

On the Difficulty of Selecting Ising Models with Approximate Recovery

Jonathan Scarlett and Volkan Cevher

Abstract—In this paper, we consider the problem of estimating the underlying graph associated with an Ising model given a number of independent and identically distributed samples. We adopt an *approximate recovery* criterion that allows for a number of missed edges or incorrectly-included edges, in contrast with the widely-studied exact recovery problem. Our main results provide information-theoretic lower bounds on the sample complexity for graph classes imposing constraints on the number of edges, maximal degree, and other properties. We identify a broad range of scenarios where, either up to constant factors or logarithmic factors, our lower bounds match the best known lower bounds for the exact recovery criterion, several of which are known to be tight or near-tight. Hence, in these cases, approximate recovery has a similar difficulty to exact recovery in the minimax sense.

Our bounds are obtained via a modification of Fano’s inequality for handling the approximate recovery criterion, along with suitably-designed ensembles of graphs that can broadly be classed into two categories: (i) Those containing graphs that contain several isolated edges or cliques and are thus difficult to distinguish from the empty graph; (ii) Those containing graphs for which certain groups of nodes are highly correlated, thus making it difficult to determine precisely which edges connect them. We support our theoretical results on these ensembles with numerical experiments.

Index Terms—Graphical model selection, Ising model, Gaussian graphical models, Markov random fields, information-theoretic limits, lower bounds, Fano’s inequality.

I. INTRODUCTION

Graphical models are a widely-used tool for providing compact representations of the conditional independence relations between random variables, and arise in areas such as image processing [1], statistical physics [2], computational biology [3], natural language processing [4], and social network analysis [5]. The problem of *graphical model selection* consists of recovering the graph structure given a number of independent samples from the underlying distribution.

While this fundamental problem is NP-hard in general [6], there exist a variety of methods guaranteeing *exact* recovery with high probability on *restricted* classes of graphs, such as bounded degree and bounded number of edges. Existing works have focused primarily on Ising models and Gaussian models, and our focus in this paper is on the former.

In particular, we focus in the problem of *approximate recovery*, in which one can tolerate some number of missed edges or incorrectly-included edges. The motivation for such a study is that the exact recovery criterion is very restrictive,

and not something that one would typically expect to achieve in practice. In particular, if the number of samples required for exact recovery is very large, it is of significant interest to know the potential savings by allowing for approximate recovery. The answer is unclear *a priori*, since this can lead to vastly improved scaling laws in some inference and learning problems [7] and virtually no gain in others [8].

Our main focus is on algorithm-independent lower bounds for Ising models, revealing the number of measurements required for approximate recovery regardless of the computational complexity. We extend Fano’s inequality [9], [10] to the case of approximate recovery, and apply it to restricted sets of graphs that prove the difficulty of approximate recovery.

Our main results reveal a broad range of graph classes for which the approximate recovery lower bounds exhibit the same scalings as the best-known exact recovery lower bounds [9], [10], which are known to be tight or near-tight in many cases of interest. This indicates that, at least for the classes that we consider, the approximate recovery problem is not much easier than the exact recovery problem in the minimax sense.

A. Problem Statement

The ferromagnetic Ising model [11] is specified by a graph $G = (V, E)$ with vertex set $V = \{1, \dots, p\}$ and edge set E . Each vertex is associated with a binary random variable $X_i \in \{-1, 1\}$, and the corresponding joint distribution is

$$P_G(x) = \frac{1}{Z} \exp\left(\sum_{i,j} \lambda_{ij} x_i x_j\right), \quad (1)$$

where

$$\lambda_{ij} = \begin{cases} \lambda & (i, j) \in E \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

and Z is a normalizing constant called the partition function. Here $\lambda > 0$ is a parameter to the distribution, sometimes called the inverse temperature.

Let $\mathbf{X} \in \{0, 1\}^{n \times p}$ be a matrix of n independent samples from this distribution, each row corresponding to one such sample of the p variables. Given \mathbf{X} , an *estimator* or *decoder* constructs an estimate \hat{G} of the graph G , or equivalently, an estimate \hat{E} of the edge set E .

Recovery Criterion: Given some class \mathcal{G} of graphs, the widely-studied exact recovery criterion seeks to characterize

$$P_e := \max_{G \in \mathcal{G}} \mathbb{P}[\hat{E} \neq E]. \quad (3)$$

We instead consider the following approximate recovery criterion, for some maximum number of errors $q_{\max} \geq 0$:

$$P_e(q_{\max}) := \max_{G \in \mathcal{G}} \mathbb{P}[|E \Delta \hat{E}| > q_{\max}], \quad (4)$$

The authors are with the Laboratory for Information and Inference Systems (LIONS), École Polytechnique Fédérale de Lausanne (EPFL) (e-mail: {jonathan.scarlett, volkan.cevher}@epfl.ch).

This work was supported in part by the European Commission under Grant ERC Future Proof, SNF 200021-146750 and SNF CRSII2-147633, and EPFL Fellows Horizon2020 grant 665667.

where $E\Delta\hat{E} = (E\setminus\hat{E}) \cup (\hat{E}\setminus E)$, so that $|E\Delta\hat{E}|$ denotes the *edit distance*, i.e., the number of edge insertions and deletions required to transform one graph to another. In this definition, q_{\max} does not depend on G , and hence, the number of allowed edge errors does not depend on the graph itself. We consider graph classes with a maximum number of edges equal to some value k , and set $q_{\max} = \theta^*k$ for some constant $\theta^* \in (0, 1)$ not scaling with the problem size. Note that $\theta^* = 1$ would trivially give $P_e(q_{\max}) = 0$.

Graph Classes: We consider the following three nested classes of graphs $\mathcal{G}_k \supseteq \mathcal{G}_{k,d} \supseteq \mathcal{G}_{k,d,\eta,\gamma}$:

- (*Edge bounded class \mathcal{G}_k*) This class contains all graphs with at most k edges.
- (*Edge and degree bounded class $\mathcal{G}_{k,d}$*) This class contains the graphs in \mathcal{G}_k such that each node has degree (i.e., number of edges it is involved in) at most d .
- (*Sparse separator class $\mathcal{G}_{k,d,\eta,\gamma}$*) This class contains the graphs in $\mathcal{G}_{k,d}$ satisfying the (η, γ) -separation condition [12]: For any two non-connected vertices in the graph, one can simultaneously block all paths of length γ or less by blocking at most η nodes.

The restriction on the number of edges is motivated by the fact that real-world graphs are often sparse. The restriction on the degree is also relevant in applications, and is particularly commonly-assumed in the statistical physics literature. The sparse separation condition is somewhat more technical, but it is of interest since it is known to permit polynomial-time exact recovery in many cases [12], [13]. Moreover, it is known to hold with high probability for several interesting random graphs; see [12] for some examples.

Generalized Edge Weights: A generalization of the above Ising model allows λ_{ij} to take different non-zero values for each $(i, j) \in E$, some of which may be negative. Previous works considering model selection for this generalized model have sought minimax bounds with respect to the graph class *and* these parameters subject to $\lambda_{\min} \leq |\lambda_{ij}| \leq \lambda_{\max}$ for some λ_{\min} and λ_{\max} . The lower bounds derived in this paper immediately imply corresponding lower bounds for this generalized setting, provided that our parameter λ in (2) lies in the range $[\lambda_{\min}, \lambda_{\max}]$.

Notation and Terminology: Throughout the paper, we let \mathbb{P}_G and \mathbb{E}_G denote probabilities and expectations with respect to P_G (e.g., $\mathbb{P}_G[X_i = X_j]$, $\mathbb{E}[X_i X_j]$). We denote the floor function by $\lfloor \cdot \rfloor$, and the ceiling function by $\lceil \cdot \rceil$. We use the standard terminology that the *degree* of a node $v \in V$ is the number of edges in E containing v , and that a *clique* is a subset $C \subset V$ of size at least two within which all pairs of nodes have an edge between them.

B. Related Work

A variety of algorithms with varying levels of computational efficiency have been proposed for selecting Ising models with rigorous guarantees, including conditional independence tests for candidate neighborhoods [14], correlation tests in the presence of sparse separators [12], [15], greedy techniques [16]–[19], convex optimization approaches [20], elementary

estimators [21], and intractable information-theoretic techniques [9].

These works have made various assumptions on the underlying model, including incoherence assumptions [20], [21] and long-range correlation assumptions [12], [15]. A notable recent work avoiding these is [19], which provides recovery guarantees using an algorithm whose complexity is only quadratic in the number of nodes for a fixed maximum degree, thus resolving an open question posed in [22].

Early works providing algorithm-independent lower bounds used only graph-theoretic properties [12], [14], [23]; the resulting bounds are loose in general, since they do not capture the effects of the parameters of the joint distribution (e.g., λ). Several refined bounds were given in [9] for graphs with a bounded degree or a bounded number of edges. Additional classes were considered in [10], including the bounded girth class and a class related to the separation criterion of [12] (and hence related to $\mathcal{G}_{k,d,\eta,\gamma}$ defined above). While our techniques build on those of [9], [10], we must consider significantly different ensembles, since those in [9], [10] contain graphs that differ only by one or two edges, thus making approximate recovery trivial.

To our knowledge, the only other work giving an approximate recovery bound for the Ising model is [24], where the degree-bounded class is considered. The effect of edge weights is not considered therein, and the bound is proved by counting graphs rather than constructing restricted ensembles. Consequently, only an $\Omega(d \log p)$ necessary condition is shown, in contrast with our bounds containing a d^2 or $e^{\lambda d}$ term (cf., Table I). The necessary conditions for list decoding [25] bear some similarity to approximate recovery, but the problem and its analysis are in fact much more similar to exact recovery, allowing the ensembles from [9], [10] to be applied directly.

Beyond Ising models, several works have provided necessary and sufficient conditions for recovering Gaussian graphical models [13], [26]–[29]. In this context, a necessary condition for approximate recovery was given in [13, Cor. 7], but the corresponding assumptions and techniques used were vastly different to ours: The random Erdős-Rényi model was considered instead of a deterministic class, and an additional walk-summability condition specific to the Gaussian model was imposed.

C. Contributions

Our main results, and the corresponding existing results for exact recovery, are summarized in Table I, where we provide necessary scaling laws on the number of samples needed to obtain a vanishing probability of error $P_e(q_{\max})$. Note that some of the exact recovery conditions given in the final column were not explicitly given in [9], [10], but they can easily be inferred from the proofs therein; see Section II for further discussion. We also observe that our analysis requires handling more cases separately compared to [9], [10]; in those works, the final three rows corresponding to \mathcal{G}_k in Table I are all a single case giving $\Omega(k \log p)$ scaling, and similarly for $\mathcal{G}_{k,d}$.

