

Optimal Estimation and Completion of Matrices with Biclustering Structures

Chao Gao

CHAO.GAO@YALE.EDU

Yu Lu

YU.LU@YALE.EDU

Yale University

Zongming Ma

ZONGMING@WHARTON.UPENN.EDU

University of Pennsylvania

Harrison H. Zhou

HUIBIN.ZHOU@YALE.EDU

Yale University

Editor: Edo Airoldi

Abstract

Biclustering structures in data matrices were first formalized in a seminal paper by John Hartigan (Hartigan, 1972) where one seeks to cluster cases and variables simultaneously. Such structures are also prevalent in block modeling of networks. In this paper, we develop a theory for the estimation and completion of matrices with biclustering structures, where the data is a partially observed and noise contaminated matrix with a certain underlying biclustering structure. In particular, we show that a constrained least squares estimator achieves minimax rate-optimal performance in several of the most important scenarios. To this end, we derive unified high probability upper bounds for all sub-Gaussian data and also provide matching minimax lower bounds in both Gaussian and binary cases. Due to the close connection of graphon to stochastic block models, an immediate consequence of our general results is a minimax rate-optimal estimator for sparse graphons.

Keywords: Biclustering, graphon, matrix completion, missing data, stochastic block models, sparse network

1. Introduction

In a range of important data analytic scenarios, we encounter matrices with biclustering structures. For instance, in gene expression studies, one can organize the rows of a data matrix to correspond to individual cancer patients and the columns to transcripts. Then the patients are expected to form groups according to different cancer subtypes and the genes are also expected to exhibit clustering effect according to the different pathways they belong to. Therefore, after appropriate reordering of the rows and the columns, the data matrix is expected to have a biclustering structure contaminated by noises (Lee et al., 2010). Here, the observed gene expression levels are real numbers. In a different context, such a biclustering structure can also be present in network data. For example, stochastic block model (SBM for short) (Holland et al., 1983) is a popular model for exchangeable networks. In SBMs, the graph nodes are partitioned into k disjoint communities and the probability that any pair of nodes are connected is determined entirely by the community memberships of the nodes. Consequently, if one rearranges the nodes from the same communities together in

the graph adjacency matrix, then the mean adjacency matrix, where each off-diagonal entry equals the probability of an edge connecting the nodes represented by the corresponding row and column, also has a biclustering structure.

The goal of the present paper is to develop a theory for the estimation (and completion when there are missing entries) of matrices with biclustering structures. To this end, we propose to consider the following general model

$$X_{ij} = \theta_{ij} + \epsilon_{ij}, \quad i \in [n_1], j \in [n_2], \quad (1)$$

where for any positive integer m , we let $[m] = \{1, \dots, m\}$. Here, for each (i, j) , $\theta_{ij} = \mathbb{E}[X_{ij}]$ and ϵ_{ij} is an independent piece of mean zero sub-Gaussian noise. Moreover, we allow entries to be missing completely at random (Rubin, 1976). Thus, let E_{ij} be i.i.d. Bernoulli random variables with success probability $p \in (0, 1]$ indicating whether the (i, j) th entry is observed, and define the set of observed entries

$$\Omega = \{(i, j) : E_{ij} = 1\}. \quad (2)$$

Our final observations are

$$X_{ij}, \quad (i, j) \in \Omega. \quad (3)$$

To model the biclustering structure, we focus on the case where there are k_1 row clusters and k_2 column clusters, and the values of $\{\theta_{ij}\}$ are taken as constant if the rows and the columns belong to the same clusters. The goal is then to recover the signal matrix $\theta \in \mathbb{R}^{n_1 \times n_2}$ from the observations (3). To accommodate most interesting cases, especially the case of undirected networks, we shall also consider the case where the data matrix X is symmetric with zero diagonals. In such cases, we also require $X_{ij} = X_{ji}$ and $E_{ij} = E_{ji}$ for all $i \neq j$.

Main contributions In this paper, we propose a unified estimation procedure for partially observed data matrix generated from model (1) – (3). We establish high probability upper bounds for the mean squared errors of the resulting estimators. In addition, we show that these upper bounds are minimax rate-optimal in both the continuous case and the binary case by providing matching minimax lower bounds. Furthermore, SBM can be viewed as a special case of the symmetric version of (1). Thus, an immediate application of our results is the network completion problem for SBMs. With partially observed network edges, our method gives a rate-optimal estimator for the probability matrix of the whole network in both the dense and the sparse regimes, which further leads to rate-optimal graphon estimation in both regimes.

Connection to the literature If only a low rank constraint is imposed on the mean matrix θ , then (1) – (3) becomes what is known in the literature as the matrix completion problem (Recht et al., 2010). An impressive list of algorithms and theories have been developed for this problem, including but not limited to Candès and Recht (2009); Keshavan et al. (2009); Candès and Tao (2010); Candès and Plan (2010); Cai et al. (2010); Keshavan et al. (2010); Recht (2011); Koltchinskii et al. (2011). In this paper, we investigate an alternative biclustering structural assumption for the matrix completion problem, which was first proposed by John Hartigan (Hartigan, 1972). Note that a biclustering structure

automatically implies low-rankness. However, if one applies a low rank matrix completion algorithm directly in the current setting, the resulting estimator suffers an inferior error bound to the minimax rate-optimal one. Thus, a full exploitation of the biclustering structure is necessary, which is the focus of the current paper.

The results of our paper also imply rate-optimal estimation for sparse graphons. Previous results on graphon estimation include Airoidi et al. (2013); Wolfe and Olhede (2013); Olhede and Wolfe (2014); Borgs et al. (2015); Choi (2015) and the references therein. The minimax rates for dense graphon estimation were derived by Gao et al. (2015a). During the time when this paper was written, we became aware of an independent result on optimal sparse graphon estimation by Klopp et al. (2015).

There are also an interesting line of works on biclustering (Flynn and Perry, 2012; Rohe et al., 2012; Choi and Wolfe, 2014). While these papers aim to recover the clustering structures of rows and columns, the goal of the current paper is to estimate the underlying mean matrix with optimal rates.

Organization After a brief introduction to notation, the rest of the paper is organized as follows. In Section 2, we introduce the precise formulation of the problem and propose a constrained least squares estimator for the mean matrix θ . In Section 3, we show that the proposed estimator leads to minimax optimal performance for both Gaussian and binary data. Section 4 presents some extensions of our results to sparse graphon estimation and adaptation. Implementation and simulation results are given in Section 5. In Section 6, we discuss the key points of the paper and propose some open problems for future research. The proofs of the main results are laid out in Section 7, with some auxiliary results deferred to the appendix.

Notation For a vector $z \in [k]^n$, define the set $z^{-1}(a) = \{i \in [n] : z(i) = a\}$ for $a \in [k]$. For a set S , $|S|$ denotes its cardinality and $\mathbf{1}_S$ denotes the indicator function. For a matrix $A = (A_{ij}) \in \mathbb{R}^{n_1 \times n_2}$, the ℓ_2 norm and ℓ_∞ norm are defined by $\|A\| = \sqrt{\sum_{ij} A_{ij}^2}$ and $\|A\|_\infty = \max_{ij} |A_{ij}|$, respectively. The inner product for two matrices A and B is $\langle A, B \rangle = \sum_{ij} A_{ij} B_{ij}$. Given a subset $\Omega \in [n_1] \times [n_2]$, we use the notation $\langle A, B \rangle_\Omega = \sum_{(i,j) \in \Omega} A_{ij} B_{ij}$ and $\|A\|_\Omega = \sqrt{\sum_{(i,j) \in \Omega} A_{ij}^2}$. Given two numbers $a, b \in \mathbb{R}$, we use $a \vee b = \max(a, b)$ and $a \wedge b = \min(a, b)$. The floor function $\lfloor a \rfloor$ is the largest integer no greater than a , and the ceiling function $\lceil a \rceil$ is the smallest integer no less than a . For two positive sequences $\{a_n\}, \{b_n\}$, $a_n \lesssim b_n$ means $a_n \leq C b_n$ for some constant $C > 0$ independent of n , and $a_n \asymp b_n$ means $a_n \lesssim b_n$ and $b_n \lesssim a_n$. The symbols \mathbb{P} and \mathbb{E} denote generic probability and expectation operators whose distribution is determined from the context.

2. Constrained least squares estimation

Recall the generative model defined in (1) and also the definition of the set Ω in (2) of the observed entries. As we have mentioned, throughout the paper, we assume that the ϵ_{ij} 's are independent sub-Gaussian noises with sub-Gaussianity parameter uniformly bounded from above by $\sigma > 0$. More precisely, we assume

$$\mathbb{E} e^{\lambda \epsilon_{ij}} \leq e^{\lambda^2 \sigma^2 / 2}, \quad \text{for all } i \in [n_1], j \in [n_2] \text{ and } \lambda \in \mathbb{R}. \quad (4)$$

We consider two types of biclustering structures. One is rectangular and asymmetric, where we assume that the mean matrix belongs to the following parameter space

$$\Theta_{k_1 k_2}(M) = \left\{ \theta = (\theta_{ij}) \in \mathbb{R}^{n_1 \times n_2} : \theta_{ij} = Q_{z_1(i)z_2(j)}, z_1 \in [k_1]^{n_1}, z_2 \in [k_2]^{n_2}, \right. \\ \left. Q \in [-M, M]^{k_1 \times k_2} \right\}. \quad (5)$$

In other words, the mean values within each bicluster is homogenous, i.e., $\theta_{ij} = Q_{ab}$ if the i th row belongs to the a th row cluster and the j th column belong to the b th column cluster. The other type of structures we consider is the square and symmetric case. In this case, we impose symmetry requirement on the data generating process, i.e., $n_1 = n_2 = n$ and

$$X_{ij} = X_{ji}, E_{ij} = E_{ji}, \text{ for all } i \neq j. \quad (6)$$

Since the case is mainly motivated by undirected network data where there is no edge linking any node to itself, we also assume $X_{ii} = 0$ for all $i \in [n]$. Finally, the mean matrix is assumed to belong to the following parameter space

$$\Theta_k^s(M) = \left\{ \theta = (\theta_{ij}) \in \mathbb{R}^{n \times n} : \theta_{ii} = 0, \theta_{ij} = \theta_{ji} = Q_{z(i)z(j)} \text{ for } i > j, z \in [k]^n, \right. \\ \left. Q = Q^T \in [-M, M]^{k \times k} \right\}. \quad (7)$$

We proceed by assuming that we know the parameter space Θ which can be either $\Theta_{k_1 k_2}(M)$ or $\Theta_k^s(M)$ and the rate p of an independent entry being observed. The issues of adaptation to unknown numbers of clusters and unknown observation rate p are addressed later in Section 4.1 and Section 4.2. Given Θ and p , we propose to estimate θ by the following program

$$\min_{\theta \in \Theta} \left\{ \|\theta\|^2 - \frac{2}{p} \langle X, \theta \rangle_\Omega \right\}. \quad (8)$$

If we define

$$Y_{ij} = X_{ij} E_{ij} / p, \quad (9)$$

then (8) is equivalent to the following constrained least squares problem

$$\min_{\theta \in \Theta} \|Y - \theta\|^2, \quad (10)$$

and hence the name of our estimator. When the data is binary, $\Theta = \Theta_k^s(1)$ and $p = 1$, the problem specializes to estimating the mean adjacency matrix in stochastic block models, and the estimator defined as the solution to (10) reduces to the least squares estimator in Gao et al. (2015a).

3. Main results

In this section, we provide theoretical justifications of the constrained least squares estimator defined as the solution to (10). Our first result is the following universal high probability upper bounds.

Theorem 1. *For any global optimizer of (10) and any constant $C' > 0$, there exists a constant $C > 0$ only depending on C' such that*

$$\|\hat{\theta} - \theta\|^2 \leq C \frac{M^2 \vee \sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2),$$

with probability at least $1 - \exp(-C'(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2))$ uniformly over $\theta \in \Theta_{k_1 k_2}(M)$ and all error distributions satisfying (4). For the symmetric parameter space $\Theta_k^s(M)$, the bound is simplified to

$$\|\hat{\theta} - \theta\|^2 \leq C \frac{M^2 \vee \sigma^2}{p} (k^2 + n \log k),$$

with probability at least $1 - \exp(-C'(k^2 + n \log k))$ uniformly over $\theta \in \Theta_k^s(M)$ and all error distributions satisfying (4).

