
Provable Overlapping Community Detection in Weighted Graphs

Jimit Majmudar
University of Waterloo
jmajmuda@uwaterloo.ca

Stephen Vavasis
University of Waterloo
vavasis@uwaterloo.ca

Abstract

Community detection is a widely-studied unsupervised learning problem in which the task is to group similar entities together based on observed pairwise entity interactions. This problem has applications in diverse domains such as social network analysis and computational biology. There is a significant amount of literature studying this problem under the assumption that the communities do not overlap. When the communities are allowed to overlap, often a *pure nodes* assumption is made, i.e. each community has a node that belongs exclusively to that community. This assumption, however, may not always be satisfied in practice. In this paper, we provide a provable method to detect overlapping communities in weighted graphs without explicitly making the pure nodes assumption. Moreover, contrary to most existing algorithms, our approach is based on convex optimization, for which many useful theoretical properties are already known. We demonstrate the success of our algorithm on artificial and real-world datasets.

1 Introduction

Given a graph, determining subsets of vertices that are closely related in some sense is a problem of interest to many researchers. The two most common titles, in unsupervised learning, for problems of such variety are “community detection in graphs” and “graph clustering”. These problems, due to their fundamental nature, arise in more than one domains. Some examples are: determining social circles in a social network (Du et al. [2007], Mishra et al. [2007], Bedi and Sharma [2016]), identifying functional modules in biological networks such as protein-protein interaction networks (Nepusz et al. [2012]), and finding groups of webpages on the World Wide Web that have content on similar topics (Dourisboure et al. [2009]).

The *Stochastic Block Model (SBM)* is a common mathematical framework for community detection, an in-depth survey of which can be found in Abbe [2017]. The SBM literature has a vast number of recovery guarantees such as those by Rohe et al. [2011], Lei et al. [2015], Li et al. [2018]. However, an obvious shortcoming of SBM is that it allows the nodes to belong to exactly one community. In practice, such an assumption is rarely satisfied. For example, in social network analysis, it is expected that some agents belong to multiple social circles or interest groups. Similarly, in the problem of clustering webpages, it is plausible that some webpages span multiple topics. Airoldi et al. [2008] proposed an extension of SBM, called the *Mixed Membership Stochastic Blockmodel (MMSB)*, in which nodes are allowed to have memberships in multiple communities. MMSB generalizes the traditional SBM by positing that each node may have fractional memberships in the different communities. If n and k denote the number of nodes and the number of communities respectively, matrix $\Theta \in [0, 1]^{n \times k}$, called the *node-community distribution matrix*, is generated such that each of its rows is drawn from the Dirichlet distribution with parameters $\alpha \in \mathbb{R}^k$. Then the $n \times n$ *probability matrix* is given as

$$P = \Theta B \Theta^T \tag{1}$$

where B is a $k \times k$ *community interaction matrix*. Lastly, a random graph according to MMSB is generated on n nodes by placing an edge between nodes i and j with probability P_{ij} . MMSB has been shown to be effective in many real-world settings, but the recovery guarantees regarding it are very limited compared to the SBM.

For a theoretical analysis of the MMSB, it is usually assumed that the user has access to only an unweighted random graph generated according to the model. While this assumption may be necessary in some settings, it makes the analysis difficult without much advantage. Indeed in many settings of practical interest, the user does have access to a similarity measure between node pairs, and this motivates us to work with weighted graphs generated according to the MMSB. For example, in the context of social network analysis, one may define a *communication graph* as an unweighted graph in which edge ij exists if and only if agents i and j exchanged messages in a certain fixed time window. Then the weighted adjacency matrix for the social network may be obtained by averaging the adjacency matrices of multiple observed communication graphs. On the other hand, we make the problem difficult in a more realistic manner; we remove a common assumption in literature which is quite unrealistic if not mathematically problematic. This assumption requires each community in the input graph to contain a node which belongs exclusively to that community. Such nodes are called *pure nodes* in the literature. The notion of pure nodes in community detection is related to that of *separability* in nonnegative matrix factorization in the sense that they both induce a simplicial structure on the data. Although we do not make the pure nodes assumption, the Dirichlet distribution naturally generates increasingly better approximations to pure nodes as n , the number of nodes, gets large, and we use this fact in our analysis. As far as we know, this exact setup has not been studied before.

Our Contributions: We provide a simple provable algorithm for the recovery of Θ in the MMSB without explicitly requiring the communities to have pure nodes. Moreover, unlike most existing methods, our algorithm enjoys the merit of being rooted in linear programming. Indeed, multiple convex relaxations exist for SBM, but that is not the case for MMSB. As a byproduct of our analysis, we provide concentration results for some key random variables associated with the MMSB. We also demonstrate the applicability of MMSB, and consequently of our algorithm, to a problem of significant consequence in computational biology which is that of *protein complex detection* via experimental results using real-world datasets.

Existing Provable Methods: Zhang et al. [2014] propose the so-called *Overlapping Communities Community Assignment Model (OCCAM)* which only slightly differs from MMSB; in OCCAM each row of Θ has unit ℓ_2 -norm as opposed to unit ℓ_1 -norm in MMSB. They provide a provable algorithm for learning the OCCAM parameters in which one performs k -medians clustering on the rows of the $n \times k$ matrix corresponding to k largest eigenvectors of the adjacency matrix corresponding to the observed unweighted random graph. However, their assumptions may be difficult to verify in practice. Indeed for their k -medians clustering to succeed, they assume that the ground-truth community structure provides the unique global optimal solution of their chosen k -medians loss function, which is also required to satisfy a special curvature condition around this minimum.

A moment-based tensor spectral approach to recover the MMSB parameters Θ and B from an unweighted random graph generated according to the model was shown by Anandkumar et al. [2014]. Their approach, however, is not very straightforward to implement and involves multiple tuning parameters. Indeed one of the tuning parameters must be close to the sum of the k Dirichlet parameters, which are not known in advance.

In a series of works, Mao et al. [2017a,b, 2018] have also tackled the problem of learning the parameters in MMSB from a single random graph generated by the model. However, they require the pure node assumption. Additionally, they cast the MMSB recovery problem as problems that are nonconvex. Consequently, to get around the nonconvexity, more assumptions on the model parameters are required. For instance, in Mao et al. [2017a], the MMSB recovery problem is formulated as a symmetric nonnegative matrix factorization (SNMF) problem, which is both nonconvex and NP-hard. Then to ensure the uniqueness of the global optimal solution for the SNMF problem, they require B to be a diagonal matrix. In contrast, not only does our approach directly tackle the factorization in (1) to recover Θ , we also do so using linear programming.

Recently Huang and Fu [2019] have also proposed a linear programming-based algorithm for recovery in MMSB. However, the connection between their proposed linear programs and ours is unclear, and they require the pure nodes assumption for their method to provably recover the communities.

Notation: For any matrix M , we use \mathbf{m}_i and \mathbf{m}^i to denote its column i and the transpose of its row i respectively, and M_{ij} to denote its entry ij ; for any set $\mathcal{R} \subseteq [n]$, $M(\mathcal{R}, :)$ denotes the submatrix of M containing all columns but only the rows indexed by \mathcal{R} . We use $\|\cdot\|$ to denote the ℓ_2 -norm for vectors and the spectral norm (largest singular value) for matrices. For a matrix, $\|\cdot\|_{\max}$ denotes its largest absolute value. I denotes the identity matrix whose dimension will be clear from the context. For any positive integer i , \mathbf{e}_i denotes column i of the identity matrix and \mathbf{e} denotes the vector with each entry set to one; the dimension of these vectors will be clear from context.

2 Proposed Work

2.1 Problem Formulation

We ask the following question for the MMSB described by (1):

Given P , how can we efficiently obtain a matrix $\hat{\Theta} \in [0, 1]^{n \times k}$ such that $\hat{\Theta} \approx \Theta$?

Typically, one imposes the pure nodes assumption on Θ which greatly simplifies the above posed problem. That is, one assumes that for each $j \in [k]$, there exists $i \in [n]$ such that $\theta^i = \mathbf{e}_j$, i.e. node i belongs exclusively to community j . In other words, the rows of Θ contain all corners of unit simplex in \mathbb{R}^k . However, such an assumption is mathematically problematic and/or practically unrealistic. Indeed if the rows of Θ are sampled from the Dirichlet distribution, then the probability of sampling even one pure node is zero. Moreover, even from a practical standpoint such an assumption may not always be satisfied since in real-world networks, such as protein-protein interaction networks, one encounters communities with no pure nodes. Lastly note that we are interested in recovering only Θ and not B since the former contains the community membership information of each node which is usually what a user of such methods is interested in.

We provide an answer to posed question without making an explicit assumption regarding the presence of pure nodes. To that effect, we propose a novel simple and efficient convex optimization-based method to approximate Θ entrywise under a very natural condition that just requires n to be sufficiently large. Such a condition is often satisfied in practice since real-world graphs in application settings such as social network analysis are usually large-scale.

Identifiability: It is known that having pure nodes for each community is both necessary and sufficient for identifiability of MMSB. Since we make no assumption regarding the presence of pure nodes in this work, we cannot necessarily expect MMSB to be identifiable. However, as a result of our analysis, we are able to show that for sufficiently large graphs, if there exist two distinct sets of parameters for the MMSB which yield the same probability matrix P , then their corresponding node-community distribution matrices are sufficiently close to each other with high probability. This notion, which is formalized in Corollary 3.2, may be interpreted as *near identifiability*.