Table I reveals the following facts:

- 1) In all of the known cases where exact recovery is known to be difficult, i.e., exponential in a quantity that

Graph Class	Parameters	Necessary for approximate recovery (this paper)	Best known necessary for exact recovery [9], [10]
Bounded edge \mathcal{G}_k Distortion $q_{\max} < \frac{k}{4}$ (Theorems 1 and 2)	$\lambda = \omega\left(\frac{1}{\sqrt{k}}\right)$	Exponential in $\lambda\sqrt{k}$	Exponential in $\lambda\sqrt{k}$
	$\lambda = O\left(\frac{1}{\sqrt{k}}\right)$ $1 \ll k \ll p$	$\Omega(k \log p)$	$\Omega(k \log p)$
	$\lambda = O\left(\frac{1}{\sqrt{k}}\right)$ $p \ll k \ll p^{\frac{4}{3}}$	$\Omega(k)$	$\Omega(k \log p)$
	$\lambda = O\left(\frac{1}{\sqrt{k}}\right)$ $p^{\frac{4}{3}} \ll k \ll p^2$	$\Omega\left(\frac{p^2}{\sqrt{k}}\right)$ (between $\Omega(p)$ and $\Omega(k)$)	$\Omega(k \log p)$
Bounded edge and degree $\mathcal{G}_{k,d}$ Distortion $q_{\max} < \frac{k}{4} \frac{d-2}{d}$ (Theorems 3 and 4)	$\lambda = \omega\left(\frac{1}{d}\right)$	Exponential in λd	Exponential in λd
	$\lambda = O\left(\frac{1}{d}\right)$ $d^2 \ll k \ll p$	$\Omega(d^2 \log p)$	$\Omega(d^2 \log p)$
	$\lambda = O\left(\frac{1}{d}\right)$ $p \ll k \ll p\sqrt{d}$	$\Omega(d^2)$	$\Omega(d^2 \log p)$
	$\lambda = O\left(\frac{1}{d}\right)$ $p\sqrt{d} \ll k \leq \frac{pd}{2}$	$\Omega\left(\frac{d^3 p^2}{k^2}\right)$ (between $\Omega(d)$ and $\Omega(d^2)$)	$\Omega(d^2 \log p)$
Bounded edge and degree with sparse separators $\mathcal{G}_{k,d,\eta,\gamma}$ Distortion $q_{\max} < \frac{(c\eta-1)^2 k}{2c\eta(2\eta+m(\gamma+1))}$ ($c \in (0, 1)$, $m \in \{0, \dots, \frac{d}{2} - \eta\}$) (Theorem 5)	$\lambda = \omega\left(\min\left\{\frac{1}{\sqrt{\eta}}, \frac{1}{m^{\frac{1}{1+\gamma}}}\right\}\right)$ $\lambda = O(1)$ $k \leq \frac{p}{4}$	Exponential in $\max\{\lambda^2 \eta, \lambda^{\gamma+1} m\}$	Exponential in $\max\{\lambda^2 \eta, \lambda^{\gamma+1} d\}$
	$\lambda = O\left(\min\left\{\frac{1}{\sqrt{\eta}}, \frac{1}{m^{\frac{1}{1+\gamma}}}\right\}\right)$ $\lambda = O(1)$ $k \leq \frac{p}{4}$	$\Omega\left(\max\left\{\eta, m^{\frac{2}{\gamma+1}}\right\} \log p\right)$	$\Omega\left(\max\left\{\eta, d^{\frac{2}{\gamma+1}}\right\} \log p\right)$

Table I: Summary of main results on partial recovery, and comparisons to the best known necessary conditions for exact recovery. Each entry shows the necessary scaling law for the number of samples required to achieve a vanishing error probability.

increases in the problem dimension, the same difficulty is observed for approximate recovery, at least for the values of q_{\max} shown. For \mathcal{G}_k and $\mathcal{G}_{k,d}$, this is true even when we allow for up to a quarter of the edges to be in error. Note that we did not seek to optimize this fraction in our analysis, and we expect similar difficulties to arise even when higher proportions of errors are allowed. In fact, by a simple variation of our analysis outlined in Remark 1 in Section IV-C, we can already increase this fraction from $\frac{1}{4}$ to $\frac{1}{2}$.

- 2) In many of the cases where the necessary conditions for exact recovery lack exponential terms, the corresponding necessary conditions for approximate recovery are identical or near-identical; in particular, see the second and third rows corresponding to \mathcal{G}_k , the second and third rows corresponding to $\mathcal{G}_{k,d}$, and the second row corresponding to $\mathcal{G}_{k,d,\eta,\gamma}$ with $m = \frac{d}{2} - \eta$. While there are logarithmic terms missing in some cases (e.g., k vs. $k \log p$), these are typically insignificant in the regimes considered (e.g., $k = \Omega(p)$).
- 3) In contrast, there are some cases where significant gaps remain between the best-known conditions for exact recovery and approximate recovery. The two most extreme cases are as follows: (i) If $k = \Theta(p^{2-\epsilon})$ for some small $\epsilon > 0$, the necessary conditions for \mathcal{G}_k are $\Omega(p^{2-\epsilon} \log p)$ and $\Omega(p^{1+\epsilon/2})$, respectively; (ii) If $k = \Theta(pd)$, then the necessary conditions for $\mathcal{G}_{k,d}$ are $\Omega(d^2 \log p)$ and $\Omega(d \log p)$, respectively. It remains an open problem as to whether this behavior is fundamental, or due to a weakness in the analysis.

The starting point of our results is a modification of Fano's inequality for the purpose of handling approximate recovery. To obtain the above results, we apply this bound to ensembles

of graphs that can be broadly classed into two categories. The first considers graphs with a large number of isolated edges, or more generally, isolated cliques. We characterize how difficult each graph is to distinguish from the empty graph, and use this to derive the results given in item 2) above. On the other hand, the results on the exponential terms discussed in item 1) arise from considering ensembles in which several groups of nodes are always highly correlated due to the presence of a large number of edges among them, thus making it difficult to determine precisely which edges these are.

Both of these categories help in providing bounds that match those for exact recovery. For example, the $\Omega(k \log p)$ behavior for $\lambda = O\left(\frac{1}{k}\right)$ in [9] is proved by considering graphs with a single isolated edge, and our analysis extends this to approximate recovery by considering graphs with k isolated edges. Analogously, the exponential behavior (e.g., in $\lambda\sqrt{k}$) in [9] is proved by considering cliques with one edge removed, and our analysis reveals that the same exponential behavior arises even if a constant fraction of the edges are removed.

We provide numerical results on our ensembles in Section VI supporting our theoretical findings. Specifically, we implement optimal or near-optimal decoding rules in a variety of cases, and find that while approximate recovery can be easier than exact recovery, the general behavior of the two is similar.

II. MAIN RESULTS

In this section, we present our main results, namely, algorithm-independent necessary conditions for the criterion in (4) with all $\lambda_{ij} = \lambda$. Our conditions are written in terms of asymptotic $o(1)$ terms for clarity, but purely non-asymptotic variants can be inferred from the proofs. Throughout the section, we make use of the binary entropy function in nats,

$H_2(\theta) := -\theta \log \theta - (1-\theta) \log(1-\theta)$. Here and subsequently, all logarithms have base e .

All proofs are deferred to later sections; some preliminary results are presented in Section III, a number of ensembles are presented and analyzed in Section IV, and the resulting theorems are deduced in Section V.

A. Bounded Number of Edges Class \mathcal{G}_k

We first consider the class \mathcal{G}_k of graphs with at most k edges. It will prove convenient to treat two cases separately depending on how k scales with p .

Theorem 1. (Class \mathcal{G}_k with $k \leq p/4$) *For any number of edges such that $k \rightarrow \infty$ and $k \leq p/4$, and any distortion level $q_{\max} = \lfloor \theta k \rfloor$ for some $\theta \in (0, \frac{1}{4})$, it is necessary that*

$$n \geq \max \left\{ \frac{e^{\lambda(\sqrt{k/2}-2)/2} (\log 2 - H_2(2\theta))}{6\lambda k}, \frac{2(1-\theta) \log p}{\lambda \tanh \lambda} \right\} (1 - \delta - o(1)) \quad (5)$$

in order to have $P_e(q_{\max}) \leq \delta$ for all $G \in \mathcal{G}_k$.

We proceed by considering two cases as in [9]. In the case that $\lambda\sqrt{k} \rightarrow \infty$ at any rate faster than logarithmic in p (or even logarithmic with a constant that is not too small), the sample complexity is dominated by the exponential term $e^{\lambda(\sqrt{k}-2)/2}$, and any recovery procedure requires a huge number of samples. Thus, in this case, even the approximate recovery problem is very difficult. On the other hand, if $\lambda = O(\frac{1}{\sqrt{k}})$ then the second condition in (5) gives a sample complexity of $\Omega(k \log p)$, since $\tanh \lambda = O(\lambda)$ as $\lambda \rightarrow 0$.

These observations are the same as those made for exact recovery in [9], where the best known necessary conditions for \mathcal{G}_k were given. Thus, we have reached similar conclusions even allowing for nearly a quarter of the edges to be in error.

Theorem 2. (Class \mathcal{G}_k with $k = \Omega(p)$) *For any number of edges of the form $k = \lfloor cp^{1+\nu} \rfloor$ for constants $c > 0$ and $\nu \in [0, 1)$, and any distortion level $q_{\max} = \lfloor \theta k \rfloor$ for some $\theta \in (0, \frac{1}{4})$, it is necessary that*

$$n \geq \max \left\{ \frac{e^{\lambda(\sqrt{k/2}-2)/2} (\log 2 - H_2(2\theta))}{6\lambda k}, \frac{\log 2 - H_2(\theta)}{\lambda \frac{e^{2\lambda} \cosh(4\lambda cp^\nu) - 1}{e^{2\lambda} \cosh(4\lambda cp^\nu) + 1}} \right\} (1 - \delta - o(1)) \quad (6)$$

in order to have $P_e(q_{\max}) \leq \delta$ for all $G \in \mathcal{G}_k$.

As above, the sample complexity is exponential in $\lambda\sqrt{k}$ due to the first term in (6). On the other hand, we claim that when $\lambda = O(\frac{1}{\sqrt{k}})$, the second term in (6) leads to the sample complexity $O(\min\{k, p^2/\sqrt{k}\})$. To see this, we choose k as in the theorem statement and note that $\lambda p^\nu = O(p^{-\frac{1}{2}(1+\nu)+\nu}) = O(p^{-\frac{1}{2}(1-\nu)})$; since $\cosh \zeta = 1 + O(\zeta^2)$ as $\zeta \rightarrow 0$, this implies that $\cosh(2c\lambda p^\nu) = 1 + O(p^{-(1-\nu)})$. We thus have $e^{2\lambda} \cosh(2c\lambda p^\nu) = (1 + O(p^{-(1-\nu)})) (1 + O(p^{-(1-\nu)}))$, which finally yields $\frac{e^{2\lambda} \cosh(2c\lambda p^\nu) - 1}{e^{2\lambda} \cosh(2c\lambda p^\nu) + 1} = O(\max\{p^{-\frac{1}{2}(1+\nu)}, p^{-(1-\nu)}\}) = O(\max\{1/\sqrt{k}, k/p^2\})$.

When $k = \Omega(p)$ and $k = O(p^{4/3})$, we have $\min\{k, p^2/\sqrt{k}\} = k$, and hence, these observations are again the same as those made for exact recovery in [9], except that our growth rates do not include a $\log p$ term; this logarithmic factor is insignificant compared to the leading term $k = \Omega(p)$. In contrast, the gap is more significant when $k \gg p^{4/3}$; in the extreme case, when $k = \Theta(p^{2-\epsilon})$ for some small $\epsilon > 0$, we obtain a scaling of $\Omega(p^{1+\epsilon/2})$, as opposed to $\Omega(k \log p) = \Omega(p^{2-\epsilon} \log p)$.

B. Bounded Degree Class $\mathcal{G}_{k,d}$

Next, we consider the class $\mathcal{G}_{k,d}$ of graphs such that every node has degree at most d , and the total number of edges does not exceed k .

Theorem 3. (Class $\mathcal{G}_{k,d}$ with $k \leq p/4$) *For any maximal degree $d > 2$ and number of edges k such that $k = \omega(d^2)$ and $k \leq p/4$, and any distortion level $q_{\max} = \lfloor \theta k \rfloor$ for some $\theta \in (0, \frac{1}{4} \frac{d-2}{d})$, it is necessary that*

$$n \geq \max \left\{ \frac{e^{\lambda(d-2)/4} (\log 2 - H_2(\frac{d}{d-2} \cdot 2\theta))}{3\lambda d^2}, \frac{2(1-\theta) \log p}{\lambda \tanh \lambda} \right\} (1 - \delta - o(1)) \quad (7)$$

in order to have $P_e(q_{\max}) \leq \delta$ for all $G \in \mathcal{G}_{k,d}$.

The first term in (7) reveals that the sample complexity is exponential in λd . On the other hand, if $\lambda = O(\frac{1}{d})$ then the second term gives a sample complexity of $\Omega(d^2 \log p)$.

We cannot directly compare Theorem 3 to [9], since there k was assumed to be unrestricted for the degree-bounded ensemble. However, the analysis therein is easily extended to $\mathcal{G}_{k,d}$, and doing so recovers the nearly identical observations to those above, as summarized in Table I. In this sense, Theorem 3 matches the best known necessary conditions for exact recovery even when nearly a quarter of the edges may be in error.

