

Information-Theoretic Limits of Selecting Binary Graphical Models in High Dimensions

Narayana P. Santhanam, *Member, IEEE*, and Martin J. Wainwright, *Senior Member, IEEE*

Abstract—The problem of graphical model selection is to estimate the graph structure of a Markov random field given samples from it. We analyze the information-theoretic limitations of the problem of graph selection for binary Markov random fields under high-dimensional scaling, in which the graph size p and the number of edges k , and/or the maximal node degree d , are allowed to increase to infinity as a function of the sample size n . For pairwise binary Markov random fields, we derive both necessary and sufficient conditions for correct graph selection over the class $\mathcal{G}_{p,k}$ of graphs on p vertices with at most k edges, and over the class $\mathcal{G}_{p,d}$ of graphs on p vertices with maximum degree at most d . For the class $\mathcal{G}_{p,k}$, we establish the existence of constants c and c' such that if $n < ck \log p$, any method has error probability at least $1/2$ uniformly over the family, and we demonstrate a graph decoder that succeeds with high probability uniformly over the family for sample sizes $n > c'k^2 \log p$. Similarly, for the class $\mathcal{G}_{p,d}$, we exhibit constants c and c' such that for $n < cd^2 \log p$, any method fails with probability at least $1/2$, and we demonstrate a graph decoder that succeeds with high probability for $n > c'd^3 \log p$.

Index Terms—High dimensional inference, KL divergence between Ising models, Markov random fields, sample complexity, structure of Ising models.

I. INTRODUCTION

MARKOV random fields, also known as undirected graphical models, provide a structured representation of the joint distributions of families of random variables. They are used in various application domains, among them image processing [5], [14], social network analysis [27], [30], and computational biology [1], [12], [20]. Any Markov random field is associated with an underlying graph that describes conditional independence properties associated with the joint distribution of the random variables. The problem of *graphical model selection* is to recover this unknown graph using samples from the distribution.

Given its relevance in many domains, the graph selection problem has attracted a great deal of attention. The naive approach of searching exhaustively over the space of all graphs is

computationally intractable, since there are $2^{\binom{p}{2}}$ distinct graphs over p vertices. If the underlying graph is known to be tree-structured, then the graph selection problem can be reduced to a maximum-weight spanning tree problem and solved in polynomial time [9]. On the other hand, for general graphs with cycles, the problem is known to be difficult in a complexity-theoretic sense [8]. Nonetheless, a variety of methods have been proposed, including constraint-based approaches [20], [26], thresholding methods [6], and ℓ_1 -based relaxations [13], [21], [22], [24], [33]. Other researchers [11], [19] have analyzed graph selection methods based on penalized forms of pseudolikelihood.

Given a particular procedure for graph selection, classical analysis studies the algorithm's behavior for a fixed graph as the sample size n is increased. In this paper, as with an evolving line of contemporary statistical research, we address the graph selection problem in the *high-dimensional setting*, meaning that we allow the graph size p as well as other structural parameters, such as the number of edges k or the maximum vertex degree d , to scale with the sample size n . We note that a line of recent work has established some high-dimensional consistency results for various graph selection procedures, including methods based on ℓ_1 -regularization for Gaussian models [21], [23], [24], ℓ_1 -regularization for binary discrete Markov random fields [22], thresholding methods for discrete models [6], and variants of the PC algorithm for directed graphical models [20]. All of these methods are practically appealing given their low-computational cost.

Of complementary interest—and the focus of the paper—is the information-theoretic limitations of graphical model selection. More concretely, consider a graph $G = (V, E)$, consisting of a vertex set V with cardinality p , and an edge set E . In this paper, we consider both the class $\mathcal{G}_{p,k}$ of all graphs with $|E| \leq k$ edges, as well as the class $\mathcal{G}_{p,d}$ of all graphs with maximum vertex degree d . Suppose that we are allowed to collect n independent and identically distributed (i.i.d.) samples from a Markov random field defined by some graph $G \in \mathcal{G}_{p,k}$ (or $\mathcal{G}_{p,d}$). Remembering that the graph size p and structural parameters (k, d) are allowed to scale with the sample size, we thereby obtain sequences of statistical inference problems, indexed by the triplet (n, p, k) for the class $\mathcal{G}_{p,k}$, and by the triplet (n, p, d) for the class $\mathcal{G}_{p,d}$. The goal of this paper is to address questions of the following type. First, under what scalings of the triplet (n, p, k) (or correspondingly, the triplet (n, p, d)) is it possible to recover the correct graph with high probability? Conversely, under what scalings of these triplets, does any method fail most of the time?

Although our methods are somewhat more generally applicable, so as to bring sharp focus to these issues, we limit the

Manuscript received November 04, 2009; revised September 11, 2011; accepted January 29, 2012. Date of publication March 21, 2012; date of current version June 12, 2012. The work of M. J. Wainwright was supported in part by the National Science Foundation under Grants CAREER-0545862 and DMS-0528488. The work of N. P. Santhanam was supported by a Startup Grant from the University of Hawaii, Honolulu. This paper was presented in part at the 2008 IEEE International Symposium on Information Theory.

N. P. Santhanam is with the Department of Electrical Engineering, University of Hawaii, Honolulu, HI 96822 USA (e-mail: nsanthan@hawaii.edu).

M. J. Wainwright is with the Department of Electrical Engineering and Computer Science and the Department of Statistics, University of California, Berkeley, CA 94720 USA (e-mail: wainwrig@eecs.berkeley.edu).

Communicated by A. Krzyżak, Associate Editor for Pattern Recognition, Statistical Learning, and Inference.

Digital Object Identifier 10.1109/TIT.2012.2191659

analysis of this paper to the case of pairwise binary Markov random fields, also known as the Ising model. The Ising model is a classical model from statistical physics [4], [18], where it is used to model physical phenomena such as crystal structure and magnetism; more recently, it has been used in image analysis [5], [14], social network modeling [3], [27], and gene network analysis [1], [25].

At a high level, then, the goal of this paper is to understand the information-theoretic capacity of Ising model selection.¹ Our perspective is not unrelated to a line of statistical work in nonparametric estimation [15], [17], [31], [32], in that we view the observation process as a channel communicating information about graphs to the statistician. In contrast to nonparametric estimation, the spaces of possible “codewords” are not function spaces but rather classes of graphs. Accordingly, part of the analysis in this paper involves developing ways in which to measure distances between graphs, and to relate these distances to the Kullback–Leibler (KL) divergence known to control error rates in statistical testing.

We note that understanding of the graph selection capacity can be practically useful in two different ways. On one hand, it can clarify when computationally efficient algorithms achieve information-theoretic limits and, hence, are optimal up to constant factors. On the other hand, it can reveal regimes in which the best known methods to date are suboptimal, thereby motivating the search for new and possibly better methods. Indeed, the analysis of this paper has consequences of both types.

Our main contributions are to establish necessary and sufficient conditions for two classes of graphical models: the class $\mathcal{G}_{p,k}$ of bounded edge cardinality models, and the class $\mathcal{G}_{p,d}$ of bounded vertex degree models. Proofs of the necessary conditions (see Theorems 1 and 2) use indirect methods, based on a version of Fano’s lemma applied to carefully constructed subfamilies of graphs. On the other hand, our proof of the sufficient conditions (see Theorems 3 and 4) is based on direct analysis of explicit “graph decoders.”

The remainder of this paper is organized as follows. We begin in Section II with background on Markov random fields, the classes of graphs considered in this paper, and a precise statement of the graphical model selection problem. In Section III, we state our main results and explore some of their consequences. Section IV is devoted to proofs of the necessary conditions on the sample size (see Theorems 1 and 2), whereas Section V is devoted to proofs of the sufficient conditions. We conclude with a discussion in Section VI.

Notation: For the convenience of the reader, we summarize here notation to be used throughout the paper. We use the following standard notation for asymptotics: we write $f(n) = \mathcal{O}(g(n))$ if $f(n) \leq cg(n)$ for some constant $c < \infty$, and $f(n) = \Omega(g(n))$ if $f(n) \geq c'g(n)$ for some constant $c' > 0$. The notation $f(n) = \Theta(g(n))$ means that $f(n) = \mathcal{O}(g(n))$ and $f(n) = \Omega(g(n))$.

¹In this paper, we assume that the data is drawn from some Ising model in the class $\mathcal{G}_{p,k}$ and $\mathcal{G}_{p,d}$, thus we study the probability of recovering the exact model. However, similar analysis can be applied to the problem of best approximating an arbitrary distribution using an Ising model from class $\mathcal{G}_{p,k}$ or $\mathcal{G}_{p,d}$.

II. BACKGROUND AND PROBLEM FORMULATION

We begin with some background on Markov random fields, and then provide a precise formulation of the problem.

A. Markov Random Fields and Ising Models

An undirected graph $G = (V, E)$ consists a collection $V = \{1, 2, \dots, p\}$ of vertices joined by a collection E of undirected edges.² The neighborhood of any node $s \in V$ is the subset $\mathcal{N}(s) \subset V$

$$\mathcal{N}(s) := \{t \in V \mid (s, t) \in E\} \quad (1)$$

and the degree of vertex s is given by $d_s := |\mathcal{N}(s)|$, corresponding to the cardinality of this neighbor set. We use $d = \max_{s \in V} d_s$ to denote the maximum vertex degree, and $k = |E|$ to denote the total number of edges.

A Markov random field is obtained by associating a random variable X_s to each vertex $s \in V$, and then specifying a joint distribution \mathbb{P} over the random vector (X_1, \dots, X_p) that respects the graph structure in a specific way. In the special case of the Ising models, we consider each random variable X_s takes values $\{-1, +1\}$ and the probability mass function has the form

$$\mathbb{P}_\theta(x_1, \dots, x_p) = \frac{1}{Z(\theta)} \exp \left\{ \sum_{(s,t) \in E} \theta_{st} x_s x_t \right\}. \quad (2)$$

Here the normalization constant $Z(\theta)$, also known as the partition function, is given by

$$Z(\theta) := \sum_{x \in \{-1, +1\}^p} \exp \left\{ \sum_{(s,t) \in E} \theta_{st} x_s x_t \right\}. \quad (3)$$

To be clear, we view the parameter vector θ as an element of $\mathbb{R}^{\binom{p}{2}}$ with the understanding that $\theta_{st} = 0$ for all pairs $(s, t) \notin E$. So as to emphasize the graph-structured nature of the parameter θ , we often use the notation $\theta(G)$.

The edge weight θ_{st} captures the conditional dependence between X_s and X_t , given fixed values for all vertices $\{X_u, u \neq s, t\}$. In particular, a little calculation shows that the conditional distribution takes the form

$$\mathbb{P}_\theta(x_s, x_t \mid x_{V \setminus \{s, t\}}) \propto \exp \left(\theta_{st} x_s x_t + \sum_{u \in \mathcal{N}(s) \setminus t} \theta_{us} x_u x_s + \sum_{u \in \mathcal{N}(t) \setminus s} \theta_{ut} x_u x_t \right).$$

As mentioned earlier, the Ising model (2) has its origins in statistical physics [4], [18], and it has also been used as a simple model in image processing [5], [14], gene network analysis [1], [25], and in modeling social networks [3], [27]. For instance, Banerjee *et al.* [3] use this model to describe the voting behaviors of p politicians, where X_s represents whether politician s voted for ($X_s = +1$) or against ($X_s = -1$) a particular bill. In this case, a positive edge weight $\theta_{st} > 0$ would mean that conditioned on the other politicians’ votes, politician s and t are

²In this paper, we forbid self-loops in the graph, meaning that $(s, s) \notin E$ for all $s \in V$.




Fig. 1. Illustration of the family $\mathcal{G}_{p,k}$ for $p = 3$ and $k = 2$; note that there are three distinct graphs G with $p = 3$ vertices and $k = 2$ edges. Setting the edge parameter $\theta(G) = [\theta \ \theta \ 0]$ induces a family of three Markov random fields. As the edge weight parameter θ increases, the associated distributions $\mathbb{P}_{\theta(G)}$ become arbitrarily difficult to separate.

more likely to agree in their voting (i.e., $X_s = X_t$) than to disagree ($X_s \neq X_t$), whereas a negative edge weight means that they are more likely to disagree.

B. Classes of Graphical Models

In this paper, we consider two different classes of Ising models (2), depending on the condition that we impose on the edge set E . In particular, we consider the two classes of graphs:

- 1) the collection $\mathcal{G}_{p,d}$ of graphs such that each vertex has degree at most d for some $d \geq 1$, and
- 2) the collection $\mathcal{G}_{p,k}$ of graphs G with $|E| \leq k$ edges for some $k \geq 1$.

In addition to the structural properties of the graphs, the difficulty of graph selection also depends on properties of the vector of edge weights $\theta(G) \in \mathbb{R}^{\binom{p}{2}}$. Naturally, one important property is the minimum value over the edges, captured by

$$\lambda^*(\theta(G)) := \min_{(s,t) \in E} |\theta_{st}|. \quad (4)$$

The interpretation of the parameter λ is clear: as in any signal detection problem, it is obviously difficult to detect an interaction θ_{st} if it is extremely close to zero.

Estimation of the graphical structure also turns out to be hard if the edge parameters θ_{st} are large, since large values of edge parameters can mask interactions on other edges. The following example illustrates this point.

Example 1: Consider the family $\mathcal{G}_{p,k}$ of graphs with $p = 3$ vertices and $k = 2$ edges; note that there are a total of 3 such graphs. For each of these three graphs, consider the parameter vector

$$\theta(G) = [\theta \ \theta \ 0]$$

where the single zero corresponds to the single distinct pair $s \neq t$ not in the graph's edge set, as illustrated in Fig. 1.

In the limiting case $\theta = +\infty$, for any choice of graph with two edges, the Ising model distribution enforces the “hard-core” constraint that (X_1, X_2, X_3) must all be equal; that is, for any graph G , the distribution $\mathbb{P}_{\theta(G)}$ places mass $1/2$ on the configuration $[+1 \ +1 \ +1]$ and mass $1/2$ on the configuration $[-1 \ -1 \ -1]$. Of course, this hard-core limit is an extreme case, in which the models are not actually identifiable. Nonetheless, it shows that if the edge weight θ is finite but very large, the models will not be identical, but will be very hard to distinguish.

Motivated by this example, we define the *maximum neighborhood weight*

$$\omega^*(\theta(G)) := \max_{s \in V} \sum_{t \in \mathcal{N}(s)} |\theta_{st}|. \quad (5)$$

Our analysis shows that the number of samples n required to distinguish graphs typically grows exponentially in ω^* .

In this paper, we study classes of Markov random fields that are parameterized by a lower bound λ on the minimum edge weight, and an upper bound ω on the maximum neighborhood weight.

Definition 1 (Classes of Graphical Models):

- a) Given a pair (λ, ω) of positive numbers, the set $\mathcal{G}_{p,d}(\lambda, \omega)$ consists of all distributions $\mathbb{P}_{\theta(G)}$ of the form (2) such that 1) the underlying graph $G = (V, E)$ is a member of the family $\mathcal{G}_{p,d}$ of graphs on p vertices with vertex degree at most d ; 2) the parameter vector $\theta = \theta(G)$ respects the structure of G , meaning that $\theta_{st} \neq 0$ only when $(s, t) \in E$; and 3) the minimum edge weight and maximum neighborhood satisfy the bounds

$$\lambda^*(\theta(G)) \geq \lambda \quad \text{and} \quad \omega^*(\theta(G)) \leq \omega. \quad (6)$$

- b) The set $\mathcal{G}_{p,k}(\lambda, \omega)$ is defined in an analogous manner, with the graph G belonging to the class $\mathcal{G}_{p,k}$ of graphs with p vertices and at most k edges.