2.2 SP+LP Recovery Algorithm

We may think of our recovery procedure, *Successive Projection followed by Linear Programming* (SP+LP), as divided into two stages. First, via a preprocessing step, called Successive Projection, we obtain a set $\mathcal{J} \subseteq [n]$ of cardinality k such that $\Theta(\mathcal{J}, :)$ is entrywise close to I up to a permutation of the rows. Then we use the nodes in \mathcal{J} to recover approximations to the k columns of Θ , the *community characteristic vectors*, using exactly k linear programs (LPs). Note that SP+LP has no tuning parameters other than the number of communities, which is also a parameter for most other community detection algorithms.

The time complexity of Successive Projection is $\mathcal{O}(n^2)$. Moreover, based on the results in Megiddo [1984], each LP in SP+LP can be solved in time $\mathcal{O}(n)$. Therefore the time complexity of SP+LP is given as $\mathcal{O}(n^2)$.

3 Theoretical Guarantees

Let Z be a $k \times k$ submatrix of Θ such that for each $j \in [k]$, there exists $i \in [k]$ satisfying

$$\|\mathbf{z}^i - \mathbf{e}_j\|_{\infty} \leq \|\boldsymbol{\theta}^p - \mathbf{e}_j\|_{\infty} \quad (2)$$

Algorithm 1 SP+LP

Input: Matrix P generated according to MMSB, number of communities k **Output:** Estimated characteristic vectors $\hat{\theta}_1, \dots, \hat{\theta}_k \in [0, 1]^n$

- 1: $\mathcal{J} = \text{SuccessiveProjection}(P)$
 - 2: **for** $i \in [k]$ **do**
 - 3: $(\mathbf{x}^*, \mathbf{y}^*) = \arg \min_{(\mathbf{x}, \mathbf{y})} \mathbf{e}^T \mathbf{x}$ s.t. $\mathbf{x} \geq \mathbf{0}, x_{\mathcal{J}(i)} \geq 1, \mathbf{x} = P\mathbf{y}$
 - 4: $\hat{\theta}_i = \mathbf{x}^* / \|\mathbf{x}^*\|_\infty$
 - 5: **end for**
-

Algorithm 2 SuccessiveProjection

Input: Matrix P generated according to MMSB, number of communities k **Output:** Estimated set of almost pure nodes $\mathcal{J} \subseteq [n]$

- 1: $\mathcal{J} = \{\}, R = P, j = 1$
 - 2: **while** $R \neq 0$ and $j \in [k]$ **do**
 - 3: $s' = \arg \max_{s \in [n]} \|\mathbf{p}_s\|^2$
 - 4:
 - 5: $R = \left(I - \frac{\mathbf{p}_{s'} \mathbf{p}_{s'}^T}{\|\mathbf{p}_{s'}\|^2} \right) R$
 - 6: $\mathcal{J} = \mathcal{J} \cup \{s'\}$
 - 7: $j = j + 1$
 - 8: **end while**
-

for any $p \in [n]$. The indices of the rows of Z in Θ , denoted by the set \mathcal{A} , correspond to nodes which we may intuitively interpret as *almost pure nodes*. Indeed the rows of Z do not exactly correspond to the corners of the unit simplex; they are, however, the best entrywise approximations of the corners that can be obtained among the rows of Θ . Note that without loss of generality, through appropriate relabelling of the nodes, we may assume that indices i and j in (2) are identical. Define the $k \times k$ matrix $\Delta := Z - I$.

Define $\mathbf{c} := \Theta^T \mathbf{e}$, and let c_{\min} and c_{\max} denote the smallest and largest entries in \mathbf{c} respectively.

Let κ and κ_0 denote the condition numbers of B and ΘB respectively, associated with the ℓ_2 -norm.

Now we state our main result, which provides complete theoretical justification for the success of SP+LP in approximately recovering the k community vectors.

Theorem 3.1. *Suppose $k \geq 2$, B is full-rank, and all k parameters of the Dirichlet distribution are equal to $\alpha \in \mathbb{R}$. Let $w := 8\kappa\sqrt{\alpha k + 1}$ and define*

$$\epsilon_1 := \min \left(\frac{1}{\sqrt{k-1}}, \frac{1}{2} \right) \frac{1}{2\sqrt{2}w(1+80w^2)}$$

$$\epsilon_2 := \frac{7}{3520\sqrt{2}kw^2}$$

If $n > \frac{\log(p/k)}{\log I_{1-\epsilon}(\alpha, (k-1)\alpha)}$ for some $p \in (0, 1)$ and $\epsilon \in (0, \min\{\epsilon_1, \epsilon_2\})$, then there exists a permutation π of the set $[k]$ such that vectors $\hat{\theta}_1, \dots, \hat{\theta}_k$ returned by SP+LP satisfy

$$\max_{j \in [k]} \|\hat{\theta}_j - \theta_{\pi(j)}\|_\infty = \mathcal{O}(\alpha k^2 \kappa^2 \epsilon) \quad (3)$$

with probability at least $1 - p - c_1 e^{-c_2 n}$ where c_1, c_2 are constants that depend on α, k, κ .

(Here $I_x(y, z)$ denotes the regularized incomplete beta function.)

We note that even though our main result is stated for an equal parameter Dirichlet distribution, our proof techniques simply extend, in principle, to a setting in which the Dirichlet parameters are different but not too far from each other. Doing so, however, adds only incremental value but makes the analysis significantly tedious.

Before presenting further results about our recovery procedure, we present a result regarding the near identifiability of the MMSB without the presence of pure nodes for each community.

Corollary 3.2. *Let (n, k, α, B) and $(n, k, \bar{\alpha}, \bar{B})$ be two distinct sets of parameters for the MMSB such that κ and $\bar{\kappa}$ denote the condition numbers of B and \bar{B} respectively, and ϵ_1, ϵ_2 and $\bar{\epsilon}_1, \bar{\epsilon}_2$ are defined respectively for the two sets as in Theorem 3.1. Moreover, suppose that these two sets satisfy the conditions of Theorem 3.1 for some p, ϵ and $\bar{p}, \bar{\epsilon}$ such that $p, \bar{p} \in (0, 1)$, $\epsilon \in (0, \min\{\epsilon_1, \epsilon_2\})$ and $\bar{\epsilon} \in (0, \min\{\bar{\epsilon}_1, \bar{\epsilon}_2\})$. If the respective node-community distribution matrices are Θ and $\bar{\Theta}$ such that $\Theta B \Theta^T = \bar{\Theta} \bar{B} \bar{\Theta}^T$, then there exists a permutation π of the set $[k]$ such that*

$$\max_{j \in [k]} \|\bar{\theta}_j - \theta_{\pi(j)}\|_{\infty} = \mathcal{O}(k^2(\alpha\kappa^2\epsilon + \bar{\alpha}\bar{\kappa}^2\bar{\epsilon})) \quad (4)$$

with probability at least $1 - p - \bar{p} - c_1 e^{-c_2 n} - \bar{c}_1 e^{-\bar{c}_2 n}$ where c_1, c_2 and \bar{c}_1, \bar{c}_2 are constants that depend on α, k, κ and $\bar{\alpha}, k, \bar{\kappa}$ respectively.

In line with our algorithm description, we divide the theoretical analysis also in two parts: one for analysis of the preprocessing Successive Projection subroutine, and another for analysis of the LPs in the main algorithm.

Successive Projection Algorithm was first studied by Gillis and Vavasis [2013] in far more generality than what is used here. Adopting their main recovery theorem to our setup yields the following theorem.

Theorem 3.3 (Gillis and Vavasis [2013]). *Suppose that*

$$\|\Delta\|_{\max} < \min\left(\frac{1}{\sqrt{k}-1}, \frac{1}{2}\right) \frac{1}{2\sqrt{2}\kappa_0(1+80\kappa_0^2)} \quad (5)$$

and let \mathcal{J} be the index set of cardinality k extracted by Algorithm 2. Then there exists a $k \times k$ permutation matrix Π such that

$$\|\Pi\Theta(\mathcal{J}, :) - I\|_{\max} \leq 40\sqrt{2}\kappa_0^2 \|\Delta\|_{\max}. \quad (6)$$

Theorem 3.3 provides theoretical justification for the success of the subroutine highlighted in Algorithm 2. To this end, our contribution is to show that the condition in (5) is satisfied in MMSB with high probability provided the number of nodes in the graph is sufficiently large. This involves deriving concentration bounds for the smallest and largest singular values of Θ and ΘB . The following result provides theoretical guarantee for the performance of the LP in Algorithm 1.

Theorem 3.4. *Assume $k \geq 2$, B is full-rank, and $\frac{c_{\min}}{c_{\max}} > \frac{1}{2}$. Suppose for each $s \in [k]$, there exists $p \in [n]$ such that $\|\theta^p - \mathbf{e}_s\|_{\infty} \leq \eta$ for some $0 \leq \eta < \frac{1}{4k} \left(\frac{c_{\min}}{c_{\max}} - \frac{1}{2}\right)$.*

Let $i \in [n]$ such that $\|\theta^i - \mathbf{e}_j\|_{\infty} \leq \eta$ for some $j \in [k]$. Then the LP

$$\begin{aligned} \min \quad & \mathbf{e}^T \mathbf{x} \\ \text{s.t.} \quad & \mathbf{x} \geq \mathbf{0} \\ & x_i \geq 1 \\ & \mathbf{x} = P\mathbf{y} \end{aligned} \quad (\text{P})$$

has an optimal solution, and if \mathbf{x}^ is an optimal solution then*

$$\left\| \frac{\mathbf{x}^*}{\|\mathbf{x}^*\|_{\infty}} - \theta_j \right\|_{\infty} \leq 4\eta(2\sqrt{2}k + 1). \quad (7)$$

Combining Theorems 3.3 and 3.4 yields Theorem 3.1, which provides entrywise error bounds for the k community characteristic vectors returned by SP+LP.

