

Estimating network edge probabilities by neighborhood smoothing

BY YUAN ZHANG

*Department of Statistics, Ohio State University
404 Cockins Hall, 1958 Neil Avenue, Columbus, Ohio, U.S.A
yzhanghf@stat.osu.edu*

ELIZAVETA LEVINA AND JI ZHU

*Department of Statistics, University of Michigan
311 West Hall, 1085 South University Avenue, Ann Arbor, Michigan, U.S.A.
elevina@umich.edu, jizhu@umich.edu*

SUMMARY

The estimation of probabilities of network edges from the observed adjacency matrix has important applications to predicting missing links and network denoising. It has usually been addressed by estimating the graphon, a function that determines the matrix of edge probabilities, but this is ill-defined without strong assumptions on the network structure. Here we propose a novel computationally efficient method, based on neighborhood smoothing to estimate the expectation of the adjacency matrix directly, without making the structural assumptions that graphon estimation requires. The neighborhood smoothing method requires little tuning, has a competitive mean-squared error rate, and outperforms many benchmark methods on link prediction in simulated and real networks.

Some key words: Graphon estimation; network analysis; nonparametric statistics.

1. INTRODUCTION

Statistical network analysis spans a wide range of disciplines, including network science, statistics, physics, computer science and sociology, and an equally wide range of applications and analysis tasks such as community detection and link prediction. In this paper, we study the problem of inferring the generative mechanism of an undirected network based on a single realization of the network. The data consist of the network adjacency matrix $A \in \{0, 1\}^{n \times n}$, where n is the number of nodes, and $A_{ij} = A_{ji} = 1$ if there is an edge between nodes i and j . We assume the observed adjacency matrix A is generated from an underlying probability matrix P , so that for $i \leq j$, A_{ij} 's are independent Bernoulli(P_{ij}) trials, and the P_{ij} are edge probabilities.

It is impossible to estimate P from a single realization of A unless one assumes some form of structure in P . When the network is expected to have communities, arguably the most popular assumption is that of the stochastic block model, where each node belongs to one of K blocks and the probability of an edge between two nodes is determined by the block to which the nodes belong. In this case, the $n \times n$ matrix P is parametrized by the $K \times K$ matrix of within- and between-block edge probabilities, and thus it is possible to estimate P from a single realization. The main challenge in fitting the stochastic block model is in estimating the blocks themselves, and that has been the focus of the literature, see for example Bickel & Chen (2009), Rohe et al. (2011), Amini et al. (2013), Saade et al. (2014) and Guédon & Vershynin (2016). Once the blocks are estimated, P can be estimated efficiently by a plug-in moment estimator. Many extensions

and alternatives to the stochastic block model have been proposed to model networks with communities, including those of Hoff (2008), Airoldi et al. (2008), Karrer & Newman (2011), Cai & Li (2015) and Zhang et al. (arXiv:1412.3432), but their properties are generally only known under the correctly specified model with communities. Here we are interested in estimating P for more general networks.

A general representation for the matrix P for unlabeled exchangeable networks goes back to Aldous (1981) and the 1979 preprint by D. N. Hoover entitled “Relations on probability spaces and arrays of random variables”. Formally, a network is exchangeable if for any permutation π of the set $\{1, \dots, n\}$, the distribution of edges remains the same. That is, if the adjacency matrix $A = [A_{ij}]$ is drawn from the probability matrix P , which we write as $A \sim P$, then for any permutation π ,

$$[A_{\pi(i)\pi(j)}] \sim P. \quad (1)$$

Aldous (1981) and Hoover showed that an exchangeable network always admits the following Aldous–Hoover representation:

DEFINITION 1. *For any network satisfying (1), there exists a function $f : [0, 1] \times [0, 1] \rightarrow [0, 1]$ and a set of independent and identically distributed random variables $\xi_i \sim \text{Uniform}[0, 1]$, such that*

$$P_{ij} = f(\xi_i, \xi_j). \quad (2)$$

Following the literature, we call f the graphon function. Unfortunately, as pointed out in Diaconis & Jason (arXiv 0712.2749), f in this representation is neither unique nor identifiable, since for any measure-preserving one-to-one transformation $\sigma : [0, 1] \rightarrow [0, 1]$, both $f\{\sigma(u), \sigma(v)\}$ and $f(u, v)$ yield the same distribution of A . An identifiable and unique canonical representation can be defined if one requires $g(u) = \int_0^1 f(u, v)dv$ to be non-decreasing (Bickel & Chen, 2009). Chan & Airoldi (2014) show that f and ξ_i ’s are jointly identifiable when $g(u)$, which can be interpreted as expected node degree, is strictly monotone. This assumption is strong and excludes the stochastic block model.

In practice, the main purpose of estimating f is to estimate P , and thus identifiability of f or lack thereof may not matter if P itself can be estimated. The preprint by Hoover and Diaconis & Jason (arXiv 0712.2749) showed that the measure-preserving map σ is the only source of non-identifiability. Wolfe and Olhede (arXiv:1309.5936) and Choi & Wolfe (2014) proposed estimating f up to a measure-preserving transformation σ via step-function approximations based on fitting the stochastic block model with a larger number of blocks K . This approximation does not assume that the network itself follows the block model, and some theoretical guarantees have been obtained under more general models. In related work, Olhede & Wolfe (2014) proposed to approximate the graphon with so-called network histograms, that is, stochastic block models with many blocks of equal size, akin to histogram bins. Another method to compute a network histogram was proposed by Amini & Levina (2017), as an application of their semi-definite programming approach to fitting block models with equal size blocks. Recently, Gao et al. (2015) established the minimax error rate for estimating P and proposed a least squares type estimator to achieve this rate, which obtains the estimated probability P by averaging the adjacency matrix elements within a given block partition. A similar estimator was proposed in Choi (2017), applicable also to non-smooth graphons. However, these methods are in principle computationally infeasible since they require an exhaustive enumeration of all possible block partitions. Cai et al. (arXiv:1412.2129) proposed an iterative algorithm to fit a stochastic blockmodel and approximate the graphon, but its error rate is unknown for general graphons. A Bayesian approach

using block priors proposed by Gao et al. (arXiv:1506.02174) achieves the minimax error rate adaptively, but it still requires the evaluation of the posterior likelihood over all possible block partitions to obtain the posterior mode or the expectation for the probability matrix.

Other recent efforts on graphon estimation focus on the case of monotone node degrees, which make the graphon identifiable. The sort and smooth methods of Yang et al. (2014) and Chan & Airoldi (2014) estimate the graphon under this assumption by first sorting nodes by their degrees and then smoothing the matrix A locally to estimate edge probabilities. The monotone degree assumption is crucial for the success of these methods, and as we later show, the sort and smooth methods perform poorly when it does not hold. Finally, general matrix denoising methods can be applied to this problem if one considers A to be a noisy version of its expectation P ; a good general representative of this class of methods is the universal singular value thresholding approach of Chatterjee (2015). Since this is a general method, we cannot expect its error rate to be especially competitive for this specific problem, and indeed its mean squared error rate is slower than the cubic root of the minimax rate.

In this paper, we propose a novel computationally efficient method for edge probability matrix estimation based on neighborhood smoothing, for piecewise Lipschitz graphon functions. The key to this method is adaptive neighborhood selection, which allows us to avoid making strong assumptions about the graphon. A node’s neighborhood consists of nodes with similar rows in the adjacency matrix, which intuitively correspond to nodes with similar values of the latent node positions ξ_i . To the best of our knowledge, our estimator achieves the best error rate among existing computationally feasible methods; it allows easy parallelization. The size of the neighborhood is controlled by a tuning parameter, similar to bandwidth in nonparametric regression; the rate of this bandwidth parameter is determined by theory, and we show empirically that the method is robust to the choice of the constant. Experiments on synthetic networks demonstrate that our method performs very well under a wide range of graphon models, including those of low rank and full rank, with and without monotone degrees. We also test its performance on the link prediction problem, using both synthetic and real networks.

2. THE NEIGHBORHOOD SMOOTHING ESTIMATOR AND ITS ERROR RATE

2.1. Neighborhood smoothing for edge probability estimation

Our goal is to estimate the probabilities P_{ij} from the observed network adjacency matrix A , where each A_{ij} is independently drawn from $\text{Bernoulli}(P_{ij})$. While $P_{ij} = f(\xi_i, \xi_j)$, where ξ_i ’s are latent, our goal is to estimate P for the single realization of ξ_i ’s that gave rise to the data, rather than the function f . We think of f as a fixed unknown smooth function on $[0, 1]^2$, with formal smoothness assumptions to be stated later. Let $e_{ij} = e_{ij}(P_{ij})$ denote the Bernoulli error and omit its dependence on P . We can then write

$$A_{ij} = P_{ij} + e_{ij} = f(\xi_i, \xi_j) + e_{ij}. \quad (3)$$

Formulation (3) resembles a nonparametric regression problem, except that the ξ_i are not observed. This has important consequences: for example, assuming further smoothness in f beyond order one does not improve the minimax error rate when estimating P (Gao et al., 2015). Our approach is to apply neighborhood smoothing, which would be natural had the latent variables ξ_i ’s been observed. Intuitively, if we had a set \mathcal{N}_i of neighbors of a node i , in the sense that $\mathcal{N}_i = \{i' : P_{i'} \approx P_i\}$, where P_i represents the i -th row of P , then we could estimate P_i by averaging $A_{i'}$ over $i' \in \mathcal{N}_i$. Postponing the question of how to select \mathcal{N}_i until Section 2.2, we

define a general neighborhood smoothing estimator by

$$\tilde{P}_{ij} = \frac{\sum_{i' \in \mathcal{N}_i} A_{i'j}}{|\mathcal{N}_i|}. \quad (4)$$

When the network is symmetric, we instead use a symmetric estimator

$$\hat{P} = (\tilde{P} + \tilde{P}^T) / 2. \quad (5)$$

For simplicity, we focus on undirected networks. A natural alternative is to average over $\mathcal{N}_i \times \mathcal{N}_j$, but (4) and (5) allow vectorization and are thus more computationally efficient. Our estimator can also be viewed as a relaxation of step function approximations such as Olhede & Wolfe (2014). In step function approximations, the neighborhood for each node is the nodes from its block, so the neighborhoods for two nodes from the same block are very similar, and the blocks have to be estimated first; in contrast, neighborhood smoothing provides for more flexible neighborhoods that differ from node to node, and an efficient way to select the neighborhood, which we will discuss next.