When $(M^2 \vee \sigma^2)$ is bounded, the rate in Theorem 1 is $(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2)/p$ which can be decomposed into two parts. The part involving $k_1 k_2$ reflects the number of parameters in the biclustering structure, while the part involving $(n_1 \log k_1 + n_2 \log k_2)$ results from the complexity of estimating the clustering structures of rows and columns. It is the price one needs to pay for not knowing the clustering information. In contrast, the minimax rate for matrix completion under low rank assumption would be $(n_1 \vee n_2)(k_1 \wedge k_2)/p$ (Koltchinskii et al., 2011; Ma and Wu, 2015), since without any other constraint the biclustering assumption implies that the rank of the mean matrix θ is at most $k_1 \wedge k_2$. Therefore, we have $(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2)/p \ll (n_1 \vee n_2)(k_1 \wedge k_2)/p$ as long as both $n_1 \vee n_2$ and $k_1 \wedge k_2$ tend to infinity. Thus, by fully exploiting the biclustering structure, we obtain a better convergence rate than only using the low rank assumption.

In the rest of this section, we discuss two most representative cases, namely the Gaussian case and the symmetric Bernoulli case. The latter case is also known in the literature as stochastic block models.

The Gaussian case Specializing Theorem 1 to Gaussian random variables, we obtain the following result.

Corollary 2. *Assume $\epsilon_{ij} \stackrel{iid}{\sim} N(0, \sigma^2)$ and $M \leq C_1 \sigma$ for some constant $C_1 > 0$. For any constant $C' > 0$, there exists some constant C only depending on C_1 and C' such that*

$$\|\hat{\theta} - \theta\|^2 \leq C \frac{\sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2),$$

with probability at least $1 - \exp(-C'(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2))$ uniformly over $\theta \in \Theta_{k_1 k_2}(M)$. For the symmetric parameter space $\Theta_k^s(M)$, the bound is simplified to

$$\|\hat{\theta} - \theta\|^2 \leq C \frac{\sigma^2}{p} (k^2 + n \log k),$$

with probability at least $1 - \exp(-C'(k^2 + n \log k))$ uniformly over $\theta \in \Theta_k^s(M)$.

We now present a rate matching lower bound in the Gaussian model to show that the result of Corollary 2 is minimax optimal. To this end, we use $\mathbb{P}_{(\theta, \sigma^2, p)}$ to indicate the probability distribution of the model $X_{ij} \stackrel{iid}{\sim} N(\theta_{ij}, \sigma^2)$ with observation rate p .

Theorem 3. Assume $\frac{\sigma^2}{p} \left(\frac{k_1 k_2}{n_1 n_2} + \frac{\log k_1}{n_2} + \frac{\log k_2}{n_1} \right) \lesssim M^2$. There exist some constants $C, c > 0$, such that

$$\inf_{\hat{\theta}} \sup_{\theta \in \Theta_{k_1 k_2}(M)} \mathbb{P}_{(\theta, \sigma^2, p)} \left(\|\hat{\theta} - \theta\|^2 > C \frac{\sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2) \right) > c,$$

when $\log k_1 \asymp \log k_2$, and

$$\inf_{\hat{\theta}} \sup_{\theta \in \Theta_k^s(M)} \mathbb{P}_{(\theta, \sigma^2, p)} \left(\|\hat{\theta} - \theta\|^2 > C \frac{\sigma^2}{p} (k^2 + n \log k) \right) > c.$$

The symmetric Bernoulli case When the observed matrix is symmetric with zero diagonal and Bernoulli random variables as its super-diagonal entries, it can be viewed as the adjacency matrix of an undirected network and the problem of estimating its mean matrix with missing data can be viewed as a network completion problem. Given a partially observed Bernoulli adjacency matrix $\{X_{ij}\}_{(i,j) \in \Omega}$, one can predict the unobserved edges by estimating the whole mean matrix θ . Also, we assume that each edge is observed independently with probability p .

Given a symmetric adjacency matrix $X = X^T \in \{0, 1\}^{n \times n}$ with zero diagonals, the stochastic block model (Holland et al., 1983) assumes $\{X_{ij}\}_{i>j}$ are independent Bernoulli random variables with mean $\theta_{ij} = Q_{z(i)z(j)} \in [0, 1]$ with some matrix $Q \in [0, 1]^{k \times k}$ and some label vector $z \in [k]^n$. In other words, the probability that there is an edge between the i th and the j th nodes only depends on their community labels $z(i)$ and $z(j)$. The following class then includes all possible mean matrices of stochastic block models with n nodes and k clusters and with edge probabilities uniformly bounded by ρ :

$$\Theta_k^+(\rho) = \left\{ \theta \in [0, 1]^{n \times n} : \theta_{ii} = 0, \theta_{ij} = \theta_{ji} = Q_{z(i)z(j)}, Q = Q^T \in [0, \rho]^{k \times k}, z \in [k]^n \right\}. \quad (11)$$

By the definition in (7), $\Theta_k^+(\rho) \subset \Theta_k^s(\rho)$.

Although the tail probability of Bernoulli random variables does not satisfy the sub-Gaussian assumption (4), a slightly modification of the proof of Theorem 1 leads to the following result. The proof of Corollary 4 will be given in Section A in the appendix.

Corollary 4. Consider the optimization problem (10) with $\Theta = \Theta_k^s(\rho)$. For any global optimizer $\hat{\theta}$ and any constant $C' > 0$, there exists a constant $C > 0$ only depending on C' such that

$$\|\hat{\theta} - \theta\|^2 \leq C \frac{\rho}{p} (k^2 + n \log k),$$

with probability at least $1 - \exp(-C' (k^2 + n \log k))$ uniformly over $\theta \in \Theta_k^s(\rho) \supset \Theta_k^+(\rho)$.

When $\rho = p = 1$, Corollary 4 implies Theorem 2.1 in Gao et al. (2015a). A rate matching lower bound is given by the following theorem. We denote the probability distribution of a stochastic block model with mean matrix $\theta \in \Theta_k^+(\rho)$ and observation rate p by $\mathbb{P}_{(\theta, p)}$.

Theorem 5. For stochastic block models, we have

$$\inf_{\hat{\theta}} \sup_{\theta \in \Theta_k^+(\rho)} \mathbb{P}_{(\theta, p)} \left(\|\hat{\theta} - \theta\|^2 > C \left(\frac{\rho (k^2 + n \log k)}{p} \wedge \rho^2 n^2 \right) \right) > c,$$

for some constants $C, c > 0$.

The lower bound is the minimum of two terms. When $\rho \geq \frac{k^2+n \log k}{pn^2}$, the rate becomes $\frac{\rho(k^2+n \log k)}{p} \wedge \rho^2 n^2 \asymp \frac{\rho(k^2+n \log k)}{p}$. It is achieved by the constrained least squares estimator according to Corollary 4. When $\rho < \frac{k^2+n \log k}{pn^2}$, the rate is dominated by $\rho^2 n^2$. In this case, a trivial zero estimator achieves the minimax rate.

In the case of $p = 1$, a comparable result has been found independently by Klopp et al. (2015). However, our result here is more general as it accommodates missing observations. Moreover, the general upper bounds in Theorem 1 even hold for networks with weighted edges.

4. Extensions

In this section, we extend the estimation procedure and the theory in Sections 2 and 3 toward three directions: adaptation to unknown observation rate, adaptation to unknown model parameters, and sparse graphon estimation.

4.1 Adaptation to unknown observation rate

The estimator (10) depends on the knowledge of the observation rate p . When p is not too small, such a knowledge is not necessary for achieving the desired rates. Define

$$\hat{p} = \frac{\sum_{i=1}^{n_1} \sum_{j=1}^{n_2} E_{ij}}{n_1 n_2} \quad (12)$$

for the asymmetric and

$$\hat{p} = \frac{\sum_{1 \leq i < j \leq n} E_{ij}}{\frac{1}{2}n(n-1)} \quad (13)$$

for the symmetric case, and redefine

$$Y_{ij} = X_{ij} E_{ij} / \hat{p} \quad (14)$$

where the actual definition of \hat{p} is chosen between (12) and (13) depending on whether one is dealing with the asymmetric or symmetric parameter space. Then we have the following result for the solution to (10) with Y redefined by (14).

Theorem 6. For $\Theta = \Theta_{k_1 k_2}(M)$, suppose for some absolute constant $C_1 > 0$,

$$p \geq C_1 \frac{[\log(n_1 + n_2)]^2}{k_1 k_2 + n_1 \log k_1 + n_2 \log k_2}.$$

Let $\hat{\theta}$ be the solution to (10) with Y defined as in (14). Then for any constant $C' > 0$, there exists a constant $C > 0$ only depending on C' and C_1 such that

$$\|\hat{\theta} - \theta\|^2 \leq C \frac{M^2 \vee \sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2),$$

with probability at least $1 - (n_1 n_2)^{-C'}$ uniformly over $\theta \in \Theta$ and all error distributions satisfying (4).

For $\Theta = \Theta_k^s(M)$, the same result holds if we replace n_1 and n_2 with n and k_1 and k_2 with k in the foregoing statement.

4.2 Adaptation to unknown model parameters

We now provide an adaptive procedure for estimating θ without assuming the knowledge of the model parameters k_1 , k_2 and M . The procedure can be regarded as a variation of a 2-fold cross validation (Wold, 1978). We give details on the procedure for the asymmetric parameter spaces $\Theta_{k_1 k_2}(M)$, and that for the symmetric parameter spaces $\Theta_k^s(M)$ can be obtained similarly.

To adapt to k_1 , k_2 and M , we split the data into two halves. Namely, sample i.i.d. T_{ij} from Bernoulli($\frac{1}{2}$). Define $\Delta = \{(i, j) \in [n_1] \times [n_2] : T_{ij} = 1\}$. Define $Y_{ij}^\Delta = 2X_{ij}E_{ij}T_{ij}/p$ and $Y_{ij}^{\Delta^c} = 2X_{ij}E_{ij}(1 - T_{ij})/p$ for all $(i, j) \in [n_1] \times [n_2]$. Then, for some given (k_1, k_2, M) , the least squares estimators using Y^Δ and Y^{Δ^c} are given by

$$\hat{\theta}_{k_1 k_2 M}^\Delta = \underset{\theta \in \Theta_{k_1 k_2}(M)}{\operatorname{argmin}} \|Y^\Delta - \theta\|^2, \quad \hat{\theta}_{k_1 k_2 M}^{\Delta^c} = \underset{\theta \in \Theta_{k_1 k_2}(M)}{\operatorname{argmin}} \|Y^{\Delta^c} - \theta\|^2.$$

Select the parameters by

$$(\hat{k}_1, \hat{k}_2, \hat{M}) = \underset{(k_1, k_2, M) \in [n_1] \times [n_2] \times \mathcal{M}}{\operatorname{argmin}} \|\hat{\theta}_{k_1 k_2 M}^\Delta - Y^{\Delta^c}\|_{\Delta^c}^2,$$

where $\mathcal{M} = \left\{ \frac{h}{n_1 + n_2} : h \in [(n_1 + n_2)^6] \right\}$, and define $\hat{\theta}^\Delta = \hat{\theta}_{\hat{k}_1 \hat{k}_2 \hat{M}}^\Delta$. Similarly, we can also define $\hat{\theta}^{\Delta^c}$ by validate the parameters using Y^Δ . The final estimator is given by

$$\hat{\theta}_{ij} = \begin{cases} \hat{\theta}_{ij}^{\Delta^c}, & (i, j) \in \Delta; \\ \hat{\theta}_{ij}^\Delta, & (i, j) \in \Delta^c. \end{cases}$$

Theorem 7. *Assume $(n_1 + n_2)^{-1} \leq M \leq (n_1 + n_2)^5 - (n_1 + n_2)^{-1}$. For any constant $C' > 0$, there exists a constant $C > 0$ only depending on C' such that*

$$\|\hat{\theta} - \theta\|^2 \leq C \frac{M^2 \vee \sigma^2}{p} \left(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2 + \frac{\log(n_1 + n_2)}{p} \right),$$

with probability at least $1 - \exp(-C'(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2)) - (n_1 n_2)^{-C'}$ uniformly over $\theta \in \Theta_{k_1 k_2}(M)$ and all error distributions satisfying (4).

Compared with Theorem 1, the rate given by Theorem 7 has an extra $p^{-1} \log(n_1 + n_2)$ term. A sufficient condition for this extra term to be inconsequential is $p \gtrsim \frac{\log(n_1 + n_2)}{n_1 \wedge n_2}$. Theorem 7 is adaptive for all $(k_1, k_2) \in [n_1] \times [n_2]$ and for $(n_1 + n_2)^{-1} \leq M \leq (n_1 + n_2)^5 - (n_1 + n_2)^{-1}$. In fact, by choosing a larger \mathcal{M} , we can extend the adaptive region for M to $(n_1 + n_2)^{-a} \leq M \leq (n_1 + n_2)^b$ for arbitrary constants $a, b > 0$.