We note that for any parameter vector $\theta(G)$, we always have the inequality

$$\omega^*(\theta(G)) \geq \max_{s \in V} |\mathcal{N}(s)| \lambda^*(\theta(G)) \quad (7)$$

so that the families $\mathcal{G}_{p,k}(\lambda, \omega)$ and $\mathcal{G}_{p,d}(\lambda, \omega)$ are only well defined for suitable pairs (λ, ω) .

C. Graph Decoders and Error Criterion

For a given graph class \mathcal{G} (either $\mathcal{G}_{p,d}$ or $\mathcal{G}_{p,k}$) and positive weights (λ, ω) , suppose that nature chooses some member $\mathbb{P}_{\theta(G)}$ from the associated family $\mathcal{G}(\lambda, \omega)$ of Markov random fields. Assume that the statistician observes n samples $\mathbf{X}_1^n := \{X^{(1)}, \dots, X^{(n)}\}$ drawn in an i.i.d. manner from the distribution $\mathbb{P}_{\theta(G)}$. Note that by definition of the Markov random field, each sample $X^{(i)}$ belongs to the discrete set $\mathcal{X} := \{-1, +1\}^p$, so that the overall dataset \mathbf{X}_1^n belongs to the Cartesian product space \mathcal{X}^n .

We assume that the goal of the statistician is to use the data \mathbf{X}_1^n to infer the underlying graph $G \in \mathcal{G}$, which we refer to as the problem of *graphical model selection*. More precisely, we consider functions $\phi : \mathcal{X}^n \rightarrow \mathcal{G}$, which we refer to as graph decoders. We measure the quality of a given graph decoder ϕ using the 0–1 loss function $\mathbb{I}[\phi(\mathbf{X}_1^n) \neq G]$, which takes value 1 when $\phi(\mathbf{X}_1^n) \neq G$ and takes the value 0 otherwise, and we define associated 0–1 risk

$$\mathbb{P}_{\theta(G)}[\phi(\mathbf{X}_1^n) \neq G] = \mathbb{E}_{\theta(G)} [\mathbb{I}[\phi(\mathbf{X}_1^n) \neq G]]$$

corresponding to the probability of incorrect graph selection. Here, the probability (and expectation) are taken with the respect the product distribution of $\mathbb{P}_{\theta(G)}$ over the n i.i.d. samples.

The main purpose of this paper is to study the scaling of the sample sizes n —more specifically, as a function of the graph size p , number of edges k , maximum degree d , minimum edge weight λ , and maximum neighborhood weight ω —that are either sufficient for *some* graph decoder ϕ to output the correct graph with high probability, or conversely, are necessary for *any* graph decoder to output the correct graph with a prespecified probability.

We study two variants of the graph selection problem, depending on whether the values of the edge weights θ are known or unknown. In the *known edge weight* variant, the task of the decoder is to distinguish between graphs, where for any candidate graph $G = (V, E)$, the decoder knows the numerical values of the parameters $\theta(G)$. (Recall that by definition, $[\theta(G)]_{uv} = 0$ for all $(u, v) \notin E$, so that the additional information being provided are the values $[\theta(G)]_{st}$ for all $(s, t) \in E$.) In the *unknown edge weight* variant, both the graph structure and the numerical values of the edge weights are unknown. Clearly, the unknown edge variant is more difficult than the known edge variant. We prove necessary conditions (lower bounds on sample size) for the known edge variant, which are then also valid for the unknown variant. In terms of sufficiency, we provide separate sets of conditions for the known and unknown variants.

III. MAIN RESULTS AND SOME CONSEQUENCES

In this section, we state our main results and then discuss some of their consequences. We begin with statement and discussion of necessary conditions in Section III-A, followed by sufficient conditions in Section III-B.

A. Necessary Conditions

We begin with stating some necessary conditions on the sample size n that any decoder must satisfy for recovery over the families $\mathcal{G}_{p,d}$ and $\mathcal{G}_{p,k}$. Recall (6) for the definitions of λ and ω used in the theorems to follow.

Theorem 1 (Necessary Conditions for $\mathcal{G}_{p,d}$): Consider the family $\mathcal{G}_{p,d}(\lambda, \omega)$ of Markov random fields for some $\omega \geq 1$. If the sample size is upper bounded as

$$n \leq (1 - \delta) \max \left\{ \frac{\log p}{2\lambda \tanh(\lambda)}, \frac{\exp(\omega) \log(\frac{pd}{4} - 1)}{4\lambda d \exp(\lambda)}, \frac{d}{8} \log \frac{p}{8d} \right\} \quad (8)$$

then for any graph decoder $\phi : \mathbf{X}_1^n \rightarrow \mathcal{G}_{p,d}$, whether the edge weights are known or not,

$$\max_{\theta(G) \in \mathcal{G}_{p,d}(\lambda, \omega)} \mathbb{P}_{\theta(G)} [\phi(\mathbf{X}_1^n) \neq G] \geq \delta - \frac{1}{\log p}. \quad (9)$$

Remarks: Let us make some comments regarding the interpretation and consequences of Theorem 1. First, suppose that both the maximum degree d and the minimum edge weight λ remain bounded (i.e., do not increase with the problem sequences). In this case, the necessary conditions (8) can be summarized more compactly as requiring that for some constant c , a sample size $n > \frac{c \log p}{\lambda \tanh(\lambda)}$ is required for bounded degree

graphs. The observation of $\log p$ scaling has also been made in independent work [6], although the dependence on the minimum value λ given here is more refined. Indeed, note that if the minimum edge weight decreases to zero as the sample size increases, then since $\lambda \tanh(\lambda) = \mathcal{O}(\lambda^2)$ as $\lambda \rightarrow 0$, we conclude that a sample size $n > \frac{c' \log p}{\lambda^2}$ is required, for some constant c' .

Some interesting phenomena arise in the case of growing maximum degree d . Observe that in the family $\mathcal{G}_{p,d}$, we necessarily have $\omega \geq \lambda d$. Therefore, in the case of growing maximum degree $d \rightarrow +\infty$, if bound (8) is not to grow exponentially in d , it is necessary to constrain $\lambda = \mathcal{O}(\frac{1}{d})$. But as observed previously, since $\lambda \tanh(\lambda) = \mathcal{O}(\lambda^2)$ as $\lambda \rightarrow 0$, we obtain the following corollary of Theorem 1.

Corollary 1: For the family $\mathcal{G}_{p,d}(\lambda, \omega)$ with increasing maximum degree d , there is a constant $c > 0$ such that in a worst case sense, any method requires at least $n > c \max\{d^2, \lambda^{-2}\} \log p$ samples to recover the correct graph with probability at least $1/2$.

We note that Ravikumar *et al.* [22] have shown that under certain incoherence assumptions (roughly speaking, control on the Fisher information matrix of the distributions $\mathbb{P}_{\theta(G)}$) and assuming that $\lambda = \Omega(d^{-1})$, a computationally tractable method using ℓ_1 -regularization can recover graphs over the family $\mathcal{G}_{p,d}$ using $n > c' d^3 \log p$ samples, for some constant c' ; consequently, Corollary 1 shows concretely that this scaling is within a factor d of the information-theoretic limit.

We now turn to some analogous necessary conditions over the family $\mathcal{G}_{p,k}$ of graphs on p vertices with at most k edges. In doing so, we assume that there is some constant $\gamma \in (0, 2)$ such that $k \geq 8p^\gamma$. This is a mild condition, excluding only graphs in which all except a logarithmic number of the vertices are completely isolated.

Theorem 2 (Necessary Conditions for $\mathcal{G}_{p,k}$): Consider the family $\mathcal{G}_{p,k}(\lambda, \omega)$ of Markov random fields for some $\omega \geq 1$. If for any $\nu \in (0, 1)$ and universal constant c_ν , the sample size is upper bounded as

$$n \leq (1 - \delta) \max \left\{ \frac{\log p}{2\lambda \tanh(\lambda)}, \frac{\exp(\omega) \log(k/8)}{16\omega \exp(2\lambda) \sinh(\lambda)}, c_\nu \gamma k^{1-\nu} \log p \right\} \quad (10)$$

then for any graph decoder $\phi : \mathbf{X}_1^n \rightarrow \mathcal{G}_{p,k}$, whether edge weights are known or not,

$$\max_{\theta(G) \in \mathcal{G}_{p,k}(\lambda, \omega)} \mathbb{P}_{\theta(G)} [\phi(\mathbf{X}_1^n) \neq G] \geq \delta - \frac{1}{\log p}. \quad (11)$$

Remarks: Again, we make some comments about the consequences of Theorem 2. First, suppose that both the number of edges k and the minimum edge weight λ remain bounded (i.e., do not increase with n, p , or k). In this case, the necessary conditions (10) can be summarized more compactly as requiring that for some constant c , a sample size $n > \frac{c \log p}{\lambda \tanh(\lambda)}$ is required for graphs with a constant number of edges. Again, note that if the minimum edge weight decreases to zero as the sample size increases, then since $\lambda \tanh(\lambda) = \mathcal{O}(\lambda^2)$ for $\lambda \rightarrow 0$, we

conclude that a sample size $n > \frac{c' \log p}{\lambda^2}$ is required, for some constant c' .

The behavior is more subtle in the case of graph sequences in which the number of edges k increases with the sample size. As shown in the proof of Theorem 2, it is possible to construct a parameter vector $\theta(G)$ over a graph G with k edges such that $\omega^*(\theta(G)) \geq \lambda \lfloor \sqrt{k} \rfloor$. (More specifically, the construction is based on forming a completely connected subgraph on $\lfloor \sqrt{k} \rfloor$ vertices, which has a total of $\binom{\lfloor \sqrt{k} \rfloor}{2} \leq k$ edges, and then placing a uniform edge weight λ on each edge.) Therefore, if we wish to avoid the exponential growth from the term $\exp(\omega)$, we require that $\lambda = \mathcal{O}(k^{-1/2})$ as the graph size increases. Therefore, we obtain the following corollary of Theorem 2.

Corollary 2: For the family $\mathcal{G}_{p,k}(\lambda, \omega)$ with increasing number of edges k , there is a constant $c > 0$ such that in a worst case sense, any method requires at least $n > c \max\{k, \lambda^{-2}\} \log p$ samples to recover the correct graph with probability at least $1/2$.

To clarify a nuance about comparing Theorems 1 and 2, consider a graph $G \in \mathcal{G}_{p,d}$, with homogeneous degree d at each node. Note that such a graph has a total of $k = dp/2$ edges. Consequently, one might be misled into thinking Corollary 2 implies that $n > c \frac{pd}{2} \log p$ samples would be required in this case. However, as shown in our development of sufficient conditions for the class $\mathcal{G}_{p,d}$ (see Theorem 3), this is not true for sufficiently small degrees d .

To understand the difference, it should be remembered that our necessary conditions are for identifying a graph among all possible choices from the specified graph families. As mentioned, the necessary conditions of Theorem 2 and hence of Corollary 2 are obtained by constructing a graph G that contains a completely connected graph, $K_{\sqrt{k}}$, with uniform degree \sqrt{k} . But $K_{\sqrt{k}}$ is not a member of $\mathcal{G}_{p,d}$ unless $d \geq \sqrt{k}$. On the other hand, for the case when $d \geq \sqrt{k}$, the necessary conditions of Corollary 1 amount to $n > ck \log p$ samples being required, which matches the scaling given in Corollary 2.

B. Sufficient Conditions

We now turn to stating and discussing sufficient conditions (lower bounds on the sample size) for graph recovery over the families $\mathcal{G}_{p,d}$ and $\mathcal{G}_{p,k}$. These conditions provide complementary insight to the necessary conditions discussed so far.

Theorem 3 (Sufficient Conditions for $\mathcal{G}_{p,d}$):

a) Suppose that for some $\delta \in (0, 1)$, the sample size n satisfies

$$n \geq \frac{3[3 \exp(2\omega) + 1]}{\sinh^2(\frac{\lambda}{2})} d \left(3 \log p + \log(2d) + \log \frac{1}{\delta} \right). \quad (12)$$

Then, if the edge weights are known *a priori*, there exists a graph decoder $\phi^* : \mathbf{X}_1^n \rightarrow \mathcal{G}_{p,d}$ whose worst case error probability satisfies

$$\max_{\theta(G) \in \mathcal{G}_{p,d}(\lambda, \omega)} \mathbb{P}_{\theta(G)}[\phi^*(\mathbf{X}_1^n) \neq G] \leq \delta. \quad (13)$$

b) In the case of unknown edge weights, suppose that the sample size satisfies

$$n > \left[\frac{\omega (3 \exp(2\omega) + 1)}{\sinh^2(\lambda/4)} \right]^2 (16 \log p + 4 \log(2/\delta)). \quad (14)$$

Then, there exists a graph decoder $\phi^\dagger : \mathbf{X}_1^n \rightarrow \mathcal{G}_{p,d}$ that has worst case error probability at most δ .

Remarks: It is worthwhile comparing the sufficient conditions provided by Theorem 3 to the necessary conditions from Theorem 1. First, consider the case of finite degree graphs. In this case, condition (12) reduces to the statement that for some constant c , it suffices to have $n > c \lambda^{-2} \log(p/\delta)$ samples in order to obtain the correct graph with probability at least $1 - \delta$. Comparing with the necessary conditions (see the discussion following Theorem 1), we see that for known edge weights and bounded degrees, the information-theoretic capacity scales as $\lambda^{-2} \log p$. For unknown edge weights, conditions (14) provide a weaker guarantee, namely that $n > c' \lambda^{-4} \log p$ samples are required, but we suspect that this guarantee could be improved by a more careful analysis.

Like the necessary conditions in Theorem 1, the sample size specified by the sufficient condition (12) scales *exponentially* in the parameter ω —compare to the middle term in the lower bound (8). Recalling that ω is always at least $d\lambda$, we conclude that if we wish not to incur such exponential growth, we necessarily must have that $\lambda = \mathcal{O}(1/d)$. Consequently, it becomes interesting to compare the dependence on the minimum edge weight λ , especially under the scaling $\lambda \rightarrow 0^+$. The bound (12) depends on λ via the term $\sinh^{-2}(\lambda)^2 = \frac{4}{(\exp(\lambda) - \exp(-\lambda))^2}$. For small λ , we have $\sinh^{-2}(\lambda) \asymp \lambda^{-2}$. Similarly, the necessary condition (8) includes a term involving $(\lambda \tanh(\lambda))^{-1}$, which exhibits the same λ^{-2} scaling as $\lambda \rightarrow 0^+$. Thus, in the case that $\lambda = \mathcal{O}(1/d)$, we obtain the following consequence of Theorem 3.

Corollary 3: Suppose that $\lambda \leq \frac{c_0}{d}$ for some universal constant c_0 . Then, for the graph family $\mathcal{G}_{p,d}(\lambda, \omega)$, there exists a graph decoder that succeeds with probability at least $1 - \delta$ using $n > c_1 \max\{d^2, \lambda^{-2}\} d \log(p/\delta)$ samples.