2.2. Neighborhood selection

Selecting the neighborhood \mathcal{N}_i in (5) is the core of our method. Since we estimate P_i by averaging over $A_{i'}$ for $i' \in \mathcal{N}_i$, good neighborhood candidates i' should have $f(\xi_{i'}, \cdot)$ close to $f(\xi_i, \cdot)$, which implies $P_{i'}$ close to P_i . We use the ℓ_2 distance between graphon slices to quantify this, defining

$$d(i, i') = \|f(\xi_i, \cdot) - f(\xi_{i'}, \cdot)\|_2 = \left\{ \int_0^1 |f(\xi_i, v) - f(\xi_{i'}, v)|^2 dv \right\}^{1/2}. \quad (6)$$

While one may consider more general ℓ_p or other distances, the ℓ_2 distance is particularly easy to work with theoretically. For the purpose of neighborhood selection, it is not necessary to estimate $d(i, i')$; it suffices to provide a tractable upper bound. For integrable functions g_1 and g_2 defined on $[0, 1]$, define $\langle g_1, g_2 \rangle = \int_0^1 g_1(u)g_2(u)du$. Then we can write

$$d^2(i, i') = \langle f(\xi_i, \cdot), f(\xi_i, \cdot) \rangle + \langle f(\xi_{i'}, \cdot), f(\xi_{i'}, \cdot) \rangle - 2\langle f(\xi_i, \cdot), f(\xi_{i'}, \cdot) \rangle. \quad (7)$$

The third term in (7) can be estimated by $2\langle A_i, A_{i'} \rangle / n$, where A_i and $A_{i'}$ are nearly independent up to a single duplicated entry due to symmetry. The first two terms in (7) are more difficult, since $\langle A_i, A_i \rangle / n$ is not a good estimator for $\langle f(\xi_i, \cdot), f(\xi_i, \cdot) \rangle$. Here we present the intuition and provide a full theoretical justification in Theorem 1. For simplicity, assume for now f is Lipschitz with a Lipschitz constant of 1. The idea is to use nodes with graphon slices similar to i and i' to make the terms in the inner product distinct graphon slices. With high probability, for each i , we can find $\tilde{i} \neq i$ such that $|\xi_{\tilde{i}} - \xi_i| \leq e_n$, where the sequence e_n is a function of n and represents the error rate to be specified later. Then $\|f(\xi_i, \cdot) - f(\xi_{\tilde{i}}, \cdot)\|_2 \leq e_n$, and we can approximate $\langle f(\xi_i, \cdot), f(\xi_i, \cdot) \rangle$ by $\langle f(\xi_i, \cdot), f(\xi_{\tilde{i}}, \cdot) \rangle$, where the latter can now be estimated by $\langle A_i, A_{\tilde{i}} \rangle / n$. The same technique can be used to approximate the second term in (7), but all these approximations depend on the unknown ξ 's. To deal with this, we rearrange the terms in (7) as follows:

$$\begin{aligned} d^2(i, i') &= \langle f(\xi_i, \cdot) - f(\xi_{i'}, \cdot), f(\xi_i, \cdot) \rangle - \langle f(\xi_i, \cdot) - f(\xi_{i'}, \cdot), f(\xi_{i'}, \cdot) \rangle \\ &\leq |\langle f(\xi_i, \cdot) - f(\xi_{i'}, \cdot), f(\xi_{\tilde{i}}, \cdot) \rangle| + |\langle f(\xi_i, \cdot) - f(\xi_{i'}, \cdot), f(\xi_{\tilde{i}'}', \cdot) \rangle| + 2e_n \\ &\leq 2 \max_{k \neq i, i'} |\langle f(\xi_i, \cdot) - f(\xi_{i'}, \cdot), f(\xi_k, \cdot) \rangle| + 2e_n. \end{aligned} \quad (8)$$

The inner product on the right side of (8) can be estimated by

$$\tilde{d}^2(i, i') = \max_{k \neq i, i'} |\langle A_{i \cdot} - A_{i' \cdot}, A_{k \cdot} \rangle| / n. \quad (9)$$

Intuitively, the neighborhood \mathcal{N}_i should consist of i 's with small $\tilde{d}(i, i')$. To formalize this, let $q_i(h)$ denote the h -th sample quantile of the set $\{\tilde{d}(i, i') : i' \neq i\}$, where h is a tuning parameter, and set

$$\mathcal{N}_i = \left\{ i' \neq i : \tilde{d}(i, i') \leq q_i(h) \right\} \quad (10)$$

where for notational simplicity we suppress the dependence of \mathcal{N}_i on h . Thresholding at a quantile rather than at some absolute value is convenient since real networks vary in their average node degrees and other parameters, which leads to very different values and distributions of \tilde{d} . Empirically, thresholding at a quantile shows significant advantage in stability and performance compared to an absolute threshold. The choice of h will be guided by both the theory in Section 2.3, which suggests the order of h , and empirical performance which suggests the constant factor. More details are included in the Supplementary Material.

An important feature of this definition is that the neighborhood admits nodes with similar graphon slices, but not necessarily similar ξ 's. For example, in the stochastic block model, all nodes from the same block would be equally likely to be included in each other's neighborhoods, regardless of their ξ 's. Even though we use ξ_i and $\xi_{i'}$ to motivate (8), we always work with the function values $f(\xi_i, \xi_j)$'s and never attempt to estimate the ξ_i or f by themselves. This contrasts with the approaches of Chan & Airolidi (2014) and Yang et al. (2014), and gives us a substantial computational advantage as well as much more flexibility in assumptions.

2.3. Consistency of the neighborhood smoothing estimator

We study the theoretical properties of our estimator for a family of piecewise Lipschitz graphon functions, defined as follows.

DEFINITION 2 (PIECEWISE LIPSCHITZ GRAPHON FAMILY). *For any $\delta, L > 0$, let $\mathcal{F}_{\delta;L}$ denote a family of piecewise Lipschitz graphon functions $f : [0, 1]^2 \rightarrow [0, 1]$ such that (i) there exists an integer $K \geq 1$ and a sequence $0 = x_0 < \dots < x_K = 1$ satisfying $\min_{0 \leq s \leq K-1} (x_{s+1} - x_s) \geq \delta$, and (ii) both $|f(u_1, v) - f(u_2, v)| \leq L|u_1 - u_2|$ and $|f(u, v_1) - f(u, v_2)| \leq L|v_1 - v_2|$ hold for all $u, u_1, u_2 \in [x_s, x_{s+1}]$, $v, v_1, v_2 \in [x_t, x_{t+1}]$ and $0 \leq s, t \leq K - 1$.*

For any $P, Q \in \mathbb{R}^{m \times m}$, define $d_{2,\infty}$, the normalized $2, \infty$ matrix norm, by

$$d_{2,\infty}(P, Q) = m^{-1/2} \|P - Q\|_{2,\infty} = \max_i m^{-1/2} \|P_{i \cdot} - Q_{i \cdot}\|_2.$$

Then we have the following error rate bound.

THEOREM 1. *Assume that L is a global constant and $\delta = \delta(n)$ depends on n , satisfying $\lim_{n \rightarrow \infty} \delta / (n^{-1} \log n)^{1/2} \rightarrow \infty$. Then the estimator \tilde{P} defined in (5), with neighborhood \mathcal{N}_i defined in (10) and $h = C(n^{-1} \log n)^{1/2}$ for any global constant $C \in (0, 1]$, satisfies*

$$\max_{f \in \mathcal{F}_{\delta;L}} \text{pr} \left\{ d_{2,\infty}(\tilde{P}, P)^2 \geq C_1 \left(\frac{\log n}{n} \right)^{1/2} \right\} \leq n^{-C_2} \quad (11)$$

where C_1 and C_2 are positive global constants.

Since for any $P, Q \in \mathbb{R}^{m \times m}$, we have $d_{2,\infty}(P, Q) \geq m^{-1} \|P - Q\|_F$, Theorem 1 yields

COROLLARY 1. *Under conditions of Theorem 1,*

$$\max_{f \in \mathcal{F}_{\delta;L}} \text{pr} \left\{ \frac{1}{n^2} \|\tilde{P} - P\|_F^2 \geq C_1 \left(\frac{\log n}{n} \right)^{1/2} \right\} \leq n^{-C_2}. \quad (12)$$

The bound (12) continues to hold if we replace \tilde{P} by \hat{P} , but (11) may not hold. Next, we show that under the $(2, \infty)$ norm, our estimator \hat{P} is nearly rate-optimal, up to a $\log n$ factor.

THEOREM 2. *Under conditions of Theorem 1, we have*

$$\inf_{\hat{P}} \sup_{f \in \mathcal{F}_{\delta;L}} E \left\{ d_{2,\infty}^2(\hat{P}, P) \right\} \geq C (n \log n)^{-1/2} \quad (13)$$

for some global constant $C > 0$.

To the best of our knowledge, result (11) is the only $(2, \infty)$ error rate available for polynomial time graphon estimation methods. Most previous work focused on the mean squared error and only considered the special case $\delta = 1$. For $\delta = 1$, the minimax error rate $\log n/n$ established by Gao et al. (2015) has so far only been achieved by methods that require combinatorial optimization or evaluation, including Gao et al. (2015) and Klopp et al. (2017). The rate $(\log n/n)^{1/2}$ was previously achieved by combinatorial methods, including Wolfe and Olhede (arXiv:1309.5936) and Olhede & Wolfe (2014). Among computationally efficient methods, singular value thresholding (Chatterjee (2015), Theorem 2.7) achieves $n^{-1/3}$. Additionally, the sort-and-smooth method proposed by Chan & Airoidi (2014) achieves the minimax error rate under the strong assumption that f has strictly monotone expected node degrees $d_f(v) = \int_0^1 f(u, v) du$. An anonymous referee of this manuscript sent us a proof that thresholding the leading k singular values of the matrix A achieves the mean squared error of $k/n + k^{-2}$, where the variance k/n is due to Candès & Plan (2011) and k^{-2} is the bias. Taking $k = n^{1/3}$ gives the best known mean squared error rate of $n^{-2/3}$ for a computationally efficient algorithm. For the graphon family $f \in \mathcal{F}_{\delta;L}$ where $\delta/(n^{-1} \log n)^{1/2} \rightarrow \infty$ that we study, the $n^{-1/3}$ singular value thresholding method and our method achieve the same mean squared error rate.

For the case of general δ , we can show that the minimax rate of $\log n/n$ established by Gao et al. (2015) still holds for the family $\mathcal{F}_{\delta;L}$, in Proposition 1; see the Supplementary Material.

PROPOSITION 1. *Under conditions of Theorem 1, when $\delta/(\log n/n)^{1/2} \rightarrow \infty$, there exists a global constant $C_3 > 0$ such that*

$$\inf_{\hat{P}} \max_{f \in \mathcal{F}_{\delta;L}} E \left\{ \frac{1}{n^2} \|\hat{P} - P\|_F^2 \right\} \asymp \frac{\log n}{n}.$$

Whether this minimax error rate can be achieved by a computationally efficient method remains an open question.

