

# Privately Learning Markov Random Fields

Huanyu Zhang\*    Gautam Kamath<sup>†¶</sup>    Janardhan Kulkarni<sup>‡¶</sup>    Zhiwei Steven Wu<sup>§¶</sup>

August 13, 2020

## Abstract

We consider the problem of learning Markov Random Fields (including the prototypical example, the Ising model) under the constraint of differential privacy. Our learning goals include both *structure learning*, where we try to estimate the underlying graph structure of the model, as well as the harder goal of *parameter learning*, in which we additionally estimate the parameter on each edge. We provide algorithms and lower bounds for both problems under a variety of privacy constraints – namely pure, concentrated, and approximate differential privacy. While non-privately, both learning goals enjoy roughly the same complexity, we show that this is not the case under differential privacy. In particular, only structure learning under approximate differential privacy maintains the non-private logarithmic dependence on the dimensionality of the data, while a change in either the learning goal or the privacy notion would necessitate a polynomial dependence. As a result, we show that the privacy constraint imposes a strong separation between these two learning problems in the high-dimensional data regime.

## 1 Introduction

Graphical models are a common structure used to model high-dimensional data, which find a myriad of applications in diverse research disciplines, including probability theory, Markov Chain Monte Carlo, computer vision, theoretical computer science, social network analysis, game theory, and computational biology [LPW09, Cha05, Fel04, DMR11, GG86, Ell93, MS10]. While statistical tasks involving general distributions over  $p$  variables often run into the curse of dimensionality (i.e., an exponential sample complexity in  $p$ ), Markov Random Fields (MRFs) are a particular family of undirected graphical models which are parameterized by the “order”  $t$  of their interactions. Restricting the order of interactions allows us to capture most distributions which may naturally arise, and also avoids this severe dependence on the dimension (i.e., we often pay an exponential dependence on  $t$  instead of  $p$ ). An MRF is defined as follows, see Section 2 for more precise definitions and notations we will use in this paper.

**Definition 1.1.** *Let  $k, t, p \in \mathbb{N}$ ,  $G = (V, E)$  be a graph on  $p$  nodes, and  $C_t(G)$  be the set of cliques of size at most  $t$  in  $G$ . A Markov Random Field with alphabet size  $k$  and  $t$ -order interactions is a*

---

\*Cornell University. [hz388@cornell.edu](mailto:hz388@cornell.edu). Supported by NSF #1815893 and by NSF #1704443. This work was partially done while the author was an intern at Microsoft Research Redmond.

†University of Waterloo. [g@csail.mit.edu](mailto:g@csail.mit.edu). Supported by a University of Waterloo startup grant. Part of this work was done while supported as a Microsoft Research Fellow, as part of the Simons-Berkeley Research Fellowship program, and while visiting Microsoft Research Redmond.

‡Microsoft Research Redmond. [jakul@microsoft.com](mailto:jakul@microsoft.com).

§University of Minnesota. [zsw@umn.edu](mailto:zsw@umn.edu). Supported in part by the NSF FAI Award #1939606, a Google Faculty Research Award, a J.P. Morgan Faculty Award, a Facebook Research Award, and a Mozilla Research Grant.

¶These authors are in alphabetical order.

distribution  $\mathcal{D}$  over  $[k]^p$  such that

$$\Pr_{X \sim \mathcal{D}} [X = x] \propto \exp \left( \sum_{I \in \mathcal{C}_t(G)} \psi_I(x) \right),$$

where  $\psi_I : [k]^p \rightarrow \mathbb{R}$  depends only on variables in  $I$ .

The case when  $k = t = 2$  corresponds to the prototypical example of an MRF, the Ising model [Isi25] (Definition 2.1). More generally, if  $t = 2$ , we call the model *pairwise* (Definition 2.2), and if  $k = 2$  but  $t$  is unrestricted, we call the model a *binary MRF* (Definition 2.4). In this paper, we mainly look at these two special cases of MRFs.

Given the wide applicability of these graphical models, there has been a great deal of work on the problem of graphical model estimation [RWL10, SW12, Bre15, VMLC16, KM17, HKM17, RH17, LVMC18, WSD19]. That is, given a dataset generated from a graphical model, can we infer properties of the underlying distribution? Most of the attention has focused on two learning goals.

1. *Structure learning* (Definition 2.9): Recover the set of non-zero edges in  $G$ .
2. *Parameter learning* (Definition 2.10): Recover the set of non-zero edges in  $G$ , as well as  $\psi_I$  for all cliques  $I$  of size at most  $t$ .

It is clear that structure learning is easier than parameter learning. Nonetheless, the sample complexity of both learning goals is known to be roughly equivalent. That is, both can be performed using a number of samples which is only *logarithmic* in the dimension  $p$  (assuming a model of bounded “width”  $\lambda^1$ ), thus facilitating estimation in very high-dimensional settings.

However, in modern settings of data analysis, we may be running our algorithms on datasets which are sensitive in nature. For instance, graphical models are often used to model medical and genetic data [FLNP00, LAFH01] – if our learning algorithm reveals too much information about individual datapoints used to train the model, this is tantamount to releasing medical records of individuals providing their data, thus violating their privacy. In order to assuage these concerns, we consider the problem of learning graphical models under the constraint of *differential privacy* (DP) [DMNS06], considered by many to be the gold standard of data privacy. Informally, an algorithm is said to be differentially private if its distribution over outputs is insensitive to the addition or removal of a single datapoint from the dataset (a more formal definition is provided in Section 2). Differential privacy has enjoyed widespread adoption, including deployment in Apple [Dif17], Google [EPK14], Microsoft [DKY17], and the US Census Bureau for the 2020 Census [DLS+17].

Our goal is to design algorithms which guarantee both:

- **Accuracy:** With probability greater than  $2/3$ , the algorithm learns the underlying graphical model;
- **Privacy:** The algorithm satisfies differential privacy, even when the dataset is not drawn from a graphical model.

Thematically, we investigate the following question: how much additional data is needed to learn Markov Random Fields under the constraint of differential privacy? As mentioned before, absent privacy constraints, the sample complexity is logarithmic in  $p$ . Can we guarantee privacy with comparable amounts of data? Or if more data is needed, how much more?

---

<sup>1</sup>This is a common parameterization of the problem, which roughly corresponds to the graph having bounded-degree, see Section 2 for more details.

## 1.1 Results and Techniques

We proceed to describe our results on privately learning Markov Random Fields. In this section, we will assume familiarity with some of the most common notions of differential privacy: pure  $\epsilon$ -differential privacy,  $\rho$ -zero-concentrated differential privacy, and approximate  $(\epsilon, \delta)$ -differential privacy. In particular, one should know that these are in (strictly) decreasing order of strength (i.e., an algorithm which satisfies pure DP gives more privacy to the dataset than concentrated DP), formal definitions appear in Section 2. Furthermore, in order to be precise, some of our theorem statements will use notation which is defined later (Section 2) – these may be skipped on a first reading, as our prose will not require this knowledge.

**Upper Bounds.** Our first upper bounds are for parameter learning. First, we have the following theorem, which gives an upper bound for parameter learning pairwise graphical models under concentrated differential privacy, showing that this learning goal can be achieved with  $O(\sqrt{p})$  samples. In particular, this includes the special case of the Ising model, which corresponds to an alphabet size  $k = 2$ . Note that this implies the same result if one relaxes the learning goal to structure learning, or the privacy notion to approximate DP, as these modifications only make the problem easier. Further details are given in Section 3.3.

**Theorem 1.2.** *There exists an efficient  $\rho$ -zCDP algorithm which learns the parameters of a pairwise graphical model to accuracy  $\alpha$  with probability at least  $2/3$ , which requires a sample complexity of*

$$n = O\left(\frac{\lambda^2 k^5 \log(pk) e^{O(\lambda)}}{\alpha^4} + \frac{\sqrt{p} \lambda^2 k^{5.5} \log^2(pk) e^{O(\lambda)}}{\sqrt{\rho} \alpha^3}\right)$$

This result can be seen as a private adaptation of the elegant work of [WSD19] (which in turn builds on the structural results of [KM17]). Wu, Sanghavi, and Dimakis [WSD19] show that  $\ell_1$ -constrained logistic regression suffices to learn the parameters of all pairwise graphical models. We first develop a private analog of this method, based on the private Franke-Wolfe method of Talwar, Thakurta, and Zhang [TTZ14, TTZ15], which is of independent interest. This method is studied in Section 3.1.

**Theorem 1.3.** *If we consider the problem of private sparse logistic regression, there exists an efficient  $\rho$ -zCDP algorithm that produces a parameter vector  $w^{priv}$ , such that with probability at least  $1 - \beta$ , the empirical risk*

$$\mathcal{L}(w^{priv}; D) - \mathcal{L}(w^{erm}; D) = O\left(\frac{\lambda^{\frac{4}{3}} \log(\frac{np}{\beta})}{(n\sqrt{\rho})^{\frac{2}{3}}}\right).$$

We note that Theorem 1.3 avoids a polynomial dependence on the dimension  $p$  in favor of a polynomial dependence on the “sparsity” parameter  $\lambda$ . The greater dependence on  $p$  which arises in Theorem 1.2 is from applying Theorem 1.3 and then using composition properties of concentrated DP.

We go on to generalize the results of [WSD19], showing that  $\ell_1$ -constrained logistic regression can also learn the parameters of binary  $t$ -wise MRFs. This result is novel even in the non-private setting. Further details are presented in Section 4.

The following theorem shows that we can learn the parameters of binary  $t$ -wise MRFs with  $\tilde{O}(\sqrt{p})$  samples.

**Theorem 1.4.** *Let  $\mathcal{D}$  be an unknown binary  $t$ -wise MRF with associated polynomial  $h$ . Then there exists a  $\rho$ -zCDP algorithm which, with probability at least  $2/3$ , learns the maximal monomials of  $h$  to accuracy  $\alpha$ , given  $n$  i.i.d. samples  $Z^1, \dots, Z^n \sim \mathcal{D}$ , where*

$$n = O\left(\frac{e^{5\lambda t} \sqrt{p} \log^2(p)}{\sqrt{\rho} \alpha^{\frac{9}{2}}} + \frac{t \lambda^2 \sqrt{p} \log p}{\sqrt{\rho} \alpha^2} + \frac{e^{6\lambda t} \log(p)}{\alpha^6}\right).$$

To obtain the rate above, our algorithm uses the Private Multiplicative Weights (PMW) method by [HR10] to estimate all parity queries of all orders no more than  $t$ . The PMW method runs in time exponential in  $p$ , since it maintains a distribution over the data domain. We can also obtain an *oracle-efficient* algorithm that runs in polynomial time when given access to an empirical risk minimization oracle over the class of parities. By replacing PMW with such an oracle-efficient algorithm SEPFEM in [VTB<sup>+</sup>19], we obtain a slightly worse sample complexity

$$n = O\left(\frac{e^{5\lambda t} \sqrt{p} \log^2(p)}{\sqrt{\rho} \alpha^{\frac{9}{2}}} + \frac{t \lambda^2 p^{5/4} \log p}{\sqrt{\rho} \alpha^2} + \frac{e^{6\lambda t} \log(p)}{\alpha^6}\right).$$

For the special case of structure learning under approximate differential privacy, we provide a significantly better algorithm. In particular, we can achieve an  $O(\log p)$  sample complexity, which improves exponentially on the above algorithm's sample complexity of  $O(\sqrt{p})$ . The following is a representative theorem statement for pairwise graphical models, though we derive similar statements for binary MRFs of higher order.

**Theorem 1.5.** *There exists an efficient  $(\varepsilon, \delta)$ -differentially private algorithm which, with probability at least  $2/3$ , learns the structure of a pairwise graphical model, which requires a sample complexity of*

$$n = O\left(\frac{\lambda^2 k^4 \exp(14\lambda) \log(pk) \log(1/\delta)}{\varepsilon \eta^4}\right).$$

This result can be derived using stability properties of non-private algorithms. In particular, in the non-private setting, the guarantees of algorithms for this problem recover the entire graph *exactly* with constant probability. This allows us to derive private algorithms at a multiplicative cost of  $O(\log(1/\delta)/\varepsilon)$  samples, using either the propose-test-release framework [DL09] or stability-based histograms [KKMN09, BNSV15]. Further details are given in Section 6.

**Lower Bounds.** We note the significant gap between the aforementioned upper bounds: in particular, our more generally applicable upper bound (Theorem 1.2) has a  $O(\sqrt{p})$  dependence on the dimension, whereas the best known lower bound is  $\Omega(\log p)$  [SW12]. However, we show that our upper bound is tight. That is, even if we relax the privacy notion to approximate differential privacy, *or* relax the learning goal to structure learning, the sample complexity is still  $\Omega(\sqrt{p})$ . Perhaps surprisingly, if we perform both relaxations simultaneously, this falls into the purview of Theorem 1.5, and the sample complexity drops to  $O(\log p)$ .

First, we show that even under approximate differential privacy, learning the parameters of a graphical model requires  $\Omega(\sqrt{p})$  samples. The formal statement is given in Section 5.

**Theorem 1.6 (Informal).** *Any algorithm which satisfies approximate differential privacy and learns the parameters of a pairwise graphical model with probability at least  $2/3$  requires  $\text{poly}(p)$  samples.*

This result is proved by constructing a family of instances of binary pairwise graphical models (i.e., Ising models) which encode product distributions. Specifically, we consider the set of graphs

formed by a perfect matching with edges  $(2i, 2i+1)$  for  $i \in [p/2]$ . In order to estimate the parameter on every edge, one must estimate the correlation between each such pair of nodes, which can be shown to correspond to learning the mean of a particular product distribution in  $\ell_\infty$ -distance. This problem is well-known to have a gap between the non-private and private sample complexities, due to methods derived from fingerprinting codes [BUV14, DSS<sup>+</sup>15, SU17], and differentially private Fano’s inequality [ASZ20b].

Second, we show that learning the structure of a graphical model, under either pure or concentrated differential privacy, requires  $\text{poly}(p)$  samples. The formal theorem appears in Section 7.

**Theorem 1.7** (Informal). *Any algorithm which satisfies pure or concentrated differential privacy and learns the structure of a pairwise graphical model with probability at least  $2/3$  requires  $\text{poly}(p)$  samples.*

We derive this result via packing arguments [HT10, BBKN14, ASZ20b], by showing that there exists a large number (exponential in  $p$ ) of different binary pairwise graphical models which must be distinguished. The construction of a packing of size  $m$  implies lower bounds of  $\Omega(\log m)$  and  $\Omega(\sqrt{\log m})$  for learning under pure and concentrated differential privacy, respectively.

### 1.1.1 Summary and Discussion

We summarize our findings on privately learning Markov Random Fields in Table 1, focusing on the specific case of the Ising model. We note that qualitatively similar relationships between problems also hold for general pairwise models as well as higher-order binary Markov Random Fields. Each cell denotes the sample complexity of a learning task, which is a combination of an objective and a privacy constraint. Problems become harder as we go down (as the privacy requirement is tightened) and to the right (structure learning is easier than parameter learning).

The top row shows that both learning goals require only  $\Theta(\log p)$  samples to perform absent privacy constraints, and are thus tractable even in very high-dimensional settings or when data is limited. However, if we additionally wish to guarantee privacy, our results show that this logarithmic sample complexity is only achievable when one considers structure learning under approximate differential privacy. If one changes the learning goal to parameter learning, *or* tightens the privacy notion to concentrated differential privacy, then the sample complexity jumps to become polynomial in the dimension, in particular  $\Omega(\sqrt{p})$ . Nonetheless, we provide algorithms which match this dependence, giving a tight  $\Theta(\sqrt{p})$  bound on the sample complexity.