4 Experiments

In this section, we compare the performance of SP+LP on both synthetic and real-world graphs with other popular algorithms. In practice, the user has access to the adjacency matrix, called A , of the

observed weighted graph which is an approximation of P . Matrix A may even be full-rank, and so for implementation we have to slightly modify the constraint $\mathbf{x} = P\mathbf{y}$ in the LP in SP+LP. (Indeed note that if A is full-rank, then the optimal solution to the LP is $\mathbf{e}_{\mathcal{J}(i)}$.) Specifically, we replace that constraint with $\mathbf{x} = V\mathbf{y}$ where V is an $n \times k$ matrix whose columns contain the eigenvectors of A corresponding to its k largest eigenvalues. The intuition behind this is that we expect the range of V to approximate the k -dimensional subspace of \mathbb{R}^n which is the range of P . In terms of efficient computation, one may employ the Lanczos method, for instance, to compute the k largest eigenvalues of A .

4.1 Synthetic Graphs

We demonstrate the performance of SP+LP on artificial graphs generated according to the MMSB. In practice, the weighted adjacency matrix available is only approximately equal to P . Therefore for our experiments, we compute a weighted adjacency matrix by averaging s number of 0, 1-adjacency matrices, each of which is sampled according to P . That is, entry ij of a sampled adjacency matrix is a Bernoulli random variable with parameter P_{ij} . The diagonal entries in these adjacency matrices are all set to 1.

Evaluation Metrics: We evaluate SP+LP in terms of the entrywise error in the predicted columns of Θ and the wall-clock running time (Figure 1). The entrywise error is defined as $\min_{\Pi} \|\hat{\Theta} - \Theta\Pi\|_{\max}$ over all $k \times k$ permutation matrices Π , where $\hat{\Theta} := [\hat{\theta}_1 \ \dots \ \hat{\theta}_k]$ contains as columns the predicted community characteristic vectors. For each plot, each point is determined by averaging the results over 10 samples and the error bars represent one standard deviation.

We compare our results with the GeoNMF algorithm which has been shown in Mao et al. [2017a] to computationally outperform popular methods such as Stochastic Variational Inference (SVI) by Gopalan and Blei [2013], a Bayesian variant of SNMF by Psorakis et al. [2011], the OCCAM algorithm by Zhang et al. [2014], and the SAAC algorithm by Kaufmann et al. [2016]. We use the implementation of GeoNMF that is made available by the authors without any modification and also the provided default values for the tuning parameters.

Parameter Settings: Unless otherwise stated, the default parameter settings are $n = 5000, k = 3, \alpha = 0.5, s = \sqrt{n}$. Figures 1(d) and 1(f) show the performance of the SP+LP for community interaction matrices B with higher off-diagonal elements. More specifically, for those plots, we set $B = (1 - \delta) \cdot I + \delta \cdot \mathbf{e}\mathbf{e}^T$. For Figures 1(a), 1(b), 1(c), 1(e), we set $B = 0.5 \cdot I + 0.5 \cdot R$ where R is a $k \times k$ diagonal matrix whose each diagonal entry is generated from a uniform distribution over $[0, 1]$. One reason for choosing these parameter settings is to have a fair comparison. Indeed GeoNMF has already been shown to perform well over these parameter choices.

Figures 1(a), 1(b), 1(c), 1(d) demonstrate that SP+LP outperforms GeoNMF in terms of the entrywise error in the recovered MMSB communities with increasing n, k, α and δ . In particular, this implies that, compared to GeoNMF, SP+LP can handle larger graphs, more number of communities, more overlap among the communities, and a more general community interaction matrix B . However, Figure 1(e) shows that SP+LP is slower compared to GeoNMF and that opens up possibilities for future work to expedite SP+LP. On the other hand, Figure 1(f) shows that for a more general B , the time performances of GeoNMF and SP+LP are quite comparable.

4.2 Real-world Graphs

For practical application of SP+LP, we consider a well-studied problem in computational biology: that of clustering functionally similar proteins together based on protein-protein interaction (PPI) observations (see Nepusz et al. [2012]). In the language of our problem setup, each node in the weighted graph represents a protein, and the weights represent the reliability with which any two proteins interact. The communities or clusters of similar proteins are called *protein complexes* in biology literature.

It is important to highlight that the PPI networks typically contain a large number of communities compared to the number of nodes and therefore our theory does not necessarily guarantee that SP+LP will succeed with high probability. Despite that, we observe that on some datasets, SP+LP matches or even outperforms commonly-used protein complex detection heuristics. Additionally, protein complex detection is a very well-studied problem in biology and there exist a vast number

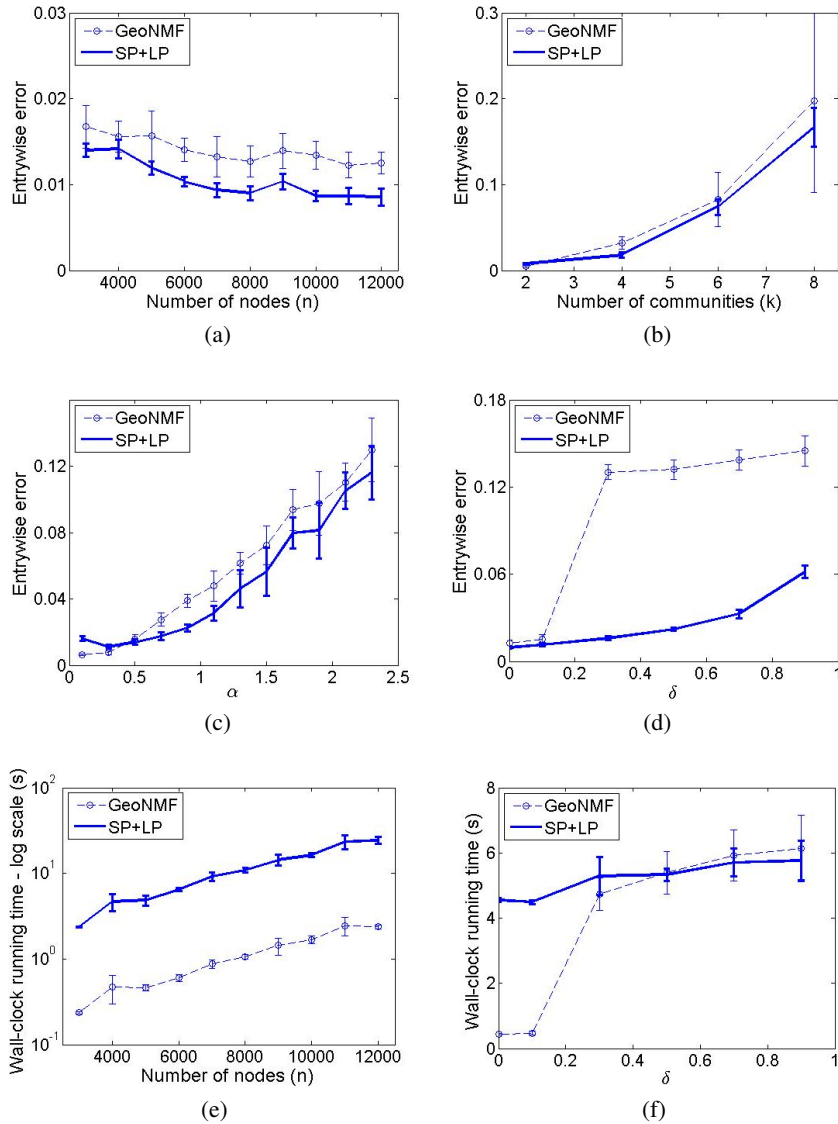


Figure 1: Performance of SP+LP on synthetic MMSB weighted graphs compared with GeoNMF.

of heuristics which are tailored for this specific problem. For instance, recent works of Yu et al. [2013], Yu et al. [2014], and Yu et al. [2015] incorporate existing ground truth knowledge of protein complexes in the algorithm to obtain a supervised learning-based approach. Our goal in this paper is not to design a fine-tuned method specifically for protein complex detection. We are focused on studying the general purpose MMSB with minimal assumptions and demonstrating its applicability to a real-world problem of immense consequence. The connection of MMSB with protein complex detection was also made in Airolidi et al. [2006]; however, their theoretical and experimental results are quite preliminary compared to ours.

Datasets: We consider PPI datasets provided by Krogan et al. [2006] and Collins et al. [2007], which are very popular among the biological community for the protein complex detection problem. The former contains two weighted graph datasets, which are referred to as Krogan core and Krogan extended. The weighted graph dataset in the latter is referred to as Collins. The ground truth validation sets used are two standard repositories of protein complexes, which also appear to be the benchmarks in the biological community. These repositories are Munich Information Centre for Protein Sequence

Table 1: Comparison of SP+LP with ClusterONE on Krogan core, Krogan extended, and Gavin datasets using SGD repository as validation set.

Validation set	Metric	Krogan core		Krogan extended		Collins	
		SP+LP	ClusterONE	SP+LP	ClusterONE	SP+LP	ClusterONE
SGD	MMR	0.389	0.418	0.428	0.364	0.372	0.532
	frac	0.598	0.667	0.632	0.594	0.557	0.828
	GA	0.525	0.663	0.542	0.628	0.504	0.731
	Score	1.512	1.748	1.602	1.586	1.433	2.091

Table 2: Comparison of SP+LP with ClusterONE on Krogan core, Krogan extended, and Gavin datasets using MIPS repository as validation set.

