do this since a WMI problem defined on continuous and Boolean variables of the form $\operatorname{WMI}(\Delta, w ; \mathcal{X}, \mathbf{B})$ can always be reduced in polytime to a new WMI problem $\operatorname{WMI}\left(\Delta^{\prime}, w^{\prime} ; \mathcal{X}^{\prime}\right)$ on continuous variables only, by properly introducing auxiliary variables in $\mathcal{X}^{\prime}$ to account for Boolean variables $\mathbf{B}$ without increasing the problem size [ZB19].

### 5.3 Related Work

WMI generalizes weighted model counting (WMC) [SBK05] to hybrid domains [BPB15]. WMC is one of the state-of-the-art approaches for inference in many discrete probabilistic models. Existing exact WMI solvers for arbitrarily structured problems include DPLL-based search with numerical [BPB15, MPS17, MPS19] or symbolic integration [SOG16] and compilation-based algorithms [KMS18, ZDD19].

Motivated by their success in WMC, [BBP16] present a caching scheme for WMI that allows reusing computations at the cost of not supporting algebraic constraints between variables. Different from usual, [MAD17] adopt Gaussian distributions, while [ZDD19] fixed univariate parametric assumptions for weight functions. Many recent efforts in WMI converged in the pywmi [KMZ19] python framework.

Among the exact WMI solvers, the majority ignores the problem structure to be as generalpurpose as possible [BPB15, MPS17, MPS19, KMS18]. However, by doing so they are unable to scale beyond tens of variables in practice. Conversely, our recent efficient alternatives such as SMI [ZB19] and MP-WMI [ZMY20a] can greatly scale but only on WMI problems amenable to tractable inference. We further leverage the strengths of the latter to efficiently solve iterative integration problems in building approximate WMI solvers.

So far, most approximate WMI solvers rely on sampling, and as such inherit all the classical issues of Monte Carlo approaches like poor scalability and convergence [CC96]. Among these, SAMPO [ZDD19] employs Gibbs sampling but does not support generic polynomial weights. A very recent alternative is a fully polynomial randomized approximation scheme [ACD20]. However, it can only operate on DNF SMT formulas, and it is not applicable to our CNF representation as a conversion into DNF can blow up the problem size. Other MCMC variants [ASA15, AD15, ASW16] operating with algebraic constraints, while more effective, cannot be readily used for WMI inference problems. The only alternative to sampling schemes is the hashing-based WMI algorithm [BBP15a] which is known to perform poorly on non-trivial problems due the hardness of calibrating the tilt [CFM14].

Research on learning WMI distributions from data is at its early stages. Parameter learning for piecewise constant densities has been addressed in [BPB15]. Recently, an approach for jointly learning the structure and parameters of a WMI problem has been proposed in [MKT20]. Developing faster inference algorithms is thus beneficial in learning scenarios as, typically, learning a full model requires numerous calls to an inference procedure. WMI inference is closely related to probabilistic program inference, where complex arithmetic and logical constraints are induced by the program structure or its abstraction [HMB17, HBM18].

## CHAPTER 6

## Exact Inference Over Constraints

Despite the appealing features of WMI, current state-of-the-art WMI solvers are far from being applicable to high-dimensional real-world scenarios. This is due to the fact that most solvers ignore the dependency structure of the problem, here expressible through the notion of a primal or factor graph of an SMT formula [DM07]. Thus, their practical utility is limited by their inability to scale up the WMI inference. This chapter presents two proposed exact WMI solvers that are capable of leveraging the dependency structure and thus achieve superior inference performance than the existing solvers.

### 6.1 Search-based WMI Solver

As a first approach to leverage structure, we propose a search-based inference procedure for exact model integration that leverages decomposition to speed up inference. We demonstrate how local structure encoded in SMT theories gives rise to context-specific decomposition during search, reducing the number of models to be generated and integrated over. The integration problem is decomposed into sub-problems by instantiating shared variables and recursing independently on the resulting simplified SMT theories. We show how to choose finitely many values to instantiate continuous variables with, and subsequently do polynomial interpolation to recover exact answers to WMI problems. Our complexity analysis proves the first tractability result for a nontrivial class of WMI problems. Moreover, our experimental evaluation shows that our solver is drastically faster than existing solvers on WMI problems with sparse, tree-shaped primal graphs.


Figure 6.1: WMI runtime on independent model in Example 6.

### 6.1.1 Structure in WMI Problems

This section shows how to reduce WMI to model integration (MI) problems whose structural independence properties can be captured by graph abstractions.

### 6.1.1.1 Independence

We begin by motivating why we want to exploit independence structure during probabilistic reasoning.

Example 6. Consider $n$ houses, and conjoin the theory $\Gamma_{i}$ from Example $2 n$ times, once for each house, into a larger SMT theory $\Gamma=\wedge_{i=1}^{n} \Gamma_{i}$. The $n$ houses are independent since no formula in $\Gamma$ connects properties of different houses. Thus, the WMI of $\Gamma$ can be computed by multiplying the WMI of each individual theory $\Gamma_{i}$.

Figure 6.1 takes a trivial weight function and compares existing WMI solvers on this simple problem. None is able to exploit the extreme independence structure in $\Gamma$. Our proposed method SMI, however, runs in linear time, as expected by the trivial factorization.

This explosion in runtime is due to the fact that existing solvers ignore independence between variables in the $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ theory. However, in discrete graphical models and WMC, leveraging independence to decompose problems is at the core of all exact inference methods, and search-based algorithms in particular [Dar09, DM07]. Specifically, exact discrete inference methods create independence even when it is not immediately present, by performing a case analysis on selected discrete variables, instantiating them to all values, and simplifying the model.

(a) Primal graph

(b) Discrete And/Or Search Tree

Figure 6.2: Primal graph and search tree for $\left(y \vee x_{1}\right) \wedge\left(y \vee x_{2}\right)$.

Through this process, search-based inference algorithms induce and exploit context-specific independence [BFG96]. The decompositions afforded by (conditional and context-specific) independence vastly reduce the computational cost of inference. Example 6 illustrated that this intuition carries over to WMI problems.

In what follows, we first describe the graph abstraction of SMT theories that characterizes dependencies between variables. These form the basis of our algorithm. Second, we show how WMI in hybrid domains can be reduced to unweighted MI in real domains. Hence, the solver we develop in this section will target MI problems.

### 6.1.1.2 Graph Abstarction of SMT

Primal graphs are often used to characterize variable dependencies. For the example Boolean CNF formula $\theta_{B}=\left(y \vee x_{1}\right) \wedge\left(y \vee x_{2}\right)$ the primal graph is shown in Figure 6.2a. Its edges encode that variable pairs $\left(y, x_{1}\right)$ and ( $y, x_{2}$ ) appear in the same clause, while ( $x_{1}, x_{2}$ ) never appear together, and are thus independent given $y$. Similarly, we will use primal graphs for SMT theories to capture variable dependency information as follows.

Definition 3. (Primal graph of SMT) The primal graph of an $S M T(\mathcal{L R} \mathcal{A})$ CNF is an undirected graph whose vertices are all variables and whose edges connect any two variables that appear in the same clause.

(a) Primal graph of theory $\theta_{n}$

(b) Feasible region for $\theta_{2}$

Figure 6.3: Primal graph and feasible region from Example 7.

Example 7. Consider the following theory $\theta_{n}$.

$$
\theta_{n}=\left\{\begin{array}{l}
(-1 \leq y \leq 1) \wedge\left(-0.5 \leq x_{1}, \cdots, x_{n} \leq 0.5\right) \\
\left(x_{i}+1 \leq y\right) \vee\left(y \leq x_{i}-1\right), \text { for all } i \in[n]
\end{array}\right.
$$

Figure 6.3 shows its primal graph and solution space.

While there are many flavors of search-based exact inference, including recursive conditioning [Dar01], DPLL model counting [SBK05], knowledge compilation [CD08], and SumProd algorithms [BDP09], we use the And/Or-search framework to illustrate the required concepts [Nil82, DM07].

The And/Or search algorithm for WMC problems recursively simplifies a discrete counting problem by alternating between two steps. The first (OR) step selects a Boolean variable and tries to instantiate it to both true and false (we will later see how to choose the variable). The second (AND) step finds ways of partitioning the WMC problem into independent sub-problems that can be solved separately. Such sub-problems are introduced by instantiating variables in the OR step in a way that creates independence. The OR step is also called the Shannon expansion. The AND step is also referred to as component caching [SBK05] or detecting decomposability [CD08].

This process is illustrated in Figure 6.2 b for the earlier Boolean CNF $\theta_{B}$. Circles denote ORstep variables whose square-node children are its instantiations. After instatiating $y$, the search tree creates independent problems for $x_{1}$ and $x_{2}$. This independence can be read off directly
from the primal graph in Figure 6.2a. Search-based algorithms (with caching) are known to run efficiently on WMC problems with a tree or tree-like primal graph [Dar09, BDP09].

### 6.1.1.3 WMI to Model Integration Reduction

This section casts hybrid WMI problems into MI problems over only real variables. We consider the case where per-literal weight functions are monomials - functions of the form $\beta x_{1}^{\alpha_{1}} \cdots x_{n}^{\alpha_{n}}$ over real variables $x_{i}$ where $\beta \in \mathbb{R}$ and $\alpha_{i} \in \mathbb{N}$. We further assume that literals in $\mathcal{L}$ also appear in the theory, and that literals and their weights range over the same real variables.

We first show that any WMI problem with Boolean variables can be reduced to a WMI problem without Booleans. Then we show that WMI problems with per-literal weights can be reduced to an unweighted MI problem where the weight function is 1 .

Proposition 15. For each problem $\operatorname{WMI}(\theta, w \mid \boldsymbol{x}, \boldsymbol{b})$ there exists an equivalent problem $\operatorname{WMI}\left(\theta^{\prime}, w^{\prime} \mid\right.$ $\boldsymbol{x}^{\prime}$ ) without Boolean variables $\boldsymbol{b}$ such that

$$
\mathrm{WMI}(\theta, w \mid \boldsymbol{x}, \boldsymbol{b})=\operatorname{WMI}\left(\theta^{\prime}, w^{\prime} \mid \boldsymbol{x}^{\prime}\right)
$$

and the primal graphs of $\theta$ and $\theta^{\prime}$ are isomorphic.

This reduction encodes Boolean variables using fresh real variables and replaces each Boolean atom and its negation by two exclusive $\mathcal{L} \mathcal{R} \mathcal{A}$ atoms over those real variables. Proposition 15 allows us to focus on WMI problems without Boolean variables involved. Certain weight functions can also be reduced, as we show next.

Proposition 16. For each problem $\operatorname{WMI}(\theta, w \mid \boldsymbol{x})$ with per-literal weights $w$ as defined in this section, there exists an equivalent unweighted problem $\operatorname{MI}\left(\theta^{\prime} \mid \boldsymbol{x}^{\prime}\right)$ s.t.

$$
\operatorname{WMI}(\theta, w \mid \boldsymbol{x})=\operatorname{MI}\left(\theta^{\prime} \mid \boldsymbol{x}^{\prime}\right)
$$

Moreover, when weights $w$ are defined over univariate literals, theories $\theta$ and $\theta^{\prime}$ have identical primal graph treewidth [RS86].

This reduction encodes weights using auxiliary parameter variables. For each literal over which a weight is defined, two set of clauses will be appended such that if the literal holds, the MI over the auxiliary variables equals the monomial weight function; otherwise, it equals 1.

Crucially, both reductions can be constructed in polynomial time. Similar efficient reductions exist for arbitrary polynomial weight functions, but can slightly increase treewidth.

Example 8. Consider $S M T(\mathcal{L R} \mathcal{A})$ theory $(b \vee \neg b) \wedge \Gamma_{i}$ with literal set $\mathcal{L}$ and per-literal weight functions $\mathcal{P}$ as defined in Example 4. There exists an equivalent MI problem $\operatorname{MI}\left(\Delta \mid \boldsymbol{x} \cup\left\{\lambda_{b}, z_{b}, z_{i}^{(1)}, z_{i}^{(2)}\right\}\right)$ with a weight function of 1 and without Boolean variables. Its $S M T(\mathcal{L R A})$ theory $\Delta$ is shown below. Note that its primal graph remains a tree.

$$
\Delta=\left\{\begin{array}{l}
\Gamma_{i} \wedge\left(-1<\lambda_{b}<1\right) \wedge_{j=1,2}\left(0<z_{i}^{(j)}<\text { price }_{i}\right) \\
\quad \lambda_{b}>0 \Rightarrow\left(0<z_{b}<1.5\right) \\
\neg\left(\lambda_{b}>0\right) \Rightarrow\left(0<z_{b}<1\right) .
\end{array}\right.
$$

### 6.1.2 Method

The goal of our work is to take advantage of the independence structure in $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ theories to reduce the computational cost of model integration. Our solution is to exploit context-specific independence by search.

One obstacle is that to introduce independence in discrete search, we instantiate a variable with all values in its domain. Unfortunately, when the variable has a real domain (e.g., $y \in[0,1]$ ), we cannot instantiate it with every value in its domain, since there are uncountably many (see Figure 6.4a). This basic limitation has precluded the use of search-based inference in continuous graphical models.

We overcome this problem by observing that MI is an integration over a piecewise polynomial, which can be fully recovered from a finite number of points. Specifically, for real variable $y$ in theory $\theta$, if we instantiate the variable $y$ with a value $\alpha$, then the MI of theory $\theta \wedge(y=\alpha)$ is the density of $\operatorname{WMI}(\theta, w)$ at $y=\alpha$. Recall that a polynomial function $p(y)$ with degree $d$ defined over


Figure 6.4: Continuous search trees for $\theta_{2}$ from Example 7.
an interval $I$ is uniquely defined by its values at $d+1$ distinct points in $I$, and that a closed-form expression for $p(y)$ can be recovered exactly and efficiently.

Consider again the theory $\Gamma_{i}$ from Example 2. As shown in Figure 5.2, function $f(\alpha)=$ $\operatorname{MI}\left(\Gamma_{i} \wedge\left(s q f t_{i}=\alpha\right)\right)$ is a piecewise polynomial with three intervals. We can recover all three polynomials from a finite number of points, and thus obtain the integration of $f(\alpha)$, that is, the model integration $\operatorname{MI}\left(\Gamma_{i}\right)$. This motivates the search-based model integration algorithm we develop next.

### 6.1.2.1 Variable Instantiation

We first show that when per-literal weight functions $\mathcal{P}$ are polynomials, WMI of theory $\theta$ can be obtained by doing search with finite instantiations on real variables.

Proposition 17. Let $y$ be a real variable in $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ theory $\theta$ with a tree primal graph. If perliteral weight functions $\mathcal{P}$ are polynomials, the WMI is an integration over a univariate piecewise polynomial $p(y)$, that is,

$$
\begin{equation*}
\operatorname{WMI}(\theta, w \mid \boldsymbol{x}, \boldsymbol{b})=\int_{I} p(y) d y \tag{6.1}
\end{equation*}
$$

where piecewise polynomial $p(y)$ is integrated over set $I=\left\{y^{*} \mid \exists \hat{\boldsymbol{x}}^{*}, \exists \boldsymbol{b}^{*}\right.$ s.t. $\theta\left(y^{*}, \hat{\boldsymbol{x}}^{*}, \boldsymbol{b}^{*}\right)$ is $\left.S A T\right\}$ with $\hat{\boldsymbol{x}}$ being the remaining real variables.


Figure 6.5: Piecewise polynomial $p(y)$ as defined in Proposition 17 for theory $\theta_{2}$ from Example 7, whose integration is $\operatorname{MI}\left(\theta_{2}\right)$. The two polynomials $p_{-}(y)$ and $p_{+}(y)$ are unknown, but we can recover them from a finite number of points.

The set $I$ is a union of disjoint supports for piecewise polynomial $p(y)$. We refer to these intervals as "pieces". To describe our MI algorithm, we first assume in this section that these intervals and their polynomial degrees are given. Our method to explicitly find these intervals and degrees will be given in Section 6.1.2.2.

Although Proposition 17 holds for WMI problems with polynomial per-literal weight functions in general, we use the insights from Section 6.1.1.3 to only focus on MI problems. For interval set $I$ defined in Proposition 17, suppose we are given the interval pieces $[l, u] \in I$ and degrees $d$ of their associated polynomials. If we instantiate variable $y$ with $d+1$ distinct values in each piece $[l, u]$ of degree $d$, and solve any sub-problems recursively, we can recover polynomial $p_{l, u}(y)$ defined on interval $[l, u]$ by performing interpolation on $d+1$ points. Finally, MI of the full theory $\theta$ can be computed as follows.

$$
\begin{equation*}
\operatorname{MI}(\theta, w \mid \boldsymbol{x}, \boldsymbol{b})=\sum_{[l, u] \in I} \int_{l}^{u} p_{l, u}(y) d y \tag{6.2}
\end{equation*}
$$

For example, consider theory $\theta_{2}$ from Example 7. We can interpret $\mathrm{MI}\left(\theta_{2}\right)$ as an integration over piecewise polynomial $p(y)$ whose intervals $[-1,-0.5]$ and $[0.5,1]$ both have associated degree two. After instantiating $y$ to three values in each interval, we get two independent sub-MI problems that contain variable $x_{1}$ and variable $x_{2}$ respectively. By solving these sub-problems, we obtain three points fitted by each polynomial $p_{-}(y)$ and $p_{+}(y)$ as shown in Figure 6.5. There-

```
Algorithm 6 SMI: Search-Based Model Integration
Input: \(T\) : pseudo tree, \(\theta: \operatorname{SMT}(\mathcal{L R} \mathcal{A})\) theory
Output: \(p\) : MI of theory \(\theta\)
    if \(T\) is a forest of trees \(T^{\prime}\) then
        \(\theta^{\prime} \leftarrow\) sub-theories containing variables in \(T^{\prime}\)
        return \(\prod_{T^{\prime}} \operatorname{SMI}\left(T^{\prime}, \theta^{\prime}\right)\)
    \(p=0, y=\operatorname{root}(T), S T_{y}=\) set of subtrees below \(y\)
    \(I=\operatorname{PE} \_\operatorname{NODE}(\theta, y)\)
    for all polynomial piece \(\{[l, u], d\} \in I\) do
        select \(d+1\) distinct values \(\alpha_{i}\) 's in \([l, u]\)
        \(p_{i} \leftarrow \operatorname{SMI}\left(S T_{r},\left.\theta\right|_{\left(y=\alpha_{i}\right)}\right)\)
        \(p_{l, u}(y) \leftarrow\) polynomial interpolation on \(\left(\alpha_{i}, p_{i}\right)\) 's
        \(p \leftarrow p+\int_{l}^{u} p_{l, u}(y) d y\)
    return \(p\)
```

fore, we can recover both by polynomial interpolation and can obtain $\operatorname{MI}\left(\theta_{2}\right)$ by Equation 6.2. Figure 6.4b depicts the search space of our algorithm on interval $[0.5,1]$.

The above discussion has shown that for MI problems, we can instantiate a real variable to finitely many values, decompose the problem into independent parts, and then solve the subproblems recursively. Algorithm 6 follows exactly this strategy for search-based model integration. The role of pseudo trees will be explained in Section 6.2.2.4. The remaining problem is how to exactly obtain pieces $[l, u]$ and their associated degrees $d$ in function PE_NODE. We address this problem next.

### 6.1.2.2 Finding Pieces via Critical Points

Recall that by Proposition 17, WMI of $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ theory $\theta$ can be rewritten as $\operatorname{WMI}(\theta, w \mid$ $\boldsymbol{x}, \boldsymbol{b})=\int_{I} p(y) d y$ where $p(y)$ is a piecewise polynomial, set $I$ is a union of disjoint support of
polynomials in $p(y)$, and each piece $[l, u] \in I$ is associated with a polynomial degree $d$. We hope that when a real variable $y$ in theory $\theta$ is chosen to be instantiated, we can exactly find all pieces and their associated degrees for piecewise polynomial $p(y)$.

It turns out that this can be achieved. While integrating over satisfying assignments with respect to a certain variable given an $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ theory, integration upper bounds and lower bonds are defined by its literals. Changes in integration bounds give rise to different pieces of integration and therefore result in the piecewise nature of the polynomial in Proposition 17. In our method we determine these pieces by collecting points where certain bounds meet. Further, by propagating polynomial piece and degree information in a bottom-up manner along the primal graph, we can obtain the pieces and degree for the chosen piecewise polynomial.

We will first describe our method in a basic case where there are only two real variables in the theory. Then we extend this approach to theories with tree primal graphs.

### 6.1.2.3 Base Case: Pieces of Two Real Variables

First we investigate a simple case where there are only two real variables $x$ and $y$ in $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ theory $\theta$. Recall that we are solving an unweighted MI problem. We would like to find pieces and associated degrees for variable $y$ such that we can instantiate $y$ as in Section 6.1.2.1:

$$
p(y)=\int_{\theta(x, y)} 1 d x=\sum_{[l(y), u(y)] \in I(y)} \int_{l(y)}^{u(y)} 1 d x=\sum_{[l(y), u(y)] \in I(y)} u(y)-l(y)
$$

where set $I(y)$ is defined as

$$
\begin{equation*}
\{[l(y), u(y)] \mid \forall x \in[l(y), u(y)], \theta(x, y) \text { is SAT }\} . \tag{6.3}
\end{equation*}
$$

That is, for any fixed value $y^{*}$, the set $I\left(y^{*}\right)$ consists of intervals of consistent values for variable $x$. For any $[l(y), u(y)] \in I(y)$, it gives a pair of integration bounds for variable $x$. Further by integrating over $x$ we can obtain a polynomial with respect to variable $y$.

Each piece $[l, u]$ corresponds to a certain class of values that gives the same symbolic integration bounds to variable $x$. The two values $y=l$ and $y=u$ are endpoints of the piece only if
integration bound set $I(y)$ changes at these points, since the piecewise polynomial $p(y)$ is defined by these bounds. That is, for arbitrarily small $\epsilon$, we have $I(l-\epsilon) \neq I(l+\epsilon)$, and it also holds at point $y=u$. We formally define critical points below.

Definition 4. (Critical Point) Let $\theta$ be an $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ theory with two real variables, and denote one of the real variables by $y$. Let $I(y)$ be an integration bound set as defined in Equation 6.3. Then $y=\alpha$ is a critical point if for arbitrarily small $\epsilon$, it holds that $I(\alpha-\epsilon) \neq I(\alpha+\epsilon)$.

Remark. The comparison of set $I(y)$ is done symbolically. That is, for two distinct values $\alpha, \beta$, we say $I(\alpha)=I(\beta)$ if they have the same set of symbolic integration bounds. For example, if at $y=\alpha, I(y)=\{[1, y]\}$ and at $y=\beta \neq \alpha, I(x)=\{[1, y]\}$, it holds that $I(\alpha)=I(\beta)$. However, if at $y=\alpha, I(y)=\{[1, y]\}$ and at $y=\beta, I(y)=\{[y, 2]\}$, then we say $I(\alpha) \neq I(\beta)$.

Our idea is that, if we can find all critical points $y=\alpha$ where the set $I(y)$ changes, then we can partition real domains of $y$ into disjoint intervals, such that any support of piecewise polynomial $p(y)$ is either one of these intervals or a union of some intervals. For the resulting interval $[l, u]$, we can apply an $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ solver to $\theta^{\prime}=\theta \wedge(l<y<u)$ to check whether it is a satisfiable piece of function $p(y)$; if this is true, we can obtain the polynomial degree of $p_{l, u}(y)$ defined over this piece by simply traversing theory $\theta^{\prime}$.

### 6.1.2.4 General Case: Pieces of Tree Structures

Given an $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ theory $\theta$ with a tree-shape primal graph $G$, our goal is to enumerate pieces and their associated degrees for the root variable $y$, building on the algorithm we developed in the base case above. This can be done in a bottom-up manner with tree primal graphs.

Specifically, we first partition theory $\theta$ into sub-theories $\theta_{r, c}$ and $\theta_{G_{c}}$ for each $c$, such that $\theta=\bigwedge_{c}\left(\theta_{r, c} \wedge \theta_{G_{c}}\right)$, where variables $c$ are the child variables of root $r$, and graph $G_{c}$ is the subtree rooted at variable $c$. Each theory $\theta_{r, c}$ contains only variables $r$ and $c$, on which we can apply the enumeration for the base case above, and each theory $\theta_{G_{c}}$ contains only variables in sub-tree $G_{c}$. This is possible provided that the primal graph of theory $\theta$ has a tree structure, which is why our algorithm is restricted to $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ theories with tree-shaped primal graphs.

For each child variable $c$, we first obtain its pieces with respect to theory $\theta_{G_{c}}$ in a recursive way. Then we can apply our enumeration algorithm for two-variable theory PE_EDGE to theory $\theta_{r, c}$ with the given pieces of variable $c$. What we would get are sets of pieces for each child variable $c$. To be consistent with theory $\theta$, we need to take intersections of these sets which we refer to as the shattering operation. Finally, the resulting intersections are pieces and polynomial degrees for root variable $r$.

As described above, our piece enumeration algorithm is applicable to MI problems for theories with tree primal graphs. Moreover, it is also applicable to WMI problems whose SMT theory has a tree primal graph and whose per-literal weights are monomials over univariate literals as described in Section 6.1.1.3, since our reduction process can preserve the tree structure of the primal graph.

### 6.1.2.5 Complexity Analysis

Inference over networks involving real variables raises considerable challenges for inference, and network structures that are tractable in the discrete case, such as polytrees, give rise to NP-hard inference problems in the hybrid case [KF09]. We show that the complexity of our algorithm is mainly exponentially bounded by the tree height of the primal graph.

Our search algorithm for MI needs to choose which variables to instantiate first. This choice can be based on a tree data structure that orders the variables. Such a tree characterizes the computational complexity as it does for discrete And/Or search algorithms. We first formally defined the tree that helps guide our search.

Definition 5. (Pseudo Tree) Given an undirected graph $G$ with vertices and edges ( $V, E_{G}$ ), a pseudo tree for $G$ is a directed rooted tree $T$ with vertices and edges $\left(V, E_{T}\right)$, such that any edge $e$ that is in $G$ but not in $T$ must connect a vertex in $T$ to one of its ancestors.

That is, edge $e=\left(v_{1}, v_{2}\right)$ such that $e \in E_{G}$ and $e \notin E_{T}$ implies that either vertex $v_{1}$ is an ancestor of vertex $v_{2}$ in $T$ or vertex $v_{2}$ is an ancestor of vertex $v_{1}$ in $T$. Note that the pseudo tree

(a) Runtime on star graphs.

(b) Runtime on 3-ary tree graphs.

(e) 3-ary tree with nodes $n=12$.
(x)

(c) Runtime on path graphs.

(d) Star graph with nodes $n=8$.
(f) Path graph with nodes $n=5$.

Figure 6.6: (a)-(c) MI execution time on $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ with tree primal graphs. (d)-(f) Example tree primal graphs.
has the same set of vertices as $G$. Such a pseudo tree guides SMI (Algorithm 6) in deciding which variable to instantiate, and when to decompose.

Next, we analyze the complexity of SMI. Since our algorithm performs search, its time and space complexity is characterized by the size of its search space. Our analysis does not take caching improvements into consideration.

Theorem 4. (Size of Search Space) Consider an $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ theory $\theta$ with a tree-shaped primal graph with height $h_{p}$, and a pseudo tree $T$ with $l$ leaves and height $h_{t}$. Let $m$ be the number of $\mathcal{L R} \mathcal{A}$ literals in $\theta$, and $n$ be the number of real variables. Then the size of the SMI search space is $\mathcal{O}\left(l \cdot\left(n^{3} \cdot m^{h_{p}}\right)^{h_{t}}\right)$.

Hence, we can conclude that the complexity of our algorithm is bounded exponentially by tree heights of both the primal graph and pseudo tree. In fact, for any tree-shaped primal graph, we can always choose a pseudo tree whose height $h_{t}$ is $\mathcal{O}(\log n)$ to guide the search [DM07]. Moreover, the number of leaves $l$ in pseudo tree $T$ is no larger than the number of nodes $n$. Thus, we have the following corollary.

Corollary 2. Following the notation in Theorem 4, with properly chosen pseudo tree $T$ whose tree height $h_{t}$ is $\mathcal{O}(\log n)$, the size of the search space generated by SMI is $\mathcal{O}\left(n^{1+3 \log n+h_{p} \log m}\right)$.

Therefore, the complexity of our algorithm is mainly decided by tree heights of primal graphs $h_{p}$. In the worst case when tree primal graphs have height $\mathcal{O}(n)$ - for instance path graphs, whose tree height is $n$ when rooted at the start node - then the worst-case complexity of our algorithm is $\mathcal{O}\left(n^{n \log m}\right)$ by Corollary 2 . That is, the time complexity is worst-case super-exponential.

In cases when the tree primal graph has tree height of size $\mathcal{O}(\log n)$, the complexity of our algorithm is $\mathcal{O}\left(n^{1+(3+\log m) \log n}\right)$ which is of quasi-polynomial complexity, and considered to be efficient. Trees with tree height in $\mathcal{O}(\log n)$ are a general class of trees used in various models. Balancing trees like AVL trees and full k-ary trees are of tree height $\mathcal{O}(\log n)$. Another example is a star graph, which has one internal node and all other nodes as leaves. This graph corresponds to the well-known naive Bayes structure for directed graphical models. It is the primal graph of a theory modeling independent variables predicting one and the same dependent (class) variable. The tree height of star graphs is constant 1 when choosing the internal node as root. Hence, our algorithm runs efficiently on such WMI problems.

### 6.1.3 Empirical Evaluation

We analyze the performance of our search-based MI algorithm on $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ theories with tree primal graphs. First, we show that our algorithm is efficient for theories whose primal graphs have constant tree heights, or tree heights of $\log$ scale w.r.t. the number of real variables $n$. For theories whose primal graph has tree height in $\mathcal{O}(n)$ - the cases where our algorithm has superexponential worst-case complexity in theory - empirical results show that our algorithm still runs efficiently. We also consider a more complex house price model where house sizes are dependent, as opposed to those in Example 6. Moreover, the house price model has non-trivial weight functions that our algorithm first reduces to a MI problem as outlined in Section 6.1.1.3. We compare our algorithm to alternative WMI solvers and conclude that it significantly outperforms existing solvers on these benchmarks.

Benchmarks We compare our algorithm (SMI) with other WMI solvers. The block-clause-strategy-based solver (BC) [BPB15] iteratively generates new models by adding the negation of the latest model to the formula for the following iteration. The all-satisfying-assignments-based solver (ALLSMT) [BBP16] first generates the set of all $\mathcal{L R} \mathcal{A}$-satisfiable total truth assignments on atoms that propositionally satisfy the theory. The implementation of [SOG16] (PRAiSE) is a variable-elimination-based solver. The predicate-abstraction-based solver (PA) [MPS17] exploits the power of SMT-based predicate abstraction to reduce the number of models to be integrated over. Both the extended algebraic-decision-diagram-based solver (XADD) [KMS18] and sentential-decision-diagram-based solver (Symbo) [ZDD19] use circuit-based compilation languages and exploits the circuit structures.