	<b>Structure Learning</b>	<b>Parameter Learning</b>
<b>Non-private</b>	$\Theta(\log p)$ (folklore)	$\Theta(\log p)$ (folklore)
<b>Approximate DP</b>	$\Theta(\log p)$ (Theorems 6.3)	$\Theta(\sqrt{p})$ (Theorems 3.8 and 5.1)
<b>Zero-concentrated DP</b>	$\Theta(\sqrt{p})$ (Theorems 3.8 and 7.1)	$\Theta(\sqrt{p})$ (Theorems 3.8 and 5.1)
<b>Pure DP</b>	$\Omega(p)$ (Theorem 7.1)	$\Omega(p)$ (Theorem 7.1)

Table 1: Sample complexity (dependence on  $p$ ) of privately learning an Ising model.

## 1.2 Related Work

As mentioned before, there has been significant work in learning the structure and parameters of graphical models, see, e.g., [CL68, CT06, AKN06, RWL10, JJR11, JRVS11, SW12, BGS14, Bre15,

VMLC16, KM17, HKM17, RH17, LVMC18, WSD19]. Perhaps a turning point in this literature is the work of Bresler [Bre15], who showed for the first time that general Ising models of bounded degree can be learned in polynomial time. Since this result, following works have focused on both generalizing these results to broader settings (including MRFs with higher-order interactions and non-binary alphabets) as well as simplifying existing arguments. There has also been work on learning, testing, and inferring other statistical properties of graphical models [BM16, MdCCU16, DDK17, MMY18, Bha19]. In particular, learning and testing Ising models in statistical distance have also been explored [DDK18, GLP18, DMR18, DDK19, BBC<sup>+</sup>19], and are interesting questions under the constraint of privacy.

Recent investigations at the intersection of graphical models and differential privacy include [BMS<sup>+</sup>17, CRJ19, MSM19]. Bernstein et al. [BMS<sup>+</sup>17] privately learn graphical models by adding noise to the sufficient statistics and use an expectation-maximization based approach to recover the parameters. However, the focus is somewhat different, as they do not provide finite sample guarantees for the accuracy when performing parameter recovery, nor consider structure learning at all. Chowdhury, Rekatsinas, and Jha [CRJ19] study differentially private learning of Bayesian Networks, another popular type of graphical model which is incomparable with Markov Random Fields. McKenna, Sheldon, and Miklau [MSM19] apply graphical models in place of full contingency tables to privately perform inference.

Graphical models can be seen as a natural extension of product distributions, which correspond to the case when the order of the MRF  $t$  is 1. There has been significant work in differentially private estimation of product distributions [BDMN05, BUV14, DMNS06, SU17, KLSU19, CWZ19, BKSU19]. Recently, this investigation has been broadened into differentially private distribution estimation, including sample-based estimation of properties and parameters, see, e.g., [NRS07, Smi11, BNSV15, DHS15, KV18, AKSZ18, KLSU19, BKSU19]. For further coverage of differentially private statistics, see [KU20].

## 2 Preliminaries

Given an integer  $n$ , we let  $[n] := \{1, 2, \dots, n\}$ . Given a set of points  $X^1, \dots, X^n$ , we use superscripts, i.e.,  $X^i$  to denote the  $i$ -th datapoint. Given a vector  $X \in \mathbb{R}^p$ , we use subscripts, i.e.,  $X_i$  to denote its  $i$ -th coordinate. We also use  $X_{-i}$  to denote the vector after deleting the  $i$ -th coordinate, i.e.  $X_{-i} = [X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_p]$ .

### 2.1 Markov Random Field Preliminaries

We first introduce the definition of the Ising model, which is a special case of general MRFs when  $k = t = 2$ .

**Definition 2.1.** *The  $p$ -variable Ising model is a distribution  $\mathcal{D}(A, \theta)$  on  $\{-1, 1\}^p$  that satisfies*

$$\Pr(Z = z) \propto \exp \left( \sum_{1 \leq i < j \leq p} A_{i,j} z_i z_j + \sum_{i \in [p]} \theta_i z_i \right),$$

where  $A \in \mathbb{R}^{p \times p}$  is a symmetric weight matrix with  $A_{ii} = 0, \forall i \in [p]$  and  $\theta \in \mathbb{R}^p$  is a mean-field vector. The dependency graph of  $\mathcal{D}(A, \theta)$  is an undirected graph  $G = (V, E)$ , with vertices  $V = [p]$

and edges  $E = \{(i, j) : A_{i,j} \neq 0\}$ . The width of  $\mathcal{D}(A, \theta)$  is defined as

$$\lambda(A, \theta) = \max_{i \in [p]} \left( \sum_{j \in [p]} |A_{i,j}| + |\theta_i| \right).$$

Let  $\eta(A, \theta)$  be the minimum edge weight in absolute value, i.e.,  $\eta(A, \theta) = \min_{i,j \in [p]: A_{i,j} \neq 0} |A_{i,j}|$ .

We note that the Ising model is supported on  $\{-1, 1\}^p$ . A natural generalization is to generalize its support to  $[k]^p$ , and maintain pairwise correlations.

**Definition 2.2.** *The  $p$ -variable pairwise graphical model is a distribution  $\mathcal{D}(\mathcal{W}, \Theta)$  on  $[k]^p$  that satisfies*

$$\Pr(Z = z) \propto \exp \left( \sum_{1 \leq i < j \leq p} W_{i,j}(z_i, z_j) + \sum_{i \in [p]} \theta_i(z_i) \right),$$

where  $\mathcal{W} = \{W_{i,j} \in \mathbb{R}^{k \times k} : i \neq j \in [p]\}$  is a set of weight matrices satisfying  $W_{i,j} = W_{j,i}^T$ , and  $\Theta = \{\theta_i \in \mathbb{R}^k : i \in [p]\}$  is a set of mean-field vectors. The dependency graph of  $\mathcal{D}(\mathcal{W}, \Theta)$  is an undirected graph  $G = (V, E)$ , with vertices  $V = [p]$  and edges  $E = \{(i, j) : W_{i,j} \neq 0\}$ . The width of  $\mathcal{D}(\mathcal{W}, \Theta)$  is defined as

$$\lambda(\mathcal{W}, \Theta) = \max_{i \in [p], a \in [k]} \left( \sum_{j \in [p] \setminus i} \max_{b \in [k]} |W_{i,j}(a, b)| + |\theta_i(a)| \right).$$

Define  $\eta(\mathcal{W}, \Theta) = \min_{(i,j) \in E} \max_{a,b} |W_{i,j}(a, b)|$ .

Both the models above only consider pairwise interactions between nodes. In order to capture higher-order interactions, we examine the more general model of Markov Random Fields (MRFs). In this paper, we will restrict our attention to MRFs over a binary alphabet (i.e., distributions over  $\{\pm 1\}^p$ ). In order to define binary  $t$ -wise MRFs, we first need the following definition of multilinear polynomials, partial derivatives and maximal monomials.

**Definition 2.3.** *Multilinear polynomial is defined as  $h : \mathbb{R}^p \rightarrow \mathbb{R}$  such that  $h(x) = \sum_I \bar{h}(I) \prod_{i \in I} x_i$  where  $\bar{h}(I)$  denotes the coefficient of the monomial  $\prod_{i \in I} x_i$  with respect to the variables  $(x_i : i \in I)$ . Let  $\partial_i h(x) = \sum_{J: i \notin J} \bar{h}(J \cup \{i\}) \prod_{j \in J} x_j$  denote the partial derivative of  $h$  with respect to  $x_i$ . Similarly, for  $I \subseteq [p]$ , let  $\partial_I h(x) = \sum_{J: J \cap I = \emptyset} \bar{h}(J \cup I) \prod_{j \in J} x_j$  denote the partial derivative of  $h$  with respect to the variables  $(x_i : i \in I)$ . We say  $I \subseteq [p]$  is a maximal monomial of  $h$  if  $\bar{h}(J) = 0$  for all  $J \supset I$ .*

Now we are able to formally define binary  $t$ -wise MRFs.

**Definition 2.4.** *For a graph  $G = (V, E)$  on  $p$  vertices, let  $C_t(G)$  denotes all cliques of size at most  $t$  in  $G$ . A binary  $t$ -wise Markov random field on  $G$  is a distribution  $\mathcal{D}$  on  $\{-1, 1\}^p$  which satisfies*

$$\Pr_{Z \sim \mathcal{D}}(Z = z) \propto \exp \left( \sum_{I \in C_t(G)} \varphi_I(z) \right),$$

and each  $\varphi_I : \mathbb{R}^p \rightarrow \mathbb{R}$  is a multilinear polynomial that depends only on the variables in  $I$ .

We call  $G$  the dependency graph of the MRF and  $h(x) = \sum_{I \in C_t(G)} \varphi_I(x)$  the factorization polynomial of the MRF. The width of  $\mathcal{D}$  is defined as  $\lambda = \max_{i \in [p]} \|\partial_i h\|_1$ , where  $\|h\|_1 := \sum_I |\bar{h}(I)|$ .

Now we introduce the definition of  $\delta$ -unbiased distribution and its properties. The proof appears in [KM17].

**Definition 2.5** ( $\delta$ -unbiased). *Let  $S$  be the alphabet set, e.g.,  $S = \{1, -1\}$  for binary  $t$ -pairwise MRFs and  $S = [k]$  for pairwise graphical models. A distribution  $\mathcal{D}$  on  $S^p$  is  $\delta$ -unbiased if for  $Z \sim \mathcal{D}$ ,  $\forall i \in [p]$ , and any assignment  $x \in S^{p-1}$  to  $Z_{-i}$ ,  $\min_{z \in S} \Pr(Z_i = z | Z_{-i} = x) \geq \delta$ .*

The marginal distribution of a  $\delta$ -unbiased distribution also satisfies  $\delta$ -unbiasedness.

**Lemma 2.6.** *Let  $\mathcal{D}$  be a  $\delta$ -unbiased on  $S^p$ , with alphabet set  $S$ . For  $X \sim \mathcal{D}$ ,  $\forall i \in [p]$ , the distribution of  $X_{-i}$  is also  $\delta$ -unbiased.*

The following lemmas provide  $\delta$ -unbiased guarantees for various graphical models.

**Lemma 2.7.** *Let  $\mathcal{D}(\mathcal{W}, \Theta)$  be a pairwise graphical model with alphabet size  $k$  and width  $\lambda(\mathcal{W}, \Theta)$ . Then  $\mathcal{D}(\mathcal{W}, \Theta)$  is  $\delta$ -unbiased with  $\delta = e^{-2\lambda(\mathcal{W}, \Theta)}/k$ . In particular, an Ising model  $\mathcal{D}(A, \theta)$  is  $e^{-2\lambda(A, \theta)}/2$ -unbiased.*

**Lemma 2.8.** *Let  $\mathcal{D}$  be a binary  $t$ -wise MRFs with width  $\lambda$ . Then  $\mathcal{D}$  is  $\delta$ -unbiased with  $\delta = e^{-2\lambda}/2$ .*

Finally, we define two possible goals for learning graphical models. First, the easier goal is *structure learning*, which involves recovering the set of non-zero edges.

**Definition 2.9.** *An algorithm learns the structure of a graphical model if, given samples  $Z_1, \dots, Z_n \sim \mathcal{D}$ , it outputs a graph  $\hat{G} = (V, \hat{E})$  over  $V = [p]$  such that  $\hat{E} = E$ , the set of edges in the dependency graph of  $\mathcal{D}$ .*

The more difficult goal is *parameter learning*, which requires the algorithm to learn not only the location of the edges, but also their parameter values.

**Definition 2.10.** *An algorithm learns the parameters of an Ising model (resp. pairwise graphical model) if, given samples  $Z_1, \dots, Z_n \sim \mathcal{D}$ , it outputs a matrix  $\hat{A}$  (resp. set of matrices  $\hat{W}$ ) such that  $\max_{i,j \in [p]} |A_{i,j} - \hat{A}_{i,j}| \leq \alpha$  (resp.  $|W_{i,j}(a,b) - \hat{W}_{i,j}(a,b)| \leq \alpha$ ,  $\forall i \neq j \in [p], \forall a, b \in [k]$ ).*

**Definition 2.11.** *An algorithm learns the parameters of a binary  $t$ -wise MRF with associated polynomial  $h$  if, given samples  $X^1, \dots, X^n \sim \mathcal{D}$ , it outputs another multilinear polynomial  $u$  such that for all maximal monomial  $I \subseteq [p]$ ,  $|\hat{h}(I) - \bar{u}(I)| \leq \alpha$ .*

## 2.2 Privacy Preliminaries

A dataset  $X = (X^1, \dots, X^n) \in \mathcal{X}^n$  is a collection of points from some universe  $\mathcal{X}$ . We say that two datasets  $X$  and  $X'$  are neighboring, which are denoted as  $X \sim X'$  if they differ in exactly one single point. In our work we consider a few different variants of differential privacy. The first is the standard notion of differential privacy.

**Definition 2.12** (Differential Privacy (DP) [DMNS06]). *A randomized algorithm  $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{S}$  satisfies  $(\epsilon, \delta)$ -differential privacy  $((\epsilon, \delta)$ -DP) if for every pair of neighboring datasets  $X, X' \in \mathcal{X}^n$ , and any event  $S \subseteq \mathcal{S}$ ,*

$$\Pr(\mathcal{A}(X) \in S) \leq e^\epsilon \Pr(\mathcal{A}(X') \in S) + \delta.$$

The second is *concentrated differential privacy* [DR16]. In this work, we specifically consider its refinement *zero-mean concentrated differential privacy* [BS16].



**Definition 2.13** (Concentrated Differential Privacy (zCDP) [BS16]). *A randomized algorithm  $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{S}$  satisfies  $\rho$ -zCDP if for every pair of neighboring datasets  $X, X' \in \mathcal{X}^n$ ,*

$$\forall \alpha \in (1, \infty) \quad D_\alpha (M(X) || M(X')) \leq \rho \alpha,$$

where  $D_\alpha (M(X) || M(X'))$  is the  $\alpha$ -Rényi divergence between  $M(X)$  and  $M(X')$ .

The following lemma quantifies the relationships between  $(\varepsilon, 0)$ -DP,  $\rho$ -zCDP and  $(\varepsilon, \delta)$ -DP.

**Lemma 2.14** (Relationships Between Variants of DP [BS16]). *For every  $\varepsilon \geq 0$ ,*

1. *If  $\mathcal{A}$  satisfies  $(\varepsilon, 0)$ -DP, then  $\mathcal{A}$  is  $\frac{\varepsilon^2}{2}$ -zCDP.*
2. *If  $\mathcal{A}$  satisfies  $\frac{\varepsilon^2}{2}$ -zCDP, then  $\mathcal{A}$  satisfies  $(\frac{\varepsilon^2}{2} + \varepsilon \sqrt{2 \log(\frac{1}{\delta})}, \delta)$ -DP for every  $\delta > 0$ .*

Roughly speaking, pure DP is stronger than zero-concentrated DP, which is stronger than approximate DP.