3. PROBABILITY MATRIX ESTIMATION ON SYNTHETIC NETWORKS

In this section we evaluate the performance of our symmetric estimator (5) on estimating the probability matrix for synthetic networks. We generate the networks from the four graphons listed in Table 1, selected to have different features in different combinations (monotone degrees, low rank, etc). The corresponding probability matrices are pictured in the first column of Figure 1 (lower triangular half). All networks have $n = 2000$ nodes.

Additional empirical results in the Supplementary Material show that our method is robust to the choice of the constant factor C in the bandwidth h , for simplicity, we set $C = 1$ for the rest

Table 1: Synthetic graphons

Graphon	Function $f(u, v)$	Monotone degrees	Rank	Local structure
1	$k/(K + 1)$ if $u, v \in ((k - 1)/K, k/K)$, $0.3/(K + 1)$ otherwise; $K = \lfloor \log n \rfloor$	Yes	$\lfloor \log n \rfloor$	No
2	$\sin \{5\pi(u + v - 1) + 1\} / 2 + 0.5$	No	3	No
3	$1 - \left[1 + \exp \{15(0.8 u - v)^{4/5} - 0.1\}\right]^{-1}$	No	Full	No
4	$(u^2 + v^2) / 3 \cos \{1 / (u^2 + v^2)\} + 0.15$	No	Full	Yes

of this paper. Here we focus on comparing to benchmarks. From the general matrix denoising methods, we include the widely used method of universal singular value thresholding (Chatterjee, 2015) and the $n^{1/3}$ leading singular value thresholding method suggested by a referee. We also compare to the sort and smooth methods of Chan & Airolidi (2014) and Yang et al. (2014). These methods differ only in that the latter one employs singular value thresholding to denoise the network as a pre-processing step. Due to space constraints, we present both methods in Table 2 but only Chan & Airolidi (2014) in figures, since they are visually very similar.

We also include two approximations based on fitting a stochastic block model, called network histograms by Olhede & Wolfe (2014). One is the oracle stochastic block model, where the blocks are based on the true values of the latent ξ_i 's. This cannot be done in practice, but we use it as the gold standard for a step-function approximation. The feasible version of this is an approximation based on a stochastic block model with estimated blocks; we fit it by regularized spectral clustering (Chaudhuri et al., 2012). Any other algorithm for fitting the stochastic blockmodel can be used to estimate the blocks; for example, Olhede & Wolfe (2014) used a local updating algorithm initialized with spectral clustering to compute their network histograms. Here we chose regularized spectral clustering because of its speed and good empirical performance. For both approximations, we set the number of blocks to $n^{1/2}$, as in Olhede & Wolfe (2014).

A recent as yet unpublished method kindly shared with us by E. Airolidi proposes a stochastic block model approximation, adapting the method of Airolidi et al. (2013) to work with a single adjacency matrix. It uses a dissimilarity measure $\sum_{k \neq i, i'} |\langle A_{i \cdot} - A_{k \cdot}, A_{k \cdot} \rangle|$, which we considered before choosing (9) because it leads to a better guaranteed error rate. Airolidi's method then builds blocks by starting with one not-yet-clustered node i and including all nodes whose dissimilarity from i is below a threshold Δ as neighbors. We found that our strategy of thresholding by quantile instead of a fixed threshold is more efficient and stable, and the theoretical error rate is better for our method.

We present the heatmaps of results for a single realization in Figure 1, and the root mean squared errors and the mean absolute errors of \hat{P} in Table 2. While these two errors mostly agree on method ranking, the few cases where they disagree indicate whether the errors are primarily coming from a small number of poorly estimated entries or are more uniformly distributed throughout the matrix.

Graphon 1 has $K = \lfloor \log n \rfloor = 7$ blocks with different within-block edge probabilities, which all dominate the low between-block probability. The best results are obtained by our method, singular value thresholding, spectral clustering, and the oracle stochastic blockmodel approximation, which is expected given that the data are generated from a stochastic block model. The oracle uses $n^{1/2}$ blocks rather than the true K , and thus makes substantial errors on the block boundaries, but not anywhere else. The method of Chan & Airolidi (2014) correctly estimates the main blocks because they have different expected degrees, but suffers from boundary effects due to smoothing over the entire region. In contrast, our method, which determines smoothing




Fig. 1: Estimated probability matrices for graphons 1–4, shown in rows 1–4. Column 1: true P (lower) and our method (upper). Column 2: Chan & Airol di (2014) (lower) and $n^{1/3}$ singular value thresholding (upper). Column 3: Block model oracle (lower) and spectral clustering (upper). Column 4: Chatterjee (2015) (lower) and Airol di’s method (upper).

neighborhoods based on similarities of graphon slices, does not suffer from boundary effects. Chatterjee (2015) does a good job on denser blocks but thresholds away sparser blocks. Airol di’s method captures tightly connected communities, but does not do as well on weaker communities.

Graphon 2 lacks node degree monotonicity, and thus the method of Chan & Airol di (2014) does not work here. Spectral clustering also performs poorly, likely because it uses too many ($n^{1/2}$) eigenvectors which add noise. Airol di’s method and the stochastic block model oracle give grainy but reasonable approximations to P , and the best results are obtained by our method, Chatterjee (2015), and singular value thresholding with $n^{1/3}$ eigenvalues. The latter two are expected to work well since this is a low-rank matrix.

Graphon 3 is a diagonal-dominated matrix, and our method is the best among computationally efficient methods. The method of Chatterjee (2015) does not perform well because this is not a low-rank matrix; spectral clustering, on the other hand, does fine, because there are many non-zero eigenvalues and the $n^{1/2}$ eigenvectors contain enough information. The $n^{1/3}$ singular value thresholding does better than Chatterjee (2015) and provides a lower-resolution denoising. Airol di’s method only roughly shows the structure, likely due to the similarity measure it uses. The method of Chan & Airol di (2014) fails since all node expected degrees are almost the same.

Graphon 4 is difficult to estimate for all methods. It is full rank, with structure at different scales. This makes it a difficult setting for low-rank approximations, among which the $n^{1/3}$ singular value thresholding alone uses enough eigenvalues to produce a reasonable result, albeit with boundary effects. This graphon is not a block matrix, and thus spectral clustering does not perform well. The expected node degrees are not the same, but their ordering does not match

the ordering of the latent node positions, so this graphon is also difficult for the sort-and-smooth method of Chan & Airol di (2014). Our method successfully picks up the global structure and the curvature. While visually it is fairly similar to the result of $n^{1/3}$ singular value thresholding, our method has significantly better errors in Table 2. Overall, this example illustrates a limitation of all global methods when there are subtle local differences.

Table 2 shows the mean squared errors and the mean absolute errors of all methods on the four graphons averaged over 2000 replications. The results generally agree with those shown in the figures. The few relative discrepancies between RMSE and L1 errors occur when there is a small number of large errors, such as the boundary effects for the oracle for graphon 1, which affect RMSE more than the L1 error.

For graphon 1, our method and the spectral clustering perform best. For graphon 2, our method is only outperformed by universal singular value thresholding, whereas $n^{1/3}$ leading eigenvalue thresholding selects fewer eigenvalues than needed. For graphon 3, our method is comparable to the $n^{1/3}$ leading eigenvalue thresholding, and they are both better than other methods, not counting the oracle. For graphon 4, our method has shows significant advantage over all other methods except for the oracle. Thus in all cases, our method shows very competitive performance compared to benchmarks.

Table 2: Root mean squared errors and mean absolute errors with standard errors, all multiplied by 10^2 , averaged over 2000 replications. The largest relative error is less than 4%.

	Graphon 1		Graphon 2		Graphon 3		Graphon 4	
	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE
Our method	1.92	1.33	3.06	2.25	3.00	1.41	3.55	2.76
Chan & Airol di (2014)	8.78	3.09	34.17	30.16	11.27	8.04	4.46	3.58
Yang et al. (2014)	9.56	4.14	34.18	30.19	11.47	8.59	5.67	4.87
$n^{1/3}$ singular value	2.99	2.25	4.74	3.59	3.16	1.79	5.86	4.33
Blockmodel spectral	1.72	0.75	33.06	28.80	3.98	1.78	9.08	6.64
Blockmodel oracle	5.48	1.42	5.11	3.80	1.62	0.75	1.06	0.83
Chatterjee (2015)	4.09	2.25	1.89	1.47	6.39	3.81	5.67	4.87
Airol di’s method	15.94	8.92	15.82	9.23	9.40	5.74	4.60	3.16

RMSE, root mean squared error; MAE, mean absolute error.

Overall, the results in this section show that various previously proposed methods can perform very well under their assumptions, which may be monotone degrees or low-rank or an underlying block model, but they fail when these assumptions are not satisfied. Our method is the only one among those compared that performs well in a large range of scenarios, because it learns the structure from data via neighborhood selection instead of imposing a priori structural assumptions. The $n^{1/3}$ singular value thresholding method also shows consistent performance across all graphons, although in all cases somewhat worse than ours. It performs very well in the low-rank case, but if the leading singular values decay slowly, our method performs better.

4. APPLICATION TO LINK PREDICTION

Evaluating probability matrix estimation methods on real networks directly is difficult, since the true probability matrix is unknown. We assess the practical utility of our method by applying it to link prediction, a task that relies on estimating the probability matrix. Here we think of the true adjacency matrix A^{true} as unobserved, with binary edges drawn independently with probabilities given by P , also unobserved. Instead we observe $A_{ij}^{\text{obs}} = M_{ij}A_{ij}^{\text{true}}$, where unobserved

M_{ij} 's are independent Bernoulli($1 - p$), and p is unknown. Therefore $A_{ij}^{\text{obs}} = 1$ is always a true edge, but $A_{ij}^{\text{obs}} = 0$ could be either a true 0 or a false negative. This setting is different from and perhaps more realistic than the link prediction setting in Gao et al. (2016), who assumed that M_{ij} 's are observed. Under their setting, the missing rate p can be estimated by the empirical missing rate \hat{p} , and all estimators can be corrected for missingness simply by dividing them by $1 - \hat{p}$.

A link prediction method usually outputs a nonnegative score matrix \hat{A} , with scores giving the estimated propensity of a node pair to form an edge. For methods that estimate the probability matrix, \hat{A} can be taken to be \hat{P} ; other link prediction methods construct a binary \hat{A} by working directly on A . Both types of methods essentially output a ranked list of most likely missing links, useful in practice for follow-up confirmatory analysis.