### 6.1.3.1 Tree Primal Graphs

We investigate the performance of our algorithm on $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ theories with three types of tree primal graphs: 1) star graphs, consisting of one center node connected to all other nodes, and no other connections; 2) full 3-ary trees, whose non-leaf vertices have exactly three children and all levels are full except for some rightmost position of the bottom level; 3) path graphs, consisting of linearly connected nodes. These structural constraints arise naturally in data and many probabilistic graphical modeling problems.

For each graph type, given a number of nodes $n$, we introduce $n$ real variables which we denote by $\boldsymbol{x}=\left\{x_{0}, x_{1}, \cdots, x_{n-1}\right\}$ with bounded domains $\forall i,\left(-1 \leq x_{i} \leq 1\right)$. Denote the graph by $G=(V, E)$ where $V=\{0,1, \cdots, n-1\}$ is the vertex set and $E=\{(i, j), i, j \in V\}$ the edge set. We perform MI for the following theories and increasing $n$.

$$
\theta(\boldsymbol{x})=\left\{\begin{array}{l}
\bigwedge_{i \in V}\left(-1 \leq x_{i} \leq 1\right) \\
\bigwedge_{(i, j) \in E}\left(\left(x_{i}+1 \leq x_{j}\right) \vee\left(x_{j} \leq x_{i}-1\right)\right)
\end{array}\right.
$$

Figure 6.6 shows example primal graphs and the execution time of experiments comparing SMI with baselines.


Figure 6.7: Runtime and primal graph for house price model.

For MI over theories with all three types of tree primal graphs, our algorithm significantly outperform other WMI solvers in terms of execution time. The runtime curves of other solvers grow seemingly exponentially while our curve grows slowly with the number of real variables. For theories with star graphs and full 3-ary trees as primal graphs, the time curves of SMI are consistent with the complexity analysis in Section 6.2.2.4 stating that our algorithm has quasipolynomial complexity. For theories with path graphs as primal graphs, which are still sparse graphs, we perform caching and the runtime curve grows slowly, even though our worst-case analysis allows for a super-exponential time complexity.

### 6.1.3.2 House Price $\operatorname{SMT}(\mathcal{L} \mathcal{R} \mathcal{A})$ Model

In Example 6 we performed MI for multiple houses based on extreme independence assumptions. Now we consider a more complicated case where houses are not independent and there are Boolean variables in the $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ model. Moreover, we choose non-trivial per-literal weight functions in order to evaluate our algorithm for reducing WMI to unweighted MI problems.

Specifically, we consider $n$ houses that are located along a street. Each house $i$ has its price and square footage model as in Example 2. Also, we enforce the constraint that square footage between two neighboring houses should not vary too much and we use a Boolean variable $b$ to indicate whether or not these houses are located in an urban area. This gives the following SMT
theory.

$$
\Gamma_{\text {street }}=\left\{\begin{array}{l}
(b \vee \neg b) \wedge \bigwedge_{i=1}^{n} \Gamma_{i} \\
\bigwedge_{i=1}^{n-1}\left(s q f t_{i} \leq s q f t_{i+1}+o f f s e t\right)
\end{array}\right.
$$

with offset a constant characterizing maximum difference in square footage between two neighboring houses. For weights $w$, consider the set of literals $\mathcal{L}=\{b\} \cup\left\{0<\right.$ price $_{i}<3000, i=$ $1, \cdots, n\}$ and per-literal weight functions $\mathcal{P}=\left\{p_{b}\right\} \cup\left\{p_{\left(0<\text { price }_{i}<3000\right)}, i=1, \cdots, n\right\}$, with $p_{b}(\boldsymbol{x})=1.5$ and $p_{\left(0<\text { price }_{i}<3000\right)}(\boldsymbol{x})=$ price $_{i}^{2}$ for all $i$. Then, in worlds where all literals in $\mathcal{L}$ are satisfied, our weight function is $1.5 \prod_{i=1}^{n}$ price $e_{i}^{2}$. In worlds where $b$ is false but other literals are satisfied, the weight function is $\prod_{i=1}^{n}$ price $_{i}^{2}$. Figure 6.7 shows an example primal graph and WMI runtime for this house price model.

### 6.2 Message-Passing WMI Solver

While SMI directly exploits the problem structure encoded in primal graphs, in order to perform a tractable reduction, SMI is limited to a restricted set of weights, and hence a very narrow set of WMI problems. To build more powerful exact WMI solvers, we make the following contributions. First, we theoretically trace the boundaries for the classes of tractable WMI problems known in the literature. Second, we expand these boundaries by devising a polytime algorithm for exact WMI inference on a class that is strictly larger than the class previously known to be tractable. Our proposed WMI solver, called MP-WMI, adopts a novel message-passing scheme for WMI problems. It is able to exactly compute all the variable marginal densities at once. By doing so, we are able to scale inference beyond the capabilities of all current exact WMI solvers. Moreover, we can amortize inference inter-queries for rich SMT queries that conform to the problem structure.

### 6.2.1 Tractability Analysis

The major efforts in advancing WMI inference have been so far concentrated on devising sophisticated WMI solvers to deliver exact inference routines for general scenarios without investigating the effect of the structure of a WMI problems on its complexity. Little to no attention has gone
to formally understand which classes of WMI problems can be guaranteed to be solved exactly and in polynomial time, that is, tractably.

One notable exception can be found in [ZB19] where the search-based MI (SMI) solver is introduced. WMI problems for which SMI guarantees polytime exact inference constitute the first class of tractable WMI. Intuitively, SMI solves MI problems by using search to leverage the conditional independence among variables.

As in [ZB19] we characterize the structure of an SMT formula via its primal graph.

Definition 6. (Primal graph of SMT) The primal graph of an SMT formula $\Delta$ is an undirected graph $\mathcal{G}_{\Delta}$ whose vertices are variables in formula $\Delta$ and whose edges connect any two variables that appear in a same clause in the formula $\Delta$.

An example primal graph of the SMT formula in Example 1 is shown in Figure 5.1. The SMI solver guarantees polynomial time execution for the class of MI problems with certain treeshaped primal graphs, which we denote as treeMI.

Definition 7. (treeMI Problem Class) Let treeMI be the set of all MI problems over real variables whose SMT formula $\Delta$ induces a primal graph $\mathcal{G}_{\Delta}$ with treewidth one and with bounded diameter d. Problems in treeMI can be solved in polytime via SMI [ZB19].

Note that in Definition 7 the primal graph diameter here plays the role of a constant since, otherwise, SMI complexity can be worst-case exponential in diameter $d$. In the following we will try to answer if larger classes than tree MI are still amenable to tractable inference. We start by demonstrating a novel result that states the hardness of a larger class of MI problems, still focusing on dependencies between two variables, but allowing for non-tree-shaped primal graphs.

Definition 8. (2MI Problem Class) Let 2MI be the set of all MI problems over real variables whose SMT formula $\Delta$ is a conjunction of clauses comprising at most two variables.

Note that a clause comprising at most two variables can be a conjunction of arbitrarily many literals. Moreover, when there are more than two variables in a clause, in the primal graph there


Figure 6.8: The current landscape of classes of WMI problems. We enlarge the boundaries of tractable WMI inference from tree MI to treeWMI and prove the hardness of 2 MI and 2 WMI .
must be a loop and thus the treewidth of the primal graph is larger than one. Hence all MI problems with tree-shaped primal graph must be in the class 2MI.

Theorem 5. (Hardness of 2MI) Given an MI problem in 2 MI with an SMT formula $\Delta$, computing $\operatorname{MI}(\Delta)$ is \#P-hard.

Sketch of Proof. The proof is done by reducing the \#P-complete problem \#2SAT to an MI problem in 2 MI with an SMT formula $\Delta$ such that counting the number of satisfying assignments to the \#2SAT problem equates to the MI of formula $\Delta$.

From Theorem 5 it follows that the problem class $2 \mathrm{WMI}(\Omega)$, i.e., the WMI problems with SMT formulas being a conjunction of clauses comprising at most two variables, and with per-literal weights in weight function family $\Omega$, is also hard since class 2 MI is a sub-class of $2 \mathrm{WMI}(\Omega)$. We revert our attention to WMI problems exhibiting a dependency tree structure. Notice that for WMI problems, the tractability not only depends on the logical structure defined by the SMT formulas, but also the statistical structure defined by weight functions. Next in our analysis, we take into consideration the weight function families. Analogously to what Definition 7 states, we introduce the notion of $\operatorname{treeWMI}(\Omega)$ with the associated weight function family specified as follows.

Definition 9. (treeWMI $(\boldsymbol{\Omega})$ Problem Class) Let $\operatorname{treeWMI}(\boldsymbol{\Omega})$ be the set of all WMI problems over real variables whose SMT formula $\Delta$ induces a primal graph $\mathcal{G}_{\Delta}$ with treewidth one and with bounded diameter $d$, and whose per-literal weights are in a function family $\Omega$.
[ZB19] propose a WMI-to-MI reduction such that some treeWMI $(\boldsymbol{\Omega})$ problems with polynomial weights are reduced in polynomial time to treeMI problems amenable to tractable inference by the SMI solver. Intuitively, the reduction process introduces auxiliary continuous variables and SMT formulas over these variables to encode the polynomial weight functions. We refer the readers to [ZB19] for a detailed description of the reduction. However, as shown next, the set of treeWMI problems that can be reduced to treeMI is rather limited.

Definition 10. ( $\Omega^{\text {SMI }}$ Weight Function Family) Let $\Omega^{\text {SMI }}$ be the family of per-literal weight functions that are monomials associated with either (i) univariate literals or (ii) a literal that appears exclusively in a unit clause, i.e., a clause consisting of a single literal.

Theorem 6. Let $\rho$ be the polytime WMI-to-MI reduction for treeWMI $(\boldsymbol{\Omega})$ problems as defined in $[Z B 19]$. Then the image $\{\rho(\nu) \mid \nu \in \operatorname{treeWMI}(\Omega)\} \subset$ treeMI if-and-only-if $\boldsymbol{\Omega} \subset \boldsymbol{\Omega}^{\text {SMI }}$.

Sketch of Proof. The necessary condition can be proved by the reduction process and the sufficient one can be proved by contradiction.

Therefore, the SMI solver is limited to a rather restricted subset of treeWMI $(\Omega)$ since from the definition of $\Omega^{\text {SMI }}$ we can tell that it is a strict subset of monomial per-literal weights. In order to enlarge the tractable class of WMI problems, next we will define a rich family of weight functions.

Definition 11. (Tractable Weight Conditions) Let $\Omega$ be a family of per-literal weight functions. We say that the tractable weight conditions (TWC) hold for $\Omega$ if we have:
(i) closedness under product: $\forall f, g \in \boldsymbol{\Omega}, f \cdot g \in \boldsymbol{\Omega}$;
(ii) tractable symbolic integration: $\forall f \in \Omega$, the symbolic antiderivative of function $f$ can be tractably computed by symbolic integration;
(iii) closedness under definite integration: $\forall f \in \Omega$ with its antiderivative denoted by $F$, given integration bounds $l(x), u(x)$ in $\mathcal{L R} \mathcal{A}$ with $x \in \mathcal{X}, F(u(x))-F(l(x)) \in \boldsymbol{\Omega}$.

Some example weight function families that satisfy TWC include the polynomial family, exponentiated linear function family and the function family resulting from their product. Moreover note that piecewise function families, when pieces belong to the above families, also satisfy TWC. It turns out that the weight function families that satisfy TWC subsume and extend all the parametric weight functions adopted in the WMI literature so far. The following proposition is a direct result from the fact that the piecewise polynomial weight family $\Omega^{p}$ is a strict superset of the family $\Omega^{\text {SMI }}$.

Proposition 18. Let $\Omega^{P}$ be the piecewise polynomial weight function family. The WMI problem class treeWMI $\left(\boldsymbol{\Omega}^{\mathrm{P}}\right)$ is a strict superset of problem class treeWMI $\left(\boldsymbol{\Omega}^{\mathrm{SMI}}\right)$.

Theorem 7. If a weight function family $\boldsymbol{\Omega}$ satisfies TWC as in Definition 11, WMI problems in class $\operatorname{treeWMI}(\boldsymbol{\Omega})$ are tractable, i.e., they can be solved in polynomial time.

The proof to the above theorem is provided in the next two sections by construction where in Section 6.2.2 we proposed our WMI solver, called MP-WMI, operating on WMI problems in treeWMI $(\boldsymbol{\Omega})$ with its complexity analysis in Section 6.2.2.4. A summary of the WMI problem classes is shown in Figure 6.8.

### 6.2.2 Method

Message passing on tree-structured graphs has achieved remarkable attention in the PGM literature [Pea88, KFL01]. Its classical formulation and efficiency relies on compact factor representations allowing easy computations. However, adapting existing message-passing algorithms to WMI inference is non-trivial. This is due to the fact that inference is computed in a hybrid structured space with logical and arithmetic constraints. We present our message-passing scheme by first deriving a factorized representation of WMI problems.

### 6.2.2.1 Factor Graph Representation of WMI

In the literature of WMC, incidence graphs are proposed to characterize the structure of problems defined by Boolean CNF formulas [SS10]. Incidence graphs are bipartite graphs with clause nodes and variable nodes, where a clause and a variable node are joined by an edge if the variable occurs in the clause. We derive the analogous representation for the more general SMT formulas, which we then turn into a factor graph of WMI problems.

Recall that for the joint distribution represented by a WMI problem, the support is defined by the logical constraints and the unnormalized density is defined by weight functions. In the following, we first factorize the SMT formula $\Delta$ of a WMI problem $\mathrm{WMI}(\Delta, w)$ in the class treeWMI:

$$
\begin{equation*}
\Delta=\bigwedge_{i \in \mathcal{V}} \Delta_{i} \wedge \bigwedge_{i, j \in \mathbb{E}} \Delta_{i j} \tag{6.4}
\end{equation*}
$$

where the set $\mathcal{V}$ is the index set of variables and the set $\mathbb{E}$ is the index pairs of variables in the same clause. Then a WMI problem can be conveniently represented as a bipartite graph, known as factor graph, that includes two sets of nodes: variable nodes $X_{i}$, and factor nodes $f_{\mathcal{S}}$, where $\mathcal{S}$ denotes a factor scope, i.e., the set of indices of the variables appearing in it. A variable node $X_{i}$ is connected to a factor node $f_{\mathcal{S}}$ if and only if $i \in \mathcal{S}$. Specifically, the factors are defined as follows:

$$
\begin{equation*}
f_{\mathcal{S}}\left(\boldsymbol{x}_{\mathcal{S}}\right)=\prod_{\Gamma \in \operatorname{CLS}\left(\Delta_{\mathcal{S}}\right)} \llbracket \boldsymbol{x}_{\mathcal{S}} \vDash \Gamma \rrbracket \prod_{\ell \in L I T S} w_{\ell)} w_{\ell}\left(\boldsymbol{x}_{\mathcal{S}}\right)^{\llbracket \boldsymbol{x}_{\mathcal{S}} \mid \ell \rrbracket} \tag{6.5}
\end{equation*}
$$

where $\boldsymbol{x}_{\mathcal{S}}$ denotes the restriction of $\boldsymbol{x}$ to the variables in factor $f_{\mathcal{S}}$ and analogously $\Delta_{\mathcal{S}}$ is the restriction of formula $\Delta$ to the clauses over the variables in $\mathcal{S}$. Here, the set of clauses in the SMT formula $\Delta$ is denoted by $\operatorname{CLS}(\Delta)$, and the set of literals in a clause $\Gamma$ is denoted by $\operatorname{LITS}(\Gamma)$. Intuitively, the factors include the parameterized densities as in the classic PGM literature, here represented by the per-literal weights, but also the structure enforced by the logical constraints in the SMT formula, via the indicator functions. Figure 6.9 shown an example of a factor graph.

As in every tree-shaped factor graphs, we define an unnormalized joint distribution corre-


Figure 6.9: Factor graph (left) of formula $\Gamma$ with two players and piecewise polynomial messages (right) sent from the three factor nodes to variable node $X_{\mathcal{T}}$ when solving the WMI in Example 3 by MP-WMI.
sponding to the WMI problem in the form of a product of factors as follows.

$$
\begin{equation*}
p(\boldsymbol{x})=\prod_{\mathcal{S}} f_{\mathcal{S}}\left(\boldsymbol{x}_{\mathcal{S}}\right)=\prod_{i \in \mathcal{V}} f_{i}\left(X_{i}\right) \prod_{i, j \in \mathbb{E}} f_{i j}\left(X_{i}, X_{j}\right) \tag{6.6}
\end{equation*}
$$

By construction, it is easy to see that the normalization constant of such a distribution equals computing the corresponding weighted model integral.

Proposition 19. Given a problem $\mathrm{WMI}(\Delta, w)$ in treeWMI, let $p(\boldsymbol{x})$ being the unnormalized joint distribution as defined in Equation 6.6. Then the partition function of distribution $p(\boldsymbol{x})$ is equal to $\operatorname{WMI}(\Delta, w)$.

### 6.2.2.2 Message-Passing Scheme

Deriving a message-passing scheme for WMI poses unique and considerable challenges. First, different from discrete domains, on continuous or hybrid domains one generally does not have universal and compact representations for messages, and logical constraints in WMI make it even harder to derive such representations. Moreover, marginalization over real variables requires integration over polytopes, which is already \#P-hard [DF88]. The integration poses the problem of whether the messages defined are integrable and how hard it is to perform the integration. In the following part, we will present our solutions to these challenges by first describing a general message-passing scheme for WMI and then investigating of which form the messages are, given certain weight families.

Given the factorized representation of WMI in Section 6.2.2.1, our message-passing scheme, called MP-WMI and summarized in Algorithm 7, comprises exchanging messages between nodes in the factor graph. Before the message passing starts, we choose an arbitrary node in the factor graph as root and orient all edges away from the root to define the message sending order. MPWMI operates in two phases: an upward pass and a downward one. First, we send messages up from the leaves to the root (upward pass) such that each node has all information from its children and then we incorporate messages from the root down to the leaves (downward pass) such that each node also has information from its parent. The messages are formulated as follows.

Proposition 20. Both messages $\mathrm{m}_{f_{i j} \rightarrow X_{i}}$ from factor node to variable node and messages $\mathrm{m}_{X_{i} \rightarrow f_{i j}}$ from variable node to factor node have iterative formulations as follows.
(i) $\mathrm{m}_{f_{i j} \rightarrow X_{i}}\left(x_{i}\right)=\int f_{i j}\left(x_{i}, x_{j}\right) \cdot \mathrm{m}_{X_{j} \rightarrow f_{i j}}\left(x_{j}\right) d x_{j}$;
(ii) $\mathrm{m}_{X_{i} \rightarrow f_{\mathcal{S}}}\left(x_{i}\right)=\prod_{f_{\mathcal{S}^{\prime}} \in \operatorname{neigh}\left(X_{i}\right) \backslash f_{\mathcal{S}}} \mathrm{m}_{f_{\mathcal{S}^{\prime}} \rightarrow X_{i}}\left(x_{i}\right)$.

For the start of sending messages, when a leaf node is a variable node $X_{i}$, the message that it sends along its one and only edge to a factor $f_{\mathcal{S}}$ is $\mathrm{m}_{X_{i} \rightarrow f_{\mathcal{S}}}\left(c_{i}\right)=1$; in the case when a leaf node is a factor node $f_{i}$, the message from the factor node $f_{i}$ to a variable node $X_{i}$ is $\mathrm{m}_{f_{i} \rightarrow X_{i}}\left(x_{i}\right)=f_{i}\left(x_{i}\right)$. Even though the weight function family is not specified here, it can be shown that when the integration in Proposition 20 is well-defined, i.e., the integrands are integrable, then the messages are univariate piecewise functions, which is a striking difference with classical message-passing schemes.

Proposition 21. For any problem in treeWMI, the messages as in Proposition 20 are univariate piecewise functions.

The specific form of messages also depends on the chosen weight function family as mentioned in Section 6.2.1. For example, when the weight functions are chosen to be polynomials, the messages are piecewise polynomials, as in the example in Figure 6.9. We show how to compute the piecewise polynomial messages in Algorithm 8 with functions critical-points and get-msgpieces as subroutines to compute the numeric and symbolic integration bounds for the message

```
Algorithm 7 MP-WMI( \(\Delta\) )
    \(: \mathbf{V}_{\text {up }} \leftarrow\) sort variable nodes in factor graph
    for [ doupward pass] each \(X_{i} \in \mathbf{V}_{\text {up }}\)
        send-message \(\left(X_{i}, f_{i, \mathrm{pa}(i)}\right)\)
        \(\operatorname{send-message}\left(f_{i, \mathrm{pa}(i)}, X_{\mathrm{pa}(i)}\right)\)
    \(\mathbf{V}_{\text {down }} \leftarrow\) sort nodes in set \(\mathbf{V}_{\text {up }}\) in reverse order
    for [ dodownward pass] each \(X_{i} \in \mathbf{V}_{\text {down }}\)
        for each \(X_{c} \in \operatorname{ch}\left(X_{i}\right)\) do
        send-message \(\left(X_{i}, f_{i c}\right)\)
        send-message \(\left(f_{i c}, X_{c}\right)\)
    return \(\left\{\mathrm{m}_{X_{i} \rightarrow f_{s}}, \mathrm{~m}_{f_{s} \rightarrow X_{i}}\right\}_{\left(x_{i}, f_{s}\right) \in \mathbb{E}}\)
```

pieces. Both of them can be efficiently implemented, see [ZB19] for details. The actual integration of the polynomial pieces can be efficiently performed symbolically, as supported by many scientific computing packages.

When MP-WMI terminates, the information stored in the obtained messages is sufficient to compute the unnormalized marginals for each variable and it is independent of the choice of root. Moreover, the integration of unnormalized marginals equals to $\operatorname{WMI}(\Delta, w)$.

Proposition 22. Let $\Delta$ be an SMT formula with a tree factor graph and with per-literal weights $w$. For any variable $X_{i}$, the unnormalized marginal $p\left(x_{i}\right)$ is

$$
p\left(x_{i}\right)=\prod_{f_{s} \in \operatorname{neigh}\left(X_{i}\right)} \mathrm{m}_{f_{\mathcal{S}} \rightarrow X_{i}}\left(x_{i}\right) .
$$

Moreover, the partition function for any $x_{i}$ is the WMI of SMT formula $\Delta$, i.e., $\mathrm{WMI}(\Delta, w)=$ $\int_{x_{i}} p\left(x_{i}\right) d x_{i}$.

### 6.2.2.3 Amortization

We will show that by leveraging the messages pre-computed in MP-WMI, we are able to speed up (amortize) inference time over multiple queries on formula $\Delta$. More specifically, when answering

```
Algorithm 8 send-message \((s, t)\)
    if \(s=X_{i}\) and \(t=f_{i j}\) then
        Return \(\prod_{f_{s^{\prime}} \in \operatorname{neigh}\left(X_{i}\right) \backslash f_{i j}} \mathrm{~m}_{f_{s^{\prime}} \rightarrow X_{i}}\)
    else if \(s=f_{i j}\) and \(t=X_{i}\) then
        \(\mathcal{P} \leftarrow\) critical-points \(\left(\mathrm{m}_{X_{j} \rightarrow f_{i j}}, \Delta_{i j}\right)\)
        \(\mathcal{I} \leftarrow\) intervals-from-points \((\mathcal{P})\)
        for interval \(I \in \mathcal{I}\) consistent with formula \(\Delta_{i j}\) do
            \(\left\langle l_{s}, u_{s}, p\right\rangle \leftarrow\) get-msg-pieces \(\left(\mathrm{m}_{X_{j} \rightarrow f_{i j}}, I, w\right)\)
            \(p^{\prime}\left(x_{i}\right) \leftarrow \int_{l_{s}}^{u_{s}} p\left(x_{i}, x_{j}\right) d x_{j}\)
            \(\mathrm{m}_{f_{i j} \rightarrow X_{i}} \leftarrow \mathrm{~m}_{f_{i j} \rightarrow X_{i}}+\llbracket x_{i} \in I \rrbracket \cdot p^{\prime}\left(x_{i}\right)\)
    return \(\mathrm{m}_{s \rightarrow t}\)
```

queries that do not change the tree structure in the factor graph of formula $\Delta$, we only need to update messages that are related to the queries while other messages are pre-computed. Some examples are SMT queries on a node variable or queries over a pair of variables that are connected by an edge in the factor graph, since these queries either add leaf nodes or do not change existing nodes. Thus we can reuse the local information encoded in messages.

Proposition 23. Let $\mathrm{WMI}(\Delta, w)$ be a problem in treeWMI, and $\Phi$ be an SMT query over a factor $f_{s}^{*}$ involving a variable $X_{i} \in \mathcal{X}$. Given pre-computed messages $\left\{\mathrm{m}_{f_{\mathcal{S}} \rightarrow X_{i}}\right\}_{f_{\mathcal{S}} \in \text { neigh }\left(X_{i}\right)}$,

$$
\begin{aligned}
& \mathrm{WMI}(\Delta \wedge \Phi)=\int_{\mathbb{R}} \mathrm{m}_{f_{s}^{*} \rightarrow X_{i}}^{*}\left(x_{i}\right) \\
& \prod_{f_{s} \in \operatorname{neigh}\left(X_{i}\right) \backslash f_{s}^{*}} \mathrm{~m}_{f_{s} \rightarrow X_{i}}\left(x_{i}\right) d x_{i}
\end{aligned}
$$

with message $\mathrm{m}_{f_{s}^{*} \rightarrow X_{i}}^{*}$ computed over factor $f_{s}^{*}\left(\boldsymbol{x}_{s}\right):=f_{s}\left(\boldsymbol{x}_{s}\right) \cdot \llbracket \boldsymbol{x}_{s}=\Phi \rrbracket$ as in Proposition 20.

Pre-computing messages can dramatically speed up inference by amortization, as we will show in our experiments, especially when traversing the factor graph is expensive or the number of queries is large.

### 6.2.2.4 Complexity Analysis

This section provides a complexity analysis of our proposed WMI solver MP-WMI. Given the SMT formula $\Delta$ with a tree factor graph with a chosen root node, each factor node would be traversed exactly once in each phase of the message-passing scheme. We denote the set of directed factor nodes by $\mathcal{F}:=\left\{\overrightarrow{f_{s}}\right\}=\left\{f_{s}^{+}, f_{s}^{-} \mid f_{s} \in \mathcal{V}\right\}$ where $f_{s}^{+}$denotes the factor node $f_{s}$ visited in the upward pass and $f_{s}^{-}$denotes the one visited in downward pass respectively.

To characterize the message-passing scheme, we define a nilpotent matrix $A$ as follows. The matrix $A \in \mathbb{N}^{|\mathcal{F}| \times|\mathcal{F}|}$ has both its columns and rows denoted by the factor nodes in set $\mathcal{F}$. At each column denoted by $\overrightarrow{f_{s}}$, only entries at rows denoted by factor nodes visited right after $\overrightarrow{f_{s}}$ are non-zero.

Proposition 24. The nilpotent matrix $A$ as described above has its order at most the diameter of the factor graph.

Next we show how to define the non-zero entries in matrix $A$ with parameters about the SMT formulas in WMI problems.

Proposition 25. Suppose that the two variables $X_{i}$ and $X_{j}$ are connected in the factor graph by a factor $f_{i j}$ associated with a sub-formula $\Delta_{i j}$ of size $c$, then in MP-WMI:
(i) the number of pieces in message $\mathrm{m}_{X_{i} \rightarrow f_{i j}}$ is bounded by $\sum m_{s}$, where $m_{s}$ is the number of pieces in message $\mathrm{m}_{f_{s} \rightarrow X_{i}}$ with $f_{s} \in \operatorname{neigh}\left(X_{i}\right) \backslash f_{i j}$;
(ii) the number of pieces in message $\mathrm{m}_{f_{i j} \rightarrow X_{j}}$ is bounded by $2 m c+c^{2}$ with $m$ being the number of pieces in message $\mathrm{m}_{X_{i} \rightarrow f_{i j}}$.

Now we show how to use the matrix $A$ to bound the number of pieces in messages. We define the non-zero entries in the nilpotent matrix $A$ to be $2 c$ with $c$ being a constant that bounds the size of sub-formulas associated to factors. Define a vector $v^{(t)} \in \mathbb{N}|\overrightarrow{\mathbb{E}}|$ for the state of the message-passing scheme at step $t$ - by state it means that each entry in vector $v^{(t)}$ is denoted by a factor node in set $\mathcal{F}$ and the entry denoted by $\overrightarrow{f_{s}}$ bounds the number of pieces in the message
sent to $f_{s}$ in the MP-WMI. For the initial state vector $v^{(0)}$, it has all non-zero entries to be $c$, the constant bounding the sub-formula size, and these entries are those denoted by $\overrightarrow{f_{s}}=f_{s}^{+}$with factor node $f_{s}$ connected with a leaf.

Proposition 26. Let $A$ be the nilpotent matrix and $v$ the initial state vector as described above. Also let $v^{(t)}:=A v^{(t-1)}+c^{2} \cdot \operatorname{sgn}\left(A v^{(t-1)}\right)$ with sgn being the sign function. Then each entry in vector $v^{(t)}$ denoted by $\overrightarrow{f_{s}}$ bounds the number of pieces in the message $\mathrm{m}_{X_{i} \rightarrow f_{s}}$ received by factor $f_{s}$ from some variable node $X_{i}$ at step $t$ in MP-WMI.

Proposition 27. Let $A$ be the nilpotent matrix and $v^{(t)}$ the state vectors as described above. The total number of pieces in the all the messages is bounded by $\left\|\sum_{t=0}^{d} v^{(t)}\right\|_{1}$ with $d$ being the diameter of the factor graph. Moreover, it holds that $\left\|\sum_{t=0}^{d} v^{(t)}\right\|_{1}$ is of $\mathcal{O}\left((4 n c)^{2 d+2}\right)$.

This gives the worst-case total number of message pieces in MP-WMI. From Proposition 27, it holds that the problems in class treeWMI $(\Omega)$ with the weight function family $\boldsymbol{\Omega}$ satisfying TWC are tractable to MP-WMI, since the complexity of MP-WMI is the total number of message pieces multiplied by the symbolic integration cost of each piece, which is tractable for functions in family $\Omega$ by definition. This finishes the constructive proof for Theorem 7 in Section 6.2.1. Notice the complexity of WMI problems depends on the graph structures. In our experiments, we will compare solvers on WMI problems with three representative problem classes with different factor graph diameters.

### 6.2.3 Empirical Evaluation

In this Section, we aim to answer the following research questions: ${ }^{1}$ Q1) Can we effectively scale WMI inference with MP-WMI? Q2) How beneficial is inter-query amortization with MP-WMI?

To answer Q1, we generated a benchmark of WMI problems with tree-shaped primal graphs of different diameters: star-shaped graphs (STAR), complete ternary trees (SNOW) and linear

[^2]

Figure 6.10: Results of the comparison between MP-WMI, WMI-PA and F-XSDD on WMI problems with tree dependencies. In this setting, MP-WMI remarkably scales to problems having up to 60 variables on STAR, while solving SNOW and PATH problems having up to 90 variables, considerably "raising the bar" for the size of tractable WMI inference problems.