Comparing to Corollary 1 of Theorem 1, we see that no method has vanishing error probability if $n < c_2 \max\{d^2, \lambda^{-2}\} \log p$, for some constant c_2 . Therefore, together Theorems 1 and 3 provide upper and lower bounds on the sample complexity of graph selection that are matching to within a factor of d . We note that under the condition $\lambda \geq \frac{c_3}{d}$, the results of Ravikumar *et al.* [22] also guarantee correct recovery with high probability for $n > c_4 d^3 \log p$ using ℓ_1 -regularized logistic regression; however, their method requires additional (somewhat restrictive) incoherence assumptions that are not imposed here. In terms of the correct scaling, we conjecture that the scaling $n \asymp \max\{d^2, \lambda^{-2}\} \log p$ provided by our lower bounds in Theorem 1 is the optimal one, since our analysis in this case seems relatively sharp.

Finally, we state sufficient conditions for the class $\mathcal{G}_{p,k}$ in the case of known edge weights.

Theorem 4 (Sufficient Conditions for $\mathcal{G}_{p,k}$):

- a) Suppose that for some $\delta \in (0, 1)$, the sample size n satisfies

$$n > \frac{3 \exp(2\omega) + 1}{\sinh^2(\frac{\lambda}{4})} \left((k+1) \log p + \log \frac{1}{\delta} \right). \quad (15)$$

Then, for known edge weights, there exists a graph decoder $\phi^* : \mathbf{X}_1^n \rightarrow \mathcal{G}_{p,k}$ such that

$$\max_{\theta(G) \in \mathcal{G}_{p,k}(\lambda, \omega)} \mathbb{P}_{\theta(G)} [\phi^*(\mathbf{X}_1^n) \neq G] \leq \delta. \quad (16)$$

- b) For unknown edge weights, there also exists a graph decoder that succeeds with probability at least $1 - \delta$ under condition (14).

Remarks: It is again interesting to compare Theorem 4 with the necessary conditions from Theorem 2. To begin, let the number of edges k remain bounded. In this case, for $\lambda = o(1)$, condition (15) states that for some constant c , it suffices to have $n > \frac{c \log p}{\lambda^2}$ samples, which matches (up to constant factors) the lower bound implied by Theorem 2. In the more general setting of $k \rightarrow +\infty$, we begin by noting that like in Theorem 2, the sample size in Theorem 4 grows exponentially unless the parameter ω stays controlled. As with the discussion following Theorem 2, one interesting scaling is to require that $\lambda \asymp k^{-1/2}$, a choice which controls the worst case construction that leads to the factor $\exp(\omega)$ in the proof of Theorem 2. With this scaling, we have the following consequence.

Corollary 4: Suppose that the minimum value λ scales with the number of edges k as $\lambda \asymp k^{-1/2}$. Then, in the case of known edge weights, there exists a decoder that succeeds with high probability using $n > ck^2 \log p$ samples.

Note that these sufficient conditions are within a factor of k of the necessary conditions from Corollary 2, which show that unless $n > c' \max\{k, \lambda^{-2}\} \log p$, then any graph estimator fails at least half of the time.

IV. PROOFS OF NECESSARY CONDITIONS

In the following two sections, we provide the proofs of our main theorems. We begin by introducing some background on distances between distributions, as well as some results on the cardinalities of our model classes. We then provide proofs of the necessary conditions (see Theorems 1 and 2) in this section, followed by the proofs of the sufficient conditions (see Theorems 3 and 4) in Section V.

A. Preliminaries

We begin with some preliminary definitions and results concerning “distance” measures between different models, and some estimates of the cardinalities of different model classes.

1) *Distance Measures:* In order to quantify the distinguishability of different models, we begin by defining some useful “distance” measures. Given two parameters θ and θ' in $\mathbb{R}^{\binom{p}{2}}$, we let $D(\theta \parallel \theta')$ denote the KL divergence [10] between the two

distributions \mathbb{P}_θ and $\mathbb{P}_{\theta'}$. For the special case of the Ising model distributions (2), this KL divergence takes the form

$$D(\theta \parallel \theta') := \sum_{x \in \{-1, +1\}^p} \mathbb{P}_\theta(x) \log \frac{\mathbb{P}_\theta(x)}{\mathbb{P}_{\theta'}(x)}. \quad (17)$$

Note that the KL divergence is not symmetric in its arguments (i.e., $D(\theta \parallel \theta') \neq D(\theta' \parallel \theta)$ in general).

Our analysis also makes use of two other closely related divergence measures, both of which are symmetric. First, we define the *symmetrized KL divergence*, defined in the natural way as

$$S(\theta \parallel \theta') := D(\theta \parallel \theta') + D(\theta' \parallel \theta). \quad (18)$$

Second, given two parameter vectors θ and θ' , we consider the model $\mathbb{P}_{\frac{\theta + \theta'}{2}}$ specified by their average. Using this averaged model, we define another type of divergence as

$$J(\theta \parallel \theta') := D\left(\frac{\theta + \theta'}{2} \parallel \theta\right) + D\left(\frac{\theta + \theta'}{2} \parallel \theta'\right). \quad (19)$$

Note that this divergence is also symmetric in its arguments. A straightforward calculation shows that this divergence measure can be expressed in terms of the partition function (3) associated with the Ising family as

$$J(\theta \parallel \theta') = \log \frac{Z(\theta)Z(\theta')}{Z^2\left(\frac{\theta + \theta'}{2}\right)}. \quad (20)$$

Useful in our analysis are representations of these distance measures in terms of the vector of *mean parameters* $\mu(\theta) \in \mathbb{R}^{\binom{p}{2}}$, where element μ_{st} is given by

$$\mu_{st} := \mathbb{E}_\theta[X_s X_t] = \sum_{x \in \{-1, +1\}^p} \mathbb{P}_\theta(x) x_s x_t. \quad (21)$$

It is well known from the theory of exponential families [7], [28] that there is a bijection between the canonical parameters θ and the mean parameters μ .

Using this notation, a straightforward calculation shows that the symmetrized KL divergence between \mathbb{P}_θ and $\mathbb{P}_{\theta'}$ is equal to

$$S(\theta \parallel \theta') = \frac{1}{2} \sum_{s, t \in V, s \neq t} (\theta_{st} - \theta'_{st}) (\mu_{st} - \mu'_{st}) \quad (22)$$

where μ_{st} and μ'_{st} denote the edge-based mean parameters under θ and θ' respectively.

2) *Cardinalities of Graph Classes:* In addition to these divergence measures, we require some estimates of the cardinalities of the graph classes $\mathcal{G}_{p,d}$ and $\mathcal{G}_{p,k}$, as summarized in the following.

Lemma 1:

- a) For $k \leq \binom{p}{2}/2$, the cardinality of $\mathcal{G}_{p,k}$ is bounded as

$$\binom{\binom{p}{2}}{k} \leq |\mathcal{G}_{p,k}| \leq k \binom{\binom{p}{2}}{k} \quad (23)$$

and hence $\log |\mathcal{G}_{p,k}| = \Theta(k \log \frac{p}{\sqrt{k}})$.

b) For $d \leq \frac{p-1}{2}$, the cardinality of $\mathcal{G}_{p,d}$ is bounded as

$$\left\lfloor \frac{p}{d+1} \right\rfloor!^{\frac{d(d+1)}{2}} \leq |\mathcal{G}_{p,d}| \leq \frac{pd}{2} \binom{\frac{p}{2}}{\frac{pd}{2}} \quad (24)$$

and hence $\log |\mathcal{G}_{p,d}| = \Theta\left(pd \log \frac{p}{d}\right)$.

Proof:

- a) For the bounds (23) on $|\mathcal{G}_{p,k}|$, we observe that there are $\binom{p}{\ell}$ graphs with exactly ℓ edges, and that for $k \leq \frac{p}{2}/2$, we have $\binom{p}{\ell} \leq \binom{p}{k}$ for all $\ell = 1, 2, \dots, k$.
- b) Turning to the bounds (24) on $|\mathcal{G}_{p,d}|$, observe that every model in $\mathcal{G}_{p,d}$ has at most $\frac{pd}{2}$ edges. Note that $d \leq \frac{p-1}{2}$ ensures that

$$\frac{pd}{2} \leq \binom{p}{2}/2.$$

Therefore, following the argument in part (a), we conclude that $|\mathcal{G}_{p,d}| \leq \frac{pd}{2} \binom{\frac{p}{2}}{\frac{pd}{2}}$, as claimed.

In order to establish the lower bound (24), we first group the p vertices into $d+1$ groups of size $\lfloor \frac{p}{d+1} \rfloor$, discarding any remaining vertices. We consider a subset of $\mathcal{G}_{p,d}$: graphs with maximum degree d each edge of the graph straddling vertices in two different groups.

To construct one such graph, we pick a permutation of $\left\lfloor \frac{p}{d+1} \right\rfloor$ and form an bijection between vertices of group 1 and group 2 representing that permutation. Similarly, we form an bijection from groups 1 to 3, and so on up until $d+1$. Note that use d permutations to complete this procedure, and at the end of this round, every vertex in group 1 has degree d , vertices in all other groups have degree 1.

Similarly, in the next round, we use $d-1$ permutations to connect group 2 to group 3 through $d+1$. In general, for $i = 1, \dots, d$, in round i , we use $d+1-i$ permutations to connect group i with groups $i+1, \dots, d+1$. Each choice of these permutations yields a distinct graph in $\mathcal{G}_{p,d}$. Note that we use a total of

$$\sum_{i=1}^d (d+1-i) = \sum_{\ell=1}^d \ell = \frac{d(d+1)}{2}$$

permutations over $\lfloor \frac{p}{d+1} \rfloor$ elements, from which the stated claim (24) follows. ■

3) *Fano's Lemma and Variants:* We provide some background on Fano's lemma and its variants needed in our arguments. Consider a family of M models indexed by the parameter vectors $\{\theta^{(1)}, \theta^{(2)}, \dots, \theta^{(M)}\}$. Suppose we choose a model index K uniformly at random from $\{1, \dots, M\}$; given $K = k$, we obtain a dataset \mathbf{X}_1^n of n samples drawn in an i.i.d. manner according to a distribution $\mathbb{P}_{\theta^{(k)}}$. In this setting, Fano's lemma and related inequalities provide lower bounds on the probability of error of any classification function $\phi : \mathcal{X}^n \rightarrow \{1, \dots, M\}$, specified in terms of the mutual information

$$I(\mathbf{X}_1^n; K) = H(\mathbf{X}_1^n) - H(\mathbf{X}_1^n | K) \quad (25)$$

Authorized licensed use limited to: Universitaet Linz. Downloaded on April 20, 2025 at 19:57:26 UTC from IEEE Xplore. Restrictions apply.

between the data \mathbf{X}_1^n and the random model index K . We say that a decoder $\phi : \mathcal{X}^n \rightarrow \{1, \dots, M\}$ is δ -unreliable over the family $\{\theta^{(1)}, \dots, \theta^{(M)}\}$ if

$$\max_{k=1, \dots, M} \mathbb{P}_{\theta^{(k)}} [\phi(\mathbf{X}_1^n) \neq k] \geq \delta - \frac{1}{\log M}. \quad (26)$$

In this paper, we use the following forms of Fano's inequality.

Lemma 2: Any of the following upper bounds on the sample size imply that any decoder is δ -unreliable over the family $\{\theta^{(1)}, \dots, \theta^{(M)}\}$.

- a) The sample size n is upper bounded as

$$n < \frac{(1-\delta) \log(M)}{I(X^{(1)}; K)}. \quad (27)$$

- b) The sample size n is upper bounded as

$$n < (1-\delta) \frac{\log(M)}{\frac{2}{M^2} \sum_{k=1}^M \sum_{\ell=k+1}^M S(\theta^{(k)} || \theta^{(\ell)})}. \quad (28)$$

We note that these variants of Fano's inequality are standard and widely used in the nonparametric statistics literature (see e.g., [15], [17], [31], and [32]).

B. Key Separation Result

In order to exploit (28), one needs to construct families of models with relatively large cardinality (M large) such that the models are all relatively close in symmetrized KL divergence. Recalling definition (21) of the mean parameters and the form of the symmetrized KL divergence (22), we see that control of the divergence between \mathbb{P}_{θ} and $\mathbb{P}_{\theta'}$ can be achieved by ensuring that their respective mean parameters μ_{st} stay relatively close for all edges (s, t) where the models differ.

In this section, we state and prove a key technical lemma that allows us to construct such a class of models. As shown in the proofs of Theorems 1 and 2 to follow, this lemma allows us to gain good control on the symmetrized KL divergences between pairs of models. Our construction of the model class, which applies to any integer $m \geq 2$, is based on the following procedure. We begin with the complete graph on $m+1$ vertices, denoted by K_{m+1} . We then form a set of $\binom{m+1}{2}$ graphs, each of which is a subgraph of K_{m+1} , by removing a particular edge. Now denoting by G^{st} the subgraph with edge (s, t) removed, we define the Ising model distribution $\mathbb{P}_{\theta(G^{st})}$ by setting $[\theta(G^{st})]_{uv} = \lambda$ for all edges $(u, v) \neq (s, t)$, and $[\theta(G^{st})]_{st} = 0$. Thus, we are considering a graph ensemble with uniform edge weights.

The following lemma shows that the mean parameter $\mu_{st} = \mathbb{E}_{\theta(G^{st})}[X_s X_t]$ approaches its maximum value 1 exponentially quickly in the parameter $\omega = \lambda m$.

Lemma 1: For some integer $m, 1 \leq m \leq p-1$, consider the ensemble of $\binom{m+1}{2}$ graphs with uniform edge weights previously defined. Suppose that $\omega = \lambda m \geq \frac{2}{m-1}$. Then, the likelihood ratio on edge (s, t) is lower bounded as

$$\frac{\mathbb{P}_{\theta(G^{st})}[X_s X_t = +1]}{\mathbb{P}_{\theta(G^{st})}[X_s X_t = -1]} = \frac{q_{st}}{1 - q_{st}} \geq \frac{\exp(\omega - \lambda)}{m} \quad (29)$$

and moreover, the mean parameter over the pair (s, t) is lower bounded as

$$\mathbb{E}_{\theta(G^{st})}[X_s X_t] \geq 1 - \frac{2m \exp(\lambda)}{\exp(\omega) + m \exp(\lambda)}. \quad (30)$$

Proof: Let us introduce the convenient shorthand $q_{st} = \mathbb{P}_{\theta(G^{st})}[X_s X_t = 1]$. We begin by observing that the bound (29) implies the bound (30). To see this, suppose that (29) holds, or equivalently that $q_{st} \geq \frac{b}{1+b}$ where $b = \frac{\exp(\omega-2\lambda)}{m+1}$. Observing that $\mathbb{E}_{\theta(G^{st})}[X_s X_t] = 2q_{st} - 1$, we see that

$$\mathbb{E}_{\theta(G^{st})}[X_s X_t] \geq \frac{2b}{1+b} - 1 = 1 - \frac{2}{1+b}.$$

from which (30) follows.