We compare various link prediction methods via their receiver operating characteristic curves. For each $t > 0$, we define the false positive and the true positive rates by

$$r_{\text{FP}}(t) = \sum_{ij} 1 \left(\hat{A}_{ij} > t, A_{ij}^{\text{true}} = 0, M_{ij} = 0 \right) / \sum_{ij} 1 \left(A_{ij}^{\text{true}} = 0, M_{ij} = 0 \right)$$

$$r_{\text{TP}}(t) = \sum_{ij} 1 \left(\hat{A}_{ij} > t, A_{ij}^{\text{true}} = 1, M_{ij} = 0 \right) / \sum_{ij} 1 \left(\hat{A}_{ij} = 1, M_{ij} = 0 \right).$$

Then by varying t we obtain the receiver operating characteristic curve. In practice, t is often selected to output a fixed number of most likely links.

In this section we include three additional benchmark methods that produce score matrices rather than estimated probability matrices. One standard score is the Jaccard index $\langle A_{i\cdot}, A_{j\cdot} \rangle / \{(\sum_k A_{ik})(\sum_k A_{jk})\}$, see for example Lichtenwalter et al. (2010). The method by Zhao et al. (2017) computes scores so that similar node pairs to have similar predicted scores. The PropFlow algorithm of Lichtenwalter et al. (2010) uses an expected random walk distance between nodes as the score.

We first compare all methods on simulated networks generated from the graphons in Table 1. We set $n = 500$ due to the computational cost of some of the benchmarks, and set $p = 10\%$. All experiments are repeated 1000 times/ Figure 2 in the Supplementary Material shows the receiver operating characteristic curves for four graphons. Most differences between the methods can be inferred from Figure 1. Overall, the methods based on graphon estimation outperform score-based methods. Our method outperforms all other methods on this task, producing a receiver operating characteristic curve very close to that based on the true probability matrix P .

We also applied our method and others to the political blogs network (Adamic & Glance, 2005). This network consists of 1222 manually labelled blogs, 586 liberal and 636 conservative. The network clearly shows two communities, with heterogeneous node degrees (there are hubs). We removed 10% of edges at random and calculated the receiver operating characteristic curve for predicting the missing links, shown in Figure 2. Again, methods based on estimating the probability matrix performed much better than the scoring methods, and our method performs best overall. Sort and smooth methods slightly outperformed spectral clustering and Chatterjee (2015), perhaps due to the presence of hubs.

5. DISCUSSION

The strength of our method is the adaptive neighborhood choice which works well under many different conditions; it is also computationally efficient, easy to implement, and essentially tuning free. Its main limitation is the piecewise Lipschitz condition, which occasionally leads

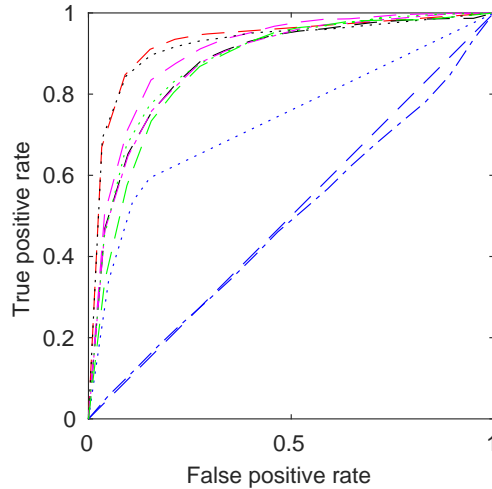


Fig. 2: Receiver operating characteristic curves for link prediction on the political blogs network. 10% of edges are missing at random. Red dashed curve: our method; black dotted curve: $n^{1/3}$ singular value thresholding; blue dashed curve: Zhao et al. (2017); blue dash dotted curve: the Jaccard index; blue dotted curve: PropFlow; black dashed curve: Chatterjee (2015); magenta dashed curve: Chan & Airoldi (2014); magenta dashed dotted curve: Yang et al. (2014); green dashed curve: block model with spectral clustering.

to over-smoothing. Our method does not achieve the minimax error rate, and its rate cannot be improved; whether the minimax rate can be achieved by any polynomial time method is, to the best of our knowledge, an open problem. Another major future challenge is relaxing the assumption of independent edges to better fit real-world networks.

ACKNOWLEDGMENTS

The authors thank an associate editor and two anonymous referees for very helpful suggestions, and E. M. Airoldi of Harvard University for sharing a copy of his manuscript and code and great comments. E.L. is supported by grants NSF DMS 1521551 and ONR N000141612910. J.Z. is supported by grants NSF DMS 1407698, KLAS 130026507, and KLAS 130028612.

SUPPLEMENTARY MATERIAL

The supplementary material includes numerical results on the bandwidth constant in Theorem 1, $(2, \infty)$ -norm errors and comparisons to benchmarks on link prediction for synthetic graphons from Section 3, and the proofs of Theorems 1 and 2 and Proposition 1.

6. CHOOSING THE CONSTANT FACTOR FOR THE BANDWIDTH

First, we need to choose the quantile cut-off parameter h which controls neighborhood selection. Theorem 1 gives the order of h , and the following numerical experiments empirically justify our choice of the constant factor. Figure 3 shows the mean squared error curves for networks with $n = 2000$ nodes generated from the four graphons in Table 1, with the constant factor C varying in the range $\{2^{-3}, 2^{-2}, \dots, 2^3\}$.

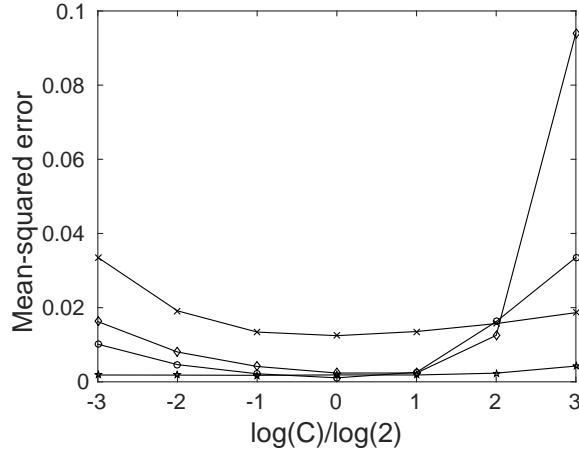


Fig. 3: Mean squared error of our method as a function of the constant C in the tuning parameter $h = C(n^{-1} \log n)^{1/2}$. Graphons 1–4 are marked with a circle, an asterisk, a plus and a cross, respectively.

Figure 3 demonstrates that C in the range from 2^{-2} to 2 works equally well for all these very different graphons. This suggests empirically that the method is robust to the choice of C , and therefore we set $C = 1$ in all numerical results in the paper.

7. RECEIVER OPERATING CHARACTERISTIC CURVES FOR LINK PREDICTION SIMULATIONS IN SECTION 4

The results are presented in Figure 4. The link prediction results generally agree with how the methods performed on the task of estimating the probability matrix for these synthetic graphons, shown in Section 3 of the main manuscript.

8. ROOT MEAN SQUARED $2, \infty$ ERRORS FOR SIMULATIONS IN SECTION 3

In Table 3, we report $\|\hat{P} - P\|_{2, \infty} / n^{1/2}$ for \hat{P} estimated by all methods in Section 3. The performances of our method is slightly better than but comparable to $n^{1/3}$ singular value thresholding and Chatterjee (2015), and these three methods are generally significantly better than other practical methods. These results suggest there may be a $(2, \infty)$ error bound applicable to the low rank methods, even though they were not designed to control this type of errors; investigating this is outside the scope of this manuscript.

9. PROOFS

Proof of Theorem 1. For convenience, we start with summarizing notation and assumptions made in the main paper. Let $0 = x_0 < x_1 < \dots < x_K = 1$, $I_k := [x_{k-1}, x_k)$ for $1 \leq k \leq K - 1$ and $I_K = [x_{K-1}, X_K]$. Assume the graphon f is a bi-Lipschitz function on each of $I_k \times I_\ell$ for $1 \leq k, \ell \leq K$. Let L denote the maximum piece-wise bi-Lipschitz constant.

Assumption 1. The number of pieces K may grow with n , as long as $\min_k |I_k| / (n^{-1} \log n)^{1/2} \rightarrow \infty$.

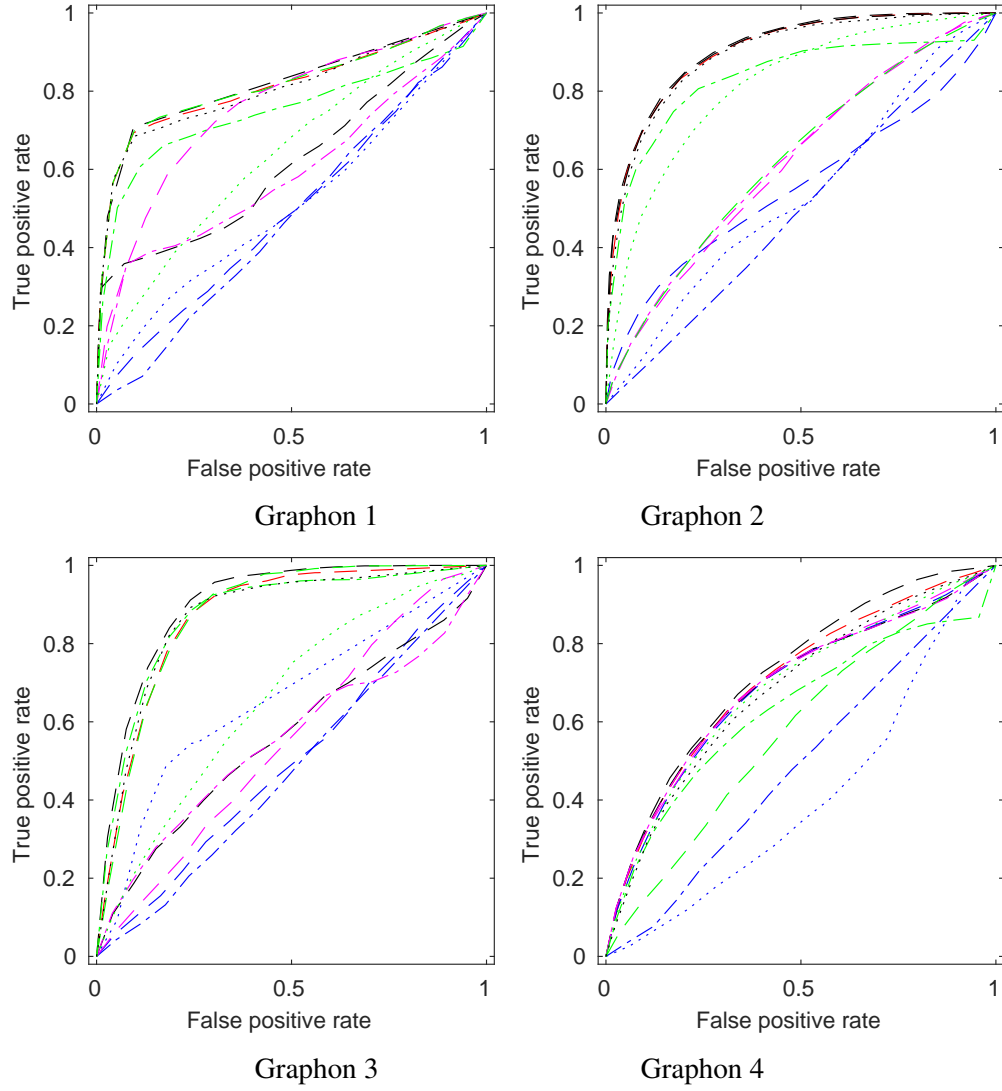


Fig. 4: Receiver operating characteristic curves for link prediction by different methods for graphons 1 to 4. Black dashed curve: the true probability matrix; red dashed curve: our method; black dotted curve: $n^{1/3}$ singular value thresholding; blue dashed curve: Zhao et al. (2017); blue dash dotted curve: the Jaccard index; blue dotted curve: PropFlow; black dash dotted curve: Chatterjee (2015); magenta dashed curve: Chan & Airolidi (2014); magenta dash dotted curve: Yang et al. (2014); green dashed curve: block model with spectral clustering; green dash dotted curve: oracle block model approximation.