A crucial property of all the variants of differential privacy is that they can be composed adaptively. By adaptive composition, we mean a sequence of algorithms  $\mathcal{A}_1(X), \dots, \mathcal{A}_T(X)$  where the algorithm  $\mathcal{A}_t(X)$  may also depend on the outcomes of the algorithms  $\mathcal{A}_1(X), \dots, \mathcal{A}_{t-1}(X)$ .

**Lemma 2.15** (Composition of zero-concentrated DP [BS16]). *If  $\mathcal{A}$  is an adaptive composition of differentially private algorithms  $\mathcal{A}_1, \dots, \mathcal{A}_T$ , and  $\mathcal{A}_1, \dots, \mathcal{A}_T$  are  $\rho_1, \dots, \rho_T$ -zCDP respectively, then  $\mathcal{A}$  is  $\rho$ -zCDP for  $\rho = \sum_t \rho_t$ .*

## 3 Parameter Learning of Pairwise Graphical Models

### 3.1 Private Sparse Logistic Regression

As a subroutine of our parameter learning algorithm, we consider the following problem: given a training data set  $D$  consisting of  $n$  pairs of data  $D = \{d^j\}_{j=1}^n = \{(x^j, y^j)\}_{j=1}^n$ , where  $x^j \in \mathbb{R}^p$  and  $y^j \in \mathbb{R}$ , a constraint set  $\mathcal{C} \in \mathbb{R}^p$ , and a loss function  $\ell : \mathcal{C} \times \mathbb{R}^{p+1} \rightarrow \mathbb{R}$ , we want to find  $w^{erm} = \arg \min_{w \in \mathcal{C}} \mathcal{L}(w; D) = \arg \min_{w \in \mathcal{C}} \frac{1}{n} \sum_{j=1}^n \ell(w; d^j)$  with a zCDP constraint. This problem was previously studied in [TTZ14]. Before stating their results, we need the following two definitions. The first definition is regarding Lipschitz continuity.

**Definition 3.1.** *A function  $\ell : \mathcal{C} \rightarrow \mathbb{R}$  is  $L_1$ -Lipschitz with respect to  $\ell_1$  norm, if the following holds.*

$$\forall w_1, w_2 \in \mathcal{C}, |\ell(w_1) - \ell(w_2)| \leq L_1 \|w_1 - w_2\|_1.$$

The performance of the algorithm also depends on the ‘‘curvature’’ of the loss function, which is defined below, based on the definition of [Cla10, Jag13]. A side remark is that this is a strictly weaker constraint than smoothness [TTZ14].

**Definition 3.2** (Curvature constant). *For  $\ell : \mathcal{C} \rightarrow \mathbb{R}$ ,  $\Gamma_\ell$  is defined as*

$$\Gamma_\ell = \sup_{w_1, w_2 \in \mathcal{C}, \gamma \in (0, 1], w_3 = w_1 + \gamma(w_2 - w_1)} \frac{2}{\gamma^2} (\ell(w_3) - \ell(w_1) - \langle w_3 - w_1, \nabla \ell(w_1) \rangle).$$

Now we are able to introduce the algorithm and its theoretical guarantees.

<b>Algorithm 1:</b> $\mathcal{A}_{PFW}(D, \mathcal{L}, \rho, \mathcal{C})$ : Private Frank-Wolfe Algorithm	
<b>Input:</b> Data set: $D = \{d^1, \dots, d^n\}$ , loss function: $\mathcal{L}(w; D) = \frac{1}{n} \sum_{j=1}^n \ell(w; d^j)$ (with Lipschitz constant $L_1$ ), privacy parameters: $\rho$ , convex set: $\mathcal{C} = \text{conv}(S)$ with $\ \mathcal{C}\ _1 := \max_{s \in S} \ s\ _1$ , iteration times: $T$	
1	Initialize $w$ from an arbitrary point in $\mathcal{C}$
2	<b>For</b> $t = 1$ to $T - 1$
3	$\forall s \in S, \alpha_s \leftarrow \langle s, \nabla \mathcal{L}(w; D) \rangle + \text{Lap}\left(0, \frac{L_1 \ \mathcal{C}\ _1 \sqrt{T}}{n\sqrt{\rho}}\right)$
4	$\tilde{w}_t \leftarrow \arg \min_{s \in S} \alpha_s$
5	$w_{t+1} \leftarrow (1 - \mu_t)w_t + \mu_t \tilde{w}_t$ , where $\mu_t = \frac{2}{t+2}$
<b>Output:</b> $w^{\text{priv}} = w_T$	

**Lemma 3.3** (Theorem 5.5 from [TTZ14]). *Algorithm 1 satisfies  $\rho$ -zCDP. Furthermore, let  $L_1, \|\mathcal{C}\|_1$  be defined as in Algorithm 1. Let  $\Gamma_\ell$  be an upper bound on the curvature constant for the loss function  $\ell(\cdot; d)$  for all  $d$  and  $|S|$  be the number of extreme points in  $S$ . If we set  $T = \frac{\Gamma_\ell^{\frac{2}{3}} (n\sqrt{\rho})^{\frac{2}{3}}}{L_1 \|\mathcal{C}\|_1^{\frac{2}{3}}}$ , then with probability at least  $1 - \beta$  over the randomness of the algorithm,*

$$\mathcal{L}(w^{\text{priv}}; D) - \mathcal{L}(w^{\text{erm}}; D) = O\left(\frac{\Gamma_\ell^{\frac{1}{3}} (L_1 \|\mathcal{C}\|_1)^{\frac{2}{3}} \log\left(\frac{n|S|}{\beta}\right)}{(n\sqrt{\rho})^{\frac{2}{3}}}\right).$$

*Proof.* The utility guarantee is proved in [TTZ14]. Therefore, it is enough to prove the algorithm satisfies  $\rho$ -zCDP. According to the definition of the Laplace mechanism, every iteration of the algorithm satisfies  $(\sqrt{\frac{\rho}{T}}, 0)$ -DP, which naturally satisfies  $\frac{\rho}{T}$ -zCDP by Lemma 2.14. Then, by the composition theorem of zCDP (Lemma 2.15), the algorithm satisfies  $\rho$ -zCDP.  $\square$

If we consider the specific problem of sparse logistic regression, we will get the following corollary.

**Corollary 3.4.** *If we consider the problem of sparse logistic regression, i.e.,  $\mathcal{L}(w; D) = \frac{1}{n} \sum_{j=1}^n \log(1 + e^{-y^j \langle w, x^j \rangle})$ , with the constraint that  $\mathcal{C} = \{w : \|w\|_1 \leq \lambda\}$ , and we further assume that  $\forall j, \|x^j\|_\infty \leq 1, y^j \in \{\pm 1\}$ , let  $T = \lambda^{\frac{2}{3}} (n\sqrt{\rho})^{\frac{2}{3}}$ , then with probability at least  $1 - \beta$  over the randomness of the algorithm,*

$$\mathcal{L}(w^{\text{priv}}; D) - \mathcal{L}(w^{\text{erm}}; D) = O\left(\frac{\lambda^{\frac{4}{3}} \log\left(\frac{np}{\beta}\right)}{(n\sqrt{\rho})^{\frac{2}{3}}}\right).$$

*Furthermore, the time complexity of the algorithm is  $O(T \cdot (np + p^2)) = O\left(n^{\frac{2}{3}} \cdot (np + p^2)\right)$ .*

*Proof.* First let we show  $L_1 \leq 2$ . If we fix sample  $d = (x, y)$ , then for any  $w_1, w_2 \in \mathcal{C}$ ,

$$|\ell(w_1; d) - \ell(w_2; d)| \leq \max_w \|\nabla_w(\ell(w; d))\|_\infty \cdot \|w_1 - w_2\|_1.$$

Since  $\nabla_w(\ell(w; d)) = (\sigma(\langle w, x \rangle) - y) \cdot x$ , we have  $\|\nabla_w(\ell(w; d))\|_\infty \leq 2$ .

Next, we wish to show  $\Gamma_\ell \leq \lambda^2$ . We use the following lemma from [TTZ14].

**Lemma 3.5** (Remark 4 in [TTZ14]). *For any  $q, r \geq 1$  such that  $\frac{1}{q} + \frac{1}{r} = 1$ ,  $\Gamma_\ell$  is upper bounded by  $\alpha \|\mathcal{C}\|_q^2$ , where  $\alpha = \max_{w \in \mathcal{C}, \|v\|_q=1} \|\nabla^2 \ell(w) \cdot v\|_q$ .*

If we take  $q = 1, r = +\infty$ , then  $\Gamma_\ell \leq \alpha\lambda^2$ , where

$$\alpha = \max_{w \in \mathcal{C}, \|v\|_1=1} \|\nabla^2 \ell(w; d) \cdot v\|_\infty \leq \max_{i,j \in [p]} (\nabla^2 \ell(w; d))_{i,j}.$$

We have  $\alpha \leq 1$ , since  $\nabla^2 \ell(w; d) = \sigma(\langle w, x \rangle)(1 - \sigma(\langle w, x \rangle)) \cdot xx^T$ , and  $\|x\|_\infty \leq 1$ ,

Finally given  $\mathcal{C} = \{w : \|w\|_1 \leq 1\}$ , the number of extreme points of  $S$  equals  $2p$ . By replacing all these parameters in Lemma 3.3, we have proved the loss guarantee in the corollary.

With respect to the time complexity, we note that the time complexity of each iteration is  $O(np + p^2)$  and there are  $T$  iterations in total.  $\square$

Now if we further assume the data set  $D$  is drawn i.i.d. from some underlying distribution  $P$ , the following lemma from learning theory relates the true risk and the empirical risk, which shall be heavily used in the following sections.

**Theorem 3.6.** *If we consider the same problem setting and assumptions as in Corollary 3.4, and we further assume that the training data set  $D$  is drawn i.i.d. from some unknown distribution  $P$ , then with probability at least  $1 - \beta$  over the randomness of the algorithm and the training data set,*

$$\mathbb{E}_{(X,Y) \sim P} [\ell(w^{priv}; (X, Y))] - \mathbb{E}_{(X,Y) \sim P} [\ell(w^*; (X, Y))] = O\left(\frac{\lambda^{\frac{4}{3}} \log(\frac{np}{\beta})}{(n\sqrt{\rho})^{\frac{2}{3}}} + \frac{\lambda \log(\frac{1}{\beta})}{\sqrt{n}}\right),$$

where  $w^* = \arg \min_{w \in \mathcal{C}} \mathbb{E}_{(X,Y) \sim P} [\ell(w; (X, Y))]$ .

*Proof.* By triangle inequality,

$$\begin{aligned} & \mathbb{E}_{(X,Y) \sim P} [\ell(w^{priv}; (X, Y))] - \mathbb{E}_{(X,Y) \sim P} [\ell(w^*; (X, Y))] \\ & \leq \left| \mathbb{E}_{(X,Y) \sim P} [\ell(w^{priv}; (X, Y))] - \frac{1}{n} \sum_{m=1}^n \ell(w^{priv}; d^m) \right| + \left| \frac{1}{n} \sum_{m=1}^n \ell(w^{priv}; d^m) - \frac{1}{n} \sum_{m=1}^n \ell(w^{erm}; d^m) \right| \\ & + \left( \frac{1}{n} \sum_{m=1}^n \ell(w^{erm}; d^m) - \frac{1}{n} \sum_{m=1}^n \ell(w^*; d^m) \right) + \left| \mathbb{E}_{(X,Y) \sim P} [\ell(w^*; (X, Y))] - \frac{1}{n} \sum_{m=1}^n \ell(w^*; d^m) \right| \end{aligned}$$

Now we need to bound each term. We firstly bound the first and last term simultaneously. By the generalization error bound (Lemma 7 from [WSD19]), they are bounded by  $O\left(\frac{\lambda \log(\frac{1}{\beta})}{\sqrt{n}}\right)$  simultaneously, with probability greater than  $1 - \frac{2}{3}\beta$ . Then we turn to the second term, by Corollary 3.4, with probability greater than  $1 - \frac{1}{3}\beta$ , it is bounded by  $O\left(\frac{\lambda^{\frac{4}{3}} \log(\frac{np}{\beta})}{(n\sqrt{\rho})^{\frac{2}{3}}}\right)$ . Finally we bound the third term. According to the definition of  $w^{erm}$ , the third term should be smaller than 0. Therefore, by union bound,  $\mathbb{E}_{(X,Y) \sim P} [\ell(w^{priv}; (X, Y))] - \mathbb{E}_{(X,Y) \sim P} [\ell(w^*; (X, Y))] = O\left(\frac{\lambda^{\frac{4}{3}} \log(\frac{np}{\beta})}{(n\sqrt{\rho})^{\frac{2}{3}}} + \frac{\lambda \log(\frac{1}{\beta})}{\sqrt{n}}\right)$ , with probability greater than  $1 - \beta$ .  $\square$

## 3.2 Privately Learning Ising Models

We first consider the problem of estimating the weight matrix of the Ising model. To be precise, given  $n$  i.i.d. samples  $\{z^1, \dots, z^n\}$  generated from an unknown distribution  $\mathcal{D}(A, \theta)$ , our goal is to design an  $\rho$ -zCDP estimator  $\hat{A}$  such that with probability at least  $\frac{2}{3}$ ,  $\max_{i,j \in [p]} |A_{i,j} - \hat{A}_{i,j}| \leq \alpha$ .

An observation of the Ising model is that for any node  $Z_i$ , the probability of  $Z_i = 1$  conditioned on the values of the remaining nodes  $Z_{-i}$  follows from a sigmoid function. The next lemma comes from [KM17], which formalizes this observation.

**Lemma 3.7.** *Let  $Z \sim \mathcal{D}(A, \theta)$  and  $Z \in \{-1, 1\}^p$ , then  $\forall i \in [p], \forall x \in \{-1, 1\}^{[p] \setminus \{i\}}$ ,*

$$\Pr(Z_i = 1 | Z_{-i} = x) = \sigma \left( \sum_{j \neq i} 2A_{i,j}x_j + 2\theta_i \right) = \sigma(\langle w, x' \rangle).$$

where  $w = 2[A_{i,1}, \dots, A_{i,i-1}, A_{i,i+1}, \dots, A_{i,p}, \theta_i] \in \mathbb{R}^p$ , and  $x' = [x, 1] \in \mathbb{R}^p$ .

*Proof.* The proof is from [KM17], and we include it here for completeness. According to the definition of the Ising model,

$$\begin{aligned} \Pr(Z_i = 1 | Z_{-i} = x) &= \frac{\exp \left( \sum_{j \neq i} A_{i,j}x_j + \sum_{j \neq i} \theta_j + \theta_i \right)}{\exp \left( \sum_{j \neq i} A_{i,j}x_j + \sum_{j \neq i} \theta_j + \theta_i \right) + \exp \left( \sum_{j \neq i} -A_{i,j}x_j + \sum_{j \neq i} \theta_j - \theta_i \right)} \\ &= \sigma \left( \sum_{j \neq i} 2A_{i,j}x_j + 2\theta_i \right). \end{aligned}$$

□

By Lemma 3.7, we can estimate the weight matrix by solving a logistic regression for each node, which is utilized in [WSD19] to design non-private estimators. Our algorithm uses the private Frank-Wolfe method to solve the per-node logistic regression problem, achieving the following theoretical guarantee.