Theorem 4. (Class $\mathcal{G}_{k,d}$ with $k = \Omega(p)$) *For any maximal degree $d > 2$ and number of edges k such that $k = \omega(d^2)$ and $k \leq \frac{1}{2}p(d-1)$ for some $d' \leq d$, and any distortion level $q_{\max} = \lfloor \theta k \rfloor$ for some $\theta \in (0, \frac{1}{4} \frac{d'-2}{d'})$, it is necessary that*

$$n \geq \max \left\{ \frac{e^{\lambda(d-2)/4} (\log 2 - H_2(\frac{d}{d-2} \cdot 2\theta))}{3\lambda d^2}, \frac{\log 2 - H_2(\theta)}{\lambda \frac{e^{2\lambda} \cosh(2\lambda d') - 1}{e^{2\lambda} \cosh(2\lambda d') + 1}} \right\} (1 - \delta - o(1)) \quad (8)$$

in order to have $P_e(q_{\max}) \leq \delta$ for all $G \in \mathcal{G}_{k,d}$.

The sample complexity remains exponential in λd . By some standard asymptotic expansions similar to those following Theorem 2, we have $\frac{e^{2\lambda} \cosh(2\lambda d') - 1}{e^{2\lambda} \cosh(2\lambda d') + 1} = O(\max\{\frac{1}{d'}, (\frac{d'}{d})^2\})$ whenever $\lambda = O(\frac{1}{d})$; hence, the second condition in (8) becomes $n = \Omega(d \min\{d, (\frac{d'}{d})^2\})$. Thus, if $d' = O(\sqrt{d})$ then we again get the desired $n = O(d^2 \log p)$ behavior; this means that we can allow for k up to $O(p\sqrt{d})$. More generally, we instead get the possibly weaker scaling law

$n = \Omega(\min\{d^2, d^3/(d')^2\})$, which is equivalent to $n = \Omega(\min\{d^2, \frac{d^3 p^2}{k^2}\})$ when $k = \Theta(pd')$. In the extreme case, when $k = \Theta(pd)$ (the highest growth rate possible given the degree constraint alone), this only recovers $\Omega(d \log p)$ scaling.

C. Sparse Separator Class $\mathcal{G}_{k,d,\eta,\gamma}$

We now consider the class $\mathcal{G}_{k,d,\eta,\gamma}$ of graphs in $\mathcal{G}_{k,d}$ that satisfy the (η, γ) -separation condition [12]. We focus on the case $k \leq p/4$, since the main graph ensemble that we consider for this class is not suited to the case that $k = \omega(p)$.

Theorem 5. (Class $\mathcal{G}_{k,d,\eta,\gamma}$ with $k \leq p/4$) Fix any parameters (d, k, η, γ) with $k \leq p/4$ and $\eta \leq \lfloor \frac{d}{2} \rfloor$, and let m be an integer in $\{0, \dots, \lfloor \frac{d}{2} \rfloor - \eta\}$. For any distortion level $q_{\max} = \lfloor \theta \frac{(c\eta-1)^2 k}{2c\eta(2\eta+m(\gamma+1))} \rfloor$ for some $\theta \in (0, \frac{1}{2})$ and $c \in (\frac{1}{\eta}, 1]$, it is necessary that

$$n \geq \max \left\{ \frac{\left(1 + (\cosh(2\lambda))^{(1-c)\eta-1} \left(\frac{1+(\tanh \lambda)^{\gamma+1}}{1-(\tanh \lambda)^{\gamma+1}}\right)^m\right)}{2\lambda c\eta} \right. \\ \left. \times (\log 2 - H_2(\theta)), \frac{2(k - q_{\max}) \log p}{k\lambda \tanh \lambda} \right\} (1 - \delta - o(1)) \quad (9)$$

in order to have $P_e(q_{\max}) \leq \delta$ for all $G \in \mathcal{G}_{k,d,\eta,\gamma}$.

We proceed by considering only the case $\lambda = O(1)$, though simplifications of Theorem 5 for $\lambda \rightarrow \infty$ are also possible. With $\lambda = O(1)$, we have $(\cosh(2\lambda))^{(1-c)\eta} = e^{\zeta \lambda^2 (1-c)\eta}$ for some $\zeta = \Theta(1)$, and similarly $\left(\frac{1+(\tanh \lambda)^{\gamma+1}}{1-(\tanh \lambda)^{\gamma+1}}\right)^m = e^{\zeta' m \lambda^{\gamma+1}}$ for some $\zeta' = \Theta(1)$ [10, Sec. 5]. These identities reveal that the sample complexity is exponential in both $\lambda^2 \eta$ and $\lambda^{\gamma+1} m$. On the other hand, if $\lambda = O(\frac{1}{\sqrt{\eta}})$ and $\lambda = O(\frac{1}{m^{\frac{2}{\gamma+1}}})$ then the second term in (9) gives $n = \Omega(\max\{\eta, m^{\frac{2}{\gamma+1}}\} \log p)$.

Due to the choice $q_{\max} = \lfloor \theta \frac{(c\eta-1)^2 k}{2c\eta(2\eta+m(\gamma+1))} \rfloor$, if we set $m = d/2 - \eta$ then we are only in the regime of a constant fraction of errors if $d\gamma = \Theta(\eta)$. This is true, for example, if $\eta = \Theta(d)$ so that the separator set size is a fixed fraction of the maximum degree, and $\gamma = \Theta(1)$ so that the separation is with respect to paths of a bounded length.

More generally, to handle larger values of q_{\max} , one can choose a smaller value of m , thus leading to a larger value of q_{\max} but with a less stringent condition on the number of measurements in (9). In the extreme case, $m = 0$, and then we are always in the regime of a constant proportion of errors; however, this yields a necessary condition $\Omega(\eta \log p)$ not depending on d or γ .

The graph family studied in [10, Thm. 2] is somewhat different from $\mathcal{G}_{k,d,\eta,\gamma}$, in particular not putting any constraints on the maximal degree nor the number of edges. Nevertheless, by choosing the parameters in the proof therein to meet these constraints,¹ one again obtains similar conditions to those above, as summarized in Table I. In particular, for any choice of m that grows as $\Theta(d)$, the scaling laws for exact recovery and approximate recovery coincide.

¹Specifically, in [10, Sec. 9.2], one can set $t_\nu = d - \eta$ to satisfy the degree constraint, and then choose $\alpha = \lfloor \frac{k}{t_\nu(\gamma+1)+2\eta-1} \rfloor$ to ensure there are at most k edges in total.

III. AUXILIARY RESULTS

In this section, we provide a number of auxiliary results that will be used to prove the theorems in Section II. We first present a general form of Fano's depending on both the Kullback-Leibler (KL) divergence and edit distance between graphs, and then provide a number of properties of Ising models that will be useful for characterizing the KL divergence and edit distance in specific scenarios.

A. Fano's Inequality for Approximate Recovery

As is common in studies of algorithm-independent lower bounds in learning problems, we make use of bounds based on Fano's inequality [30, Sec. 2.10]. We first briefly outline the most relevant results for the exact recovery problem.

Recall the definitions of P_e and $P_e(q_{\max})$ in (3)–(4) with respect to a given graph class \mathcal{G} . It is known that for any subset $\mathcal{T} \subseteq \mathcal{G}$, and any covering set $C_{\mathcal{T}}(\epsilon)$ such that any graph $G \in \mathcal{T}$ has an “ ϵ -close” graph $G' \in C_{\mathcal{T}}(\epsilon)$ satisfying $D(P_G \| P_{G'}) \leq \epsilon$, we have [10]

$$P_e \geq 1 - \frac{\log |C_{\mathcal{T}}(\epsilon)| + n\epsilon + \log 2}{\log |\mathcal{T}|}. \quad (10)$$

In particular, if $C_{\mathcal{T}}(\epsilon)$ is a singleton, solving for n gives the necessary condition

$$n \geq \frac{\log |\mathcal{T}|}{\epsilon} \left(1 - \delta - \frac{\log 2}{\log |\mathcal{T}|}\right) \quad (11)$$

in order to have $P_e \leq \delta$.

For approximate recovery, we consider ensembles (i.e., choices of \mathcal{T}) for which the decoder's outputs may lie in some set \mathcal{T}' without loss of optimality; in most cases we will have $\mathcal{T} = \mathcal{T}'$, but in general, \mathcal{T}' need not even be a subset of the graph class \mathcal{G} . We use the following generalization of (11).

Lemma 1. Suppose that the decoder minimizing the average error probability with respect to a distortion level q_{\max} , averaged over a graph uniformly drawn from a set $\mathcal{T} \subseteq \mathcal{G}$, always outputs a graph in some set \mathcal{T}' . Moreover, suppose that there exists a graph G' such that $D(P_G \| P_{G'}) \leq \epsilon$ for all $G \in \mathcal{T}$, and that there are at most $A(q_{\max})$ graphs in \mathcal{T}' within an edit distance q_{\max} of any given graph $G \in \mathcal{T}$. Then it is necessary that

$$n \geq \frac{\log |\mathcal{T}| - \log A(q_{\max})}{\epsilon} \left(1 - \delta - \frac{\log 2}{\log |\mathcal{T}'|}\right) \quad (12)$$

in order to have $P_e(q_{\max}) \leq \delta$.

Proof. See Appendix A. \square

B. Properties of Ferromagnetic Ising Models

We will use a number of useful results on ferromagnetic Ising models, each of which is either self-evident or can be found in [9] or [10]. We start with some basic properties.

Lemma 2. For any graphs G and G' with edge sets E and E' respectively, we have the following:

(i) For any pair (i, j) , we have [9]

$$\mathbb{E}_G[X_i X_j] = 2\mathbb{P}_G[X_i = X_j] - 1. \quad (13)$$

(ii) The divergence between the corresponding distributions satisfies [10, Eq. (4)]

$$D(P_G \| P_{G'}) \leq \sum_{(i,j) \in E \setminus E'} \lambda (\mathbb{E}_G[X_i X_j] - \mathbb{E}_{G'}[X_i X_j]) + \sum_{(i,j) \in E' \setminus E} \lambda (\mathbb{E}_{G'}[X_i X_j] - \mathbb{E}_G[X_i X_j]) \quad (14)$$

$$\leq \sum_{(i,j) \in E \setminus E'} \lambda (1 - \mathbb{E}_{G'}[X_i X_j]) + \sum_{(i,j) \in E' \setminus E} \lambda (1 - \mathbb{E}_G[X_i X_j]). \quad (15)$$

(iii) If $E' \subset E$, then we have for any pair (i, j) that [10, Eq. (13)]

$$\mathbb{E}_G[X_i X_j] \geq \mathbb{E}_{G'}[X_i X_j]. \quad (16)$$

(iv) Let (V_1, \dots, V_K) be a partition of V into K disjoint non-empty subsets. If G and G' are such that there are no edges between nodes in V_i and V_j when $i \neq j$, then

$$D(P_G \| P_{G'}) = \sum_{i=1}^K D(P_{G_i} \| P_{G'_i}), \quad (17)$$

where $G_i = (V, E_i)$, with E_i containing the edges in E between nodes in V_i (and analogously for G'_i).

The remaining properties concern the probabilities, expectations and divergences associated with more specific graphs.