4.3 Sparse graphon estimation

Consider a random graph with adjacency matrix $\{X_{ij}\} \in \{0, 1\}^{n \times n}$, whose sampling procedure is determined by

$$(\xi_1, \dots, \xi_n) \sim \mathbb{P}_\xi, \quad X_{ij} | (\xi_i, \xi_j) \sim \text{Bernoulli}(\theta_{ij}), \quad \text{where } \theta_{ij} = f(\xi_i, \xi_j). \quad (15)$$

For $i \in [n]$, $X_{ii} = \theta_{ii} = 0$. Conditioning on (ξ_1, \dots, ξ_n) , $X_{ij} = X_{ji}$ is independent across $i > j$. The function f on $[0, 1]^2$, which is assumed to be symmetric, is called a graphon. The concept of graphon is originated from graph limit theory (Hoover, 1979; Lovász and Szegedy, 2006; Diaconis and Janson, 2007; Lovász, 2012) and the studies of exchangeable arrays (Aldous, 1981; Kallenberg, 1989). It is the underlying nonparametric object that generates the random graph. Statistical estimation of graphon has been considered by Wolfe and Olhede (2013); Olhede and Wolfe (2014); Gao et al. (2015a,b); Lu and Zhou (2015) for dense networks. Using Corollary 4, we present a result for sparse graphon estimation.

Let us start with specifying the function class of graphons. Define the derivative operator by

$$\nabla_{jk}f(x, y) = \frac{\partial^{j+k}}{(\partial x)^j(\partial y)^k}f(x, y),$$

and we adopt the convention $\nabla_{00}f(x, y) = f(x, y)$. The Hölder norm is defined as

$$\|f\|_{\mathcal{H}_\alpha} = \max_{j+k \leq \lfloor \alpha \rfloor} \sup_{x, y \in \mathcal{D}} |\nabla_{jk}f(x, y)| + \max_{j+k = \lfloor \alpha \rfloor} \sup_{(x, y) \neq (x', y') \in \mathcal{D}} \frac{|\nabla_{jk}f(x, y) - \nabla_{jk}f(x', y')|}{\|(x - x', y - y')\|^{\alpha - \lfloor \alpha \rfloor}},$$

where $\mathcal{D} = \{(x, y) \in [0, 1]^2 : x \geq y\}$. Then, the sparse graphon class with Hölder smoothness α is defined by

$$\mathcal{F}_\alpha(\rho, L) = \{0 \leq f \leq \rho : \|f\|_{\mathcal{H}_\alpha} \leq L\sqrt{\rho}, f(x, y) = f(y, x) \text{ for all } x \in \mathcal{D}\},$$

where $L > 0$ is the radius of the class, which is assumed to be a constant. As argued in Gao et al. (2015a), it is sufficient to approximate a graphon with Hölder smoothness by a piecewise constant function. In the random graph setting, a piecewise constant function is the stochastic block model. Therefore, we can use the estimator defined by (10). Using Corollary 4, a direct bias-variance tradeoff argument leads to the following result. An independent finding of the same result is also made by Klopp et al. (2015).

Corollary 8. *Consider the optimization problem (10) where $Y_{ij} = X_{ij}$ and $\Theta = \Theta_k^s(M)$ with $k = \lceil n^{\frac{1}{1+\alpha \wedge 1}} \rceil$ and $M = \rho$. Given any global optimizer $\hat{\theta}$ of (10), we estimate f by $\hat{f}(\xi_i, \xi_j) = \hat{\theta}_{ij}$. Then, for any constant $C' > 0$, there exists a constant $C > 0$ only depending on C' and L such that*

$$\frac{1}{n^2} \sum_{i, j \in [n]} \left(\hat{f}(\xi_i, \xi_j) - f(\xi_i, \xi_j) \right)^2 \leq C\rho \left(n^{-\frac{2\alpha}{\alpha+1}} + \frac{\log n}{n} \right),$$

with probability at least $1 - \exp(-C'(n^{\frac{1}{\alpha+1}} + n \log n))$ uniformly over $f \in \mathcal{F}_\alpha(\rho, L)$ and \mathbb{P}_ξ .

Corollary 8 implies an interesting phase transition phenomenon. When $\alpha \in (0, 1)$, the rate becomes $\rho(n^{-\frac{2\alpha}{\alpha+1}} + \frac{\log n}{n}) \asymp \rho n^{-\frac{2\alpha}{\alpha+1}}$, which is the typical nonparametric rate times a sparsity index of the network. When $\alpha \geq 1$, the rate becomes $\rho(n^{-\frac{2\alpha}{\alpha+1}} + \frac{\log n}{n}) \asymp \frac{\rho \log n}{n}$, which does not depend on the smoothness α . Corollary 8 extends Theorem 2.3 of Gao et al. (2015a) to the case $\rho < 1$. In Wolfe and Olhede (2013), the graphon f is defined in a different way. Namely, they considered the setting where (ξ_1, \dots, ξ_n) are i.i.d. Unif $[0, 1]$ random variables under \mathbb{P}_ξ . Then, the adjacency matrix is generated with

Bernoulli random variables having means $\theta_{ij} = \rho f(\xi_i, \xi_j)$ for a nonparametric graphon f satisfying $\int_0^1 \int_0^1 f(x, y) dx dy = 1$. For this setting, with appropriate smoothness assumption, we can estimate f by $\hat{f}(\xi_i, \xi_j) = \hat{\theta}_{ij}/\rho$. The rate of convergence would be $\rho^{-1}(n^{-\frac{2\alpha}{\alpha+1}} + \frac{\log n}{n})$.

Using the result of Theorem 7, we present an adaptive version for Corollary 8. The estimator we consider is a symmetric version of the one introduced in Section 4.2. The only difference is that we choose the set \mathcal{M} as $\{m/n : m \in [n+1]\}$. The estimator is fully data driven in the sense that it does not depend on α or ρ .

Corollary 9. *Assume $\rho \geq n^{-1}$. Consider the adaptive estimator $\hat{\theta}$ introduced in Theorem 7, and we set $\hat{f}(\xi_i, \xi_j) = \hat{\theta}_{ij}$. Then, for any constant $C' > 0$, there exists a constant $C > 0$ only depending on C' and L such that*

$$\frac{1}{n^2} \sum_{i,j \in [n]} \left(\hat{f}(\xi_i, \xi_j) - f(\xi_i, \xi_j) \right)^2 \leq C\rho \left(n^{-\frac{2\alpha}{\alpha+1}} + \frac{\log n}{n} \right),$$

with probability at least $1 - n^{-C'}$ uniformly over $f \in \mathcal{F}_\alpha(\rho, L)$ and \mathbb{P}_ξ .

5. Numerical Studies

To introduce an algorithm solving (8) or (10), we write (10) in an alternative way,

$$\min_{z_1 \in [k_1]^{n_1}, z_2 \in [k_2]^{n_2}, Q \in [l, u]^{k_1 \times k_2}} L(Q, z_1, z_2),$$

where l and u are the lower and upper constraints of the parameters and

$$L(Q, z_1, z_2) = \sum_{(i,j) \in [n_1] \times [n_2]} (Y_{ij} - Q_{z_1(i)z_2(j)})^2.$$

For biclustering, we set $l = -M$ and $u = M$. For SBM, we set $l = 0$ and $u = \rho$. We do not impose symmetry for SBM to gain computational convenience without losing much statistical accuracy. The simple form of $L(Q, z_1, z_2)$ indicates that we can iteratively optimize over (Q, z_1, z_2) with explicit formulas. This motivates the following algorithm.

The iteration steps (16), (17) and (18) of Algorithm 1 can be equivalently expressed as

$$\begin{aligned} Q &= \operatorname{argmin}_{Q \in [l, u]^{k_1 \times k_2}} L(Q, z_1, z_2); \\ z_1 &= \operatorname{argmin}_{z_1 \in [k_1]^{n_1}} L(Q, z_1, z_2); \\ z_2 &= \operatorname{argmin}_{z_2 \in [k_2]^{n_2}} L(Q, z_1, z_2). \end{aligned}$$

Thus, each iteration will reduce the value of the objective function. It is worth noting that Algorithm 1 can be viewed as a two-way extension for the ordinary k -means algorithm. Since the objective function is non-convex, one cannot guarantee convergence to global optimum. We initialize the algorithm via a spectral clustering step using multiple random starting points, which worked well on simulated datasets we present below.

Now we present some numerical results to demonstrate the accuracy of the error rate behavior suggested by Theorem 1 on simulated data.

Algorithm 1: A Biclustering Algorithm

Input : $\{X_{ij}\}_{(i,j) \in \Omega}$, the numbers of column and row clusters (k_1, k_2) , lower and upper constraints (l, u) and the number of random starting points m .

Output: An $n_1 \times n_2$ matrix $\hat{\theta}$ with $\hat{\theta}_{ij} = Q_{z_1(i)z_2(j)}$.

Preprocessing: Let $X_{ij} = 0$ for $(i, j) \notin \Omega$, $\hat{p} = |\Omega|/(n_1 n_2)$ and $Y = X/\hat{p}$.

1 Initialization Step

Apply singular value decomposition and obtain $Y = UDV^T$.

Run k -means algorithm on the first k_1 columns of U with m random starting points to get z_1 .

Run k -means algorithm on the first k_2 columns of V with m random starting points to get z_2 .

while not converge do

2 Update Q : for each $(a, b) \in [k_1] \times [k_2]$,

$$Q_{ab} = \frac{1}{|z_1^{-1}(a)||z_2^{-1}(b)|} \sum_{i \in z_1^{-1}(a)} \sum_{j \in z_2^{-1}(b)} Y_{ij}. \quad (16)$$

If $Q_{ab} > u$, let $Q_{ab} = u$. If $Q_{ab} < l$, let $Q_{ab} = l$.

3 Update z_1 : for each $i \in [n_1]$,

$$z_1(i) = \operatorname{argmin}_{a \in [k_1]} \sum_{j=1}^{n_2} (Q_{az_2(j)} - A_{ij})^2. \quad (17)$$

4 Update z_2 : for each $j \in [n_2]$,

$$z_2(j) = \operatorname{argmin}_{b \in [k_2]} \sum_{i=1}^{n_1} (Q_{z_1(i)b} - A_{ij})^2. \quad (18)$$

Bernoulli case. Our theoretical result indicates the rate of recovery is $\sqrt{\frac{\rho}{p} \left(\frac{k^2}{n^2} + \frac{\log k}{n} \right)}$ for the root mean squared error (RMSE) $\frac{1}{n} \|\hat{\theta} - \theta\|$. When k is not too large, the dominating term is $\sqrt{\frac{\rho \log k}{pn}}$. We are going to confirm this rate by simulation. We first generate our data from SBM with the number of blocks $k \in \{2, 4, 8, 16\}$. The observation rate $p = 0.5$. For every fixed k , we use four different $Q = 0.5\mathbf{1}_k \mathbf{1}_k^T + 0.1t\mathbf{I}_k$ with $t = 1, 2, 3, 4$ and generate the community labels z uniformly on $[k]$. Then we calculate the error $\frac{1}{n} \|\hat{\theta} - \theta\|$. Panel (a) of Figure 1 shows the error versus the sample size n . In Panel (b), we rescale the x-axis to $N = \sqrt{\frac{pn}{\log k}}$. The curves for different k align well with each other and the error decreases at the rate of $1/N$. This confirms our theoretical results in Theorem 1.

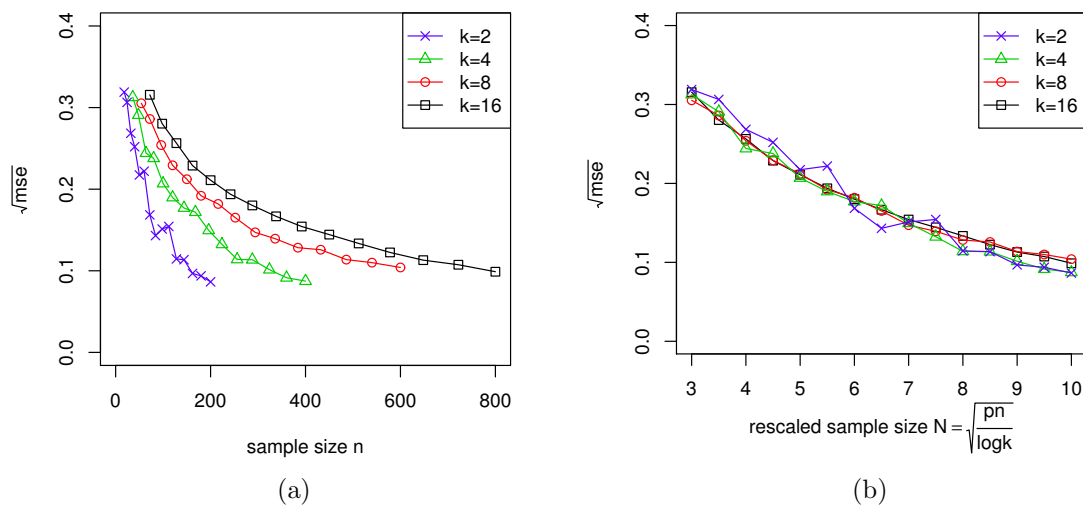


Figure 1: Plots of $\frac{1}{n} \|\hat{\theta} - \theta\|$ when using our algorithm on SBM. Each curve corresponds to a different sample size n with a fixed k . (a) Plots of error against the raw sample size n . (b) Plots of the same error against rescaled sample size $\sqrt{pn/\log k}$.