**Algorithm 2:** Privately Learning Ising Models

**Input:**  $n$  samples  $\{z^1, \dots, z^n\}$ , where  $z^m \in \{\pm 1\}^p$  for  $m \in [n]$ ; an upper bound on  $\lambda(A, \theta) \leq \lambda$ , privacy parameter  $\rho$

1 **For**  $i = 1$  to  $p$

2      $\forall m \in [n], x^m \leftarrow [z_{-i}^m, 1], y^m \leftarrow z_i^m$

3      $w^{priv} \leftarrow \mathcal{A}_{PFW}(D, \mathcal{L}, \rho', \mathcal{C})$ , where  $\rho' = \frac{\rho}{p}$ ,  $D = \{(x^m, y^m)\}_{m=1}^n$ ,

$\mathcal{L}(w; D) = \frac{1}{n} \sum_{m=1}^n \log(1 + e^{-y^m \langle w, x^m \rangle})$ ,  $\mathcal{C} = \{\|w\|_1 \leq 2\lambda\}$

4      $\forall j \in p, \hat{A}_{i,j} \leftarrow \frac{1}{2} w_j^{priv}$ , where  $\tilde{j} = j$  when  $j < i$  and  $\tilde{j} = j - 1$  if  $j > i$

**Output:**  $\hat{A} \in \mathbb{R}^{p \times p}$

**Theorem 3.8.** *Let  $\mathcal{D}(A, \theta)$  be an unknown  $p$ -variable Ising model with  $\lambda(A, \theta) \leq \lambda$ . There exists an efficient  $\rho$ - $z$ CDP algorithm which outputs a weight matrix  $\hat{A} \in \mathbb{R}^{p \times p}$  such that with probability greater than  $2/3$ ,  $\max_{i,j \in [p]} |A_{i,j} - \hat{A}_{i,j}| \leq \alpha$  if the number of i.i.d. samples satisfies*

$$n = \Omega \left( \frac{\lambda^2 \log(p) e^{12\lambda}}{\alpha^4} + \frac{\sqrt{p} \lambda^2 \log^2(p) e^{9\lambda}}{\sqrt{\rho} \alpha^3} \right).$$

*Proof.* We first prove that Algorithm 2 satisfies  $\rho$ -zCDP. Notice that in each iteration, the algorithm solves a private sparse logistic regression under  $\frac{\rho}{p}$ -zCDP. Therefore, Algorithm 2 satisfies  $\rho$ -zCDP by composition (Lemma 2.15).

For the accuracy analysis, we start by looking at the first iteration ( $i = 1$ ) and showing that  $|A_{1,j} - \hat{A}_{1,j}| \leq \alpha, \forall j \in [p]$ , with probability greater than  $1 - \frac{1}{10p}$ .

Given a random sample  $Z \sim \mathcal{D}(A, \theta)$ , we let  $X = [Z_{-1}, 1], Y = Z_1$ . From Lemma 3.7,  $\Pr(Y = 1 | X = x) = \sigma(\langle w^*, x \rangle)$ , where  $w^* = 2[A_{1,2}, \dots, A_{1,p}, \theta_1]$ . We also note that  $\|w^*\|_1 \leq 2\lambda$ , as a consequence of the width constraint of the Ising model.

For any  $n$  i.i.d. samples  $\{z^m\}_{m=1}^n$  drawn from the Ising model, let  $x^m = [z_{-1}^m, 1]$  and  $y^m = z_1^m$ , it is easy to check that each  $(x^m, y^m)$  is the realization of  $(X, Y)$ . Let  $w^{priv}$  be the output of  $\mathcal{A}(D, \mathcal{L}, \frac{\rho}{p}, \{w : \|w\|_1 \leq 2\lambda\})$ , where  $D = \{(x^m, y^m)\}_{m=1}^n$ . By Lemma 3.6, when  $n = O\left(\frac{\sqrt{p}\lambda^2 \log^2(p)}{\sqrt{p}\gamma^{\frac{3}{2}}} + \frac{\lambda^2 \log(p)}{\gamma^2}\right)$ , with probability greater than  $1 - \frac{1}{10p}$ ,  $\mathbb{E}_{Z \sim \mathcal{D}(A, \theta)} [\ell(w^{priv}; (X, Y))] - \mathbb{E}_{Z \sim \mathcal{D}(A, \theta)} [\ell(w^*; (X, Y))] \leq \gamma$ .

We will use the following lemma from [WSD19]. Roughly speaking, with the assumption that the samples are generated from an Ising model, any estimator  $w^{priv}$  which achieves a small error in the loss  $\mathcal{L}$  guarantees an accurate parameter recovery in  $\ell_\infty$  distance.

**Lemma 3.9.** *Let  $P$  be a distribution on  $\{-1, 1\}^{p-1} \times \{-1, 1\}$ . Given  $u_1 \in \mathbb{R}^{p-1}, \theta_1 \in \mathbb{R}$ , suppose  $\Pr(Y = 1 | X = x) = \sigma(\langle u_1, x \rangle + \theta_1)$  for  $(X, Y) \sim P$ . If the marginal distribution of  $P$  on  $X$  is  $\delta$ -unbiased, and  $\mathbb{E}_{(X, Y) \sim P} [\log(1 + e^{-Y(\langle u_1, X \rangle + \theta_1)})] - \mathbb{E}_{(X, Y) \sim P} [\log(1 + e^{-Y(\langle u_2, X \rangle + \theta_2)})] \leq \gamma$  for some  $u_2 \in \mathbb{R}^{p-1}, \theta_2 \in \mathbb{R}$ , and  $\gamma \leq \delta e^{-2\|u_1\|_1 - 2\|\theta_1\|_1 - 6}$ , then  $\|u_1 - u_2\|_\infty = O(e^{\|u_1\|_1 + \|\theta_1\|_1} \cdot \sqrt{\gamma/\delta})$ .*

By Lemma 2.6, Lemma 2.7 and Lemma 3.9, if  $\mathbb{E}_{Z \sim \mathcal{D}(A, \theta)} [\ell(w^{priv}; (X, Y))] - \mathbb{E}_{Z \sim \mathcal{D}(A, \theta)} [\ell(w^*; (X, Y))] \leq O(\alpha^2 e^{-6\lambda})$ , we have  $\|w^{priv} - w^*\|_\infty \leq \alpha$ . By replacing  $\gamma = \alpha^2 e^{-6\lambda}$ , we prove that  $|A_{1,j} - \hat{A}_{1,j}|_\infty \leq \alpha$  with probability greater than  $1 - \frac{1}{10p}$ . Noting that similar argument works for the other iterations and non-overlapping part of the matrix is recovered in different iterations. By union bound over  $p$  iterations, we prove that with probability at least  $\frac{2}{3}$ ,  $\max_{i,j \in [p]} |A_{i,j} - \hat{A}_{i,j}| \leq \alpha$ .

Finally, we note that the time complexity of the algorithm is  $poly(n, p)$  since the private Frank-Wolfe algorithm is time efficient by Corollary 3.4.  $\square$

### 3.3 Privately Learning Pairwise Graphical Models

Next, we study parameter learning for pairwise graphical models over general alphabet. Given  $n$  i.i.d. samples  $\{z^1, \dots, z^n\}$  drawn from an unknown distribution  $\mathcal{D}(\mathcal{W}, \Theta)$ , we want to design an  $\rho$ -zCDP estimator  $\hat{\mathcal{W}}$  such that with probability at least  $\frac{2}{3}$ ,  $\forall i \neq j \in [p], \forall u, v \in [k], |W_{i,j}(u, v) - \widehat{W}_{i,j}(u, v)| \leq \alpha$ . To facilitate our presentation, we assume that  $\forall i \neq j \in [p]$ , every row (and column) vector of  $W_{i,j}$  has zero mean.<sup>2</sup>

Analogous to Lemma 3.7 for the Ising model, a pairwise graphical model has the following property, which can be utilized to recover its parameters.

**Lemma 3.10** (Fact 2 of [WSD19]). *Let  $Z \sim \mathcal{D}(\mathcal{W}, \Theta)$  and  $Z \in [k]^p$ . For any  $i \in [p]$ , any  $u \neq v \in [k]$ , and any  $x \in [k]^{p-1}$ ,*

<sup>2</sup>The assumption that  $W_{i,j}$  is centered is without loss of generality and widely used in the literature [KM17, WSD19]. We present the argument here for completeness. Suppose the  $a$ -th row of  $W_{i,j}$  is not centered, i.e.,  $\sum_b W_{i,j}(a, b) \neq 0$ , we can define  $W'_{i,j}(a, b) = W_{i,j}(a, b) - \frac{1}{k} \sum_b W_{i,j}(a, b)$  and  $\theta'_i(a) = \theta_i(a) + \frac{1}{k} \sum_b W_{i,j}(a, b)$ , and the probability distribution remains unchanged.

$$\Pr(Z_i = u | Z_i \in \{u, v\}, Z_{-i} = x) = \sigma \left( \sum_{j \neq i} (W_{i,j}(u, x_j) - W_{i,j}(v, x_j)) + \theta_i(u) - \theta_i(v) \right).$$

Now we introduce our algorithm. Without loss of generality, we consider estimating  $W_{1,j}$  for all  $j \in [p]$  as a running example. We fix a pair of values  $(u, v)$ , where  $u, v \in [k]$  and  $u \neq v$ . Let  $S_{u,v}$  be the samples where  $Z_1 \in \{u, v\}$ . In order to utilize Lemma 3.10, we perform the following transformation on the samples in  $S_{u,v}$ : for the  $m$ -th sample  $z^m$ , let  $y^m = 1$  if  $z_1^m = u$ , else  $y^m = -1$ . And  $x^m$  is the one-hot encoding of the vector  $[z_{-1}^m, 1]$ , where  $\text{OneHotEncode}(s)$  is a mapping from  $[k]^p$  to  $\mathbb{R}^{p \times k}$ , and the  $i$ -th row is the  $t$ -th standard basis vector given  $s_i = t$ . Then we define  $w^* \in \mathbb{R}^{p \times k}$  as follows:

$$\begin{aligned} w^*(j, \cdot) &= W_{1,j+1}(u, \cdot) - W_{1,j+1}(v, \cdot), \forall j \in [p-1]; \\ w^*(p, \cdot) &= [\theta_1(u) - \theta_1(v), 0, \dots, 0]. \end{aligned}$$

Lemma 3.10 implies that  $\forall t, \Pr(Y^t = 1) = \sigma(\langle w^*, X^t \rangle)$ , where  $\langle \cdot, \cdot \rangle$  is the element-wise multiplication of matrices. According to the definition of the width of  $\mathcal{D}(\mathcal{W}, \Theta)$ ,  $\|w^*\|_1 \leq \lambda k$ . Now we can apply the sparse logistic regression method of Algorithm 3 to the samples in  $S_{u,v}$ .

Suppose  $w_{u,v}^{priv}$  is the output of the private Frank-Wolfe algorithm, we define  $U_{u,v} \in \mathbb{R}^{p \times k}$  as follows:  $\forall b \in [k]$ ,

$$\begin{aligned} U_{u,v}(j, b) &= w_{u,v}^{priv}(j, b) - \frac{1}{k} \sum_{a \in [k]} w_{u,v}^{priv}(j, a), \forall j \in [p-1]; \\ U_{u,v}(p, b) &= w_{u,v}^{priv}(p, b) + \frac{1}{k} \sum_{j \in [p-1]} \sum_{a \in [k]} w_{u,v}^{priv}(j, a). \end{aligned} \quad (1)$$

$U_{u,v}$  can be seen as a ‘‘centered’’ version of  $w_{u,v}^{priv}$  (for the first  $p-1$  rows). It is not hard to see that  $\langle U_{u,v}, x \rangle = \langle w_{u,v}^{priv}, x \rangle$ , so  $U_{u,v}$  is also a minimizer of the sparse logistic regression.

For now, assume that  $\forall j \in [p-1], b \in [k], U_{u,v}(j, b)$  is a ‘‘good’’ approximation of  $(W_{1,j+1}(u, b) - W_{1,j+1}(v, b))$ , which we will show later. If we sum over  $v \in [k]$ , it can be shown that  $\frac{1}{k} \sum_{v \in [k]} U_{u,v}(j, b)$  is also a ‘‘good’’ approximation of  $W_{1,j+1}(u, b)$ , for all  $j \in [p-1]$ , and  $u, b \in [k]$ , because of the centering assumption of  $\mathcal{W}$ , i.e.,  $\forall j \in [p-1], b \in [k], \sum_{v \in [k]} W_{1,j+1}(v, b) = 0$ . With these considerations in mind, we are able to introduce our algorithm.

**Algorithm 3:** Privately Learning Pairwise Graphical Model

**Input:** alphabet size  $k$ ,  $n$  i.i.d. samples  $\{z^1, \dots, z^n\}$ , where  $z^m \in [k]^p$  for  $m \in [n]$ ; an upper bound on  $\lambda(\mathcal{W}, \Theta) \leq \lambda$ , privacy parameter  $\rho$

1 **For**  $i = 1$  to  $p$

2     **For** each pair  $u \neq v \in [k]$

3          $S_{u,v} \leftarrow \{z^m, m \in [n] : z_i^m \in \{u, v\}\}$

4          $\forall z^m \in S_{u,v}, x^m \leftarrow \text{OneHotEncode}([z_{-i}^m, 1]), y^m \leftarrow 1$  if  $z_i^m = u; y^m \leftarrow -1$  if  $z_i^m = v$

5          $w_{u,v}^{priv} \leftarrow \text{APFW}(D, \mathcal{L}, \rho', \mathcal{C})$ , where  $\rho' = \frac{\rho}{k^2 p}$ ,  $D = \{(x^m, y^m) : z^m \in S_{u,v}\}$ ,

$$\mathcal{L}(w; D) = \frac{1}{|S_{u,v}|} \sum_{m=1}^{|S_{u,v}|} \log(1 + e^{-y^m \langle w, x^m \rangle}), \mathcal{C} = \{\|w\|_1 \leq 2\lambda k\}$$

6         Define  $U_{u,v} \in \mathbb{R}^{p \times k}$  by centering the first  $p-1$  rows of  $w_{u,v}^{priv}$ , as in Equation 1

7     **For**  $j \in [p] \setminus i$  and  $u \in [k]$

8          $\widehat{W}_{i,j}(u, \cdot) \leftarrow \frac{1}{k} \sum_{v \in [k]} U_{u,v}(\tilde{j}, \cdot)$ , where  $\tilde{j} = j$  when  $j < i$  and  $\tilde{j} = j-1$  when  $j > i$

**Output:**  $\widehat{W}_{i,j} \in \mathbb{R}^{k \times k}$  for all  $i \neq j \in [p]$

The following theorem is the main result of this section. Its proof is structurally similar to that of Theorem 3.8.

**Theorem 3.11.** *Let  $\mathcal{D}(\mathcal{W}, \Theta)$  be an unknown  $p$ -variable pairwise graphical model distribution, and we suppose that  $\mathcal{D}(\mathcal{W}, \Theta)$  has width  $\lambda(\mathcal{W}, \Theta) \leq \lambda$ . There exists an efficient  $\rho$ -zCDP algorithm which outputs  $\widehat{W}$  such that with probability greater than  $2/3$ ,  $|W_{i,j}(u, v) - \widehat{W}_{i,j}(u, v)| \leq \alpha$ ,  $\forall i \neq j \in [p], \forall u, v \in [k]$  if the number of i.i.d. samples satisfy*

$$n = \Omega\left(\frac{\lambda^2 k^5 \log(pk) e^{O(\lambda)}}{\alpha^4} + \frac{\sqrt{p} \lambda^2 k^{5.5} \log^2(pk) e^{O(\lambda)}}{\sqrt{\rho} \alpha^3}\right).$$

*Proof.* We consider estimating  $W_{1,j}$  for all  $j \in [p]$  as an example. Fixing one pair  $(u, v)$ , let  $S_{u,v}$  be the samples whose first element is either  $u$  or  $v$ , and  $n^{u,v}$  be the number of samples in  $S_{u,v}$ . We perform the following transformation on the samples in  $S_{u,v}$ : for the sample  $Z$ , let  $Y = 1$  if  $Z_1 = u$ , else  $Y = -1$ , and let  $X$  be the one-hot encoding of the vector  $[Z_{-1}, 1]$ .