For any $\xi \in [0, 1]$, let $I(\xi)$ denote the I_k that contains ξ . Let $S_i(\Delta) = [\xi_i - \Delta, \xi_i + \Delta] \cap I(\xi_i)$ denote the neighborhood of ξ_i in which $f(x, y)$ is Lipschitz in $x \in S_i(\Delta)$ for any fixed y . Finally, recall our estimator is defined by

$$\tilde{P}_{ij} = \frac{\sum_{i' \in \mathcal{N}_i} A_{i'j}}{|\mathcal{N}_i|}$$

Table 3: Root mean squared errors in the $(2, \infty)$ norm, followed by standard errors in parenthesis, all multiplied by 10^2 , averaged over 2000 replications.

	Graphon 1	Graphon 2	Graphon 3	Graphon 4
Our method	4.84(0.12)	5.97(0.08)	5.60(0.11)	7.13(0.11)
Chan & Airoldi (2014)	38.87(0.35)	50.44(0.80)	13.74(0.09)	9.43(0.14)
Yang et al. (2014)	38.87(0.41)	49.77(0.74)	13.11(0.04)	7.75(0.07)
$n^{1/3}$ singular value thresholding	6.55(0.18)	8.34(0.15)	5.22(0.11)	13.67(0.24)
Blockmodel spectral	16.65(1.38)	45.34(0.39)	13.00(0.57)	29.50(0.53)
Blockmodel oracle	35.36(0.04)	9.47(0.02)	2.53(0.02)	1.73(0.03)
Chatterjee (2015)	8.77(0.01)	3.93(0.10)	8.42(0.07)	7.73(0.06)
Airoldi's method	34.74(0.46)	66.97(0.12)	16.79(0.07)	30.47(0.74)

To prove the main theorem, it suffices to show that with high probability

$$\frac{1}{n} \sum_j \left(\tilde{P}_{ij} - P_{ij} \right)^2 \leq \left(\frac{\log n}{n} \right)^{1/2}$$

holds for all i . We begin the proof with the following decomposition of the error term:

$$\begin{aligned} \frac{1}{n} \sum_j \left(\tilde{P}_{ij} - P_{ij} \right)^2 &= \frac{1}{n} \sum_j \left\{ \frac{\sum_{i' \in \mathcal{N}_i} (A_{i'j} - P_{i'j})}{|\mathcal{N}_i|} \right\}^2 \\ &= \frac{1}{n} \sum_j \left[\frac{\sum_{i' \in \mathcal{N}_i} \{ (A_{i'j} - P_{i'j}) + (P_{i'j} - P_{ij}) \}}{|\mathcal{N}_i|} \right]^2. \end{aligned}$$

We can bound the summand by

$$\begin{aligned} &\left[\frac{\sum_{i' \in \mathcal{N}_i} \{ (A_{i'j} - P_{i'j}) + (P_{i'j} - P_{ij}) \}}{|\mathcal{N}_i|} \right]^2 \\ &\leq 2 \left\{ \frac{\sum_{i' \in \mathcal{N}_i} (A_{i'j} - P_{i'j})}{|\mathcal{N}_i|} \right\}^2 + 2 \left\{ \frac{\sum_{i' \in \mathcal{N}_i} (P_{i'j} - P_{ij})}{|\mathcal{N}_i|} \right\}^2 = 2J_1(i, j) + 2J_2(i, j). \quad (14) \end{aligned}$$

Our goal is to bound $n^{-1} \sum_i \{ J_1(i, j) + J_2(i, j) \}$. First, we prove a lemma which estimates the proportion of nodes in a diminishing neighborhood of ξ_i 's.

LEMMA 1. For arbitrary global constants $C_1, \tilde{C}_1 > 0$, define $\Delta_n = \left\{ C_1 + (\tilde{C}_1 + 4)^{1/2} \right\} (n^{-1} \log n)^{1/2}$. For n large enough so that $\left\{ (\tilde{C}_1 + 4) \log n / n \right\}^{1/2} \leq 1$ and $\Delta_n < \min_k |I_k|/2$, we have

$$\text{pr} \left\{ \min_i \frac{|\{i' \neq i : \xi_{i'} \in S_i(\Delta_n)\}|}{n-1} \geq C_1 (n^{-1} \log n)^{1/2} \right\} \geq 1 - 2n^{-\tilde{C}_1/4}. \quad (15)$$

Proof of Lemma 1. For any $0 < \epsilon \leq 1$ and n large enough to satisfy the assumptions, by Bernstein's inequality we have, for any i ,

$$\text{pr} \left\{ \left| \frac{|\{i' \neq i : \xi_{i'} \in S_i(\Delta_n)\}|}{n-1} - |S_i(\Delta_n)| \right| \geq \epsilon \right\} \leq 2 \exp \left\{ -\frac{\frac{1}{2}(n-1)\epsilon^2}{1 + \frac{1}{3}\epsilon} \right\} \leq 2 \exp \left(-\frac{1}{4}n\epsilon^2 \right).$$

Taking a union bound over all i 's gives

$$\text{pr} \left\{ \max_i \left| \frac{|\{i' \neq i : \xi_{i'} \in S_i(\Delta_n)\}|}{n-1} - |S_i(\Delta_n)| \right| \geq \epsilon \right\} \leq 2n \exp \left(-\frac{1}{4} n \epsilon^2 \right).$$

Letting $\epsilon = \left\{ (\tilde{C}_1 + 4) \log n / n \right\}^{1/2}$, we have

$$\text{pr} \left[\max_i \left| \frac{|\{i' \neq i : \xi_{i'} \in S_i(\Delta_n)\}|}{n-1} - |S_i(\Delta_n)| \right| \geq \left\{ \frac{(\tilde{C}_1 + 4) \log n}{n} \right\}^{1/2} \right] \leq 2n^{-\tilde{C}_1/4}. \quad (16)$$

Next we claim that either $[\xi_i - \Delta_n, \xi_i] \subseteq I(\xi_i)$ or $[\xi_i, \xi_i + \Delta_n] \subseteq I(\xi_i)$ holds for all i . If for some i the claim does not hold, by the definition of $I(\xi_i)$, we have $I(\xi_i) \subset [\xi_i - \Delta_n, \xi_i + \Delta_n]$. So we have $|I(\xi_i)| \leq 2\Delta_n$, but this contradicts the condition $\Delta_n < \min_k |I_k|/2$. The claim yields that $|S_i(\Delta_n)| \geq \Delta_n$. Finally, by (16), with probability $1 - 2n^{-\tilde{C}_1/4}$, we have

$$\begin{aligned} \min_i \frac{|\{i' \neq i : \xi_{i'} \in S_i(\Delta_n)\}|}{n-1} &\geq |S_i(\Delta_n)| - \left\{ \frac{(\tilde{C}_1 + 4) \log n}{n} \right\}^{1/2} \\ &\geq \Delta_n - \left\{ \frac{(\tilde{C}_1 + 4) \log n}{n} \right\}^{1/2} \geq C_1 (n^{-1} \log n)^{1/2}. \end{aligned}$$

This completes the proof of Lemma 1. \square

We now continue with the proof of Theorem 1. Recall that we defined a measure of closeness of adjacency matrix slices in Section 2 as

$$\tilde{d}(i, i') = \max_{k \neq i, i'} |\langle A_{i.} - A_{i'.}, A_{k.} \rangle| / n = \max_{k \neq i, i'} |(A^2/n)_{ik} - (A^2/n)_{jk}|.$$

The neighborhood \mathcal{N}_i of node i consists of nodes (i') 's with $\tilde{d}(i, i')$ below the h -th quantile of $\{\tilde{d}(i, k)\}_{k \neq i}$. The next lemma shows two key properties of \mathcal{N}_i .

LEMMA 2. *Suppose we select the neighborhood \mathcal{N}_i by thresholding at the lower h -th quantile of $\{\tilde{d}(i, k)\}_{k \neq i}$, where we set $h = C_0 (n^{-1} \log n)^{1/2}$ with an arbitrary global constant C_0 satisfying $0 < C_0 \leq C_1$ for the C_1 from Lemma 1. Let $C_2, \tilde{C}_2 > 0$ be arbitrary global constants and assume $n \geq 6$ is large enough so that (i) All conditions on n in Lemma 1 are satisfied; (ii) $\{(C_2 + 2) \log n / n\}^{1/2} \leq 1$; (iii) $C_1 (n \log n)^{1/2} \geq 4$; and (iv) $4/n \leq \left\{ (C_2 + \tilde{C}_2 + 2)^{1/2} - (C_2 + 2)^{1/2} \right\} (n^{-1} \log n)^{1/2}$. Then the neighborhood \mathcal{N}_i has the following properties:*

1. $|\mathcal{N}_i| \geq C_0 (n \log n)^{1/2}$.
2. With probability $1 - 2n^{-\tilde{C}_1/4} - 2n^{-\tilde{C}_2/4}$, for all i and $i' \in \mathcal{N}_i$, we have

$$\|P_{i'} - P_i\|_2^2 / n \leq \left[6L \left\{ C_1 + (\tilde{C}_2 + 4)^{1/2} \right\}^{1/2} + 8 (C_2 + \tilde{C}_2 + 2)^{1/2} \right] (n^{-1} \log n)^{1/2}$$