Gaussian case. We simulate data with Gaussian noise under four different settings of k_1 and k_2 . For each $(k_1, k_2) \in \{(4, 4), (4, 8), (8, 8), (8, 12)\}$, the entries of matrix Q are independently and uniformly generated from $\{1, 2, 3, 4, 5\}$. The cluster labels z_1 and z_2 are uniform on $[k_1]$ and $[k_2]$ respectively. After generating Q , z_1 and z_2 , we add an $N(0, 1)$ noise to the data and observe X_{ij} with probability $p = 0.1$. For each number of rows n_1 , we set the number of columns as $n_2 = n_1 \log k_1 / \log k_2$. Panel (a) of Figure 2 shows the error versus n_1 . In Panel (b), we rescale the x-axis by $N = \sqrt{\frac{pn_1}{\log k_2}}$. Again, the plots for different (k_1, k_2) align fairly well and the error decreases roughly at the rate of $1/N$.

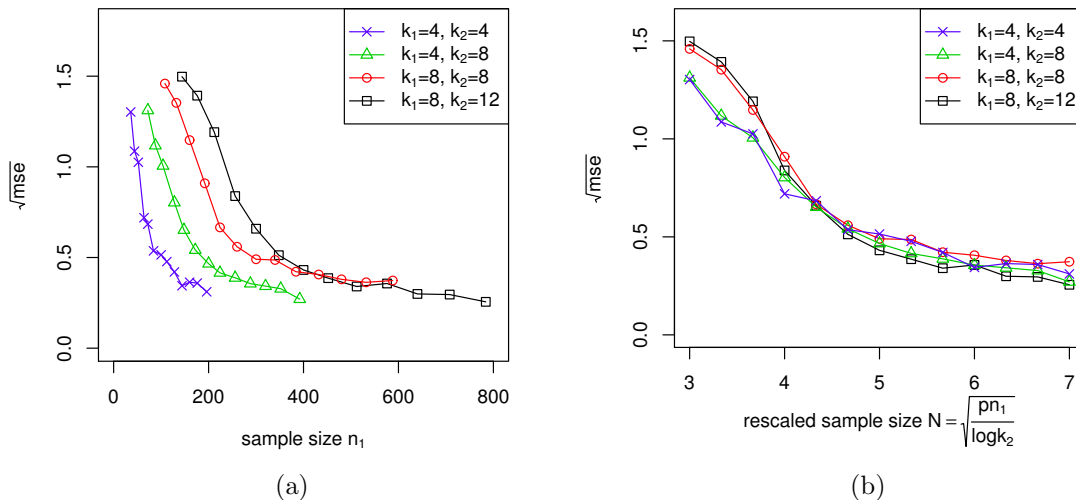


Figure 2: Plots of error $\frac{1}{\sqrt{n_1 n_2}} \|\hat{\theta} - \theta\|$ when using our algorithm on biclustering data with gaussian noise. Each curve corresponds to a fixed (k_1, k_2) . (a) Plots of error against n_1 . (b) Plots of the same error against $\sqrt{\rho n_1 / \log k_2}$.

Sparse Bernoulli case. We also study recovery of sparse SBMs. We do the same simulation as the Bernoulli case except that we choose $Q = 0.021_k \mathbf{1}_k^T + 0.05t \mathbf{I}_k$ for $t = 1, 2, 3, 4$. The results are shown in Figure 3.

Adaptation to unknown parameters. We use the 2-fold cross validation procedure proposed in Section 4.2 to adaptively choose the unknown number of clusters k and the sparsity level ρ . We use the setting of sparse SBM with the number of block $k \in \{4, 6\}$ and $Q = 0.05 \mathbf{1}_k \mathbf{1}_k^T + 0.1t \mathbf{I}_k$ for $t = 1, 2, 3, 4$. When running our algorithms, we search over all the (k, ρ) pair for $k \in \{2, 3, \dots, 8\}$ and $\rho \in \{0.2, 0.3, 0.4, 0.5\}$. In Table 1, we report the errors for different configurations of Q . The first row is the error obtained by our adaptive procedure and the second row is the error using the true k and ρ . Consistent with our Theorem 7, the error from the adaptive procedure is almost the same as the oracle error.

rescaled sample size	6	12	18	24
(k=4) adaptive \sqrt{mse}	0.084	0.066	0.058	0.058
oracle \sqrt{mse}	0.085	0.069	0.060	0.053
(k=6) adaptive \sqrt{mse}	0.074	0.061	0.051	0.050
oracle \sqrt{mse}	0.078	0.067	0.056	0.048

Table 1: Errors of the adaptive procedure versus the oracle.

The effect of constraints. The optimization (10) and Algorithm 1 involves the constraint $Q \in [l, u]^{k_1 \times k_2}$. It is curious whether this constraint really helps reduce the error or

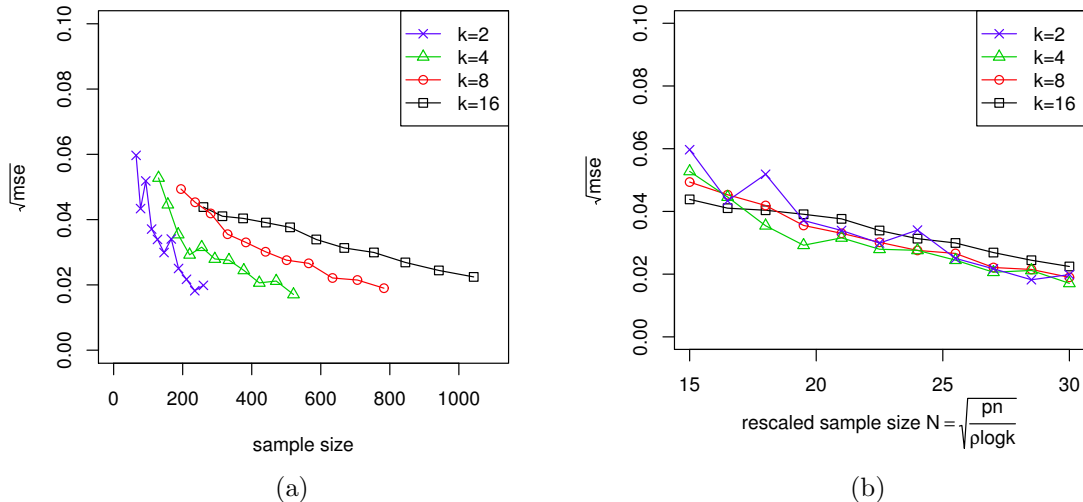


Figure 3: Plots of error $\frac{1}{n}\|\hat{\theta} - \theta\|$ when using our algorithm on sparse SBM. Each curve corresponds to a fixed k . (a) Plots of error against the raw sample size n . (b) Plots of the same error against rescaled sample size $\sqrt{pn/\rho \log k}$.

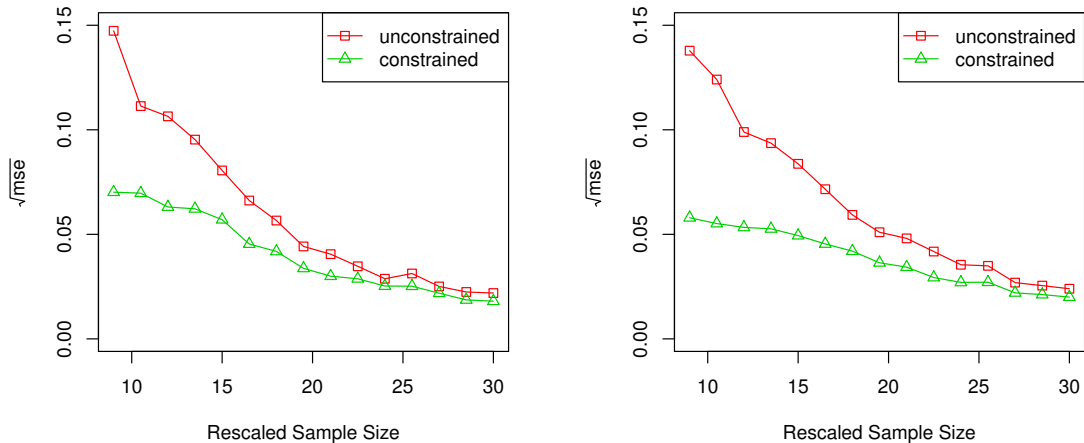
merely an artifact of the proof. We investigate the effect of this constraint on simulated data by comparing Algorithm 1 with its variation without the constraint for both Gaussian case and sparse Bernoulli case. Panel (a) of Figure 4 shows the plots of sparse SBM with 8 communities. Panel (b) is the plots of Gaussian case with $(k_1, k_2) = (4, 8)$. For both panels, when the rescaled sample size is small, the effect of constraint is significant, while as the rescaled sample size increases, the performance of two estimators become similar.

6. Discussion

This paper studies the optimal rates of recovering a matrix with biclustering structure. While the recent progresses in high-dimensional estimation mainly focus on sparse and low rank structures, the study of biclustering structure does not gain much attention. This paper fills in the gap. In what follows, we discuss some key points of the paper and some possible future directions of research.

Difference from low-rankness. A biclustering structure is implicitly low-rank. Therefore, we show that by exploring the stronger biclustering assumption, one can achieve better rates of convergence in estimation and completion. The minimax rates derived in this paper precisely characterize how much one can gain by taking advantage of this structure.

Relation to other structures. A natural question to investigate is whether there is similarity between the biclustering structure and the well-studied sparsity structure. The paper Gao et al. (2015b) gives a general theory of structured estimation in linear models that puts both sparse and biclustering structures in a unified theoretical framework. According


 (a) Sparse SBM with $k = 8$

 (b) Gaussian biclustering with $(k_1, k_2) = (4, 8)$

Figure 4: Constrained versus unconstrained least-squares

to this general theory, the $k_1 k_2$ part in the minimax rate is the complexity of parameter estimation and the $n_1 \log k_1 + n_2 \log k_2$ part is the complexity of structure estimation.

Open problems. The optimization problem (10) is not convex, thus causing difficulty in devising a provably optimal polynomial-time algorithm. An open question is whether there is a convex relaxation of (10) that can be solved efficiently without losing much statistical accuracy. Another interesting problem for future research is whether the objective function in (10) can be extended beyond the least squares framework.

7. Proofs

7.1 Proof of Theorem 1

Below, we focus on the proof for the asymmetric parameter space $\Theta_{k_1 k_2}(M)$. The result for the symmetric parameter space $\Theta_k^s(M)$ can be obtained by letting $k_1 = k_2$ and by taking care of the diagonal entries. Since $\hat{\theta} \in \Theta_{k_1 k_2}(M)$, there exists $\hat{z}_1 \in [k_1]^{n_1}$, $\hat{z}_2 \in [k_2]^{n_2}$ and $\hat{Q} \in [-M, M]^{k_1 \times k_2}$ such that $\hat{\theta}_{ij} = \hat{Q}_{\hat{z}_1(i)\hat{z}_2(j)}$. For this (\hat{z}_1, \hat{z}_2) , we define a matrix $\tilde{\theta}$ by

$$\tilde{\theta}_{ij} = \frac{1}{|\hat{z}_1^{-1}(a)||\hat{z}_2^{-1}(b)|} \sum_{(i,j) \in \hat{z}_1^{-1}(a) \times \hat{z}_2^{-1}(b)} \theta_{ij},$$

for any $(i, j) \in \hat{z}_1^{-1}(a) \times \hat{z}_2^{-1}(b)$ and any $(a, b) \in [k_1] \times [k_2]$. To facilitate the proof, we need to following three lemmas, whose proofs are given in the supplementary material.

Lemma 10. *For any constant $C' > 0$, there exists a constant $C_1 > 0$ only depending on C' , such that*

$$\|\hat{\theta} - \tilde{\theta}\|^2 \leq C_1 \frac{M^2 \vee \sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2),$$

with probability at least $1 - \exp(-C'(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2))$.