Suppose the underlying joint distribution of  $X$  and  $Y$  is  $P$ , i.e.,  $(X, Y) \sim P$ , then by Theorem 3.6, when  $n^{u,v} = O\left(\frac{\lambda^2 k^2 \log^2(pk)}{\gamma^2} + \frac{\sqrt{d} \lambda^2 k^3 \log^2(pk)}{\gamma^{\frac{3}{2}} \sqrt{\rho}}\right)$ , with probability greater than  $1 - \frac{1}{10pk^2}$ ,

$$\mathbb{E}_{(X,Y) \sim P} [\ell(U_{u,v}; (X, Y))] - \mathbb{E}_{(X,Y) \sim P} [\ell(w^*; (X, Y))] \leq \gamma.$$

The following lemma appears in [WSD19], which is analogous to Lemma 3.9 for the Ising model.

**Lemma 3.12.** *Let  $\mathcal{D}$  be a  $\delta$ -unbiased distribution on  $[k]^{p-1}$ . For  $Z \sim \mathcal{D}$ ,  $X$  denotes the one-hot encoding of  $Z$ . Let  $u_1, u_2 \in \mathbb{R}^{(p-1) \times k}$  be two matrices where  $\sum_a u_1(i, a) = 0$  and  $\sum_a u_2(i, a) = 0$  for all  $i \in [p-1]$ . Let  $P$  be a distribution such that given  $u_1, \theta_1 \in \mathbb{R}$ ,  $\Pr(Y = 1 | X = X) = \sigma(\langle u_1, x \rangle + \theta_1)$  for  $(X, Y) \sim P$ . Suppose  $\mathbb{E}_{(X,Y) \sim P} [\log(1 + e^{-Y(\langle u_1, X \rangle + \theta_1)})] - \mathbb{E}_{(X,Y) \sim P} [\log(1 + e^{-Y(\langle u_2, X \rangle + \theta_2)})] \leq \gamma$  for  $u_2 \in \mathbb{R}^{(p-1) \times k}, \theta_2 \in \mathbb{R}$ , and  $\gamma \leq \delta e^{-2\|u_1\|_{\infty,1} - 2\|\theta_1\|_1 - 6}$ , then  $\|u_1 - u_2\|_{\infty} = O(e^{\|u_1\|_{\infty,1} + \|\theta_1\|_1} \sqrt{\gamma/\delta})$ .*

By Lemma 2.6, Lemma 2.7 and Lemma 3.12, if we substitute  $\gamma = \frac{e^{-6\lambda} \alpha^2}{k}$ , when  $n^{u,v} = O\left(\frac{\lambda^2 k^4 \log(pk) e^{O(\lambda)}}{\alpha^4} + \frac{\sqrt{p} \lambda^2 k^{4.5} \log^2(pk) e^{O(\lambda)}}{\sqrt{\rho} \alpha^3}\right)$ ,

$$|W_{1,j}(u, b) - W_{1,j}(v, b) - U^{u,v}(j, b)| \leq \alpha, \forall j \in [p-1], \forall b \in [k]. \quad (2)$$

By a union bound, Equation (2) holds for all  $(u, v)$  pairs simultaneously with probability greater than  $1 - \frac{1}{10p}$ . If we sum over  $v \in [k]$  and use the fact that  $\forall j, b, \sum_{v \in [k]} W_{1,j}(v, b) = 0$ , we have

$$\left| W_{1,j}(u, b) - \frac{1}{k} \sum_{v \in [k]} U_{u,v}(j, b) \right| \leq \alpha, \forall j \in [p-1], \forall u, b \in [k].$$

Note that we need to guarantee that we obtain  $n^{u,v}$  samples for each pair  $(u, v)$ . Since  $\mathcal{D}(\mathcal{W}, \Theta)$  is  $\delta$ -unbiased, given  $Z \sim \mathcal{D}(\mathcal{W}, \Theta)$ , for all  $u \neq v$ ,  $\Pr(Z \in S_{u,v}) \geq 2\delta$ . By Hoeffding's inequality, when  $n = O\left(\frac{n^{u,v}}{\delta} + \frac{\log(pk^2)}{\delta^2}\right)$ , with probability greater than  $1 - \frac{1}{10p}$ , we have enough samples for all  $(u, v)$  pairs simultaneously. Substituting  $\delta = \frac{e^{-6\lambda}}{k}$ , we have

$$n = O\left(\frac{\lambda^2 k^5 \log(pk) e^{O(\lambda)}}{\alpha^4} + \frac{\sqrt{p} \lambda^2 k^{5.5} \log^2(pk) e^{O(\lambda)}}{\sqrt{\rho} \alpha^3}\right).$$

The same argument holds for other entries of the matrix. We conclude the proof by a union bound over  $p$  iterations.

Finally, we note that the time complexity of the algorithm is  $\text{poly}(n, p)$  since the private Frank-Wolfe algorithm is time efficient by Corollary 3.4.  $\square$

## 4 Privately Learning Binary $t$ -wise MRFs

Let  $\mathcal{D}$  be a  $t$ -wise MRF on  $\{1, -1\}^p$  with underlying dependency graph  $G$  and factorization polynomial  $h(x) = \sum_{I \in \mathcal{C}_t(G)} h_I(x)$ . We assume that the width of  $\mathcal{D}$  is bounded by  $\lambda$ , i.e.,  $\max_{i \in [p]} \|\partial_i h\|_1 \leq \lambda$ , where  $\|h\|_1 := \sum_{I \in \mathcal{C}_t(G)} |\bar{h}(I)|$ . Similar to [KM17], given  $n$  i.i.d. samples  $\{z^1, \dots, z^n\}$  generated from an unknown distribution  $\mathcal{D}$ , we consider the following two related learning objectives, under the constraint of  $\rho$ -zCDP:

1. find a multilinear polynomial  $u$  such that with probability greater than  $\frac{2}{3}$ ,  $\|h - u\|_1 := \sum_{I \in \mathcal{C}_t(G)} |\bar{h}(I) - \bar{u}(I)| \leq \alpha$ ;
2. find a multilinear polynomial  $u$  such that with probability greater than  $\frac{2}{3}$ , for every maximal monomial  $I$  of  $h$ ,  $|\bar{h}(I) - \bar{u}(I)| \leq \alpha$ .

We note that our first objective can be viewed as parameter estimation in  $\ell_1$  distance, where only an average performance guarantee is provided. In the second objective, the algorithm recovers every maximal monomial, which can be viewed as parameter estimation in  $\ell_\infty$  distance. These two objectives are addressed in Sections 4.1 and 4.2, respectively.

### 4.1 Parameter Estimation in $\ell_1$ Distance

The following property of MRFs, from [KM17], plays a critical role in our algorithm. The proof is similar to that of Lemma 3.7.

**Lemma 4.1** (Lemma 7.6 of [KM17]). *Let  $\mathcal{D}$  be a  $t$ -wise MRF on  $\{1, -1\}^p$  with underlying dependency graph  $G$  and factorization polynomial  $h(x) = \sum_{I \in \mathcal{C}_t(G)} h_I(x)$ , then*

$$\Pr(Z_i = 1 | Z_{-i} = x) = \sigma(2\partial_i h(x)), \forall i \in [p], \forall x \in \{1, -1\}^{[p] \setminus i}.$$

Lemma 4.1 shows that, similar to pairwise graphical models, it also suffices to learn the parameters of binary  $t$ -wise MRF using sparse logistic regression.

#### Algorithm 4: Private Learning binary $t$ -wise MRF in $\ell_1$ distance

**Input:**  $n$  i.i.d. samples  $\{z^1, \dots, z^n\}$ , where  $z^m \in \{\pm 1\}^p$  for  $m \in [n]$ ; an upper bound  $\lambda$  on  $\max_{i \in [p]} \|\partial_i h\|_1$ , privacy parameter  $\rho$

- 1 **For**  $i = 1$  **to**  $p$
- 2      $\forall m \in [n], x_m \leftarrow \left[ \prod_{j \in I} z_j^m : I \subset [p] \setminus i, |I| \leq t - 1 \right], y_m \leftarrow z_i^m$
- 3      $w^{\text{priv}} \leftarrow \mathcal{A}(D, \mathcal{L}, \rho', \mathcal{C})$  where  $D = \{(x_m, y_m)\}_{m=1}^n$ ,  $\ell(w; d) = \log(1 + e^{-y\langle w, x \rangle})$ ,  $\mathcal{C} = \{\|w\|_1 \leq 2\lambda\}$ , and  $\rho' = \frac{\rho}{p}$
- 4     **For**  $I \subset [p] \setminus i$  **with**  $|I| \leq t - 1$
- 5          $\bar{u}(I \cup \{i\}) = \frac{1}{2} w^{\text{priv}}(I)$ , when  $\arg \min(I \cup i) = i$

**Output:**  $\bar{u}(I) : I \in \mathcal{C}_t(K_p)$ , where  $K_p$  is the  $p$ -dimensional complete graph



**Theorem 4.2.** *There exists a  $\rho$ -zCDP algorithm which, with probability at least  $2/3$ , finds a multilinear polynomial  $u$  such that  $\|h - u\|_1 \leq \alpha$ , given  $n$  i.i.d. samples  $Z^1, \dots, Z^n \sim \mathcal{D}$ , where*

$$n = O\left(\frac{(2t)^{O(t)} e^{O(\lambda t)} \cdot p^{4t} \cdot \log(p)}{\alpha^4} + \frac{(2t)^{O(t)} e^{O(\lambda t)} \cdot p^{3t+\frac{1}{2}} \cdot \log^2(p)}{\sqrt{\rho} \alpha^3}\right).$$

*Proof.* Similar to the previous proof, we start by fixing  $i = 1$ . Given a random sample  $Z \sim \mathcal{D}$ , let  $X = \left[\prod_{j \in I} Z^j : I \subset [p] \setminus 1, |I| \leq t-1\right]$  and  $Y = Z_i$ . According to Lemma 4.1, we know that  $\mathbb{E}[Y|X] = \sigma(\langle w^*, X \rangle)$ , where  $w^* = (2 \cdot \overline{\partial_1 h}(I) : I \subset [p] \setminus 1, |I| \leq t-1)$ . Furthermore,  $\|w^*\|_1 \leq 2\lambda$  by the width constraint. Now, given  $n$  i.i.d. samples  $\{z^m\}_{m=1}^n$  drawn from  $\mathcal{D}$ , it is easy to check that for any given  $z^m$ , its corresponding  $(x^m, y^m)$  is one realization of  $(X, Y)$ . Let  $w^{priv}$  be the output of  $\mathcal{A}\left(D, \mathcal{L}, \frac{\rho}{p}, \{w : \|w\|_1 \leq 2\lambda\}\right)$ , where  $D = \{(x^m, y^m)\}_{m=1}^n$  and  $\ell(w; (x, y)) = \log(1 + e^{-y\langle w, x \rangle})$ . By Lemma 3.6,  $\mathbb{E}_{Z \sim \mathcal{D}(A, \theta)}[\ell(w^{priv}; (X, Y))] - \mathbb{E}_{Z \sim \mathcal{D}(A, \theta)}[\ell(w^*; (X, Y))] \leq \gamma$  with probability greater than  $1 - \frac{1}{10p}$ , assuming  $n = \Omega\left(\frac{\sqrt{p}\lambda^2 \log^2(p)}{\sqrt{\rho}\gamma^{\frac{3}{2}}} + \frac{\lambda^2 \log(p)}{\gamma^2}\right)$ .

Now we need the following lemma from [KM17], which is analogous to Lemma 3.7 for the Ising model.

**Lemma 4.3** (Lemma 6.4 of [KM17]). *Let  $P$  be a distribution on  $\{-1, 1\}^{p-1} \times \{-1, 1\}$ . Given multilinear polynomial  $u_1 \in \mathbb{R}^{p-1}$ ,  $\Pr(Y = 1|X = x) = \sigma(u_1(X))$  for  $(X, Y) \sim P$ . Suppose the marginal distribution of  $P$  on  $X$  is  $\delta$ -unbiased, and  $\mathbb{E}_{(X, Y) \sim P}[\log(1 + e^{-Y(u_1(X))})] - \mathbb{E}_{(X, Y) \sim P}[\log(1 + e^{-Y(u_2(X))})] \leq \gamma$  for another multilinear polynomial  $u_2$ , where  $\gamma \leq \delta^t e^{-2\|u_1\|_1 - 6}$ , then  $\|u_1 - u_2\|_1 = O\left((2t)^t e^{\|u_1\|_1} \cdot \sqrt{\gamma/\delta^t} \cdot \binom{p}{t}\right)$ .*

By substituting  $\gamma = e^{-O(\lambda t)} \cdot (2t)^{-O(t)} \cdot p^{-3t} \cdot \alpha^2$ , we have that with probability greater than  $1 - \frac{1}{10p}$ ,  $\sum_{I: \arg \min_{I=1} |\bar{u}(I) - h(I)|} |\bar{u}(I) - h(I)| \leq \frac{\alpha}{p}$ . We note that the coefficients of different monomials are recovered in each iteration. Therefore, by a union bound over  $p$  iterations, we prove the desired result.  $\square$

## 4.2 Parameter Estimation in $\ell_\infty$ Distance

In this section, we introduce a slightly modified version of the algorithm in the last section.

**Algorithm 5:** Private Learning binary  $t$ -wise MRF in  $\ell_\infty$  distance

**Input:**  $n = n_1 + n_2$  i.i.d. samples  $\{z^1, \dots, z^n\}$ , where  $z^m \in \{\pm 1\}^p$  for  $m \in [n]$ ; an upper bound  $\lambda$  on  $\max_{i \in [p]} \|\partial_i h\|_1$ , privacy parameter  $\rho$

- 1 **For** each  $I \subset [p]$  with  $|I| \leq t$
- 2     Let  $Q(I) := \frac{1}{n_2} \sum_{m=n_1}^{n_1+n_2} \prod_{j \in I} z_j^m$
- 3     Compute  $\hat{Q}(I)$ , an estimate of  $Q(I)$  through an  $\rho/2$ -zCDP query release algorithm (PMW [HR10] or sepFEM [VTB<sup>+</sup>19])
- 4 **For**  $i = 1$  to  $p$
- 5      $\forall m \in [n_1]$ ,  $x_m \leftarrow \left[\prod_{j \in I} z_j^m : I \subset [p] \setminus i, |I| \leq t-1\right]$ ,  $y_m \leftarrow z_i^m$
- 6      $w^{priv} \leftarrow \mathcal{A}(D, \mathcal{L}, \rho', \mathcal{C})$  where  $D = \{(x_m, y_m)\}_{m=1}^n$ ,  $\ell(w; d) = \log(1 + e^{-y\langle w, x \rangle})$ ,  $\mathcal{C} = \{\|w\|_1 \leq 2\lambda\}$ , and  $\rho' = \frac{\rho}{2p}$
- 7     Define a polynomial  $v_i : \mathbb{R}^{p-1} \rightarrow \mathbb{R}$  by setting  $\bar{v}_i(I) = \frac{1}{2} w^{priv}(I)$  for all  $I \subset [p] \setminus i$
- 8     **For** each  $I \subset [p] \setminus i$  with  $|I| \leq t-1$
- 9          $\bar{u}(I \cup \{i\}) = \sum_{I' \subset [p]} \bar{\partial}_{I'} v_i(I') \cdot \hat{Q}(I')$ , when  $\arg \min(I \cup i) = i$

**Output:**  $\bar{u}(I) : I \in \mathcal{C}_t(K_p)$ , where  $K_p$  is the  $p$ -dimensional complete graph

We first show that if the estimates  $\hat{Q}$  for the parity queries  $Q$  are sufficiently accurate, Algorithm 5 solves the  $\ell_\infty$  estimation problem, as long as the sample size  $n_1$  is large enough.