Lemma 3. (i) If G' is obtained from G by removing a single edge (i, j) , then [10, Eq. (19)]

$$\frac{\mathbb{P}_G[X_i = X_j]}{1 - \mathbb{P}_G[X_i = X_j]} = e^{2\lambda} \frac{\mathbb{P}_{G'}[X_i = X_j]}{1 - \mathbb{P}_{G'}[X_i = X_j]} \quad (18)$$

and [10, Lemma 4]

$$D(P_G \| P_{G'}) \leq \lambda \tanh \lambda. \quad (19)$$

(ii) Let G contain a clique on $m \geq 2$ nodes and no other edges, and let G' be obtained from G by removing a single edge (i, j) . Then, defining $\bar{m} := m - 1$, we have [9, Eq. (31)]

$$\frac{\mathbb{P}_{G'}[X_i = X_j]}{1 - \mathbb{P}_{G'}[X_i = X_j]} = \frac{\sum_{j=0}^{\bar{m}} \binom{\bar{m}}{j} \exp\left(\frac{\lambda}{2}(2j - \bar{m})^2\right) \exp(2\lambda(2j - \bar{m}))}{\sum_{j=0}^{\bar{m}} \binom{\bar{m}}{j} \exp\left(\frac{\lambda}{2}(2j - \bar{m})^2\right)}. \quad (20)$$

Moreover, we have [9, Lemma 1]

$$\mathbb{P}_{G'}[X_i = X_j] \geq 1 - \frac{\bar{m}}{\bar{m} + e^{\bar{m}\lambda/2}} \quad (21)$$

and

$$\mathbb{E}_{G'}[X_i X_j] \geq 1 - \frac{2\bar{m}e^\lambda}{e^{\bar{m}\lambda} + \bar{m}e^\lambda}. \quad (22)$$

(iii) Suppose that for some edge $(i, j) \in E \Delta E'$, there exist at least m node-disjoint paths of length ℓ between i and j in G . Then [10, Lemma 3]

$$\mathbb{E}_G[X_i X_j] \geq 1 - \frac{2}{1 + \left(\frac{1 + (\tanh \lambda)^\ell}{1 - (\tanh \lambda)^\ell}\right)^m}. \quad (23)$$

If the same is true in both G and G' for all $(i, j) \in E \Delta E'$, then [10, Cor. 3]

$$D(P_G \| P_{G'}) \leq \frac{2\lambda |E \Delta E'|}{1 + \left(\frac{1 + (\tanh \lambda)^\ell}{1 - (\tanh \lambda)^\ell}\right)^m}. \quad (24)$$

(iv) More generally, if there exist at least m_l node-disjoint paths of length ℓ_l between (i, j) for $l = 1, \dots, L$, where the values of ℓ_l are all distinct, then

$$\mathbb{E}_G[X_i X_j] \geq 1 - \frac{2}{1 + \prod_{l=1}^L \left(\frac{1 + (\tanh \lambda)^{\ell_l}}{1 - (\tanh \lambda)^{\ell_l}}\right)^{m_l}}. \quad (25)$$

IV. GRAPH ENSEMBLES AND LOWER BOUNDS ON THEIR SAMPLE COMPLEXITIES

In this section, we provide necessary conditions for the approximate recovery of a number of ensembles, making use of the tools from the previous section. In particular, we seek choices of \mathcal{T} , \mathcal{T}' and $A(q_{\max})$ for substitution into Fano's inequality in Lemma 1. In Section V, we use these to establish our main theorems.

A. Ensemble 1: Many Isolated Edges

This ensemble contains numerous isolated edges, such that if λ is small then it is difficult to determine precisely which ones are present. It is constructed as follows with some integer parameter $\alpha \leq p/4$:

Ensemble1(α) [Isolated edges ensemble]:

- Each graph in \mathcal{T} is obtained by forming exactly α node-disjoint edges that may otherwise be arbitrary.

For this ensemble, we have the following properties:

- The number of graphs is $|\mathcal{T}| = \prod_{i=0}^{\alpha} \binom{p-2i}{2} \geq \binom{\lfloor p/2 \rfloor}{2}^\alpha$, since $p - 2\alpha \geq \frac{p}{2}$ by the assumption $\alpha \leq p/4$.
- The maximum degree of each graph is one.
- For this ensemble, it suffices to trivially let \mathcal{T}' contain all graphs.
- The number of graphs within an edit distance q_{\max} of any single graph is upper bounded as $A(q_{\max}) \leq \sum_{q=0}^{q_{\max}} \sum_{q'=0}^{q_{\max}-q} \binom{\alpha}{q} \binom{p}{q'} \leq (1 + q_{\max})^2 \binom{\alpha}{\lfloor \alpha/2 \rfloor} \binom{p}{2}^{q_{\max}}$. Here the term $\binom{\alpha}{q}$ corresponds to choosing q edges to remove, and the term $\binom{p}{2}^{q'}$ upper bounds the number of ways to add $q' \leq q_{\max} - q$ new edges. We have also used the fact that $\binom{\alpha}{q}$ is maximized at $q = \lfloor \alpha/2 \rfloor$.
- From (19), the KL divergence from a single-edge graph to the empty graph is upper bounded by $\lambda \tanh \lambda$. Using this fact along with (17), any graph in \mathcal{T} has a KL divergence to the empty graph of at most $\epsilon = \alpha \lambda \tanh \lambda$.

Combining these with (12) gives the necessary condition

$$n \geq \frac{\alpha \log \binom{\lfloor p/2 \rfloor}{2} - \log \left((1 + q_{\max})^2 \binom{\alpha}{\lfloor \alpha/2 \rfloor} \binom{p}{2}^{q_{\max}} \right)}{\alpha \lambda \tanh \lambda} \times \left(1 - \delta - \frac{\log 2}{|\mathcal{T}|} \right) \quad (26)$$

in order to have $P_e(q_{\max}) \leq \delta$.

Simplifying both $\log \binom{\lfloor p/2 \rfloor}{\alpha}$ and $\log \binom{p}{2}$ to $(2 \log p)(1 + o(1))$, and writing $\log \binom{\alpha}{\lfloor \alpha/2 \rfloor} \leq \alpha \log 2 = o(\alpha \log p)$ as well as $\log(1 + q_{\max})^2 \leq 2 \log(1 + \alpha) = o(\alpha \log p)$, we can simplify (26) to

$$n \geq \frac{2\alpha \log p - 2q_{\max} \log p}{\alpha \lambda \tanh \lambda} (1 - \delta - o(1)), \quad (27)$$

provided that $\alpha \rightarrow \infty$ and $q_{\max} \leq (1 - \Omega(1))\alpha$. Letting $q_{\max} = \lfloor \theta_1 \alpha \rfloor$ for some $\theta_1 \in (0, 1)$, this becomes

$$n \geq \frac{2(1 - \theta_1) \log p}{\lambda \tanh \lambda} (1 - \delta - o(1)). \quad (28)$$

B. Ensemble 2: Many Isolated Groups of Nodes

As an alternative to Ensemble 1, this ensemble allows for significantly more edges, in particular permitting $k = \omega(p)$. It is constructed as follows with integer parameters m and α :

Ensemble2(m, α) [Isolated cliques ensemble]:

- Form α fixed groups of nodes, each containing m nodes.
- Each graph in \mathcal{T} is formed by forming arbitrarily many edges within each group, but no edges between the groups.

For this ensemble, we have the following:

- The number of nodes forming these groups is $m\alpha$.
- The total number of possible edges is $\alpha \binom{m}{2}$, and hence the total number of graphs is $|\mathcal{T}| = 2^{\alpha \binom{m}{2}}$.
- The maximal degree of each graph is at most $m - 1$.
- The decoder can output an element of \mathcal{T} without loss of optimality, since any inter-group edges declared to be present are guaranteed to be wrong. Thus, we may set $\mathcal{T}' = \mathcal{T}$.
- The number of graphs within an edit distance q_{\max} of any single graph is $A(q_{\max}) = \sum_{q=0}^{q_{\max}} \binom{\alpha \binom{m}{2}}{q} \leq 1 + q_{\max} \binom{\alpha \binom{m}{2}}{q_{\max}}$, assuming $q_{\max} \leq \frac{1}{2} \alpha \binom{m}{2}$.
- In Lemma 4 below, we show that the KL divergence of the graph associated with one group to the corresponding empty graph is upper bounded by $\binom{m}{2} \lambda \frac{e^{2\lambda} \cosh(2\lambda m) - 1}{e^{2\lambda} \cosh(2\lambda m) + 1}$. Hence, the KL divergence of any $G \in \mathcal{T}$ to the empty graph is upper bounded by $\epsilon = \alpha \binom{m}{2} \lambda \frac{e^{2\lambda} \cosh(2\lambda m) - 1}{e^{2\lambda} \cosh(2\lambda m) + 1}$ due to (17).

Substituting these into (12), setting $q_{\max} = \lfloor \theta_2 \alpha \binom{m}{2} \rfloor$ for some $\theta_2 \in (0, \frac{1}{2})$, and applying some simplifications, we obtain the following necessary condition for $P_e(q_{\max}) \leq \delta$:

$$n \geq \frac{\log 2 - H_2(\theta_2)}{\lambda \frac{e^{2\lambda} \cosh(2\lambda m) - 1}{e^{2\lambda} \cosh(2\lambda m) + 1}} (1 - \delta - o(1)), \quad (29)$$

whenever $\alpha \binom{m}{2} \rightarrow \infty$. Note that the binary entropy function arises from the identity $\binom{N}{\lfloor \theta N \rfloor} = e^{n H_2(\theta)(1+o(1))}$ as $N \rightarrow \infty$.

It remains to prove the claim on the KL divergence, formalized as follows.

Lemma 4. *Let G denote an arbitrary graph with edges connected to at most $m \geq 2$ nodes, and let G' be the empty graph. Then, it holds that*

$$D(P_G \| P_{G'}) \leq \binom{m}{2} \lambda \frac{e^{2\lambda} \cosh(2\lambda m) - 1}{e^{2\lambda} \cosh(2\lambda m) + 1}. \quad (30)$$

Proof. We prove the claim for the case that G contains a single m -clique; the general case then follows in a similar fashion using (16).

Let \bar{G} be obtained from G by removing a single edge, say indexed by (i, j) . Defining $q(G) := \mathbb{P}_G[X_i = X_j]$ and $\bar{m} := m - 1$, we have from (18) that

$$\frac{q(G)}{1 - q(G)} = e^{2\lambda} \frac{q(\bar{G})}{1 - q(\bar{G})}, \quad (31)$$

and from (20) that

$$\frac{q(\bar{G})}{1 - q(\bar{G})} = \frac{\sum_{j=0}^{\bar{m}} \binom{\bar{m}}{j} \exp\left(\frac{\lambda}{2}(2j - \bar{m})^2\right) \exp(2\lambda(2j - \bar{m}))}{\sum_{j=0}^{\bar{m}} \binom{\bar{m}}{j} \exp\left(\frac{\lambda}{2}(2j - \bar{m})^2\right)}. \quad (32)$$

Noting the symmetry of the summands with respect to j and $\bar{m} - j$, we obtain the following when \bar{m} is odd (the case that \bar{m} is even is handled similarly, leading to the same conclusion):

$$\begin{aligned} & \frac{q(\bar{G})}{1 - q(\bar{G})} \\ &= \frac{\sum_{j=0}^{\lfloor \bar{m}/2 \rfloor} \binom{\bar{m}}{j} \exp\left(\frac{\lambda}{2}(2j - \bar{m})^2\right) \cdot 2 \cosh(2\lambda(2j - \bar{m}))}{2 \sum_{j=0}^{\lfloor \bar{m}/2 \rfloor} \binom{\bar{m}}{j} \exp\left(\frac{\lambda}{2}(2j - \bar{m})^2\right)} \end{aligned} \quad (33)$$

$$\leq \max_{j=0, \dots, \lfloor \bar{m}/2 \rfloor} \cosh(2\lambda(2j - \bar{m})) \quad (34)$$

$$= \cosh(2\lambda \bar{m}) \quad (35)$$

$$\leq \cosh(2\lambda m). \quad (36)$$

Substituting (36) into (31), solving for $q(G)$, and converting from probability to expectation via (13), we obtain

$$\mathbb{E}_G[X_i X_j] \leq \frac{e^{2\lambda} \cosh(2\lambda m) - 1}{e^{2\lambda} \cosh(2\lambda m) + 1}. \quad (37)$$

The proof is concluded by substituting into (14) and noting that $\mathbb{E}_{G'}[X_i X_j] = 0$, $|E \setminus E'| = \binom{m}{2}$, and $|E' \setminus E| = 0$. \square

C. Ensemble 3: Large Inter-Connected Cliques

This ensemble involves cliques with numerous edges between them, making it difficult to determine precisely which inter-clique connections are present, particularly for large cliques and large values of λ . It is constructed as follows with integer parameters m and α :

Ensemble3(m, α) [Inter-connected cliques ensemble]:

- Construct a fixed “building block” as follows: Take an arbitrary subset of the p vertices of size $2m$, split the $2m$ vertices into two sets of size m each, fully connect each of those sets, and then put m extra edges between the two sets in a fixed but arbitrary one-to-one fashion.
- Form α disjoint copies of this building block to obtain a base graph G_0 .
- Each graph in \mathcal{T} is formed by taking G_0 and adding an arbitrary number of additional edges between each pair of partially-connected cliques. Thus, G_0 itself contains the fewest edges within \mathcal{T} , and the union of α cliques of size $2m$ contains the most edges.