Lemma 11. For any constant $C' > 0$, there exists a constant $C_2 > 0$ only depending on C' , such that the inequality $\|\tilde{\theta} - \theta\|^2 \geq C_2(M^2 \vee \sigma^2)(n_1 \log k_1 + n_2 \log k_2)/p$ implies

$$\left| \left\langle \frac{\tilde{\theta} - \theta}{\|\tilde{\theta} - \theta\|}, Y - \theta \right\rangle \right| \leq \sqrt{C_2 \frac{M^2 \vee \sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2)},$$

with probability at least $1 - \exp(-C'(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2))$.

Lemma 12. For any constant $C' > 0$, there exists a constant $C_3 > 0$ only depending on C' , such that

$$\left| \left\langle \hat{\theta} - \tilde{\theta}, Y - \theta \right\rangle \right| \leq C_3 \frac{M^2 \vee \sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2),$$

with probability at least $1 - \exp(-C'(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2))$.

Proof of Theorem 1. Applying union bound, the results of Lemma 10-12 hold with probability at least $1 - 3 \exp(-C'(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2))$. We consider the following two cases.

Case 1: $\|\tilde{\theta} - \theta\|^2 \leq C_2(M^2 \vee \sigma^2)(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2)/p$.

Then we have

$$\|\hat{\theta} - \theta\|^2 \leq 2\|\hat{\theta} - \tilde{\theta}\|^2 + 2\|\tilde{\theta} - \theta\|^2 \leq 2(C_1 + C_2) \frac{M^2 \vee \sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2)$$

by Lemma 10.

Case 2: $\|\tilde{\theta} - \theta\|^2 > C_2(M^2 \vee \sigma^2)(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2)/p$.

By the definition of the estimator, we have $\|\hat{\theta} - Y\|^2 \leq \|\theta - Y\|^2$. After rearrangement, we have

$$\begin{aligned} \|\hat{\theta} - \theta\|^2 &\leq 2 \left\langle \hat{\theta} - \theta, Y - \theta \right\rangle \\ &= 2 \left\langle \hat{\theta} - \tilde{\theta}, Y - \theta \right\rangle + 2 \left\langle \tilde{\theta} - \theta, Y - \theta \right\rangle \\ &\leq 2 \left\langle \hat{\theta} - \tilde{\theta}, Y - \theta \right\rangle + 2\|\tilde{\theta} - \theta\| \left| \left\langle \frac{\tilde{\theta} - \theta}{\|\tilde{\theta} - \theta\|}, Y - \theta \right\rangle \right| \\ &\leq 2 \left\langle \hat{\theta} - \tilde{\theta}, Y - \theta \right\rangle + 2(\|\tilde{\theta} - \hat{\theta}\| + \|\hat{\theta} - \theta\|) \left| \left\langle \frac{\tilde{\theta} - \theta}{\|\tilde{\theta} - \theta\|}, Y - \theta \right\rangle \right| \\ &\leq 2(C_2 + C_3 + \sqrt{C_1 C_2}) \frac{M^2 \vee \sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2) + \frac{1}{2} \|\hat{\theta} - \theta\|^2, \end{aligned}$$

which leads to the bound

$$\|\hat{\theta} - \theta\|^2 \leq 4(C_2 + C_3 + \sqrt{C_1 C_2}) \frac{M^2 \vee \sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2).$$

Combining the two cases, we have

$$\|\hat{\theta} - \theta\|^2 \leq C \frac{M^2 \vee \sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2),$$

with probability at least $1 - 3 \exp(-C'(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2))$ for $C = 4(C_2 + C_3 + \sqrt{C_1 C_2}) \vee 2(C_1 + C_2)$. \square

7.2 Proof of Theorem 7

We first present a lemma for the tail behavior of sum of independent products of sub-Gaussian and Bernoulli random variables. Its proof is given in the supplementary material.

Lemma 13. *Let $\{X_i\}$ be independent sub-Gaussian random variables with mean $\theta_i \in [-M, M]$ and $\mathbb{E}e^{\lambda(X_i - \theta_i)} \leq e^{\lambda^2 \sigma^2 / 2}$. Let $\{E_i\}$ be independent Bernoulli random variables with mean p . Assume $\{X_i\}$ and $\{E_i\}$ are all independent. Then for $|\lambda| \leq p/(M \vee \sigma)$ and $Y_i = X_i E_i / p$, we have*

$$\mathbb{E}e^{\lambda(Y_i - \theta_i)} \leq 2e^{(M^2 + 2\sigma^2)\lambda^2 / p}.$$

Moreover, for $\sum_{i=1}^n c_i^2 = 1$,

$$\mathbb{P} \left\{ \left| \sum_{i=1}^n c_i (Y_i - \theta_i) \right| \geq t \right\} \leq 4 \exp \left\{ - \min \left(\frac{pt^2}{4(M^2 + 2\sigma^2)}, \frac{pt}{2(M \vee \sigma) \|c\|_\infty} \right) \right\} \quad (19)$$

for any $t > 0$.

Proof of Theorem 7. Consider the mean matrix θ that belongs to the space $\Theta_{k_1 k_2}(M)$. By the definition of $(\hat{k}_1, \hat{k}_2, \hat{M})$, we have $\|\hat{\theta}_{\hat{k}_1 \hat{k}_2 \hat{M}}^\Delta - Y^{\Delta^c}\|_{\Delta^c}^2 \leq \|\hat{\theta}_{k_1 k_2 m}^\Delta - Y^{\Delta^c}\|_{\Delta^c}^2$, where k_1 and k_2 are true numbers of row and column clusters and m is chosen to be the smallest element in \mathcal{M} that is no smaller than M . After rearrangement, we have

$$\begin{aligned} & \|\hat{\theta}_{\hat{k}_1 \hat{k}_2 \hat{M}}^\Delta - \theta\|_{\Delta^c}^2 \\ & \leq \|\hat{\theta}_{k_1 k_2 m}^\Delta - \theta\|_{\Delta^c}^2 + 2\|\hat{\theta}_{\hat{k}_1 \hat{k}_2 \hat{M}}^\Delta - \hat{\theta}_{k_1 k_2 m}^\Delta\|_{\Delta^c} \left\langle \frac{\hat{\theta}_{\hat{k}_1 \hat{k}_2 \hat{M}}^\Delta - \hat{\theta}_{k_1 k_2 m}^\Delta}{\|\hat{\theta}_{\hat{k}_1 \hat{k}_2 \hat{M}}^\Delta - \hat{\theta}_{k_1 k_2 m}^\Delta\|_{\Delta^c}}, Y^{\Delta^c} - \theta \right\rangle_{\Delta^c} \\ & \leq \|\hat{\theta}_{k_1 k_2 m}^\Delta - \theta\|_{\Delta^c}^2 + 2\|\hat{\theta}_{\hat{k}_1 \hat{k}_2 \hat{M}}^\Delta - \hat{\theta}_{k_1 k_2 m}^\Delta\|_{\Delta^c} \max_{(l_1, l_2) \in [n_1] \times [n_2]} \left| \left\langle \frac{\hat{\theta}_{l_1 l_2 h}^\Delta - \hat{\theta}_{k_1 k_2 m}^\Delta}{\|\hat{\theta}_{l_1 l_2 h}^\Delta - \hat{\theta}_{k_1 k_2 m}^\Delta\|_{\Delta^c}}, Y^{\Delta^c} - \theta \right\rangle_{\Delta^c} \right|. \end{aligned}$$

By Lemma 13 and the independence structure, we have

$$\max_{(l_1, l_2, h) \in [n_1] \times [n_2] \times \mathcal{M}} \left| \left\langle \frac{\hat{\theta}_{l_1 l_2 h}^\Delta - \hat{\theta}_{k_1 k_2 m}^\Delta}{\|\hat{\theta}_{l_1 l_2 h}^\Delta - \hat{\theta}_{k_1 k_2 m}^\Delta\|_{\Delta^c}}, Y^{\Delta^c} - \theta \right\rangle_{\Delta^c} \right| \leq C(M \vee \sigma) \frac{\log(n_1 + n_2)}{p},$$

with probability at least $1 - (n_1 n_2)^{-C'}$. Using triangle inequality and Cauchy-Schwarz inequality, we have

$$\|\hat{\theta}_{\hat{k}_1 \hat{k}_2 \hat{M}}^\Delta - \theta\|_{\Delta^c}^2 \leq \frac{3}{2} \|\hat{\theta}_{k_1 k_2 m}^\Delta - \theta\|_{\Delta^c}^2 + \frac{1}{2} \|\hat{\theta}_{\hat{k}_1 \hat{k}_2 \hat{M}}^\Delta - \theta\|_{\Delta^c}^2 + 4C^2(M^2 \vee \sigma^2) \left(\frac{\log(n_1 + n_2)}{p} \right)^2.$$

By rearranging the above inequality, we have

$$\|\hat{\theta}_{\hat{k}_1 \hat{k}_2 \hat{M}}^\Delta - \theta\|_{\Delta^c}^2 \leq 3\|\hat{\theta}_{k_1 k_2 m}^\Delta - \theta\|_{\Delta^c}^2 + 8C^2(M^2 \vee \sigma^2) \left(\frac{\log(n_1 + n_2)}{p} \right)^2.$$

A symmetric argument leads to

$$\|\hat{\theta}_{\hat{k}_1 \hat{k}_2 \hat{M}}^{\Delta^c} - \theta\|_{\Delta}^2 \leq 3\|\hat{\theta}_{k_1 k_2 m}^{\Delta^c} - \theta\|_{\Delta}^2 + 8C^2(M^2 \vee \sigma^2) \left(\frac{\log(n_1 + n_2)}{p} \right)^2.$$

Summing up the above two inequalities, we have

$$\|\hat{\theta} - \theta\|^2 \leq 3\|\hat{\theta}_{k_1 k_2 m}^\Delta - \theta\|^2 + 3\|\hat{\theta}_{k_1 k_2 m}^c - \theta\|^2 + 16C^2(M^2 \vee \sigma^2) \left(\frac{\log(n_1 + n_2)}{p} \right)^2. \quad (20)$$

Using Theorem 1 to bound $\|\hat{\theta}_{k_1 k_2 m}^\Delta - \theta\|^2$ and $\|\hat{\theta}_{k_1 k_2 m}^c - \theta\|^2$ can be bounded by $C \frac{m^2 \vee \sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2)$. Given that $m = M \left(1 + \frac{m-M}{M}\right) \leq 2M$ by the choice of m , the proof is complete. \square

7.3 Proof of Theorem 6

Recall the augmented data $Y_{ij} = X_{ij} E_{ij} / p$. Define $\mathcal{Y}_{ij} = X_{ij} E_{ij} / \hat{p}$. Let us give two lemmas to facilitate the proof.

Lemma 14. *Assume $p \gtrsim \frac{\log(n_1 + n_2)}{n_1 n_2}$. For any $C' > 0$, there is some constant $C > 0$ such that*

$$\|Y - \mathcal{Y}\|^2 \leq C [M^2 + \sigma^2 \log(n_1 + n_2)] \frac{\log(n_1 + n_2)}{p^2},$$

with probability at least $1 - (n_1 n_2)^{-C'}$.

Lemma 15. *The inequalities in Lemma 10-12 continue to hold with bounds*

$$C_1 \frac{M^2 \vee \sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2) + 2\|Y - \mathcal{Y}\|^2,$$

$$\sqrt{C_2 \frac{M^2 \vee \sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2) + \|Y - \mathcal{Y}\|},$$

and

$$C_3 \frac{M^2 \vee \sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2) + \|\hat{\theta} - \tilde{\theta}\| \|Y - \mathcal{Y}\|,$$

respectively.

Proof of Theorem 6. The proof is similar to that of Theorem 1. We only need to replace Lemma 10-12 by Lemma 14 and Lemma 15 to get the desired result. \square

7.4 Proofs of Theorem 3 and Theorem 5

This section gives proofs of the minimax lower bounds. We first introduce some notation. For any probability measures \mathbb{P}, \mathbb{Q} , define the Kullback–Leibler divergence by $D(\mathbb{P}||\mathbb{Q}) = \int \left(\log \frac{d\mathbb{P}}{d\mathbb{Q}} \right) d\mathbb{P}$. The chi-squared divergence is defined by $\chi^2(\mathbb{P}||\mathbb{Q}) = \int \left(\frac{d\mathbb{P}}{d\mathbb{Q}} \right) d\mathbb{P} - 1$. The main tool we will use is the following proposition.