Figure 6.11: Log-log plot of cumulative time (seconds, y-axis) for MP-WMI (orange, red, brown) and SMI (blue, green, purple) over STAR, SNOW and PATH primal graphs (see text) with 10, 20 and 30 variables for increasing numbers of univariate and bivariate queries (x-axis). For every class, MP-WMI takes up to two order of magnitude less time when amortizing 100 queries, while being faster than SMI on a single query.
chains (PATH). These structures were originally investigated by the authors of SMI and are prototypical of tree shapes that can be encountered in real-world scenarios such as phylogenetic trees [NK00], hierarchies in file and networks systems [VGR81], and natural language grammars [PBT06].

We sampled random SMT formulas with $N$ variables with the tree structures described above and polynomial weights mapping a subset of literals to a random non-negative polynomials. We generated problems with $N$ ranging from 2 to 20 with step size 2, and from 20 to 100 with
step size 10. We compared our MP-WMI python implementation against the following baselines: WMI-PA [MPS19], a solid general-purpose WMI solver exploiting SMT-based predicate abstraction techniques that is less sensitive to the problem structure; and F-XSDD(BR) [ZKD19], a compilation-based solver achieving state-of-the-art results in several WMI benchmarks.

Fig. 6.10 shows that, with timeout being an hour, our proposed solver MP-WMI is able to scale up to 60 variables for STAR problems and up to 90 variables for SNOW and PATH problems, while the other two solvers stop at problem size 20 for all three classes. Note that the results are in line with those reported in [ZKD19]. This answers Q1 affirmatively, raising the bar of the size of WMI problems that can be solved exactly up to 100 variables.

We tackle Q2 by comparing MP-WMI with SMI [ZB19] on tree-structured MI problems. SMI is a search-based MI solver that has been shown to be efficient for such problems. WMI-PA, FXSDD and the SGDPLL(T) [SOG16] solver are not included in the comparison since they were already shown in [ZB19] to not be competitive on such problems. The synthetic SMT formulas range over $n \in\{10,20,30\}$ variables with tree factor graphs being STAR, SNOW and PATH. We generate 100 univariate or bivariate random queries for each MI problem.

Figure 6.11 shows the cumulative runtime of answering random queries by both solvers. As expected, MP-WMI takes a fraction of the time of SMI (up to two order of magnitudes) to answer 100 univariate or bivariate queries in all experimental scenarios, since it is able to amortize inference inter-query. More surprisingly, by looking at the first point of each curve, we can tell that MP-WMI is even faster than SMI to compute a single query. This is because SMI solves polynomial integration numerically, by reconstructing the univariate polynomials before the numeric integration via interpolation, e.g., Lagrange interpolation; while in MP-WMI we adopt symbolic integration. Hence the complexity of the former is always quadratic in the degree of the polynomial, while for the latter the average case is linear in the number of monomials in the polynomial to integrate, which in practice might be much less than the degree of the polynomial.

### 6.3 Discussion

In this chapter, we first introduce a search-based WMI algorithm that exploits structural independence properties to improve efficiency. For WMI on $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ theories with tree primal graphs and piecewise polynomial weight functions, it decomposes WMI problems during search. A complexity analysis showed that for balanced tree primal graphs, it yields quasi-polynomial complexity. By further theoretically tracing the boundaries of tractable WMI inferece, we introduce another exact WMI solver based on message-passing, MP-WMI, which is efficient on a rich class of tractable WMI problems with tree-shaped factor graphs, the largest known so far. Furthermore, MP-WMI dramatically reduces the answering time of multiple queries by amortizing local computations and allows to compute all marginals and moments simultaneously.

## CHAPTER 7

## Approximate Inference Over Constraints

Beyond our effort on exact WMI solvers, we advance the WMI framework on two fronts. First, we deepen the theoretical understanding of the complexity of WMI inference on real-world problems by proving hardness results. Second, we deliver an efficient and accurate approximate WMI solver as a practical algorithmic solution to deploy WMI inference at a larger scale.

### 7.1 On the hardness of WMI

While the general formulation of WMI we have provided in the previous section is elegant and appealing for advanced probabilistic reasoning, it is, however, not practical in general. In fact, it requires solving an arbitrarily complex integral, which is a \#P-hard problem [BBD11].

To fill this gap, recent works have started looking for classes of tractable WMI problems, i.e., problems for which a solution can be computed exactly in polytime [ZB19, ZMY20a]. These classes of problems can be characterized by two parameters: the treewidth and the diameter of the primal graph of the SMT formulas considered, where the latter is generally expressed as a function of the number of variables in the problem. Note that this is strikingly different from classical discrete probabilistic graphical models, where most of the complexity results are stated in terms of the treewidth alone [KF09, Rot96b].

Definition 12. $\mathcal{W} \mathcal{M I}(\Omega, \delta, t)$ Problem Class) Let $\mathcal{W M} \mathcal{M}(\Omega, \delta, t)$ be the class of WMI problems over models of the form $(\Delta, \mathcal{W})$ on real domains, having primal graph $\mathcal{G}_{\Delta}$ with diameter of $\mathcal{O}(\delta(n))$ and treewidth $t$, where $n$ is the number of variables in the formula $\Delta$; and having per-literal weights $\mathcal{W}$ in a function family $\Omega$.

The largest tractable WMI class known so far has been introduced in our previous work [ZMY20a] as $\mathcal{W M} \mathcal{I}(\Omega, \log (n), 1)$, i.e., the class of problems over $n$ real variables whose primal graph is tree-shaped (treewidth 1 ) and has diameter of length logarithmic in $n$, and whose weight functions belong to a function family $\Omega$ satisfying some conditions called tractable weight conditions (TWCs).

Definition 13. (TWCs) Given a parametric weight function family $\Omega$, it satisfies the TWCs iff
i) it is closed under product, i.e., $\forall f, g \in \Omega, f \cdot g \in \boldsymbol{\Omega}$;
ii) it is closed under definite integration, i.e., $\forall f \in \boldsymbol{\Omega}, F(u(x))-F(l(x)) \in \boldsymbol{\Omega}$ where $F$ is the antiderivative of $f$, and $l(x), u(x)$ are $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ integration bounds for any $x \in \mathcal{X}$;
iii) the symbolic antiderivative of any $f \in \boldsymbol{\Omega}$ can be tractably computed by symbolic integration.

Examples of weight functions in family $\Omega$ include the largely adopted family of (piecewise) polynomials [BPB15], the family of exponentiated linear functions and the family of their products. In the following analysis, we will restrict our attention to weight function families satisfying the TWCs.

In [ZMY20a] the tractability of problem class $\mathcal{W M} \mathcal{M}(\Omega, \log (n), 1)$ is demonstrated by construction, where they introduce a message passing scheme, named MP-WMI, that runs in polytime on tree-shaped and diameter-bounded primal graphs. That is, some sufficient conditions for tractable WMI classes are provided. Here we provide a finer charting of the "tractable islands" of WMI problems by questioning the necessity of the above conditions while looking for larger tractable classes. We prove that unless $P=N P$, larger classes are not tractable. We begin by proving that increasing the diameter of a tree-shaped problem structure makes it hard.

Theorem 8. Let $\mathcal{W} \mathcal{M I}(\Omega, n, 1)$ be the class of $W M I$ problems whose weight function family $\Omega$ satisfies the TWCs. Then inference in $\mathcal{W} \mathcal{M I}(\Omega, n, 1)$ is $\# P$-hard.

Proof. We prove our complexity result by reducing a \#P-complete variant of the subset sum problem [GJ02] to an MI problem over an $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ formula $\Delta$ with tree primal graph whose diameter is $n$. This problem is a counting version of subset sum problem saying that given a set of
positive integers $S=\left\{s_{1}, s_{2}, \cdots, s_{n}\right\}$, and a positive integer $L$, the goal is to count the number of subsets $S^{\prime} \subseteq S$ such that the sum of all the integers in the subset $S^{\prime}$ equals to $L$. Notice that our proof can be applied to rational numbers as well and we assume binary representations for numbers.

First, we reduce the counting subset sum problem in polynomial time to a model integration problem by constructing the following $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ formula $\Delta$ on real variables $\mathcal{X}$ whose primal graph is shown in Figure 7.1:


Figure 7.1: Primal graph $\mathcal{G}_{\Delta}$ used for the \#P-hardness reduction in Theorem 8. We construct the corresponding formula $\Delta$ such that $\mathcal{G}_{\Delta}$ has maximum diameter (it is a chain). We graphically augment graph $\mathcal{G}_{\Delta}$ by introducing blue nodes to indicate that integers $s_{i}$ in set $S$ are contained in clauses between two variables.
$\Delta=\left\{\begin{array}{l}\underbrace{s_{1}-\frac{1}{2 n}<x_{1}<s_{1}+\frac{1}{2 n}}_{\ell(1,0)} \vee \underbrace{-\frac{1}{2 n}<x_{1}<\frac{1}{2 n}}_{\ell(1,1)} \\ \underbrace{x_{i-1}+s_{i}-\frac{1}{2 n}<x_{i}<x_{i-1}+s_{i}+\frac{1}{2 n}}_{\ell(i, 0)} \vee \underbrace{x_{i-1}-\frac{1}{2 n}<x_{i}<x_{i-1}+\frac{1}{2 n}}_{\ell(i, 1)}, \quad i=2, \cdots n\end{array}\right.$
For brevity, we denote the first and the second literal in the $i$-th clause by $\ell(i, 0)$ and $\ell(i, 1)$ respectively as shown above. Also We choose two constants $l=L-\frac{1}{2}$ and $u=L+\frac{1}{2}$.

In the following, we prove that $n^{n} \operatorname{MI}\left(\Delta \wedge\left(l<X_{n}<u\right)\right)$ equals to the number of subset $S^{\prime} \subseteq S$ whose element sum equals to $L$, which indicates that WMI problem whose tree primal graph has diameter $\mathcal{O}(n)$ is \#P-hard.

Let $\boldsymbol{a}^{k}=\left(a_{1}, a_{2}, \cdots, a_{k}\right)$ be some assignment to Boolean variables $\left(A_{1}, A_{2}, \cdots, A_{k}\right)$ with $a_{i} \in\{0,1\}, i \in[k]$. Given an assignment $\boldsymbol{a}^{k}$, we define subset sums to be $S\left(\boldsymbol{a}^{k}\right) \triangleq \sum_{i=1}^{k} a_{i} s_{i}$, and formulas $\Delta_{\boldsymbol{a}^{k}} \triangleq \bigwedge_{i=1}^{k} \ell\left(i, a_{i}\right)$.

Claim 9. The model integration for formula $\Delta_{\boldsymbol{a}^{k}}$ with an given assignment $\boldsymbol{a}^{k} \in\{0,1\}^{k}$ is $\operatorname{MI}\left(\Delta_{\boldsymbol{a}^{k}}\right)=$ $\left(\frac{1}{n}\right)^{k}$. Moreover, for each variable $X_{i}$ in $\Delta_{\boldsymbol{a}^{k}}$, its satisfying assignments consist of the interval [ $\sum_{j=1}^{i} a_{j} s_{j}-\frac{i}{2 n}, \sum_{j=1}^{i} a_{j} s_{j}+\frac{i}{2 n}$ ]. Specifically, the satisfying assignments for variable $X_{n}$ in formula $\Delta_{\boldsymbol{a}^{n}}$ can be denoted by the interval $\left[S\left(\boldsymbol{a}^{n}\right)-\frac{1}{2}, S\left(\boldsymbol{a}^{n}\right)+\frac{1}{2}\right]$.

Proof. (Claim 9) First we prove that $\operatorname{MI}\left(\Delta_{a^{k}}\right)=\left(\frac{1}{n}\right)^{k}$. For brevity, denote $a_{i} s_{i}$ by $\hat{s_{i}}$. By definition of model integration and the fact that the integral is absolutely convergent (since we are integrating a constant function, i.e., one, over finite volume regions), we have the following equation.

$$
\operatorname{MI}\left(\Delta_{\boldsymbol{a}^{k}}\right)=\int_{\left(x_{1}, \cdots, x_{k}\right) \neq \Delta_{\boldsymbol{a}^{k}}} 1 d x_{1} \cdots d x_{k}=\int_{\hat{s}_{1}-\frac{1}{2 n}}^{\hat{s}_{1}+\frac{1}{2 n}} d x_{1} \cdots \int_{x_{k-2}+\hat{s}_{k-1}-\frac{1}{2 n}}^{x_{k-2}+\hat{s}_{k-1}+\frac{1}{2 n}} d x_{k-1} \int_{x_{k-1}+\hat{s}_{k}-\frac{1}{2 n}}^{x_{k-1}+\hat{s}_{k}+\frac{1}{2 n}} 1 d x_{k}
$$

Observe that for the most inner integration over variable $x_{k}$, the integration result is $\frac{1}{n}$. By doing this iteratively, we have that $\operatorname{MI}\left(\Delta_{a^{k}}\right)=\left(\frac{1}{n}\right)^{k}$.

Next we prove that satisfying assignments for variable $X_{i}$ in formula $\Delta_{a^{k}}$ is the interval [ $\sum_{j=1}^{i} a_{j} s_{j}-\frac{i}{2 n}, \sum_{j=1}^{i} a_{j} s_{j}+\frac{i}{2 n}$ ] by mathematical induction. For $i=1$, since $X_{1}$ is in interval $\left[a_{1} s_{1}-\frac{1}{2 n}, a_{1} s_{1}+\frac{1}{2 n}\right]$, the statement holds in this case. Suppose that the statement holds for $i=m$, i.e. variable $X_{m}$ has its satisfying assignments in interval $\left[\sum_{j=1}^{m} a_{j} s_{j}-\frac{m}{2 n}, \sum_{j=1}^{m} a_{j} s_{j}+\frac{m}{2 n}\right.$ ]. Since variable $X_{m+1}$ has its satisfying assignments in interval $\left[X_{m}+a_{m+1} s_{m+1}-\frac{1}{2 n}, X_{m}+a_{m+1} s_{m+1}+\right.$ $\frac{1}{2 n}$ ], then its satisfying assignments consist interval [ $\sum_{j=1}^{m+1} a_{j} s_{j}-\frac{m+1}{2 n}, \sum_{j=1}^{m+1} a_{j} s_{j}+\frac{m+1}{2 n}$ ], that is, the statement also holds for $i=m+1$. Thus the claim holds.

The above claim shows how to compute the model integration of formula $\Delta_{a^{k}}$. We will show in the next claim how to compute the model integration of formula $\Delta_{a^{n}}$ conjoined with a query $l<X_{n}<u$.

Claim 10. For each assignment $\boldsymbol{a}^{n} \in\{0,1\}^{n}$, the model integration offormula $\Delta_{\boldsymbol{a}^{n}} \wedge\left(l<X_{n}<u\right)$ falls into one of the following cases:
i) If $S\left(\boldsymbol{a}^{n}\right)<L$ or $S\left(\boldsymbol{a}^{n}\right)>L$, it holds that $\operatorname{MI}\left(\Delta_{\boldsymbol{a}^{n}} \wedge\left(l<X_{n}<u\right)\right)=0$.
ii) If $S\left(\boldsymbol{a}^{n}\right)=L$, it holds that $\operatorname{MI}\left(\Delta_{a^{n}} \wedge\left(l<X_{n}<u\right)\right)=\left(\frac{1}{n}\right)^{n}$.

Proof. (Claim 10) From the previous Claim 9, it is shown that variable $X_{n}$ has its satisfying assignments in interval $\left[S\left(\boldsymbol{a}^{n}\right)-\frac{1}{2}, S\left(\boldsymbol{a}^{n}\right)+\frac{1}{2}\right]$ in formula $\Delta_{\boldsymbol{a}^{n}}$ for each $\boldsymbol{a}^{n} \in\{0,1\}^{n}$. If $S\left(\boldsymbol{a}^{n}\right)<L$, given that $S\left(\boldsymbol{a}^{n}\right)$ is a sum of positive integers, then it holds that $S\left(\boldsymbol{a}^{n}\right)+\frac{1}{2} \leq(L-1)+\frac{1}{2}=$ $L-\frac{1}{2}=l$ and therefore, $\operatorname{MI}\left(\Delta_{\boldsymbol{a}^{n}} \wedge\left(l<X_{n}<u\right)\right)=0$; similarly, if $S\left(\boldsymbol{a}^{n}\right)>L$, then it holds that $S\left(\boldsymbol{a}^{n}\right)-\frac{1}{2} \geq u$ and therefore, $\operatorname{MI}\left(\Delta_{\boldsymbol{a}^{n}} \wedge\left(l<X_{n}<u\right)\right)=0$. If $S\left(\boldsymbol{a}^{n}\right)=L$, by Claim 9 we have that the satisfying assignment interval is inside the interval $[l, u]$ and thus it holds that $\operatorname{MI}\left(\Delta_{a^{n}} \wedge\left(l<X_{n}<u\right)\right)=\operatorname{MI}\left(\Delta_{a^{n}}\right)=\left(\frac{1}{n}\right)^{n}$.

In the next claim, we show how to compute the model integration of formula $\Delta$ as well as for formula $\Delta$ conjoined with query $l<X_{n}<u$ based on the already proven Claim 9 and Claim 10.

Claim 11. The following two equations hold:
i) $\operatorname{MI}(\Delta)=\sum_{a^{n}} \operatorname{MI}\left(\Delta_{a^{n}}\right)$.
ii) $\operatorname{MI}\left(\Delta \wedge\left(l<X_{n}<u\right)\right)=\sum_{\boldsymbol{a}^{n}} \operatorname{MI}\left(\Delta_{\boldsymbol{a}^{n}} \wedge\left(l<X_{n}<u\right)\right)$.

Proof. (Claim 11) Observe that for each clause in $\Delta$, literals are mutually exclusive since each $s_{i}$ is a positive integer. Then we have that formulas $\Delta_{a^{n}}$ are mutually exclusive and meanwhile $\Delta=$ $\bigvee_{\boldsymbol{a}^{n}} \Delta_{\boldsymbol{a}^{n}}$. Thus it holds that $\operatorname{MI}(\Delta)=\sum_{\boldsymbol{a}^{n}} \operatorname{MI}\left(\Delta_{\boldsymbol{a}^{n}}\right)$. Similarly, we have formulas $\left(\Delta_{\boldsymbol{a}^{n}} \wedge(l<\right.$ $\left.X_{n}<u\right)$ )'s are mutually exclusive and meanwhile $\Delta \wedge\left(l<X_{n}<u\right)=\bigvee_{\boldsymbol{a}^{n}} \Delta_{\boldsymbol{a}^{n}} \wedge\left(l<X_{n}<u\right)$. Thus the second equation holds.

From the above claims, we can conclude that $\operatorname{MI}\left(\Delta \wedge\left(l<X_{n}<u\right)\right)=t\left(\frac{1}{n}\right)^{n}$ where $t$ is the number of assignments $\boldsymbol{a}^{n}$ s.t. $S\left(\boldsymbol{a}^{n}\right)=L$. Notice that for each $\boldsymbol{a}^{n} \in\{0,1\}^{n}$, there is a one-toone correspondance to a subset $S^{\prime} \subseteq S$ by defining $\boldsymbol{a}^{n}$ as $a_{i}=1$ if and only if $s_{i} \in S^{\prime}$; and $S\left(\boldsymbol{a}^{n}\right)$ equals to $L$ if and only if the sum of elements in $S^{\prime}$ is $L$. Therefore $n^{n} \operatorname{MI}\left(\Delta \wedge\left(l<X_{n}<u\right)\right)$ equals to the number of subset $S^{\prime} \subseteq S$ whose element sum equals to $L$. This finishes the proof for the statement that inference in $\mathcal{W} \mathcal{M I}(\Omega, n, 1)$ is \#P-hard.

Next, we turn our attention to another class of WMI problems, the class $\mathcal{W} \mathcal{M I}(\Omega, \log (n), 2)$, having logarithmic diameter but treewidth 2. This class is also supposed to be "easy" in the sense


Figure 7.2: Primal graph used for \#P-hardness reduction in Theorem 7. We also put blue nodes to indicate that integer $s_{i}$ 's in set $S$ are contained in some clauses and that model integration over some cliques is the sum of some $s_{i}$ 's.
that it extends the tractable class $\mathcal{W M} \mathcal{I}(\Omega, \log (n), 1)$ by slightly increasing the treewidth by one. Unfortunately, inference in $\mathcal{W M} \mathcal{M}(\Omega, \log (n), 2)$ is also hard.

Theorem 12. Let $\mathcal{W M} \mathcal{M}(\Omega, \log (n), 2)$ be the class of WMI problems whose parametric weight function family $\Omega$ satisfies the TWCs. Then inference in $\mathcal{W M} \mathcal{M}(\Omega, \log (n), 2)$ is $\# P$-hard.

Proof. Again we prove our complexity result by reducing the \#P-complete variant of the subset sum problem [GJ02] to an MI problem over an $\operatorname{SMT}(\mathcal{L} \mathcal{R} \mathcal{A})$ formula $\Delta$ with primal graph whose diameter is $\mathcal{O}(\log n)$ and treewidth two. In the $\# \mathrm{P}$-complete subset sum problem, we are given a set of positive integers $S=\left\{s_{1}, s_{2}, \cdots, s_{n}\right\}$, and a positive integer $L$. Notice that our proof can be applied to rational numbers as well and we assume binary representations for numbers. The goal is to count the number of subsets $S^{\prime} \subseteq S$ such that the sum of all the integers in $S^{\prime}$ equals $L$.

First, we reduce this problem in polynomial time to a model integration problem with the following $\operatorname{SMT}(\mathcal{L R} \mathcal{A})$ formula $\Delta$ where variables are real and $u$ and $l$ are two constants. Its primal graph is shown in Figure 7.2. Consider $n=2^{k}, n, k \in \mathbb{N}$.

$$
\begin{aligned}
& \Delta=\bigwedge_{i \in[n]}\left(-\frac{1}{4 n}<X_{k+1, i}<\frac{1}{4 n} \vee-\frac{1}{4 n}+s_{i}<X_{k+1, i}<\frac{1}{4 n}+s_{i}\right) \bigwedge \Delta_{t} \\
& \text { where } \Delta_{t}=\bigwedge_{j \in[k], i \in\left[2^{j}\right]}-\frac{1}{4 n}+X_{j+1,2 i-1}+X_{j+1,2 i}<X_{j, i}<\frac{1}{4 n}+X_{j+1,2 i-1}+X_{j+1,2 i}
\end{aligned}
$$

For brevity, we denote all the variables by $\mathcal{X}$ and denote the literal $-\frac{1}{4 n}<X_{k+1, i}<\frac{1}{4 n}$ by $\ell(i, 0)$ and literal $-\frac{1}{4 n}+s_{i}<X_{k+1, i}<\frac{1}{4 n}+s_{i}$ by $\ell(i, 1)$ respectively. Also We choose two constants $l=L-\frac{1}{2}$ and $u=L+\frac{1}{2}$. In the following, we prove that $(2 n)^{2 n-1} \mathrm{MI}\left(\Delta \wedge\left(l<X_{1,1}<\right.\right.$ $u)$ ) equals to the number of subset $S^{\prime} \subseteq S$ whose element sum equals to $L$, which indicates that model integration problem with primal graph whose diameter is $\mathcal{O}(\log n)$ and treewidth two is \#P-hard.

Let $\boldsymbol{a}^{n}=\left(a_{1}, a_{2}, \cdots, a_{n}\right) \in\{0,1\}^{n}$ be some assignment to Boolean variables $\left(A_{1}, A_{2}, \cdots, A_{n}\right)$. Given assignment $\boldsymbol{a}^{n}$, define sum as $S\left(\boldsymbol{a}^{n}\right) \triangleq \sum_{i=1}^{n} a_{i} s_{i}$, and formula as $\Delta_{\boldsymbol{a}^{n}} \triangleq \bigwedge_{i=1}^{n} \ell\left(i, a_{i}\right) \wedge \Delta_{t}$. Claim 13. The model integration for formula $\Delta_{\boldsymbol{a}^{n}}$ with given $\boldsymbol{a}^{n} \in\{0,1\}^{n}$ is $\operatorname{MI}\left(\Delta_{\boldsymbol{a}^{n}}\right)=\left(\frac{1}{2 n}\right)^{2 n-1}$. Moreover, for each variable $X_{j, i}$ in formula $\Delta_{a^{n}}$, its satisfying assignments consist of the interval $\left[\sum_{l} a_{l} s_{l}-\frac{2^{k-j+2}-1}{4 n}, \sum_{l} a_{l} s_{l}+\frac{2^{k-j+2}-1}{4 n}\right]$ where $l \in\left\{l \mid X_{k+1, l}\right.$ is a descendant of $\left.X_{j, i}\right\}$. Specifically, the satisfying assignments for the root variable $X_{1,1}$ can be denoted the interval $\left[S\left(\boldsymbol{a}^{n}\right)-\right.$ $\left.\frac{2 n-1}{4 n}, S\left(\boldsymbol{a}^{n}\right)+\frac{2 n-1}{4 n}\right] \subset\left[S\left(\boldsymbol{a}^{n}\right)-\frac{1}{2}, S\left(\boldsymbol{a}^{n}\right)+\frac{1}{2}\right]$.

Proof. (Claim 13) First we prove that $\operatorname{MI}\left(\Delta_{\boldsymbol{a}^{n}}\right)=\left(\frac{1}{2 n}\right)^{2 n-1}$. For brevity, denote $a_{i} s_{i}$ by $\hat{s_{i}}$. By definition of model integration and the fact that the integral is absolutely convergent (since we are integrating a constant function, i.e., one, over finite volume regions), we have the following

$$
\begin{aligned}
& \operatorname{MI}\left(\Delta_{\boldsymbol{a}^{n}}\right)=\int_{\boldsymbol{x} \vDash \Delta_{\boldsymbol{a}^{n}}} 1 d \mathcal{X} \\
& =\int_{-\frac{1}{4 n}+\hat{s}_{n}}^{\frac{1}{4 n}+\hat{s}_{n}} d x_{k+1, n} \cdots \int_{-\frac{1}{4 n}+\hat{s}_{1}}^{\frac{1}{4 n}+\hat{s}_{1}} d x_{k+1,1} \int_{-\frac{1}{4 n}+x_{k+1, n-1}+x_{k+1, n}}^{\frac{1}{4 n}+x_{k+1, n-1}+x_{k+1, n}} d x_{k, 2^{k-1} \cdots \int_{-\frac{1}{4 n}+x_{2,1}+x_{2,2}}^{\frac{1}{4 n}+x_{2,1}+x_{2,2}} 1 d x_{1,1} .} .
\end{aligned}
$$

Observe that for the most inner integration over variable $x_{1,1}$, the integration result is $\frac{1}{2 n}$. By doing this iteratively, we have that $\operatorname{MI}\left(\Delta_{\boldsymbol{a}^{k}}\right)=\left(\frac{1}{2 n}\right)^{2 n-1}$ where the $2 n-1$ comes from the
number of variables.
Then we prove that satisfying assignments for variable $X_{j, i}$ in formula $\Delta_{\boldsymbol{a}^{n}}$ lie in the inter$\operatorname{val}\left[\sum_{l} a_{l} s_{l}-\frac{2^{k-j+2}-1}{4 n}, \sum_{l} a_{l} s_{l}+\frac{2^{k-j+2}-1}{4 n}\right]$ where $l \in\left\{l \mid X_{k+1, l}\right.$ is a descendant of $\left.X_{j, i}\right\}$ by performing mathematical induction in a bottom-up way.

For $j=1$, any variable $X_{k+2-j, i}$ with $i \in\left[2^{k+2-j}\right]$ has satisfying assignments consisting of the interval $\left[a_{i} s_{i}-\frac{1}{4 n}, a_{i} s_{i}+\frac{1}{4 n}\right]$. Thus the statement holds for this case.

Suppose that the statement holds for $j=m$, that is, for any $i \in\left[2^{k+2-m}\right]$, any variable $X_{k+2-m, i}$ has satisfying assignments consisting interval $\left[\sum_{l} a_{l} s_{l}-\frac{2^{m}-1}{4 n}, \sum_{l} a_{l} s_{l}+\frac{2^{m}-1}{4 n}\right]$ where $l \in\left\{l \mid X_{k+1, l}\right.$ is a descendant of $\left.X_{k+2-m, i}\right\}$.

Then for $j=m+1$ and any $i \in\left[2^{k+1-m}\right]$, the variable $X_{k+1-m, i}$ has two descendants, variable $X_{k+2-m, 2 i-1}$ and variable $X_{k+2-m, 2 i}$. Moreover, we have that $-\frac{1}{4 n}+X_{k+2-m, 2 i-1}+X_{k+2-m, 2 i}<$ $X_{k+1-m, i}<\frac{1}{4 n}+X_{k+2-m, 2 i-1}+X_{k+2-m, 2 i}$. Then the lower bound of the interval for variable $X_{k+1-m, i}$ is $-\frac{1}{4 n}+\sum_{l} a_{l} s_{l}-2 \frac{2^{m}-1}{4 n}=\sum_{l} a_{l} s_{l}-\frac{2^{m+1}-1}{4 n}$; similarly the upper bound of the interval is $\sum_{l} a_{l} s_{l}+\frac{2^{m+1}-1}{4 n}$, where $l \in\left\{l \mid X_{k+1, l}\right.$ is a descendant of $\left.X_{k+1-m, i}\right\}$. That is, the statement also holds for $j=m+1$ which finishes our proof.

The above claim shows what the model integration of formula $\Delta_{a^{k}}$ is like. We'll show in the next claim what the model integration of formula $\Delta_{a^{n}}$ conjoined with a query $l<X_{1,1}<u$ is like.

Claim 14. For each assignments $\boldsymbol{a}^{n} \in\{0,1\}^{n}$, the model integration of $\Delta_{\boldsymbol{a}^{n}} \wedge\left(l<X_{1,1}<u\right)$ falls into one of the following cases:
i) If $S\left(\boldsymbol{a}^{n}\right)<L$ or $S\left(\boldsymbol{a}^{n}\right)>L$, then $\operatorname{MI}\left(\Delta_{a^{n}} \wedge\left(l<X_{1,1}<u\right)\right)=0$.
ii) If $S\left(\boldsymbol{a}^{n}\right)=L$, then $\operatorname{MI}\left(\Delta_{\boldsymbol{a}^{n}} \wedge\left(l<X_{1,1}<u\right)\right)=\left(\frac{1}{2 n}\right)^{2 n-1}$.

Proof. (Claim 14) From previous Claim 13, it is shown that variable $X_{1,1}$ has its satisfying assignments in the interval $\left[S\left(\boldsymbol{a}^{n}\right)-\frac{2 n-1}{4 n}, S\left(\boldsymbol{a}^{n}\right)+\frac{2 n-1}{4 n}\right]$ in formula $\Delta_{\boldsymbol{a}^{n}}$ for each $\boldsymbol{a}^{n} \in\{0,1\}^{n}$.