Proof of Lemma 2. The first claim follows immediately from the choice of h and the definition of \mathcal{N}_i . To show the second claim, we start with concentration results. For any i, j such that $i \neq j$,

we have

$$\begin{aligned}
& \left| (A^2/n)_{ij} - (P^2/n)_{ij} \right| = \left| \sum_k (A_{ik}A_{kj} - P_{ik}P_{kj}) \right| / n \\
& \leq \frac{|\sum_{k \neq i,j} (A_{ik}A_{kj} - P_{ik}P_{kj})|}{n-2} \cdot \frac{n-2}{n} + \frac{|(A_{ii} + A_{jj})A_{ij}| + |(P_{ii} + P_{jj})P_{ij}|}{n} \\
& \leq \frac{|\sum_{k \neq i,j} (A_{ik}A_{kj} - P_{ik}P_{kj})|}{n-2} + \frac{4}{n} \tag{17}
\end{aligned}$$

By Bernstein's inequality, for any $0 < \epsilon \leq 1$ and $n \geq 3$ we have

$$\text{pr} \left\{ \frac{|\sum_{k \neq i,j} (A_{ik}A_{kj} - P_{ik}P_{kj})|}{n-2} \geq \epsilon \right\} \leq 2 \exp \left\{ -\frac{\frac{1}{2}(n-2)\epsilon^2}{1 + \frac{1}{3}\epsilon} \right\} \leq 2 \exp \left(-\frac{1}{4}n\epsilon^2 \right).$$

Taking a union bound over all $i \neq j$, we have

$$\text{pr} \left\{ \max_{i,j:i \neq j} \frac{|\sum_{k \neq i,j} (A_{ik}A_{kj} - P_{ik}P_{kj})|}{n-2} \geq \epsilon \right\} \leq 2n^2 \exp \left(-\frac{1}{4}n\epsilon^2 \right).$$

Then setting $\epsilon = \{(C_2 + 2) \log n/n\}^{1/2}$ with n large enough so that $\epsilon \leq 1$, we have

$$\text{pr} \left\{ \max_{i,j:i \neq j} \frac{|\sum_{k \neq i,j} (A_{ik}A_{kj} - P_{ik}P_{kj})|}{n-2} \geq \left\{ \frac{(C_2 + 2) \log n}{n} \right\}^{1/2} \right\} \leq 2n^{-\tilde{C}_2/4} \tag{18}$$

Combining (17) and (18), with probability $1 - 2n^{-\tilde{C}_2/4}$, the following holds

$$\max_{i,j:i \neq j} \left| (A^2/n)_{ij} - (P^2/n)_{ij} \right| \leq \left\{ \frac{(C_2 + 2) \log n}{n} \right\}^{1/2} + \frac{4}{n} \leq \left\{ \frac{(C_2 + \tilde{C}_2 + 2) \log n}{n} \right\}^{1/2} \tag{19}$$

for n large enough to satisfy (iv). Next, we prove a useful inequality. For all i and any \tilde{i} such that $\xi_{\tilde{i}} \in S_i(\Delta_n)$, we have

$$\left| (P^2/n)_{ik} - (P^2/n)_{\tilde{i}k} \right| = |\langle P_{i\cdot}, P_k \rangle - \langle P_{\tilde{i}\cdot}, P_k \rangle| / n \leq \|P_{i\cdot} - P_{\tilde{i}\cdot}\|_2 \|P_k\|_2 / n \leq L\Delta_n \tag{20}$$

for all k , where the last inequality follows from

$$|P_{i'\ell} - P_{i\ell}| = |f(\xi_{i'}, \xi_\ell) - f(\xi_i, \xi_\ell)| \leq L|\xi_{i'} - \xi_i| \leq L\Delta_n$$

for all ℓ , and $\|P_k\|_2 \leq n^{1/2}$ for all k . Note that this holds for all k , including $k = i$ or $k = \tilde{i}$.

We are now ready to upper bound $\tilde{d}(i, i')$ for $i' \in \mathcal{N}_i$. We bound $\tilde{d}(i, i')$ via bounding $\tilde{d}(i, \tilde{i})$ for \tilde{i} with $\xi_{\tilde{i}} \in S_i(\Delta_n)$. By (19) and (20), with probability $1 - 2n^{-\tilde{C}_2/4}$, we have

$$\begin{aligned}
\tilde{d}(i, \tilde{i}) &= \max_{k \neq i, \tilde{i}} |(A^2/n)_{ik} - (A^2/n)_{\tilde{i}k}| \\
&\leq \max_{k \neq i, \tilde{i}} |(P^2/n)_{ik} - (P^2/n)_{\tilde{i}k}| + 2 \max_{i,j:i \neq j} |(A^2/n)_{ij} - (P^2/n)_{ij}| \\
&\leq L\Delta_n + 2 \left\{ \frac{(C_2 + \tilde{C}_2 + 2) \log n}{n} \right\}^{1/2} \tag{21}
\end{aligned}$$

Now since the fraction of nodes contained in $|\{\tilde{i} : \xi_{\tilde{i}} \in S_i(\Delta_n)\}|$ is at least h , this upper bounds $\tilde{d}(i, i')$ for $i' \in \mathcal{N}_i$, since nodes in \mathcal{N}_i have the lowest h fraction of values in $\{\tilde{d}(i, k)\}_k$.

Setting Δ_n as in Lemma 1, by Lemma 1 and (19), with probability $1 - 2n^{-\tilde{C}_1/4} - 2n^{-\tilde{C}_2/4}$, for all i , at least $C_1(n^{-1} \log n)^{1/2}$ fraction of nodes $\tilde{i} \neq i$ satisfy both $\xi_{\tilde{i}} \in S_i(\Delta_n)$ and

$$\tilde{d}(i, \tilde{i}) \leq L\Delta_n + 2 \left\{ \frac{(C_2 + \tilde{C}_2 + 2) \log n}{n} \right\}^{1/2}. \quad (22)$$

Recall that $i' \in \mathcal{N}_i$ have the smallest $h = C_0(n^{-1} \log n)^{1/2} \leq C_1(n^{-1} \log n)^{1/2}$ fraction of $\tilde{d}(i, i')$'s. Then (22) yields that

$$\tilde{d}(i, i') \leq L\Delta_n + 2 \left\{ \frac{(C_2 + \tilde{C}_2 + 2) \log n}{n} \right\}^{1/2} \quad (23)$$

holds for all i and all $i' \in \mathcal{N}_i$ simultaneously with probability $1 - 2n^{-\tilde{C}_1/4} - 2n^{-\tilde{C}_2/4}$.

We are now ready to complete the proof of the second claim of Lemma 2. By Lemma 1, (19), (20) and (23), with probability $1 - 2n^{-\tilde{C}_1/4} - 2n^{-\tilde{C}_2/4}$, the following holds. For n large enough such that $\min_i |\{i' : \xi_{i'} \in S_i(\Delta_n)\}| \geq C_1(n \log n)^{1/2} \geq 4$ (by Lemma 1), for all i and $i' \in \mathcal{N}_i$ we can find $\tilde{i} \in S_i(\Delta_n)$ and $\tilde{i}' \in S_{i'}(\Delta_n)$ such that i, i', \tilde{i} and \tilde{i}' are different from each other. Then we have

$$\begin{aligned} & \|P_{i\cdot} - P_{i'\cdot}\|_2^2/n = (P^2/n)_{ii} - (P^2/n)_{i'i} + (P^2/n)_{i'i'} - (P^2/n)_{ii'} \\ & \leq |(P^2/n)_{ii} - (P^2/n)_{i'i}| + |(P^2/n)_{i'i'} - (P^2/n)_{ii'}| \\ & \leq |(P^2/n)_{\tilde{i}\tilde{i}} - (P^2/n)_{i'\tilde{i}}| + |(P^2/n)_{i'\tilde{i}'} - (P^2/n)_{\tilde{i}\tilde{i}'}| + 4L\Delta_n \\ & \leq |(A^2/n)_{\tilde{i}\tilde{i}} - (A^2/n)_{i'\tilde{i}}| + |(A^2/n)_{i'\tilde{i}'} - (A^2/n)_{\tilde{i}\tilde{i}'}| + 4 \left\{ \frac{(C_2 + \tilde{C}_2 + 2) \log n}{n} \right\}^{1/2} + 4L\Delta_n \\ & \leq 2 \max_{k \neq i, i'} |(A^2/n)_{ik} - (A^2/n)_{i'k}| + 4 \left\{ \frac{(C_2 + \tilde{C}_2 + 2) \log n}{n} \right\}^{1/2} + 4L\Delta_n \\ & = 2\tilde{d}(i, i') + 4 \left\{ \frac{(C_2 + \tilde{C}_2 + 2) \log n}{n} \right\}^{1/2} + 4L\Delta_n \leq 8 \left\{ \frac{(C_2 + \tilde{C}_2 + 2) \log n}{n} \right\}^{1/2} + 6L\Delta_n \\ & = \left[6L \left\{ C_1 + (\tilde{C}_2 + 4)^{1/2} \right\}^{1/2} + 8(C_2 + \tilde{C}_2 + 2)^{1/2} \right] (n^{-1} \log n)^{1/2}. \end{aligned}$$

This completes the proof of Lemma 2. \square

We are now ready to bound $n^{-1} \sum_j \{J_1(i, j) + J_2(i, j)\}$, which will complete the proof of Theorem 1. Note that we cannot simply bound each individual $J_1(i, j)$'s by Bernstein's inequality since $A_{i'j}$ is not independent of the event $i' \in \mathcal{N}_i$. Instead, we work with the sum $n^{-1} \sum_j J_1(i, j)$ and decompose it as follows.

$$\begin{aligned} \frac{1}{n} \sum_j J_1(i, j) &= \frac{1}{n|\mathcal{N}_i|^2} \sum_j \left\{ \sum_{i' \in \mathcal{N}_i} (A_{i'j} - P_{i'j}) \right\}^2 \\ &= \frac{1}{n|\mathcal{N}_i|^2} \sum_j \left\{ \sum_{i' \in \mathcal{N}_i} (A_{i'j} - P_{i'j})^2 + \sum_{i' \in \mathcal{N}_i} \sum_{i'' \neq i', i'' \in \mathcal{N}_i} (A_{i'j} - P_{i'j})(A_{i''j} - P_{i''j}) \right\}. \quad (24) \end{aligned}$$