Validation set	Metric	Krogan core		Krogan extended		Collins	
		SP+LP	ClusterONE	SP+LP	ClusterONE	SP+LP	ClusterONE
MIPS	MMR	0.285	0.317	0.319	0.282	0.275	0.418
	frac	0.537	0.669	0.576	0.573	0.547	0.782
	GA	0.331	0.438	0.336	0.422	0.397	0.555
	Score	1.153	1.424	1.231	1.277	1.219	1.755

(MIPS) and *Saccharomyces* Genome Database (SGD). These repositories are manually curated and therefore are independent of the PPI datasets.

Evaluation Metrics: The success of a protein complex detection algorithm is typically measured via a composite score which is the sum of three quantities: maximum matching ratio (MMR), fraction of detected complexes (frac), and geometric accuracy (GA). The definitions of these quantities are non-trivial and it is beyond the scope of this paper to define them; the reader may refer to Nepusz et al. [2012] for an excellent in-depth discussion about these domain-specific quantities. A higher score corresponds to better performance and the highest possible scores for MMR and frac are one each.

The validation sets have binary memberships for the protein complexes, i.e. each protein is either present in a complex or it is not. The memberships determined via MMSB, on the other hand, are fractional. However, the former can be easily binarized by rounding all entries that are at least 0.5 to 1 and rounding the remaining entries to 0. Additionally, we have performed another post-processing step after binarizing the result of SP+LP which appears quite commonly in the domain literature. Any pair of complexes that overlap significantly (as determined by a user-defined threshold) are merged. Tables 1 and 2 show the performance of SP+LP for protein complex detection, and we compare our results with one of the most popular problem-specific heuristics called ClusterONE.

5 Conclusions

In this work, we show how to detect potentially overlapping communities in a setup that is more plausible in real-world applications, i.e. in weighted graphs without assuming the presence of pure nodes. Our method uses linear programming, which is a relatively principled approach since the literature on the theory of convex optimization is quite rich. We show that our method performs excellently on synthetic datasets. Additionally, we also show that our method succeeds in solving an important problem in computational biology without any major domain-specific modifications to the algorithm. This work entails interesting future work directions such as developing specialized linear programming solvers for our proposed algorithm, and possibly employing semidefinite programming techniques to denoise the input graph.

References

Nan Du, Bin Wu, Xin Pei, Bai Wang, and Liutong Xu. Community detection in large-scale social networks. In *Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining*

- and social network analysis, pages 16–25, 2007.
- Nina Mishra, Robert Schreiber, Isabelle Stanton, and Robert E Tarjan. Clustering social networks. In *International Workshop on Algorithms and Models for the Web-Graph*, pages 56–67. Springer, 2007.
- Punam Bedi and Chhavi Sharma. Community detection in social networks. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 6(3):115–135, 2016.
- Tamás Nepusz, Haiyuan Yu, and Alberto Paccanaro. Detecting overlapping protein complexes in protein-protein interaction networks. *Nature methods*, 9(5):471, 2012.
- Yon Dourisboure, Filippo Geraci, and Marco Pellegrini. Extraction and classification of dense implicit communities in the web graph. *ACM Transactions on the Web (TWEB)*, 3(2):1–36, 2009.
- Emmanuel Abbe. Community detection and stochastic block models: recent developments. *The Journal of Machine Learning Research*, 18(1):6446–6531, 2017.
- Karl Rohe, Sourav Chatterjee, Bin Yu, et al. Spectral clustering and the high-dimensional stochastic blockmodel. *The Annals of Statistics*, 39(4):1878–1915, 2011.
- Jing Lei, Alessandro Rinaldo, et al. Consistency of spectral clustering in stochastic block models. *The Annals of Statistics*, 43(1):215–237, 2015.
- Xiaodong Li, Yudong Chen, and Jiaming Xu. Convex relaxation methods for community detection. *arXiv preprint arXiv:1810.00315*, 2018.
- Edoardo M Airoldi, David M Blei, Stephen E Fienberg, and Eric P Xing. Mixed membership stochastic blockmodels. *Journal of machine learning research*, 9(Sep):1981–2014, 2008.
- Yuan Zhang, Elizaveta Levina, and Ji Zhu. Detecting overlapping communities in networks using spectral methods. *arXiv preprint arXiv:1412.3432*, 2014.
- Animashree Anandkumar, Rong Ge, Daniel Hsu, and Sham M Kakade. A tensor approach to learning mixed membership community models. *The Journal of Machine Learning Research*, 15(1):2239–2312, 2014.
- Xueyu Mao, Purnamrita Sarkar, and Deepayan Chakrabarti. On mixed memberships and symmetric nonnegative matrix factorizations. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2324–2333. JMLR. org, 2017a.
- Xueyu Mao, Purnamrita Sarkar, and Deepayan Chakrabarti. Estimating mixed memberships with sharp eigenvector deviations. *arXiv preprint arXiv:1709.00407*, 2017b.
- Xueyu Mao, Purnamrita Sarkar, and Deepayan Chakrabarti. Overlapping clustering models, and one (class) svm to bind them all. In *Advances in Neural Information Processing Systems*, pages 2126–2136, 2018.
- Kejun Huang and Xiao Fu. Detecting overlapping and correlated communities without pure nodes: Identifiability and algorithm. In *International Conference on Machine Learning*, pages 2859–2868, 2019.
- Nimrod Megiddo. Linear programming in linear time when the dimension is fixed. *Journal of the ACM (JACM)*, 31(1):114–127, 1984.
- Nicolas Gillis and Stephen A Vavasis. Fast and robust recursive algorithms for separable nonnegative matrix factorization. *IEEE transactions on pattern analysis and machine intelligence*, 36(4):698–714, 2013.
- Prem K Gopalan and David M Blei. Efficient discovery of overlapping communities in massive networks. *Proceedings of the National Academy of Sciences*, 110(36):14534–14539, 2013.
- Ioannis Psorakis, Stephen Roberts, Mark Ebdon, and Ben Sheldon. Overlapping community detection using bayesian non-negative matrix factorization. *Physical Review E*, 83(6):066114, 2011.

- Emilie Kaufmann, Thomas Bonald, and Marc Lelarge. A spectral algorithm with additive clustering for the recovery of overlapping communities in networks. In *International Conference on Algorithmic Learning Theory*, pages 355–370. Springer, 2016.
- Yang Yu, Xiaolong Wang, Lei Lin, Chengjie Sun, and Xuan Wang. A supervised approach to detect protein complex by combining biological and topological properties. *International journal of data mining and bioinformatics*, 8(1):105–121, 2013.
- Feng Ying Yu, Zhi Hao Yang, Nan Tang, Hong Fei Lin, Jian Wang, and Zhi Wei Yang. Predicting protein complex in protein interaction network—a supervised learning based method. *BMC systems biology*, 8(S3):S4, 2014.
- Feng Ying Yu, Zhi Hao Yang, Xiao Hua Hu, Yuan Yuan Sun, Hong Fei Lin, and Jian Wang. Protein complex detection in ppi networks based on data integration and supervised learning method. *BMC bioinformatics*, 16(12):S3, 2015.
- Edoardo M Airoldi, David M Blei, Stephen E Fienberg, Eric P Xing, and Tommi Jaakkola. Mixed membership stochastic block models for relational data with application to protein-protein interactions. In *Proceedings of the international biometrics society annual meeting*, volume 15, 2006.
- Nevan J Krogan, Gerard Cagney, Haiyuan Yu, Gouqing Zhong, Xinghua Guo, Alexandr Ignatchenko, Joyce Li, Shuye Pu, Nira Datta, Aaron P Tikuisis, et al. Global landscape of protein complexes in the yeast *saccharomyces cerevisiae*. *Nature*, 440(7084):637–643, 2006.
- Sean R Collins, Kyle M Miller, Nancy L Maas, Assen Roguev, Jeffrey Fillingham, Clement S Chu, Maya Schuldiner, Marinella Gebbia, Judith Recht, Michael Shales, et al. Functional dissection of protein complexes involved in yeast chromosome biology using a genetic interaction map. *Nature*, 446(7137):806–810, 2007.

Appendix

A LP Analysis

Let $\eta \in (0, 1)$ and assume for now that for each $j \in [k]$, there exists $i \in [n]$ such that $\|\boldsymbol{\theta}^i - \mathbf{e}_j\|_\infty \leq \eta$. Moreover, assume, without loss of generality, that for each $i \in [k]$

$$\|\boldsymbol{\theta}^i - \mathbf{e}_i\|_\infty \leq \eta. \quad (8)$$

Indeed such a property can always be satisfied with appropriate relabelling of the nodes. Define $I' := \Theta([k], \cdot)$.