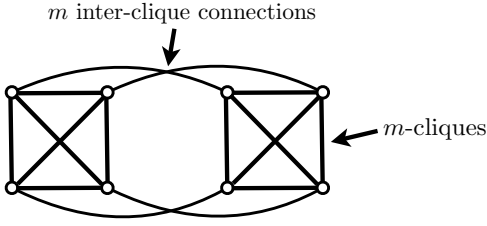


Figure 1: Building block for Ensemble 3 with $m = 4$.

An illustration of one building block is given in Figure 1. For this ensemble, we have the following:

- The number of nodes forming these groups is $2m\alpha$, and the number of edges in each graph is upper bounded by $\alpha \binom{2m}{2} \leq 2\alpha m^2$.
- The number of potential edges between two m -cliques is m^2 , and m of them are always there in each building block. Hence, the number of ways of adding edges to one building block is $2^{m(m-1)}$, and the total number of graphs is $2^{\alpha m(m-1)}$.
- The maximal degree of each graph is at most $2m - 1$.
- Similarly to Ensemble 2, the decoder can output an element of \mathcal{T} without loss of optimality, so that $\mathcal{T}' = \mathcal{T}$.
- The number of graphs within an edit distance q_{\max} of any single graph is $A(q_{\max}) = \sum_{q=0}^{q_{\max}} \binom{\alpha m(m-1)}{q} \leq 1 + q_{\max} \binom{\alpha m(m-1)}{q_{\max}}$, assuming $q_{\max} \leq \frac{1}{2} \alpha m(m-1)$.
- In Lemma 5 below, we show that the KL divergence of the graph associated with one group to the $2m$ -clique graph is upper bounded by $12\lambda m^4 e^{-\lambda(m-1)/2}$. Thus, the KL divergence from any $G \in \mathcal{T}$ to the union of α $2m$ -cliques is upper bounded by $\epsilon = 12\lambda \alpha m^4 e^{-\lambda(m-1)/2}$ due to (17).

Substituting these into (12), setting $q_{\max} = \lfloor \theta_3 \alpha m(m-1) \rfloor$ for some $\theta_3 \in (0, \frac{1}{2})$, and simplifying, we obtain

$$n \geq \frac{e^{\lambda(m-1)/2} (\log 2 - H_2(\theta_3))}{12\lambda m^2} (1 - \delta - o(1)), \quad (38)$$

whenever $\alpha m(m-1) \rightarrow \infty$.

It remains to prove the claim on the KL divergence, formalized as follows.

Lemma 5. *Let G denote the graph corresponding to a single group in Ensemble 3, and let G' be the corresponding graph containing a $2m$ -clique. Then*

$$D(P_G \| P_{G'}) \leq 12\lambda m^4 e^{-\lambda(m-1)/2}. \quad (39)$$

Proof. We focus on the case that G is the building block obtained by forming two cliques of size m and connecting m edges between them; the case that further edges are present is handled similarly using (16).

From (16) and (21), we have for any (i, j) within either of the two m -cliques that

$$\mathbb{P}_G[X_i = X_j] \geq 1 - \frac{\bar{m}}{\bar{m} + e^{\bar{m}\lambda/2}} \quad (40)$$

$$\geq 1 - \frac{m}{m + e^{\bar{m}\lambda/2}}, \quad (41)$$

where $\bar{m} := m - 1$. By taking an arbitrary node from each clique and applying the union bound over the $2(m-1) \leq 2m$ events corresponding to other nodes in the clique having the same value as that node, we find that the probability that each of the cliques have nodes that all take the same value satisfies the following:

$$\mathbb{P}_G[\text{all nodes same within each clique}] \geq 1 - \frac{2m^2}{m + e^{\bar{m}\lambda/2}}. \quad (42)$$

Next, we consider the probabilities of the two cliques taking a common value vs. two different values. Letting $A_{\nu, \sigma}$ be the event that the ν -th clique has values all equal to $\sigma \in \{+1, -1\}$, we have from (1) that

$$\mathbb{P}_G[A_{1,+} \cap A_{2,+}] = \frac{1}{Z} \exp\left(\lambda \left(2 \binom{m}{2} + m\right)\right) \quad (43)$$

$$\mathbb{P}_G[A_{1,+} \cap A_{2,-}] = \frac{1}{Z} \exp\left(\lambda \left(2 \binom{m}{2} - m\right)\right). \quad (44)$$

Taking the ratio between the two gives

$$\frac{\mathbb{P}_G[A_{1,+} \cap A_{2,+}]}{\mathbb{P}_G[A_{1,+} \cap A_{2,-}]} = e^{2m\lambda}. \quad (45)$$

By the same argument, this is also the ratio between any analogous events with the same signs in the numerator and differing signs in the denominator. The same argument also applies when we condition on each of the two cliques having common-valued nodes; in this case, the left-hand side of (45) simply amounts to $\frac{\psi}{1-\psi}$, where ψ is the conditional probability that all of the $2m$ nodes making up the two cliques take the same value. Equating $\frac{\psi}{1-\psi} = e^{2m\lambda}$ in accordance with (45) and solving for ψ , we obtain the following:

$$\begin{aligned} \mathbb{P}_G[\text{all nodes same} | \text{all nodes same within each clique}] \\ = 1 - \frac{1}{1 + e^{2m\lambda}}, \end{aligned} \quad (46)$$

where ‘‘all nodes’’ refers to the $2m$ nodes making up the two cliques. Multiplying this with (42) gives

$$\mathbb{P}_G[\text{all nodes same}] \geq 1 - \frac{2m^2}{m + e^{\bar{m}\lambda/2}} - \frac{1}{1 + e^{2m\lambda}}. \quad (47)$$

Using this fact along with (13), we have for all (i, j) , even in different cliques, that

$$\mathbb{E}_G[X_i X_j] \geq 1 - \frac{4m^2}{m + e^{\bar{m}\lambda/2}} - \frac{2}{1 + e^{2m\lambda}}. \quad (48)$$

Finally, the number of edges that are in the complete graph G' but not in G is trivially upper bounded by $\binom{2m}{2} \leq 2m^2$, and thus substitution into (15) yields

$$D(P_G \| P_{G'}) \leq 2\lambda m^2 \left(\frac{4m^2}{m + e^{\lambda(m-1)/2}} + \frac{2}{1 + e^{2\lambda m}} \right). \quad (49)$$

The proof is concluded by writing

$$\frac{4m^2}{m + e^{\lambda(m-1)/2}} + \frac{2}{1 + e^{2\lambda m}} \leq \frac{4m^2}{e^{\lambda(m-1)/2}} + \frac{2}{e^{2\lambda m}} \quad (50)$$

$$\leq \frac{6m^2}{e^{\lambda(m-1)/2}}. \quad (51)$$

□

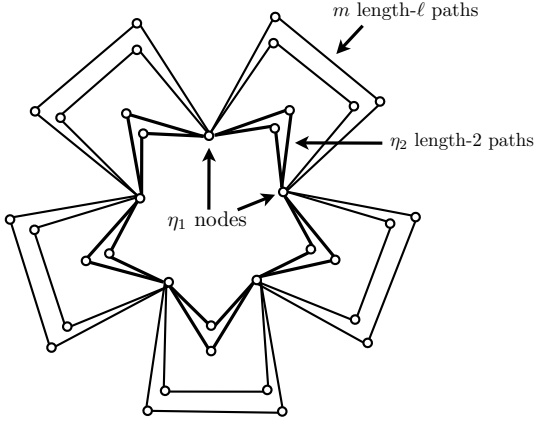


Figure 2: Building block for Ensemble 4 with $\eta_1 = 5$, $\eta_2 = 2$, $m = 2$, and $\ell = 3$.

Remark 1. In this ensemble, there are αm^2 edges known with certainty, and a possible further $\alpha m(m-1)$ that are unknown. Thus, slightly more than half of the potential edges are known. This limits the values of q_{\max} that are meaningful when applying this ensemble, and is the reason for the constraints on q_{\max} (e.g., $q_{\max} \leq k/4$) in Theorems 1–4. However, one can generalize this ensemble by considering *more than two* groups of m -cliques such that each pair has m inter-clique connections. With this extension, the fraction of potential edges that are known can be made arbitrarily close to zero, and similar results to those shown in Table I for \mathcal{G}_k (respectively, $\mathcal{G}_{k,d}$) can be obtained even when $q_{\max} = \lfloor \theta \frac{k}{2} \rfloor$ (respectively, $q_{\max} = \lfloor \theta \frac{k}{2} \frac{d-2}{d} \rfloor$) for some $\theta \in (0, 1)$.

D. Ensemble 4: Many Node-Disjoint Paths

This ensemble is based on forming a large number of node-disjoint paths between pairs of nodes, making it difficult to determine whether or not *direct* edges also exist between those nodes [10]. It is constructed as follows with integer parameters $\eta_1, \eta_2, m, \ell, \alpha$:

Ensemble4($\eta_1, \eta_2, m, \ell, \alpha$) [Disjoint paths ensemble]:

- Take an arbitrary subset of the p vertices of size η_1 and label them $1, 2, \dots, \eta_1$. For each consecutive pair of these nodes, including the wrapped-around pair $(\eta_1, 1)$, form η_2 node-disjoint paths of length two between them, and also form m node-disjoint paths of length ℓ between them.
- Form a base graph G_0 by taking α copies of this graph.
- Each graph in \mathcal{T} is formed by taking G_0 and adding arbitrarily many edges among the η_1 “center” nodes of each building block. Thus, G_0 itself has the fewest edges, whereas the graph with $\alpha \binom{\eta_1}{2}$ additional center edges contains the most edges.

An illustration of one building block is shown in Figure 2. For this ensemble, have the following:

- The number of nodes within each building block is $\eta_1(1 + \eta_2 + m(\ell - 1))$, and hence the total number of nodes is $\alpha \eta_1(1 + \eta_2 + m(\ell - 1))$.

- Within each building block, there are up to $\binom{\eta_1}{2}$ edges in the center, as well as $2\eta_1\eta_2$ further edges forming paths of length two, and $m\eta_1\ell$ edges forming paths of length ℓ . Hence, the total number of edges is between $\alpha\eta_1(2\eta_2 + m\ell)$ and $\alpha\eta_1((\eta_1 - 1)/2 + 2\eta_2 + m\ell)$.
- The total number of graphs is $|\mathcal{T}| = 2^{\alpha \binom{\eta_1}{2}}$.
- The maximal degree is less than $\eta_1 + 2\eta_2 + 2m$.
- Similarly to Ensembles 2 and 3, we may set $\mathcal{T}' = \mathcal{T}$.
- The number of graphs within an edit distance q_{\max} of any given graph is $A(q_{\max}) = \sum_{q=0}^{q_{\max}} \binom{\alpha \binom{\eta_1}{2}}{q} \leq 1 + q_{\max} \binom{\alpha \binom{\eta_1}{2}}{q_{\max}}$, assuming $q_{\max} \leq \frac{1}{2} \alpha \binom{\eta_1}{2}$.
- Using Lemma 6 below, along with (17), the KL divergence from any graph in \mathcal{T} to the corresponding graph with all centers connected is upper bounded by $\epsilon = \frac{2\lambda\alpha\eta_1 \binom{\eta_1}{2}}{1 + (\cosh(2\lambda))^{\eta_2} \left(\frac{1 + (\tanh \lambda)^\ell}{1 - (\tanh \lambda)^\ell} \right)^m}$.