The remainder of our proof is devoted to proving the lower bound (29). Our first step is to establish the relation

$$\frac{q_{st}}{1 - q_{st}} = \frac{\sum_{j=0}^{m-1} \binom{m-1}{j} \exp\left(\frac{\lambda}{2} [(2j - m + 3)^2 - 4]\right)}{\sum_{j=0}^{m-1} \binom{m-1}{j} \exp\left(\frac{\lambda}{2} [(2j - m + 1)^2]\right)}. \quad (31)$$

In order to prove (31), let us consider a fixed configuration $X \in \{-1, +1\}^p$ such that $X_s X_t = 1$. Define the subsets

$$S_+ := \{u \in \{1, \dots, p\} \setminus \{s, t\} \mid X_u = +1\} \quad \text{and} \\ S_- := \{u \in \{1, \dots, p\} \setminus \{s, t\} \mid X_u = -1\}.$$

If the set S_+ has j elements in total, then the contribution to the exponential weight from variables in S_+ alone is $\lambda \binom{j}{2}$, from variables in S_- alone is $\lambda \binom{m-1-j}{2}$, and the cross terms between S_+ and S_- is $-\lambda j(m-1-j)$. Depending on whether X_s and X_t are both $+1$ or -1 , the contribution from terms involving either X_s or X_t is $2\lambda(2j - m + 1)$ or $-2\lambda(2j - m + 1)$. Overall, the probability of configurations for which $X_s X_t = +1$ and $|S_+| = j$, denoted by $r_{st} = \mathbb{P}[X_s X_t = 1, |S_+| = j]$, is given by

$$r_{st} = \frac{1}{Z} \exp\left(\lambda \binom{j}{2} + \lambda \binom{m-1-j}{2} - \lambda j(m-1-j)\right) \\ \left[e^{(2\lambda(2j-m+1))} + e^{(-2\lambda(2j-m+1))} \right] \\ = \frac{1}{Z} \exp\left(\lambda \binom{j}{2} + \lambda \binom{m-1-j}{2} - \lambda j(m-1-j)\right) \\ \left[e^{(2\lambda(2j-m+1))} + e^{(2\lambda(2(m-1-j)-m+1))} \right]$$

where Z is the partition function. Noting that

$$\binom{j}{2} + \binom{m-1-j}{2} - j(m-1-j) = \frac{1}{2} ((2j - m + 1)^2 - (m - 1))$$

and that there are $\binom{m-1}{j}$ ways of choosing j variables other than X_s and X_t to be $+1$, we find that

$$q_{st} = \frac{1}{Z} \sum_{j=0}^{m-1} \binom{m-1}{j} \exp\left(\frac{\lambda}{2} ((2j - m + 1)^2 - m + 1)\right) \\ \left[e^{(2\lambda(2j-m+1))} + e^{(2\lambda(2(m-1-j)-m+1))} \right] \\ = \frac{2}{Z} \sum_{j=0}^{m-1} \binom{m-1}{j} \exp\left(\frac{\lambda}{2} ((2j - m + 1)^2 - m + 1)\right) \\ \left[e^{(2\lambda(2j-m+1))} \right] \\ = \frac{2}{Z} \sum_{j=0}^{m-1} \binom{m-1}{j} \exp\left(\frac{\lambda}{2} ((2j - m + 3)^2 - m - 3)\right).$$

Similarly, we obtain

$$1 - q_{st} = \frac{1}{Z} \sum_{j=0}^{m-1} \binom{m-1}{j} \exp\left(\frac{\lambda}{2} ((2j - m + 1)^2 - m + 1)\right) \\ (1 + 1) \\ = \frac{2}{Z} \sum_{j=0}^{m-1} \binom{m-1}{j} \exp\left(\frac{\lambda}{2} ((2j - m + 1)^2 - m + 1)\right)$$

and combining the pieces yields the bound (31).

It remains to use the representation (31) to prove the stated lower bound. We lower bound the ratio in (31) by identifying the largest terms in the denominator. To do so, note that the ratio between adjacent terms in the denominator is

$$\frac{\binom{m-1}{j+1} \exp\left(\frac{\lambda}{2} [(2j - m + 3)^2]\right)}{\binom{m-1}{j} \exp\left(\frac{\lambda}{2} [(2j - m + 1)^2]\right)} \\ = \frac{(m-1-j)}{j+1} \exp\left(\frac{\lambda}{2} [2(4j - 2m + 4)]\right).$$

Consequently, if we write $j = (m-1)/2 - (m-1)\alpha - \frac{1}{2}$, then we have $4j - 2m + 4 = -4(m-1)\alpha$. Hence, if $\lambda(m-1) \geq 2$, the ratio between $j+1$ th and j th terms is

$$\frac{1 + 2\alpha + \frac{1}{(m-1)}}{1 - 2\alpha - \frac{1}{(m-1)}} e^{-4\lambda(m-1)\alpha} \leq \frac{1 + 2\alpha + \frac{1}{(m-1)}}{1 - 2\alpha - \frac{1}{(m-1)}} e^{-8\alpha}.$$

Note that if $\alpha \neq 0$, then $\alpha \geq \frac{1}{m-1}$. If $2\alpha + \frac{1}{(m-1)} \leq 1/2$ and $\alpha \neq 0$

$$\frac{1 + 2\alpha + \frac{1}{(m-1)}}{1 - 2\alpha - \frac{1}{(m-1)}} \leq e^{(2\alpha + \frac{1}{(m-1)})(1+2 \ln 2) - 8\alpha} < 1.$$

Consequently, the maximum term in the summation cannot occur when $2\alpha + \frac{1}{(m-1)} \leq 1/2$ nor when $(m-1)/4 + 1 \leq j \leq (m-1)/2 - 1$. By inspection, the $\frac{(m-1)}{2}$ th term is smaller than the $\frac{(m-1)}{4}$ th term whenever

$\lambda(m-1) \geq 2$. Furthermore, by symmetry, we also have a maximum when $j \geq 3(m-1)/4$.

Picking a maximum $j^* > 3(m-1)/4$, using the representation (31), and noting that all terms in the numerator are nonnegative, we find that

$$\begin{aligned} \frac{q_{st}}{1-q_{st}} &= \frac{\sum_{j=0}^{m-1} \binom{m-1}{j} \exp\left(\frac{\lambda}{2} [(2j-m+3)^2-4]\right)}{\sum_{j=0}^{m-1} \binom{m-2}{j} \exp\left(\frac{\lambda}{2} [(2j-m+1)^2]\right)} \\ &= \frac{\binom{m-1}{j^*} \exp\left(\frac{\lambda}{2} [(2j^*-m+3)^2-4]\right)}{m \binom{m-1}{j^*} \exp\left(\frac{\lambda}{2} [(2j^*-m+1)^2]\right)} \\ &\geq \frac{\exp(\lambda[(4j^*-2m+4)-2])}{m} \\ &\geq \frac{\exp(\lambda(m-1))}{m} \\ &= \frac{m}{\exp(\omega-\lambda)} \end{aligned}$$

which completes the proof. \blacksquare

C. Proof of Theorem 1

We begin with necessary conditions for the bounded degree family $\mathcal{G}_{p,d}$. The proof is based on applying Fano's inequality to three ensembles of graphical models, each contained within the family $\mathcal{G}_{p,d}(\lambda, \omega)$. Since the worst case performance is no better than the performance over any subfamily, any such ensemble provides a lower bound on the performance.

Ensemble A: In this ensemble, we consider the set of $\binom{p}{2}$ graphs, each of which contains a single edge. For each such graph—say the one containing edge (s, t) , which we denote by H_{st} —we set $[\theta(H_{st})]_{st} = \lambda$, and all other entries equal to zero. Clearly, the resulting Markov random fields $\mathbb{P}_{\theta(H)}$ all belong to the family $\mathcal{G}_{p,d}(\lambda, \omega)$. (Note that by definition, we must have $\omega \geq \lambda$ for the family to be nonempty.) Let us compute the symmetrized KL divergence between the MRFs indexed by $\theta(G_{st})$ and $\theta(G_{uv})$. Using the representation (22), we have

$$\begin{aligned} S(\theta(H_{st}) \parallel \theta(H_{uv})) &= \lambda \left\{ (\mathbb{E}_{\theta(H_{st})}[X_s X_t] \mathbb{E}_{\theta(H_{uv})}[X_s X_t]) - \right. \\ &\quad \left. - (\mathbb{E}_{\theta(H_{st})}[X_u X_v] - \mathbb{E}_{\theta(H_{uv})}[X_u X_v]) \right\} \\ &= 2\lambda \mathbb{E}_{\theta(H_{st})}[X_s X_t] \end{aligned}$$

since $\mathbb{E}_{\theta(H_{st})}[X_u X_v] = 0$ for all $(u, v) \neq (s, t)$, and $\mathbb{E}_{\theta(H_{uv})}[X_u X_v] = \mathbb{E}_{\theta(H_{st})}[X_s X_t]$. Finally, by definition of the distribution $\mathbb{P}_{\theta(H_{st})}$, we have

$$\mathbb{E}_{\theta(H_{st})}[X_s X_t] = \frac{\exp(\lambda) - \exp(-\lambda)}{\exp(\lambda) + \exp(-\lambda)} = \tanh(\lambda)$$

so that we conclude that the symmetrized KL divergence is equal to $2\lambda \tanh(\lambda)$ for each pair.

Using the bound (28) from Lemma 2 with $M = \binom{p}{2}$, we conclude that the graph recovery error is at least $\delta - \frac{1}{\log\left(\frac{p}{2}\right)}$ if the sample size is upper bounded as

$$n < \frac{(1-\delta) \log\left(\frac{p}{2}\right)/4}{\lambda \tanh(\lambda)}. \quad (32)$$

Ensemble B: In order to form this graph ensemble, we begin with a grouping of the p vertices into $\lfloor \frac{p}{d+1} \rfloor$ groups, each with $d+1$ vertices. We then consider the graph \bar{G} obtained by fully connecting each subset of $d+1$ vertices, namely \bar{G} is a graph that contains $\lfloor \frac{p}{d+1} \rfloor$ cliques of size $d+1$. Using this base graph, we form a collection of graphs by beginning with \bar{G} , and then removing a single edge (u, v) . We denote the resulting graph by G^{uv} . Note that if $p \geq 2(d+1)$, then we can form

$$\lfloor \frac{p}{d+1} \rfloor \binom{d+1}{2} \geq \frac{pd}{4}$$

such graphs. For each graph G^{uv} , we form an associated Markov random field $\mathbb{P}_{\theta(G^{uv})}$ by setting $[\theta(G^{uv})]_{ab} = \lambda > 0$ for all (a, b) in the edge set of G^{uv} , and setting the parameter to zero otherwise.

A central component of the argument is the following bound on the symmetrized KL divergence between these distributions

Lemma 2: For all distinct pairs of models $\theta(G^{st}) \neq \theta(G^{uv})$ in ensemble B and for all $\lambda \geq 1/d$, the symmetrized KL divergence is upper bounded as

$$S(\theta(G^{st}) \parallel \theta(G^{uv})) \leq \frac{4\lambda d \exp(\lambda)}{\exp(\lambda d)}.$$

Proof: Note that any pair of distinct parameter vectors $\theta(G^{st}) \neq \theta(G^{uv})$ differ in exactly two edges. Consequently, by the representation (22), and the definition of the parameter vectors,

$$\begin{aligned} S(\theta(G^{st}) \parallel \theta(G^{uv})) &= \lambda (\mathbb{E}_{\theta(G^{uv})}[X_s X_t] - \mathbb{E}_{\theta(G^{st})}[X_s X_t]) \\ &\quad + \lambda (\mathbb{E}_{\theta(G^{st})}[X_u X_v] - \mathbb{E}_{\theta(G^{uv})}[X_u X_v]) \\ &\leq \lambda (1 - \mathbb{E}_{\theta(G^{st})}[X_s X_t]) + \lambda (1 - \mathbb{E}_{\theta(G^{uv})}[X_u X_v]) \end{aligned}$$

where the inequality uses the fact that $\lambda > 0$, and the edge-based mean parameters are upper bounded by 1.

Since the model $\mathbb{P}_{\theta(G^{st})}$ factors as a product of separate distributions over the $\lfloor \frac{p}{d+1} \rfloor$ cliques, we can now apply the separation result (30) from Lemma 1 with $m = d$ to conclude that

$$\begin{aligned} S(\theta(G^{st}) \parallel \theta(G^{uv})) &\leq 2\lambda \frac{2d \exp(\lambda)}{\exp(\lambda d) + d \exp(\lambda)} \\ &\leq \frac{4\lambda d \exp(\lambda)}{\exp(\lambda d)} \end{aligned}$$

as claimed. \blacksquare

Using Lemma 2 and applying the bound (28) from Lemma 2 with $M = \frac{pd}{4}$ yields that for probability of error below δ and $\lambda d \geq 2$, we require at least

$$n > (1-\delta) \frac{\log\left(\frac{pd}{4} - 1\right)}{S(\theta(G^{st}) \parallel \theta(G^{uv}))} \geq \frac{\exp(d\lambda) \log\left(\frac{pd}{4} - 1\right)}{4d\lambda \exp(\lambda)}$$

samples. Since $\omega = d\lambda$ in this construction, we conclude that

$$n > (1-\delta) \frac{\exp(\omega) \log\left(\frac{pd}{4} - 1\right)}{4\lambda d \exp(\lambda)}$$

samples are required, in order to obtain error below $\delta - \frac{1}{\log(pd/4)}$, as claimed in Theorem 1.

Ensemble C: Finally, we prove the third component in the bound (8). In this case, we consider the ensemble consisting of all graphs in $\mathcal{G}_{p,d}$. From Lemma 1(b), we have

$$\begin{aligned} \log |\mathcal{G}_{p,d}| &\geq \frac{d(d+1)}{2} \log \left\lfloor \frac{p}{d+1} \right\rfloor! \\ &\geq \frac{d(d+1)}{2} \left\lfloor \frac{p}{d+1} \right\rfloor \log \frac{\left\lfloor \frac{p}{d+1} \right\rfloor}{e} \\ &\geq \frac{dp}{4} \log \frac{p}{8d}. \end{aligned}$$

For this ensemble, it suffices to use a trivial upper bound on the mutual information (25), namely

$$I(\mathbf{X}_1^n; G) \leq H(\mathbf{X}_1^n) \leq np$$

where the second bound follows since \mathbf{X}_1^n is a collection of np binary variables, each with entropy at most 1. Therefore, from the Fano bound (27), we conclude that the error probability stays above $\delta - \frac{1}{\log |\mathcal{G}_{p,d}|}$ if the sample size n is upper bounded as $n < (1 - \delta) \frac{d}{8} \log \frac{p}{8d}$, as claimed.

D. Proof of Theorem 2

We now turn to the proof of necessary conditions for the graph family $\mathcal{G}_{p,k}$ with at most k edges. As with the proof of Theorem 2, it is based on applying Fano's inequality to three ensembles of Markov random fields contained in $\mathcal{G}_{p,k}(\lambda, \omega)$.

Ensemble A: Note that the ensemble (A) previously constructed in the proof of Theorem 1 is also valid for the family $\mathcal{G}_{p,k}(\lambda, \omega)$, and hence the bound (32) is also valid for this family.

Ensemble B: For this ensemble, we choose the largest integer m such that $k+1 \geq \binom{m}{2}$. Note that we have

$$m \geq \lfloor \sqrt{k} \rfloor \geq \frac{\sqrt{k}}{2}.$$

We then form a family of $\binom{m}{2}$ graphs as follows: 1) first form the complete graph K_m on a subset of m vertices, and 2) for each $(s, t) \in K_m$, form the graph G^{st} by removing edge (s, t) from K_m . We form Markov random fields on these graphs by setting $[\theta(G^{st})]_{wz} = \lambda$ on every edge (wz) of G^{st} .