Lemma A.1. *Suppose M is a $k \times k$ matrix whose rows belong to the unit simplex. If*

$$\|\mathbf{m}^i - \mathbf{e}_i\|_\infty \leq \delta \quad (9)$$

for each $i \in [k]$ and for some $\delta \in \left[0, \frac{1}{2\sqrt{2k}}\right]$, then

$$\|M^{-T} - I\|_\infty \leq 2\sqrt{2}\delta k. \quad (10)$$

Proof. Since each row of M belongs to the unit simplex and satisfies (9), we note that ℓ_2 -norm of each column of $M - I$ is bounded above by $\delta\sqrt{2}$. This implies that

$$\|M - I\| \leq \delta\sqrt{2k}. \quad (11)$$

Moreover

$$\begin{aligned} \left| \|M^{-1}\| - 1 \right| &\leq \|M^{-1} - I\| && \text{(using reverse triangle inequality)} \\ &= \|(M - I)M^{-1}\| \\ &\leq \|M - I\| \|M^{-1}\| \end{aligned}$$

which implies that

$$\|M^{-1}\| \leq \frac{1}{1 - \|M - I\|}. \quad (12)$$

Then, we have

$$\begin{aligned} \|M^{-T} - I\|_\infty &\leq \sqrt{k} \|M^{-T} - I\| \\ &= \sqrt{k} \|M^{-1} - I\| \\ &\leq \sqrt{k} \|M - I\| \|M^{-1}\| \\ &\leq \frac{\sqrt{k} \|M - I\|}{1 - \|M - I\|} && \text{(using (12))} \\ &\leq \frac{\sqrt{2}\delta k}{1 - \delta\sqrt{2k}} \\ &\leq 2\sqrt{2}\delta k. && \text{(by assumption on } \delta) \end{aligned}$$

□

For any $i \in [k]$, consider the LP

$$\begin{aligned} \min \quad & \mathbf{c}^T \mathbf{y} \\ \text{s.t.} \quad & \Theta \mathbf{y} \geq \mathbf{0} \\ & \mathbf{y}^T \boldsymbol{\theta}^i \geq 1. \end{aligned} \quad (\text{Pi})$$

and its dual

$$\begin{aligned} \max \quad & \beta \\ \text{s.t.} \quad & \beta \boldsymbol{\theta}^i + \Theta^T \mathbf{u} = \mathbf{c} \\ & \beta, \mathbf{u} \geq 0. \end{aligned} \tag{Di}$$

Note that both (Pi) and (Di) are feasible optimization problems. Thus let \mathbf{y}^* and (β^*, \mathbf{u}^*) be a (Pi)-(Di) optimal solution pair.

Lemma A.2. Suppose $\eta \leq \frac{1}{2\sqrt{2}k} \frac{c_{\min}}{c_{\max}}$.

Then

$$c_i - 2\sqrt{2}\eta k c_{\max} \leq \beta^* \leq \frac{c_i}{1 - \eta}. \tag{13}$$

Proof. The upper bound follows from observing that $\beta^* = \mathbf{c}^T \mathbf{y}^*$ due to Strong Duality and that $\mathbf{e}_i / \theta_{ii}$ is a feasible solution for (Pi), combined with the fact that $\theta_{ii} \geq 1 - \eta$.

For the lower bound we construct a feasible solution for (Di). Define \mathbf{z} as the solution of the system $I'^T \mathbf{z} = \mathbf{c}$. Note that the rows of I' belong to the unit simplex and for any $i \in [k]$, we have

$$\begin{aligned} \|I'(i, \cdot) - \mathbf{e}^i\|_{\infty} &\leq \eta \\ &\leq \frac{1}{2\sqrt{2}k}. \end{aligned} \quad (\text{by assumption on } \eta)$$

Therefore using Lemma A.1, we conclude that $\|I'^{-T} - I\|_{\infty} \leq 2\sqrt{2}\eta k$.

Then for any $s \in [k]$, we have

$$\begin{aligned} |z_s - c_s| &\leq \|\mathbf{z} - \mathbf{c}\|_{\infty} \\ &\leq \|I'^{-T} - I\|_{\infty} c_{\max} \\ &\leq 2\sqrt{2}\eta k c_{\max}. \end{aligned}$$

Moreover since $\eta \leq \frac{1}{2\sqrt{2}k} \frac{c_{\min}}{c_{\max}}$, we conclude that $\mathbf{z} \geq \mathbf{0}$. Now define the point (β', \mathbf{u}') such that

$$\beta' := z_i$$

and

$$\mathbf{u}'_s := \begin{cases} z_s, & \text{if } s \in [k] \setminus \{i\} \\ 0, & \text{otherwise.} \end{cases}$$

Note that (β', \mathbf{u}') is feasible for (Di) with objective value

$$\beta' \geq c_i - 2\sqrt{2}\eta k c_{\max}.$$

□

Define the vector $\mathbf{r} := \Theta^T \mathbf{u}^* / 2$. We shall prove some bounds on the entries of \mathbf{r} which will be used for subsequent proofs.

Lemma A.3. Suppose $\eta \leq \frac{1}{2\sqrt{2}k} \frac{c_{\min}}{c_{\max}}$.

Then we have the following inequalities.

1.

$$0 \leq r_i \leq 2k\eta c_{\max}$$

2. For any $s \in [k] \setminus \{i\}$

$$c_{\min} - \frac{\eta}{1 - \eta} c_{\max} \leq r_s \leq \frac{c_{\max}}{2}.$$

Proof. First note that $\mathbf{r} \geq \mathbf{0}$ by definition and therefore the lower bound on r_i follows. From the feasibility of (β^*, \mathbf{u}^*) for (Di), we have for any $s \in [k]$

$$r_s = \frac{c_s - \beta^* \theta_{is}}{2}. \quad (14)$$

The upper bound on r_i follows from (14), and using the lower bound on β^* from Lemma A.2 and the fact that $\theta_{ii} \geq 1 - \eta$. Indeed, we have

$$\begin{aligned} r_i &= \frac{c_i - \beta^* \theta_{ii}}{2} \\ &\leq \frac{c_i - [(c_i - 2\sqrt{2}\eta k c_{\max})(1 - \eta)]}{2} \\ &= \frac{\eta c_i + 2\sqrt{2}\eta(1 - \eta)k c_{\max}}{2} \\ &\leq \frac{\eta c_{\max}[1 + 2\sqrt{2}(1 - \eta)k]}{2} \\ &\leq \eta c_{\max} \left(\frac{1 + 3k}{2} \right) \\ &\leq 2k\eta c_{\max}. \end{aligned} \quad (\because k \geq 2)$$

For any $s \in [k] \setminus \{i\}$, the upper bound on r_s follows from (14), and noting that β^* and θ_{is} are nonnegative and $c_s \leq c_{\max}$.

For any $s \in [k] \setminus \{i\}$, the lower bound on r_s follows from (14), and using the upper bound on β^* from Lemma A.2, the fact that $c_s \geq c_{\min}$ and the fact that $\theta_{is} \leq \eta$. \square

Lemma A.4. *Suppose $\eta \leq \frac{1}{3k} \frac{c_{\min}}{c_{\max}}$. Then $\|\mathbf{r}\|_{\infty} \leq \frac{c_{\max}}{2}$.*

Proof. We prove this statement by proving that $\|\mathbf{r}\|_{\infty}$ is attained at some index in $[k] \setminus \{i\}$. It suffices to show that $r_i \leq r_s$ for any $s \in [k] \setminus \{i\}$. Note that by assumption $\eta \leq \frac{1}{3k} \frac{c_{\min}}{c_{\max}} \leq \frac{1}{2\sqrt{2}} \frac{c_{\min}}{c_{\max}}$ and therefore the entries of \mathbf{r} are bounded according to Lemma A.3.

We have

$$\begin{aligned} c_{\min} &\geq 2\eta c_{\max} \frac{3k}{2} && \text{(by assumption on } \eta) \\ &\geq 2\eta c_{\max}(k + 1) && (\because k \geq 2) \\ &= 2\eta c_{\max} + 2k\eta c_{\max} \\ &\geq \frac{\eta}{1 - \eta} c_{\max} + 2k\eta c_{\max} && (\because \eta \leq 1/2) \end{aligned}$$

which is equivalent to

$$2k\eta c_{\max} \leq c_{\min} - \frac{\eta}{1 - \eta} c_{\max}.$$

Therefore using Lemma A.3, we conclude that $r_i \leq r_s$ for any $s \in [k] \setminus \{i\}$. \square

Lemma A.5. *Suppose $\frac{c_{\min}}{c_{\max}} > \frac{1}{2}$ and $\eta < \frac{1}{3k} \left(\frac{c_{\min}}{c_{\max}} - \frac{1}{2} \right)$. Then for any $s \in [k] \setminus \{i\}$, if y_s^* is positive, we have*

$$y_s^* < 2\sqrt{2}\eta k. \quad (15)$$

Proof. Pick any $s \in [k] \setminus \{i\}$ such that $y_s^* > 0$. Consider the auxiliary LP

$$\begin{aligned} \min \quad & \mathbf{c}^T \mathbf{y} \\ \text{s.t.} \quad & \Theta \mathbf{y} \geq \mathbf{0} \\ & \mathbf{y}^T \boldsymbol{\theta}^i \geq 1 \\ & y_s \geq 2\sqrt{2}\eta k \end{aligned} \tag{Pi-aux}$$

and its dual

$$\begin{aligned} \max \quad & \beta + (2\sqrt{2}\eta k)\gamma \\ \text{s.t.} \quad & \beta \boldsymbol{\theta}^i + \gamma \mathbf{e}_s + \Theta^T \mathbf{u} = \mathbf{c} \\ & \beta, \gamma, \mathbf{u} \geq 0. \end{aligned} \tag{Di-aux}$$

If we show that \mathbf{y}^* is not an optimal solution to (Pi-aux), then we can conclude that $y_s^* < 2\sqrt{2}\eta k$. Therefore our goal is to show that the optimal value of (Pi-aux) is greater than $\mathbf{c}^T \mathbf{y}^*$. Equivalently, we may also show that the optimal value of (Di-aux) is greater than β^* . We do so by constructing a feasible solution for (Di-aux) at which the objective value is greater than β^* .

Now define \bar{I} to be identical to I' except the s^{th} row which is set to be \mathbf{e}_s^T . Let \mathbf{z}^* be the solution to the system

$$\bar{I}^T \mathbf{z} = \mathbf{r} \tag{16}$$

where recall that $\mathbf{r} = \Theta^T \mathbf{u}^* / 2$.