Substituting these into (12) and setting $q_{\max} = \lfloor \theta_4 \alpha \binom{\eta_1}{2} \rfloor$ for some $\theta_4 \in (0, \frac{1}{2})$ gives

$$n \geq \frac{1 + (\cosh(2\lambda))^{\eta_2} \left(\frac{1 + (\tanh \lambda)^\ell}{1 - (\tanh \lambda)^\ell} \right)^m}{2\lambda\eta_1} (\log 2 - H_2(\theta_4)) \times (1 - \delta - o(1)) \quad (52)$$

provided that $\alpha \binom{\eta_1}{2} \rightarrow \infty$.

It remains to prove the claim on the KL divergence, formalized as follows.

Lemma 6. *Let G denote the graph corresponding to a single group in the construction in Ensemble 4, and let G' be the corresponding building block with all of the center nodes connected. Then*

$$D(P_G \| P_{G'}) \leq \frac{2\lambda\eta_1 \binom{\eta_1}{2}}{1 + (\cosh(2\lambda))^{\eta_2} \left(\frac{1 + (\tanh \lambda)^\ell}{1 - (\tanh \lambda)^\ell} \right)^m}. \quad (53)$$

Proof. We focus on the case that G is the building block described above; the case that further edges are present is handled similarly using (16).

We know from (25) that the joint distribution between any two consecutive nodes in the center satisfies

$$\mathbb{E}_G[X_i X_j] \geq 1 - \frac{2}{1 + (\cosh(2\lambda))^{\eta_2} \left(\frac{1 + (\tanh \lambda)^\ell}{1 - (\tanh \lambda)^\ell} \right)^m}, \quad (54)$$

since $\frac{1 + (\tanh \lambda)^\ell}{1 - (\tanh \lambda)^\ell} = \cosh(2\lambda)$. Using (13), this implies

$$\mathbb{P}_G[X_i = X_j] \geq 1 - \frac{1}{1 + (\cosh(2\lambda))^{\eta_2} \left(\frac{1 + (\tanh \lambda)^\ell}{1 - (\tanh \lambda)^\ell} \right)^m}. \quad (55)$$

Thus, by applying the union bound over (i, j) pairs of the form $(1, 2), (2, 3), \dots, (\eta_1 - 1, \eta_1), (\eta_1, 1)$, the probability that all η_1 of the center nodes take the same value satisfies

$$\begin{aligned} \mathbb{P}_G[\text{all center nodes same}] \\ \geq 1 - \frac{\eta_1}{1 + (\cosh(2\lambda))^{\eta_2} \left(\frac{1 + (\tanh \lambda)^\ell}{1 - (\tanh \lambda)^\ell} \right)^m}. \end{aligned} \quad (56)$$

Again using (13), this implies for any pair of center nodes (i, j) , including non-adjacent pairs, that

$$\mathbb{E}_G[X_i X_j] \geq 1 - \frac{2\eta_1}{1 + (\cosh(2\lambda))^{\eta_2} \left(\frac{1 + (\tanh \lambda)^\ell}{1 - (\tanh \lambda)^\ell} \right)^m}. \quad (57)$$

Observing that the corresponding edge sets E and E' satisfy $|E' \setminus E| \leq \binom{n_1}{2}$ and $|E \setminus E'| = 0$, (53) follows from (15). \square

V. APPLICATIONS TO GRAPH FAMILIES

Finally, we prove our main results by applying the ensembles from the previous section to the graph families introduced in Section I-A. All of the necessary conditions on n stated in this section are those needed to obtain $P_e(q_{\max}) \leq \delta$, where the graph class defining $P_e(\cdot)$ will be clear from the context.

A. Proofs of Theorems 1–2: Bounded Edges Ensemble

For the class \mathcal{G}_k of graphs with at most k edges, we have the following:

- If $k \leq p/4$, then using Ensemble 1 with $\alpha = k$, we obtain from (28) that

$$n \geq \frac{2(1 - \theta_1) \log p}{\lambda \tanh \lambda} (1 - \delta - o(1)) \quad (58)$$

provided that $q_{\max} \leq \lfloor \theta_1 k \rfloor$ for some $\theta_1 \in (0, 1)$.

- If $k = \lfloor cp^{1+\nu} \rfloor$ for some $c > 0$ and $\nu \in [0, 1)$, then we use Ensemble 2 with $m = \lfloor 2cp^\nu \rfloor$ and $\alpha = \lfloor p/m \rfloor = \frac{1}{2c} p^{1-\nu} (1 + o(1))$, chosen so that $m\alpha \leq p$ nodes are used in the construction. The number of possible edges is $\alpha \binom{m}{2} \leq \frac{1}{2} \alpha m^2 \leq \frac{1}{2} pm \leq cp^{1+\nu}$, as desired. We obtain from (29) that

$$n \geq \frac{\log 2 - H_2(\theta_2)}{\lambda \frac{e^{2\lambda} \cosh(2\lambda cp^\nu) - 1}{e^{2\lambda} \cosh(2\lambda cp^\nu) + 1}} (1 - \delta - o(1)) \quad (59)$$

provided that $q_{\max} \leq \lfloor \theta_2 \alpha \binom{m}{2} \rfloor$ for some $\theta_2 \in (0, \frac{1}{2})$. Substituting the choices of m and α into the latter expression, we find that q_{\max} can be as large as $\theta_2 k (1 + o(1))$.

- We use Ensemble 3 with $\alpha = 1$ and $m = \lfloor \sqrt{k/2} \rfloor$, chosen so that the number of edges does not exceed $2\alpha m^2 \leq k$. With these choices, we obtain from (38), along with the identity $\lfloor m \rfloor \geq m - 1$, that

$$n \geq \frac{e^{\lambda(\sqrt{k/2}-2)/2} (\log 2 - H_2(\theta_3))}{6\lambda k} (1 - \delta - o(1)), \quad (60)$$

provided that $q_{\max} \leq \lfloor \theta_3 \alpha m(m-1) \rfloor$ for some $\theta_3 \in (0, \frac{1}{2})$. Substituting the choices of m and α into the latter expression, we find that q_{\max} can be as large as $\theta_3 \frac{k}{2} (1 + o(1))$, provided that $k \rightarrow \infty$. Note that this construction uses $2m\alpha \leq \sqrt{2k}$ nodes, which is asymptotically less than p since $k = o(p^2)$.

We obtain Theorem 1 from (58) and (60), and Theorem 2 from (59) and (60). Specifically, we set $q_{\max} = \lfloor \theta k \rfloor$ for some $\theta \in (0, \frac{1}{4})$, and by equating this with the above upper bounds on q_{\max} we see that we may set $\theta_1 = \theta$, $\theta_2 = \theta(1 + o(1))$ and $\theta_3 = 2\theta(1 + o(1))$.

B. Proofs of Theorems 3–4: Bounded Degree Ensemble

For the class $\mathcal{G}_{k,d}$ of graphs such that every node has degree at most d , and the total number of edges does not exceed k , we have the following:

- If $k \leq p/4$, then using Ensemble 1 with $\alpha = k$, we obtain from (28) that

$$n \geq \frac{2(1 - \theta_1) \log p}{\lambda \tanh \lambda} (1 - \delta - o(1)), \quad (61)$$

provided that $q_{\max} \leq \lfloor \theta_1 k \rfloor$ for some $\theta_1 \in (0, 1)$.

- In the case that $k = \Omega(p)$, we use Ensemble 2 with the following parameters:

- 1) $m = d' \leq d$, chosen so that the maximal degree $m - 1$ does not exceed d ;
- 2) $\alpha = \lfloor k / \binom{d'}{2} \rfloor$, chosen so that the number of edges $\alpha \binom{m}{2}$ does not exceed k .

With these choices, we obtain from (29) that

$$n \geq \frac{\log 2 - H_2(\theta_2)}{\lambda \frac{e^{2\lambda} \cosh(2\lambda d') - 1}{e^{2\lambda} \cosh(2\lambda d') + 1}} (1 - \delta - o(1)), \quad (62)$$

whenever $q_{\max} \leq \lfloor \theta_2 \alpha \binom{d'}{2} \rfloor$ for some $\theta_2 \in (0, \frac{1}{2})$. Substituting the choice of α , we find that q_{\max} can be as large as $\theta_2 k (1 + o(1))$. Note also that the number of nodes used is upper bounded as $\alpha m \leq \frac{k}{\binom{d'}{2}} d' = \frac{2k}{d'-1}$, which is upper bounded by p provided that $k \leq \frac{1}{2} p (d' - 1)$.

- We use Ensemble 3 with the following parameters:
 - 1) $m = \lceil d/2 \rceil$, chosen so that each block has nodes with degree not exceeding $2m - 1 \leq d$;
 - 2) $\alpha = \lfloor \frac{k}{\binom{2m}{2}} \rfloor$, chosen to ensure that the number of edges does not exceed $\alpha \binom{2m}{2} \leq k$.

With these choices, we obtain from (38) that

$$n \geq \frac{e^{\lambda(d-2)/4} (\log 2 - H_2(\theta_3))}{3\lambda d^2} (1 - \delta - o(1)), \quad (63)$$

when $q_{\max} \leq \lfloor \theta_3 \alpha m(m-1) \rfloor$ for some $\theta_3 \in (0, \frac{1}{2})$. Substituting the choice of α to obtain $\alpha m(m-1) = k \frac{m-1}{2m-1} (1 + o(1))$, and then writing $d/2 \leq m \leq (d+1)/2$, we find that the latter condition holds provided that $q_{\max} \leq \frac{d/2-1}{d} \theta_3 k (1 + o(1))$. The number of nodes used is $2m\alpha \leq \frac{2mk}{m(2m-1)} = \frac{2k}{d-1} \leq \frac{2k}{d-1}$, which is upper bounded by p provided that $k \leq \frac{1}{2} p (d-1)$.

We obtain Theorem 3 from (61) and (63), and Theorem 4 from (62) and (63). Similarly to the previous subsection, we set $\theta_1 = \theta$, $\theta_2 = \theta(1 + o(1))$, and $\theta_3 = \frac{d}{d-2} \cdot 2\theta(1 + o(1))$.

C. Proofs of Theorem 5: Sparse Separator Ensemble

For the class $\mathcal{G}_{k,d,\eta,\gamma}$ (cf. Section II-C), we have the following:

- If $k \leq p/4$, then again using Ensemble 1 with $\alpha = k$, we obtain from (27) that

$$n \geq \frac{2(k - q_{\max}) \log p}{k\lambda \tanh \lambda} (1 - \delta - o(1)). \quad (64)$$

- We use Ensemble 4 with the following parameters:
 - 1) $\eta_1 = \lfloor c\eta \rfloor$ and $\eta_2 = \lfloor (1-c)\eta \rfloor$ for some $c \in (\frac{1}{\eta}, 1]$, thus ensuring that $\eta_1 \geq 1$;
 - 2) $\ell = \gamma + 1$, chosen to ensure that the (η, γ) -separation condition is satisfied;
 - 3) $m \leq d/2 - \eta$, chosen so that the maximal degree is upper bounded by $\eta_1 + 2\eta_2 + 2m \leq 2\eta + 2m \leq d$;

- 4) $\alpha = \lfloor \frac{k}{c\eta(c\eta/2+2(1-c)\eta+m(\gamma+1))} \rfloor$, chosen to ensure the total number of edges $\alpha\eta_1((\eta_1-1)/2+2\eta_2+m\ell)$ does not exceed k .