The first term in (24) satisfies

$$\sum_j (A_{i'j} - P_{i'j})^2/n = \|A_{i'} - P_{i'}\|_2^2/n \leq 1 \quad (25)$$

where the inequality is due to $|A_{i'j} - P_{i'j}| \leq 1$ for all j . The second term in (24) can be bounded by

$$\begin{aligned} & \frac{1}{n|\mathcal{N}_i|^2} \sum_j \sum_{i' \in \mathcal{N}_i} \sum_{i'' \neq i', i'' \in \mathcal{N}_i} (A_{i'j} - P_{i'j})(A_{i''j} - P_{i''j}) \leq \\ & \leq \frac{1}{|\mathcal{N}_i|^2} \sum_{i', i'' \in \mathcal{N}_i: i' \neq i''} \left| \frac{1}{n} \sum_j (A_{i'j} - P_{i'j})(A_{i''j} - P_{i''j}) \right| \\ & \leq \frac{1}{|\mathcal{N}_i|^2} \sum_{i', i'' \in \mathcal{N}_i: i' \neq i''} \left\{ \frac{1}{n-2} \left| \sum_{j \neq i', i''} (A_{i'j} - P_{i'j})(A_{i''j} - P_{i''j}) \right| \cdot \frac{n-2}{n} \right. \\ & \quad \left. + \frac{|(A_{i'i''} - P_{i'i''})| |(A_{i'i'} - P_{i'i'} + A_{i''i''} - P_{i''i''})|}{n} \right\} \\ & \leq \frac{1}{|\mathcal{N}_i|^2} \sum_{i', i'' \in \mathcal{N}_i: i' \neq i''} \left\{ \frac{1}{n-2} \left| \sum_{j \neq i', i''} (A_{i'j} - P_{i'j})(A_{i''j} - P_{i''j}) \right| + \frac{2}{n} \right\}. \quad (26) \end{aligned}$$

To bound the first term in (26), for any $i_1 \neq i_2$ and $\epsilon > 0$, by Bernstein's inequality we have

$$\text{pr} \left\{ \frac{1}{n-2} \left| \sum_{j \neq i_1, i_2} (A_{i_1j} - P_{i_1j})(A_{i_2j} - P_{i_2j}) \right| \geq \epsilon \right\} \leq 2 \exp \left\{ -\frac{\frac{1}{2}(n-2)\epsilon^2}{1 + \frac{1}{3}\epsilon} \right\} \leq 2n^2 e^{-n\epsilon^2/4}.$$

Let $C_3, \tilde{C}_3 > 0$ be arbitrary global constants and let n be large enough so that $1/\{C_0(n \log n)^{1/2}\} + 2/n \leq \left\{ (C_3 + \tilde{C}_3 + 8)^{1/2} - (C_3 + 8)^{1/2} \right\} (n^{-1} \log n)^{1/2}$. First, taking $\epsilon = \{(C_3 + 8) \log n/n\}^{1/2}$ and a union bound over all $i_1 \neq i_2$, we have

$$\text{pr} \left[\max_{i_1, i_2, i_1 \neq i_2} \frac{1}{n-2} \left| \sum_{j \neq i_1, i_2} (A_{i_1j} - P_{i_1j})(A_{i_2j} - P_{i_2j}) \right| \geq \left\{ \frac{(C_3 + 8) \log n}{n} \right\}^{1/2} \right] \leq 2n^{-\tilde{C}_3/4}. \quad (27)$$

Then plugging (25), (26) and (27) into (24) and combining with claim 1 of Lemma 2, with probability $1 - 2n^{-\tilde{C}_1/4} - 2n^{-\tilde{C}_2/4} - 2n^{-\tilde{C}_3/4}$, for all i simultaneously, we have

$$\begin{aligned} & \frac{1}{n} \sum_j J_1(i, j) \leq \frac{1}{|\mathcal{N}_i|^2} \sum_{i' \in \mathcal{N}_i} \left[1 + (|\mathcal{N}_i| - 1) \left(\left\{ \frac{(C_3 + 8) \log n}{n} \right\}^{1/2} + \frac{2}{n} \right) \right] \\ & \leq \frac{1}{|\mathcal{N}_i|} + \left\{ \frac{(8 + C_3) \log n}{n} \right\}^{1/2} + \frac{2}{n} \leq \frac{1}{C_0(n \log n)^{1/2}} + \frac{2}{n} + \left\{ \frac{(C_3 + 8) \log n}{n} \right\}^{1/2} \\ & \leq \left\{ \frac{(C_3 + \tilde{C}_3 + 8) \log n}{n} \right\}^{1/2}. \quad (28) \end{aligned}$$

We now bound $n^{-1} \sum_j J_2(i, j)$. By Lemma 2, with probability $1 - 2n^{-\tilde{C}_1/4} - 2n^{-\tilde{C}_2/4}$, for all i simultaneously, we have

$$\begin{aligned} \frac{1}{n} \sum_j J_2(i, j) &= \frac{1}{n} \sum_j \left\{ \frac{\sum_{i' \in \mathcal{N}_i} (P_{i'j} - P_{ij})}{|\mathcal{N}_i|} \right\}^2 \\ &\leq \frac{\sum_{i' \in \mathcal{N}_i} \sum_j (P_{i'j} - P_{ij})^2 / n}{|\mathcal{N}_i|} = \frac{\sum_{i' \in \mathcal{N}_i} \|P_{i'} - P_i\|_2^2 / n}{|\mathcal{N}_i|} \\ &\leq \left[6L \left\{ C_1 + (\tilde{C}_2 + 4)^{1/2} \right\}^{1/2} + 8 \left(C_2 + \tilde{C}_2 + 2 \right)^{1/2} \right] (n^{-1} \log n)^{1/2}, \end{aligned} \quad (29)$$

where the first inequality is the Cauchy-Schwartz inequality and the second inequality follows from claim 2 of Lemma 2.

Combining (28) and (29) completes the proof of Theorem 1.

Proof of Proposition 1. We only need to prove the upper bound on the error rate; the lower bound follows by setting $\delta = 1$ and applying Theorem 2.3 in Gao et al. (2015). We show that the least squares estimator defined in (2.4) of Gao et al. (2015), which we shall denote as \hat{P}_{LS} here, has the same error rate for graphons in our family $\mathcal{F}_{\delta, L}$. The proof is almost identical to the proof of Theorem 2.3 in Gao et al. (2015), except we need to choose a θ^* that respects the partition of $[0, 1]$ into intervals of continuity with a non-essential adaptation of their Lemma 2.1. Referring to the proof of Lemma 2.1, instead of choosing z^* using $U_a = [(a-1)/k, a/k]$, we now set

$$(z^*)^{-1}(a) = \{i \in [n] : \xi_i \in \tilde{U}_a\}$$

where \tilde{U}_a 's are defined as follows. Recall Definition 2 of our paper which specifies the sequence x_0, \dots, x_K . We set $\delta = \min_{0 \leq s \leq K-1} (x_{s+1} - x_s) \succ (n^{-1} \log n)^{1/2}$ and consequently $K \prec (n/\log n)^{1/2}$. Like in Gao et al. (2015), we use a stochastic block model with $k = n^{1/2}$ equal-sized communities to approximate the true probability matrix. This corresponds to using a piece-wise constant graphon function, with pieces of equal size $1/k = n^{-1/2}$ to approximate the true graphon. Thus for large enough n , we have the following properties: at most one x_i , $1 \leq i \leq K-1$ may fall in any U_a for $1 \leq a \leq k$; no x_i may fall in U_1, U_2, U_{n-1} or U_k ; and if an x_i falls in U_a , no x_i will fall in $U_{a-2}, U_{a-1}, U_{a+1}$ or U_{a+2} , because for large enough n , we have

$$\begin{aligned} x_i - x_{i-1} &\geq \delta \asymp (n^{-1} \log n)^{1/2} \succ n^{-1/2} \asymp 3/k = |U_{a-2} \cup U_{a-1} \cup U_a|, \\ x_{i+1} - x_i &\geq \delta \asymp (n^{-1} \log n)^{1/2} \succ n^{-1/2} \asymp 3/k = |U_a \cup U_{a+1} \cup U_{a+2}|. \end{aligned}$$

Then we define \tilde{U} as follows. First, we set $\tilde{U}_1 = U_1$ and $\tilde{U}_n = U_n$. For all $2 \leq a \leq n-1$ such that no x_i , $1 \leq i \leq K-1$ falls in any of U_{a-1}, U_a and U_{a+1} , we let $\tilde{U}_a = U_a$. Lastly for all $2 \leq a \leq n-1$ such that an x_i , $1 \leq i \leq K-1$ falls in U_a , if $x_1 \leq (2a-1)/(2k)$, we set $\{\tilde{U}_{a-1}, \tilde{U}_a, \tilde{U}_{a+1}\} = \{U_{a-1} \cup [(a-1)/k, x_i], [x_i, a/k], U_{a+1}\}$, and otherwise we set $\{\tilde{U}_{a-1}, \tilde{U}_a, \tilde{U}_{a+1}\} = \{U_{a-1}, [(a-1)/k, x_i], [x_i, a/k] \cup U_{a+1}\}$.

Let P^* denote the probability matrix for the stochastic block model approximation to the true probability matrix P corresponding to the partition \tilde{U} . That is, for all i, j such that $\xi_i \in \tilde{U}_a$ and $\xi_j \in \tilde{U}_b$, define:

$$P_{ij}^* = \frac{1}{|\{i' : \xi_{i'} \in \tilde{U}_a\}| |\{j' : \xi_{j'} \in \tilde{U}_b\}|} \sum_{i' \in \tilde{U}_a, j' \in \tilde{U}_b} f(\xi_{i'}, \xi_{j'})$$

We can apply the reasoning in the proof of Lemma 2.1 in Gao et al. (2015), which yields:

$$\begin{aligned}
|P_{ij}^* - P_{ij}| &\leq \frac{1}{|\{i' : \xi_{i'} \in \tilde{U}_a\}| |\{j' : \xi_{j'} \in \tilde{U}_b\}|} \sum_{i' \in \tilde{U}_a, j' \in \tilde{U}_b} |f(\xi_{i'}, \xi_{j'}) - f(\xi_i, \xi_j)| \\
&\leq \frac{1}{|\{i' : \xi_{i'} \in \tilde{U}_a\}| |\{j' : \xi_{j'} \in \tilde{U}_b\}|} \sum_{i' \in \tilde{U}_a, j' \in \tilde{U}_b} L (|\xi_i - \xi_{i'}| + |\xi_j - \xi_{j'}|) \\
&\leq 4L/k \asymp 1/k = n^{-1/2}.
\end{aligned} \tag{30}$$

We can then apply the argument in Equations (4.1) through (4.5) in Gao et al. (2015) and combine our result (30) with Lemmas 4.1, 4.3 and 4.4 in Gao et al. (2015) to obtain the following result:

$$\|\hat{P}_{\text{LS}} - P\|_F^2 \leq 2\|\hat{P}_{\text{LS}} - P^*\|_F^2 + 2\|P^* - P\|_F^2 \leq C(k^2 + k \log n + n^2 k^{-2}) \leq C' n \log n,$$

where we plugged in $k = n^{1/2}$. \square

Proof of Theorem 2. In this proof, we first construct a ‘‘baseline’’ network to be a stochastic block model with $K = 3$ communities, which we label 0, 1, and 2, of sizes m , ℓ , and ℓ , respectively. Thus $n = m + 2\ell$. Letting $Z^{(0)} \in \{0, 1\}^{n \times K}$ denote the membership matrix of the n nodes, we can, without loss of generality, set $Z_{i1} = 1$ for $1 \leq i \leq m$, $Z_{i2} = 1$ for $m + 1 \leq i \leq m + \ell$ and $Z_{i3} = 1$ for $m + \ell + 1 \leq i \leq n$. For all other (i, k) , we set $Z_{ik} = 0$. Next, we define the $K \times K$ probability matrix B as:

$$B = \begin{pmatrix} 1/2 & 1/2 + \phi & 1/2 \\ 1/2 + \phi & 1/2 & 1/2 \\ 1/2 & 1/2 & 1/2 \end{pmatrix}$$

where ϕ is a small positive value to be determined later. The probability matrix is then $EA = P = ZBZ^T$.