Note that the rows of \bar{I} belong to the unit simplex and for any $i \in [k]$, we have

$$\begin{aligned} \|\bar{I}(i, \cdot) - \mathbf{e}^i\|_\infty &\leq \eta \\ &\leq \frac{1}{2\sqrt{2}k}. \end{aligned} \tag{by assumption on } \eta$$

Therefore using Lemma A.1, we conclude that

$$\|\bar{I}^{-T} - I\|_\infty \leq 2\sqrt{2}\eta k. \tag{17}$$

Define the point

$$\begin{bmatrix} \bar{\beta} \\ \bar{\gamma} \\ \bar{\mathbf{u}} \end{bmatrix} := \begin{bmatrix} \beta^* \\ 0 \\ \mathbf{u}^*/2 \end{bmatrix} + \begin{bmatrix} \beta' \\ \gamma' \\ \mathbf{u}' \end{bmatrix} \tag{18}$$

where $\beta' := z_i^*$, $\gamma' := z_s^*$ and

$$u'_p := \begin{cases} z_p^* & \text{if } p \in [k] \setminus \{i, s\} \\ 0 & \text{otherwise.} \end{cases}$$

First we argue that $(\bar{\beta}, \bar{\gamma}, \bar{\mathbf{u}})$ is feasible for (Di-aux). From (18), we have

$$\begin{aligned} \bar{\beta} \boldsymbol{\theta}^i + \bar{\gamma} \mathbf{e}_s + \Theta^T \bar{\mathbf{u}} &= \beta^* \boldsymbol{\theta}^i + \Theta^T \mathbf{u}^* / 2 + \beta' \boldsymbol{\theta}^i + \gamma' \mathbf{e}_s + \Theta^T \mathbf{u}' \\ &= \mathbf{c} - \mathbf{r} + \beta' \boldsymbol{\theta}^i + \gamma' \mathbf{e}_s + \Theta^T \mathbf{u}' && (\because (\beta^*, \mathbf{u}^*) \text{ is feasible for (Di)}) \\ &= \mathbf{c} - \mathbf{r} + \bar{I}_k^T \mathbf{z}^* && (\text{using the definition of } (\beta', \gamma', \mathbf{u}')) \\ &= \mathbf{c}. && (\text{using (16)}) \end{aligned}$$

To argue about the nonnegativity of $(\bar{\beta}, \bar{\gamma}, \bar{\mathbf{u}})$, it suffices to argue that

1. $z_i^* + \beta^* \geq 0$
2. $\mathbf{z}^*([k] \setminus \{i\}) \geq \mathbf{0}$.

Note that our assumption on η implies $\eta < \frac{1}{2\sqrt{2}k} \frac{c_{\min}}{c_{\max}}$ and therefore Lemmas A.2 and A.3 apply.

We have

$$\begin{aligned}
z_i^* &= \bar{I}^{-T}(i, i)r_i + \sum_{p \in [k] \setminus \{i\}} \bar{I}^{-T}(i, p)r_p \\
&\geq 0 + \sum_{p \in [k] \setminus \{i\}} \bar{I}^{-T}(i, p)r_p && (\because \bar{I}^{-T}(i, i) \geq 0, r_i \geq 0) \\
&\geq -2\sqrt{2}\eta k \frac{c_{\max}}{2}. && (\text{using (17) and Lemma A.3})
\end{aligned} \tag{19}$$

Combining the lower bound on z_i^* with the lower bound on β^* in Lemma A.2 we get

$$\begin{aligned}
z_i^* + \beta^* &\geq c_i - 3\sqrt{2}\eta k c_{\max} \\
&\geq c_{\min} - 3\sqrt{2}\eta k c_{\max} \\
&> 0.
\end{aligned}$$

The last inequality above follows from our assumption on η . Indeed, we have

$$\begin{aligned}
\eta &< \frac{1}{3k} \left(\frac{c_{\min}}{c_{\max}} - \frac{1}{2} \right) \\
&< \frac{1}{3\sqrt{2}k} \frac{c_{\min}}{c_{\max}}. && \left(\because \frac{c_{\min}}{c_{\max}} \leq 1 \right)
\end{aligned}$$

Similarly, for any $t \in [k] \setminus \{i\}$ we have

$$\begin{aligned}
z_t^* &\geq r_t - \|\bar{I}^{-T} - I\|_{\infty} \|\mathbf{r}\|_{\infty} && (\text{using (16)}) \\
&\geq r_t - 2\sqrt{2}\eta k \frac{c_{\max}}{2} && (\text{using (17) and Lemma A.4}) \\
&\geq c_{\min} - \frac{\eta}{1-\eta} c_{\max} - 2\sqrt{2}\eta k \frac{c_{\max}}{2}. && (\text{using Lemma A.3})
\end{aligned} \tag{20}$$

Our assumption on η yields a positive lower bound on the above expression. Indeed, we have

$$\begin{aligned}
c_{\min} &> \frac{c_{\max}}{2} + 3k\eta c_{\max} && (\text{by assumption on } \eta) \\
&\geq \frac{c_{\max}}{2} + 2(k+1)\eta c_{\max} && (\because k \geq 2) \\
&= \frac{c_{\max}}{2} + 2\eta c_{\max} + 2\eta k c_{\max} \\
&\geq \frac{c_{\max}}{2} + \frac{\eta}{1-\eta} c_{\max} + \sqrt{2}\eta k c_{\max} && (\because \eta \leq 1/2)
\end{aligned}$$

Using the above in (20), we get

$$z_t^* > c_{\max}/2. \tag{21}$$

Therefore $(\bar{\beta}, \bar{\gamma}, \bar{\mathbf{u}})$ is feasible for (Di-aux).

Now we argue that the objective value of (Di-aux) at $(\bar{\beta}, \bar{\gamma}, \bar{\mathbf{u}})$ is greater than β^* . Indeed note that

$$\begin{aligned}
\beta' + (2\sqrt{2}\eta k)\gamma' &= z_i^* + (2\sqrt{2}\eta k)z_s^* \\
&> -\sqrt{2}\eta k c_{\max} + 2\sqrt{2}\eta k \frac{c_{\max}}{2} && (\text{using (19) and (21)}) \\
&= 0.
\end{aligned}$$

That is, $\beta' + (2\sqrt{2}\eta k)\gamma' > 0$ or equivalently, $\bar{\beta} + (2\sqrt{2}\eta k)\bar{\gamma} > \beta^*$ thereby concluding the proof. \square

Lemma A.6. Suppose $\frac{c_{\min}}{c_{\max}} > \frac{1}{2}$ and $\eta < \frac{1}{4k} \left(\frac{c_{\min}}{c_{\max}} - \frac{1}{2} \right)$. Then for any $s \in [k] \setminus \{i\}$, if y_s^* is negative, we have

$$y_s^* > -4\sqrt{2}\eta k. \quad (22)$$

Proof. Pick any $s \in [k] \setminus \{i\}$ such that $y_s^* < 0$. Consider the auxiliary LP

$$\begin{aligned} \min \quad & \mathbf{c}^T \mathbf{y} \\ \text{s.t.} \quad & \Theta \mathbf{y} \geq \mathbf{0} \\ & \mathbf{y}^T \boldsymbol{\theta}^i \geq 1 \\ & y_s \leq -4\sqrt{2}\eta k \end{aligned} \quad (\text{Pi-aux})$$

and its dual

$$\begin{aligned} \max \quad & \beta + (4\sqrt{2}\eta k)\gamma \\ \text{s.t.} \quad & \beta \boldsymbol{\theta}^i - \gamma \mathbf{e}_s + \Theta^T \mathbf{u} = \mathbf{c} \\ & \beta, \gamma, \mathbf{u} \geq 0. \end{aligned} \quad (\text{Di-aux})$$

If we show that \mathbf{y}^* is not an optimal solution to (Pi-aux), then we can conclude that $y_s^* > -4\sqrt{2}\eta k$. Therefore our goal is to show that the optimal value of (Pi-aux) is greater than $\mathbf{c}^T \mathbf{y}^*$. Equivalently, we may also show that the optimal value of (Di-aux) is greater than β^* . We do so by constructing a feasible solution for (Di-aux) at which the objective value is greater than β^* .

Let \mathbf{z}^* be the solution to the system

$$I'^T \mathbf{z} = \mathbf{r} + \frac{c_{\max}}{2} \mathbf{e}_s \quad (23)$$

where recall that $\mathbf{r} = \Theta^T \mathbf{u}^*/2$.

Note that the rows of I' belong to the unit simplex and for any $i \in [k]$, we have

$$\begin{aligned} \|I'(i, \cdot) - \mathbf{e}^i\|_\infty &\leq \eta \\ &\leq \frac{1}{2\sqrt{2}k}. \end{aligned} \quad (\text{by assumption on } \eta)$$

Therefore using Lemma A.1, we conclude that

$$\|I'^{-T} - I\|_\infty \leq 2\sqrt{2}\eta k. \quad (24)$$

Define the point

$$\begin{bmatrix} \bar{\beta} \\ \bar{\gamma} \\ \bar{\mathbf{u}} \end{bmatrix} := \begin{bmatrix} \beta^* \\ 0 \\ \mathbf{u}^*/2 \end{bmatrix} + \begin{bmatrix} \beta' \\ c_{\max}/2 \\ \mathbf{u}' \end{bmatrix} \quad (25)$$

where $\beta' := z_i^*$ and

$$u'_p := \begin{cases} z_p^* & \text{if } p \in [k] \setminus \{i\} \\ 0 & \text{otherwise.} \end{cases}$$

First we argue that $(\bar{\beta}, \bar{\gamma}, \bar{\mathbf{u}})$ is feasible for (Di-aux). From (25), we have

$$\begin{aligned} \bar{\beta} \boldsymbol{\theta}^i - \bar{\gamma} \mathbf{e}_s + \Theta^T \bar{\mathbf{u}} &= \beta^* \boldsymbol{\theta}^i + \Theta^T \mathbf{u}^*/2 + \beta' \boldsymbol{\theta}^i - c_{\max} \mathbf{e}_s/2 + \Theta^T \mathbf{u}' \\ &= \mathbf{c} - \mathbf{r} + \beta' \boldsymbol{\theta}^i - c_{\max} \mathbf{e}_s/2 + \Theta^T \mathbf{u}' && (\because (\beta^*, \mathbf{u}^*) \text{ is feasible for (Di)}) \\ &= \mathbf{c} - \mathbf{r} + I'^T \mathbf{z}^* - c_{\max} \mathbf{e}_s/2 && (\text{using the definition of } (\beta', \mathbf{u}')) \\ &= \mathbf{c}. && (\text{using (23)}) \end{aligned}$$

To argue about the nonnegativity of $(\bar{\beta}, \bar{\gamma}, \bar{\mathbf{u}})$, it suffices to argue that

1. $z_i^* + \beta^* \geq 0$
2. $\mathbf{z}^*([k] \setminus \{i\}) \geq \mathbf{0}$.