Proposition 16. *Let (Ξ, ℓ) be a metric space and $\{\mathbb{P}_\xi : \xi \in \Xi\}$ be a collection of probability measures. For any totally bounded $T \subset \Xi$, define the Kullback–Leibler diameter and the chi-squared diameter of T by*

$$d_{KL}(T) = \sup_{\xi, \xi' \in T} D(\mathbb{P}_\xi || \mathbb{P}_{\xi'}), \quad d_{\chi^2}(T) = \sup_{\xi, \xi' \in T} \chi^2(\mathbb{P}_\xi || \mathbb{P}_{\xi'}).$$

Then

$$\inf_{\hat{\xi}} \sup_{\xi \in \Xi} \mathbb{P}_{\xi} \left\{ \ell^2(\hat{\xi}(X), \xi) \geq \frac{\epsilon^2}{4} \right\} \geq 1 - \frac{d_{KL}(T) + \log 2}{\log \mathcal{M}(\epsilon, T, \ell)}, \quad (21)$$

$$\inf_{\hat{\xi}} \sup_{\xi \in \Xi} \mathbb{P}_{\xi} \left\{ \ell^2(\hat{\xi}(X), \xi) \geq \frac{\epsilon^2}{4} \right\} \geq 1 - \frac{1}{\mathcal{M}(\epsilon, T, \ell)} - \sqrt{\frac{d_{\chi^2}(T)}{\mathcal{M}(\epsilon, T, \ell)}}, \quad (22)$$

for any $\epsilon > 0$, where the packing number $\mathcal{M}(\epsilon, T, \ell)$ is the largest number of points in T that are at least ϵ away from each other.

The inequality (21) is the classical Fano's inequality. The version we present here is by Yu (1997). The inequality (22) is a generalization of the classical Fano's inequality by using chi-squared divergence instead of KL divergence. It is due to Guntuboyina (2011).

The following proposition bounds the KL divergence and the chi-squared divergence for both Gaussian and Bernoulli models.

Proposition 17. *For the Gaussian model, we have*

$$D(\mathbb{P}_{(\theta, \sigma^2, p)} \| \mathbb{P}_{(\theta', \sigma^2, p)}) \leq \frac{p}{2\sigma^2} \|\theta - \theta'\|^2, \quad \chi^2(\mathbb{P}_{(\theta, \sigma^2, p)} \| \mathbb{P}_{(\theta', \sigma^2, p)}) \leq \exp\left(\frac{p}{\sigma^2} \|\theta - \theta'\|^2\right) - 1.$$

For the Bernoulli model with any $\theta, \theta' \in [\rho/2, 3\rho/4]^{n_1 \times n_2}$, we have

$$D(\mathbb{P}_{(\theta, p)} \| \mathbb{P}_{(\theta', p)}) \leq \frac{8p}{\rho} \|\theta - \theta'\|^2, \quad \chi^2(\mathbb{P}_{(\theta, p)} \| \mathbb{P}_{(\theta', p)}) \leq \exp\left(\frac{8p}{\rho} \|\theta - \theta'\|^2\right) - 1.$$

Finally, we need the following Varshamov–Gilbert bound. The version we present here is due to (Massart, 2007, Lemma 4.7).

Lemma 18. *There exists a subset $\{\omega_1, \dots, \omega_N\} \subset \{0, 1\}^d$ such that*

$$H(\omega_i, \omega_j) \triangleq \|\omega_i - \omega_j\|^2 \geq \frac{d}{4}, \quad \text{for any } i \neq j \in [N], \quad (23)$$

for some $N \geq \exp(d/8)$.

Proof of Theorem 3. We focus on the proof for the asymmetric parameter space $\Theta_{k_1 k_2}(M)$. The result for the symmetric parameter space $\Theta_k^s(M)$ can be obtained by letting $k_1 = k_2$ and by taking care of the diagonal entries. Let us assume n_1/k_1 and n_2/k_2 are integers without loss of generality. We first derive the lower bound for the nonparametric rate $\sigma^2 k_1 k_2 / p$. Let us fix the labels by $z_1(i) = \lceil ik_1/n_1 \rceil$ and $z_2(j) = \lceil jk_2/n_2 \rceil$. For any $\omega \in \{0, 1\}^{k_1 \times k_2}$, define

$$Q_{ab}^{\omega} = c \sqrt{\frac{\sigma^2 k_1 k_2}{pn_1 n_2}} \omega_{ab}. \quad (24)$$

By Lemma 18, there exists some $T \subset \{0, 1\}^{k_1 k_2}$ such that $|T| \geq \exp(k_1 k_2 / 8)$ and $H(\omega, \omega') \geq k_1 k_2 / 4$ for any $\omega, \omega' \in T$ and $\omega \neq \omega'$. We construct the subspace

$$\Theta(z_1, z_2, T) = \left\{ \theta \in \mathbb{R}^{n_1 \times n_2} : \theta_{ij} = Q_{z_1(i)z_2(j)}^{\omega}, \omega \in T \right\}.$$

By Proposition 17, we have

$$\sup_{\theta, \theta' \in \Theta(z_1, z_2, T)} \chi^2(\mathbb{P}_{(\theta, \sigma^2, p)} \| \mathbb{P}_{(\theta', \sigma^2, p)}) \leq \exp(c^2 k_1 k_2).$$

For any two different θ and θ' in $\Theta(z_1, z_2, T)$ associated with $\omega, \omega' \in T$, we have

$$\|\theta - \theta'\|^2 \geq \frac{c^2 \sigma^2}{p} H(\omega, \omega') \geq \frac{c^2 \sigma^2}{4p} k_1 k_2.$$

Therefore, $\mathcal{M}\left(\sqrt{\frac{c^2 \sigma^2}{4p} k_1 k_2}, \Theta(z_1, z_2, T), \|\cdot\|\right) \geq \exp(k_1 k_2 / 8)$. Using (22) with an appropriate c , we have obtained the rate $\frac{\sigma^2}{p} k_1 k_2$ in the lower bound.

Now let us derive the clustering rate $\sigma^2 n_2 \log k_2 / p$. Let us pick $\omega_1, \dots, \omega_{k_2} \in \{0, 1\}^{k_1}$ such that $H(\omega_a, \omega_b) \geq \frac{k_1}{4}$ for all $a \neq b$. By Lemma 18, this is possible when $\exp(k_1 / 8) \geq k_2$. Then, define

$$Q_{*a} = c \sqrt{\frac{\sigma^2 n_2 \log k_2}{p n_1 n_2}} \omega_a. \quad (25)$$

Define z_1 by $z_1(i) = \lceil i k_1 / n_1 \rceil$. Fix Q and z_1 and we are going to let z_2 vary. Select a set $\mathcal{Z}_2 \subset [k_2]^{n_2}$ such that $|\mathcal{Z}_2| \geq \exp(C n_2 \log k_2)$ and $H(z_2, z'_2) \geq \frac{n_2}{6}$ for any $z_2, z'_2 \in \mathcal{Z}_2$ and $z_2 \neq z'_2$. The existence of such \mathcal{Z}_2 is proved by Gao et al. (2015a). Then, the subspace we consider is

$$\Theta(z_1, \mathcal{Z}_2, Q) = \{\theta \in \mathbb{R}^{n_1 \times n_2} : \theta_{ij} = Q_{z_1(i) z_2(j)}, z_2 \in \mathcal{Z}_2\}.$$

By Proposition 17, we have

$$\sup_{\theta, \theta' \in \Theta(z_1, \mathcal{Z}_2, Q)} D(\mathbb{P}_{(\theta, \sigma^2, p)} \| \mathbb{P}_{(\theta', \sigma^2, p)}) \leq c^2 n_2 \log k_2.$$

For any two different θ and θ' in $\Theta(z_1, \mathcal{Z}_2, Q)$ associated with $z_2, z'_2 \in \mathcal{Z}_2$, we have

$$\|\theta - \theta'\|^2 = \sum_{j=1}^{n_2} \|\theta_{*z_2(j)} - \theta'_{*z'_2(j)}\|^2 \geq H(z_2, z'_2) \frac{c^2 \sigma^2 n_2 \log k_2}{p n_1 n_2} \frac{n_1}{4} \geq \frac{c^2 \sigma^2 n_2 \log k_2}{24p}.$$

Therefore, $\mathcal{M}\left(\sqrt{\frac{c^2 \sigma^2 n_2 \log k_2}{24p}}, \Theta(z_1, \mathcal{Z}_2, Q), \|\cdot\|\right) \geq \exp(C n_2 \log k_2)$. Using (21) with some appropriate c , we obtain the lower bound $\frac{\sigma^2 n_2 \log k_2}{p}$.

A symmetric argument gives the rate $\frac{\sigma^2 n_1 \log k_1}{p}$. Combining the three parts using the same argument in Gao et al. (2015a), the proof is complete. \square

Proof of Theorem 5. The proof is similar to that of Theorem 3. The only differences are (24) replaced by

$$Q_{ab}^\omega = \frac{1}{2} \rho + \left(c \sqrt{\frac{\rho k^2}{p n^2}} \wedge \frac{1}{2} \rho \right) \omega_{ab}$$

and (25) replaced by

$$Q_{*a} = \frac{1}{2} \rho + \left(c \sqrt{\frac{\rho \log k}{p n}} \wedge \frac{1}{2} \rho \right) \omega_a.$$

It is easy to check that the constructed subspaces are subsets of $\Theta_k^+(\rho)$. Then, a symmetric modification of the proof of Theorem 3 leads to the desired conclusion. \square

7.5 Proofs of Corollary 8 and Corollary 9

The result of Corollary 8 can be derived through a standard bias-variance trade-off argument by combining Corollary 4 and Lemma 2.1 in Gao et al. (2015a). The result of Corollary 9 follows Theorem 7. By studying the proof of Theorem 7, (20) holds for all k . Choosing the best k to trade-off bias and variance gives the result of Corollary 9. We omit the details here.

Acknowledgments

The researches of CG, YL and HHZ are supported in part by NSF grant DMS-1507511. The research of ZM is supported in part by NSF Career grant DMS-1352060.

Appendix A. Proofs of auxiliary results

In this section, we give proofs of Lemma 10-14. We first introduce some notation. Define the set

$$\mathcal{Z}_{k_1 k_2} = \{z = (z_1, z_2) : z_1 \in [k_1]^{n_1}, z_2 \in [k_2]^{n_2}\}.$$

For a matrix $G \in \mathbb{R}^{n_1 \times n_2}$ and some $z = (z_1, z_2) \in \mathcal{Z}_{k_1 k_2}$, define

$$\bar{G}_{ab}(z) = \frac{1}{|z_1^{-1}(a)||z_2^{-1}(b)|} \sum_{(i,j) \in z_1^{-1}(a) \times z_2^{-1}(b)} G_{ij},$$

for all $a \in [k_1], b \in [k_2]$. To facilitate the proof, we need the following two results.

Proposition 19. *For the estimator $\hat{\theta}_{ij} = \hat{Q}_{\hat{z}_1(i)\hat{z}_2(j)}$, we have*

$$\hat{Q}_{ab} = \text{sign}(\bar{Y}_{ab}(\hat{z})) (|\bar{Y}_{ab}(\hat{z})| \wedge M),$$

for all $a \in [k_1], b \in [k_2]$.

Lemma 20. *Under the setting of Lemma 13, define $S = \frac{1}{\sqrt{n}} \sum_{i=1}^n (Y_i - \theta_i)$ and $\tau = 2(M^2 + 2\sigma^2)/(M \vee \sigma)$. Then we have the following results:*

- Let $T = S \mathbf{1}\{|S| \leq \tau\sqrt{n}\}$, then $\mathbb{E}e^{pT^2/(8(M^2+2\sigma^2))} \leq 5$;
- Let $R = \tau\sqrt{n}|S| \mathbf{1}\{|S| > \tau\sqrt{n}\}$, then $\mathbb{E}e^{pR/(8(M^2+2\sigma^2))} \leq 9$.