**Lemma 4.4.** *Suppose that the estimates  $\hat{Q}$  satisfies  $|\hat{Q}(I) - Q(I)| \leq \alpha/(8\lambda)$  for all  $I \subset [p]$  such that  $|I| \leq t$  and  $n_2 = \Omega(\lambda^2 t \log(p)/\alpha^2)$ . Then with probability at least  $3/4$ , Algorithm 5 outputs a multilinear polynomial  $u$  such that for every maximal monomial  $I$  of  $h$ ,  $|\bar{h}(I) - \bar{u}(I)| \leq \alpha$ , given  $n$  i.i.d. samples  $Z^1, \dots, Z^n \sim \mathcal{D}$ , as long as*

$$n_1 = \Omega\left(\frac{e^{5\lambda t} \cdot \sqrt{p} \log^2(p)}{\sqrt{\rho} \alpha^{\frac{9}{2}}} + \frac{e^{6\lambda t} \cdot \log(p)}{\alpha^6}\right).$$

*Proof.* We will condition on the event that  $\hat{Q}$  is a “good” estimate of  $Q$ :  $|\hat{Q}(I) - Q(I)| \leq \alpha/(8\lambda)$  for all  $I \subset [p]$  such that  $|I| \leq t$ . Let us fix  $i = 1$ . Let  $X = \left[\prod_{j \in I} Z^j : I \subset [p] \setminus \{1\}, |I| \leq t - 1\right]$ ,  $Y = Z_i$ , and we know that  $\mathbb{E}[Y|X] = \sigma(\langle w^*, X \rangle)$ , where  $w^* = (2 \cdot \partial_1 \bar{h}(I) : I \subset [p] \setminus \{1\}, |I| \leq t - 1)$ . Now given  $n_1$  i.i.d. samples  $\{z^m\}_{m=1}^{n_1}$  drawn from  $\mathcal{D}$ , let  $w^{priv}$  be the output of  $\mathcal{A}\left(D, \mathcal{L}, \frac{\ell}{p}, \{w : \|w\|_1 \leq 2\lambda\}\right)$ , where  $D = \{(x^m, y^m)\}_{m=1}^{n_1}$  and  $\ell(w; (x, y)) = \log(1 + e^{-y\langle w, x \rangle})$ . Similarly, with probability at least  $1 - \frac{1}{10p}$ ,

$$\mathbb{E}_{Z \sim \mathcal{D}(A, \theta)} [\ell(w^{priv}; (X, Y))] - \mathbb{E}_{Z \sim \mathcal{D}(A, \theta)} [\ell(w^*; (X, Y))] \leq \gamma$$

as long as  $n_1 = \Omega\left(\frac{\sqrt{p} \lambda^2 \log^2(p)}{\sqrt{\rho} \gamma^{\frac{3}{2}}} + \frac{\lambda^2 \log(p)}{\gamma^2}\right)$ .

Now we utilize Lemma 6.4 from [KM17], which states that if  $\mathbb{E}_{Z \sim \mathcal{D}(A, \theta)} [\ell(w^{priv}; (X, Y))] - \mathbb{E}_{Z \sim \mathcal{D}(A, \theta)} [\ell(w^*; (X, Y))] \leq \gamma$ , given a random sample  $X$ , for any maximal monomial  $I \subset [p] \setminus \{1\}$  of  $\partial_1 h$ ,

$$\Pr\left(|\bar{\partial}_1 \bar{h}(I) - \partial_I v_1(X)| \geq \frac{\alpha}{4}\right) < O\left(\frac{\gamma \cdot e^{3\lambda t}}{\alpha^2}\right).$$

By replacing  $\gamma = \frac{e^{-3\lambda t} \cdot \alpha^3}{8\lambda}$ , we have  $\Pr\left(|\bar{\partial}_1 \bar{h}(I) - \partial_I v_1(X)| \geq \frac{\alpha}{4}\right) < \frac{\alpha}{8\lambda}$ , as long as  $n_1 = \Omega\left(\frac{\sqrt{p} e^{5\lambda t} \log^2(p)}{\sqrt{\rho} \alpha^{\frac{9}{2}}} + \frac{e^{6\lambda t} \log(p)}{\alpha^6}\right)$ . Accordingly, for any maximal monomial  $I$ ,  $|\mathbb{E}[\partial_I v_1(X)] - \bar{\partial}_1 \bar{h}(I)| \leq \mathbb{E}[|\partial_I v_1(X) - \bar{\partial}_1 \bar{h}(I)|] \leq \frac{\alpha}{4} + 2\lambda \cdot \frac{\alpha}{8\lambda} = \frac{\alpha}{2}$ . By Hoeffding inequality, given  $n_2 = \Omega\left(\frac{\lambda^2 t \log p}{\alpha^2}\right)$ , for each maximal monomial  $I$ , with probability greater than  $1 - \frac{1}{p^t}$ ,  $\left|\frac{1}{n_2} \sum_{m=1}^{n_2} \partial_I v_1(X_m) - \mathbb{E}[\partial_I v_1(X)]\right| \leq \frac{\alpha}{4}$ . Note that  $|Q(I) - \hat{Q}(I)| \leq \frac{\alpha}{8\lambda}$ , then  $\left|\frac{1}{n_2} \sum_{m=1}^{n_2} \partial_I v_1(X_m) - \sum_{I' \subset [p]} \bar{\partial}_I \bar{v}_1(I') \cdot \hat{Q}(I')\right| \leq \frac{\alpha}{8}$ . Therefore,

$$\begin{aligned} & \left| \sum_{I' \subset [p]} \bar{\partial}_I \bar{v}_1(I') \cdot \hat{Q}(I') - \bar{\partial}_1 \bar{h}(I) \right| \\ & \leq \left| \sum_{I' \subset [p]} \bar{\partial}_I \bar{v}_1(I') \cdot \hat{Q}(I') - \frac{1}{n_2} \sum_{m=1}^{n_2} \partial_I v_1(X_m) \right| + \left| \frac{1}{n_2} \sum_{m=1}^{n_2} \partial_I v_1(X_m) - \mathbb{E}[\partial_I v_1(X)] \right| + |\mathbb{E}[\partial_I v_1(X)] - \bar{\partial}_1 \bar{h}(I)| \\ & \leq \frac{\alpha}{8} + \frac{\alpha}{4} + \frac{\alpha}{2} = \frac{7\alpha}{8}. \end{aligned}$$

Finally, by a union bound over  $p$  iterations and all the maximal monomials, we prove the desired results.  $\square$

We now consider two private algorithms for releasing the parity queries. The first algorithm is called Private Multiplicative Weights (PMW) [HR10], which provides a better accuracy guarantee but runs in time exponential in the dimension  $p$ . The following theorem can be viewed as a zCDP version of Theorem 4.3 in [Vad17], by noting that during the analysis, every iteration satisfies  $\varepsilon_0$ -DP, which naturally satisfies  $\varepsilon_0^2$ -zCDP, and by replacing the strong composition theorem of  $(\varepsilon, \delta)$ -DP by the composition theorem of zCDP (Lemma 2.15).

**Lemma 4.5** (Sample complexity of PMW, modification of Theorem 4.3 of [Vad17]). *The PMW algorithm satisfies  $\rho$ -zCDP and releases  $\hat{Q}$  such that with probability greater than  $\frac{19}{20}$ , for all  $I \subset [p]$  with  $|I| \leq t$ ,  $|\hat{Q}(I) - Q(I)| \leq \frac{\alpha}{8\lambda}$  as long as the size of the data set*

$$n_2 = \Omega\left(\frac{t\lambda^2 \cdot \sqrt{p} \log p}{\sqrt{\rho}\alpha^2}\right).$$

The second algorithm sepFEM (Separator-Follow-the-perturbed-leader with exponential mechanism) has slightly worse sample complexity, but runs in polynomial time when it has access to an optimization oracle  $\mathcal{O}$  that does the following: given as input a weighted dataset  $(I_1, w_1), \dots, (I_m, w_m) \in 2^{[p]} \times \mathbb{R}$ , find  $x \in \{\pm 1\}^p$ ,

$$\max_{x \in \{\pm 1\}^p} \sum_{i=1}^m w_i \prod_{j \in I_i} x_j.$$

The oracle  $\mathcal{O}$  essentially solves cost-sensitive classification problems over the set of parity functions [ZLA03], and it can be implemented with an integer program solver [VTB+19, GAH+14].

**Lemma 4.6** (Sample complexity of sepFEM, [VTB+19]). *The sepFEM algorithm satisfies  $\rho$ -zCDP and releases  $\hat{Q}$  such that with probability greater than  $\frac{19}{20}$ , for all  $I \subset [p]$  with  $|I| \leq t$ ,  $|\hat{Q}(I) - Q(I)| \leq \frac{\alpha}{8\lambda}$  as long as the size of the data set*

$$n_2 = \Omega\left(\frac{t\lambda^2 \cdot p^{5/4} \log p}{\sqrt{\rho}\alpha^2}\right)$$

*The algorithm runs in polynomial time given access to the optimization oracle  $\mathcal{O}$  defined above.*

Now we can combine Lemmas 4.4, 4.5, and 4.6 to state the formal guarantee of Algorithm 5.

**Theorem 4.7.** *Algorithm 5 is a  $\rho$ -zCDP algorithm which, with probability at least  $2/3$ , finds a multilinear polynomial  $u$  such that for every maximal monomial  $I$  of  $h$ ,  $|\bar{h}(I) - \bar{u}(I)| \leq \alpha$ , given  $n$  i.i.d. samples  $Z^1, \dots, Z^n \sim \mathcal{D}$ , and*

1. *if it uses PMW for releasing  $\hat{Q}$ ; it has a sample complexity of*

$$n = O\left(\frac{e^{5\lambda t} \cdot \sqrt{p} \log^2(p)}{\sqrt{\rho}\alpha^{\frac{9}{2}}} + \frac{t\lambda^2 \cdot \sqrt{p} \log p}{\sqrt{\rho}\alpha^2} + \frac{e^{6\lambda t} \cdot \log(p)}{\alpha^6}\right)$$

*and a runtime complexity that is exponential in  $p$ ;*

2. *if it uses sepFEM for releasing  $\hat{Q}$ , it has a sample complexity of*

$$n = \tilde{O}\left(\frac{e^{5\lambda t} \cdot \sqrt{p} \log^2(p)}{\sqrt{\rho}\alpha^{\frac{9}{2}}} + \frac{t\lambda^2 \cdot p^{5/4} \log p}{\sqrt{\rho}\alpha^2} + \frac{e^{6\lambda t} \cdot \log(p)}{\alpha^6}\right)$$

*and runs in polynomial time whenever  $t = O(1)$ .*

## 5 Lower Bounds for Parameter Learning

The lower bound for parameter estimation is based on mean estimation in  $\ell_\infty$  distance.

**Theorem 5.1.** *Suppose  $\mathcal{A}$  is an  $(\varepsilon, \delta)$ -differentially private algorithm that takes  $n$  i.i.d. samples  $Z^1, \dots, Z^n$  drawn from any unknown  $p$ -variable Ising model  $\mathcal{D}(A, \theta)$  and outputs  $\hat{A}$  such that  $\mathbb{E} \left[ \max_{i,j \in [p]} |A_{i,j} - \hat{A}_{i,j}| \right] \leq \alpha \leq 1/50$ . Then  $n = \Omega\left(\frac{\sqrt{p}}{\alpha\varepsilon}\right)$ .*

*Proof.* Consider a Ising model  $\mathcal{D}(A, 0)$  with  $A \in \mathbb{R}^{p \times p}$  defined as follows: for  $i \in [\frac{p}{2}]$ ,  $A_{2i-1, 2i} = A_{2i, 2i-1} = \eta_i \in [-\ln(2), \ln(2)]$ , and  $A_{ll'} = 0$  for all other pairs of  $(l, l')$ . This construction divides the  $p$  nodes into  $\frac{p}{2}$  pairs, where there is no correlation between nodes belonging to different pairs. It follows that

$$\begin{aligned} \Pr(Z_{2i-1} = 1, Z_{2i} = 1) &= \Pr(Z_{2i-1} = -1, Z_{2i} = -1) = \frac{1}{2} \frac{e^{\eta_i}}{e^{\eta_i} + 1}, \\ \Pr(Z_{2i-1} = 1, Z_{2i} = -1) &= \Pr(Z_{2i-1} = -1, Z_{2i} = 1) = \frac{1}{2} \frac{1}{e^{\eta_i} + 1}. \end{aligned}$$

For each observation  $Z$ , we obtain an observation  $X \in \{\pm 1\}^{p/2}$  such that  $X_i = Z_{2i-1}Z_{2i}$ . Then each observation  $X$  is distributed according to a product distribution in  $\{\pm 1\}^{(p/2)}$  such that the mean of each coordinate  $j$  is  $(e^{\eta_i} - 1)/(e^{\eta_i} + 1) \in [-1/3, 1/3]$ .

Suppose that an  $(\varepsilon, \delta)$ -differentially private algorithm takes  $n$  observations drawn from any such Ising model distribution and output a matrix  $\hat{A}$  such that  $\mathbf{E} \left[ \max_{i,j \in [p]} |A_{i,j} - \hat{A}_{i,j}| \right] \leq \alpha$ . Let  $\hat{\eta}_i = \min\{\max\{\hat{A}_{2i-1, 2i}, -\ln(2)\}, \ln(2)\}$  be the value of  $A_{2i-1, 2i}$  rounded into the range of  $[-\ln(2), \ln(2)]$ , and so  $|\eta_i - \hat{\eta}_i| \leq \alpha$ . It follows that

$$\begin{aligned} \left| \frac{e^{\eta_i} - 1}{e^{\eta_i} + 1} - \frac{e^{\hat{\eta}_i} - 1}{e^{\hat{\eta}_i} + 1} \right| &= 2 \left| \frac{e^{\eta_i} - e^{\hat{\eta}_i}}{(e^{\eta_i} + 1)(e^{\hat{\eta}_i} + 1)} \right| \\ &< 2 \left| e^{\eta_i} - e^{\hat{\eta}_i} \right| \leq 4 \left( e^{|\eta_i - \hat{\eta}_i|} - 1 \right) \leq 8|\eta_i - \hat{\eta}_i| \end{aligned}$$

where the last step follows from the fact that  $e^a \leq 1 + 2a$  for any  $a \in [0, 1]$ . Thus, such private algorithm also can estimate the mean of the product distribution accurately:

$$\mathbf{E} \left[ \sum_{i=1}^{p/2} \left| \frac{e^{\eta_i} - 1}{e^{\eta_i} + 1} - \frac{e^{\hat{\eta}_i} - 1}{e^{\hat{\eta}_i} + 1} \right|^2 \right] \leq 32p\alpha^2$$

Now we will use the following sample complexity lower bound on private mean estimation on product distributions.