We then construct N probability matrices $\mathcal{P} = \{P^{(1)}, \dots, P^{(N)}\}$, where N is a natural number to be determined later. Each $P^{(i)}$ is a stochastic block model with $K = 3$ communities, its own community membership assignment, and the same B as the baseline network. That is, $P^{(i)} = Z^{(i)} B \{Z^{(i)}\}^T$.

Now we construct $Z^{(i)}$. Without loss of generality, we assume ℓ is an even number. Otherwise, we can switch one node from communities 1 and 2 each to community 0, and this will not change the lower bound on the error rate we are going to establish. To proceed, we use the Gilbert-Varshamov bound, which was also used by Gao et al. (2015) and Klopp et al. (2017), to construct N indicator vectors $w^{(1)}, \dots, w^{(N)}$:

LEMMA 3. *For any positive integer ℓ , there exists $\{w^{(1)}, \dots, w^{(N)}\}$ for some $N \geq \exp(\ell/8)$, where $w^{(i)} \in \{0, 1\}^\ell$, such that for any $i \neq j$, we have*

$$\|w^{(i)} - w^{(j)}\|_H \geq \ell/4$$

where $\|x - y\|_H = \sum_s 1(x_s \neq y_s)$ is the Hamming distance.

We construct $Z^{(i)}$ as follows. For $1 \leq j \leq m + \ell/2$ and $m + \ell + 1 \leq j \leq m + 3\ell/2$, we set $Z_j^{(i)} = Z_j^{(0)}$. That is, for all nodes in community 0 and the first half of nodes in communities 1 and 2, their community memberships match the baseline network. For $1 \leq j' \leq \ell/2$, set $Z_{m+\ell/2+j'}^{(i)} = (0, 1, 0)$ if $w_{j'}^{(i)} = 0$, and set $Z_{m+\ell/2+j'}^{(i)} = (0, 0, 1)$ if $w_{j'}^{(i)} = 1$. For $\ell/2 + 1 \leq j' \leq \ell$, set $Z_{m+\ell+j'}^{(i)} = (0, 0, 1)$ if $w_{j'}^{(i)} = 0$, and set $Z_{m+\ell+j'}^{(i)} = (0, 1, 0)$ if $w_{j'}^{(i)} = 1$. That is, for

the second half of community 1 and 2, the community membership matches the baseline network if the corresponding element in w is 0; otherwise 1 and 2 are switched. Therefore, for any $i \neq j \in \{1, \dots, N\}$, we have

$$d_{2,\infty}^2 \{P^{(i)}, P^{(j)}\} \geq \|P_1^{(i)} - P_1^{(j)}\|_2^2/n = \phi^2 \|w^{(i)} - w^{(j)}\|_H/n \geq \frac{\phi^2 \ell}{4n}$$

We will choose ϕ small enough for a sufficiently large n such that $\phi \leq 1/4$, so that every edge probability lies in $(1/2, 3/4)$, which enables us to apply Proposition 4.2 from Gao et al. (2015). Then noticing that for $i \neq j$, matrices $P^{(i)}$ and $P^{(j)}$ can only differ by ϕ in at most $2mn$ elements by definition, we have

$$D \{P^{(i)} \| P^{(j)}\} \leq 8 \|P^{(i)} - P^{(j)}\|_F^2 \leq 8 \cdot 2mn \cdot \phi^2 = 16mn\phi^2$$

where $D(f \| g)$ denotes the Kullback-Leibler divergence between distributions f and g .

Now we are ready to complete the proof. By Lemma 3 of Yu (1997), we have

$$\begin{aligned} \max_i E_{P^{(i)}} d_{2,\infty}^2 \{\hat{P}, P^{(i)}\} &\geq \frac{\min_{i,j} d_{2,\infty}^2 \{P^{(i)}, P^{(j)}\}}{4} \left[1 - \frac{\max_{i,j} D \{P^{(i)} \| P^{(j)}\} + 2}{\log |\mathcal{P}|} \right]^2 \\ &\geq \frac{\phi^2 \ell}{16n} \left(1 - \frac{16mn\phi^2 + 2}{\ell/8} \right)^2 \geq \frac{\phi^2 \ell}{16n} \left(\frac{1}{2} - \frac{128mn\phi^2}{\ell} \right)^2 \end{aligned} \quad (31)$$

The we maximize the right hand side of (31) by setting $\phi^2 = \ell/(257mn)$ and $\ell \geq \max\{n/3, 32\}$. Then

$$\frac{\phi^2 \ell}{16n} \left(\frac{1}{2} - \frac{128mn\phi^2}{\ell} \right)^2 = \frac{\ell^2}{8224mn^2} \geq \frac{1}{74016m}$$

The matrices $\mathcal{P} = \{P^{(1)}, \dots, P^{(N)}\}$ need to correspond to stochastic block models in our piecewise bi-Lipschitz graphon space $F_{\delta,L}$ such that $\delta / (\log n/n)^{1/2} \rightarrow \infty$. That is, the size of the smallest block must grow to infinity faster than $(n \log n)^{1/2}$. Therefore,

$$\max_{f \in F_{\delta,L}: \delta / (\log n/n)^{1/2} \rightarrow \infty} 1/m = (n \log n)^{-1/2}$$

This completes the proof. \square

REFERENCES

- ADAMIC, L. A. & GLANCE, N. (2005). The political blogosphere and the 2004 U.S. Election: divided they blog. In *Proceedings of the 3rd International Workshop on Link Discovery*, LinkKDD '05. New York: ACM, pp. 36–43.
- AIROLDI, E. M., BLEI, D. M., FIENBERG, S. E. & XING, E. P. (2008). Mixed membership stochastic blockmodels. *Journal of Machine Learning Research* **9**, 1981–2014.
- AIROLDI, E. M., COSTA, T. B. & CHAN, S. H. (2013). Stochastic blockmodel approximation of a graphon: Theory and consistent estimation. In *Advances in Neural Information Processing Systems 26*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani & K. Q. Weinberger, eds. Red Hook, NY: Curran Associates, Inc., pp. 692–700.
- ALDOUS, D. J. (1981). Representations for partially exchangeable arrays of random variables. *Journal of Multivariate Analysis* **11**, 581–598.
- AMINI, A. A., CHEN, A., BICKEL, P. J. & LEVINA, E. (2013). Pseudo-likelihood methods for community detection in large sparse networks. *The Annals of Statistics* **41**, 2097–2122.
- AMINI, A. A. & LEVINA, E. (2017). On semidefinite relaxations for the block model. *The Annals of Statistics* To appear.
- BICKEL, P. J. & CHEN, A. (2009). A nonparametric view of network models and Newman–Girvan and other modularities. *Proceedings of the National Academy of Sciences* **106**, 21068–21073.

- CAI, T. T. & LI, X. (2015). Robust and computationally feasible community detection in the presence of arbitrary outlier nodes. *The Annals of Statistics* **43**, 1027–1059.
- CANDES, E. J. & PLAN, Y. (2011). Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Transactions on Information Theory* **57**, 2342–2359.
- CHAN, S. H. & AIROLDI, E. (2014). A consistent histogram estimator for exchangeable graph models. *Journal of Machine Learning Research Workshop and Conference Proceedings* **32**, 208–216.
- CHATTERJEE, S. (2015). Matrix estimation by universal singular value thresholding. *The Annals of Statistics* **43**, 177–214.
- CHAUDHURI, K., CHUNG, F. & TSIATAS, A. (2012). Spectral clustering of graphs with general degrees in the extended planted partition model. *Journal of Machine Learning Research* **2012**, 1–23.
- CHOI, D. (2017). Co-clustering of nonsmooth graphons. *The Annals of Statistics* **45**, 1488–1515.
- CHOI, D. & WOLFE, P. J. (2014). Co-clustering separately exchangeable network data. *The Annals of Statistics* **42**, 29–63.
- GAO, C., LU, Y., MA, Z. & ZHOU, H. H. (2016). Optimal estimation and completion of matrices with biclustering structures. *Journal of Machine Learning Research* **17**, 1–29.
- GAO, C., LU, Y. & ZHOU, H. H. (2015). Rate-optimal graphon estimation. *The Annals of Statistics* **43**, 2624–2652.
- GUÉDON, O. & VERSHYNIN, R. (2016). Community detection in sparse networks via Grothendieck’s inequality. *Probability Theory and Related Fields* **165**, 1025–1049.
- HOFF, P. (2008). Modeling homophily and stochastic equivalence in symmetric relational data. *Advances in Neural Information Processing Systems* **20**, 657–664.
- KARRER, B. & NEWMAN, M. E. (2011). Stochastic blockmodels and community structure in networks. *Physical Review E* **83**, 016107–1–10.
- KLOPP, O., TSYBAKOV, A. B. & VERZELEN, N. (2017). Oracle inequalities for network models and sparse graphon estimation. *The Annals of Statistics* **45**, 316–354.
- LICHTENWALTER, R. N., LUSSIER, J. T. & CHAWLA, N. V. (2010). New perspectives and methods in link prediction. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’10. New York: ACM, pp. 243–252.
- OLHEDE, S. C. & WOLFE, P. J. (2014). Network histograms and universality of blockmodel approximation. *Proceedings of the National Academy of Sciences* **111**, 14722–14727.
- ROHE, K., CHATTERJEE, S. & YU, B. (2011). Spectral clustering and the high-dimensional stochastic blockmodel. *The Annals of Statistics* **39**, 1878–1915.
- SAADE, A., KRZAKALA, F. & ZDEBOROVÁ, L. (2014). Spectral clustering of graphs with the bethe hessian. In *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence & K. Q. Weinberger, eds. Red Hook, NY: Curran Associates, Inc., pp. 406–414.
- YANG, J. J., HAN, Q. & AIROLDI, E. M. (2014). Nonparametric estimation and testing of exchangeable graph models. *Journal of Machine Learning Research, Workshop and Conference Proceedings* **33**, 1060–1067.
- YU, B. (1997). Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam*. New York: Springer, pp. 423–435.
- ZHAO, Y., WU, Y.-J., LEVINA, E. & ZHU, J. (2017). Link prediction for partially observed networks. *Journal of Computational and Graphical Statistics To appear*.