If $S\left(\boldsymbol{a}^{n}\right)<L$, given that $S\left(\boldsymbol{a}^{n}\right)$ is a sum of positive integers, then it holds that $S\left(\boldsymbol{a}^{n}\right)+\frac{1}{2} \leq$
$(L-1)+\frac{2 n-1}{4 n}<L-\frac{1}{2}=l$ and therefore, $\operatorname{MI}\left(\Delta_{a^{n}} \wedge\left(l<X_{1,1}<u\right)\right)=0$; similarly, if $S\left(\boldsymbol{a}^{n}\right)>L$, then it holds that $S\left(\boldsymbol{a}^{n}\right)-\frac{1}{2}>u$ and therefore, $\operatorname{MI}\left(\Delta_{\boldsymbol{a}^{n}} \wedge\left(l<X_{1,1}<u\right)\right)=0$. If $S\left(\boldsymbol{a}^{n}\right)=L$, then by Claim 13 we have that the satisfying assignment interval is inside the interval $[l, u]$ and thus it holds that $\operatorname{MI}\left(\Delta_{\boldsymbol{a}^{n}} \wedge\left(l<X_{1,1}<u\right)\right)=\operatorname{MI}\left(\Delta_{\boldsymbol{a}^{n}}\right)=\left(\frac{1}{2 n}\right)^{2 n-1}$.

Claim 15. The following two equations hold:
i) $\operatorname{MI}(\Delta)=\sum_{a^{n}} \operatorname{MI}\left(\Delta_{a^{n}}\right)$.
ii) $\operatorname{MI}\left(\Delta \wedge\left(l<X_{1,1}<u\right)\right)=\sum_{\boldsymbol{a}^{n}} M I\left(\Delta_{\boldsymbol{a}^{n}} \wedge\left(l<X_{1,1}<u\right)\right)$.

Proof. (Claim 15) Observe that for each pair of literals $\ell(i, 0)$ and $\ell(i, 1), i \in[n]$, literals are mutually exclusive since each $s_{i}$ is a positive integer. Then we have that formulas $\Delta_{\boldsymbol{a}^{n}}$ are mutually exclusive and meanwhile formula $\Delta=\bigvee_{a^{n}} \Delta_{a^{n}}$. Thus it holds that $\operatorname{MI}(\Delta)=\sum_{a^{n}} \operatorname{MI}\left(\Delta_{a^{n}}\right)$. Similarly, we have formulas $\left(\Delta_{\boldsymbol{a}^{n}} \wedge\left(l<X_{1,1}<u\right)\right)$ 's are mutually exclusive and meanwhile $\Delta \wedge\left(l<X_{1,1}<u\right)=\bigvee_{\boldsymbol{a}^{n}} \Delta_{\boldsymbol{a}^{n}} \wedge\left(l<X_{1,1}<u\right)$. Thus the second equation holds.

From the above claims, we can conclude that $\operatorname{MI}\left(\Delta \wedge\left(l<X_{1,1}<u\right)\right)=t\left(\frac{1}{2 n}\right)^{2 n-1}$ where $t$ is the number of assignments $\boldsymbol{a}^{n}$ s.t. $S\left(\boldsymbol{a}^{n}\right)=L$. Notice that for each $\boldsymbol{a}^{n} \in\{0,1\}^{n}$, there is a one-toone correspondence to a subset $S^{\prime} \subseteq S$ by defining $\boldsymbol{a}^{n}$ as $a_{i}=1$ if and only if $s_{i} \in S^{\prime}$; and $S\left(\boldsymbol{a}^{n}\right)$ equals to $L$ if and only if the sum of elements in $S^{\prime}$ is $L$. Therefore $(2 n)^{2 n-1} \mathrm{MI}\left(\Delta \wedge\left(l<X_{1,1}<u\right)\right)$ equals to the number of subset $S^{\prime} \subseteq S$ whose element sum equals to $L$. This finishes the proof for the statement that inference in $\mathcal{W M} \mathcal{M}(\Omega, \log (n), 2)$ is $\# \mathrm{P}$-hard.

Note that our result differs from the one presented in [ZMY20a] for the hardness of the class $2 \mathrm{WMI}(\Omega)$, containing WMI problems with SMT formulas being conjunctions of clauses comprising at most two variables. In fact, $\mathcal{W M} \mathcal{M}(\boldsymbol{\Omega}, \log (n), 2)$ is contained in $2 \mathrm{WMI}(\boldsymbol{\Omega})$. As such, we trace the tractablity boundaries of WMI inference with higher precision, as the next corollary states. Its proof follows from Theorems 8 and 12 and from the sufficiency as demonstrated in [ZMY20a].

Corollary 3. Let $\mathcal{W M} \mathcal{M}(\Omega, \log (n), t)$ be the class of WMI problems whose parametric weight function family $\Omega$ satisfies the TWCs. Then $\mathcal{W} \mathcal{M I}(\Omega, \log (n), t)$ is a tractable WMI class for inference if-and-only-if treewidth $t=1$.

These complexity results set the standard for the solver complexity: every exact WMI solver that aims to be efficient, needs to operate in the regime of Corollary 3. However, real-world problems do not always conform to the structural desiderata for primal graphs stated in it. This implies that efficient approximations might not only be useful in these scenarios, but needed. In the next section we fill this gap, by introducing our approximate WMI solver that navigates the tractable islands in WMI problems by performing efficient inference on a relaxed version of intractable WMI problems.

### 7.2 ReCoin: Approximate WMI Solver

Our algorithm to approximate WMI inference comprises three phases: i) RElaxing an intractable WMI model into a simpler one amenable to exact inference by removing dependencies from it; then ii) introduce certain literals and weights to COmpensate for the dependency structure lost in this way and iii) optimize them by solving a series of exact INtegration problems. We name it ReCoIn. With ReCoIn we can navigate a spectrum of approximations - with the original primal graph $\mathcal{G}_{\Delta}$ on one end, and a fully disconnected version on the other - by removing more and more edges. As such, ReCoIn can be viewed as extending the relax-compensate-recover (RCR) framework [CD06, CD10, CD12] for approximate inference on discrete probabilistic models to continuous representations and in presence of algebraic constraints.

### 7.2.1 Relaxation: introducing and then "breaking" equivalence constraints

The aim of the relaxation step is to obtain a new SMT formula $\Delta^{\text {rel }}$ such that its associated primal graph $\mathcal{G}_{\Delta^{\text {rel }}}$, serves as the simplification of the original $\mathcal{G}_{\Delta}$ by removing a given set of edges. We will show that the removal of any edge can be formulated as the removal of an equivalence edge [CD12]. This process consists of two steps. First, we create an augmented formula $\Delta^{\text {aug }}$ by
introducing new variables to $\Delta$ and enforcing them to act as copies of certain original variables by explicitly adding equivalence constraints. Second, we deliver the relaxed $\mathcal{G}_{\Delta^{\text {rel }}}$ by removing these equivalence constraints.

### 7.2.1.1 Augmentation.

The detailed process of distilling a new augmented model ( $\left.\Delta^{\text {aug }}, \mathcal{W}^{\text {aug }}\right)$ from $(\Delta, \mathcal{W})$, given a subset of edges $\mathbb{E}_{d} \subseteq \mathbb{E}$ in $\mathcal{G}_{\Delta}$ to remove, is listed in Algorithm 9. At its core, there are routines for copying one variable and adding the corresponding equivalence constraints and compensating literals. For each edge $X_{i}-X_{j} \in \mathbb{E}_{d}$ to be removed, one of its variables is arbitrarily selected, say $X_{i}$. Then a variable $X_{i}^{c}$, as a copy of the chosen $X_{i}$, is introduced in $\Delta^{\text {aug }}$ as well as one equivalence constraint between the two as the literal $\hat{\ell}:\left(X_{i}^{c}=X_{i}\right)$ with associated weight function $\delta\left(X_{i}, X_{i}^{c}\right)$ where $\delta$ is the Dirac delta function. Then we properly rename all occurrences of $X_{i}$ by $X_{i}^{c}$ in the literals appearing in the clauses of $\Delta^{\text {aug }}$ that also contain $X_{j}$ and introduce copied literals for the univariate clauses over $X_{i}$ only. These steps cause the primal graph $\mathcal{G}_{\Delta^{\text {aug }}}$ to now contain the dependency $X_{i}-X_{i}^{c}-X_{j}$ but not $X_{i}-X_{j}$.

Note that the augmented WMI model ( $\left.\Delta^{\text {aug }}, \mathcal{W}^{\text {aug }}\right)$ now contains more variables than the original one. Specifically, for each variable $X_{i} \in \mathcal{G}_{\Delta}$ we might have introduced $C_{i}$ different copies in $\mathcal{G}_{\Delta^{\text {aug }}}$, denoted as $X_{i}^{1}, \ldots, X_{i}^{C_{i}}$, if we removed $C_{i}$ edges over $X_{i}$. We will denote the original $X_{i}$ as $X_{i}^{0}$ for notation consistency. Even if the dimensionality of the augmented WMI problem is increased by augmentation, the next propositions are guaranteeing that we are not altering the partition function and the marginal distributions of $\mathrm{Pr}_{\Delta}$, and that introducing equivalence constraints does not alter the induced distribution.

Proposition 28. Let $\Delta$ be an SMT formula with primal graph $\mathcal{G}_{\Delta}$ and per-literal weight functions $\mathcal{W}$, and let $\Delta^{\text {aug }}$ and $\mathcal{W}^{\text {aug }}$ be the output of Algorithm 9 when applied to $\Delta$ and $\mathcal{G}_{\Delta}$ given a certain subset of edges in $\mathcal{G}_{\Delta}$. Then it holds that $\operatorname{WMI}(\Delta, \mathcal{W})=\operatorname{WMI}\left(\Delta^{\text {aug }}, \mathcal{W}^{\text {aug }}\right)$. Moreover, for any $X_{i}$ in $\mathcal{G}_{\Delta}$ and univariate literal $\ell$ over $X_{i}$, it holds that $\operatorname{Pr}_{\Delta}(\ell)=\operatorname{Pr}_{\Delta^{\text {aug }}}(\ell)$.

Algorithm 9 augment $\operatorname{Model}\left(\Delta, \mathcal{W}, \mathbb{E}_{d}\right)$
Input: a WMI model with SMT formula $\Delta$ and per-literal weights $\mathcal{W}$ and a set $\mathbb{E}_{d}$ of edges to be deleted

Output: augmented WMI model ( $\left.\Delta^{\text {aug }}, \mathcal{W}^{\text {aug }}\right)$ and equivalence constraint set $\mathcal{L}$
$: \Delta^{\mathrm{aug}} \leftarrow \operatorname{copy}(\Delta)$
$\mathcal{W}^{\text {aug }} \leftarrow \operatorname{copy}(\mathcal{W})$
$\mathcal{L} \leftarrow\}$
: for edge $X_{i}-X_{j} \in \mathbb{E}_{d}$ do
$X_{i}^{c} \leftarrow \operatorname{copy}\left(X_{i}\right) \quad / /$ Assume to copy $X_{i}$
$\hat{\ell} \leftarrow\left(X_{i}=X_{i}^{c}\right)$
$\mathcal{L} \leftarrow \mathcal{L} \cup\{\hat{\ell}\}$
$\Delta^{\prime} \leftarrow \Delta^{\text {aug }} \wedge \hat{\ell}$,
$w_{\hat{\ell}}:=\delta\left(X_{i}, X_{i}^{c}\right)$
$\mathcal{W}^{\text {aug }} \leftarrow \mathcal{W}^{\text {aug }} \cup\left\{w_{\hat{\ell}}\right\}$
for clause $\Gamma \in \Delta_{i, j}$ do // Rename edges
$\Gamma^{\prime} \leftarrow \Gamma\left\{X_{i}: X_{i}^{c}\right\}$
$\Delta^{\prime} \leftarrow \Delta^{\prime}\left\{\Gamma: \Gamma^{\prime}\right\}$
for each literal $\ell \in \Gamma$ do

$$
\begin{aligned}
& \ell^{\prime} \leftarrow \ell\left\{X_{i}: X_{i}^{c}\right\} \\
& w_{\ell^{\prime}} \leftarrow \operatorname{copy}\left(w_{\ell}\right) \\
& \mathcal{W}^{\text {aug }} \leftarrow \mathcal{W}^{\text {aug }} \cup\left\{w_{\ell^{\prime}}\right\} \backslash\left\{w_{\ell}\right\}
\end{aligned}
$$

for clause $\Gamma \in \Delta_{i}$ do // Copy and rename bounding-box literals
$\Gamma^{\prime} \leftarrow \operatorname{copy}(\Gamma)$
$\Delta^{\prime} \leftarrow \Delta^{\prime} \wedge \Gamma^{\prime}\left\{X_{i}: X_{i}^{c}\right\}$
$\Delta^{\text {aug }} \leftarrow \Delta^{\prime}$
return $\Delta^{\text {aug }}, \mathcal{W}^{\text {aug }}, \mathcal{L}$

```
Algorithm 10 relaxModel \(\left(\Delta^{\text {aug }}, \mathcal{W}^{\text {aug }}, \mathcal{L}\right)\)
Input: an augmented WMI model ( \(\left.\Delta^{\text {aug }}, \mathcal{W}^{\text {aug }}\right), \mathcal{L}\) : equivalence constraints to be relaxed
Output: a relaxed WMI model \(\left(\Delta^{\text {rel }}, \mathcal{W}^{\text {rel }}\right)\), and its "remaining-part" model ( \(\left.\Delta^{\text {rem }}, \mathcal{W}^{\text {rem }}\right)\).
    \(\Delta^{\mathrm{rem}} \leftarrow \top\)
    \(\mathcal{W}^{\text {rem }} \leftarrow\{ \}\)
    \(\Delta^{\text {rel }} \leftarrow \operatorname{copy}\left(\Delta^{\text {aug }}\right)\)
    \(\mathcal{W}^{\text {rel }} \leftarrow \operatorname{copy}\left(\mathcal{W}^{\text {aug }}\right)\)
    for each \(\ell^{*}:\left(X_{i}=X_{i}^{c}\right) \in \mathcal{L}\) do
    for clause \(\Gamma \in \Delta_{i}\) do
        \(\Delta^{\mathrm{rem}} \leftarrow \Delta^{\mathrm{rem}} \wedge \Gamma \wedge \Gamma\left\{X_{i}: X_{i}^{c}\right\}\)
        for each literal \(\ell \in \Gamma\) do
            \(\ell^{\prime} \leftarrow \ell\left\{X_{i}: X_{i}^{c}\right\}\)
            \(w_{\ell^{\prime}} \leftarrow \operatorname{copy}\left(w_{\ell}\right)\)
            \(\mathcal{W}^{\text {rel }} \leftarrow \mathcal{W}^{\text {rel }} \cup\left\{w_{\ell^{\prime}}\right\}\)
            \(\mathcal{W}^{\text {rem }} \leftarrow \mathcal{W}^{\text {rem }} \cup\left\{w_{\ell}, w_{\ell^{\prime}}\right\}\)
    \(\Delta^{\mathrm{rel}} \leftarrow \Delta^{\mathrm{rel}}\left\{\ell^{*}: \top\right\} \quad / /\) disconnect \(X_{i}\) and copy \(X_{i}^{c}\)
    \(\mathcal{W}^{\text {rel }} \leftarrow \mathcal{W}^{\text {rel }} \backslash\left\{w_{\ell^{*}}\right\}\)
    \(\Delta^{\mathrm{rem}} \leftarrow \Delta^{\mathrm{rem}} \wedge \ell^{*}\)
    \(\mathcal{W}^{\mathrm{rem}} \leftarrow \mathcal{W}^{\mathrm{rem}} \cup\left\{w_{\ell^{*}}\right\}\)
return \(\left(\Delta^{\text {rel }}, \mathcal{W}^{\text {rel }}\right),\left(\Delta^{\text {rem }}, \mathcal{W}^{\text {rem }}\right)\)
```


### 7.2.1.2 Removing equivalence constraints.

Given an augmented model ( $\left.\Delta^{\text {aug }}, \mathcal{W}^{\text {aug }}\right)$, we remove equivalence constraints introduced at the augmentation step to obtain the relaxed model $\left(\Delta^{\text {rel }}, \mathcal{W}^{\text {rel }}\right)$. As a result, each original variable in $\mathcal{G}_{\Delta^{\text {rel }}}$ will be detached from its copies, thus ignoring the dependencies encoded by the edges $\mathbb{E}_{d}$ that were marked to be removed. Algorithm 10 details this procedure. Note that relaxation "breaks" the augmented formula $\Delta^{\text {aus }}$ into a relaxed part $\Delta^{\text {rel }}$ and a "remaining part" $\Delta^{\text {rem }}$, which
contains the equivalence constraints just removed.

### 7.2.1.3 Which edges to relax?

After relaxing enough constraints, we can obtain a WMI problem amenable to exact inference, for example, one whose primal graph $\mathcal{G}_{\Delta^{\text {rel }}}$ has treewidth one and logarithmic diameter. Running an exact WMI solver on such a problem would already deliver a cheap way to perform approximate inference. However, the quality of such an approximation can be greatly improved if we compensate for the relaxed constraints. We will discuss this in the next section.

A question remains: how to select the set of edges $\mathbb{E}_{d}$ to relax? Note that the more edges we remove from $\Delta$, the easier it is to perform inference on $\Delta^{\text {rel }}$ given fewer dependencies, but the lower the approximation quality, and the harder to compensate for them all, since it would differ from the augmented model more, and meanwhile from the original model as Proposition 28 indicates. For example, removing all edges in $\mathcal{G}_{\Delta}$ will yield a fully disconnected $\mathcal{G}_{\Delta^{\text {rel }}}$ where performing exact inference on each component is going to be embarrassingly parallelizable. This would correpond to perform a loopy version of the MP-WMI algorithm. Analogous to its discrete counterpart, loopy belief propagation, it would be susceptible to poor converge rates [KF09, CD06]. Therefore we propose a simple strategy for selecting the edges to be removed, which is to retrieve a spanning tree of the original primal graph. In Section 7.3 we demonstrate its practical effectiveness on a range of inference problems of increasing complexity. Devising and evaluating alternative relaxing strategies is an interesting topic for future work.

### 7.2.1.4 Compensation

The aim of the compensation phase is to recover the relaxed equivalence constraints and hence, make the distribution $\operatorname{Pr}_{\Delta^{\text {rel }}}$ better approximate $\operatorname{Pr}_{\Delta^{\text {aug }}}$ and thus better approximate $\operatorname{Pr}_{\Delta}$ as $\operatorname{Propo-}$ sition 28 suggests. In order to do so, we introduce new literals, named compensating literals, to the variables and their copies in the relaxed formula $\Delta^{\text {rel }}$ and equip them with parameterized weights, named compensating weights, and further we optimize them in order to synchronize the variable marginals among a copied variable and its copies.

For each variable $X_{i}=X_{i}^{0}$ and its $C_{i}$ copies $X_{i}^{1}, \ldots, X_{i}^{C_{i}}$ in formula $\Delta^{\text {rel }}$, we generate $K$ different univariate literals of the form $\ell_{i, k}^{c}:\left(X_{i}^{(c)} \leq \sigma_{i, k} \cdot \tau_{i, k}\right)$ for $k=1, \ldots, K$ and $c=$ $0,1, \ldots, C_{i}$ where each $\sigma_{i, k}$ and $\tau_{i, k}$ are respectively drawn at uniform from $\{+1,-1\}$ and the support of $X_{i}$ as encoded in $\Delta_{i}^{\text {rel }}$. Note that the $\sigma_{i, k}, \tau_{i, k}$ are shared across all the copies. Each compensating literal $\ell_{i, k}^{c}$ is therefore responsible for a portion of the support of the marginal distribution of $X_{i}^{c}$, and also for the (unnormalized) marginal density of $X_{i}^{c}$ by equipping it with a parameterized weight $w_{\ell_{i, k}}$.

To retain tractable inference, the parametric function family chosen for each $w_{\ell_{i, k}}$ should satisfy the TWCs as discussed in section 7.1. Striving for simplicity, we employ constant weights of the form $w_{\ell_{i, k}^{c}}:=\exp \left(\theta_{i, k}^{c}\right)$. Therefore, our induced marginal density takes the form of a piecewise constant approximation. As such, by increasing the number of compensating literals $K$ one could obtain a finer approximation, however at the price of introducing more parameters to optimize for. We empirically investigate the effect of increasing $K$ in our experiments in section 7.3.

### 7.2.1.5 Iterative integration

Instead of matching marginal density functions we settle for the weaker condition of matching the marginal probabilities of the newly introduced compensating literals. This in turn can be stated by the following set of equivalence constraints for each variable $X_{i}$ :

$$
\begin{equation*}
\operatorname{Pr}_{\Delta^{\mathrm{rem}}}\left(\bigwedge_{c=0}^{C_{i}} \ell_{k, i}^{c}\right)=\operatorname{Pr}_{\Delta^{\mathrm{rel}}}\left(\ell_{k, i}^{0}\right)=\operatorname{Pr}_{\Delta \mathrm{rel}}\left(\ell_{k, i}^{1}\right)=\cdots=\operatorname{Pr}_{\Delta^{\mathrm{rel}}}\left(\ell_{k, i}^{C_{i}}\right), \text { for } k=1, \cdots, K \tag{7.1}
\end{equation*}
$$

where the first term $\operatorname{Pr}_{\Delta^{\text {rem }}}\left(\bigwedge_{j=0}^{C_{i}} \ell_{k, i}^{c}\right)$ is the probability of the compensating literals in the remaining WMI model $\left(\Delta^{\text {rem }}, \mathcal{W}^{\text {rem }}\right)$ and $\operatorname{Pr}_{\Delta^{\text {rel }}}\left(\ell_{k, i}^{c}\right)$ are the probabilities of compensating literals in the relaxed formula $\Delta^{\text {rel }}$. Intuitively, for a single equivalence constraint that has been relaxed, there exists a set of parameters $\theta$ for the compensating weights that exactly match the probabilities in Equation 7.1 and hence guarantee exact marginal recovery [CD06]. The next theorem better formalizes it.

Algorithm 11 ReCoIn $(\Delta, \mathcal{W}, K)$
Input: a WMI model $(\Delta, \mathcal{W}), K$ number of compensating literals
Output: $\left(\Delta^{\text {rel }}, \mathcal{W}^{\text {rel }}\right):$ a relaxed and compensated WMI model
$: \mathbb{E}_{d} \leftarrow \operatorname{initStrategy}(\Delta, \mathcal{W}) \quad / /$ Select edges to remove
$\Delta^{\text {aug }}, \mathcal{W}^{\text {aug }}, \mathcal{L} \leftarrow \operatorname{augment} \operatorname{Model}\left(\Delta, \mathcal{W}, \mathbb{E}_{d}\right)$
$\left(\Delta^{\text {rel }}, \mathcal{W}^{\text {rel }}\right),\left(\Delta^{\text {rem }}, \mathcal{W}^{\text {rem }}\right) \leftarrow \operatorname{relaxModel}\left(\Delta^{\text {aug }}, \mathcal{W}^{\text {aug }}, \mathcal{L}\right)$
$\Delta^{\text {rel }}, \mathcal{W}^{\text {rel }} \leftarrow$ addingCompensations $\left(\Delta^{\text {rel }}, \mathcal{W}^{\text {rel }}, \mathcal{L}, K\right)$
while not converged do

$$
\text { for } X_{i} \in \operatorname{copiedNodes}\left(\Delta^{\text {rel }}\right) \text { do }
$$

$$
\text { for } k=1, \ldots, K \mathbf{d o}
$$

$$
r^{k} \leftarrow \mathrm{WMI}\left(\Delta^{\mathrm{rem}}, \mathcal{W}^{\mathrm{rem}}\right) / \mathrm{WMI}\left(\Delta^{\mathrm{rem}} \wedge \bigwedge_{c=0}^{C_{i}} \ell_{k, i}^{c}, \mathcal{W}^{\mathrm{rem}}\right)-1
$$

$$
\text { for } c=0,1, \ldots, C_{i} \text { do }
$$

$$
\theta_{k, i}^{c,(t+1)} \leftarrow \log \left(r^{k} \alpha_{k, \sigma(c)}\right)-\log \left(1-\alpha_{k, \sigma(c)}\right)-\sum_{c^{\prime} \neq c} \theta_{k, i}^{c,(t)}
$$

11: $\operatorname{Return}\left(\Delta^{\text {rel }}, \mathcal{W}^{\text {rel }}\right)$

Theorem 16. Suppose that a relaxed model $\left(\Delta^{\mathrm{rel}}, \mathcal{W}^{\mathrm{rel}}\right)$ and a remaining model $\left(\Delta^{\mathrm{rem}}, \mathcal{W}^{\mathrm{rem}}\right)$ are obtained by relaxing a single equivalence constraint $\left(X_{i}=X_{i}^{c}\right)$ from an augmented model $\Delta^{\text {aug, }}$, and that the primal graph of $\Delta^{\text {rel }}$ is split into two disconnected components by the relaxation. Let $\left(\ell_{i, k}, \ell_{i, k}^{c}\right)$ for $k=1, \ldots, K$ be the $K$ pairs of compensating literals introduced, and $\theta_{k, i}, \theta_{k, i}^{c}$, for $k=1, \ldots, K$, be the parameters attached to the compensating weights. Then Equation 7.1 holds when the compensating weight parameters satisfy the following equalities.

$$
\theta_{k, i}=\log \frac{r^{k} \alpha_{k, c}}{1-\alpha_{k, c}}-\theta_{k, i}^{c}, \quad \theta_{k, i}^{c}=\log \frac{r^{k} \alpha_{k}}{1-\alpha_{k}}-\theta_{k, i} \quad \text { for } k=1, \ldots, K
$$

where
$r^{k}=\frac{\operatorname{WMI}\left(\Delta^{\text {rem }} \bigwedge \neg \ell_{k, i} \bigwedge \neg \ell_{k, i}^{c}, \mathcal{W}^{\text {rem }}\right)}{\operatorname{WMI}\left(\Delta^{\text {rem }} \bigwedge \ell_{k, i} \bigwedge \ell_{k, i}^{c}, \mathcal{W}^{\text {rem }}\right)}, \alpha_{k}=\operatorname{Pr}_{\Delta_{\text {rel }}}\left(\ell_{i, k}\right), \alpha_{k, c}=\operatorname{Pr}_{\Delta^{\mathrm{rel}}}\left(\ell_{i, k}^{c}\right)$, for $k=1, \ldots, K$.
Theorem 16 suggests an iterative optimization scheme to find the fixed point solutions for all the compensating parameters introduced to compensate multiple relaxed equivalence constraints. Specifically, starting from a random initialization of the parameters of the compensating weights, ${ }^{1}$ at each iteration $t+1$, we can update each parameter $\theta_{k, i}^{c,(t+1)}$ as

$$
\begin{equation*}
\theta_{k, i}^{c,(t+1)} \leftarrow \log \left(r^{k} \alpha_{k, \pi(c)}\right)-\log \left(1-\alpha_{k, \pi(c)}\right)-\sum_{c^{\prime} \neq c} \theta_{k, i}^{c^{\prime},(t)}, \tag{7.2}
\end{equation*}
$$

where $\pi$ is a permutation over the copies and each $\alpha_{k, \pi(c)}$ is computed as the probability of $\ell_{k, i}^{\pi(c)}$ according to the relaxed model.

Therefore, at each iteration $t$, we need to solve $2 K$ integration problems for computing the $r^{k}$ terms and $C_{i} \cdot K$ integrations for $\operatorname{Pr}_{\Delta^{\text {rel }}}\left(\ell_{k, i}^{\pi(c)}\right)$ for each pair of variable and its copies. While in principle we could use any exact WMI solver to solve these problems, we adopt MP-WMI [ZMY20a] because it is the fastest solver yet for tree-shaped and bounded diameter problems, and even more importantly, it allows to amortize inference across queries. That is, we can compute all the $C_{i} \cdot K$ literal probabilities in a single message-passing step with it.

[^3]

Figure 7.3: Average integrated absolute errors (left) and times in seconds (right) for 5 problems of increasing size ( $n$, x -axis) for ReCoIn and competitors. Number of compensating literals (2-4) or samples used are in parentheses. Mean values per problem size are connected by a line.

From this perspective, our algorithm ReCoIn generates a sequence of induced distributions $\operatorname{Pr}_{\Delta^{\text {rel }}}^{(1)}, \ldots, \operatorname{Pr}_{\Delta^{\text {rel }}}^{(2)}, \operatorname{Pr}_{\Delta^{\text {rel }}}^{(t)}$, that should converge to a fixed-point distribution. In practice to check for convergence, one can monitor the quality of the literal probability approximations and stop when a threshold $\epsilon$ is met before a certain number of iterations are done. We choose the threshold to be the maximum $L-\infty$ norm of compensation literal probability differences. To ease convergence, we apply dampening, that is, we smooth each parameter update at iteration $t+1$ by a factor $\lambda>0: \theta_{k, i}^{c,(t+1)} \leftarrow(1-\lambda) \cdot \theta_{k, i}^{c,(t+1)}+\lambda \cdot \theta_{k, i}^{c,(t+1)}$. This completes the steps in our ReCoIn solver. Algorithm 11 recaps them.

### 7.3 Empirical Evaluation

We aim to answer the following questions: (Q1) how fast and scalable is ReCoIn?, (Q2) how accurate are its approximations?, (Q3) what is the effect of increasing the number of compensating literals $K$ ?

We generate WMI problems whose primal graphs are random Watts-Strogatz graphs [WS98] with increasing size $n=1, \ldots, 11$, with two additional neighbor connections and probability of rewiring 0.5 , to which we attach randomly generated clauses of length 2 and piecewise constant densities. For each setting we generate 5 independent problems.

We run ReCoIn for up to 20 iterations, employing a dampening coefficient $\lambda=0.5$ in two settings that differ by the number of compensating literals $K=2,4$. We compare it against the fastest sampling scheme available, the rejection sampler (REJ) implemented in [KMZ19] and the hybrid solver XSDD(Sampling) [ZKD19] employing sophisticated knowledge-compilation [DM02] techniques [KMS18] to guide sampling. For both REJ and XSDD we employ 100 thousand samples per query.

To compare the quality of approximations for a problem, we compute for a model $\mathcal{M}$ the mean integral absolute error (IAE) as $\frac{1}{|\mathcal{X}|} \sum_{i=1}^{|\mathcal{X}|} \sum_{j=1}^{B}\left|\operatorname{Pr}_{\mathrm{G}}\left(X_{i} \in b_{j}\right)-\operatorname{Pr}_{M}\left(X_{i} \in b_{j}\right)\right|$ where we partition the support for each marginal $i=1, \ldots, n$ into $B$ equal-widths bins $b_{j}$ for $j=1, \ldots, B$ and compare the probability $\operatorname{Pr}_{M}$ according to model $M$ against the ground truth $\operatorname{Pr}_{G}$, which we compute using PA [MPS17]. We employ PA as it is so far the most reliable general-purpose exact WMI solver [ZMY20a]. Note that as such REJ and XSDD are bounded to solve $|\mathcal{X}| \cdot B$ independent WMI problems, while ReCoIn can naturally amortize $|\mathcal{X}| \cdot B$ queries after a single run of optimization (cf. Section 7.2.1.5). We impose a timeout of 1 hour.