With these choices, we obtain from (52) that

$$n \geq \frac{1 + (\cosh(2\lambda))^{(1-c)\eta-1} \left(\frac{1+(\tanh \lambda)^{\gamma+1}}{1-(\tanh \lambda)^{\gamma+1}} \right)^m}{2\lambda c\eta} \times (\log 2 - H_2(\theta_4))(1 - \delta - o(1)) \quad (65)$$

provided that $q_{\max} \leq \lfloor \theta_4 \alpha (c\eta - 1)^2 / 2 \rfloor$ for some $\theta_4 \in (0, \frac{1}{2})$. Here we have used $\zeta - 1 \leq \lfloor \zeta \rfloor \leq \zeta$ and $\binom{c\eta}{2} \geq (c\eta - 1)^2 / 2$. Note that the graph in this ensemble with the most edges has at least as many edges as nodes, since each node is connected to at least two edges. Thus, since we have assumed $k \leq p/4$ and we have already chosen the parameters to ensure there are at most k edges, we have also ensured that less than p nodes are used. Substituting the above choice of α into the upper bound on q_{\max} , we find that q_{\max} can be as large as

$$\left\lfloor \theta_4 \frac{(c\eta - 1)^2 k}{2c\eta(c\eta/2 + 2(1-c)\eta + m(\gamma + 1))} \right\rfloor \geq \left\lfloor \theta_4 \frac{(c\eta - 1)^2 k}{2c\eta(2\eta + m(\gamma + 1))} \right\rfloor, \quad (66)$$

since $c\eta/2 + 2(1-c)\eta \leq 2\eta$ for $c \in [0, 1]$.

We obtain Theorem 5 by combining (64), (65) and (66), and renaming θ_4 as θ .

VI. NUMERICAL RESULTS

In this section, we simulate the graph learning problem for some of the ensembles presented in Section IV, as well as the analogous ensembles used for exact recovery in [9], [10]. Before proceeding, we discuss the optimal decoding techniques for the two recovery criteria.

Suppose that the graph G is uniformly drawn from some class \mathcal{G} . In the case of exact recovery, the optimal decoder is the maximum-likelihood (ML) rule

$$\hat{G} = \arg \max_{G \in \mathcal{G}} \mathbb{P}_G[\mathbf{X}], \quad (67)$$

where $\mathbb{P}_G[\mathbf{X}]$ is the probability of observing the samples $\mathbf{X} \in \{0, 1\}^{n \times p}$ when the true graph is G . In contrast, the optimal rule for approximate recovery is given by

$$\hat{G} = \arg \max_{G \in \mathcal{G}} \sum_{G' : |E \Delta E'| \leq q_{\max}} \mathbb{P}_{G'}[\mathbf{X}], \quad (68)$$

where E and E' are the edge sets of G and G' respectively. Both (67) and (68) are, in general, computationally intractable, requiring a search over the entire space \mathcal{G} . However, in the examples below, we are able to apply (67) by using various tricks such as symmetry arguments. While we need to consider relatively small graph sizes for Ensembles 3 and 4, these will still be adequate for generating results that support the theory.

Unfortunately, we found the implementation of (68) much more difficult, and we therefore also use (67) for approximate recovery even though, in general, it is only optimal for exact recovery. Nevertheless, even with approximate recovery, we

expect ML to provide a benchmark that that is unlikely to be beaten by any practical methods.

In all of the experiments, the error probabilities are obtained by evaluating the empirical average over 5000 trials.

A. A Variant of Ensemble 1 and a Counterpart from [9]

It was shown in [9] that if one considers all graphs with a single edge, then it is difficult to distinguish each of these from the empty graph if λ is small, thus making exact recovery difficult. In Figure 3, we simulate the performance of this ensemble with $p = 100$. Since the partition function Z (see (1)) is the same for all graphs in this ensemble, the ML rule (67) simply amounts to declaring the single edge to be the pair (i, j) among the $\binom{p}{2}$ possibilities such that $X_i = X_j$ in the highest number of samples.

Our Ensemble 1 is analogous to the single-edge ensemble from [9]; however, in order to facilitate the computation, we consider a slight variant defined as follows:

Ensemble1a(α) [Isolated edges ensemble]:

- Group the p vertices into $p/2$ fixed pairs in an arbitrary manner.
- Each graph in \mathcal{T} is obtained by connecting exactly α of those $p/2$ pairs.

Note that Ensemble 1a can be interpreted as a genie-aided version of Ensemble 1, where the decoder is given information narrowing the $\prod_{i=0}^{\alpha} \binom{p-2i}{2}$ possible graphs down to a smaller set of size $\binom{p/2}{\alpha}$. For this reason, the performance under Ensemble 1a is an optimistic estimate of the performance under Ensemble 1, and moving to the latter should only narrow the gaps seen in our comparisons to [9].

Figure 3 plots the approximate recovery error probability for Ensemble 1a with $p = 100$ and $\alpha = 12$, setting $q_{\max} = 3$ so that up to a quarter of the edges may be in error. The maximum-likelihood rule (67) is simple to implement: Since all graphs have the same partition function, the most likely graph corresponds to choosing the α edges among the $p/2$ potential edges, such that the corresponding pairs of nodes agree in as many observations as possible. This can be implemented by simply counting the number agreements of the $p/2$ pairs and then sorting.

In accordance with our theory, the general behavior of the error probability as a function of n is similar for Ensemble 1a (approximate recovery) and the ensemble from [9] (exact recovery). Moving to approximate recovery does provide some benefit, but it appears to be only in the constant factors. More specifically, across the range shown, the number of measurements required to achieve a given error probability in $[0.01, 0.5]$ differs for the two ensembles and recovery criteria only by a multiplicative factor in the range $[1, 2.2]$. In both cases, the learning problem becomes increasingly difficult as λ becomes smaller, since the edges are weaker and therefore more difficult to detect.

B. Ensemble 3 and a Counterpart from [9]

A counterpart to Ensemble 3 from [9] considers the $\binom{m'}{2}$ possible graphs on m' nodes obtained by removing a single

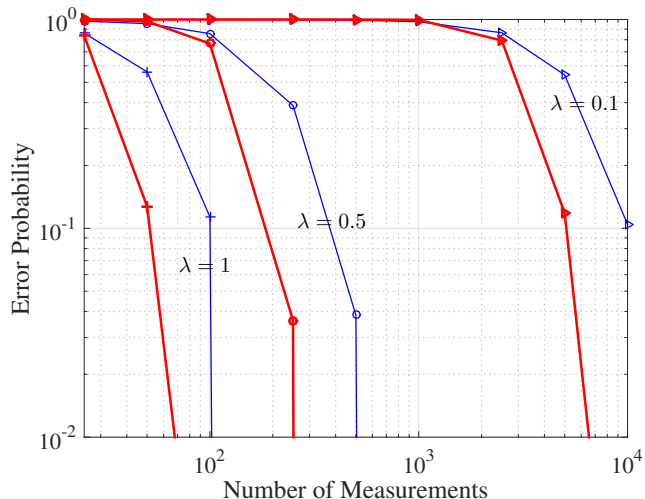


Figure 3: Empirical performance for Ensemble 1a (approximate recovery; red bold) and its counterpart from [9] (exact recovery; blue non-bold).

edge from the m' -clique. Thus, every graph is difficult to distinguish from the m' -clique, particularly as m' and λ increase, and exact recovery is difficult. In Figure 4, we plot the performance of this ensemble with $m' = 8$. In this case, ML decoding amounts to choosing the pair (i, j) such that $X_i \neq X_j$ in the highest number of samples.

For comparison, we consider Ensemble 3 with $m = 4$ and $\alpha = 1$, chosen so that the maximal number of edges and degree match those of the ensemble from [9] with $m' = 8$. We set $q_{\max} = 3$, so that up to a quarter of the 12 *unknown* edges may be in error. We perform ML decoding using a brute force search over the 2^{12} possible graphs.

Compared to the previous example, the gap between the curves for approximate recovery and exact recovery are more significant. This is because although both our results and those of [9] prove that the sample complexity is exponential in λm , the exponent in [9] is double that of ours. Intuitively, this is because we work with cliques of half the size. Despite this, the general behavior of our curves and those of [9] is similar, with the sample complexity rapidly growing large as λ increases due to higher correlations among the 8 nodes.

C. Ensemble 4 and a Counterpart from [10]

A counterpart to Ensemble 4 from [10] first constructs α disjoint building blocks, each of which connects two nodes (i, j) , and then forms η node-disjoint paths of length 2 between them. Each graph in the ensemble is then obtained by removing the direct edge from one of the α building blocks, while leaving the length-2 paths unchanged. We consider this construction with $\alpha = 4$ and $\eta = 8$, thus leading to the use of $p = 40$ nodes and $k = 68$ edges, and a maximal degree $d = 9$. Figure 5 plots the performance of the ML decoder, which amounts to counting the number agreements between the α pairs of “central” nodes (one per building block), and declaring the edge to be absent in the one with the most disagreements.

For comparison, we consider Ensemble 4 with $\eta_1 = 4$, $\eta_2 = 3$, $m = 0$ and $\alpha = 2$; this construction uses $p = 32$

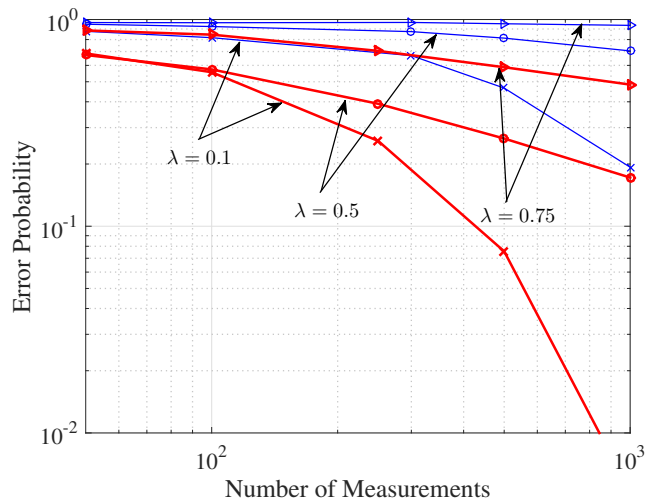


Figure 4: Empirical performance for Ensemble 3 (approximate recovery; red bold) and its counterpart from [9] (exact recovery; blue non-bold).

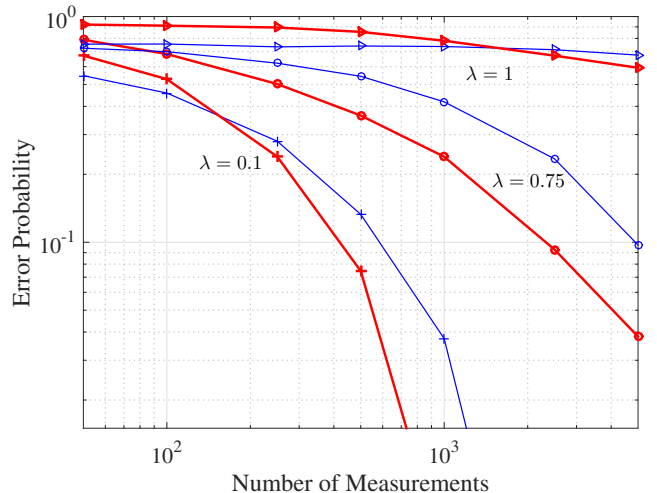


Figure 5: Empirical performance for Ensemble 4 (approximate recovery; red bold) and its counterpart from [10] (exact recovery; blue non-bold).

nodes and $k = 60$ edges, and has a maximal degree $d = 9$, thus being comparable to the above construction from [10]. We set $q_{\max} = 3$, so that up to a quarter of the 12 *unknown* edges may be in error. We perform ML decoding using a brute force search over the 2^{12} possible graphs, which simplifies to performing ML separately on the 2^6 possible graphs corresponding to each of the two building blocks.

Once again, we observe the same general behavior between our ensemble and that of [10]. While it may appear unusual that the exact recovery curves have a smaller error probability at low values of n , this occurs because even a random guess achieves an probability of exact recovery of $\frac{1}{4}$ for the ensemble in [10] with $\alpha = 4$. Despite this, we see that approximate recovery is easier for large n as expected, and that in both cases the recovery problem rapidly becomes more difficult as λ increases due to higher correlations among the nodes.