Lemma 3: For all distinct model pairs $\theta(G^{st})$ and $\theta(G^{uv})$, we have

$$S(\theta(G^{st}) \parallel \theta(G^{uv})) \leq \frac{16\omega \exp(\frac{5\lambda}{2}) \sinh(\lambda)}{\exp(\frac{\omega}{2})}. \quad (33)$$

Proof: We begin by claiming for any pair $(s, t) \neq (u, v)$, the distribution $\mathbb{P}_{\theta(G^{uv})}$ (i.e., corresponding to the subgraph that does not contain edge (u, v)) satisfies

$$\begin{aligned} \frac{\mathbb{P}_{\theta(G^{uv})}[X_s X_t = +1]}{\mathbb{P}_{\theta(G^{uv})}[X_s X_t = -1]} &\leq \exp(2\lambda) \frac{\mathbb{P}_{\theta(G^{st})}[X_s X_t = +1]}{\mathbb{P}_{\theta(G^{st})}[X_s X_t = -1]} \\ &= \exp(2\lambda) \frac{q_{st}}{1 - q_{st}} \end{aligned} \quad (34)$$

where we have reintroduced the convenient shorthand $q_{st} = \mathbb{P}_{\theta(G^{st})}[X_s X_t = 1]$ from Lemma 1.

To prove this claim, let \mathbb{P}_θ be the distribution that contains all edges in the complete subgraph K_m , each with weight λ . Let $Z(\theta)$ and $Z(\theta(G^{uv}))$ be the normalization constants associated with \mathbb{P}_θ and $\mathbb{P}_{\theta(G^{uv})}$, respectively. Now since $\lambda > 0$ by assumption, by the FKG inequality [2], we have

$$\frac{\mathbb{P}_{\theta(G^{uv})}[X_s X_t = +1]}{\mathbb{P}_{\theta(G^{uv})}[X_s X_t = -1]} \leq \frac{\mathbb{P}_\theta[X_s X_t = +1]}{\mathbb{P}_\theta[X_s X_t = -1]}.$$

We now apply the definition of \mathbb{P}_θ and expand the right-hand side of this expression, recalling the fact that the model $\mathbb{P}_{\theta(G^{st})}$ does not contain the edge (s, t) . Thus, we obtain

$$\begin{aligned} &\frac{\mathbb{P}_\theta[X_s X_t = +1]}{\mathbb{P}_\theta[X_s X_t = -1]} \\ &= \frac{\exp(\lambda) \frac{Z(\theta(G^{st}))}{Z(\theta)} \mathbb{P}_{\theta(G^{st})}[X_s X_t = +1]}{\exp(-\lambda) \frac{Z(\theta(G^{st}))}{Z(\theta)} \mathbb{P}_{\theta(G^{st})}[X_s X_t = -1]} \\ &= \exp(2\lambda) \frac{\mathbb{P}_{\theta(G^{st})}[X_s X_t = +1]}{\mathbb{P}_{\theta(G^{st})}[X_s X_t = -1]} \end{aligned}$$

which establishes the claim (34).

Finally, from the representation (22) for the symmetrized KL divergence and the definition of the models

$$\begin{aligned} S(\theta(G^{st}) \parallel \theta(G^{uv})) &= \lambda \{ \mathbb{E}_{\theta(G^{st})}[X_u X_v] - \mathbb{E}_{\theta(G^{uv})}[X_u X_v] \} \\ &\quad + \lambda \{ \mathbb{E}_{\theta(G^{uv})}[X_s X_t] - \mathbb{E}_{\theta(G^{st})}[X_s X_t] \} \\ &= 2\lambda \{ \mathbb{E}_{\theta(G^{uv})}[X_s X_t] - \mathbb{E}_{\theta(G^{st})}[X_s X_t] \} \end{aligned}$$

where we have used the symmetry of the two terms. Continuing on, we observe the decomposition $\mathbb{E}_{\theta(G^{st})}[X_s X_t] = 1 - 2\mathbb{P}_{\theta(G^{st})}[X_s X_t = -1]$, and using the analogous decomposition for the other expectation, we obtain

$$\begin{aligned} S(\theta(G^{st}) \parallel \theta(G^{uv})) &= 4\lambda \{ \mathbb{P}_{\theta(G^{st})}[X_s X_t = -1] - \mathbb{P}_{\theta(G^{uv})}[X_s X_t = -1] \} \\ &= 4\lambda \left\{ \frac{1}{\frac{\mathbb{P}_{\theta(G^{st})}[X_s X_t = +1]}{\mathbb{P}_{\theta(G^{st})}[X_s X_t = -1]} + 1} - \frac{1}{\frac{\mathbb{P}_{\theta(G^{uv})}[X_s X_t = +1]}{\mathbb{P}_{\theta(G^{uv})}[X_s X_t = -1]} + 1} \right\} \\ &\stackrel{(a)}{\leq} 4\lambda \left\{ \frac{1}{\frac{q_{st}}{1 - q_{st}} + 1} - \frac{1}{\exp(2\lambda) \frac{q_{st}}{1 - q_{st}} + 1} \right\} \\ &= \frac{4\lambda}{\frac{q_{st}}{1 - q_{st}}} \left\{ \frac{1}{1 + \frac{1 - q_{st}}{q_{st}}} - \frac{1}{\exp(2\lambda) + \frac{1 - q_{st}}{q_{st}}} \right\} \\ &= \frac{4\lambda (\exp(2\lambda) - 1)}{\frac{q_{st}}{1 - q_{st}}} \frac{1}{\left[1 + \frac{1 - q_{st}}{q_{st}} \right] \left[\exp(2\lambda) + \frac{1 - q_{st}}{q_{st}} \right]} \end{aligned}$$

where in obtaining the inequality (a), we have applied the bound (34) and recalled our shorthand notation $q_{st} = \mathbb{P}_{\theta(G^{st})}[X_s X_t = +1]$. Since $\lambda > 0$ and $(1 - q_{st})/q_{st} \geq 0$, both terms in the denominator of the second term are at least one, so that we conclude that $S(\theta(G^{st}) \parallel \theta(G^{uv})) \leq \frac{4\lambda(\exp(2\lambda) - 1)}{\frac{q_{st}}{1 - q_{st}}}$.

Finally, applying the lower bound (29) from Lemma 1 on the ratio $q_{st}/(1 - q_{st})$, we obtain that

$$\begin{aligned} S(\theta(G^{st}) \parallel \theta(G^{uv})) &\leq \frac{4\lambda m (\exp(2\lambda) - 1)}{\exp(\omega - \lambda)} \\ &\leq \frac{8\omega \exp(2\lambda) \sinh(\lambda)}{\exp(\omega)} \end{aligned}$$

where we have used the fact that $m\lambda = \omega$. \blacksquare

By combining Lemma 3 with Lemma 2(b), we conclude that for correctness with probability at least $1/2$, the sample size n must be at least

$$n > \frac{\exp(\omega) \log\left(\frac{m}{4}\right)}{16\omega \exp(2\lambda) \sinh(\lambda)} \geq \frac{\exp(\omega) \log(k/8)}{16\omega \exp(2\lambda) \sinh(\lambda)}$$

as claimed in Theorem 2.

Ensemble C: We now establish the third term in the bound (10). As shown in the construction of Ensemble B, there is an ensemble with $M \geq \frac{\sqrt{k}}{2}$ graphs such that each edge has weight λ , and $\omega = \lambda\sqrt{k}$. Combining the bounds from Ensembles B and using the bound $k/8 \geq p^\gamma$, we conclude that the decoding will be δ -unreliable if

$$\frac{n}{(1 - \delta) \log p} \leq \frac{1}{4\lambda \tanh(\lambda)} + \frac{\gamma \exp(\frac{\lambda}{2}\sqrt{k})}{128\lambda\sqrt{k} \exp(\frac{5\lambda}{2}) \sinh(\lambda)}.$$

In fact, since we must necessarily have $\gamma \leq 2$, it suffices to impose the (slightly) stronger condition

$$\begin{aligned} \frac{n}{(1 - \delta) \log p} &\leq \gamma \left\{ \underbrace{\frac{1}{8\lambda \tanh(\lambda)}}_{F_k(\lambda)} \right. \\ &\quad \left. + \underbrace{\frac{\exp(\frac{\lambda}{2}\sqrt{k})}{128\lambda\sqrt{k} \exp(\frac{5\lambda}{2}) \sinh(\lambda)}}_{G_k(\lambda)} \right\}. \end{aligned}$$

Since $e^x \leq 1 + 4x$ for $0 \leq x \leq 2$, it can be verified that $\lambda \tanh(\lambda) \leq 4\lambda^2$ for all $\lambda \in [0, 1]$, and hence

$$F_k(\lambda) \geq \frac{1}{32\lambda^2} \quad \text{for all } \lambda \in [0, 1].$$

On the other hand, we have $\sinh(\lambda) \leq \exp(\lambda)$, and hence

$$\begin{aligned} G_k(\lambda) &\geq \frac{\exp(\frac{\lambda}{4}\sqrt{k}) \exp(\frac{\lambda}{4}\sqrt{k})}{128\lambda\sqrt{k} \exp(\frac{7\lambda}{2})} \stackrel{(i)}{\geq} \frac{\exp(\frac{\lambda}{4}\sqrt{k})}{128\lambda\sqrt{k}} \\ &\geq \frac{1}{1024} \exp\left(\frac{\lambda}{8}\sqrt{k}\right) \end{aligned}$$

where inequality (i) follows as long as $\sqrt{k} \geq 14$. Combining the pieces, we see that

$$F_k(\lambda) + G_k(\lambda) \geq \underbrace{\frac{1}{4\lambda^2} + \frac{\exp(\frac{\lambda}{8}\sqrt{k})}{1024}}_{H_k(\lambda)}.$$

Note that H_k is a strictly convex function on the positive real line, and since $H_k(\lambda) \rightarrow +\infty$ as $\lambda \rightarrow 0^+$ and $\lambda \rightarrow +\infty$, it must achieve its minimum somewhere in the open interval $(0, \infty)$. Taking derivatives to find the optimum λ^* , we see that it must satisfy the equation

$$\frac{1}{16(\lambda^*)^3} = \frac{\sqrt{k} \exp(\frac{\lambda^*}{8}\sqrt{k})}{8192}. \quad (35)$$

Substituting back, we find that $H_k(\lambda^*) = \frac{1}{32(\lambda^*)^2} + \frac{4}{2\sqrt{k}(\lambda^*)^3}$, which decreases as λ^* increases. But from (35), for any $\nu \in (0, 1)$, we have $\lambda^* = o(k^{\frac{1-\nu}{2}})$, which implies that there is a universal constant c_ν such that $H_k(\lambda^*) \geq c_\nu k^{1-\nu}$, as claimed.

V. PROOFS OF SUFFICIENT CONDITIONS

We now turn to the proofs of the sufficient conditions given in Theorems 3 and 4, respectively, for the classes $\mathcal{G}_{p,d}$ and $\mathcal{G}_{p,k}$. In both cases, our method involves a direct analysis of a maximum likelihood (ML) decoder, which searches exhaustively over all graphs in the given class, and computes the model with highest likelihood. We begin by describing this ML decoder and providing a standard large deviations bound that governs its performance. The remainder of the proof involves more delicate analysis to lower bound the error exponent in the large deviations bound in terms of the minimum edge weight λ and other structural properties of the distributions.

A. ML Decoding and Large Deviations Bound

Given a collection $\mathbf{X}_1^n = \{X^{(1)}, \dots, X^{(n)}\}$ of n i.i.d. samples, its (rescaled) likelihood with respect to model \mathbb{P}_θ is given by

$$\ell_\theta(\mathbf{X}_1^n) := \frac{1}{n} \sum_{i=1}^n \log \mathbb{P}_\theta[X^{(i)}]. \quad (36)$$

For a given graph class \mathcal{G} and an associated set of graphical models $\{\theta(G) \mid G \in \mathcal{G}\}$, the *ML decoder* is the mapping $\phi^* : \mathcal{X} \rightarrow \mathcal{G}$ defined by

$$\phi^*(\mathbf{X}_1^n) = \arg \max_{G \in \mathcal{G}} \ell_{\theta(G)}(\mathbf{X}_1^n). \quad (37)$$

If the maximum is not uniquely achieved, we choose some graph G from the set of models that attains the maximum.

Suppose that the data are drawn from model $\mathbb{P}_{\theta(G)}$ for some $G \in \mathcal{G}$. Then, the ML decoder ϕ^* fails only if there exists some other $\theta(G') \neq \theta(G)$ such that $\ell_{\theta(G')}(\mathbf{X}_1^n) \geq \ell_{\theta(G)}(\mathbf{X}_1^n)$. Note that we are being conservative by declaring failure when equality holds. Consequently, by union bound, we have

$$\mathbb{P}_\theta[\phi^*(\mathbf{X}_1^n) \neq G] \leq \sum_{G' \in \mathcal{G} \setminus G} \mathbb{P}[\ell_{\theta(G')}(\mathbf{X}_1^n) \geq \ell_{\theta(G)}(\mathbf{X}_1^n)].$$

Therefore, in order to provide sufficient conditions for the error probability of the ML decoder to vanish, we need to provide an appropriate large deviations bound.

Lemma 3: Collection a collection of n i.i.d. samples $\mathbf{X}_1^n = \{X^{(1)}, \dots, X^{(n)}\}$ from the distribution $\mathbb{P}_{\theta(G)}$. Then, for any $G' \neq G$, we have

$$\mathbb{P}[\ell_{\theta(G')}(\mathbf{X}_1^n) \geq \ell_{\theta(G)}(\mathbf{X}_1^n)] \leq e^{-\frac{n}{2} J(\theta(G) \parallel \theta(G'))} \quad (38)$$

where the distance J was defined previously (19).

Proof: So as to lighten notation, let us write $\theta = \theta(G)$ and $\theta' = \theta(G')$. We apply the Chernoff bound to the random variable $V := \ell_{\theta'}(\mathbf{X}_1^n) - \ell_{\theta}(\mathbf{X}_1^n)$, thereby obtaining that

$$\begin{aligned} & \frac{1}{n} \log \mathbb{P}_{\theta}[V \geq 0] \\ & \leq \frac{1}{n} \inf_{s>0} \log \mathbb{E}_{\theta}[\exp(sV)] \\ & = \inf_{s>0} \sum_{x \in \{-1, +1\}^p} [\mathbb{P}_{\theta}(x)]^{1-s} [\mathbb{P}_{\theta'}(x)]^s \\ & \leq \log Z(\theta/2 + \theta'/2) - \frac{1}{2} \log Z(\theta) - \frac{1}{2} \log Z(\theta'), \end{aligned}$$

where $Z(\theta)$ denotes the normalization constant associated with the Markov random field \mathbb{P}_{θ} , as defined in (3). The claim then follows by applying the representation(20) of $J(\theta \parallel \theta')$. ■

B. Lower Bounds Based on Matching

In order to exploit the large deviations claim in Lemma 3, we need to derive lower bounds on the divergence $J(\theta(G) \parallel \theta(G'))$ between different models. Intuitively, it is clear that this divergence is related to the discrepancy of the edge sets of the two graph. The following lemma makes this intuition precise. We first recall that a *matching* of a graph $G = (V, E)$ is a subgraph H such that each vertex in H has degree one. The *matching number* of G is the maximum number of edges in any matching of G .