Figure 7.3 reports the IAEs and running times (in seconds) for all problems, settings and competitors. Concerning Q1 and Q2, ReCoIn is the best performer overall. The naive sampling strategy in REJ, while being the fastest as expected, cannot exploit the structure in the problem and clearly suffers from the curse of dimensionality. Conversely, XSDD can deliver accurate approximations thanks to compiling the problem structure, but on highly loopy graphs compilation cannot scale beyond $n=5$. On the other hand, ReCoIn gracefully scales to larger problem sizes and multiple queries, and delivers very low IAE scores that are close to the best by XSDD on small problem sizes. Note that while ReCoIn can solve much larger problems within our timeout, we could not retrieve a ground truth for them with PA in reasonable time (more than 24 hours per problem).

Concerning Q3, more compensating literals $(K=4)$ are achieving marginally lower IAEs at the expense of linearly increasing running times. Exploring the time-accuracy trade-off by increasing $K$ or employing different relaxation strategies is an interesting avenue to investigate
in the future. All in all, this empirical evidence candidates RECoIn as one of the best generalpurpose approximate WMI solvers in the current landscape of WMI solvers.

### 7.4 Discussion

In this chapter, we advance the WMI framework by tracing the theoretical requirements for tractable WMI inference with the highest precision so far. We introduced ReCoIn as the first solver that by exploiting our tractability insights can reliably scale approximate inference on general WMI problems. We believe these two contributions can help strengthen our theoretical understanding on the challenges and guarantees around approximate hybrid probabilistic inference and at the same time propel the construction of more efficient and scalable WMI solvers.

## CHAPTER 8

## WMI for Uncertainty Quantification

This chapter shows an application of WMI in quantifying uncertainty. Uncertainty estimation is crucial for decision making. Deep learning models, including those in safety-critical domains, tend to estimate uncertainty poorly. To overcome this issue, Bayesian deep learning obtains a posterior distribution over the model parameters hoping to improve predictions and provide reliable uncertainty estimates. Among Bayesian inference procedures with neural networks, Bayesian model averaging (BMA) is particularly compelling [Was00, FBL18, MIG19]. However, computing BMAs is distinctly challenging since it involves marginalizing over posterior parameters, which possess some unusual topological properties such as mode-connectivity [IVH21]. We show that even with simple low-dimensional approximate parameter posteriors as uniform distributions, doing BMA requires integrating over highly non-convex and multi-modal distributions with discontinuities arising from non-linear activations (cf. Figure 8.1a). Accurately approximating the BMA can achieve significant performance gains [IVH21]. Existing methods mainly focus on general-purpose MCMC, which can fail to converge, or provides inaccurate few-sample predictions [KEH22], because running longer sampling chains is computationally expensive, and variational approaches that typically use a mean-field approximation that ignores correlations induced by activations [JLB22].

In this work, we are interested in developing collapsed samplers, also known as cutset or Rao-Blackwellised samplers for BMA. A collapsed sampler improves over classical particle-based methods by limiting sampling to a subset of variables and further pairing each sample with a closed-form representation of a conditional distribution over the rest whose marginalization is
often tractable. Collapsed samplers are effective at variance reduction in graphical models [KF09], however no collapsed samplers are known for Bayesian deep learning. We believe that this is due to the lack of a closed-form marginalization technique congruous with the non-linearity in deep neural networks. Our aim is to overcome this issue and improve BMA estimation by incorporating exact marginalization over (close approximate) conditional distributions into the inference scheme. Nevertheless, scalability and efficiency are guaranteed by the sampling part of our proposed algorithm.

Marginalization is made possible by our observation that BMA reduces to weighted volume computation. Certain classes of such problems can be solved exactly by so-called weighted model integration (WMI) solvers [BPB15]. By closely approximating BMA with WMI, these solvers can provide accurate approximations to marginalization in BMA (cf. Figure 8.1b). With this observation, we propose Ciber, a collapsed sampler that uses WMI for computing conditional distributions. In the few-sample setting, Ciber delivers more accurate uncertainty estimates than the gold-standard Hamiltonian Monte Carlo (HMC) method (cf. Figure 8.2). We further evaluate the effectiveness of CIber on regression and classification benchmarks and show significant improvements over other Bayesian deep learning approaches in terms of both uncertainty estimation and accuracy.

### 8.1 Bayesian Model Averaging as Weighted Volume Computation

In Bayesian Neural Networks (BNN), given a neural network $f_{\boldsymbol{w}}$ parameterized by weights $\boldsymbol{w}$, instead of doing inference with deterministic $\boldsymbol{w}$ that optimize objectives such as cross-entropy or mean squared error, Bayesian learning infers a posterior $p(\boldsymbol{w} \mid \mathcal{D})$ over parameters $\boldsymbol{w}$ after observing data $\mathcal{D}$. During inference, this posterior distribution is then marginalized to produce final predictions. This process is called Bayesian Model Averaging ( $B M A$ ). It can be seen as learning an ensemble of an infinite number of neural nets and aggregating their results. Formally, given input $\boldsymbol{x}$, the posterior predictive distribution and the expected prediction for a regression task are

$$
\begin{equation*}
p(\boldsymbol{y} \mid \boldsymbol{x})=\int p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{w}) p(\boldsymbol{w} \mid \mathcal{D}) d \boldsymbol{w}, \quad \text { and } \quad \mathbb{E}_{p(\boldsymbol{y} \mid \boldsymbol{x})}[\boldsymbol{y}]=\int \boldsymbol{y} p(\boldsymbol{y} \mid \boldsymbol{x}) d \boldsymbol{y} \tag{8.1}
\end{equation*}
$$



（a）$p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{w})$ being Gaussian．See Example 10.
（b）$p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{w})$ being triangular．See Section 8．2．
Figure 8．1：The integral surface of（a）the expected prediction in BMA，and（b）our proposed ap－ proximation．Both are highly non－convex and multi－modal．The z －axis is the weighted prediction $\boldsymbol{y} p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{w}) p(\boldsymbol{w} \mid \mathcal{D})$ ．Integration of（a）does not admit a closed－form solution，yet integration of（b）is a close approximation that can be solved exactly and efficiently by WMI solvers．

For classification，the（most likely）prediction is the class $\arg \max _{\boldsymbol{y}} p(\boldsymbol{y} \mid \boldsymbol{x})$ ．BMA is intuitively at－ tractive because it can be risky to base inference on a single neural network model．The marginal－ ization in BMA gets around this issue by averaging over models according to a Bayesian posterior．

BMA requires approximations to compute posterior predictive distributions and expected pre－ dictions，as the integrals in Equation 8.1 are intractable in general．Deriving efficient and accu－ rate approximations remains an active research topic［IVH21］．We approach this problem by observing that the marginalization in BMA with ReLU neural networks can be cast as weighted volume computation（WVC）．Later we show that it can be generalized to any neural network when combined with sampling．In WVC，various tools exist for solving certain WVC problem classes［BBD14，KMZ19，ZMY20b］．This section reveals the connection between BMA and WVC． It opens up a new perspective for developing BMA approximations by leveraging WVC tools．

Definition 14 （WVC）．A weighted volume computation（WVC）problem is defined by a pair（四，$\phi$ ） where a region ${ }^{(1)}$ is conjunction of arithmetic constraints and weight $\phi: ⿴ 囗 \mathbb{R}$ is an integrable function assigning weights to elements in The task of WVC is to compute the integral $\int_{\text {四 }} \phi(\boldsymbol{x}) d \boldsymbol{x}$ ．


Figure 8.2: Uncertainty estimates for regression. The red line is the ground truth. The dark blue line shows the predictive mean. The shaded region is the $90 \%$ confidence interval of the predictive distribution. For the same number of samples, (b) CIBER is closer than (a) small-sample HMC to (c) a highly accurate but slow HMC with a large number of samples.

### 8.1.1 A Warm-Up Example

Consider a simple yet relevant setting where the predictive distribution $p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{w})$ is a Dirac delta distribution with zero mass everywhere except at $f_{\boldsymbol{w}}(\boldsymbol{x})$, such that $\int \boldsymbol{y} p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{w}) d \boldsymbol{y}=$ $f_{\boldsymbol{w}}(\boldsymbol{x})$.

Example 9. Assume a model $f_{\boldsymbol{w}}(x)=\operatorname{ReLU}(w \cdot x)$ with a uniform posterior over the parameter: $p(w \mid \mathcal{D})=\frac{1}{6}$ with $w \in[-3,3]$. Let the input be $x=1$. For parameter $w \in[-3,0]$, the model $f_{\boldsymbol{w}}$ always predicts 0 , and otherwise (i.e., $w \in(0,3]$ ), it predicts $w$. Thus, the expected prediction

 integral decomposes into WVC problems with different weights due to the ReLU activation.

These WVC problems have easy closed-form solutions. This is no longer the case in the following.

Example 10. Assume a model $f_{\boldsymbol{w}}(x)$ and posterior distribution $p(w \mid \mathcal{D})$ as in Example 9. Let the predictive distribution $p(y \mid x, w)$ be a Gaussian distribution $p_{\mathcal{N}}\left(y ; f_{\boldsymbol{w}}(x), 1\right)$ with mean $f_{\boldsymbol{w}}(x)$ and variance 1. Given input $x=1$, the expected prediction (Equation 8.1) is

$$
\mathbb{E}_{p(y \mid x=1)}[y]=\int_{\boldsymbol{G}_{\perp}} y \cdot p_{\mathcal{N}}(y \mid 0,1) \cdot \frac{1}{6} d y d w+\int_{\boldsymbol{G}_{\top}} y \cdot p_{\mathcal{N}}(y \mid w, 1) \cdot \frac{1}{6} d y d w
$$

It is a summation of two WVC problems with 四 ${ }_{\perp}=(-3 \leq w \leq 0) \wedge(y \in \mathbb{R})$ and四 ${ }_{\top}=(0 \leq$ $w \leq 3) \wedge(y \in \mathbb{R})$ ，whose joint integral surface is shown in Figure 8．1a．

These WVC problems do not admit closed－form solutions since they involve truncated Gaus－ sian distributions．Moreover，Figure 8．1a shows that computing BMA，even in such a low－dimensional parameter space，requires integration over non－convex and multi－modal functions．

## 8．1．2 General Reduction of BMA to WVC

Let model $f_{\boldsymbol{w}}$ be a ReLU neural net．Denote the set of inputs to its ReLU activations by $\mathcal{R}=$ $\left\{r_{i}\right\}_{i=1}^{R}$ ，where each $r_{i}$ is a linear combination of weights．For a given input $\boldsymbol{x}$ ，the parameter space is partitioned by whether each ReLU activation outputs zero or not．This gives the WVC reduction

$$
p(\boldsymbol{y} \mid \boldsymbol{x})=\sum_{\boldsymbol{B} \in\{0,1\}^{R}} \int_{⿴_{\mathcal{B}}^{B}} p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{w}) p(\boldsymbol{w} \mid \mathcal{D}) d \boldsymbol{w}
$$

where $\boldsymbol{B}$ is a binary vector．The region $山_{B}$ is defined as $\wedge_{i=1}^{R} \ell_{i}$ where arithmetic constraint $\ell_{i}$ is $r_{i} \geq 0$ if $\boldsymbol{B}_{i}=1$ and $r_{i} \leq 0$ otherwise．The expected prediction $\mathbb{E}_{p(\boldsymbol{y} \mid \boldsymbol{x})}[\boldsymbol{y}]$ is analogous but includes an additional factor and variable of integration $y$ in each WVC problem．

This general reduction，however，is undesirable since it amounts to a brute－force enumeration that implies a complexity exponential in the number of ReLU activations．Moreover，not all WVC problems resulting from this reduction are amenable to existing solvers．We will therefore appeal to a framework called weighted model integration（WMI）that allows for a compact representation of these WVC problems，and a characterization of their tractability for WMI solvers［KMZ19］． This inspires us to approximate BMA by first reducing it to WVC problems and further closely approximating those with tractable WMI problems．

## 8．2 Approximating BMA by WMI

WMI is a modeling and inference framework that supports integration in the presence of logical and arithmetic constraints［BPB15，BBP15b］．Various WMI solvers have been proposed in recent
years [KMZ19], ranging from general-purpose ones to others that assume some problem structures to gain scalability. However, even with the reduction from BMA to WVC from the previous section, WMI solvers are not directly applicable. Existing solvers have two main limitations: (i) feasible regions need to be defined by Boolean combinations of linear arithmetic constraints, and (ii) weight functions need to be polynomials. In this section, we show that these issues can be bypassed using a motivating example of how to form a close approximation to BMA using WMI.

In WMI, the feasible region is defined by satisfiability modulo theories (SMT) constraints [BMR10]: an SMT formula is a (typically quantifier-free) expression containing both propositional and theory literals connected with logical connectives; the theory literals are often restricted to linear real arithmetic, where literals are of the form $\left(\mathbf{c}^{T} \mathbf{X} \leq b\right)$ with variable $\mathbf{X}$ and constants $\mathbf{c}^{T}$ and $b$.

Example 11. The ReLU model $f_{\boldsymbol{w}}(x)$ of Example 9 can be encoded as an SMT formula (see box).

$$
\Delta_{\mathrm{ReLU}}=\left\{\begin{array}{l}
W \cdot x>0 \Rightarrow Z=W \cdot x \\
W \cdot x \leq 0 \Rightarrow Z=0
\end{array}\right.
$$

The curly bracket denotes logical conjunction, the symbol $\Rightarrow$ is a logical implication, variable $W$ is the weight, and variable $Z$ denotes the model output.

The encoding of ReLU neural networks into SMT formulas is explored in existing work to enable verification of the behavior of neural networks and provide formal guarantees [KBD17, HKW17, SFM20]. We propose to use this encoding to define the feasible region of WMI problems. Let $\boldsymbol{x} \vDash \Delta$ denote the satisfaction of an SMT formula $\Delta$ by an assignment $\boldsymbol{x}$, and $\llbracket \boldsymbol{x} \vDash \Delta \rrbracket$ be its corresponding indicator function. We formally introduce WMI next.

Definition 15. (WMI) Let $\mathcal{X}$ be a set of continuous random variables. A weighted model integration problem is a pair $\mathcal{M}=(\Delta, \Phi)$, where $\Delta$ is an SMT formula over $\mathcal{X}$ and $\Phi$ is a set of per-literal
weights defined as $\Phi=\left\{\phi_{\ell}\right\}_{\ell \in \mathcal{L}}$, where $\mathcal{L}$ is a set of SMT literals and each $\phi_{\ell}$ is a function defined over variables in literal $\ell$. The task of weighted model integration is to compute

$$
\operatorname{WMI}(\Delta, \Phi)=\int_{\boldsymbol{x} F \Delta} \prod_{\ell \in \mathcal{L}} \phi_{\ell}(\boldsymbol{x})^{\llbracket x \mid=\ell \rrbracket} d \boldsymbol{x} .
$$

That is, the task is to integrate over the weighted assignments of $\mathcal{X}$ that satisfy the SMT formula $\Delta .^{1}$

An approximation to the BMA of Example 10 can be achieved with WMI using the following four steps:

Step 1. Encoding model $f_{\boldsymbol{w}}(x)$. This has been shown as the SMT formula $\Delta_{\text {ReLU }}$ in Example 11.

Step 2. Encoding posterior distribution $p(w \mid \mathcal{D})$. The uniform distribution $p(w \mid \mathcal{D})=\frac{1}{6}$ with $w \in[-3,3]$ can be encoded as a WMI problem pair $\left(\Delta_{p o s}, \Phi_{\text {pos }}\right)$ as follows:

$$
\Delta_{p o s}=-3 \leq W \leq 3 \quad \Phi_{p o s}=\left\{\phi_{\ell}(W)=\frac{1}{6} \text { with } \ell=\text { true }\right\}
$$

Step 3. Approximate encoding of predictive distribution $p(y \mid w, x)$. Recall that $p(y \mid$ $w, x)=p_{\mathcal{N}}\left(y ; f_{\boldsymbol{w}}(x), 1\right)$ is Gaussian, which cannot be handled by existing WMI solvers. To approximate it with polynomial densities, we simply use a triangular distribution encoded as a WMI problem pair:

$$
\Delta_{\text {pred }}=\left\{\begin{array}{l}
Y \leq Z+\alpha \\
Y \geq Z-\alpha
\end{array} \quad \Phi_{\text {pred }}=\left\{\begin{array}{l}
\phi_{\ell_{1}}(Y, Z)=\frac{1-Y+Z}{\alpha} \text { with } \ell_{1}=Y \geq Z \\
\phi_{\ell_{2}}(Y, Z)=\frac{1+Y-Z}{\alpha} \text { with } \ell_{2}=Y<Z
\end{array}\right\}\right.
$$

In this encoding, $\alpha$ is a constant that defines the shape of the triangular distribution. It is obtained by minimizing the $L 2$ distance between a standard normal distribution and the symmetric triangular distribution. We visualize this approximation in Figure 8.3.

[^4]

Figure 8.3: Approximating the Gaussian distribution with a triangluar distribution.

Step 4. Approximating BMA by calling WMI solvers. With the above encodings, the predictive posterior $p(y \mid x)$ (Equation 8.1) can be computed using two calls to a WMI solver. For example, the uncertainty of a prediction $y=1$ for input $x=1$ is

$$
p(y=1 \mid x=1)=\mathrm{WMI}(\Delta \wedge(Y=1), \Phi) / \operatorname{WMI}(\Delta, \Phi)=0.164 / 1
$$

where $\Delta=\Delta_{\text {ReLU }} \wedge \Delta_{\text {pos }} \wedge \Delta_{\text {pred }}$ and $\Phi=\Phi_{\text {pos }} \cup \Phi_{\text {pred. }}$. Similarly, the expected prediction $\mathbb{E}_{p(y \mid x=1)}[y]$ (Equation 8.1) can be computed using two calls to a WMI solver:

$$
\mathbb{E}_{p(y \mid x=1)}[y]=\operatorname{WMI}\left(\Delta, \Phi^{*}\right) / \operatorname{WMI}(\Delta, \Phi)=0.752 / 1
$$

where $\Phi^{*}=\Phi \cup\left\{\phi_{\ell}(Y)=Y\right.$ with $\ell=$ true $\}$. The above formulations also work for unnormalized distributions since the WMI in the denominator serves to compute the partition function.

We visualize the integral surface of the resulting approximate BMA problem in Figure 8.1b. It is very close to the integral surface of the original BMA problem in Figure 8.1a. However, it can be exactly integrated using existing WMI solvers while the original one does not admit such solutions. Next, we show how this process can be generalized to a scalable and accurate approximation of BMA.

### 8.3 Ciber: Collapsed Inference for Bayesian Deep Learning via WMI

Given a BNN with a large number of weights, naively approximating it by WMI problems can lead to computational issues, since it involves doing integration over polytopes in arbitrarily high
dimensions and this is known to be \#P-hard [Val79, DDK12, ZMY20b]. Further, weights involved with non-ReLU activation might not be amenable to the WMI encoding. To tackle these issues, we propose to use collapsed samples to combine the strengths from two worlds: the scalability and flexibility from sampling and the accuracy from WMI solvers.

Definition 16. (Collapsed BMA) Let $\left(\mathcal{W}_{s}, \mathcal{W}_{c}\right)$ be a partition of parameters $\mathcal{W}$. A collapsed sample is a tuple $\left(\boldsymbol{w}_{s}, q\right)$, where $\boldsymbol{w}_{s}$ is an assignment to the sampled parameters $\mathcal{W}_{s}$ and $q$ is a representation of the conditional posterior $p\left(\mathcal{W}_{c} \mid \boldsymbol{w}_{s}, \mathcal{D}\right)$ over the collapsed parameter set $\mathcal{W}_{c}$. Given collapsed samples $\mathcal{S}$, collapsed BMA estimates the predictive posterior and expected prediction as

$$
\begin{align*}
p(\boldsymbol{y} \mid \boldsymbol{x}) & \approx \frac{1}{|\mathcal{S}|} \sum_{\left(\boldsymbol{w}_{s}, q\right) \in \mathcal{S}}\left[\int p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{w}) q\left(\boldsymbol{w}_{c}\right) d \boldsymbol{w}_{c}\right], \text { and } \\
\mathbb{E}_{p(\boldsymbol{y} \mid \boldsymbol{x})}[\boldsymbol{y}] & \approx \frac{1}{|\mathcal{S}|} \sum_{\left(\boldsymbol{w}_{s}, q\right) \in \mathcal{S}}\left[\int \boldsymbol{y} p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{w}) q\left(\boldsymbol{w}_{c}\right) d \boldsymbol{w}_{c} d \boldsymbol{y}\right] . \tag{8.2}
\end{align*}
$$

The size of the collapsed set $\mathcal{W}_{c}$ determines the trade-off between scalability and accuracy. The more parameters in the collapsed set, the more accurate the approximation to BMA is. The fewer parameters in $\mathcal{W}_{c}$, the more efficient the computations of the integrals are since the integration is performed in a lower-dimensional space. Later in our experiments, we choose a subset of weights at the last or second-to-last hidden layer of the neural networks to be the collapsed set. This choice is known to be effective in capturing uncertainty as shown in [KHH20, SRS15].

To develop an algorithm to compute collapsed BMA, we are faced with two main design choice questions: (Q1) how to sample $\boldsymbol{w}_{s}$ from the posterior? (Q2) what should be the representation of the conditional posterior $q$ such that the integrals in Equation 8.2 can be computed exactly? Next, we provide our answers to these two questions that together give our proposed algorithm Ciber.

### 8.3.1 Approximation to Posteriors

For (Q1), we follow [MIG19] and sample from the stochastic gradient descent (SGD) trajectory after convergence and use the information contained in SGD trajectories to efficiently approxi-
mate the posterior distribution over the parameters of the neural network, leveraging the interpretation of SGD as approximate Bayesian inference [MHB17, CLT20]. Given a set of parameter samples $\mathcal{W}$ from the SGD trajectory, the sample set is defined as $\mathcal{W}_{s}=\left\{\boldsymbol{w}_{s} \mid \boldsymbol{w} \in \mathcal{W}\right\}$. For each assignment $\boldsymbol{w}_{s}$, an approximation $q\left(\mathcal{W}_{c}\right)$ to the conditional posterior $p\left(\mathcal{W}_{c} \mid \boldsymbol{w}_{s}, \mathcal{D}\right)$ is necessary since the posteriors induced by SGD trajectories are implicit. Next, we discuss the choice of approximation to the conditional posterior that is amenable to WMI.

### 8.3.2 Encoding into WMI Problems

As shown in Section 8.2, if a BNN can be encoded as a WMI problem, the posterior predictive distribution and the expected prediction, which involve marginalization over the parameter space, can be computed exactly using WMI solvers. This inspires us to use the WMI framework as the closed-form representation for the conditional posteriors of parameters. The main challenge is how to approximate the integrand in Equation 8.2 using an SMT formula and a polynomial weight function in order to obtain a WMI problem amenable to existing solvers.

For the conditional posterior approximation $q\left(\mathcal{W}_{c}\right)$, we choose it to be a uniform distribution that can be encoded into a WMI problem as $\mathcal{M}_{\text {pos }}=\left(\Delta_{\text {pos }}, \Phi_{p o s}\right)$ with the SMT formula being $\Delta_{\text {pos }}=\wedge_{i \in c}\left(l_{i} \leq W_{i} \leq u_{i}\right)$ and weights being $\Phi_{p o s}=\left\{\phi_{\ell}\left(\mathcal{W}_{c}\right)=1 \mid \ell=\right.$ true $\}$, where $l_{i}$ and $u_{i}$ are domain lower and upper bounds for the uniform distribution respectively. While seemingly over-simplistic, this choice of approximation to the conditional posterior is sufficient to robustly deliver surprisingly strong empirical performance as shown in Section 8.4. The intuition is that uniform distributions are better than a few samples. We further illustrate this point by comparing the predictive distributions of Ciber and HMC in a few-sample setting. Figure 8.2 shows that even with the same 10 samples drawn from the posterior distribution, since Ciber further approximates the 10 samples with a uniform distribution, it yields a predictive distribution closer to the ground truth than HMC, indicating that using a uniform distribution instead of a few samples forms a better approximation.

For the choice of predictive distribution $p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{w})$, we propose to use piecewise polynomial
densities. Common predictive distributions can be approximated by polynomials up to arbitrary precision in theory by the Stone-Weierstrass theorem [De 59]. For regression, the de facto choice is Gaussian and we propose to use triangular distribution as the approximation, i.e., $p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{w})=$ $\frac{1}{r}-\frac{1}{r^{2}}\left|\boldsymbol{y}-f_{\boldsymbol{w}}(\boldsymbol{x})\right|$, with domain $\left|\boldsymbol{y}-f_{\boldsymbol{w}}(\boldsymbol{x})\right| \leq r$, and $r:=\alpha \sqrt{\sigma^{2}(\boldsymbol{x})}$ where the constant $\alpha$ parameterizes the triangular distribution as described in Section 8.2. Here, $\sigma^{2}(\boldsymbol{x})$ is the variance estimate, which can be a function of input $\boldsymbol{x}$ depending on whether the BNN is homoscedastic or heteroscedastic. Then $p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{w})$ can be encoded into WMI as:

$$
\Delta_{\text {pred }}=\left\{\begin{array}{l}
Y-f_{\boldsymbol{w}}(\boldsymbol{x}) \leq r \\
Y-f_{\boldsymbol{w}}(\boldsymbol{x}) \geq-r
\end{array} \quad \Phi_{\text {pred }}=\left\{\begin{array}{l}
\phi_{\ell_{1}}\left(Y, \mathcal{W}_{c}\right)=\frac{1}{r}-\frac{Y-f_{\boldsymbol{w}}(\boldsymbol{x})}{r^{2}} \text { with } \ell_{1}=\left(Y>f_{\boldsymbol{w}}(\boldsymbol{x})\right) \\
\phi_{\ell_{2}}\left(Y, \mathcal{W}_{c}\right)=\frac{1}{r}-\frac{f_{\boldsymbol{w}}(\boldsymbol{x})-Y}{r^{2}} \text { with } \ell_{2}=\left(f_{\boldsymbol{w}}(\boldsymbol{x})>Y\right)
\end{array}\right\}\right.
$$

Similar piecewise polynomial approximations are adopted for classification tasks when the predictive distributions are induced by softmax functions.

### 8.3.3 Exact Integration in Collapsed BMA

By encoding the collapsed BMA into WMI problems, we are ready to answer (Q2), i.e., how to perform exact computation of the integrals shown in Equation 8.2.

Proposition 29. Let the SMT formula $\Delta=\Delta_{\text {ReLU }} \wedge \Delta_{\text {pos }} \wedge \Delta_{\text {pred }}$, and the set of weights $\Phi=$ $\Phi_{\text {pos }} \cup \Phi_{\text {pred }}$ as defined in Section 8.3.2. Let the set of weights $\Phi^{*}=\Phi \cup\left\{\phi_{\ell}(Y)=Y\right.$ with $\left.\ell=t r u e\right\}$. The integrals in collapsed BMA (Equation 8.2) can be computed by WMI solvers as

$$
\begin{aligned}
\int p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{w}) q\left(\boldsymbol{w}_{c}\right) d \boldsymbol{w}_{c} & =\operatorname{WMI}(\Delta \wedge(\boldsymbol{Y}=\boldsymbol{y}), \Phi) / \operatorname{WMI}(\Delta, \Phi), \text { and } \\
\int \boldsymbol{y} p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{w}) q\left(\boldsymbol{w}_{c}\right) d \boldsymbol{w}_{c} d \boldsymbol{y} & =\operatorname{WMI}\left(\Delta, \Phi^{*}\right) / \operatorname{WMI}(\Delta, \Phi)
\end{aligned}
$$

With both questions (Q1) and (Q2) answered, we summarize our proposed algorithm Ciber in Algorithm 12. To quantitatively analyze how close the approximation delivered by Ciber is to the ground-truth BMA, we consider the following experiments with closed-form BMA.

Regression. We consider a Bayesian linear regression setting where exact sampling from the posterior distribution is available. Both the likelihood and the weight posterior are Gaussian

## Algorithm 12 CIBER

Input: input $\boldsymbol{x}$, sampled weights $\mathcal{W}$, neural network model $f_{\boldsymbol{w}}$, prediction ground truth $y^{*}$
Ouput: predictions and likelihoods
1: Choose a partition $\left(\mathcal{W}_{s}, \mathcal{W}_{c}\right)$ for network parameters
Derive approximate posterior $q\left(\boldsymbol{w}_{c}\right)$ from sampled weights $\left\{\boldsymbol{w}_{c} \mid \boldsymbol{w} \in \mathcal{W}\right\}$ // cf. Section 8.3.2

Encode posterior $q\left(\boldsymbol{w}_{c}\right)$ into WMI problem $\mathcal{M}_{\text {pos }}=\left(\Delta_{\text {pos }}, \Phi_{\text {pos }}\right) \quad / / c f$. Section 8.3.2
$\mathcal{Y} \leftarrow \emptyset, \mathcal{P} \leftarrow \emptyset \quad / /$ Initialization
for sample $\boldsymbol{w}_{s}$ in $\left\{\boldsymbol{w}_{s} \mid \boldsymbol{w} \in \mathcal{W}\right\}$ do
Encode neural network model $f_{\text {ReLU }}$ parameterized by $\left(\boldsymbol{w}_{s}, \mathcal{W}_{c}\right)$ into an SMT formula $\Delta_{f_{w}}$
Encode predictive $p\left(Y \mid \boldsymbol{x}, \boldsymbol{w}_{s}, \mathcal{W}_{c}\right)$ into a WMI problem $\mathcal{M}_{\text {pred }}=\left(\Delta_{\text {pred }}, \Phi_{\text {pred }}\right)$
SMT formula $\Delta \leftarrow \Delta_{\text {ReLU }} \wedge \Delta_{\text {pos }} \wedge \Delta_{\text {pred }}$
Weights $\Phi \leftarrow \Phi_{\text {pos }} \cup \Phi_{\text {pred }}$
Weights $\Phi^{*} \leftarrow \Phi \cup\left\{\phi_{\ell}(Y)=Y\right.$ with $\ell=$ true $\}$
Add prediction $y=\operatorname{WMI}\left(\Delta, \Phi^{*}\right) / \operatorname{WMI}(\Delta, \Phi)$ to prediction set $\mathcal{Y} \quad / / c f$. Section 8.3.3
Add likelihood $p=\operatorname{WMI}\left(\Delta \wedge\left(Y=y^{*}\right), \Phi\right) / \operatorname{WMI}(\Delta, \Phi)$ to set $\mathcal{P} \quad / /$ cf. Section 8.3.3
return $y=\operatorname{Mean}(\mathcal{Y}), p\left(y^{*} \mid \boldsymbol{x}\right)=\operatorname{Mean}(\mathcal{P})$
such that the ground-truth posterior predictive distribution is Gaussian as well. With samples drawn from the weight posterior, Ciber approximates the samples with a uniform distribution as posterior $p(\boldsymbol{w} \mid \mathcal{D})$ and further approximates the likelihood with a triangular distribution such that the integral $p(\boldsymbol{y} \mid \boldsymbol{x}, \mathcal{D})=\int p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{w}) p(\boldsymbol{w} \mid \mathcal{D}) d \boldsymbol{w}$ can be computed exactly by WMI.