Note that our assumption on η implies $\eta < \frac{1}{2\sqrt{2}k} \frac{c_{\min}}{c_{\max}}$ and therefore Lemmas A.2 and A.3 apply.

We have

$$\begin{aligned}
z_i^* &= I'^{-T}(i, i)r_i + I'^{-T}(i, s)(r_s + c_{\max}/2) + \sum_{p \in [k] \setminus \{i, s\}} I'^{-T}(i, p)r_p \\
&\geq 0 + I'^{-T}(i, s)(r_s + c_{\max}/2) + \sum_{p \in [k] \setminus \{i, s\}} I'^{-T}(i, p)r_p && (\because I'^{-T}(i, i) \geq 0, r_i \geq 0) \\
&\geq -2\sqrt{2}\eta kc_{\max}. && (\text{using (24) and Lemma A.3})
\end{aligned} \tag{26}$$

Combining the lower bound on z_i^* with the lower bound on β^* in Lemma A.2 yields

$$\begin{aligned}
z_i^* + \beta^* &\geq c_i - 4\sqrt{2}\eta kc_{\max} \\
&\geq c_{\min} - 4\sqrt{2}\eta kc_{\max} \\
&> 0.
\end{aligned}$$

The last inequality above follows from our assumption on η . Indeed, we have

$$\begin{aligned}
\eta &< \frac{1}{4k} \left(\frac{c_{\min}}{c_{\max}} - \frac{1}{2} \right) \\
&< \frac{1}{4\sqrt{2}k} \frac{c_{\min}}{c_{\max}}. && \left(\because \frac{c_{\min}}{c_{\max}} \leq 1 \right)
\end{aligned}$$

Similarly, for any $t \in [k] \setminus \{i\}$ we have

$$\begin{aligned}
z_t^* &\geq r_t + c_{\max}I(s, t)/2 - \|I'^{-T} - I\|_{\infty} \|\mathbf{r} + c_{\max}\mathbf{e}_s/2\|_{\infty} && (\text{using (23)}) \\
&\geq r_t - \|I'^{-T} - I\|_{\infty} \|\mathbf{r} + c_{\max}\mathbf{e}_s/2\|_{\infty} \\
&\geq r_t - 2\sqrt{2}\eta kc_{\max} && (\text{using (24) and Lemma A.4}) \\
&\geq c_{\min} - \frac{\eta}{1-\eta} c_{\max} - 2\sqrt{2}\eta kc_{\max}. && (\text{using Lemma A.3})
\end{aligned} \tag{27}$$

Our assumption on η yields a positive lower bound on the above expression. Indeed, we have

$$\begin{aligned}
c_{\min} &> \frac{c_{\max}}{2} + 4k\eta c_{\max} && (\text{by assumption on } \eta) \\
&\geq \frac{c_{\max}}{2} + (2 + 3k)\eta c_{\max} && (\because k \geq 2) \\
&= \frac{c_{\max}}{2} + 2\eta c_{\max} + 3\eta kc_{\max} \\
&\geq \frac{c_{\max}}{2} + \frac{\eta}{1-\eta} c_{\max} + 2\sqrt{2}\eta kc_{\max} && (\because \eta \leq 1/2)
\end{aligned}$$

Using the above in (27), we get

$$z_t^* > c_{\max}/2. \tag{28}$$

Therefore $(\bar{\beta}, \bar{\gamma}, \bar{\mathbf{u}})$ is feasible for (Di-aux).

Now we argue that the objective value of (Di-aux) at $(\bar{\beta}, \bar{\gamma}, \bar{\mathbf{u}})$ is greater than β^* . Indeed note that

$$\begin{aligned}\beta' + (4\sqrt{2}\eta k) \frac{c_{\max}}{2} &= z_i^* + (4\sqrt{2}\eta k) \frac{c_{\max}}{2} \\ &> -2\sqrt{2}\eta k c_{\max} + (4\sqrt{2}\eta k) \frac{c_{\max}}{2} \quad (\text{using (26)}) \\ &= 0.\end{aligned}$$

That is, $\beta' + (4\sqrt{2}\eta k) \frac{c_{\max}}{2} > 0$ or equivalently, $\bar{\beta} + (4\sqrt{2}\eta k) \bar{\gamma} > \beta^*$ thereby concluding the proof. \square

Lemma A.7. Suppose $\frac{c_{\min}}{c_{\max}} > \frac{1}{2}$ and $\eta < \frac{1}{4k} \left(\frac{c_{\min}}{c_{\max}} - \frac{1}{2} \right)$. Then

$$\frac{1 - 4\sqrt{2}\eta^2 k}{\theta_{ii}} \leq y_i^* \leq \frac{1 + 4\sqrt{2}\eta^2 k}{\theta_{ii}}. \quad (29)$$

Proof. We note that the constraint $\mathbf{y}^T \boldsymbol{\theta}^i \geq 1$ in (Pi) is tight at optimality. Indeed otherwise one may scale the optimal solution so as to make that constraint tight and obtain a strictly smaller objective value, thereby contradicting optimality.

Then we have

$$\begin{aligned}1 &= \mathbf{y}^{*T} \boldsymbol{\theta}^i \\ &= y_i^* \theta_{ii} + \sum_{s \in [k] \setminus \{i\}} y_s^* \theta_{is}.\end{aligned} \quad (30)$$

Moreover

$$\begin{aligned}\left| \sum_{s \in [k] \setminus \{i\}} y_s^* \theta_{is} \right| &\leq \|\mathbf{y}^*([k] \setminus \{i\})\|_{\infty} \|\boldsymbol{\theta}^i([k] \setminus \{i\})\|_1 \quad (\text{using Hölder's inequality}) \\ &\leq \eta \|\mathbf{y}^*([k] \setminus \{i\})\|_{\infty} \quad (\because \|\boldsymbol{\theta}^i([k] \setminus \{i\})\|_1 \leq \eta) \\ &\leq 4\sqrt{2}\eta^2 k. \quad (\text{using Lemmas A.5 and A.6})\end{aligned} \quad (31)$$

Using (31) in (30) yields the desired result. \square

Proof of Theorem 3.4. First note that (P) is both feasible and bounded below, which implies that it has an optimal solution. Moreover, since B is full-rank, the column range of P is equal to the column range of Θ . Therefore (P) may be rewritten as

$$\begin{aligned}\min \quad & \mathbf{c}^T \mathbf{y} \\ \text{s.t.} \quad & \Theta \mathbf{y} \geq \mathbf{0} \\ & \mathbf{y}^T \boldsymbol{\theta}^i \geq 1.\end{aligned} \quad (\text{Py})$$

Since \mathbf{x}^* is an optimal solution to (P), there exists an optimal solution to (Py), called \mathbf{y}^* , satisfying $\Theta \mathbf{y}^* = \mathbf{x}^*$. Using Lemmas A.5, A.6, and A.7, we conclude that

$$\left\| \mathbf{y}^* - \frac{\mathbf{e}_j}{\theta_{ij}} \right\|_{\infty} \leq \sqrt{2}\eta k \max\{2, 4, 4\eta/\theta_{ij}\} = 4\sqrt{2}\eta k. \quad (32)$$

The last equality above holds because $\theta_{ij} \geq 1 - \eta$ and $\eta < 1/2$. Then we have

$$\begin{aligned}\left\| \mathbf{x}^* - \frac{\boldsymbol{\theta}_j}{\theta_{ij}} \right\|_{\infty} &= \left\| \Theta \mathbf{y}^* - \Theta \frac{\mathbf{e}_j}{\theta_{ij}} \right\|_{\infty} \\ &\leq \|\Theta\|_{\infty} \left\| \mathbf{y}^* - \frac{\mathbf{e}_j}{\theta_{ij}} \right\|_{\infty} \\ &\leq 4\sqrt{2}\eta k. \quad (\|\Theta\|_{\infty} = 1 \text{ and using (32)})\end{aligned} \quad (33)$$