Lemma 4: Given two distinct graphs $G = (V, E)$ and $G' = (V, E')$, let $m(G, G')$ be the matching number of the graph with edge set

$$E \Delta E' := (E \setminus E') \cup (E' \setminus E).$$

Then, for any pair of parameter vectors $\theta(G)$ and $\theta(G')$ in $\mathcal{G}(\lambda, \omega)$, we have

$$J(\theta(G) \parallel \theta(G')) \geq \frac{m(G, G')}{3 \exp(2\omega) + 1} \sinh^2 \left(\frac{\lambda}{4} \right). \quad (39)$$

Proof: Some comments on notation before proceeding: we again adopt the shorthand notation $\theta = \theta(G)$ and $\theta' = \theta(G')$. In this proof, we use e_j to denote either a particular edge, or the set of two vertices that specify the edge, depending on the context. Given any subset $A \subset V$, we use $x_A = \{x_s, s \in A\}$ to denote the collection of variables indexed by A . Given any edge $e = (u, v)$ with $u \notin A$ and $v \notin A$, we define the conditional distribution

$$\mathbb{P}_{\theta[x_A]}^e(x_u, x_v) = \frac{\mathbb{P}_{\theta}(x_u, x_v, x_A)}{\mathbb{P}_{\theta}(x_A)} \quad (40)$$

over the random variables $x_e = (x_u, x_v)$ indexed by the edge. Finally, we use

$$\begin{aligned} J_{x_A}^e(\theta \parallel \theta') &:= D \left(\frac{\mathbb{P}_{\theta[x_A] + \theta'[x_A]}^e}{2} \parallel \mathbb{P}_{\theta[x_A]}^e \right) \\ &+ D \left(\frac{\mathbb{P}_{\theta[x_A] + \theta'[x_A]}^e}{2} \parallel \mathbb{P}_{\theta'[x_A]}^e \right) \end{aligned} \quad (41)$$

to denote the divergence (19) applied to the conditional distributions of $(X_u, X_v \mid X_A = x_A)$.

With this notation, let $M \subset E \Delta E'$ be the subset of edges in some *maximal* matching of the graph with edge set $E \Delta E'$. Let $M = \{e_1, \dots, e_m\}$, and denote by $V \setminus M$ the subset of vertices that are *not* involved in the matching. Note that since J is a combination of KL divergences, the usual chain rule for KL divergences [10] also applies to it. Consequently, we have the lower bound

$$\begin{aligned} J(\theta \parallel \theta') &\geq \sum_{\ell=1}^m \sum_{\substack{x_{V \setminus M}, \\ x_{e_1}, \dots, x_{e_{\ell-1}}}} \mathbb{P}_{\frac{\theta + \theta'}{2}}(x_{S_{\ell-1}}) J_{x_{S_{\ell-1}}}^{e_{\ell}}(\theta \parallel \theta') \end{aligned}$$

where for each ℓ , we are conditioning on the set of variables $x_{S_{\ell-1}} := (x_{V \setminus M}, x_{e_1}, \dots, x_{e_{\ell-1}})$. Finally, from Lemma 7 in Appendix A, for all $\ell = 1, \dots, m$ and all values of $x_{S_{\ell-1}}$, we have

$$J_{x_{S_{\ell-1}}}^{e_{\ell}}(\theta \parallel \theta') \geq \frac{1}{3 \exp(2\omega) + 1} \sinh^2 \left(\frac{\lambda}{4} \right)$$

from which the claim follows. ■

C. Proof of Theorem 3(a)

We first consider distributions belonging to the class $\mathcal{G}_{p,d}(\lambda, \omega)$, where λ is the minimum absolute value of any nonzero edge weight, and ω is the maximum neighborhood weight (5). Consider a pair of graphs G and G' in the class $\mathcal{G}_{p,d}$ that differ in $\ell = |E \Delta E'|$ edges. Since both graphs have maximum degree at most d , we necessarily have a matching number

$$m(G, G') \geq \frac{\ell}{4d}.$$

Note that the parameter $\ell = |E \Delta E'|$ can range from 1 all the way up to dp , since a graph with maximum degree d has at most $\frac{dp}{2}$ edges.

Now consider some fixed graph $G \in \mathcal{G}_{p,d}$ and associated distribution $\mathbb{P}_{\theta(G)} \in \mathcal{G}_{p,d}$; we upper bound the error probability $\mathbb{P}_{\theta(G)}[\phi^*(\mathbf{X}_1^n) \neq G]$. For each $\ell = 1, 2, \dots, dp$, there are at most $\binom{p}{\ell}$ models in $\mathcal{G}_{p,d}$ with mismatch ℓ from G . Therefore, applying the union bound, the large deviations bound in Lemma

3, and the lower bound in terms of matching from Lemma 4, we obtain

$$\begin{aligned}
& \mathbb{P}_{\theta(G)}[\phi^*(\mathbf{X}_1^n) \neq G] \\
& \leq \sum_{\ell=1}^{pd} \binom{\binom{p}{2}}{\ell} \exp \left\{ -n \frac{\ell/(4d)}{3 \exp(2\omega) + 1} \sinh^2\left(\frac{\lambda}{2}\right) \right\} \\
& \leq pd \max_{\ell=1, \dots, pd} \exp \left\{ \log \binom{\binom{p}{2}}{\ell} \right. \\
& \quad \left. - n \frac{\ell/(4d)}{3 \exp(2\omega) + 1} \sinh^2\left(\frac{\lambda}{2}\right) \right\} \\
& \leq \max_{\ell=1, \dots, pd} \exp \left\{ \log(pd) + \ell \log p \right. \\
& \quad \left. - n \frac{\ell/(4d)}{3 \exp(2\omega) + 1} \sinh^2\left(\frac{\lambda}{2}\right) \right\}.
\end{aligned}$$

This probability is at most δ under the given conditions on n in the statement of Theorem 3(a).

D. Proof of Theorem 4

Next, we consider the class $\mathcal{G}_{p,k}$ of graphs with at most k edges. Given some fixed graph $G \in \mathcal{G}_{p,k}$, consider some other graph $G' \in \mathcal{G}_{p,k}$ such that the set $E \Delta E'$ has cardinality m . We claim that for each $m = 1, 2, \dots, 2k$, the number of such graphs is at most $2^k p^{2m(k+1)}$.

To see this claim, recall the notion of a *vertex cover* of a set of edges, namely a subset of vertices such that each edge in the set is incident to at least one vertex of the set. Note also that the set of vertices involved in any maximal matching form a vertex cover. Consequently, any graph E' such that $E \Delta E'$ has cardinality m be described in the following (suboptimal) fashion: 1) first specify which of the k edges in E are missing in E' ; 2) describe which of the at most $2m$ vertices belong to the vertex cover corresponding to the maximal matching of $E \Delta E'$; and 3) describe the subset of at most k vertices that are connected to the vertices of the cover in Step 2). This procedure yields at most $2^k p^{2m} p^{2km} = 2^k p^{2m(k+1)}$ possibilities, as claimed.

Consequently, applying the union bound, the large deviations bound in Lemma 3, and the lower bound in terms of matching from Lemma 4, we obtain

$$\begin{aligned}
& \mathbb{P}_{\theta(G)}[\phi^*(\mathbf{X}_1^n) \neq G] \\
& \leq \sum_{m=1}^k 2^k p^{2m(k+1)} e^{-n \frac{m}{3 \exp(2\omega) + 1} \sinh^2\left(\frac{\lambda}{2}\right)} \\
& \leq k \max_{m=1, \dots, k} e^{k+2m(k+1) \log p - n \frac{m}{3 \exp(2\omega) + 1} \sinh^2\left(\frac{\lambda}{2}\right)}.
\end{aligned}$$

This probability is less than δ under the conditions of Theorem 4, which completes the proof.

E. Proof of Theorems 3(b) and 4(b)

Finally, we prove the sufficient conditions given in Theorems 3(b) and 4(b), which do not assume that the decoder knows the parameter vector $\theta(G)$ for each graph $G \in \mathcal{G}_{p,d}$. In this case, the simple ML decoder (37) cannot be applied, since it assumes knowledge of the model parameters $\theta(G)$ for each graph

$G \in \mathcal{G}_{p,d}$. A natural alternative would be the generalized likelihood ratio approach, which would maximize the likelihood over each model class, and then compare the maximized likelihoods. Our proof of Theorem 3(b) is based on minimizing the distance between the empirical and model mean parameters in the ℓ_∞ norm, which is easier to analyze.

1) *Decoding From Mean Parameters:* We begin by describing the graph decoder used to establish the sufficient conditions of Theorem 3(b). For any parameter vector $\theta \in \mathbb{R}^{\binom{p}{2}}$, let $\mu(\theta) \in \mathbb{R}^{\binom{p}{2}}$ represent the associated set of mean parameters, with element (s, t) given by $[\mu(\theta)]_{st} := \mathbb{E}_\theta[X_s X_t]$. Given a set $\mathbf{X}_1^n = \{X^{(1)}, \dots, X^{(n)}\}$ of n samples, the *empirical mean parameters* are given by

$$\hat{\mu}_{st} := \frac{1}{n} \sum_{i=1}^n X^{(i)}_s X^{(i)}_t. \quad (42)$$

For a given graph $G = (V, E)$, let $\Theta_{\lambda, \omega}(G) \subset \mathbb{R}^{\binom{p}{2}}$ be a subset of exponential parameters that respect the graph structure—viz.

- a) we have $\theta_{uv} = 0$ for all $(u, v) \notin E$;
- b) for all edges $(s, t) \in E$, we have $|\theta_{st}| \geq \lambda$;
- c) for all vertices $s \in V$, we have $\sum_{t \in \mathcal{N}(s)} |\theta_{st}| \leq \omega$.

For any graph G and set of mean parameters $\mu \in \mathbb{R}^{\binom{p}{2}}$, we define a projection-type distance via $J_G(\mu) := \min_{\theta \in \Theta_{\lambda, \omega}(G)} \|\mu - \mu(\theta)\|_\infty$.

We now have the necessary ingredients to define a graph decoder $\phi^\dagger : \mathcal{X}^n \rightarrow \mathcal{G}_{p,d}$; in particular, it is given by

$$\phi^\dagger(\mathbf{X}_1^n) := \arg \min_{G \in \mathcal{G}_{p,d}} J_G(\hat{\mu}) \quad (43)$$

where $\hat{\mu}$ are the empirical mean parameters previously defined (42). (If the minimum (43) is not uniquely achieved, then we choose some graph that achieves the minimum.)

2) *Analysis of Decoder:* Suppose that the data are sampled from $\mathbb{P}_{\theta(G)}$ for some fixed but known graph $G \in \mathcal{G}_{p,d}$, and parameter vector $\theta(G) \in \Theta_{\lambda, \omega}(G)$. Note that the graph decoder ϕ^\dagger can fail only if there exists some other graph G' such that the difference $\Delta(G'; G) := J_{G'}(\hat{\mu}) - J_G(\hat{\mu})$ is not positive. Again, we are conservative in declaring failure if there are ties.

Let θ' denote some element of $\Theta_{\lambda, \omega}(G')$ that achieves the minimum defining $J_{G'}(\hat{\mu})$, so that $J_{G'}(\hat{\mu}) = \|\hat{\mu} - \mu(\theta')\|_\infty$. Note that by the definition of J_G , we have $J_{G'}(\hat{\mu}) \leq \|\hat{\mu} - \mu(\theta(G))\|_\infty$, where $\theta(G)$ are the parameters of the true model. Therefore, by the definition of $\Delta(G'; G)$, we have

$$\begin{aligned}
\Delta(G'; G) & \geq \|\hat{\mu} - \mu(\theta')\|_\infty - \|\hat{\mu} - \mu(\theta(G))\|_\infty \\
& \geq \|\mu(\theta') - \mu(\theta(G))\|_\infty - 2\|\hat{\mu} - \mu(\theta(G))\|_\infty
\end{aligned} \quad (44)$$

where the second inequality above uses the triangle inequality.

Therefore, in order to prove that $\Delta(G'; G)$ is positive, it suffices to obtain an upper bound on $\|\hat{\mu} - \mu(\theta(G))\|_\infty$, and a lower bound on $\|\mu(\theta') - \mu(\theta(G))\|_\infty$, where θ' ranges over $\Theta_{\lambda, \omega}(G')$. With this perspective, let us state two key lemmas. We begin with the deviation between the sample and population mean parameters:

Lemma 5 (Elementwise Deviation): Given n i.i.d. samples drawn from $\mathbb{P}_{\theta(G)}$, the sample mean parameters $\hat{\mu}$ and population mean parameters $\mu(\theta(G))$ satisfy the tail bound

$$\mathbb{P}[\|\hat{\mu} - \mu(\theta(G))\|_{\infty} \geq t] \leq 2 \exp \left(-n \frac{t^2}{2} + 2 \log p \right).$$

This probability is less than δ for $t \geq \sqrt{\frac{4 \log p + \log(2/\delta)}{n}}$.

Our second lemma concerns the separation of the mean parameters of models with different graph structure:

Lemma 6 (Pairwise Separations): Consider any two graphs $G = (V, E)$ and $G' = (V, E')$, and an associated set of model parameters $\theta(G) \in \Theta_{\lambda, \omega}(G)$ and $\theta(G') \in \Theta_{\lambda, \omega}(G')$. Then, for all edges $(s, t) \in E \Delta E'$

$$\begin{aligned} \max_{u \in \{s, t\}, v \in V} |\mathbb{E}_{\theta(G)}[X_u X_v] - \mathbb{E}_{\theta(G')}[X_u X_v]| \\ \geq \frac{\sinh^2(\lambda/4)}{2\omega(3 \exp(2\omega) + 1)}. \end{aligned}$$

We provide the proofs of these two lemmas in Sections V-E3 and V-E4.

Given these two lemmas, we can complete the proofs of Theorems 3(b) and 4(b). Using the lower bound (44), with probability greater than $1 - \delta$, we have

$$\Delta(G'; G) \geq \frac{\sinh^2(\lambda/4)}{\omega(3 \exp(2\omega) + 1)} - 2 \sqrt{\frac{4 \log p + \log(2/\delta)}{n}}.$$

This quantity is positive as long as

$$n > \left[\frac{\omega(3 \exp(2\omega) + 1)}{\sinh^2(\lambda/4)} \right]^2 \{16 \log p + 4 \log(2/\delta)\}$$

which completes the proof.

It remains to prove the auxiliary lemmas used in the proof.

3) *Proof of Lemma 5:* This claim is an elementary consequence of the Hoeffding bound. By definition, for each pair (s, t) of distinct vertices, we have

$$\hat{\mu}_{st} - [\mu(\theta(G))]_{st} = \frac{1}{n} \sum_{i=1}^n X_s^{(i)} X_t^{(i)} - \mathbb{E}_{\theta(G)}[X_s X_t]$$

which is the deviation of a sample mean from its expectation. Since the random variables $\{X_s^{(i)} X_t^{(i)}\}_{i=1}^n$ are i.i.d. and lie in the interval $[-1, +1]$, an application of Hoeffding's inequality [16] yields that

$$\mathbb{P}[|\hat{\mu}_{st} - [\mu(\theta(G))]_{st}| \geq t] \leq 2 \exp(-nt^2/2).$$

The lemma follows by applying union bound over all $\binom{p}{2}$ edges of the graph, and the fact that $\log \binom{p}{2} \leq 2 \log p$.