We first evaluate the posterior predictive distribution estimated by Ciber and Monte Carlo (MC) method, using the same five samples drawn from the weight posterior. Results averaged over 10 trials are shown in Figure 8.4 where the estimations by Ciber are much closer to the ground truth posterior predictive distribution than those by the MC method. Further, the averaged KL divergence between the ground truth and Ciber is 0.069 while the one for MC estimations is 0.130 ,


Figure 8.4: Posterior predictive distributions in Bayesian linear regression. The $y$-axis shows the absolute difference between an estimated predictive distribution $p(y \mid \mathbf{x})$ and the ground-truth predictive distribution $q(y \mid \mathbf{x})$. Shaded regions are the $95 \%$ confidence interval.


Figure 8.5: KL divergence in Bayesian linear regression. The $x$-axis shows the number of samples the MC method uses for estimations, ranging from 50 to 150 . The blue curve shows the MC method, and the green dashed curve shows Ciber using 50 samples.
again indicating that CIBER yields a better BMA approximation in the few-sample setting.
We further explore the question of how many samples the MC method needs to match the performance of Ciber. The performances of both approaches are evaluated using KL divergence between the ground-truth posterior distribution and the estimated one, averaged over 10 trials. The result is shown in Figure 8.5 where the dashed green line shows the performance of Ciber with 50 samples and the blue curve shows the performance of MC with an increasing number of samples. As expected, the MC method yields lower KL divergence as the number of samples increases; however, it takes more than 100 samples to match Ciber, indicating its low sample
efficiency and that developing efficient and effective inference algorithms such as Ciber for estimating BMA is a meaningful question.

Classification. For analyzing classification performance, [KEH22] propose to compute the integral $I=\int \sigma\left(f_{*}\right) p_{\mathcal{N}}\left(f_{*}\right) d f_{*}$ with $\sigma$ being the sigmoid function and $f_{*}=f\left(\boldsymbol{x}^{*} ; \boldsymbol{w}\right)$ that amounts to the posterior predictive distribution. We consider a simple case with $f(\boldsymbol{x} ; \boldsymbol{w})=\boldsymbol{w} \cdot \boldsymbol{x}$ such that the ground-truth integral can be obtained. With a randomly chosen $\boldsymbol{x}$, the ground-truth integral is $I=0.823$. The integral estimated by Ciber is $I_{C}=0.826$ while the MC estimate is $I_{M C}=0.732$. That is, Ciber gives an estimate with a much lower error than the MC estimation error, indicating that CIber is able to deliver high-quality approximations in classification tasks.

## Related Work

Bayesian Deep Learning. Bayesian inference over deep neural networks [Mac92] is proposed to fix the issue that deep learning models give poor uncertainty estimations and suffer from overconfidence [NYC15, HAB19, MTP23, MTS21]. Some methods use samples from SGD trajectories to approximate the implicit true posteriors similar to us: [IMK20] (SI) proposes to perform Bayesian inference in a subspace of the parameter space spanned by a few vectors derived from principal component analysis (PCA+ESS(SI)) or variational inference (PCA+VI(SI)); SWAG [MIG19] proposes to approximate the full parameter space using an approximate Gaussian posterior whose mean and covariance are from a partial SGD trajectory with a modified learning rate scheduler.

Some other approaches using approximate posteriors include MC Dropout (MCD) [GG15, GG16] which is one of the Bayesian dropout methods and recently, one of its modifications called Variational Structured Dropout (VSD) [NNN21] using variational inference is proposed. Other state-of-the-art approximate BNN inference methods including deterministic variational inference (DVI) [WNM19], deep Gaussian processes (DGP) [BHH16] with Gaussian process layers and variational inference (VI) [KW13]. Closely related to DGP is the deep kernel process [AYO21] that writes DGPs as deep Wishart processes.

WMI Solvers. WMI generalizes weighted model counting (WMC) [SBK05], a state-of-theart inference approach in many discrete probabilistic models, from discrete to mixed discretecontinuous domains [BPB15, BBP15b]. Recent research on WMI includes its tractability [ZMY20b, ZMY21a, ACD20] and the advancements in WMI solver development. Existing exact WMI solvers for arbitrarily structured problems include DPLL-based search with numerical [BPB15, MPS17, MPS19] or symbolic integration [SOG16] and compilation-based algorithms [KMS18, ZDD19, DHD20] that use extended algebraic decision diagrams (XADDs) [SDD12] as a compilation target which is a powerful tool for inference on mixed domains [SA12, ZSF12]. Some exact WMI solvers aiming to improve efficiency for a certain class of models are proposed such as SMI [ZB19] and MP-WMI [ZMY20a] which are greatly scalable for WMI problems that satisfy certain structural constraints. Approximate solvers are also proposed including sampling-based ones [Zui20] and relaxation-based ones [ZMY20b, ZMY20c]. Recent WMI efforts converge in the pywmi library [KMZ19]. The SMT formulas considered in this work can be seen as distributional constraints on continuous domains. There is also plenty of work in neuro-symbolic AI exploring the integration of discrete constraints into neural networks models including the architectures [ATC22a, AZN23c] and the loss [XZF18, AWC22, ALT22].

### 8.4 Empirical Evaluation

We conduct experimental evaluations of our proposed approach CIber ${ }^{1}$ on regression and classification benchmarks and compare its performance on uncertainty estimation as well as prediction accuracy with a wide range of baseline methods.

### 8.4.1 Regression on Small and Large UCI Datasets

We experiment on 5 small UCI datasets: boston, concrete, yacht, naval and energy. We follow the setup of [IMK20] and use a fully connected network with a single hidden layer and 50 units with ReLU activations. We further experiment on 6 large UCI datasets: elevators, keggdirected,

[^5]Table 8.1: Average test log likelihood for the small UCI regression task.

|  | BOSTON | ConCRETE | YACHT | NAVAL | ENERGY |
| :--- | :--- | :--- | :--- | :--- | :--- |
| CIBER (SECOND) | $-\mathbf{- 2 . 4 7 1} \pm \mathbf{0 . 1 4 0}$ | $-2.975 \pm 0.102$ | $-0.678 \pm 0.301$ | $7.276 \pm 0.532$ | $\underline{\mathbf{- 0 . 7 1 6} \pm \mathbf{0 . 2 1 1}}$ |
| CIBER (LAST) | $\mathbf{- 2 . 4 7 1} \pm \mathbf{0 . 1 4 0}$ | $\underline{\mathbf{- 2 . 9 5 9} \pm \mathbf{0 . 1 0 9}}$ | $-0.687 \pm 0.301$ | $\underline{\mathbf{7 . 4 8 2} \pm \mathbf{0 . 1 8 8}}$ | $\underline{\mathbf{- 0 . 7 1 6} \pm \mathbf{0 . 2 1 1}}$ |
| SWAG | $-2.761 \pm 0.132$ | $-3.013 \pm 0.086$ | $-0.404 \pm 0.418$ | $6.708 \pm 0.105$ | $-1.679 \pm 1.488$ |
| PCA+ESS (SI) | $-2.719 \pm 0.132$ | $-3.007 \pm 0.086$ | $\underline{\mathbf{- 0 . 2 2 5} \pm \mathbf{0 . 4 0 0}}$ | $6.541 \pm 0.095$ | $-1.563 \pm 1.243$ |
| PCA+VI (SI) | $-2.716 \pm 0.133$ | $-2.994 \pm 0.095$ | $-0.396 \pm 0.419$ | $6.708 \pm 0.105$ | $-1.715 \pm 1.588$ |
| SGD | $-2.752 \pm 0.132$ | $-3.178 \pm 0.198$ | $-0.418 \pm 0.426$ | $6.567 \pm 0.185$ | $-1.736 \pm 1.613$ |
| DVI | $-2.410 \pm 0.020$ | $-3.060 \pm 0.010$ | $-0.470 \pm 0.030$ | $6.290 \pm 0.040$ | $-1.010 \pm 0.060$ |
| DGP | $\underline{-2.330 \pm 0.060}$ | $-3.130 \pm 0.030$ | $-1.390 \pm 0.140$ | $3.600 \pm 0.330$ | $-1.320 \pm 0.030$ |
| VI | $-2.430 \pm 0.030$ | $-3.040 \pm 0.020$ | $-1.680 \pm 0.040$ | $5.870 \pm 0.290$ | $-2.380 \pm 0.020$ |
| MCD | $-2.400 \pm 0.040$ | $-2.970 \pm 0.020$ | $-1.380 \pm 0.010$ | $4.760 \pm 0.010$ | $-1.720 \pm 0.010$ |
| VSD | $-2.350 \pm 0.050$ | $-2.970 \pm 0.020$ | $-1.140 \pm 0.020$ | $4.830 \pm 0.010$ | $-1.060 \pm 0.010$ |

Table 8.2: Average test RMSE for the small UCI regression task.

|  | Boston | ConCRETE | YACHT | NAVAL | ENERGY |
| :--- | :--- | :--- | :--- | :--- | :--- |
| CIBER (SECOND) | $3.488 \pm 1.123$ | $4.880 \pm 0.506$ | $0.828 \pm 0.241$ | $\underline{\mathbf{0 . 0 0 0} \pm \mathbf{0 . 0 0 0}}$ | $\underline{\mathbf{0 . 4 4 7} \pm \mathbf{0 . 0 8 1}}$ |
| CIBER (LAST) | $3.478 \pm 1.128$ | $\mathbf{4 . 8 5 4} \pm \mathbf{0 . 5 0 3}$ | $\mathbf{0 . 7 5 2} \pm \mathbf{0 . 2 9 4}$ | $\underline{\mathbf{0 . 0 0 0} \pm \mathbf{0 . 0 0 0}}$ | $\underline{\mathbf{0 . 4 4 7} \pm \mathbf{0 . 0 8 1}}$ |
| SWAG | $3.517 \pm 0.981$ | $5.233 \pm 0.417$ | $0.973 \pm 0.375$ | $0.001 \pm 0.000$ | $1.594 \pm 0.273$ |
| PCA+ESS (SI) | $3.453 \pm 0.953$ | $5.194 \pm 0.448$ | $0.972 \pm 0.375$ | $0.001 \pm 0.000$ | $1.598 \pm 0.274$ |
| PCA+VI (SI) | $\mathbf{3 . 4 5 7} \pm \mathbf{0 . 9 5 1}$ | $5.142 \pm 0.418$ | $0.973 \pm 0.375$ | $0.001 \pm 0.000$ | $1.587 \pm 0.272$ |
| SGD | $3.504 \pm 0.975$ | $5.194 \pm 0.446$ | $0.973 \pm 0.374$ | $0.001 \pm 0.000$ | $1.602 \pm 0.275$ |
| MCD | $2.830 \pm 0.170$ | $4.930 \pm 0.140$ | $0.720 \pm 0.050$ | $\underline{0.000 \pm 0.000}$ | $1.080 \pm 0.030$ |
| VSD | $\underline{2.640 \pm 0.170}$ | $\underline{4.720 \pm 0.110}$ | $\underline{0.690 \pm 0.060}$ | $\underline{0.000 \pm 0.000}$ | $0.470 \pm 0.010$ |

keggundirected, pol, protein and skillcraft. We use a feedforward network with five hidden layers of sizes $[1000,1000,500,50,2]$ and ReLU activations on all datasets except skillcraft. For skillcraft, a smaller architecture is adopted with four hidden layers of size [1000, 500, 50, 2]. All models have two outputs for the prediction and the heteroscedastic variance respectively.

We run Ciber with two different ways of choosing the collapsed parameter set: CIBER (last) chooses all the weights at the last layer to be the collapsed set; CIBER (second) chooses three out of all the weights at the second-to-last layer to be the collapsed set. The heuristic we use for choosing the weights is to look into the sampled weights from SGD trajectories to see which ones have the greatest variance. The intuition is that a greater variance indicates that the weight is prone to have greater uncertainty and thus one might want to perform a more accurate inference over it.

Baselines. We compare CIBER to the state-of-the-art approximate BNN inference methods. We separate these methods into two categories: those sampling from SGD trajectories as approximate posteriors, which includes SWAG [MIG19], PCA+ESS (SI) and PCA+VI (SI) [IMK20], vs.

Table 8.3: Average test log likelihood for the large UCI regression task.

|  | Elevators | KegGD | KeGGU | Protein | Skillcraft | Pol |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| CIBER (SECOND) | $-0.378 \pm 0.026$ | $\underline{\mathbf{1 . 2 4 5} \pm \mathbf{0 . 0 9 0}}$ | $\underline{\mathbf{1 . 1 2 5} \pm \mathbf{0 . 2 6 9}}$ | $-0.720 \pm 0.036$ | $-1.003 \pm 0.035$ | $\underline{\mathbf{2 . 5 5 5} \pm \mathbf{0 . 1 1 5}}$ |
| CIBER (LAST) | $-0.371 \pm 0.023$ | $1.178 \pm 0.088$ | $0.964 \pm 0.231$ | $-0.720 \pm 0.036$ | $\underline{\mathbf{- 1 . 0 0 1} \pm \mathbf{0 . 0 3 2}}$ | $2.506 \pm 0.150$ |
| SWAG | $-0.374 \pm 0.021$ | $1.080 \pm 0.035$ | $0.749 \pm 0.029$ | $\underline{\mathbf{- 0 . 7 0 0} \pm \mathbf{0 . 0 5 1}}$ | $-1.180 \pm 0.033$ | $1.533 \pm 1.084$ |
| PCA+ESS (SI) | $-0.351 \pm 0.030$ | $1.074 \pm 0.034$ | $0.752 \pm 0.025$ | $-0.734 \pm 0.063$ | $-1.181 \pm 0.033$ | $-0.185 \pm 2.779$ |
| PCA+VI (SI) | $\underline{\mathbf{- 0 . 3 2 5} \pm \mathbf{0 . 0 1 9}}$ | $1.085 \pm 0.031$ | $0.757 \pm 0.028$ | $-0.712 \pm 0.057$ | $-1.179 \pm 0.033$ | $1.764 \pm 0.271$ |
| SGD | $-0.538 \pm 0.108$ | $1.012 \pm 0.154$ | $0.602 \pm 0.224$ | $-0.854 \pm 0.085$ | $-1.162 \pm 0.032$ | $1.073 \pm 0.858$ |
| ORTHVGP | -0.448 | 1.022 | 0.701 | -0.914 | - | 0.159 |
| NL | $-0.698 \pm 0.039$ | $0.935 \pm 0.265$ | $0.670 \pm 0.038$ | $-0.884 \pm 0.025$ | $-1.002 \pm 0.050$ | $-2.840 \pm 0.226$ |

those who do not, which includes the SGD baseline, deterministic variational inference baseline (DVI) [WNM19], Deep Gaussian Processes (DGP) [BHH16], variational inference (VI) [KW13], MC Dropout (MCD) [GG15, GG16], and variational structured dropout (VSD) [NNN21]. These methods achieved state-of-the-art performance on the small UCI datasets. We also compare to baselines Bayesian final layers (NL) [RTS18], deep kernel learning (DKL) [WHS16], orthogonally decoupled variational GPs (OrthVGP) [SCB18] and Fastfood approximate kernels (FF) [YWS15], which have achieved state-of-the-art performance on the large UCI datasets.

Results. We present the test log likelihoods for small UCI datasets in Table 8.1 and those for large UCI datasets in Table 8.3. In both tables, the first block summarizes SGD-trajectory sampling-based approaches and the second summarizes the rest. Underlined results are the best among all and bold results are the best among SGD-trajectory sampling-based approaches. From the results, our Ciber has substantially better performance than all others on three out of the five small UCI datasets four out of six large UCI datasets, with comparable performance on the rest, demonstrating that Ciber provides accurate uncertainty estimation. We also present the test rooted-mean-squared error results, where CIBER outperforms all other SGD-trajectory samplingbased baselines on four out of five small UCI datasets and four out of six large UCI datasets; it

Table 8.4: Average test RMSE for the large UCI regression task.

|  | ELEVATORS | KEGGD | KeGGU | Protein | Skillcraft | Pol |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| Ciber (SECOND) | $\mathbf{0 . 0 8 8} \pm \mathbf{0 . 0 0 2}$ | $0.142 \pm 0.074$ | $\mathbf{0 . 1 1 5} \pm \mathbf{0 . 0 0 7}$ | $0.438 \pm 0.009$ | $\mathbf{0 . 2 5 1} \pm \mathbf{0 . 0 1 0}$ | $2.212 \pm 0.230$ |
| Ciber (LAST) | $\mathbf{0 . 0 8 8} \pm \mathbf{0 . 0 0 2}$ | $0.142 \pm 0.072$ | $0.118 \pm 0.012$ | $0.438 \pm 0.009$ | $\mathbf{0 . 2 5 1} \pm \mathbf{0 . 0 1 0}$ | $\underline{\mathbf{2 . 1 9 9} \pm \mathbf{0 . 1 8 2}}$ |
| SWAG | $\mathbf{0 . 0 8 8} \pm \mathbf{0 . 0 0 1}$ | $0.129 \pm 0.029$ | $0.160 \pm 0.043$ | $\underline{\mathbf{0 . 4 1 5} \pm \mathbf{0 . 0 1 8}}$ | $0.293 \pm 0.015$ | $3.110 \pm 0.070$ |
| PCA+ESS (SI) | $0.089 \pm 0.002$ | $0.129 \pm 0.028$ | $0.160 \pm 0.043$ | $0.425 \pm 0.017$ | $0.293 \pm 0.015$ | $3.755 \pm 6.107$ |
| PCA+VI (SI) | $\mathbf{0 . 0 8 8} \pm \mathbf{0 . 0 0 1}$ | $\mathbf{0 . 1 2 8} \pm \mathbf{0 . 0 2 8}$ | $0.160 \pm 0.043$ | $0.418 \pm 0.021$ | $0.293 \pm 0.015$ | $2.499 \pm 0.684$ |
| SGD | $0.103 \pm 0.035$ | $0.132 \pm 0.017$ | $0.186 \pm 0.034$ | $0.436 \pm 0.011$ | $0.288 \pm 0.014$ | $3.900 \pm 6.003$ |
| NL | $0.101 \pm 0.002$ | $0.134 \pm 0.036$ | $0.120 \pm 0.003$ | $0.447 \pm 0.012$ | $0.253 \pm 0.011$ | $4.380 \pm 0.853$ |
| DKL | $\underline{0.084 \pm 0.020}$ | $\underline{0.100 \pm 0.010}$ | $\underline{0.110 \pm 0.000}$ | $0.460 \pm 0.010$ | $\underline{0.250} \pm 0.000$ | 6.617 |
| ORTHVGP | 0.095 | 0.120 | 0.117 | 0.461 | - | $4.300 \pm 0.200$ |
| FF | $0.089 \pm 0.002$ | $0.120 \pm 0.000$ | $0.120 \pm 0.000$ | $0.470 \pm 0.010$ | $\underline{0.250 \pm 0.020}$ | - |

outperforms all baselines on two small UCI datasets and one large UCI datasets and has comparable performance on the rest. This further illustrates that exact marginalization over conditional approximate posteriors enabled by WMI solvers achieves accurate estimation of the true BMA and boosts predictive performance.

### 8.4.2 Image Classification

CIFAR datasets. We experiment with two image datasets: CIFAR-10 and CIFAR-100 [KH09b] and evaluate the test performance using three metrics: 1) negative log likelihood (NLL) that reflects the quality of both uncertainty estimation and prediction accuracy, 2) classification accuracy (ACC), and 3) expected calibration errors (ECE) [NCH15] that show the difference between predictive confidence and accuracy and should be close to zero for a well-calibrated approach.

We run Ciber by choosing the collapsed parameter set to be 10 weights and 100 weights at the last layer of the neural network models for CIFAR-10 and CIFAR-100 respectively. The weights are chosen using the same heuristic as the one for regression tasks, i.e., to choose the weights

Table 8.5: Average test performance for image classification tasks on CIFAR-10 and CIFAR-100.

| METRIC | NLL |  | ACC |  | ECE |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| DATASET | CIFAR-10 | CIFAR-100 | CIFAR-10 | CIFAR-100 | CIFAR-10 | CIFAR-100 |
| CIBER | $\mathbf{0 . 1 9 2 7} \pm \mathbf{0 . 0 0 2 9}$ | $\mathbf{0 . 9 1 9 3} \pm \mathbf{0 . 0 0 2 7}$ | $\mathbf{9 3 . 6 4} \pm \mathbf{0 . 0 9}$ | $\mathbf{7 4 . 7 1} \pm \mathbf{0 . 1 8}$ | $0.0130 \pm 0.0011$ | $\mathbf{0 . 0 1 6 8} \pm \mathbf{0 . 0 0 2 5}$ |
| SWAG | $0.2503 \pm 0.0081$ | $1.2785 \pm 0.0031$ | $93.59 \pm 0.14$ | $73.85 \pm 0.25$ | $0.0391 \pm 0.0020$ | $0.1535 \pm 0.0015$ |
| SGD | $0.3285 \pm 0.0139$ | $1.7308 \pm 0.0137$ | $93.17 \pm 0.14$ | $73.15 \pm 0.11$ | $0.0483 \pm 0.0022$ | $0.1870 \pm 0.0014$ |
| SWA | $0.2621 \pm 0.0104$ | $1.2780 \pm 0.0051$ | $93.61 \pm 0.11$ | $74.30 \pm 0.22$ | $0.0408 \pm 0.0019$ | $0.1514 \pm 0.0032$ |
| SGLD | $0.2001 \pm 0.0059$ | $0.9699 \pm 0.0057$ | $93.55 \pm 0.15$ | $74.02 \pm 0.30$ | $\mathbf{0 . 0 0 8 2} \pm \mathbf{0 . 0 0 1 2}$ | $0.0424 \pm 0.0029$ |
| KFAC | $0.2252 \pm 0.0032$ | $1.1915 \pm 0.0199$ | $92.65 \pm 0.20$ | $72.38 \pm 0.23$ | $0.0094 \pm 0.0005$ | $0.0778 \pm 0.0054$ |

Table 8.6: Average test performance for image transfer learning tasks.

| Metric | NLL |  | ACC |  | ECE |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| Model | VGG-16 | PreResNet-164 | VGG-16 | PreResNet-164 | VGG-16 | PreResNet-164 |
| CIBER | $\mathbf{0 . 9 8 6 9} \pm \mathbf{0 . 0 1 0 2}$ | $\mathbf{0 . 9 6 8 4} \pm \mathbf{0 . 0 0 7 5}$ | $72.56 \pm \mathbf{0 . 2 3}$ | $75.70 \pm 0.17$ | $\mathbf{0 . 0 9 2 5} \pm \mathbf{0 . 0 0 2 8}$ | $\mathbf{0 . 0 7 0 4} \pm \mathbf{0 . 0 0 3 1}$ |
| SWAG | $1.3425 \pm 0.0015$ | $1.3842 \pm 0.0122$ | $72.30 \pm 0.11$ | $76.30 \pm \mathbf{0 . 0 6}$ | $0.1988 \pm 0.0028$ | $0.1668 \pm 0.0006$ |
| SGD | $1.6528 \pm 0.0390$ | $1.4790 \pm 0.0000$ | $72.42 \pm 0.07$ | $75.56 \pm 0.00$ | $0.2149 \pm 0.0027$ | $0.1758 \pm 0.0000$ |
| SWA | $1.3993 \pm 0.0502$ | $1.3552 \pm 0.0000$ | $71.92 \pm 0.01$ | $76.02 \pm 0.00$ | $0.2082 \pm 0.0056$ | $0.1739 \pm 0.0000$ |

whose samples from the SGD trajectories have large variances. We compare Ciber with strong baselines including SWAG [MIG19] reproduced by their open-source implementation, standard SGD, SWA [IPG18], SGLD [WT11] and KFAC [RBB18].

Transfer from CIFAR-10 to STL-10. We further consider a transfer learning task using the model trained on CIFAR-10 to be evaluated on dataset STL-10 [CNL11]. STL-10 shares nine out of ten classes with the CIFAR-10 dataset but has a different image distribution. It is a common benchmark in transfer learning to adapt models trained on CIFAR-10 to STL-10.

Results. We present the test classification performance on dataset CIFAR-10 and CIFAR-100
in Table 8.5 and that of transfer learning in Table 8.6. The neural network models used in the classification task are VGG-16 networks and the models used in the transfer learning task are VGG-16 and PreResNet-164. With the same number of samples as SWAG, Ciber outperforms SWAG and other baselines in most evaluations and delivers comparable performance otherwise, demonstrating the effectiveness of using collapsed samples in improving uncertainty estimation as well as classification performance.

### 8.5 Discussion

We reveal the connection between BMA, a way to perform Bayesian deep learning and WVC, which inspires us to approximate BMA using the framework of WMI. To further make this approximation scalable and flexible, we combine it with collapsed samples which gives our algorithm Ciber. Ciber compares favorably to Bayesian deep learning baselines on regression and classification tasks. A future direction would be to explore what other layers can be expressed as SMT formulas and thus amenable to SMT encoding. Also, the current WMI solvers are limited to polynomial weights, and thus the reduction to WMI problems is applicable to piecewise polynomial weights. This limitation might be alleviated in the future by the development of new WMI solvers that allow various weight function families.

## CHAPTER 9

## Conclusion

Nowadays people have realized that there is a limit to what AI models can do without symbolic representations. Neurosymbolic AI is considered as the third wave of AI development, following the earlier waves dominated by symbolic AI and deep learning, and is a promising way to build next-generation AI models. This thesis summarize my research on tackling the fundamental challenges in neurosymbolic learning and reasoning to advance the field. These approaches are applied to build AI models that behave as intended to achieve trustworthiness.

For neurosymbolic learning, a key insight in our work is that constraint probability lies at the center of the differentiable learning under constraints-not only enabling the development of gradient estimators achieving lower bias and variances compared to state-of-the-art estimators but also aiding in deriving probabilistically sound training objectives. Our estimator is effective in latent space regularization for deep generative models and explainability for learning to explain tasks. It leads to direct follow-up work extending it to learning graph structure in a task-adaptive way for graph neural networks, the de facto standard models for geometry-heavy applications. Further, these techniques are extended to derive a unified framework for count-based weaklysupervised learning tasks.

At the core of scaling neurosymbolic models lies the development of efficient reasoning algorithms for highly expressive models. For the reasoning part, we focus on probabilistic reasoning over arithmetic constraints that arise in scenarios such as stochastic neural networks with piecewise linear activation functions. We use the the expressive WMI framework that allows such reasoning, enabling the modeling of structured, complex, and noisy data. Still, its expressiveness
comes at a cost: building reasoning algorithms for WMI problems is highly challenging. I develop a series of techniques to perform both exact and approximate inference, and further scale to large problems. The development of WMI solvers remains an open challenge and it will benefit tons of applications given its expressiveness in representing domain knowledge.

In summary, neurosymbolic AI plays a pivotal role in advancing the trustworthiness of AI by ensuring transparency, consistency, and reliability in decision-making processes, making it a critical research field for the sustainable development of AI technologies.