VII. CONCLUSION

We have provided information-theoretic lower bounds on Ising model selection with approximate recovery for a variety of graph classes. For a wide range of scaling regimes of the relevant parameters, we have obtained necessary conditions with the same scaling laws as the best known conditions for exact recovery, thus indicating that approximate recovery is not much easier in the minimax sense.

To this end, we presented a generalized form of Fano's inequality for handling approximate recovery, and applied it to a variety of graph ensembles. These were broadly categorized into those where it is difficult to distinguish each graph from the empty graph, and those where it is difficult to determine which edges between highly-correlated groups of nodes are present. In both cases, we required a departure from the ensembles considered for exact recovery [9], [10] in which the graphs differ in only one or two edges.

It would be interesting to determine to what extent approximate recovery can help when we move beyond the minimax performance criterion and the edit distance. For example, significant gains may be possible in the setting of random Ising model edge weights $\{\lambda_{ij}\}$, since it may become safe to "ignore" the weakest edges. As another example, since our analysis is based on constructing ensembles of graphs having a KL divergence that is close to a single graph, one may expect that under a recovery criterion based on $D(P_G \| P_{\hat{G}})$ being small, there is more to be gained. Other directions for further work include models beyond the Ising model (e.g., non-binary, Gaussian), and studies of achieving approximate recovery with *practical* algorithms.

APPENDIX A PROOF OF LEMMA 1

The proof follows standard steps in the derivation of Fano's inequality as in [10], but with suitable modifications to handle the approximate recovery criterion; see [31] for analogous modifications in the context of support recovery, and [25] for a related list decoding result. Due to the similarities to other variants, we focus primarily on the details that are specific to the approximate recovery criterion.

Let G be uniformly distributed on \mathcal{T} , let \hat{G} be the estimate of G , and let E and \hat{E} be the corresponding edge sets. Moreover, let $\bar{P}_e(q_{\max})$ be the error probability $\mathbb{P}[|E\Delta E'| > q_{\max}]$ averaged over the random graph G .

By assumption, we may consider decoders such that $\hat{G} \in \mathcal{T}'$ without loss of optimality. Defining the error indicator $\mathcal{E} := \mathbb{1}\{|E\Delta E'| > q_{\max}\}$ and applying the chain rule for entropy in two different ways, we have

$$H(\mathcal{E}, G|\hat{G}) = H(G|\hat{G}) + H(\mathcal{E}|G, \hat{G}) \quad (69)$$

$$= H(\mathcal{E}|\hat{G}) + H(G|\mathcal{E}, \hat{G}). \quad (70)$$

We have $H(\mathcal{E}|G, \hat{G}) = 0$ since \mathcal{E} is a function of (G, \hat{G}) , and

$H(\mathcal{E}|\hat{G}) \leq \log 2$ since \mathcal{E} is binary. Moreover, we have

$$\begin{aligned} H(G|\mathcal{E}, \hat{G}) &= (1 - \bar{P}_e(q_{\max}))H(G|\mathcal{E} = 0, \hat{G}) \\ &\quad + \bar{P}_e(q_{\max})H(G|\mathcal{E} = 1, \hat{G}) \quad (71) \\ &\leq (1 - \bar{P}_e(q_{\max})) \log A(q_{\max}) + \bar{P}_e(q_{\max}) \log |\mathcal{T}|, \quad (72) \end{aligned}$$

where (72) follows from the definition of $A(d_{\max})$ in the lemma statement and the fact that $\mathcal{E} = 0$ implies that G is within a distance q_{\max} of \hat{G} , and we have used the fact that the entropy is upper bounded by the logarithm of the number of elements of the support. We have now handled three of the terms in (69)–(70), and for the final one we write $H(G|\hat{G}) = -I(G; \hat{G}) + H(G) = -I(G; \hat{G}) + \log |\mathcal{T}|$, since G is uniform on \mathcal{T} .

Substituting the preceding observations into (69)–(70) and performing some simple rearrangements gives

$$\bar{P}_e(q_{\max}) \geq 1 - \frac{I(G; \hat{G}) + \log 2}{\log |\mathcal{T}| - \log A(q_{\max})}. \quad (73)$$

Finally, we bound the mutual information using the steps of [10], which are stated here without the details in order to avoid repetition: We use the data processing inequality to write $I(G; \hat{G}) \leq I(G; \mathbf{X})$, where \mathbf{X} contains the n independent samples from P_G . Using a covering argument, as well as the assumption containing G' in the lemma statement, it follows that $I(G; \mathbf{X}) \leq n\epsilon$. Substituting into (73), solving for n , and writing $P_e(q_{\max}) \geq \bar{P}_e(q_{\max})$, we obtain the desired result.

REFERENCES

- [1] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Trans. Patt. Analysis and Mach. Intel.*, no. 6, pp. 721–741, 1984.
- [2] R. J. Glauber, "Time-dependent statistics of the Ising model," *J. Math. Phys.*, vol. 4, no. 2, pp. 294–307, 1963.
- [3] R. Durbin, S. R. Eddy, A. Krogh, and G. Mitchison, *Biological sequence analysis: Probabilistic models of proteins and nucleic acids*. Cambridge Univ. Press, 1998.
- [4] C. D. Manning and H. Schütze, *Foundations of statistical natural language processing*. MIT press, 1999.
- [5] S. Wasserman and K. Faust, *Social network analysis: Methods and applications*. Cambridge Univ. Press, 1994, vol. 8.
- [6] D. M. Chickering, "Learning Bayesian networks is NP-complete," in *Learning from data*. Springer, 1996, pp. 121–130.
- [7] G. Reeves and M. Gastpar, "The sampling rate-distortion tradeoff for sparsity pattern recovery in compressed sensing," *IEEE Trans. Inf. Theory*, vol. 58, no. 5, pp. 3065–3092, May 2012.
- [8] J. Scarlett and V. Cevher, "Phase transitions in group testing," in *Proc. ACM-SIAM Symp. Disc. Alg. (SODA)*, 2016.
- [9] N. Santhanam and M. Wainwright, "Information-theoretic limits of selecting binary graphical models in high dimensions," *IEEE Trans. Inf. Theory*, vol. 58, no. 7, pp. 4117–4134, July 2012.
- [10] K. Shanmugam, R. Tandon, A. Dimakis, and P. Ravikumar, "On the information theoretic limits of learning Ising models," in *Adv. Neur. Inf. Proc. Sys. (NIPS)*, 2014.
- [11] E. Ising, "Beitrag zur theorie des ferromagnetismus," *Zeitschrift für Physik A Hadrons and Nuclei*, vol. 31, no. 1, pp. 253–258, 1925.
- [12] A. Anandkumar, V. Y. F. Tan, F. Huang, and A. S. Willsky, "High-dimensional structure estimation in Ising models: Local separation criterion," *Ann. Stats.*, vol. 40, no. 3, pp. 1346–1375, 2012.
- [13] —, "High-dimensional Gaussian graphical model selection: Walk summability and local separation criterion," *J. Mach. Learn. Res.*, vol. 13, pp. 2293–2337, 2012.
- [14] G. Bresler, E. Mossel, and A. Sly, "Reconstruction of Markov random fields from samples: Some observations and algorithms," in *Appr. Rand. and Comb. Opt. Algorithms and Techniques*. Springer Berlin Heidelberg, 2008, pp. 343–356.

- [15] R. Wu, R. Srikant, and J. Ni, "Learning loosely connected Markov random fields," *Stoch. Sys.*, vol. 3, no. 2, pp. 362–404, 2013.
- [16] A. Jalali, C. C. Johnson, and P. K. Ravikumar, "On learning discrete graphical models using greedy methods," in *Adv. Neur. Inf. Proc. Sys. (NIPS)*, 2011.
- [17] A. Ray, S. Sanghavi, and S. Shakkottai, "Greedy learning of graphical models with small girth," in *Allerton Conf. Comm., Control, and Comp.*, 2012.
- [18] G. Bresler, D. Gamarnik, and D. Shah, "Structure learning of antiferromagnetic Ising models," in *Adv. Neur. Inf. Proc. Sys. (NIPS)*, 2014.
- [19] G. Bresler, "Efficiently learning Ising models on arbitrary graphs," in *ACM Symp. Theory Comp. (STOC)*, 2015.
- [20] P. Ravikumar, M. J. Wainwright, J. D. Lafferty, and B. Yu, "High-dimensional Ising model selection using ℓ_1 -regularized logistic regression," *Ann. Stats.*, vol. 38, no. 3, pp. 1287–1319, 2010.
- [21] E. Yang, A. C. Lozano, and P. K. Ravikumar, "Elementary estimators for graphical models," in *Adv. Neur. Inf. Proc. Sys. (NIPS)*, 2014, pp. 2159–2167.
- [22] A. Montanari and J. A. Pereira, "Which graphical models are difficult to learn?" in *Adv. Neur. Inf. Proc. Sys. (NIPS)*, 2009.
- [23] R. Tandon and P. Ravikumar, "On the difficulty of learning power law graphical models," in *IEEE Int. Symp. Inf. Theory*, 2013.
- [24] A. K. Das, P. Netrapalli, S. Sanghavi, and S. Vishwanath, "Learning Markov graphs up to edit distance," in *IEEE Int. Symp. Inf. Theory*, 2012, pp. 2731–2735.
- [25] D. Vats and J. M. Moura, "Necessary conditions for consistent set-based graphical model selection," in *IEEE Int. Symp. Inf. Theory*, 2011, pp. 303–307.
- [26] N. Meinshausen and P. Bühlmann, "High-dimensional graphs and variable selection with the Lasso," *Ann. Stats.*, vol. 34, no. 3, pp. 1436–1462, June 2006.
- [27] W. Wang, M. Wainwright, and K. Ramchandran, "Information-theoretic bounds on model selection for Gaussian Markov random fields," in *IEEE Int. Symp. Inf. Theory*, 2010.
- [28] V. Jog and P.-L. Loh, "On model misspecification and KL separation for Gaussian graphical models," in *IEEE Int. Symp. Inf. Theory*, 2015.
- [29] P. Ravikumar, M. J. Wainwright, G. Raskutti, and B. Yu, "High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence," *Elec. J. Stats.*, vol. 5, pp. 935–980, 2011.
- [30] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley & Sons, Inc., 2006.
- [31] G. Reeves and M. Gastpar, "Approximate sparsity pattern recovery: Information-theoretic lower bounds," *IEEE Trans. Inf. Theory*, vol. 59, no. 6, pp. 3451–3465, June 2013.

Jonathan Scarlett (S'14 – M'15) received the B.Eng. degree in electrical engineering and the B.Sci. degree in computer science from the University of Melbourne, Australia. In 2011, he was a research assistant at the Department of Electrical & Electronic Engineering, University of Melbourne. From October 2011 to August 2014, he was a Ph.D. student in the Signal Processing and Communications Group at the University of Cambridge, United Kingdom. He is now a post-doctoral researcher with the Laboratory for Information and Inference Systems at the École Polytechnique Fédérale de Lausanne, Switzerland. His research interests are in the areas of information theory, signal processing, machine learning, and high-dimensional statistics. He received the Cambridge Australia Poynton International Scholarship, and the EPFL Fellows postdoctoral fellowship co-funded by Marie Curie.

Volkhan Cevher (SM'10) received the B.Sc. (valedictorian) in electrical engineering from Bilkent University in Ankara, Turkey, in 1999 and the Ph.D. in electrical and computer engineering from the Georgia Institute of Technology in Atlanta, GA in 2005. He was a Research Scientist with the University of Maryland, College Park from 2006-2007 and also with Rice University in Houston, TX, from 2008-2009. Currently, he is an Associate Professor at the Swiss Federal Institute of Technology Lausanne and a Faculty Fellow in the Electrical and Computer Engineering Department at Rice University. His research interests include signal processing theory, machine learning, convex optimization, and information theory. Dr. Cevher was the recipient of a Best Paper Award at SPARS in 2009, a Best Paper Award at CAMSAP in 2015, and an ERC StG in 2011.