Proof. By (19),

$$\mathbb{P}(|S| > t) \leq 4 \exp \left\{ - \min \left(\frac{pt^2}{4(M^2 + 2\sigma^2)}, \frac{\sqrt{n}pt}{2(M \vee \sigma)} \right) \right\}.$$

Then

$$\begin{aligned} \mathbb{E}e^{\lambda T^2} &= \int_0^\infty \mathbb{P}(e^{\lambda T^2} > u) du \leq 1 + \int_1^\infty \mathbb{P} \left(|T| > \sqrt{\frac{\log u}{\lambda}} \right) du \\ &= 1 + \int_1^{e^{\lambda \tau^2 n}} \mathbb{P} \left(|S| > \sqrt{\frac{\log u}{\lambda}} \right) du = 1 + 4 \int_1^{e^{\lambda \tau^2 n}} u^{-p/(4\lambda(M^2+2\sigma^2))} du. \end{aligned}$$

Choosing $\lambda = p/(8(M^2 + 2\sigma^2))$, we get $\mathbb{E}e^{pT^2/(8(M^2+2\sigma^2))} \leq 5$. We proceed to prove the second claim.

$$\begin{aligned}
 \mathbb{E}e^{\lambda R} &= \mathbb{P}(R = 0) + \mathbb{P}(R > 0)\mathbb{E}[e^{\lambda R}|R > 0] \\
 &= \mathbb{P}(R = 0) + \mathbb{P}(R > 0) \int_0^\infty \mathbb{P}(e^{\lambda R} > u|R > 0)du \\
 &= \mathbb{P}(R = 0) + \int_0^\infty \mathbb{P}(e^{\lambda R} > u, R > 0)du \\
 &\leq \mathbb{P}(R = 0) + \mathbb{P}(R > 0)e^{\lambda\tau^2 n} + \int_{e^{\lambda\tau^2 n}}^\infty \mathbb{P}(e^{\lambda R} > u)du \\
 &\leq 1 + 4e^{-p\tau^2 n/(4(M^2+2\sigma^2))+\lambda\tau^2 n} + \int_{e^{\lambda\tau^2 n}}^\infty \mathbb{P}(e^{\sqrt{n}\lambda\tau|S|} > u)du \\
 &= 1 + 4e^{-p\tau^2 n/(4(M^2+3\sigma^2))+\lambda\tau^2 n} + 4 \int_{e^{\lambda\tau^2 n}}^\infty u^{-p/(2\lambda\tau(M\vee\sigma))} du
 \end{aligned}$$

Choosing $\lambda = p/(8(M^2 + 2\sigma^2))$, we get $\mathbb{E}e^{pR/(8(M^2+2\sigma^2))} \leq 9$. \square

Proof of Lemma 10. By the definitions of $\hat{\theta}_{ij}$ and $\tilde{\theta}_{ij}$ and Proposition 19, we have

$$\hat{\theta}_{ij} - \tilde{\theta}_{ij} = \begin{cases} M - \bar{\theta}_{ab}(\hat{z}), & \text{if } \bar{Y}_{ab}(\hat{z}) \geq M; \\ \bar{Y}_{ab}(\hat{z}) - \bar{\theta}_{ab}(\hat{z}), & \text{if } -M \leq \bar{Y}_{ab}(\hat{z}) < M; \\ -M - \bar{\theta}_{ab}(\hat{z}), & \text{if } \bar{Y}_{ab}(\hat{z}) < -M \end{cases}$$

for any $(i, j) \in \hat{z}_1^{-1}(a) \times \hat{z}_2^{-1}(b)$. Define $W = Y - \theta$, and it is easy to check that

$$|\hat{\theta}_{ij} - \tilde{\theta}_{ij}| \leq |\bar{W}_{ab}(\hat{z})| \wedge 2M \leq |\bar{W}_{ab}(\hat{z})| \wedge \tau,$$

where $\hat{z} = (\hat{z}_1, \hat{z}_2)$ and τ is defined in Lemma 20. Then

$$\begin{aligned}
 \|\hat{\theta} - \tilde{\theta}\|^2 &\leq \sum_{a \in [k_1], b \in [k_2]} |\hat{z}_1^{-1}(a)| |\hat{z}_2^{-1}(b)| \left(|\bar{W}_{ab}(\hat{z})| \wedge \tau \right)^2 \\
 &\leq \max_{z \in \mathcal{Z}_{k_1 k_2}} \sum_{a \in [k_1], b \in [k_2]} |z_1^{-1}(a)| |z_2^{-1}(b)| \left(|\bar{W}_{ab}(z)| \wedge \tau \right)^2. \tag{26}
 \end{aligned}$$

For any $a \in [k_1], b \in [k_2]$ and $z_1 \in [k_1]^{n_1}, z_2 \in [k_2]^{n_2}$, define $n_1(a) = |z_1^{-1}(a)|$, $n_2(b) = |z_2^{-1}(b)|$ and

$$\begin{aligned}
 V_{ab}(z) &= \sqrt{n_1(a)n_2(b)} |\bar{W}_{ab}(z)| \mathbf{1}\{|\bar{W}_{ab}(z)| \leq \tau\}, \\
 R_{ab}(z) &= n_1(a)n_2(b)\tau |\bar{W}_{ab}(z)| \mathbf{1}\{|\bar{W}_{ab}(z)| > \tau\}.
 \end{aligned}$$

Then,

$$\|\hat{\theta} - \tilde{\theta}\|^2 \leq \max_{z \in \mathcal{Z}_{k_1 k_2}} \sum_{a \in [k_1], b \in [k_2]} (V_{ab}^2(z) + R_{ab}(z)). \tag{27}$$

By Markov's inequality and Lemma 20, we have

$$\begin{aligned} \mathbb{P} \left(\sum_{a \in [k_1], b \in [k_2]} V_{ab}^2(z) > t \right) &\leq e^{-pt/(8(M^2+2\sigma^2))} \prod_{a \in [k_1], b \in [k_2]} e^{pV_{ab}^2(z)/(8(M^2+2\sigma^2))} \\ &\leq \exp \left\{ -\frac{pt}{8(M^2+2\sigma^2)} + k_1 k_2 \log 5 \right\}, \end{aligned}$$

and

$$\begin{aligned} \mathbb{P} \left(\sum_{a \in [k_1], b \in [k_2]} R_{ab}(z) > t \right) &\leq e^{-pt/(8(M^2+2\sigma^2))} \prod_{a \in [k_1], b \in [k_2]} e^{pR_{ab}(z)/(8(M^2+2\sigma^2))} \\ &\leq \exp \left\{ -\frac{pt}{8(M^2+2\sigma^2)} + k_1 k_2 \log 9 \right\}, \end{aligned}$$

Applying union bound and using the fact that $\log |[k_1]^n| + \log |[k_2]^{n_2}| = n_1 \log k_1 + n_2 \log k_2$,

$$\mathbb{P} \left(\max_{z \in \mathcal{Z}_{k_1 k_2}} \sum_{a \in [k_1], b \in [k_2]} V_{ab}^2(z) > t \right) \leq \exp \left\{ -\frac{pt}{8(M^2+2\sigma^2)} + k_1 k_2 \log 5 + n_1 \log k_1 + n_2 \log k_2 \right\}.$$

For any given constant $C' > 0$, we choose $t = C_1 \frac{M^2 \vee \sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2)$ for some sufficiently large $C_1 > 0$ to obtain

$$\max_{z \in \mathcal{Z}_{k_1 k_2}} \sum_{a \in [k_1], b \in [k_2]} V_{ab}^2(z) \leq C_1 \frac{M^2 \vee \sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2) \quad (28)$$

with probability at least $1 - \exp(-C'(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2))$. Similarly, for some sufficiently large $C_2 > 0$, we have

$$\max_{z \in \mathcal{Z}_{k_1 k_2}} \sum_{a \in [k_1], b \in [k_2]} R_{ab}(z) \leq C_2 \frac{M^2 \vee \sigma^2}{p} (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2) \quad (29)$$

with probability at least $1 - \exp(-C'(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2))$. Plugging (28) and (29) into (27), we complete the proof. \square

Proof of Lemma 11. Note that

$$\tilde{\theta}_{ij} - \theta_{ij} = \sum_{a \in [k_1], b \in [k_2]} \bar{\theta}_{ab}(\hat{z}) \mathbf{1}\{(i, j) \in \hat{z}_1^{-1}(a) \times \hat{z}_2^{-1}(b)\} - \theta_{ij}$$

is a function of \hat{z}_1 and \hat{z}_2 . Then we have

$$\left| \sum_{ij} \frac{\tilde{\theta}_{ij} - \theta_{ij}}{\sqrt{\sum_{ij} (\tilde{\theta}_{ij} - \theta_{ij})^2}} (Y_{ij} - \theta_{ij}) \right| \leq \max_{z \in \mathcal{Z}_{k_1 k_2}} \left| \sum_{ij} \gamma_{ij}(z) (Y_{ij} - \theta_{ij}) \right|,$$

where

$$\gamma_{ij}(z) \propto \sum_{a \in [k_1], b \in [k_2]} \bar{\theta}_{ab}(z) \mathbf{1}\{(i, j) \in z_1^{-1}(a) \times z_2^{-1}(b)\} - \theta_{ij}$$

satisfies $\sum_{ij} \gamma_{ij}(z)^2 = 1$. Consider the event $\|\tilde{\theta} - \theta\|^2 \geq C_2(M^2 \vee \sigma^2)(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2)/p$ for some C_2 to be specified later, we have

$$|\gamma_{ij}(z)| \leq \frac{2M}{\|\tilde{\theta} - \theta\|} \leq \sqrt{\frac{4M^2 p}{C_2(M^2 \vee \sigma^2)(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2)}}.$$

By Lemma 13 and union bound, we have

$$\begin{aligned} & \mathbb{P} \left(\max_{z \in \mathcal{Z}_{k_1 k_2}} \left| \sum_{ij} \gamma_{ij}(z) (Y_{ij} - \theta_{ij}) \right| > t \right) \\ & \leq \sum_{z_1 \in [k_1]^{n_1}, z_2 \in [k_2]^{n_2}} \mathbb{P} \left(\left| \sum_{ij} \gamma_{ij}(z) (Y_{ij} - \theta_{ij}) \right| > t \right) \\ & \leq \exp(-C' (k_1 k_2 + n_1 \log k_1 + n_2 \log k_2)), \end{aligned}$$

by setting $t = \sqrt{C_2(M^2 \vee \sigma^2)(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2)/p}$ for some sufficiently large C_2 depending on C' . Thus, the lemma is proved. \square

Proof of Lemma 12. By definition,

$$\begin{aligned} & \left| \langle \hat{\theta} - \tilde{\theta}, Y - \theta \rangle \right| \\ & = \left| \sum_{a \in [k_1], b \in [k_2]} \left(\text{sign}(\bar{Y}_{ab}(\hat{z})) (|\bar{Y}_{ab}(\hat{z})| \wedge M) - \bar{\theta}_{ab}(\hat{z}) \right) \bar{W}_{ab}(\hat{z}) |\hat{z}_1^{-1}(a)| |\hat{z}_2^{-1}(b)| \right| \\ & \leq \max_{z \in \mathcal{Z}_{k_1 k_2}} \left| \sum_{a \in [k_1], b \in [k_2]} \left(\text{sign}(\bar{Y}_{ab}(z)) (|\bar{Y}_{ab}(z)| \wedge M) - \bar{\theta}_{ab}(z) \right) \bar{W}_{ab}(z) |z_1^{-1}(a)| |z_2^{-1}(b)| \right|. \end{aligned}$$

By definition, we have

$$\left(\text{sign}(\bar{Y}_{ab}(z)) (|\bar{Y}_{ab}(z)| \wedge M) - \bar{\theta}_{ab}(z) \right) \bar{W}_{ab}(z) \leq |\bar{W}_{ab}(z)|^2 \wedge \tau |\bar{W}_{ab}(z)|.$$

For any fixed $z_1 \in [k_1]^{n_1}, z_2 \in [k_2]^{n_2}$, define $n_1(a) = |z_1^{-1}(a)|$ for $a \in [k_1]$, $n_2(b) = |z_2^{-1}(b)|$ for $b \in [k_2]$ and $V_{ab}(z) = \sqrt{n_1(a)n_2(b)} |\bar{W}_{ab}(z)| \mathbf{1}\{|\bar{W}_{ab}(z)| \leq \tau\}$, $R_{ab}(z) = \tau n_1(a)n_2(b) |\bar{W}_{ab}(z)| \mathbf{1}\{|\bar{W}_{ab}(z)| > \tau\}$. Then

$$\left| \langle \hat{\theta} - \tilde{\theta}, Y - \theta \rangle \right| \leq \max_{z \in \mathcal{Z}_{k_1 k_2}} \left\{ \sum_{a \in [k_1], b \in [k_2]} V_{ab}^2(z) + \sum_{a \in [k_1], b \in [k_2]} R_{ab}(z) \right\}.$$

Following the same argument in the proof of Lemma 10, a choice of $t = C_3(M^2 \vee \sigma^2)(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2)/p$ for some sufficiently large $C_3 > 0$ will complete the proof. \square