## References

[AAB19] Akshay Agrawal, Brandon Amos, Shane Barratt, Stephen Boyd, Steven Diamond, and J Zico Kolter. "Differentiable convex optimization layers." Advances in neural information processing systems, 32, 2019.
[AC17] Ehsan Mohammady Ardehaly and Aron Culotta. "Co-training for demographic classification using deep learning from label proportions." In 2017 IEEE International Conference on Data Mining Workshops (ICDMW), pp. 1017-1024. IEEE, 2017.
[ACB23] Kareem Ahmed, Kai-Wei Chang, and Guy Van den Broeck. "Semantic Strengthening of Neuro-Symbolic Learning." In Proceedings of the 26th International Conference on Artificial Intelligence and Statistics (AISTATS), apr 2023.
[ACD20] Ralph Abboud, Ismail Ilkan Ceylan, and Radoslav Dimitrov. "On the Approximability of Weighted Model Integration on DNF Structures." arXiv preprint arXiv:2002.06726, 2020.
[ACG20] Ralph Abboud, Ismail Ilkan Ceylan, Martin Grohe, and Thomas Lukasiewicz. "The surprising power of graph neural networks with random node initialization." arXiv preprint arXiv:2010.01179, 2020.
[AD15] Hadi Mohasel Afshar and Justin Domke. "Reflection, refraction, and hamiltonian monte carlo." In Advances in neural information processing systems, pp. 3007-3015, 2015.
[ADC22] Ralph Abboud, Radoslav Dimitrov, and Ismail Ilkan Ceylan. "Shortest path networks for graph property prediction." In Learning on Graphs Conference, pp. 5-1. PMLR, 2022.
[ALT22] Kareem Ahmed, Tao Li, Thy Ton, Quan Guo, Kai-Wei Chang, Parisa Kordjamshidi, Vivek Srikumar, Guy Van den Broeck, and Sameer Singh. "PYLON: A PyTorch Frame-
work for Learning with Constraints." In Proceedings of the 36th AAAI Conference on Artificial Intelligence (Demo Track), feb 2022.
[AMD14] Yoann Altmann, Steve McLaughlin, and Nicolas Dobigeon. "Sampling from a multivariate Gaussian distribution truncated on a simplex: a review." In 2014 IEEE Workshop on Statistical Signal Processing (SSP), pp. 113-116. IEEE, 2014.
[APK19] Sami Abu-El-Haija, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Kristina Lerman, Hrayr Harutyunyan, Greg Ver Steeg, and Aram Galstyan. "Mixhop: Higherorder graph convolutional architectures via sparsified neighborhood mixing." In International Conference on Machine Learning, pp. 21-29. PMLR, 2019.
[ASA15] Hadi Mohasel Afshar, Scott Sanner, and Ehsan Abbasnejad. "Linear-time gibbs sampling in piecewise graphical models." In Twenty-Ninth AAAI Conference on Artificial Intelligence, 2015.
[ASW16] Hadi Mohasel Afshar, Scott Sanner, and Christfried Webers. "Closed-Form Gibbs Sampling for Graphical Models with Algebraic Constraints." In AAAI, 2016.
[ATC22a] Kareem Ahmed, Stefano Teso, Kai-Wei Chang, Guy Van den Broeck, and Antonio Vergari. "Semantic Probabilistic Layers for Neuro-Symbolic Learning." In Advances in Neural Information Processing Systems 35 (NeurIPS), dec 2022.
[ATC22b] Kareem Ahmed, Stefano Teso, Kai-Wei Chang, Guy Van den Broeck, and Antonio Vergari. "Semantic probabilistic layers for neuro-symbolic learning." Advances in Neural Information Processing Systems, 35:29944-29959, 2022.
[ATH02] Stuart Andrews, Ioannis Tsochantaridis, and Thomas Hofmann. "Support Vector Machines for Multiple-Instance Learning." In NIPS, 2002.
[AWC21] Kareem Ahmed, Eric Wang, Kai-Wei Chang, and Guy Van den Broeck. "Leverag-
ing Unlabeled Data for Entity-Relation Extraction through Probabilistic Constraint Satisfaction.", mar 2021.
[AWC22] Kareem Ahmed, Eric Wang, Kai-Wei Chang, and Guy Van den Broeck. "NeuroSymbolic Entropy Regularization." In The 38th Conference on Uncertainty in Artificial Intelligence, 2022.
[AY21] Uri Alon and Eran Yahav. "On the Bottleneck of Graph Neural Networks and its Practical Implications." In International Conference on Learning Representations, 2021.
[AYO21] Laurence Aitchison, Adam Yang, and Sebastian W Ober. "Deep kernel processes." In International Conference on Machine Learning, pp. 130-140. PMLR, 2021.
[AZN23a] Kareem Ahmed, Zhe Zeng, Mathias Niepert, and Guy Van den Broeck. "SIMPLE: A Gradient Estimator for k -subset sampling." In International Conference on Learning Representations, 2023.
[AZN23b] Kareem Ahmed, Zhe Zeng, Mathias Niepert, and Guy Van den Broeck. "SIMPLE: A Gradient Estimator for k-subset sampling." In Proceedings of the International Conference on Learning Representations (ICLR), may 2023.
[AZN23c] Kareem Ahmed, Zhe Zeng, Mathias Niepert, and Guy Van den Broeck. "SIMPLE: A Gradient Estimator for k-subset sampling." In Proceedings of the International Conference on Learning Representations (ICLR), may 2023.
[BBD11] Velleda Baldoni, Nicole Berline, Jesus De Loera, Matthias Köppe, and Michèle Vergne. "How to integrate a polynomial over a simplex." Mathematics of Computation, 80(273):297-325, 2011.
[BBD14] Velleda Baldoni, Nicole Berline, Jesús A De Loera, Brandon Dutra, Matthias Köppe, Stanislav Moreinis, Gregory Pinto, Michele Vergne, and Jianqiu Wu. "A user’s guide for LattE integrale v1. 7.2." Optimization, 22(2), 2014.
[BBP15a] Vaishak Belle, Guy Van den Broeck, and Andrea Passerini. "Hashing-Based Approximate Probabilistic Inference in Hybrid Domains." In UAI, pp. 141-150, 2015.
[BBP15b] Vaishak Belle, Guy Van den Broeck, and Andrea Passerini. "Hashing-Based Approximate Probabilistic Inference in Hybrid Domains." In Proceedings of the 31st Conference on Uncertainty in Artificial Intelligence (UAI), 2015.
[BBP16] Vaishak Belle, Guy Van den Broeck, and Andrea Passerini. "Component Caching in Hybrid Domains with Piecewise Polynomial Densities." In AAAI, pp. 3369-3375, 2016.
[BBT20] Quentin Berthet, Mathieu Blondel, Olivier Teboul, Marco Cuturi, Jean-Philippe Vert, and Francis R. Bach. "Learning with Differentiable Pertubed Optimizers." In NeurIPS, 2020.
[BD20] Jessa Bekker and Jesse Davis. "Learning from positive and unlabeled data: A survey." Machine Learning, 109:719-760, 2020.
[BDO18] Gerda Bortsova, Florian Dubost, Silas Ørting, Ioannis Katramados, Laurens Hogeweg, Laura Thomsen, Mathilde Wille, and Marleen de Bruijne. "Deep learning from label proportions for emphysema quantification." In Medical Image Computing and Computer Assisted Intervention-MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part II 11, pp. 768-776. Springer, 2018.
[BDP09] Fahiem Bacchus, Shannon Dalmao, and Toniann Pitassi. "Solving \#SAT and Bayesian inference with backtracking search." Journal of Artificial Intelligence Research, 34:391-442, 2009.
[BES80] László Babai, Paul Erdos, and Stanley M Selkow. "Random graph isomorphism." SIaM fournal on computing, 9(3):628-635, 1980.
[BFG96] Craig Boutilier, Nir Friedman, Moises Goldszmidt, and Daphne Koller. "Contextspecific independence in Bayesian networks." In Proceedings of the Twelfth international conference on Uncertainty in artificial intelligence, pp. 115-123. Morgan Kaufmann Publishers Inc., 1996.
[BFW21] Cristian Bodnar, Fabrizio Frasca, Yuguang Wang, Nina Otter, Guido F Montufar, Pietro Lio, and Michael Bronstein. "Weisfeiler and lehman go topological: Message passing simplicial networks." In International Conference on Machine Learning, pp. 1026-1037. PMLR, 2021.
[BFZ22] Giorgos Bouritsas, Fabrizio Frasca, Stefanos Zafeiriou, and Michael M Bronstein. "Improving graph neural network expressivity via subgraph isomorphism counting." IEEE Transactions on Pattern Analysis and Machine Intelligence, 45(1):657-668, 2022.
[BGS16] Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. "Importance Weighted Autoencoders.", 2016.
[BGS22] Samy Badreddine, Artur d'Avila Garcez, Luciano Serafini, and Michael Spranger. "Logic tensor networks." Artificial Intelligence, 303:103649, 2022.
[BHH16] Thang Bui, Daniel Hernández-Lobato, Jose Hernandez-Lobato, Yingzhen Li, and Richard Turner. "Deep Gaussian processes for regression using approximate expectation propagation." In International conference on machine learning, pp. 1472-1481. PMLR, 2016.
[BHM09] Armin Biere, Marijn Heule, and Hans van Maaren. Handbook of satisfiability, volume 185. IOS press, 2009.
[BK79] Laszlo Babai and Ludik Kucera. "Canonical labelling of graphs in linear average time." In 20th Annual Symposium on Foundations of Computer Science (sfcs 1979), pp. 39-46. IEEE, 1979.
[BKM20] Pablo Barceló, Egor V. Kostylev, Mikael Monet, Jorge Pérez, Juan Reutter, and Juan Pablo Silva. "The Logical Expressiveness of Graph Neural Networks." In International Conference on Learning Representations, 2020.
[BKW22] Pradeep Kr Banerjee, Kedar Karhadkar, Yu Guang Wang, Uri Alon, and Guido Montúfar. "Oversquashing in GNNs through the lens of information contraction and graph expansion." In 2022 58th Annual Allerton Conference on Communication, Control, and Computing (Allerton), pp. 1-8. IEEE, 2022.
[Blo19] Mathieu Blondel. "Structured prediction with projection oracles." Advances in neural information processing systems, 32, 2019.
[BM98] Avrim Blum and Tom Mitchell. "Combining Labeled and Unlabeled Data with CoTraining." In Proceedings of the Eleventh Annual Conference on Computational Learning Theory, COLT' 98, p. 92-100, New York, NY, USA, 1998. Association for Computing Machinery.
[BMN20] Mathieu Blondel, André FT Martins, and Vlad Niculae. "Learning with Fenchel-Young losses." F. Mach. Learn. Res., 21(35):1-69, 2020.
[BMR10] Clark Barrett, Leonardo de Moura, Silvio Ranise, Aaron Stump, and Cesare Tinelli. "The SMT-LIB initiative and the rise of SMT (HVC 2010 award talk)." In Proceedings of the 6th international conference on Hardware and software: verification and testing, pp. 3-3. Springer-Verlag, 2010.
[BMS22] Jakub Bober, Anthea Monod, Emil Saucan, and Kevin N Webster. "Rewiring Networks for Graph Neural Network Training Using Discrete Geometry." arXiv preprint arXiv:2207.08026, 2022.
[BO04] Albert-Laszlo Barabasi and Zoltan N Oltvai. "Network biology: Understanding the cell's functional organization." Nature reviews genetics, 5(2):101-113, 2004.
[BPB15] Vaishak Belle, Andrea Passerini, and Guy Van den Broeck. "Probabilistic inference in hybrid domains by weighted model integration." In Proceedings of IFCAI, pp. 27702776, 2015.
[BT18] Clark Barrett and Cesare Tinelli. "Satisfiability modulo theories." In Handbook of Model Checking, pp. 305-343. Springer, 2018.
[BTB20] Mathieu Blondel, Olivier Teboul, Quentin Berthet, and Josip Djolonga. "Fast differentiable sorting and ranking." In International Conference on Machine Learning, pp. 950-959. PMLR, 2020.
[BYS22] Rickard Brüel-Gabrielsson, Mikhail Yurochkin, and Justin Solomon. "Rewiring with positional encodings for graph neural networks." arXiv preprint arXiv:2201.12674, 2022.
[CC96] Mary Kathryn Cowles and Bradley P Carlin. "Markov chain Monte Carlo convergence diagnostics: a comparative review." Journal of the American Statistical Association, 91(434):883-904, 1996.
[CCG18] Marc-André Carbonneau, Veronika Cheplygina, Eric Granger, and Ghyslain Gagnon. "Multiple instance learning: A survey of problem characteristics and applications." Pattern Recognition, 2018.
[CCH19] Guangyong Chen, Pengfei Chen, Chang-Yu Hsieh, Chee-Kong Lee, Benben Liao, Renjie Liao, Weiwen Liu, Jiezhong Qiu, Qiming Sun, Jie Tang, et al. "Alchemy: A quantum chemistry dataset for benchmarking ai models." arXiv preprint arXiv:1906.09427, 2019.
[CCK23] Quentin Cappart, Didier Chételat, Elias B. Khalil, Andrea Lodi, Christopher Morris, and Petar Veličković. "Combinatorial Optimization and Reasoning with Graph Neural Networks." Journal of Machine Learning Research, 24(130):1-61, 2023.
[CCZ17] Yulai Cong, Bo Chen, and Mingyuan Zhou. "Fast simulation of hyperplane-truncated multivariate normal distributions." 2017.
[CD06] Arthur Choi and Adnan Darwiche. "An edge deletion semantics for belief propagation and its practical impact on approximation quality." In Proceedings of the National Conference on Artificial Intelligence, volume 21, p. 1107. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2006.
[CD08] Mark Chavira and Adnan Darwiche. "On probabilistic inference by weighted model counting." 2008.
[CD10] Arthur Choi and Adnan Darwiche. "Relax, compensate and then recover." In $7 S A I$ International Symposium on Artificial Intelligence, pp. 167-180. Springer, 2010.
[CD12] Arthur Choi and Adnan Darwiche. "Approximating the partition function by deleting and then correcting for model edges." arXiv preprint arXiv:1206.3241, 2012.
[CFI92] Jin-Yi Cai, Martin Fürer, and Neil Immerman. "An optimal lower bound on the number of variables for graph identification." Combinatorica, 12(4):389-410, 1992.
[CFM14] Supratik Chakraborty, Daniel J Fremont, Kuldeep S Meel, Sanjit A Seshia, and Moshe Y Vardi. "Distribution-aware sampling and weighted model counting for SAT." In Twenty-Eighth AAAI Conference on Artificial Intelligence, 2014.
[CLT20] Xi Chen, Jason D Lee, Xin T Tong, and Yichen Zhang. "Statistical inference for model parameters in stochastic gradient descent." The Annals of Statistics, 48(1):251-273, 2020.
[CNA20] Gonçalo M Correia, Vlad Niculae, Wilker Aziz, and André FT Martins. "Efficient Marginalization of Discrete and Structured Latent Variables via Sparsity." Advances in Neural Information Processing Systems, 2020.
[CNL11] Adam Coates, Andrew Ng, and Honglak Lee. "An analysis of single-layer networks in unsupervised feature learning." In Proceedings of the fourteenth international conference on artificial intelligence and statistics, pp. 215-223. JMLR Workshop and Conference Proceedings, 2011.
[CNM19] Gonçalo M Correia, Vlad Niculae, and André FT Martins. "Adaptively Sparse Transformers." In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International foint Conference on Natural Language Processing (EMNLP-IFCNLP), pp. 2174-2184, 2019.
[COB22] Dexiong Chen, Leslie O'Bray, and Karsten Borgwardt. "Structure-aware transformer for graph representation learning." In International Conference on Machine Learning, pp. 3469-3489. PMLR, 2022.
[CST11a] Timothee Cour, Ben Sapp, and Ben Taskar. "Learning from partial labels." The fournal of Machine Learning Research, 12:1501-1536, 2011.
[CST11b] Timothee Cour, Ben Sapp, and Ben Taskar. "Learning from Partial Labels." F. Mach. Learn. Res., 12(null):1501-1536, jul 2011.
[CSW18] Jianbo Chen, Le Song, Martin Wainwright, and Michael Jordan. "Learning to Explain: An Information-Theoretic Perspective on Model Interpretation." In Jennifer Dy and Andreas Krause, editors, Proceedings of the 35th International Conference on Machine Learning, volume 80 of Proceedings of Machine Learning Research, pp. 883-892. PMLR, 10-15 Jul 2018.
[CT19] Caio Corro and Ivan Titov. "Differentiable Perturb-and-Parse: Semi-Supervised Parsing with a Structured Variational Auto-Encoder." In International Conference on Learning Representations, 2019.
[CTV19] Marco Cuturi, Olivier Teboul, and Jean-Philippe Vert. "Differentiable ranking and
sorting using optimal transport." Advances in neural information processing systems, 32, 2019.
[DAG20] Luca Di Liello, Pierfrancesco Ardino, Jacopo Gobbi, Paolo Morettin, Stefano Teso, and Andrea Passerini. "Efficient generation of structured objects with constrained adversarial networks." Advances in neural information processing systems, 33:1466314674, 2020.
[Dar01] Adnan Darwiche. "Recursive conditioning." Artificial Intelligence, 126(1-2):5-41, 2001.
[Dar09] Adnan Darwiche. Modeling and reasoning with Bayesian networks. Cambridge university press, 2009.
[DDG99] Francesco De Comité, François Denis, Rémi Gilleron, and Fabien Letouzey. "Positive and Unlabeled Examples Help Learning." pp. 219-230, 121999.
[DDK12] Jesús A De Loera, Brandon Dutra, Matthias Koeppe, Stanislav Moreinis, Gregory Pinto, and Jianqiu Wu. "Software for exact integration of polynomials over polyhedra." ACM Communications in Computer Algebra, 45(3/4):169-172, 2012.
[De 59] Louis De Branges. "The stone-weierstrass theorem." Proceedings of the American Mathematical Society, 10(5):822-824, 1959.
[DF88] Martin E. Dyer and Alan M. Frieze. "On the complexity of computing the volume of a polyhedron." SIAM fournal on Computing, 17(5):967-974, 1988.
[DG17] Dheeru Dua and Casey Graff. "UCI Machine Learning Repository.", 2017.
[DGB23] Francesco Di Giovanni, Lorenzo Giusti, Federico Barbero, Giulia Luise, Pietro Lio, and Michael M Bronstein. "On over-squashing in message passing neural networks: The impact of width, depth, and topology." In International Conference on Machine Learning, pp. 7865-7885. PMLR, 2023.
[DGM12] Michelangelo Diligenti, Marco Gori, Marco Maggini, and Leonardo Rigutini. "Bridging logic and kernel machines." Machine learning, 86:57-88, 2012.
[DHD20] Vincent Derkinderen, Evert Heylen, Pedro Zuidberg Dos Martires, Samuel Kolb, and Luc Raedt. "Ordering variables for weighted model integration." In Conference on Uncertainty in Artificial Intelligence, pp. 879-888. PMLR, 2020.
[DJL20] Vijay Prakash Dwivedi, Chaitanya K Joshi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. "Benchmarking graph neural networks." arXiv preprint arXiv:2003.00982, 2020.
[DKT07] Luc De Raedt, Angelika Kimmig, and Hannu Toivonen. "ProbLog: A probabilistic Prolog and its application in link discovery." In IFCAI 2007, Proceedings of the 20th international joint conference on artificial intelligence, pp. 2462-2467. IJCAI-INT JOINT CONF ARTIF INTELL, 2007.
[DLL01] Thomas Dietterich, Richard Lathrop, and Tomás Lozano-Pérez. "Solving the Multiple Instance Problem with Axis-Parallel Rectangles." Artificial Intelligence, 89:31-71, 03 2001.
[DLV22] Andreea Deac, Marc Lackenby, and Petar Veličković. "Expander graph propagation." In Learning on Graphs Conference, pp. 38-1. PMLR, 2022.
[DM02] Adnan Darwiche and Pierre Marquis. "A knowledge compilation map." Journal of Artificial Intelligence Research, 17:229-264, 2002.
[DM07] Rina Dechter and Robert Mateescu. "AND/OR search spaces for graphical models." Artificial intelligence, 171(2-3):73-106, 2007.
[Dom10] Justin Domke. "Implicit Differentiation by Perturbation." In Advances in Neural Information Processing Systems 23, pp. 523-531. 2010.
[DRG22] Vijay Prakash Dwivedi, Ladislav Rampášek, Michael Galkin, Ali Parviz, Guy Wolf, Anh Tuan Luu, and Dominique Beaini. "Long range graph benchmark." Advances in Neural Information Processing Systems, 35:22326-22340, 2022.
[EK12] David Easley, Jon Kleinberg, et al. "Networks, crowds, and markets." Cambridge Books, 2012.
[FBD19] Marc Fischer, Mislav Balunovic, Dana Drachsler-Cohen, Timon Gehr, Ce Zhang, and Martin Vechev. "DL2: training and querying neural networks with logic." In International Conference on Machine Learning, pp. 1931-1941. PMLR, 2019.
[FBL18] Tiago M Fragoso, Wesley Bertoli, and Francisco Louzada. "Bayesian model averaging: A systematic review and conceptual classification." International Statistical Review, 86(1):1-28, 2018.
[FNC20] Thomas Frerix, Matthias Nießner, and Daniel Cremers. "Homogeneous linear inequality constraints for neural network activations." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pp. 748-749, 2020.
[FNP19] Luca Franceschi, Mathias Niepert, Massimiliano Pontil, and Xiao He. "Learning discrete structures for graph neural networks." In International Conference on Machine Learning, pp. 1972-1982. PMLR, 2019.
[FRE20] Fabrizio Frasca, Emanuele Rossi, Davide Eynard, Ben Chamberlain, Michael Bronstein, and Federico Monti. "Sign: Scalable inception graph neural networks." arXiv preprint arXiv:2004.11198, 2020.
[FV13] Benoît Frénay and Michel Verleysen. "Classification in the presence of label noise: a survey." IEEE transactions on neural networks and learning systems, 25(5):845-869, 2013.
[GBT23] Lorenzo Giusti, Claudio Battiloro, Lucia Testa, Paolo Di Lorenzo, Stefania Sardellitti, and Sergio Barbarossa. "Cell attention networks." In 2023 International foint Conference on Neural Networks (IfCNN), pp. 1-8. IEEE, 2023.
[GCW18] Will Grathwohl, Dami Choi, Yuhuai Wu, Geoffrey Roeder, and David Duvenaud. "Backpropagation through the void: Optimizing control variates for black-box gradient estimation." ICLR, 2018.
[GDB23] Benjamin Gutteridge, Xiaowen Dong, Michael M Bronstein, and Francesco Di Giovanni. "DRew: dynamically rewired message passing with delay." In International Conference on Machine Learning, pp. 12252-12267. PMLR, 2023.
[GG15] Yarin Gal and Zoubin Ghahramani. "Bayesian convolutional neural networks with Bernoulli approximate variational inference." arXiv preprint arXiv:1506.02158, 2015.
[GG16] Yarin Gal and Zoubin Ghahramani. "Dropout as a bayesian approximation: Representing model uncertainty in deep learning." In international conference on machine learning, pp. 1050-1059. PMLR, 2016.
[GJ02] Michael R Garey and David S Johnson. Computers and intractability, volume 29. wh freeman New York, 2002.
[GL21] Eleonora Giunchiglia and Thomas Lukasiewicz. "Multi-label classification neural networks with hard logical constraints." fournal of Artificial Intelligence Research, 72:759-818, 2021.
[GL23] Artur d'Avila Garcez and Luis C Lamb. "Neurosymbolic AI: The 3 rd wave." Artificial Intelligence Review, 56(11):12387-12406, 2023.
[GPM14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. "Generative adversarial nets." In Advances in neural information processing systems, pp. 2672-2680, 2014.
[GRC23] Lorenzo Giusti, Teodora Reu, Francesco Ceccarelli, Cristian Bodnar, and Pietro Liò. "CIN++: Enhancing Topological Message Passing." arXiv preprint arXiv:2306.03561, 2023.
[Gro17] Martin Grohe. Descriptive complexity, canonisation, and definable graph structure theory, volume 47. Cambridge University Press, 2017.
[GSR17] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. "Neural message passing for quantum chemistry." In International Conference on Machine Learning, pp. 1263-1272. PMLR, 2017.
[GTH15] Luis Galárraga, Christina Teflioudi, Katja Hose, and Fabian M Suchanek. "Fast rule mining in ontological knowledge bases with AMIE++." The VLDB fournal, 24(6):707730, 2015.
[GWG19] Johannes Gasteiger, Stefan Weißenberger, and Stephan Günnemann. "Diffusion improves graph learning." Advances in Neural Information Processing Systems, 32, 2019.
[GWS21] Saurabh Garg, Yifan Wu, Alexander J Smola, Sivaraman Balakrishnan, and Zachary Lipton. "Mixture proportion estimation and pu learning: a modern approach." Advances in Neural Information Processing Systems, 34:8532-8544, 2021.
[GWZ18] Aditya Grover, Eric Wang, Aaron Zweig, and Stefano Ermon. "Stochastic Optimization of Sorting Networks via Continuous Relaxations." In International Conference on Learning Representations, 2018.
[HAB19] Matthias Hein, Maksym Andriushchenko, and Julian Bitterwolf. "Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 41-50, 2019.
[HBM18] Steven Holtzen, Guy Van den Broeck, and Todd Millstein. "Sound Abstraction and Decomposition of Probabilistic Programs." In Proceedings of the 35th International Conference on Machine Learning (ICML), 2018.
[HCW20] Pim de Haan, Taco S. Cohen, and Max Welling. "Natural graph networks." Advances in Neural Information Processing Systems, 33:3636-3646, 2020.
[HFZ20] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. "Open graph benchmark: Datasets for machine learning on graphs." Advances in Neural Information Processing Systems, 33:2211822133, 2020.
[HG95] David Heckerman and Dan Geiger. "Learning Bayesian networks: a unification for discrete and Gaussian domains." In Proceedings of the Eleventh conference on Uncertainty in artificial intelligence, pp. 274-284. Morgan Kaufmann Publishers Inc., 1995.
[HGM18] Jiawei He, Yu Gong, Joseph Marino, Greg Mori, and Andreas Lehrmann. "Variational autoencoders with jointly optimized latent dependency structure." In International conference on learning representations, 2018.
[HHL23] Xiaoxin He, Bryan Hooi, Thomas Laurent, Adam Perold, Yann LeCun, and Xavier Bresson. "A generalization of vit/mlp-mixer to graphs." In International Conference on Machine Learning, pp. 12724-12745. PMLR, 2023.
[HKW17] Xiaowei Huang, Marta Kwiatkowska, Sen Wang, and Min Wu. "Safety verification of deep neural networks." In International conference on computer aided verification, pp. 3-29. Springer, 2017.
[HLG20] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. "Strategies for Pre-training Graph Neural Networks." In International Conference on Learning Representations, 2020.
[HMB17] Steven Holtzen, Todd Millstein, and Guy Van den Broeck. "Probabilistic Program Abstractions." In Proceedings of the 33rd Conference on Uncertainty in Artificial Intelligence (UAI), 2017.
[Hul14] Eyke Hüllermeier. "Learning from imprecise and fuzzy observations: Data disambiguation through generalized loss minimization." International fournal of Approximate Reasoning, 55(7):1519-1534, 2014.
[HYL17] William L. Hamilton, Zhitao Ying, and Jure Leskovec. "Inductive Representation Learning on Large Graphs." In Advances in Neural Information Processing Systems, pp. 1024-1034, 2017.
[IMK20] Pavel Izmailov, Wesley J Maddox, Polina Kirichenko, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. "Subspace inference for Bayesian deep learning." In Uncertainty in Artificial Intelligence, pp. 1169-1179. PMLR, 2020.
[IPG18] Pavel Izmailov, Dmitrii Podoprikhin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. "Averaging weights leads to wider optima and better generalization." arXiv preprint arXiv:1803.05407, 2018.
[ITW18] Maximilian Ilse, Jakub Tomczak, and Max Welling. "Attention-based deep multiple instance learning." In International conference on machine learning, pp. 2127-2136. PMLR, 2018.
[IVH21] Pavel Izmailov, Sharad Vikram, Matthew D Hoffman, and Andrew Gordon Gordon Wilson. "What are Bayesian neural network posteriors really like?" In International Conference on Machine Learning, pp. 4629-4640. PMLR, 2021.
[JCB17] Wengong Jin, Connor Coley, Regina Barzilay, and Tommi Jaakkola. "Predicting organic reaction outcomes with weisfeiler-lehman network." Advances in Neural Information Processing Systems, 30, 2017.
[JEP21] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. "Highly accurate protein structure prediction with AlphaFold." Nature, 596(7873):583-589, 2021.
[JGP16] Eric Jang, Shixiang Gu, and Ben Poole. "Categorical reparameterization with gumbelsoftmax." arXiv preprint arXiv:1611.01144, 2016.
[JGP17a] Eric Jang, Shixiang Gu, and Ben Poole. "Categorical Reparameterization with Gumbel-Softmax." In International Conference on Learning Representations, 2017.
[JGP17b] Eric Jang, Shixiang Gu, and Ben Poole. "Categorical Reparameterization with Gumbel-Softmax." In 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings. OpenReview.net, 2017.
[JLB22] Laurent Valentin Jospin, Hamid Laga, Farid Boussaid, Wray Buntine, and Mohammed Bennamoun. "Hands-on Bayesian neural networks-A tutorial for deep learning users." IEEE Computational Intelligence Magazine, 17(2):29-48, 2022.
[KB15] Diederik P. Kingma and Jimmy Ba. "Adam: A Method for Stochastic Optimization." In International Conference on Learning Representations, 2015.
[KBD17] Guy Katz, Clark Barrett, David L Dill, Kyle Julian, and Mykel J Kochenderfer. "Reluplex: An efficient SMT solver for verifying deep neural networks." In International conference on computer aided verification, pp. 97-117. Springer, 2017.
[KBM22] Kedar Karhadkar, Pradeep Kr Banerjee, and Guido Montúfar. "FoSR: First-order spectral rewiring for addressing oversquashing in GNNs." arXiv preprint arXiv:2210.11790, 2022.
[KEH22] Agustinus Kristiadi, Runa Eschenhagen, and Philipp Hennig. "Posterior Refinement Improves Sample Efficiency in Bayesian Neural Networks." In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, Advances in Neural Information Processing Systems, volume 35, pp. 30333-30346. Curran Associates, Inc., 2022.
[KF09] Daphne Koller and Nir Friedman. "Probabilistic Graphical Models." 2009.
[KFL01] Frank R Kschischang, Brendan J Frey, and H-A Loeliger. "Factor graphs and the sumproduct algorithm." IEEE Transactions on information theory, 47(2):498-519, 2001.
[KGG20] Johannes Klicpera, Janek Groß, and Stephan Günnemann. "Directional Message Passing for Molecular Graphs." In International Conference on Learning Representations, 2020.
[KH09a] Alex Krizhevsky and Geoffrey Hinton. "Learning multiple layers of features from tiny images." Technical Report 0, University of Toronto, Toronto, Ontario, 2009.
[KH09b] Alex Krizhevsky, Geoffrey Hinton, et al. "Learning multiple layers of features from tiny images." 2009.
[KHH20] Agustinus Kristiadi, Matthias Hein, and Philipp Hennig. "Being bayesian, even just a bit, fixes overconfidence in relu networks." In International conference on machine learning, pp. 5436-5446. PMLR, 2020.
[KMS18] Samuel Kolb, Martin Mladenov, Scott Sanner, Vaishak Belle, and Kristian Kersting. "Efficient Symbolic Integration for Probabilistic Inference." In IFCAI, pp. 5031-5037, 2018.
[KMZ19] Samuel Kolb, Paolo Morettin, Pedro Zuidberg Dos Martires, Francesco Sommavilla, Andrea Passerini, Roberto Sebastiani, and Luc De Raedt. "The pywmi Framework and Toolbox for Probabilistic Inference using Weighted Model Integration." In Proceedings
of the Twenty-Eighth International foint Conference on Artificial Intelligence, IFCAI, pp. 6530-6532, 72019.
[KMZ21] Giannis Karamanolakis, Subhabrata Mukherjee, Guoqing Zheng, and Ahmed Hassan Awadallah. "Self-training with weak supervision." arXiv preprint arXiv:2104.05514, 2021.
[KND17] Ryuichi Kiryo, Gang Niu, Marthinus C Du Plessis, and Masashi Sugiyama. "Positiveunlabeled learning with non-negative risk estimator." Advances in neural information processing systems, 30, 2017.
[KSE16] Carolyn Kim, Ashish Sabharwal, and Stefano Ermon. "Exact sampling with integer linear programs and random perturbations." In Thirtieth AAAI Conference on Artificial Intelligence, 2016.
[KSF22] Alan A. Kaptanoglu, Brian M. de Silva, Urban Fasel, Kadierdan Kaheman, Andy J. Goldschmidt, Jared Callaham, Charles B. Delahunt, Zachary G. Nicolaou, Kathleen Champion, Jean-Christophe Loiseau, J. Nathan Kutz, and Steven L. Brunton. "PySINDy: A comprehensive Python package for robust sparse system identification." Journal of Open Source Software, 7(69):3994, 2022.
[KVW19] Wouter Kool, Herke Van Hoof, and Max Welling. "Stochastic beams and where to find them: The gumbel-top-k trick for sampling sequences without replacement." In International Conference on Machine Learning, pp. 3499-3508. PMLR, 2019.