4) *Proof of Lemma 6:* The proof of this claim is more involved. Let (s, t) be an edge in $E \setminus E'$, and let C be the set of all other vertices that are adjacent to s or t in either graphs—namely, the set

$$\begin{aligned} C &:= \{u \in V \mid (u, s) \text{ or } (u, t) \in E \cup E'\} \\ &= (\mathcal{N}(s) \cup \mathcal{N}(t)) \setminus \{s, t\}. \end{aligned}$$

Our approach is to first condition on the variables $x_C = \{x_u, u \in C\}$, and then to consider the two conditional distributions over the pair (X_s, X_t) , defined by \mathbb{P}_{θ} and $\mathbb{P}_{\theta'}$,

respectively. In particular, for any subset $S \subset V$, let us define the unnormalized distribution

$$\mathbb{Q}_{\theta}(x_S) := \sum_{x_a, a \notin S} \exp \left(\sum_{(u,v) \in E} \theta_{uv} x_u x_v \right) \quad (45)$$

obtained by summing out all variables x_a for $a \notin S$. With this notation, we can write the conditional distribution of (X_s, X_t) given $\{x_C = x_C\}$ as

$$\mathbb{P}_{\theta[x_C]}(x_s, x_t) = \frac{\mathbb{Q}_{\theta}(x_s, x_t, x_C)}{\mathbb{Q}_{\theta}(x_C)}. \quad (46)$$

As reflected in our choice of notation, for each fixed x_C , the distribution (40) can be viewed as a Ising model over the pair (X_s, X_t) with exponential parameter $\theta[x_C]$. We define the unnormalized distributions $\mathbb{Q}_{\theta'[x_S]}$ and the conditional distributions $\mathbb{P}_{\theta'[x_C]}$ in an analogous manner.

Our approach now is to study the divergence $J(\theta[x_C] \parallel \theta'[x_C])$ between the conditional distributions induced by \mathbb{P}_{θ} and $\mathbb{P}_{\theta'}$. Using Lemma 7 from Appendix A, for each choice of x_C , we have $J(\theta[x_C] \parallel \theta'[x_C]) \geq \frac{\sinh^2(\frac{\lambda}{4})}{3 \exp(2\omega) + 1}$, and hence

$$\mathbb{E}_{\frac{\theta+\theta'}{2}} [J(\theta[x_C] \parallel \theta'[x_C])] \geq \frac{1}{3 \exp(2\omega) + 1} \sinh^2 \left(\frac{\lambda}{4} \right) \quad (47)$$

where the expectation is taken under the model $\mathbb{P}_{\frac{\theta+\theta'}{2}}$. Some calculation shows that

$$\begin{aligned} \mathbb{E}_{\frac{\theta+\theta'}{2}} [J(\theta[x_C] \parallel \theta'[x_C])] \\ = \mathbb{E}_{\frac{\theta+\theta'}{2}} \left[\log \frac{\mathbb{Q}_{\theta}(x_C)}{\mathbb{Q}_{\frac{\theta+\theta'}{2}}(x_C)} \right] + \mathbb{E}_{\frac{\theta+\theta'}{2}} \left[\log \frac{\mathbb{Q}_{\theta'}(x_C)}{\mathbb{Q}_{\frac{\theta+\theta'}{2}}(x_C)} \right]. \end{aligned}$$

Applying Jensen's inequality yields

$$\begin{aligned} \mathbb{E}_{\frac{\theta+\theta'}{2}} \left[\log \frac{\mathbb{Q}_{\theta}(x_C)}{\mathbb{Q}_{\frac{\theta+\theta'}{2}}(x_C)} \right] \\ \leq \log \left[\frac{1}{\sum_{x_C} \mathbb{Q}_{\frac{\theta+\theta'}{2}}(x_C)} \sum_{x_C} \mathbb{Q}_{\frac{\theta+\theta'}{2}}(x_C) \frac{\mathbb{Q}_{\theta}(x_C)}{\mathbb{Q}_{\frac{\theta+\theta'}{2}}(x_C)} \right] \\ = \log \frac{\sum_{x_C} \mathbb{Q}_{\theta}(x_C)}{\sum_{x_C} \mathbb{Q}_{\frac{\theta+\theta'}{2}}(x_C)} \end{aligned}$$

with an analogous inequality for the term involving $\mathbb{Q}_{\theta'}$. Consequently, the average $\mathbb{E}_{\frac{\theta+\theta'}{2}} [J(\theta[x_C] \parallel \theta'[x_C])]$ is upper bounded by

$$\log \frac{[\sum_{x_C} \mathbb{Q}_{\theta}(x_C)] [\sum_{x_C} \mathbb{Q}_{\theta'}(x_C)]}{\left[\sum_{x_C} \mathbb{Q}_{\frac{\theta+\theta'}{2}}(x_C) \right]^2}. \quad (48)$$

In order to exploit (48), we upper bound the quantity

$$\begin{aligned} \Delta(\theta, \theta') &:= \mathbb{E}_{\theta} \left[D \left(\theta[x_C] \parallel \left(\frac{\theta + \theta'}{2} \right) [x_C] \right) \right] \\ &\quad + \mathbb{E}_{\theta'} \left[D \left(\theta'[x_C] \parallel \left(\frac{\theta + \theta'}{2} \right) [x_C] \right) \right]. \end{aligned} \quad (49)$$

By the definition of the KL divergence, we can decompose $\Delta(\theta, \theta')$ as the sum $\sum_{j=1}^3 T_j$, where

$$T_1 := \sum_{u \in \mathcal{N}(s) \setminus t} (\mu_{su} - \mu'_{su}) (\theta_{su} - \theta'_{su}) \quad (50a)$$

$$T_2 := \sum_{v \in \mathcal{N}(t) \setminus s} (\mu_{tv} - \mu'_{tv}) (\theta_{tv} - \theta'_{tv}) \quad \text{and} \quad (50b)$$

$$T_3 := \mathbb{E}_\theta \log \frac{\mathbb{Q}_{\frac{\theta+\theta'}{2}}(X_C)}{\mathbb{Q}_\theta(X_C)} + \mathbb{E}_{\theta'} \log \frac{\mathbb{Q}_{\frac{\theta+\theta'}{2}}(X_C)}{\mathbb{Q}_{\theta'}(X_C)}. \quad (50c)$$

In the expressions for T_1 and T_2 , the quantities μ and μ' denote mean parameters computed under the distributions \mathbb{P}_θ and $\mathbb{P}_{\theta'}$, respectively. But by Jensen's inequality, we have the upper bound

$$\mathbb{E}_\theta \log \frac{\mathbb{Q}_{\frac{\theta+\theta'}{2}}(X_C)}{\mathbb{Q}_\theta(X_C)} \leq \log \frac{\sum_{x_C} \mathbb{Q}_{\frac{\theta+\theta'}{2}}(x_C)}{\sum_{x_C} \mathbb{Q}_\theta(x_C)} \quad (51)$$

with an analogous upper bound for the term involving θ' .

Combining the initial bound (48), the decomposition (50) and the last bound (51), we obtain

$$\begin{aligned} \mathbb{E}_{\frac{\theta+\theta'}{2}} [J(\theta[X_C] \parallel \theta'[X_C])] &\leq \\ &\sum_{u \in \mathcal{N}(s) \setminus t} (\mu_{su} - \mu'_{su}) (\theta_{su} - \theta'_{su}) \\ &+ \sum_{v \in \mathcal{N}(t) \setminus s} (\mu_{tv} - \mu'_{tv}) (\theta_{tv} - \theta'_{tv}). \end{aligned}$$

Finally, since $\sum_{u \in \mathcal{N}(s)} |\theta_{us}| \leq \omega$ by the definition (5) (and similarly for the neighborhood of t), we conclude that

$$\mathbb{E}_{\frac{\theta+\theta'}{2}} [J(\theta[X_C] \parallel \theta'[X_C])] \leq 2\omega \max_{u \in \{s,t\}, v \in V} |\mu_{uv} - \mu'_{uv}|.$$

Combining this upper bound with the lower bound (47) yields the claim.

VI. DISCUSSION

In this paper, we have analyzed the information-theoretic limits of binary graphical model selection in a high-dimensional framework, in which the sample size n , number of graph vertices p , number of edges k , and/or the maximum vertex degree d are allowed to grow to infinity. We proved four main results, corresponding to both necessary and sufficient conditions for inference in the class $\mathcal{G}_{p,d}$ of graphs on p vertices with maximum vertex degree d , as well as for the class $\mathcal{G}_{p,k}$ of graphs on p vertices with at most k edges. More specifically, for the class $\mathcal{G}_{p,d}$, we showed that any algorithm requires at least $n > c d^2 \log p$ samples, and we demonstrated an algorithm that succeeds using $n < c' d^3 \log p$ samples. Our two main results for the class $\mathcal{G}_{p,k}$ have a similar flavor: we show that any algorithm requires at least $n > c k \log p$ samples, and we demonstrated an algorithm that succeeds using $n < c' k^2 \log p$ samples. Thus, for graphs with constant degree d or a constant number of edges k , our bounds provide a characterization of the information-theoretic complexity of binary graphical selection

that is tight up to constant factors. For growing degrees or edge numbers, there remains a minor gap in our conditions.

In terms of open questions, one immediate issue is to close the current gap between our necessary and sufficient conditions; as summarized previously, these gaps are of order d and k for $\mathcal{G}_{p,d}$ and $\mathcal{G}_{p,k}$, respectively. We note that previous work [22] has shown that a computationally tractable method, based on ℓ_1 -regularization and logistic regression, can recover binary graphical models using $n = \Omega(d^3 \log p)$ samples, a result that is consistent with but not quite matching the $n = \Omega(d^2 \log p)$ lower bound given here. Moreover, in this paper, although we have focused exclusively on binary graphical models with pairwise interactions, many of the techniques and results (e.g., constructing ‘‘packings’’ of graph classes, Fano's lemma and variants, large deviations analysis) apply to more general classes of graphical models. Some other follow-up work [29] has used related techniques for studying the information-theoretic limits of Gaussian graphical model selection. It would be interesting to explore other extensions in this direction, including the case of discrete models with larger alphabets.

APPENDIX

A) Separation Lemma: In this appendix, we prove the following lemma, which plays a key role in the proofs of both Lemmas 4 and 6. Given an edge $e = (s, t)$ and some subset $U \subseteq V \setminus \{s, t\}$, recall that $J_{x_U}^e(\theta \parallel \theta')$ denotes the divergence (19) applied to the conditional distributions of $(X_u, X_v \mid X_A = x_A)$, as defined explicitly in (41).

Lemma 7: Consider two distinct graphs $G = (V, E)$ and $G' = (V, E')$, with associated parameter vectors θ and θ' . Given an edge $(s, t) \in E \setminus E'$ and any subset $U \subseteq V \setminus \{s, t\}$, we have

$$J_{x_U}^e(\theta \parallel \theta') \geq \frac{1}{3 \exp(2\omega) + 1} \sinh^2 \left(\frac{\theta_{uv}}{4} \right). \quad (52)$$

Proof: To lighten notation, let us define

$$\epsilon := \frac{2}{3 \exp(2\omega) + 1} \sinh^2 \left(\frac{\theta_{uv}}{4} \right) > 0.$$

Note that from the definition (5), we have $\omega \geq |\theta_{uv}|$, which implies that $\epsilon \leq 2$. For future reference, we also note the relation

$$\left[\exp \left(\frac{\theta_{uv}}{4} \right) - \exp \left(-\frac{\theta_{uv}}{4} \right) \right]^2 = 2\epsilon + 6\epsilon \exp(2\omega). \quad (53)$$

With this setup, our argument proceeds via proof by contradiction. In particular, we assume that

$$J_{x_U}^e(\theta \parallel \theta') \leq \epsilon/2 \quad (54)$$

and then derive a contradiction. Recall from (45), our notation $\mathbb{Q}_\theta(x_A)$ for the unnormalized distribution applied to the subset of variables $x_A = \{x_i, i \in A\}$. With a little bit of algebra, we find that

$$J_{x_U}^e(\theta \parallel \theta') = \log \frac{\mathbb{Q}_\theta(x_U) \mathbb{Q}_{\theta'}(x_U)}{\left(\sum_{\substack{z=(z_1, \dots, z_p) \\ z_{\mathcal{S}}=x_U}} \sqrt{\mathbb{Q}_\theta(z) \mathbb{Q}_{\theta'}(z)} \right)^2}.$$

Let us introduce some additional shorthand so as to lighten notation in the remainder of the proof. First we define $\beta(x) := \sqrt{\mathbb{Q}_\theta(x) \mathbb{Q}_{\theta'}(x)}$ and $\alpha(x) := \sqrt{\frac{\mathbb{Q}_\theta(x)}{\mathbb{Q}_{\theta'}(x)}}$. We make note of the equivalence $\alpha(x) = \exp(\Delta(x)/2)$, where

$$\Delta(x) := \sum_{(s,t) \in E} \theta_{st} x_s x_t - \sum_{(s',t') \in E'} \theta'_{s't'} x_{s'} x_{t'}.$$

Observe that Lemma 8 in Appendix B characterizes the behavior of $\Delta(x)$ under changes to x . Finally, we define the set

$$\mathcal{Y}(x_U) := \{y \in \{-1, +1\}^p \mid y_i = x_i \text{ for all } i \in U\},$$

corresponding to the subset of configurations $y \in \{-1, +1\}^p$ that agree with x_U over the subset U .

From the definitions of α and β , we observe that

$$\begin{aligned} & \left[\sum_{y \in \mathcal{Y}(x_U)} \mathbb{Q}_\theta(y) \right] \left[\sum_{y \in \mathcal{Y}(x_U)} \mathbb{Q}_{\theta'}(y) \right] \\ &= \left[\sum_{y \in \mathcal{Y}(x_U)} \alpha(y) \beta(y) \right] \left[\sum_{y \in \mathcal{Y}(x_U)} \frac{\alpha(y)}{\beta(y)} \right] \\ &\leq (1 + \epsilon)^2 \left[\sum_{y \in \mathcal{Y}(x_U)} \beta(y) \right]^2 \end{aligned} \quad (55)$$

where the inequality follows from the fact $\epsilon \leq 2$, our original assumption (54), and the elementary relations for all $z \in (0, 1]$

$$\exp(z) < (1 + 2z) < (1 + 2z)^2.$$

Now consider the set of quadratics in t , one for each $y \in \mathcal{Y}(x_U)$, given by

$$\beta(y)\alpha(y) - 2(1 + \epsilon)\beta(y)t + \frac{\beta(y)}{\alpha(y)}t^2 = 0.$$

Summing these quadratic equations over $y \in \mathcal{Y}(x_U)$ yields a new quadratic q given by

$$\begin{aligned} q(t) := & \sum_{y \in \mathcal{Y}(x_U)} \beta(y)\alpha(y) - 2t(1 + \epsilon) \sum_{y \in \mathcal{Y}(x_U)} \beta(y) \\ & + t^2 \sum_{y \in \mathcal{Y}(x_U)} \frac{\beta(y)}{\alpha(y)} \end{aligned}$$

which by which by (55) must have two real roots.