Proof of Lemma 13. When $|\lambda| \leq p/(M \vee \sigma)$, $|\lambda\theta_i/p| \leq 1$ and $\lambda^2\sigma^2/p^2 \leq 1$. Then

$$\begin{aligned}
 \mathbb{E}e^{\lambda(Y_i - \theta_i)} &= p\mathbb{E}e^{\lambda(X/p - \theta_i)} + (1-p)\mathbb{E}e^{-\lambda\theta_i} \\
 &\leq pe^{\frac{\lambda^2\sigma^2}{2p^2} + \frac{1-p}{p}\lambda\theta_i} + (1-p)e^{-\lambda\theta_i} \\
 &\leq p\left(1 + \frac{1-p}{p}\lambda\theta_i + \frac{2(1-p)^2}{p^2}\lambda^2\theta_i^2\right)\left(1 + \frac{\lambda^2\sigma^2}{p^2}\right) + (1-p)(1 - \lambda\theta_i + 2\lambda^2\theta_i^2) \\
 &\leq 1 + \frac{2(1-p)\theta_i^2 + \sigma^2}{p}\lambda^2 + \frac{1-p}{p^2}\lambda^3\theta_i\sigma^2 + \frac{2(1-p)^2}{p^3}\lambda^4\sigma^2\theta_i^2 + 2(1-p)\lambda^2\theta_i^2 \\
 &\leq 2 + (2M^2 + \sigma^2)\lambda^2/p + \lambda^3\theta_i\sigma^2/p^2 + 2\lambda^4\sigma^2\theta_i^2/p^3 \\
 &\leq 2 + (2M^2 + \sigma^2)\lambda^2/p + \lambda^2\sigma^2/p + 2\lambda^2\sigma^2/p \\
 &\leq 2 + (2M^2 + 4\sigma^2)\lambda^2/p \\
 &\leq 2e^{(M^2 + 2\sigma^2)\lambda^2/p}.
 \end{aligned}$$

The second inequality is due to the fact that $e^x \leq 1 + 2x$ for all $x \geq 0$ and $e^x \leq 1 + x + 2x^2$ for all $|x| \leq 1$. Then for $|\lambda|(M \vee \sigma)\|c\|_\infty \leq p$, Markov inequality implies

$$\mathbb{P}\left(\sum_{i=1}^n c_i(Y_i - \theta_i) \geq t\right) \leq 2 \exp\left\{-\lambda t + \frac{\lambda^2}{p}(M^2 + 2\sigma^2)\right\}.$$

By choosing $\lambda = \min\left\{\frac{pt}{2(M^2 + 2\sigma^2)}, \frac{p}{(M \vee \sigma)\|c\|_\infty}\right\}$, we get (19). \square

Proof of Corollary 4. For independent Bernoulli random variables $X_i \sim \text{Ber}(\theta_i)$ with $\theta_i \in [0, \rho]$ for $i \in [n]$. Let $Y_i = X_i E_i/p$, where $\{E_i\}$ are independent Bernoulli random variables and $\{E_i\}$ and $\{X_i\}$ are independent. Note that $\mathbb{E}Y_i = p_i$, $\mathbb{E}Y_i^2 \leq \rho/p$ and $|Y_i| \leq 1/p$. Then Bernstein's inequality (Massart, 2007, Corollary 2.10) implies

$$\mathbb{P}\left\{\left|\frac{1}{\sqrt{n}}\sum_{i=1}^n (Y_i - \theta_i)\right| \geq t\right\} \leq 2 \exp\left\{-\min\left(\frac{pt^2}{4\rho}, \frac{3\sqrt{n}pt}{4}\right)\right\} \quad (30)$$

for any $t > 0$. Let $S = \frac{1}{\sqrt{n}}\sum_{i=1}^n (Y_i - \theta_i)$, $T = \mathbf{S1}\{|S| \leq 3\rho\sqrt{n}\}$ and $R = 3\rho\sqrt{n}\mathbf{1}\{|S| > 3\rho\sqrt{n}\}$. Following the same arguments as in the proof of Lemma 20, we have $\mathbb{E}e^{pT^2/(8\rho)} \leq 5$ and $\mathbb{E}e^{pR/(8\rho)} \leq 9$. Consequently, Lemma 10, Lemma 11 and Lemma 12 hold for the Bernoulli case. Then the rest of the proof follows from the proof of Theorem 1. \square

Proof of Lemma 14. By the definitions of Y and \mathcal{Y} , we have

$$\|Y - \mathcal{Y}\|^2 \leq (\hat{p}^{-1} - p^{-1})^2 \max_{i,j} X_{ij}^2 \sum_{ij} E_{ij}.$$

Therefore, it is sufficient to bound the three terms. For the first term, we have

$$|\hat{p}^{-1} - p^{-1}| \leq |\hat{p}^{-1} - p^{-1}| \frac{|\hat{p} - p|}{p} + \frac{|\hat{p} - p|}{p^2},$$

which leads to

$$|\hat{p}^{-1} - p^{-1}| \leq \left(1 - \frac{|\hat{p} - p|}{p}\right)^{-1} \frac{|\hat{p} - p|}{p^2}. \quad (31)$$

Bernstein's inequality implies $|\hat{p} - p|^2 \leq C \frac{p \log(n_1 + n_2)}{n_1 n_2}$ with probability at least $1 - (n_1 n_2)^{-C'}$ under the assumption that $p \gtrsim \frac{\log(n_1 + n_2)}{n_1 n_2}$. Plugging the bound into (31), we get

$$(\hat{p}^{-1} - p^{-1})^2 \leq C_1 \frac{\log(n_1 + n_2)}{p^3 n_1 n_2}.$$

The second term can be bounded by a union bound with the sub-Gaussian tail assumption of each X_{ij} . That is,

$$\max_{i,j} X_{ij}^2 \leq C_2 (M^2 + \sigma^2 \log(n_1 + n_2)),$$

with probability at least $1 - (n_1 n_2)^{-C'}$. Finally, using Bernstein's inequality again, the third term is bounded as

$$\sum_{ij} E_{ij} \leq C_3 n_1 n_2 \left(p + \sqrt{\frac{p \log(n_1 + n_2)}{n_1 n_2}} \right) \leq C'_3 n_1 n_2 p,$$

with probability at least $1 - (n_1 n_2)^{-C'}$ under the assumption that $p \gtrsim \frac{\log(n_1 + n_2)}{n_1 n_2}$. Combining the three bounds, we have obtained the desired conclusion. \square

Proof of Lemma 15. For the second and the third bounds, we use

$$\left| \left\langle \frac{\tilde{\theta} - \theta}{\|\tilde{\theta} - \theta\|}, \mathcal{Y} - \theta \right\rangle \right| \leq \left| \left\langle \frac{\tilde{\theta} - \theta}{\|\tilde{\theta} - \theta\|}, Y - \theta \right\rangle \right| + \|\mathcal{Y} - Y\|,$$

and

$$\left| \left\langle \hat{\theta} - \tilde{\theta}, \mathcal{Y} - \theta \right\rangle \right| \leq \left| \left\langle \hat{\theta} - \tilde{\theta}, Y - \theta \right\rangle \right| + \|\hat{\theta} - \tilde{\theta}\| \|Y - \mathcal{Y}\|,$$

followed by the original proofs of Lemma 11 and Lemma 12. To prove the first bound, we introduce the notation $\check{\theta}_{ij} = \check{Q}_{\hat{z}_1(i)\hat{z}_2(j)}$ with $\check{Q}_{ab} = \text{sign}(\bar{Y}_{ab}(\hat{z})) (|\bar{Y}_{ab}(\hat{z})| \wedge M)$. Recall the definition of \hat{Q} in Proposition 19 with Y replaced by \mathcal{Y} . Then, we have

$$\|\hat{\theta} - \tilde{\theta}\|^2 \leq 2\|\hat{\theta} - \check{\theta}\|^2 + 2\|\check{\theta} - \tilde{\theta}\|^2.$$

Since $\|\check{\theta} - \tilde{\theta}\|$ can be bounded by the exact argument in the proof of Lemma 10, it is sufficient to bound $\|\hat{\theta} - \check{\theta}\|^2$. By Jensen inequality,

$$\|\hat{\theta} - \check{\theta}\|^2 \leq \sum_{ab} |\hat{z}^{-1}(a)| |\hat{z}^{-1}(b)| (\bar{Y}_{ab}(\hat{z}) - \bar{\mathcal{Y}}_{ab}(\hat{z}))^2 \leq \|Y - \mathcal{Y}\|^2.$$

Thus, the proof is complete. \square

References

- E. M. Airoldi, T. B. Costa, and S. H. Chan. Stochastic blockmodel approximation of a graphon: Theory and consistent estimation. In *Advances in Neural Information Processing Systems*, pages 692–700, 2013.
- D. J. Aldous. Representations for partially exchangeable arrays of random variables. *Journal of Multivariate Analysis*, 11(4):581–598, 1981.
- C. Borgs, J. T. Chayes, H. Cohn, and S. Ganguly. Consistent nonparametric estimation for heavy-tailed sparse graphs. *arXiv preprint arXiv:1508.06675*, 2015.
- J.-F. Cai, E. J. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, 2010.
- E. J. Candès and Y. Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.
- E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.
- E. J. Candès and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *Information Theory, IEEE Transactions on*, 56(5):2053–2080, 2010.
- D. Choi. Co-clustering of nonsmooth graphons. *arXiv preprint arXiv:1507.06352*, 2015.
- D. Choi and P. J. Wolfe. Co-clustering separately exchangeable network data. *The Annals of Statistics*, 42(1):29–63, 2014.
- P. Diaconis and S. Janson. Graph limits and exchangeable random graphs. *arXiv preprint arXiv:0712.2749*, 2007.
- C. J. Flynn and P. O. Perry. Consistent biclustering. *arXiv preprint arXiv:1206.6927*, 2012.
- C. Gao, Y. Lu, and H. H. Zhou. Rate-optimal graphon estimation. *The Annals of Statistics*, 43(6):2624–2652, 2015a.
- C. Gao, A. W. van der Vaart, and H. H. Zhou. A general framework for bayes structured linear models. *arXiv preprint arXiv:1506.02174*, 2015b.
- A. Guntuboyina. Lower bounds for the minimax risk using f -divergences, and applications. *Information Theory, IEEE Transactions on*, 57(4):2386–2399, 2011.
- J. A. Hartigan. Direct clustering of a data matrix. *Journal of the American Statistical Association*, 67(337):123–129, 1972.
- P. W. Holland, K. B. Laskey, and S. Leinhardt. Stochastic blockmodels: First steps. *Social Networks*, 5(2):109–137, 1983.
- D. N. Hoover. Relations on probability spaces and arrays of random variables. *Preprint, Institute for Advanced Study, Princeton, NJ*, 2, 1979.

- O. Kallenberg. On the representation theorem for exchangeable arrays. *Journal of Multivariate Analysis*, 30(1):137–154, 1989.
- R. Keshavan, A. Montanari, and S. Oh. Matrix completion from noisy entries. In *Advances in Neural Information Processing Systems*, pages 952–960, 2009.
- R. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. *Information Theory, IEEE Transactions on*, 56(6):2980–2998, 2010.
- O. Klopp, A. B. Tsybakov, and N. Verzelen. Oracle inequalities for network models and sparse graphon estimation. *arXiv preprint arXiv:1507.04118*, 2015.
- V. Koltchinskii, K. Lounici, and A. B. Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *The Annals of Statistics*, 39(5):2302–2329, 2011.
- M. Lee, H. Shen, J. Z. Huang, and J. S. Marron. Biclustering via sparse singular value decomposition. *Biometrics*, 66(4):1087–1095, 2010.
- L. Lovász. *Large Networks and Graph Limits*, volume 60. American Mathematical Society, 2012.
- L. Lovász and B. Szegedy. Limits of dense graph sequences. *Journal of Combinatorial Theory, Series B*, 96(6):933–957, 2006.
- Y. Lu and H. H. Zhou. Minimax rates for estimating matrix products. *Preprint, Yale University*, 2015.
- Z. Ma and Y. Wu. Volume ratio, sparsity, and minimaxity under unitarily invariant norms. *Information Theory, IEEE Transactions on*, 61(12):6939–6956, 2015.
- P. Massart. *Concentration Inequalities and Model Selection*, volume 1896. Springer, 2007.
- S. C. Olhede and P. J. Wolfe. Network histograms and universality of blockmodel approximation. *Proceedings of the National Academy of Sciences*, 111(41):14722–14727, 2014.
- B. Recht. A simpler approach to matrix completion. *The Journal of Machine Learning Research*, 12:3413–3430, 2011.
- B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.
- K. Rohe, T. Qin, and B. Yu. Co-clustering for directed graphs: the stochastic co-blockmodel and spectral algorithm di-sim. *arXiv preprint arXiv:1204.2296*, 2012.
- D. B. Rubin. Inference and missing data. *Biometrika*, 63(3):581–592, 1976.
- S. Wold. Cross-validatory estimation of the number of components in factor and principal components models. *Technometrics*, 20(4):397–405, 1978.
- P. J. Wolfe and S. C. Olhede. Nonparametric graphon estimation. *arXiv preprint arXiv:1309.5936*, 2013.

B. Yu. Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam*, pages 423–435. Springer, 1997.