[KW13] Diederik P Kingma and Max Welling. "Auto-encoding variational bayes." arXiv preprint arXiv:1312.6114, 2013.
[KW17] Thomas N. Kipf and Max Welling. "Semi-Supervised Classification with Graph Convolutional Networks." In International Conference on Learning Representations, 2017.
[KW22] Diederik P Kingma and Max Welling. "Auto-Encoding Variational Bayes.", 2022.
[LDG00] Fabien Letouzey, François Denis, and Rémi Gilleron. "Learning From Positive and Unlabeled examples." In Proceedings of the 11th International Conference on Algorithmic Learning Theory, ALT'00, pp. 71-85, Sydney, Australia, 2000. Springer Verlag.
[LeC98] Yann LeCun. "The MNIST database of handwritten digits." 1998.
[Lou20] Andreas Loukas. "What graph neural networks cannot learn: depth vs width." In International Conference on Learning Representations, 2020.
[LQL23] Bo Li, Peng Qi, Bo Liu, Shuai Di, Jingen Liu, Jiquan Pei, Jinfeng Yi, and Bowen Zhou. "Trustworthy ai: From principles to practices." ACM Computing Surveys, 55(9):1-46, 2023.
[LSB24] Ruoyan Li, Dipti Ranjan Sahu, Guy Van den Broeck, and Zhe Zeng. "Linear-Equality Constrained Deep Generative Models." (submitted), 2024.
[LW89] Steffen Lilholt Lauritzen and Nanny Wermuth. "Graphical models for associations between variables, some of which are qualitative and some quantitative." The annals of Statistics, pp. 31-57, 1989.
[LWJ21] Meng Liu, Zhengyang Wang, and Shuiwang Ji. "Non-local graph neural networks." IEEE Transactions on Pattern Analysis and Machine Intelligence, 44(12):10270-10276, 2021.
[LWZ23] Zhiyuan Liu, Dafei Wu, Weiwei Zhai, and Liang Ma. "SONAR enables cell type deconvolution with spatially weighted Poisson-Gamma model for spatial transcriptomics." Nature Communications, 14(1):4727, 2023.
[LZS23] Kangning Liu, Weicheng Zhu, Yiqiu Shen, Sheng Liu, Narges Razavian, Krzysztof J Geras, and Carlos Fernandez-Granda. "Multiple instance learning via iterative selfpaced supervised contrastive learning." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 3355-3365, 2023.
[LZW11] Huchuan Lu, Qiuhong Zhou, Dong Wang, and Ruan Xiang. "A co-training framework for visual tracking with multiple instance learning." In 2011 IEEE International Conference on Automatic Face \& Gesture Recognition (FG), pp. 539-544, 2011.
[MA16] Andre Martins and Ramon Astudillo. "From softmax to sparsemax: A sparse model of attention and multi-label classification." In International conference on machine learning, pp. 1614-1623. PMLR, 2016.
[Maa09] Laurens van der Maaten. "Learning a Parametric Embedding by Preserving Local Structure." In David van Dyk and Max Welling, editors, Proceedings of the Twelth International Conference on Artificial Intelligence and Statistics, volume 5 of Proceedings of Machine Learning Research, pp. 384-391, Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA, 16-18 Apr 2009. PMLR.
[Mac92] David JC MacKay. "The evidence framework applied to classification networks." Neural computation, 4(5):720-736, 1992.
[MAD17] David Merrell, Aws Albarghouthi, and Loris D'Antoni. "Weighted Model Integration with Orthogonal Transformations." Proceedings of the Twenty-Sixth International foint Conference on Artificial Intelligence, 2017.
[MB18] Arthur Mensch and Mathieu Blondel. "Differentiable dynamic programming for structured prediction and attention." In International Conference on Machine Learning, pp. 3462-3471. PMLR, 2018.
[MBR22] Hassan Maatouk, Xavier Bay, and Didier Rullière. "A note on simulating hyperplane-truncated multivariate normal distributions." Statistics \& Probability Letters, 191:109650, 2022.
[MBS19a] Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. "Provably Powerful Graph Networks." In Advances in Neural Information Processing Systems, pp. 2153-2164, 2019.
[MBS19b] Haggai Maron, Heli Ben-Hamu, Nadav Shamir, and Yaron Lipman. "Invariant and Equivariant Graph Networks." In International Conference on Learning Representations, 2019.
[MCZ18] Tom Minka, Ryan Cleven, and Yordan Zaykov. "Trueskill 2: An improved bayesian skill rating system." 2018.
[MDK18] Robin Manhaeve, Sebastijan Dumancic, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. "DeepProbLog: Neural Probabilistic Logic Programming." In NeurIPS, 2018.
[MDM24] Giuseppe Marra, Sebastijan Dumančić, Robin Manhaeve, and Luc De Raedt. "From statistical relational to neurosymbolic artificial intelligence: A survey." Artificial Intelligence, p. 104062, 2024.
[MFN23] Pasquale Minervini, Luca Franceschi, and Mathias Niepert. "Adaptive PerturbationBased Gradient Estimation for Discrete Latent Variable Models." In AAAI, 2023.
[MGM23] Luis Müller, Mikhail Galkin, Christopher Morris, and Ladislav Rampášek. "Attending to graph transformers." arXiv preprint arXiv:2302.04181, 2023.
[MHB17] Stephan Mandt, Matthew D Hoffman, and David M Blei. "Stochastic gradient descent as approximate bayesian inference." arXiv preprint arXiv:1704.04289, 2017.
[MIG19] Wesley J Maddox, Pavel Izmailov, Timur Garipov, Dmitry P Vetrov, and Andrew Gordon Wilson. "A simple baseline for bayesian uncertainty in deep learning." Advances in Neural Information Processing Systems, 32:13153-13164, 2019.
[MKB20] Christopher Morris, Nils M Kriege, Franka Bause, Kristian Kersting, Petra Mutzel, and Marion Neumann. "Tudataset: A collection of benchmark datasets for learning with graphs." arXiv preprint arXiv:2007.08663, 2020.
[MKT20] Paolo Morettin, Samuel Kolb, Stefano Teso, and Andrea Passerini. "Learning Weighted Model Integration Distributions." In AAAI, pp. 5224-5231, 2020.
[ML97] Oded Maron and Tomás Lozano-Pérez. "A Framework for Multiple-Instance Learning." In M. Jordan, M. Kearns, and S. Solla, editors, Advances in Neural Information Processing Systems, volume 10. MIT Press, 1997.
[MLJ12] Julian McAuley, J. Leskovec, and Dan Jurafsky. "Learning Attitudes and Attributes from Multi-aspect Reviews." 2012 IEEE 12th International Conference on Data Mining, pp. 1020-1025, 2012.
[MLM21] Christopher Morris, Yaron Lipman, Haggai Maron, Bastian Rieck, Nils M Kriege, Martin Grohe, Matthias Fey, and Karsten Borgwardt. "Weisfeiler and leman go machine learning: The story so far." arXiv preprint arXiv:2112.09992, 2021.
[MMS22] Eleonora Misino, Giuseppe Marra, and Emanuele Sansone. "Vael: Bridging variational autoencoders and probabilistic logic programming." Advances in Neural Information Processing Systems, 35:4667-4679, 2022.
[MMT16] Chris J Maddison, Andriy Mnih, and Yee Whye Teh. "The concrete distribution: A continuous relaxation of discrete random variables." arXiv preprint arXiv:1611.00712, 2016.
[MMT17a] Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. "The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables." In International Conference on Learning Representations, 2017.
[MMT17b] Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. "The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables." In 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings. OpenReview.net, 2017.
[MPS17] Paolo Morettin, Andrea Passerini, and Roberto Sebastiani. "Efficient weighted model integration via SMT-based predicate abstraction." In Proceedings of the 26th International foint Conference on Artificial Intelligence, pp. 720-728. AAAI Press, 2017.
[MPS19] Paolo Morettin, Andrea Passerini, and Roberto Sebastiani. "Advanced SMT techniques for weighted model integration." Artificial Intelligence, 275:1-27, 2019.
[MRF19] Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. "Weisfeiler and leman go neural: Higherorder graph neural networks." In Proceedings of the AAAI Conference on Artificial Intelligence, volume 33, pp. 4602-4609, 2019.
[MSR19] Ryan Murphy, Balasubramaniam Srinivasan, Vinayak Rao, and Bruno Ribeiro. "Relational pooling for graph representations." In International Conference on Machine Learning, pp. 4663-4673. PMLR, 2019.
[MTM14] Chris J Maddison, Daniel Tarlow, and Tom Minka. "A $*$ Sampling." In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, Advances in Neural Information Processing Systems, volume 27. Curran Associates, Inc., 2014.
[MTP23] Lassi Meronen, Martin Trapp, Andrea Pilzer, Le Yang, and Arno Solin. "Fixing Overconfidence in Dynamic Neural Networks." arXiv preprint arXiv:2302.06359, 2023.
[MTS21] Lassi Meronen, Martin Trapp, and Arno Solin. "Periodic activation functions induce stationarity." Advances in Neural Information Processing Systems, 34:1673-1685, 2021.
[MVD18] Alejandro Molina, Antonio Vergari, Nicola Di Mauro, Sriraam Natarajan, Floriana Esposito, and Kristian Kersting. "Mixed sum-product networks: A deep architecture for hybrid domains." In Thirty-second AAAI conference on artificial intelligence, 2018.
[NCH15] Mahdi Pakdaman Naeini, Gregory Cooper, and Milos Hauskrecht. "Obtaining well
calibrated probabilities using bayesian binning." In Proceedings of the AAAI conference on artificial intelligence, volume 29, 2015.
[NDR13] Nagarajan Natarajan, Inderjit S Dhillon, Pradeep K Ravikumar, and Ambuj Tewari. "Learning with Noisy Labels." In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, Advances in Neural Information Processing Systems, volume 26. Curran Associates, Inc., 2013.
[NGB16] Marion Neumann, Roman Garnett, Christian Bauckhage, and Kristian Kersting. "Propagation kernels: efficient graph kernels from propagated information." Machine learning, 102:209-245, 2016.
[Nil82] Nils J Nilsson. "Principles of Artificial Intelligence." Symbolic Computation, Berlin: Springer, 1982, 1982.
[NK00] Masatoshi Nei and Sudhir Kumar. Molecular evolution and phylogenetics. Oxford university press, 2000.
[NM20] Vlad Niculae and André F. T. Martins. "LP-SparseMAP: Differentiable Relaxed Optimization for Sparse Structured Prediction." In ICML, 2020.
[NMB18] Vlad Niculae, André F. T. Martins, Mathieu Blondel, and Claire Cardie. "SparseMAP: Differentiable Sparse Structured Inference." In ICML, 2018.
[NMF21a] Mathias Niepert, Pasquale Minervini, and Luca Franceschi. "Implicit MLE: backpropagating through discrete exponential family distributions." Advances in Neural Information Processing Systems, 34:14567-14579, 2021.
[NMF21b] Mathias Niepert, Pasquale Minervini, and Luca Franceschi. "Implicit mle: Backpropagating through discrete exponential family distributions." Advances in Neural Information Processing Systems, 34, 2021.
[NNN21] Son Nguyen, Duong Nguyen, Khai Nguyen, Khoat Than, Hung Bui, and Nhat Ho. "Structured Dropout Variational Inference for Bayesian Neural Networks." Advances in Neural Information Processing Systems, 34, 2021.
[NYC15] Anh Nguyen, Jason Yosinski, and Jeff Clune. "Deep neural networks are easily fooled: High confidence predictions for unrecognizable images." In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 427-436, 2015.
[PBK21] Felix Petersen, Christian Borgelt, Hilde Kuehne, and Oliver Deussen. "Learning with Algorithmic Supervision via Continuous Relaxations." Advances in Neural Information Processing Systems, 34, 2021.
[PBT06] Slav Petrov, Leon Barrett, Romain Thibaux, and Dan Klein. "Learning accurate, compact, and interpretable tree annotation." In Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics, pp. 433-440, 2006.
[PCT20] Max B Paulus, Dami Choi, Daniel Tarlow, Andreas Krause, and Chris J Maddison. "Gradient Estimation with Stochastic Softmax Tricks." arXiv preprint arXiv:2006.08063, 2020.
[Pea88] Judea Pearl. "Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.", 1988.
[PMF21] Pál András Papp, Karolis Martinkus, Lukas Faber, and Roger Wattenhofer. "DropGNN: Random dropouts increase the expressiveness of graph neural networks." Advances in Neural Information Processing Systems, 34:21997-22009, 2021.
[PNM19] Ben Peters, Vlad Niculae, and André FT Martins. "Sparse Sequence-to-Sequence Models." In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pp. 1504-1519, 2019.
[PNS14] Marthinus C du Plessis, Gang Niu, and Masashi Sugiyama. "Analysis of Learning from Positive and Unlabeled Data." In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, Advances in Neural Information Processing Systems, volume 27. Curran Associates, Inc., 2014.
[PNS15] Marthinus Du Plessis, Gang Niu, and Masashi Sugiyama. "Convex Formulation for Learning from Positive and Unlabeled Data." In Proceedings of the 32nd International Conference on Machine Learning, volume 37, 2015.
[PR18] Tobias Plötz and Stefan Roth. "Neural Nearest Neighbors Networks." In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada, pp. 1095-1106, 2018.
[PWC20] Hongbin Pei, Bingzhe Wei, Kevin Chen-Chuan Chang, Yu Lei, and Bo Yang. "GeomGCN: Geometric Graph Convolutional Networks." In International Conference on Learning Representations, 2020.
[QCC] Gwenolé Quellec, Guy Cazuguel, Béatrice Cochener, and Mathieu Lamard. "Multipleinstance learning for medical image and video analysis." IEEE reviews in biomedical engineering.
[QMA24] Chendi Qian, Andrei Manolache, Kareem Ahmed, Zhe Zeng, Guy Van den Broeck, Mathias Niepert, and Christopher Morris. "Probabilistically Rewired MessagePassing Neural Networks." In Proceedings of the Eleventh International Conference on Learning Representations (ICLR), 2024.
[QSC08] Novi Quadrianto, Alex Smola, Tibério Caetano, and Quoc Le. "Estimating labels from label proportions." 2008.
[RBB18] Hippolyt Ritter, Aleksandar Botev, and David Barber. "A scalable laplace approximation for neural networks." In 6th International Conference on Learning Representations, ICLR 2018-Conference Track Proceedings, volume 6. International Conference on Representation Learning, 2018.
[RGD22] Ladislav Rampášek, Michael Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and Dominique Beaini. "Recipe for a general, powerful, scalable graph transformer." Advances in Neural Information Processing Systems, 35:14501-14515, 2022.
[RHX20] Yu Rong, Wenbing Huang, Tingyang Xu, and Junzhou Huang. "DropEdge: Towards Deep Graph Convolutional Networks on Node Classification." In International Conference on Learning Representations, 2020.
[RMP20] Michal Rolinek, Vít Musil, Anselm Paulus, Marin Vlastelica, Claudio Michaelis, and Georg Martius. "Optimizing rank-based metrics with blackbox differentiation." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7620-7630, 2020.
[RNE22] Patrick Reiser, Marlen Neubert, André Eberhard, Luca Torresi, Chen Zhou, Chen Shao, Houssam Metni, Clint van Hoesel, Henrik Schopmans, Timo Sommer, et al. "Graph neural networks for materials science and chemistry." Communications Materials, 3(1):93, 2022.
[Rol17] Jason Tyler Rolfe. "Discrete Variational Autoencoders." In 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings. OpenReview.net, 2017.
[Rot96a] Dan Roth. "On the Hardness of Approximate Reasoning." Artif. Intell., 82(1-2):273302, 1996.
[Rot96b] Dan Roth. "On the hardness of approximate reasoning." Artificial Intelligence, 82(1-2):273-302, 1996.
[RS86] Neil Robertson and P.D Seymour. "Graph minors. II. Algorithmic aspects of treewidth." Fournal of Algorithms, 7(3):309 - 322, 1986.
[RSS20] Ali Raza, Arni Sturluson, Cory M Simon, and Xiaoli Fern. "Message passing neural networks for partial charge assignment to metal-organic frameworks." The Journal of Physical Chemistry C, 124(35):19070-19082, 2020.
[RSZ20] Michal Rolínek, Paul Swoboda, Dominik Zietlow, Anselm Paulus, Vít Musil, and Georg Martius. "Deep Graph Matching via Blackbox Differentiation of Combinatorial Solvers." In ECCV, 2020.
[RTS18] Carlos Riquelme, George Tucker, and Jasper Snoek. "Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling." arXiv preprint arXiv:1802.09127, 2018.
[SA12] Scott Sanner and Ehsan Abbasnejad. "Symbolic variable elimination for discrete and continuous graphical models." In Twenty-Sixth AAAI Conference on Artificial Intelligence, 2012.
[SBK05] Tian Sang, Paul Beame, and Henry A Kautz. "Performing Bayesian inference by weighted model counting." In AAAI, volume 5, pp. 475-481, 2005.
[SBK20] Sergey Shirobokov, Vladislav Belavin, Michael Kagan, Andrei Ustyuzhanin, and Atilim Gunes Baydin. "Black-box optimization with local generative surrogates." Advances in Neural Information Processing Systems, 33:14650-14662, 2020.
[SCB18] Hugh Salimbeni, Ching-An Cheng, Byron Boots, and Marc Deisenroth. "Orthogonally decoupled variational Gaussian processes." Advances in neural information processing systems, 31, 2018.
[SCQ20] Brian de Silva, Kathleen Champion, Markus Quade, Jean-Christophe Loiseau, J. Kutz, and Steven Brunton. "PySINDy: A Python package for the sparse identification of
nonlinear dynamical systems from data." Fournal of Open Source Software, 5(49):2104, 2020.
[SDC24] Mihaela Cătălina Stoian, Salijona Dyrmishi, Maxime Cordy, Thomas Lukasiewicz, and Eleonora Giunchiglia. "How Realistic Is Your Synthetic Data? Constraining Deep Generative Models for Tabular Data." arXiv preprint arXiv:2402.04823, 2024.
[SDD12] Scott Sanner, Karina Valdivia Delgado, and Leliane Nunes De Barros. "Symbolic dynamic programming for discrete and continuous state MDPs." arXiv preprint arXiv:1202.3762, 2012.
[SFM20] Aishwarya Sivaraman, Golnoosh Farnadi, Todd Millstein, and Guy Van den Broeck. "Counterexample-Guided Learning of Monotonic Neural Networks." In Advances in Neural Information Processing Systems 33 (NeurIPS), dec 2020.
[SGL24] Mihaela Cătălina Stoian, Eleonora Giunchiglia, and Thomas Lukasiewicz. "Exploiting t-norms for deep learning in autonomous driving." arXiv preprint arXiv:2402.11362, 2024.
[SGT08a] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. "Computational capabilities of graph neural networks." IEEE Transactions on Neural Networks, 20(1):81-102, 2008.
[SGT08b] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. "The graph neural network model." IEEE Transactions on Neural Networks, 20(1):61-80, 2008.
[SKP22] Hwanjun Song, Minseok Kim, Dongmin Park, Yooju Shin, and Jae-Gil Lee. "Learning from noisy labels with deep neural networks: A survey." IEEE Transactions on Neural Networks and Learning Systems, 2022.
[SOG16] Rodrigo de Salvo Braz, Ciaran O’Reilly, Vibhav Gogate, and Rina Dechter. "Probabilistic inference modulo theories." In Proceedings of the Twenty-Fifth International foint Conference on Artificial Intelligence, pp. 3591-3599. AAAI Press, 2016.
[SRM16] Casper Kaae Sønderby, Tapani Raiko, Lars Maaløe, Søren Kaae Sønderby, and Ole Winther. "Ladder Variational Autoencoders.", 2016.
[SRS15] Jasper Snoek, Oren Rippel, Kevin Swersky, Ryan Kiros, Nadathur Satish, Narayanan Sundaram, Mostofa Patwary, Mr Prabhat, and Ryan Adams. "Scalable bayesian optimization using deep neural networks." In International conference on machine learning, pp. 2171-2180. PMLR, 2015.
[SRT16] Korsuk Sirinukunwattana, Shan e Ahmed Raza, Yee Tsang, David Snead, Ian Cree, and Nasir Rajpoot. "Locality Sensitive Deep Learning for Detection and Classification of Nuclei in Routine Colon Cancer Histology Images." IEEE Transactions on Medical Imaging, 35:1-1, 022016.
[SS10] Marko Samer and Stefan Szeider. "Algorithms for propositional model counting." Journal of Discrete Algorithms, 8(1):50-64, 2010.
[SSV11] Nino Shervashidze, Pascal Schweitzer, Erik Jan Van Leeuwen, Kurt Mehlhorn, and Karsten M Borgwardt. "Weisfeiler-lehman graph kernels." Journal of Machine Learning Research, 12(9), 2011.
[SVP09] Nino Shervashidze, SVN Vishwanathan, Tobias Petri, Kurt Mehlhorn, and Karsten Borgwardt. "Efficient graphlet kernels for large graph comparison." In Artificial Intelligence and Statistics, pp. 488-495. PMLR, 2009.
[SVV23] Hamed Shirzad, Ameya Velingker, Balaji Venkatachalam, Danica J Sutherland, and Ali Kemal Sinop. "Exphormer: Sparse transformers for graphs." arXiv preprint arXiv:2303.06147, 2023.
[SW11] Prakash P Shenoy and James C West. "Inference in hybrid Bayesian networks using mixtures of polynomials." International fournal of Approximate Reasoning, 52(5):641657, 2011.
[SZ20] Clayton Scott and Jianxin Zhang. "Learning from Label Proportions: A Mutual Contamination Framework." In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, Advances in Neural Information Processing Systems, volume 33, pp. 22256-22267. Curran Associates, Inc., 2020.
[SZA23] Vinay Shukla, Zhe Zeng, Kareem Ahmed, and Guy Van den Broeck. "A Unified Approach to Count-Based Weakly-Supervised Learning." In Advances in Neural Information Processing Systems (NeurIPS), 2023.
[TDC21] Jake Topping, Francesco Di Giovanni, Benjamin Paul Chamberlain, Xiaowen Dong, and Michael M Bronstein. "Understanding over-squashing and bottlenecks on graphs via curvature." arXiv preprint arXiv:2111.14522, 2021.
[TL20] Kuen-Han Tsai and Hsuan-Tien Lin. "Learning from label proportions with consistency regularization." In Asian Conference on Machine Learning, pp. 513-528. PMLR, 2020.
[TMM17] George Tucker, Andriy Mnih, Chris J Maddison, Dieterich Lawson, and Jascha SohlDickstein. "Rebar: Low-variance, unbiased gradient estimates for discrete latent variable models." Advances in Neural Information Processing Systems, 2017.
[TWL24] Trieu H. Trinh, Yuhuai Wu, Quoc V. Le, He He, and Thang Luong. "Solving olympiad geometry without human demonstrations." Nature, 625(7995):476-482, 2024.
[Val79] Leslie G Valiant. "The complexity of enumeration and reliability problems." SIAM fournal on Computing, 8(3):410-421, 1979.
[VCC18] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. "Graph Attention Networks." In International Conference on Learning Representations, 2018.
[VGR81] William E Vesely, Francine F Goldberg, Norman H Roberts, and David F Haasl. "Fault tree handbook." Technical report, Nuclear Regulatory Commission Washington DC, 1981.
[VMP19] Antonio Vergari, Alejandro Molina, Robert Peharz, Zoubin Ghahramani, Kristian Kersting, and Isabel Valera. "Automatic Bayesian density analysis." In Proceedings of the AAAI Conference on Artificial Intelligence, volume 33, pp. 5207-5215, 2019.
[Was00] Larry Wasserman. "Bayesian model selection and model averaging." Fournal of mathematical psychology, 44(1):92-107, 2000.
[WHS16] Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P Xing. "Deep kernel learning." In Artificial intelligence and statistics, pp. 370-378. PMLR, 2016.
[WIB11] Janusz Wojtusiak, Katherine Irvin, Aybike Birerdinc, and Ancha V Baranova. "Using published medical results and non-homogenous data in rule learning." In 2011 10th International Conference on Machine Learning and Applications and Workshops, volume 2, pp. 84-89. IEEE, 2011.
[Wil92] Ronald J. Williams. "Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning." Mach. Learn., 8:229-256, 1992.
[WJW21] Zhanghao Wu, Paras Jain, Matthew Wright, Azalia Mirhoseini, Joseph E Gonzalez, and Ion Stoica. "Representing long-range context for graph neural networks with global attention." Advances in Neural Information Processing Systems, 34:13266-13279, 2021.
[WL68] Boris Weisfeiler and Andrei Leman. "The reduction of a graph to canonical form and the algebra which appears therein." nti, Series, 2(9):12-16, 1968.
[WNM19] Anqi Wu, Sebastian Nowozin, Edward Meeds, Richard E. Turner, Jose Miguel Hernandez-Lobato, and Alexander L. Gaunt. "Deterministic Variational Inference for Robust Bayesian Neural Networks." In International Conference on Learning Representations, 2019.
[WS98] Duncan J Watts and Steven H Strogatz. "Collective dynamics of 'smallworld'networks." nature, 393(6684):440, 1998.
[WT11] Max Welling and Yee W Teh. "Bayesian learning via stochastic gradient Langevin dynamics." In Proceedings of the 28th international conference on machine learning (ICML-11), pp. 681-688, 2011.
[WYT18] Xinggang Wang, Yongluan Yan, Peng Tang, Xiang Bai, and Wenyu Liu. "Revisiting multiple instance neural networks." Pattern Recognition, 74:15-24, 2018.
[XDC20] Yujia Xie, Hanjun Dai, Minshuo Chen, Bo Dai, Tuo Zhao, Hongyuan Zha, Wei Wei, and Tomas Pfister. "Differentiable top-k with optimal transport." Advances in Neural Information Processing Systems, 33:20520-20531, 2020.
[XE19a] Sang Michael Xie and Stefano Ermon. "Reparameterizable Subset Sampling via Continuous Relaxations." International foint Conference on Artificial Intelligence, 2019.
[XE19b] Sang Michael Xie and Stefano Ermon. "Reparameterizable Subset Sampling via Continuous Relaxations." In Proceedings of the Twenty-Eighth International foint Conference on Artificial Intelligence, IFCAI-19, pp. 3919-3925. International Joint Conferences on Artificial Intelligence Organization, 72019.
[XE19c] Sang Michael Xie and Stefano Ermon. "Reparameterizable subset sampling via continuous relaxations." arXiv preprint arXiv:1901.10517, 2019.
[XH18] Ming-Kun Xie and Sheng-Jun Huang. "Partial multi-label learning." In Proceedings of the AAAI conference on artificial intelligence, volume 32, 2018.
[XHL19] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. "How Powerful are Graph Neural Networks?" In International Conference on Learning Representations, 2019.
[XLT18] Keyulu Xu, Chengtao Li, Yonglong Tian, Tomohiro Sonobe, Ken-ichi Kawarabayashi, and Stefanie Jegelka. "Representation learning on graphs with jumping knowledge networks." In International Conference on Machine Learning, pp. 5453-5462. PMLR, 2018.
[XTX13] Chang Xu, Dacheng Tao, and Chao Xu. "A survey on multi-view learning." arXiv preprint arXiv:1304.5634, 2013.
[XZF18] Jingyi Xu, Zilu Zhang, Tal Friedman, Yitao Liang, and Guy Van den Broeck. "A semantic loss function for deep learning with symbolic knowledge." In International conference on machine learning, pp. 5502-5511. PMLR, 2018.
[Yar95] David Yarowsky. "Unsupervised Word Sense Disambiguation Rivaling Supervised Methods." In Proceedings of the 33rd Annual Meeting on Association for Computational Linguistics, ACL '95, p. 189-196, USA, 1995. Association for Computational Linguistics.
[YBR14] Eunho Yang, Yulia Baker, Pradeep Ravikumar, Genevera Allen, and Zhandong Liu. "Mixed graphical models via exponential families." In Artificial Intelligence and Statistics, pp. 1042-1050, 2014.
[YCK14] Felix X Yu, Krzysztof Choromanski, Sanjiv Kumar, Tony Jebara, and Shih-Fu Chang. "On learning from label proportions." arXiv preprint arXiv:1402.5902, 2014.
[YWS15] Zichao Yang, Andrew Wilson, Alex Smola, and Le Song. "A la carte-learning fast kernels." In Artificial Intelligence and Statistics, pp. 1098-1106. PMLR, 2015.
[ZB19] Zhe Zeng and Guy Van den Broeck. "Efficient Search-Based Weighted Model Integration." Proceedings of UAI, 2019.
[ZB23a] Zhe Zeng and Guy Van den Broeck. "Collapsed Inference for Bayesian Deep Learning." arXiv preprint arXiv:2306.09686, 2023.
[ZB23b] Zhe Zeng and Guy Van den Broeck. "Collapsed Inference for Bayesian Deep Learning." In Advances in Neural Information Processing Systems (NeurIPS), 2023.
[ZCN18] Muhan Zhang, Zhicheng Cui, Marion Neumann, and Yixin Chen. "An end-to-end deep learning architecture for graph classification." In Proceedings of the AAAI Conference on Artificial Intelligence, volume 32, 2018.
[ZD18] Kaja Zupanc and Jesse Davis. "Estimating rule quality for knowledge base completion with the relationship between coverage assumption." In Proceedings of the 2018 World Wide Web Conference, pp. 1073-1081, 2018.
[ZDD19] Pedro Miguel Zuidberg Dos Martires, Anton Dries, and Luc De Raedt. "Exact and Approximate Weighted Model Integration withProbability Density Functions Using Knowledge Compilation." In Proceedings of the 30th Conference on Artificial Intelligence. AAAI Press, 2019.
[ZG22] Xiaojin Zhu and Andrew B Goldberg. Introduction to semi-supervised learning. Springer Nature, 2022.
[Zho18] Zhi-Hua Zhou. "A brief introduction to weakly supervised learning." National science review, 5(1):44-53, 2018.
[Zhu05] Xiaojin Jerry Zhu. "Semi-supervised learning literature survey." 2005.
[ZKD19] Pedro Miguel Zuidberg Dos Martires, Samuel Kolb, and Luc De Raedt. "How to Exploit Structure while Solving Weighted Model Integration Problems." UAI 2019 Proceedings, 2019.
[ZMY20a] Zhe Zeng, Paolo Morettin, Fanqi Yan, Antonio Vergari, and Guy Van den Broeck. "Scaling up Hybrid Probabilistic Inference with Logical and Arithmetic Constraints via Message Passing." In Proceedings of the International Conference of Machine Learning (ICML), 2020.
[ZMY20b] Zhe Zeng, Paolo Morettin, Fanqi Yan, Antonio Vergari, and Guy Van den Broeck. "Probabilistic inference with algebraic constraints: Theoretical limits and practical approximations." Advances in Neural Information Processing Systems, 33:11564-11575, 2020.
[ZMY20c] Zhe Zeng, Paolo Morettin, Fanqi Yan, Antonio Vergari, and Guy Van den Broeck. "Relax, compensate and then integrate." In Proceedings of the ECML-PKDD Workshop on Deep Continuous-Discrete Machine Learning (DeCoDeML), 2020.
[ZMY21a] Zhe Zeng, Paolo Morettin, Fanqi Yan, Andrea Passerini, Guy Van den Broeck, et al. "Is Parameter Learning via Weighted Model Integration Tractable?" In The 4th Workshop on Tractable Probabilistic Modeling, 2021.
[ZMY21b] Zhe Zeng, Paolo Morettin, Fanqi Yan, Antonio Vergari, and Guy Van den Broeck. "Is Parameter Learning via Weighted Model Integration Tractable?" In Proceedings of the UAI Workshop on Tractable Probabilistic Modeling (TPM), jul 2021.
[ZSF12] Zahra Zamani, Scott Sanner, and Cheng Fang. "Symbolic dynamic programming for continuous state and action mdps." In Proceedings of the AAAI Conference on Artificial Intelligence, volume 26, pp. 1839-1845, 2012.
[Zui20] Pedro Miguel Zuidberg Dos Martires et al. "Monte Carlo Anti-Differentiation for Approximate Weighted Model Integration." In Proceedings of the Ninth International

Workshop on Statistical Relational AI. International Workshop on Statistical Relational AI, 2020.


[^0]:    ${ }^{1}$ Code and experiments are available at https://github.com/UCLA-StarAI/CountLoss

[^1]:    ${ }^{2}$ Publicly available at archive.ics.uci.edu/ml

[^2]:    ${ }^{1}$ Our implementation of MP-WMI and the code for reproducing our empirical evaluation can be found at https : //github.com/UCLA-StarAI/mpwmi.

[^3]:    ${ }^{1}$ Following [CD10], we initialize all parameters to 1.

[^4]:    ${ }^{1}$ In this chapter, since we only work with WMI problems over continuous variables, we ignore the discrete ones in the definition for succinctness.

[^5]:    ${ }^{1}$ Code and experiments are available at https://github.com/UCLA-StarAI/CIBER.