**Lemma 5.2** (Lemma 6.2 of [KLSU19]). *If  $M: \{\pm 1\}^{n \times d} \rightarrow [-1/3, 1/3]^d$  is  $(\varepsilon, 3/(64n))$ -differentially private, and for every product distribution  $P$  over  $\{\pm 1\}^d$  such that the mean of each coordinate  $\mu_j$  satisfies  $-1/3 \leq \mu_j \leq 1/3$ ,*

$$\mathbf{E}_{X \sim P^n} [\|M(X) - \mu\|_2^2] \leq \gamma^2 \leq \frac{d}{54},$$

then  $n \geq d/(72\gamma\varepsilon)$ .

Then our stated bound follows by instantiating  $\gamma^2 = 32p\alpha^2$  and  $d = p/2$  in Lemma 5.2.  $\square$

## 6 Structure Learning of Graphical Models

In this section, we will give an  $(\varepsilon, \delta)$ -differentially private algorithm for learning the *structure* of a Markov Random Field. The dependence on the dimension  $d$  will be only *logarithmic*, in comparison to the complexity of privately learning the parameters. As we have shown in Section 5, this dependence is necessarily polynomial in  $p$ , even under approximate differential privacy. Furthermore, as we will show in Section 7, if we wish to learn the structure of an MRF under more restrictive notions of privacy (such as pure or concentrated), the complexity also becomes polynomial in  $p$ . Thus, in very high-dimensional settings, learning the structure of the MRF under approximate differential privacy is essentially the only notion of private learnability which is tractable.

The following lemma is immediate from stability-based mode arguments (see, e.g., Proposition 3.4 of [Vad17]).

**Lemma 6.1.** *Suppose there exists a (non-private) algorithm which takes  $X = (X^1, \dots, X^n)$  sampled i.i.d. from some distribution  $\mathcal{D}$ , and outputs some fixed value  $Y$  (which may depend on  $\mathcal{D}$ ) with probability at least  $2/3$ . Then there exists an  $(\varepsilon, \delta)$ -differentially private algorithm which takes  $O\left(\frac{n \log(1/\delta)}{\varepsilon}\right)$  samples and outputs  $Y$  with probability at least  $1 - \delta$ .*

We can now directly import the following theorem from [WSD19].

**Theorem 6.2** ([WSD19]). *There exists an algorithm which, with probability at least  $2/3$ , learns the structure of a pairwise graphical model. It requires  $n = O\left(\frac{\lambda^2 k^4 e^{14\lambda} \log(pk)}{\eta^4}\right)$  samples.*

This gives us the following private learning result as a corollary.

**Corollary 6.3.** *There exists an  $(\varepsilon, \delta)$ -differentially private algorithm which, with probability at least  $2/3$ , learns the structure of a pairwise graphical model. It requires  $n = O\left(\frac{\lambda^2 k^4 e^{14\lambda} \log(pk) \log(1/\delta)}{\varepsilon \eta^4}\right)$  samples.*

For binary MRFs of higher-order, we instead import the following theorem from [KM17]:

**Theorem 6.4** ([KM17]). *There exists an algorithm which, with probability at least  $2/3$ , learns the structure of a binary  $t$ -wise MRF. It requires  $n = O\left(\frac{e^{O(\lambda t)} \log(\frac{p}{\eta})}{\eta^4}\right)$  samples.*

This gives us the following private learning result as a corollary.

**Corollary 6.5.** *There exists an  $(\varepsilon, \delta)$ -differentially private algorithm which, with probability at least  $2/3$ , learns the structure of a binary  $t$ -wise MRF. It requires*

$$n = O\left(\frac{e^{O(\lambda t)} \log(\frac{p}{\eta}) \log(1/\delta)}{\varepsilon \eta^4}\right)$$

*samples.*

## 7 Lower Bounds for Structure Learning of Graphical Models

In this section, we will prove structure learning lower bounds under pure DP or zero-concentrated DP. The graphical models we consider are the Ising models and pairwise graphical model. However, we note that all the lower bounds for the Ising model also hold for binary  $t$ -wise MRFs, since the

Ising model is a special case of binary  $t$ -wise MRFs corresponding to  $t = 2$ . We will show that under  $(\varepsilon, 0)$ -DP or  $\rho$ -zCDP, a polynomial dependence on the dimension is unavoidable in the sample complexity.

In Section 7.1, we assume that our samples are generated from an Ising model. In Section 7.2, we extend our lower bounds to pairwise graphical models.

## 7.1 Lower Bounds for Structure Learning of Ising Models

**Theorem 7.1.** *Any  $(\varepsilon, 0)$ -DP algorithm which learns the structure of an Ising model with minimum edge weight  $\eta$  with probability at least  $2/3$  requires  $n = \Omega\left(\frac{\sqrt{p}}{\eta\varepsilon} + \frac{p}{\varepsilon}\right)$  samples. Furthermore, at least  $n = \Omega\left(\sqrt{\frac{p}{\rho}}\right)$  samples are required for the same task under  $\rho$ -zCDP.*

*Proof.* Our lower bound argument is in two steps. The first step is to construct a set of distributions, consisting of  $2^{\frac{p}{2}}$  different Ising models such that any feasible structure learning algorithm should output different answers for different distributions. In the second step, we utilize the probabilistic packing argument for DP [ASZ20b], or the packing argument for zCDP [BS16] to get the desired lower bound.

To start, we would like to use the following binary code to construct the distribution set. Let  $\mathcal{C} = \{0, 1\}^{\frac{p}{2}}$ , given  $c \in \mathcal{C}$ , we construct the corresponding distribution  $\mathcal{D}(A^c, 0)$  with  $A^c \in \mathbb{R}^{p \times p}$  defined as follows: for  $i \in [\frac{p}{2}]$ ,  $A_{2i-1, 2i}^c = A_{2i, 2i-1}^c = \eta \cdot c[i]$ , and 0 elsewhere. By construction, we divide the  $p$  nodes into  $\frac{p}{2}$  different pairs, where there is no correlation between nodes belonging to different pairs. Furthermore, for pair  $i$ , if  $c[i] = 0$ , which means the value of node  $2i - 1$  is independent of node  $2i$ , it is not hard to show

$$\begin{aligned} \Pr(Z_{2i-1} = 1, Z_{2i} = 1) &= \Pr(Z_{2i-1} = -1, Z_{2i} = -1) = \frac{1}{4}, \\ \Pr(Z_{2i-1} = 1, Z_{2i} = -1) &= \Pr(Z_{2i-1} = -1, Z_{2i} = 1) = \frac{1}{4}. \end{aligned}$$

On the other hand, if  $c[i] = 1$ ,

$$\begin{aligned} \Pr(Z_{2i-1} = 1, Z_{2i} = 1) &= \Pr(Z_{2i-1} = -1, Z_{2i} = -1) = \frac{1}{2} \cdot \frac{e^\eta}{e^\eta + 1}, \\ \Pr(Z_{2i-1} = 1, Z_{2i} = -1) &= \Pr(Z_{2i-1} = -1, Z_{2i} = 1) = \frac{1}{2} \cdot \frac{1}{e^\eta + 1}. \end{aligned}$$

The Chi-squared distance between these two distributions is

$$8 \left[ \left( \frac{1}{2} \cdot \frac{e^\eta}{e^\eta + 1} - \frac{1}{4} \right)^2 + \left( \frac{1}{2} \cdot \frac{1}{e^\eta + 1} - \frac{1}{4} \right)^2 \right] = \left( 1 - \frac{2}{e^\eta + 1} \right)^2 \leq 4\eta^2.$$

Now we want to upper bound the total variation distance between  $\mathcal{D}(A^{c_1}, 0)$  and  $\mathcal{D}(A^{c_2}, 0)$  for any  $c_1 \neq c_2 \in \mathcal{C}$ . Let  $P_i$  and  $Q_i$  denote the joint distribution of node  $2i - 1$  and node  $2i$  corresponding to  $\mathcal{D}(A^{c_1}, 0)$  and  $\mathcal{D}(A^{c_2}, 0)$ . We have that

$$d_{TV}(\mathcal{D}(A^{c_1}, 0), \mathcal{D}(A^{c_2}, 0)) \leq \sqrt{2d_{KL}(\mathcal{D}(A^{c_1}, 0), \mathcal{D}(A^{c_2}, 0))} = \sqrt{2 \sum_{i=1}^{\frac{p}{2}} d_{KL}(P_i, Q_i)} \leq \min(2\eta\sqrt{p}, 1),$$

where the first inequality is by Pinsker's inequality, and the last inequality comes from the fact that the KL divergence is always upper bounded by the Chi-squared distance.

In order to attain pure DP lower bounds, we utilize the probabilistic version of the packing argument in [ASZ20a], as stated below.

**Lemma 7.2.** *Let  $\mathcal{V} = \{P_1, P_2, \dots, P_M\}$  be a set of  $M$  distributions over  $\mathcal{X}^n$ . Suppose that for any pair of distributions  $P_i$  and  $P_j$ , there exists a coupling between  $P_i$  and  $P_j$ , such that  $\mathbb{E}[d_{\text{ham}}(X^n, Y^n)] \leq D$ , where  $X^n \sim P_i$  and  $Y^n \sim P_j$ . Let  $\{S_i\}_{i \in [M]}$  be a collection of disjoint subsets of  $\mathcal{S}$ . If there exists an  $\varepsilon$ -DP algorithm  $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{S}$  such that for every  $i \in [M]$ , given  $Z^n \sim P_i$ ,  $\Pr(\mathcal{A}(Z^n) \in S_i) \geq \frac{9}{10}$ , then*

$$\varepsilon = \Omega\left(\frac{\log M}{D}\right).$$

For any  $c_1, c_2 \in C$ , we have  $d_{TV}(\mathcal{D}(A^{c_1}, 0), \mathcal{D}(A^{c_2}, 0)) \leq \min(2\eta\sqrt{p}, 1)$ . By the property of maximal coupling [dH12], there must exist some coupling between  $\mathcal{D}^n(A^{c_1}, 0)$  and  $\mathcal{D}^n(A^{c_2}, 0)$  with expected Hamming distance smaller than  $\min(2n\eta\sqrt{p}, n)$ . Therefore, we have  $\varepsilon = \Omega\left(\frac{\log|C|}{\min(n\eta\sqrt{p}, n)}\right)$ , and accordingly,  $n = \Omega\left(\frac{\sqrt{p}}{\eta\varepsilon} + \frac{p}{\varepsilon}\right)$ .

Now we move to zCDP lower bounds. We utilize a different version of the packing argument [BS16], which works under zCDP.

**Lemma 7.3.** *Let  $\mathcal{V} = \{P_1, P_2, \dots, P_M\}$  be a set of  $M$  distributions over  $\mathcal{X}^n$ . Let  $\{S_i\}_{i \in [M]}$  be a collection of disjoint subsets of  $\mathcal{S}$ . If there exists an  $\rho$ -zCDP algorithm  $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{S}$  such that for every  $i \in [M]$ , given  $Z^n \sim P_i$ ,  $\Pr(\mathcal{A}(Z^n) \in S_i) \geq \frac{9}{10}$ , then*

$$\rho = \Omega\left(\frac{\log M}{n^2}\right).$$

By Lemma 7.3, we derive  $\rho = \Omega\left(\frac{p}{n^2}\right)$  and  $n = \Omega\left(\sqrt{\frac{p}{\rho}}\right)$  accordingly.  $\square$

## 7.2 Lower Bounds for Structure Learning of Pairwise Graphical Models

Similar techniques can be used to derive lower bounds for pairwise graphical models.

**Theorem 7.4.** *Any  $(\varepsilon, 0)$ -DP algorithm which learns the structure of the  $p$ -variable pairwise graphical models with minimum edge weight  $\eta$  with probability at least  $2/3$  requires  $n = \Omega\left(\frac{\sqrt{p}}{\eta\varepsilon} + \frac{k^2 p}{\varepsilon}\right)$  samples. Furthermore, at least  $n = \Omega\left(\sqrt{\frac{k^2 p}{\rho}}\right)$  samples are required for the same task under  $\rho$ -zCDP.*

*Proof.* Similar to before, we start with constructing a distribution set consisting of  $2^{O(kp)}$  different pairwise graphical models such that any accurate structure learning algorithm must output different answers for different distributions.

Let  $\mathcal{C}$  be the real symmetric matrix with each value constrained to either 0 or  $\eta$ , i.e.,  $\mathcal{C} = \{W \in \{0, \eta\}^{k \times k} : W = W^T\}$ . Without loss of generality, we assume  $p$  is even. Given  $c = [c_1, c_2, \dots, c_p]$ , where  $c_1, c_2, \dots, c_p \in C$ , we construct the corresponding distribution  $\mathcal{D}(\mathcal{W}^c, 0)$  with  $\mathcal{W}^c$  defined as follows: for  $l \in [\frac{p}{2}]$ ,  $W_{2l-1, 2l}^c = c_l$ , and for other pairs  $(i, j)$ ,  $W_{i,j}^c = 0$ . Similarly, by this construction we divide  $p$  nodes into  $\frac{p}{2}$  different pairs, and there is no correlation between nodes belonging to different pairs.

We first prove lower bounds under  $(\varepsilon, 0)$ -DP. By Lemma 7.2,  $\varepsilon = \Omega\left(\frac{\log|C|}{n}\right)$ , since for any two  $n$ -sample distributions, the expected coupling distance can be always upper bounded by  $n$ . We also note that  $|C| = \left(2^{\frac{k(k+1)}{2}}\right)^p$ . Therefore, we have  $n = \Omega\left(\frac{k^2 p}{\varepsilon}\right)$ . At the same time,  $n = \Omega\left(\frac{\sqrt{p}}{\eta\varepsilon}\right)$  is another lower bound, inherited from the easier task of learning Ising models.

With respect to zCDP, we utilize Lemma 7.3 and obtain  $\rho = \Omega\left(\frac{k^2 p}{n^2}\right)$ . Therefore, we have  $n = \Omega\left(\sqrt{\frac{k^2 p}{\rho}}\right)$ .  $\square$

## Acknowledgments

The authors would like to thank Kunal Talwar for suggesting the study of this problem, and Adam Klivans, Frederic Koehler, Ankur Moitra, and Shanshan Wu for helpful and inspiring conversations. GK would like to thank Chengdu Style Restaurant (古月飘香) in Berkeley for inspiration in the conception of this project.