Let t^* denote the value of t at which the quadratic q achieves its minimum. By the quadratic formula, we have

$$t^* = \frac{(1 + \epsilon) \sum_{y \in \mathcal{Y}(x_U)} \beta(y)}{\sum_{y \in \mathcal{Y}(x_U)} (\beta(y)/\alpha(y))} > 0.$$

Since $q(t^*) < 0$, we obtain

$$2\epsilon t^* > \frac{\sum_{y \in \mathcal{Y}(x_U)} \beta(y) t^* \left[\sqrt{t^*/\alpha(y)} - \sqrt{\alpha(y)/t^*} \right]^2}{\sum_{y \in \mathcal{Y}(x_U)} \beta(y)}. \quad (56)$$

If we define the set $\mathcal{Y}^*(x_U)$ as

$$\left\{ y \in \mathcal{Y}(x_U) \mid \max \left\{ \frac{t^*}{\alpha(y)}, \frac{\alpha(y)}{t^*} \right\} < \exp\left(\frac{\theta_{uv}}{2}\right) \right\}$$

then we can rewrite (56) as

$$\begin{aligned} 2\epsilon t^* &> \frac{\sum_{y \in \mathcal{Y}^*(x_U)} \beta(y) t^* \left\{ \left[\sqrt{\frac{t^*}{\alpha(y)}} - \sqrt{\frac{\alpha(y)}{t^*}} \right]^2 - 2\epsilon \right\}}{\sum_{y \in \mathcal{Y}^*(x_U)} \beta(y)} \\ &\stackrel{(a)}{\geq} 2\epsilon t^* \frac{\sum_{y \notin \mathcal{Y}^*(x_U)} 3\beta(y) \exp(2\omega)}{\sum_{y \in \mathcal{Y}^*(x_U)} \beta(y)} \end{aligned} \quad (57)$$

where inequality (a) follows from the definition of $\mathcal{Y}^*(x_U)$, the monotonically increasing nature of the function $f(s) = (s - \frac{1}{s})^2$ for $s \geq 1$, and the relation (53).

From Lemma 8, for each $y \in \mathcal{Y}^*(x_U)$, we obtain a configuration $a \notin \mathcal{Y}^*(x_U)$ by flipping either u, v or both. Note that at most three configurations $y \in \mathcal{Y}^*(x_U)$ can yield the same configuration $z \notin \mathcal{Y}^*(x_U)$. Since these flips do not decrease $\beta(y)$ by more than a factor of $\exp(2\omega)$, we conclude that

$$\sum_{y \in \mathcal{Y}^*(x_U)} \beta(y) \leq 3 \exp(2\omega) \sum_{y \notin \mathcal{Y}^*(x_U)} \beta(y)$$

which is a contradiction of (57). Hence, the quadratic q cannot have two real roots, which contradicts our initial assumption (54). ■

B) Proof of a Flipping Lemma: It remains to state and prove a lemma that we exploited in the proof of Lemma 7 from Appendix A.

Lemma 8: Consider distinct models θ and θ' , and for each $x \in \{-1, +1\}^p$, define

$$\Delta(x) := \sum_{(u,v) \in E} \theta_{uv} x_u x_v - \sum_{(u',v') \in E'} \theta'_{u'v'} x_{u'} x_{v'}. \quad (58)$$

Then, for any edge $(s, t) \in E \setminus E'$ and for any configuration $x \in \{-1, +1\}^p$, flipping either x_s or x_t (or both) changes $\Delta(x)$ by at least $|\theta_{st}|$.

Proof: We use $\mathcal{N}(s)$ and $\mathcal{N}'(s)$ to denote the neighborhood sets of s in the graphs $G = (V, E)$ and $G' = (V, E')$ respectively, with analogous notation for the sets $\mathcal{N}(t)$ and $\mathcal{N}'(t)$. We then define

$$\begin{aligned} \delta_s(x) &:= \sum_{u \in \mathcal{N}(s) \setminus \mathcal{N}'(s)} \theta_{su} x_u, \quad \text{and} \\ \delta'_s(x) &:= \sum_{u \in \mathcal{N}'(s) \setminus \mathcal{N}(s)} \theta'_{su} x_u \end{aligned}$$

with analogous definitions for the quantities $\delta_t(x)$ and $\delta'_t(x)$. Similarly, we define

$$\begin{aligned}\gamma_s(x) &:= \sum_{u \in \mathcal{N}(s) \cap \mathcal{N}'(s)} \theta_{su} x_u, \quad \text{and} \\ \gamma_t(x) &:= \sum_{v \in \mathcal{N}(t) \cap \mathcal{N}'(t)} \theta_{tv} x_v.\end{aligned}$$

Finally, we define the pair

$$\begin{aligned}\xi(x) &:= \sum_{\substack{(u,v) \in E_i \\ u,v \notin \{s,t\}}} \theta_{uv} x_u x_v, \quad \text{and} \\ \xi'(x) &:= \sum_{\substack{(u,v) \in E' \\ u,v \notin \{s,t\}}} \theta_{uv} x_u x_v.\end{aligned}$$

With this notation, we first proceed via proof by contradiction to show that $\Delta(x)$ must change when (x_s, x_t) are flipped: to the contrary, suppose that for $\Delta(x)$ stays fixed for all four choices $(x_s, x_t) \in \{-1, +1\}^2$. We then show that this assumption implies that $\theta_{st} = 0$. Note that both of the terms $\delta_s(x)$ and $\delta_t(x)$ include a contribution from the edge (s, t) . When $(x_s, x_t) = (+1, +1)$, we have

$$\begin{aligned}(\delta_s(x) - \theta_{st}) + (\delta_t(x) - \theta_{st}) \\ + \theta_{st} + \xi(x) + \gamma_s(x) + \gamma_t(x) \\ = \delta'_s(x) + \delta'_t(x) + \xi'(x) + \gamma_s(x) + \gamma_t(x) + \Delta(x)\end{aligned}$$

whereas when $(x_s, x_t) = (-1, -1)$, we have

$$\begin{aligned}-(\delta_s(x) + \theta_{st}) - (\delta_t(x) + \theta_{st}) \\ + \theta_{st} + \xi(x) - \gamma_s(x) - \gamma_t(x) \\ = -\delta'_s(x) - \delta'_t(x) + \xi'(x) - \gamma_s(x) - \gamma_t(x) + \Delta(x).\end{aligned}$$

Adding these two equations together yields

$$\xi(x) - \theta_{st} = \xi'(x) + \Delta(x). \quad (59)$$

On the other hand, for $(x_s, x_t) = (-1, +1)$, we have

$$\begin{aligned}-(\delta_s(x) - \theta_{st}) + (\delta_t(x) + \theta_{st}) \\ - \theta_{st} + \xi(x) - \gamma_s(x) + \gamma_t(x) \\ = -\delta'_s(x) + \delta'_t(x) + \xi'(x) - \gamma_s(x) + \gamma_t(x) + \Delta(x)\end{aligned}$$

and for $(x_s, x_t) = (+1, -1)$, we have

$$\begin{aligned}(\delta_s(x) + \theta_{st}) - (\delta_t(x) - \theta_{st}) \\ - \theta_{st} + \xi(x) + \gamma_s(x) - \gamma_t(x) \\ = \delta'_s(x) - \delta'_t(x) + \xi'(x) + \gamma_s(x) - \gamma_t(x) + \Delta(x).\end{aligned}$$

Adding together these two equations yields

$$\xi(x) + \theta_{st} = \xi'(x) + \Delta(x). \quad (60)$$

Note that (59) and (60) cannot hold simultaneously unless $\theta_{st} = 0$, which implies that our initial assumption—namely, that $\Delta(x)$ does not change as we vary $(x_s, x_t) \in \{-1, +1\}^2$ —was false.

Finally, we show that the change in $|\Delta(x)|$ must be at least $|\theta_{st}|$. For each pair $(i, j) \in \{-1, +1\}^2$, let $\mathcal{E}_{ij} = \Delta(x \mid x_s = i, x_t = j)$ be the value of $\Delta(x)$ when $x_s = i$ and $x_t = j$. Suppose that for some constant c and $\delta > 0$, we have $\mathcal{E}_{ij} \in [c - \epsilon, c + \epsilon]$ for all (i, j) . By following the same reasoning as previously, we obtain the inequalities $\xi(x) - \theta_{st} \geq \xi'(x) + c - \epsilon$ and $\xi(x) + \theta_{st} \leq \xi'(x) + c + \epsilon$, which together imply that $\theta_{st} \leq \epsilon$. In a similar manner, we obtain the inequalities $\xi(x) + \theta_{st} \geq \xi'(x) + c - \epsilon$ and $\xi(x) - \theta_{st} \leq \xi'(x) + c + \epsilon$, which imply that $-\theta_{st} \leq \epsilon$, thereby completing the proof. ■

ACKNOWLEDGMENT

We thank the anonymous reviewers for their helpful comments on an earlier version of the paper.

REFERENCES

- [1] A. Ahmady, L. Song, and E. P. Xing, Time-varying networks: Recovering temporally rewiring genetic networks during the life cycle of drosophila melanogaster Carnegie Mellon Univ., Pittsburgh, PA, Tech. Rep. CMU-ML-08-118, 2008.
- [2] N. Alon and J. Spencer, *The Probabilistic Method*. New York: Wiley-Interscience, 2000.
- [3] O. Bannerjee, L. El Ghaoui, and A. d'Aspremont, "Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data," *J. Mach. Learning Res.*, vol. 9, pp. 485–516, Mar. 2008.
- [4] R. J. Baxter, *Exactly Solved Models in Statistical Mechanics*. New York: Academic, 1982.
- [5] J. Besag, "On the statistical analysis of dirty pictures," *J. Royal Statist. Soc., Series B*, vol. 48, no. 3, pp. 259–279, 1986.
- [6] G. Bresler, E. Mossel, and A. Sly, "Reconstruction of Markov random fields from samples: Some easy observations and algorithms," in *RANDOM*, 2008, pp. 343–356.
- [7] L. D. Brown, *Fundamentals of Statistical Exponential Families*. Hayward, CA: Inst. Math. Statist., 1986.
- [8] D. Chickering, "Learning Bayesian networks is NP-complete," in *Learning from data: AI and Statist. V*, 1996, pp. 121–130.
- [9] C. K. Chow and C. N. Liu, "Approximating discrete probability distributions with dependence trees," *IEEE Trans. Inf. Theory*, vol. IT-14, no. 3, pp. 462–467, May 1968.
- [10] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.
- [11] I. Csiszár and Z. Talata, "Consistent estimation of the basic neighborhood structure of Markov random fields," *Ann. Statist.*, vol. 34, no. 1, pp. 123–145, 2006.
- [12] R. Durbin, S. Eddy, A. Krogh, and G. Mitchison, *Biological Sequence Analysis*. Cambridge, U.K.: Cambridge Univ. Press, 1998.
- [13] J. Friedman, T. Hastie, and R. Tibshirani, "Sparse inverse covariance estimation with the graphical Lasso," *Biostatistics*, vol. 9, pp. 432–441, 2007.
- [14] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-6, no. 6, pp. 721–741, Nov. 1984.
- [15] R. Z. Has'minskii, "A lower bound on the risks of nonparametric estimates of densities in the uniform metric," *Theory Prob. Appl.*, vol. 23, pp. 794–798, 1978.
- [16] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *J. Amer. Statist. Assoc.*, vol. 58, pp. 13–30, 1963.
- [17] I. A. Ibragimov and R. Z. Has'minskii, *Statistical Estimation: Asymptotic Theory*. New York: Springer-Verlag, 1981.
- [18] E. Ising, "Beitrag zur theorie der ferromagnetismus," *Zeitschrift für Physik*, vol. 31, no. 1, pp. 253–258, 1925.
- [19] C. Ji and L. Seymour, "A consistent model selection procedure for Markov random fields based on penalized pseudolikelihood," *Ann. Appl. Prob.*, vol. 6, no. 2, pp. 423–443, 1996.
- [20] M. Kalisch and P. Bühlmann, "Estimating high-dimensional directed acyclic graphs with the PC algorithm," *J. Mach. Learning Res.*, vol. 8, pp. 613–636, 2007.

- [21] N. Meinshausen and P. Bühlmann, "High-dimensional graphs and variable selection with the Lasso," *Ann. Statist.*, vol. 34, pp. 1436–1462, 2006.
- [22] P. Ravikumar, M. J. Wainwright, and J. Lafferty, "High-dimensional Ising model selection using ℓ_1 -regularized logistic regression," *Ann. Statist.*, vol. 38, no. 3, pp. 1287–1319, 2010.
- [23] P. Ravikumar, M. J. Wainwright, G. Raskutti, and B. Yu, "High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence," *Electron. J. Statist.*, vol. 5, pp. 935–980, 2011.
- [24] A. J. Rothman, P. J. Bickel, E. Levina, and J. Zhu, "Sparse permutation invariant covariance estimation," *Electron. J. Statist.*, vol. 2, pp. 494–515, 2008.
- [25] N. Santhanam, J. Dingel, and O. Milenkovic, "On modeling gene regulatory networks using Markov random fields," presented at the presented at the IEEE Inf. Theory Workshop, Volos, Greece, 2009.
- [26] P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction and Search*. Cambridge, MA: MIT Press, 2000.
- [27] F. Vega-Redondo, *Complex Social Networks Econometric Society Monographs*. Cambridge, MA: Cambridge Univ. Press, 2007.
- [28] M. J. Wainwright and M. I. Jordan, "Graphical models, exponential families and variational inference," *Found. Trends Mach. Learn.*, vol. 1, no. 1–2, pp. 1–305, Dec. 2008.
- [29] W. Wang, M. J. Wainwright, and K. Ramchandran, "Information-theoretic bounds on model selection for Gaussian Markov random fields," presented at the presented at the IEEE Int. Symp. Inf. Theory, Austin, TX, 2010.
- [30] S. Wasserman and K. Faust, *Social Network Analysis: Methods and Applications*. New York: Cambridge Univ. Press, 1994.
- [31] Y. Yang and A. Barron, "Information-theoretic determination of minimax rates of convergence," *Ann. Statist.*, vol. 27, no. 5, pp. 1564–1599, 1999.
- [32] B. Yu, "Assouad, Fano and Le Cam," in *Festschrift for Lucien Le Cam*. Berlin, Germany: Springer-Verlag, 1997, pp. 423–435.
- [33] M. Yuan and Y. Lin, "Model selection and estimation in the Gaussian graphical model," *Biometrika*, vol. 94, no. 1, pp. 19–35, 2007.

Narayana P. Santhanam (M'10) is an Assistant Professor with the Department of Electrical Engineering at the University of Hawaii, Manoa. He received a Bachelor's degree in Electrical Engineering from the Indian Institute of Technology, Chennai, and the M.S. and Ph.D. degrees from the University of California, San Diego. His research interests include the intersection of statistics and information theory, with an emphasis on combinatorial and probabilistic approaches to problems in statistics, information theory and machine learning. He is the recipient of the IEEE Best Paper award in 2006.

Martin J. Wainwright (M'03–SM'10) is currently a professor at University of California at Berkeley, with a joint appointment between the Department of Statistics and the Department of Electrical Engineering and Computer Sciences. He received a Bachelor's degree in Mathematics from University of Waterloo, Canada, and Ph.D. degree in Electrical Engineering and Computer Science (EECS) from Massachusetts Institute of Technology (MIT). His research interests include statistical signal processing, coding and information theory, statistical machine learning, and high-dimensional statistics. He has been awarded an Alfred P. Sloan Foundation Fellowship, an NSF CAREER Award, the George M. Sprowls Prize for his dissertation research (EECS department, MIT), Best Paper Awards from the IEEE Signal Processing Society (2008) and IEEE Communications Society (2010), and several outstanding conference paper awards.