## References

- [AKN06] Pieter Abbeel, Daphne Koller, and Andrew Y. Ng. Learning factor graphs in polynomial time and sample complexity. *Journal of Machine Learning Research*, 7(Aug):1743–1788, 2006.
- [AKSZ18] Jayadev Acharya, Gautam Kamath, Ziteng Sun, and Huanyu Zhang. Inspectre: Privately estimating the unseen. In *Proceedings of the 35th International Conference on Machine Learning*, ICML ’18, pages 30–39. JMLR, Inc., 2018.
- [ASZ20a] Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Personal communication, 2020.
- [ASZ20b] Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Differentially private assouad, fano, and le cam. *arXiv preprint arXiv:2004.06830*, 2020.
- [BBC<sup>+</sup>19] Ivona Bezakova, Antonio Blanca, Zongchen Chen, Daniel Štefankovič, and Eric Vigoda. Lower bounds for testing graphical models: Colorings and antiferromagnetic Ising models. In *Proceedings of the 32nd Annual Conference on Learning Theory*, COLT ’19, pages 283–298, 2019.
- [BBKN14] Amos Beimel, Hai Brenner, Shiva Prasad Kasiviswanathan, and Kobbi Nissim. Bounds on the sample complexity for private learning and private data release. *Machine Learning*, 94(3):401–437, 2014.
- [BDMN05] Avrim Blum, Cynthia Dwork, Frank McSherry, and Kobbi Nissim. Practical privacy: The SuLQ framework. In *Proceedings of the 24th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, PODS ’05, pages 128–138, New York, NY, USA, 2005. ACM.
- [BGS14] Guy Bresler, David Gamarnik, and Devavrat Shah. Structure learning of antiferromagnetic Ising models. In *Advances in Neural Information Processing Systems 27*, NIPS ’14, pages 2852–2860. Curran Associates, Inc., 2014.
- [Bha19] Bhaswar B. Bhattacharya. A general asymptotic framework for distribution-free graph-based two-sample tests. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 81(3):575–602, 2019.
- [BKSW19] Mark Bun, Gautam Kamath, Thomas Steinke, and Zhiwei Steven Wu. Private hypothesis selection. In *Advances in Neural Information Processing Systems 32*, NeurIPS ’19. Curran Associates, Inc., 2019.



- [BM16] Bhaswar B. Bhattacharya and Sumit Mukherjee. Inference in Ising models. *Bernoulli*, 2016.
- [BMS<sup>+</sup>17] Garrett Bernstein, Ryan McKenna, Tao Sun, Daniel Sheldon, Michael Hay, and Gerome Miklau. Differentially private learning of undirected graphical models using collective graphical models. In *Proceedings of the 34th International Conference on Machine Learning, ICML '17*, pages 478–487. JMLR, Inc., 2017.
- [BNSV15] Mark Bun, Kobbi Nissim, Uri Stemmer, and Salil Vadhan. Differentially private release and learning of threshold functions. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science, FOCS '15*, pages 634–649, Washington, DC, USA, 2015. IEEE Computer Society.
- [Bre15] Guy Bresler. Efficiently learning Ising models on arbitrary graphs. In *Proceedings of the 47th Annual ACM Symposium on the Theory of Computing, STOC '15*, pages 771–782, New York, NY, USA, 2015. ACM.
- [BS16] Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Proceedings of the 14th Conference on Theory of Cryptography, TCC '16-B*, pages 635–658, Berlin, Heidelberg, 2016. Springer.
- [BUV14] Mark Bun, Jonathan Ullman, and Salil Vadhan. Fingerprinting codes and the price of approximate differential privacy. In *Proceedings of the 46th Annual ACM Symposium on the Theory of Computing, STOC '14*, pages 1–10, New York, NY, USA, 2014. ACM.
- [Cha05] Sourav Chatterjee. *Concentration Inequalities with Exchangeable Pairs*. PhD thesis, Stanford University, June 2005.
- [CL68] C.K. Chow and C.N. Liu. Approximating discrete probability distributions with dependence trees. *IEEE Transactions on Information Theory*, 14(3):462–467, 1968.
- [Cla10] Kenneth L Clarkson. Coresets, sparse greedy approximation, and the frank-wolfe algorithm. *ACM Transactions on Algorithms (TALG)*, 6(4):63, 2010.
- [CRJ19] Amrita Roy Chowdhury, Theodoros Rekatsinas, and Somesh Jha. Data-dependent differentially private parameter learning for directed graphical models. *arXiv preprint arXiv:1905.12813*, 2019.
- [CT06] Imre Csiszár and Zsolt Talata. Consistent estimation of the basic neighborhood of Markov random fields. *The Annals of Statistics*, 34(1):123–145, 2006.
- [CWZ19] T. Tony Cai, Yichen Wang, and Linjun Zhang. The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy. *arXiv preprint arXiv:1902.04495*, 2019.
- [DDK17] Constantinos Daskalakis, Nishanth Dikkala, and Gautam Kamath. Concentration of multilinear functions of the Ising model with applications to network data. In *Advances in Neural Information Processing Systems 30, NIPS '17*. Curran Associates, Inc., 2017.
- [DDK18] Constantinos Daskalakis, Nishanth Dikkala, and Gautam Kamath. Testing Ising models. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '18*, pages 1989–2007, Philadelphia, PA, USA, 2018. SIAM.

- [DDK19] Constantinos Daskalakis, Nishanth Dikkala, and Gautam Kamath. Testing Ising models. *IEEE Transactions on Information Theory*, 65(11):6829–6852, 2019.
- [dH12] Frank den Hollander. Probability theory: The coupling method. *Lecture notes available online (<http://websites.math.leidenuniv.nl/probability/lecturenotes/CouplingLectures.pdf>)*, 2012.
- [DHS15] Ilias Diakonikolas, Moritz Hardt, and Ludwig Schmidt. Differentially private learning of structured discrete distributions. In *Advances in Neural Information Processing Systems 28*, NIPS ’15, pages 2566–2574. Curran Associates, Inc., 2015.
- [Dif17] Differential Privacy Team, Apple. Learning with privacy at scale. <https://machinelearning.apple.com/docs/learning-with-privacy-at-scale/appledifferentialsystem.pdf>, December 2017.
- [DKY17] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. In *Advances in Neural Information Processing Systems 30*, NIPS ’17, pages 3571–3580. Curran Associates, Inc., 2017.
- [DL09] Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In *Proceedings of the 41st Annual ACM Symposium on the Theory of Computing*, STOC ’09, pages 371–380, New York, NY, USA, 2009. ACM.
- [DLS<sup>+</sup>17] Aref N. Dajani, Amy D. Lauger, Phyllis E. Singer, Daniel Kifer, Jerome P. Reiter, Ashwin Machanavajjhala, Simson L. Garfinkel, Scot A. Dahl, Matthew Graham, Vishesh Karwa, Hang Kim, Philip Lelerc, Ian M. Schmutte, William N. Sexton, Lars Vilhuber, and John M. Abowd. The modernization of statistical disclosure limitation at the U.S. census bureau, 2017. Presented at the September 2017 meeting of the Census Scientific Advisory Committee.
- [DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Conference on Theory of Cryptography*, TCC ’06, pages 265–284, Berlin, Heidelberg, 2006. Springer.
- [DMR11] Constantinos Daskalakis, Elchanan Mossel, and Sébastien Roch. Evolutionary trees and the Ising model on the Bethe lattice: A proof of Steel’s conjecture. *Probability Theory and Related Fields*, 149(1):149–189, 2011.
- [DMR18] Luc Devroye, Abbas Mehrabian, and Tommy Reddad. The minimax learning rate of normal and Ising undirected graphical models. *arXiv preprint arXiv:1806.06887*, 2018.
- [DR16] Cynthia Dwork and Guy N. Rothblum. Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*, 2016.
- [DSS<sup>+</sup>15] Cynthia Dwork, Adam Smith, Thomas Steinke, Jonathan Ullman, and Salil Vadhan. Robust traceability from trace amounts. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science*, FOCS ’15, pages 650–669, Washington, DC, USA, 2015. IEEE Computer Society.
- [Ell93] Glenn Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071, 1993.

- [EPK14] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. RAPPOR: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM Conference on Computer and Communications Security, CCS '14*, pages 1054–1067, New York, NY, USA, 2014. ACM.
- [Fel04] Joseph Felsenstein. *Inferring Phylogenies*. Sinauer Associates Sunderland, 2004.
- [FLNP00] Nir Friedman, Michal Linial, Iftach Nachman, and Dana Pe’er. Using Bayesian networks to analyze expression data. *Journal of Computational Biology*, 7(3-4):601–620, 2000.
- [GAH<sup>+</sup>14] Marco Gaboardi, Emilio Jesús Gallego Arias, Justin Hsu, Aaron Roth, and Zhiwei Steven Wu. Dual query: Practical private query release for high dimensional data. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, pages 1170–1178, 2014.
- [GG86] Stuart Geman and Christine Graffigne. Markov random field image models and their applications to computer vision. In *Proceedings of the International Congress of Mathematicians*, pages 1496–1517. American Mathematical Society, 1986.
- [GLP18] Reza Gheissari, Eyal Lubetzky, and Yuval Peres. Concentration inequalities for polynomials of contracting Ising models. *Electronic Communications in Probability*, 23(76):1–12, 2018.
- [HKM17] Linus Hamilton, Frederic Koehler, and Ankur Moitra. Information theoretic properties of Markov random fields, and their algorithmic applications. In *Advances in Neural Information Processing Systems 30, NIPS '17*. Curran Associates, Inc., 2017.
- [HR10] Moritz Hardt and Guy N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science, FOCS '10*, pages 61–70, Washington, DC, USA, 2010. IEEE Computer Society.
- [HT10] Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In *Proceedings of the 42nd Annual ACM Symposium on the Theory of Computing, STOC '10*, pages 705–714, New York, NY, USA, 2010. ACM.
- [Isi25] Ernst Ising. Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik A Hadrons and Nuclei*, 31(1):253–258, 1925.
- [Jag13] Martin Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *ICML (1)*, pages 427–435, 2013.
- [JJR11] Ali Jalali, Christopher C. Johnson, and Pradeep K. Ravikumar. On learning discrete graphical models using greedy methods. In *Advances in Neural Information Processing Systems 24, NIPS '11*, pages 1935–1943. Curran Associates, Inc., 2011.
- [JRVS11] Ali Jalali, Pradeep K. Ravikumar, Vishvas Vasuki, and Sujay Sanghavi. On learning discrete graphical models using group-sparse regularization. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics, AISTATS '11*, pages 378–387. JMLR, Inc., 2011.

- [KKMN09] Aleksandra Korolova, Krishnaram Kenthapadi, Nina Mishra, and Alexandros Ntoulas. Releasing search queries and clicks privately. In *Proceedings of the 18th International World Wide Web Conference, WWW '09*, pages 171–180, New York, NY, USA, 2009. ACM.
- [KLSU19] Gautam Kamath, Jerry Li, Vikrant Singhal, and Jonathan Ullman. Privately learning high-dimensional distributions. In *Proceedings of the 32nd Annual Conference on Learning Theory, COLT '19*, pages 1853–1902, 2019.
- [KM17] Adam Klivans and Raghu Meka. Learning graphical models using multiplicative weights. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science, FOCS '17*, pages 343–354, Washington, DC, USA, 2017. IEEE Computer Society.
- [KU20] Gautam Kamath and Jonathan Ullman. A primer on private statistics. *arXiv preprint arXiv:2005.00010*, 2020.
- [KV18] Vishesh Karwa and Salil Vadhan. Finite sample differentially private confidence intervals. In *Proceedings of the 9th Conference on Innovations in Theoretical Computer Science, ITCS '18*, pages 44:1–44:9, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [LAFH01] Charles Lagor, Dominik Aronsky, Marcelo Fiszman, and Peter J. Haug. Automatic identification of patients eligible for a pneumonia guideline: comparing the diagnostic accuracy of two decision support models. *Studies in Health Technology and Informatics*, 84(1):493–497, 2001.
- [LPW09] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2009.
- [LVMC18] Andrey Y. Lokhov, Marc Vuffray, Sidhant Misra, and Michael Chertkov. Optimal structure and parameter learning of Ising models. *Science Advances*, 4(3):e1700791, 2018.
- [MdCCU16] Abraham Martín del Campo, Sarah Cepeda, and Caroline Uhler. Exact goodness-of-fit testing for the Ising model. *Scandinavian Journal of Statistics*, 2016.
- [MMY18] Rajarshi Mukherjee, Sumit Mukherjee, and Ming Yuan. Global testing against sparse alternatives under Ising models. *The Annals of Statistics*, 46(5):2062–2093, 2018.
- [MS10] Andrea Montanari and Amin Saberi. The spread of innovations in social networks. *Proceedings of the National Academy of Sciences*, 107(47):20196–20201, 2010.
- [MSM19] Ryan McKenna, Daniel Sheldon, and Gerome Miklau. Graphical-model based estimation and inference for differential privacy. *arXiv preprint arXiv:1901.09136*, 2019.
- [NRS07] Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the 39th Annual ACM Symposium on the Theory of Computing, STOC '07*, pages 75–84, New York, NY, USA, 2007. ACM.
- [RH17] Philippe Rigollet and Jan-Christian Hütter. High dimensional statistics. <http://www-math.mit.edu/~rigollet/PDFs/RigNotes17.pdf>, 2017. Lecture notes.

- [RWL10] Pradeep Ravikumar, Martin J. Wainwright, and John D. Lafferty. High-dimensional Ising model selection using  $\ell_1$ -regularized logistic regression. *The Annals of Statistics*, 38(3):1287–1319, 2010.
- [Smi11] Adam Smith. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the 43rd Annual ACM Symposium on the Theory of Computing*, STOC '11, pages 813–822, New York, NY, USA, 2011. ACM.
- [SU17] Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *The Journal of Privacy and Confidentiality*, 7(2):3–22, 2017.
- [SW12] Narayana P. Santhanam and Martin J. Wainwright. Information-theoretic limits of selecting binary graphical models in high dimensions. *IEEE Transactions on Information Theory*, 58(7):4117–4134, 2012.
- [TTZ14] Kunal Talwar, Abhradeep Thakurta, and Li Zhang. Private empirical risk minimization beyond the worst case: The effect of the constraint set geometry. *arXiv preprint arXiv:1411.5417*, 2014.
- [TTZ15] Kunal Talwar, Abhradeep Thakurta, and Li Zhang. Nearly-optimal private LASSO. In *Advances in Neural Information Processing Systems 28*, NIPS '15, pages 3025–3033. Curran Associates, Inc., 2015.
- [Vad17] Salil Vadhan. The complexity of differential privacy. In Yehuda Lindell, editor, *Tutorials on the Foundations of Cryptography: Dedicated to Oded Goldreich*, chapter 7, pages 347–450. Springer International Publishing AG, Cham, Switzerland, 2017.
- [VMLC16] Marc Vuffray, Sidhant Misra, Andrey Lokhov, and Michael Chertkov. Interaction screening: Efficient and sample-optimal learning of Ising models. In *Advances in Neural Information Processing Systems 29*, NIPS '16, pages 2595–2603. Curran Associates, Inc., 2016.
- [VTB<sup>+</sup>19] Giuseppe Vietri, Grace Tian, Mark Bun, Thomas Steinke, and Zhiwei Steven Wu. New oracle efficient algorithms for private synthetic data release. *NeurIPS PriML workshop*, 2019.
- [WSD19] Shanshan Wu, Sujay Sanghavi, and Alexandros G. Dimakis. Sparse logistic regression learns all discrete pairwise graphical models. In *Advances in Neural Information Processing Systems 32*, NeurIPS '19, pages 8069–8079. Curran Associates, Inc., 2019.
- [ZLA03] Bianca Zadrozny, John Langford, and Naoki Abe. Cost-sensitive learning by cost-proportionate example weighting. In *Proceedings of the 3rd IEEE International Conference on Data Mining (ICDM 2003)*, 19-22 December 2003, Melbourne, Florida, USA, page 435, 2003.