Lastly, we have

$$\begin{aligned}
\left\| \frac{\mathbf{x}^*}{\|\mathbf{x}^*\|_\infty} - \boldsymbol{\theta}_j \right\|_\infty &\leq \left\| \frac{\mathbf{x}^*}{\|\mathbf{x}^*\|_\infty} - \mathbf{x}^* \right\|_\infty + \left\| \mathbf{x}^* - \frac{\boldsymbol{\theta}_j}{\theta_{ij}} \right\|_\infty + \left\| \frac{\boldsymbol{\theta}_j}{\theta_{ij}} - \boldsymbol{\theta}_j \right\|_\infty \\
&\hspace{15em} \text{(using triangle inequality)} \\
&= |1 - \|\mathbf{x}^*\|_\infty| + \left\| \mathbf{x}^* - \frac{\boldsymbol{\theta}_j}{\theta_{ij}} \right\|_\infty + \left\| \frac{\boldsymbol{\theta}_j}{\theta_{ij}} - \boldsymbol{\theta}_j \right\|_\infty \\
&\leq \left| 1 - \frac{\|\boldsymbol{\theta}_j\|_\infty}{\theta_{ij}} \right| + \left| \|\mathbf{x}^*\|_\infty - \frac{\|\boldsymbol{\theta}_j\|_\infty}{\theta_{ij}} \right| + \left\| \mathbf{x}^* - \frac{\boldsymbol{\theta}_j}{\theta_{ij}} \right\|_\infty + \left\| \frac{\boldsymbol{\theta}_j}{\theta_{ij}} - \boldsymbol{\theta}_j \right\|_\infty \\
&\hspace{15em} \text{(using triangle inequality)} \\
&\leq \left| 1 - \frac{\|\boldsymbol{\theta}_j\|_\infty}{\theta_{ij}} \right| + 2 \left\| \mathbf{x}^* - \frac{\boldsymbol{\theta}_j}{\theta_{ij}} \right\|_\infty + \left\| \frac{\boldsymbol{\theta}_j}{\theta_{ij}} - \boldsymbol{\theta}_j \right\|_\infty \\
&\hspace{15em} \text{(using reverse triangle inequality)} \\
&\leq \left(\frac{\|\boldsymbol{\theta}_j\|_\infty}{\theta_{ij}} - 1 \right) + 8\sqrt{2}\eta k + \left(\frac{1}{\theta_{ij}} - 1 \right) \|\boldsymbol{\theta}_j\|_\infty \\
&\leq 8\sqrt{2}\eta k + 2 \left(\frac{1}{\theta_{ij}} - 1 \right) \\
&\leq 8\sqrt{2}\eta k + \frac{2\eta}{1-\eta} \\
&< 8\sqrt{2}\eta k + 4\eta \\
&= 4\eta(2\sqrt{2}k + 1)
\end{aligned}$$

where the inequality in the fifth line from bottom follows from using (33), the inequality in the fourth line from bottom follows because $\|\boldsymbol{\theta}_j\|_\infty \leq 1$, the inequality in the third line from bottom follows because $\theta_{ij} \geq 1 - \eta$, and the inequality in the second line from bottom follows because $\eta < 1/2$. \square

B Some Concentration Properties in the MMSB

In this section, we show concentration properties of some key random variables associated with random matrices Θ and ΘB . We shall use these observations for our subsequent proofs, but they may also be of independent interest. Even though we work the equal parameter Dirichlet distribution, the proof techniques here easily extend to the case with different Dirichlet parameters.

Define $l := \sigma_{\min}(B)$ and $u := \sigma_{\max}(B)$. Suppose the k parameters of the Dirichlet distribution are all equal to α . We repeatedly use the facts that for any $i \in [n]$, $s \in [k]$,

$$\mathbb{E}[\theta_{is}] = \frac{1}{k} \tag{34}$$

and

$$\mathbb{E}[\theta_{is}^2] = \frac{\alpha + 1}{k(\alpha k + 1)}. \tag{35}$$

Moreover, if $s, t \in [k]$ such that $s \neq t$ then

$$\mathbb{E}[\theta_{is}\theta_{it}] = \frac{\alpha}{k(\alpha k + 1)}. \tag{36}$$

Lemma B.1. *For any $j \in [k]$, we have $\frac{9}{10} \frac{n}{k} \leq c_j \leq \frac{11}{10} \frac{n}{k}$ with probability at least $1 - 2 \exp\left(\frac{-n}{50k^2}\right)$.*

Proof. For any $j \in [k]$, c_j is the sum of n independent bounded random variables $\{\theta_{ij}\}_{i=1}^n$. Indeed each row of Θ is sampled independently and each entry of Θ lies in $[0, 1]$. Moreover, using (34) we

get that $\mathbb{E}[c_j] = n/k$. Thus, using Hoeffding's inequality, we have that for any $z > 0$

$$\Pr(|c_j - n/k| \geq z) \leq 2 \exp\left(\frac{-2z^2}{n}\right). \quad (37)$$

Setting $z = n/10k$ in (37) yields the desired result. \square

Corollary B.2. *We have $c_{\min}/c_{\max} \geq 9/11$ with probability at least $1 - p_1$, where $p_1 := 2k \exp\left(\frac{-n}{50k^2}\right)$.*

Proof. Lemma B.1 implies that with probability at least $1 - 2k \exp\left(\frac{-n}{50k^2}\right)$, both $c_{\min} \geq 9n/10k$ and $c_{\max} \leq 11n/10k$ hold. \square

Lemma B.3. *For any $\epsilon > 0$, $\|\Theta B\| \leq u\sqrt{\frac{2n}{k}} + \epsilon\|\Theta\|$ with probability at least $1 - \left(\frac{2u}{\epsilon} + 1\right)^k \exp\left(\frac{-2n}{k^2}\right)$.*

For proving Lemma B.3, we first prove the following statements for set $\mathcal{C} := \{\mathbf{y} \in \mathbb{R}^k : \exists \mathbf{x} \in \mathbb{R}^k \text{ such that } B\mathbf{x} = \mathbf{y}, \|\mathbf{x}\| = 1\}$ defined as the image of the unit sphere under B .

Lemma B.4. *If \mathcal{E} is an ϵ -net of \mathcal{C} of smallest possible cardinality, then $|\mathcal{E}| \leq \left(\frac{2u}{\epsilon} + 1\right)^k$.*

Proof. Let \mathcal{E}' be a maximal ϵ -separated subset of \mathcal{C} . Note that by definition of an ϵ -separated subset, for any distinct $\mathbf{x}, \mathbf{y} \in \mathcal{E}'$, we have $\|\mathbf{x} - \mathbf{y}\| > \epsilon$. Moreover, the maximality of \mathcal{E}' implies that \mathcal{E}' is also an ϵ -net of \mathcal{C} . Therefore

$$|\mathcal{E}| \leq |\mathcal{E}'|. \quad (38)$$

We also have that the union of $|\mathcal{E}'|$ disjoint balls $\bigcup_{\mathbf{x} \in \mathcal{E}'} \mathcal{B}(\mathbf{x}, \epsilon/2) \subseteq \mathcal{C} + \mathcal{B}(\mathbf{0}, \epsilon/2) \subseteq \mathcal{B}(\mathbf{0}, u + \epsilon/2)$.

Therefore

$$\text{vol}\left(\bigcup_{\mathbf{x} \in \mathcal{E}'} \mathcal{B}(\mathbf{x}, \epsilon/2)\right) \leq \text{vol}(\mathcal{B}(\mathbf{0}, u + \epsilon/2)) \quad (39)$$

which implies that $|\mathcal{E}'|(\epsilon/2)^k \leq (u + \epsilon/2)^k$ which yields the desired result when combined with (38). \square

Lemma B.5. *Suppose $\mathbf{y} \in \mathcal{C}$. For any $i \in [n]$:*

1. $0 \leq \langle \boldsymbol{\theta}^i, \mathbf{y} \rangle^2 \leq u^2$
2. $\frac{l^2}{k(\alpha k + 1)} \leq \mathbb{E}[\langle \boldsymbol{\theta}^i, \mathbf{y} \rangle^2] \leq \frac{u^2}{k}$

Proof. Let $\mathbf{y} = B\mathbf{x}$ such that $\|\mathbf{x}\| = 1$. Then $l \leq \|\mathbf{y}\| \leq u$.

1. We have

$$\begin{aligned} \langle \boldsymbol{\theta}^i, \mathbf{y} \rangle^2 &\leq \|\boldsymbol{\theta}^i\|^2 \|\mathbf{y}\|^2 && \text{(using Cauchy-Schwarz inequality)} \\ &\leq u^2 && (\|\boldsymbol{\theta}^i\| \leq 1). \end{aligned}$$

2. We have

$$\begin{aligned}
\mathbb{E}[\langle \boldsymbol{\theta}^i, \mathbf{y} \rangle^2] &= \mathbb{E}[\theta_{i1}^2 y_1^2 + \cdots + \theta_{ik}^2 y_k^2] + \mathbb{E} \left[\sum_{\substack{s,t \in [k]: \\ s \neq t}} \theta_{is} \theta_{it} y_s y_t \right] \\
&= \frac{\alpha + 1}{k(\alpha k + 1)} \|\mathbf{y}\|^2 + \mathbb{E} \left[\sum_{\substack{s,t \in [k]: \\ s \neq t}} \theta_{is} \theta_{it} y_s y_t \right] && \text{(using (35))} \\
&= \frac{\alpha + 1}{k(\alpha k + 1)} \|\mathbf{y}\|^2 + \frac{\alpha}{k(\alpha k + 1)} \sum_{\substack{s,t \in [k]: \\ s \neq t}} y_s y_t && \text{(using (36))} \\
&= \frac{1}{k(\alpha k + 1)} \|\mathbf{y}\|^2 + \frac{\alpha}{k(\alpha k + 1)} (\mathbf{e}^T \mathbf{y})^2 && \text{(re-arranging terms).}
\end{aligned}$$

Now noting the second term on the right hand side above is nonnegative yields the desired lower bound.

Similarly noting that $\mathbf{e}^T \mathbf{y} \leq u\sqrt{k}$ (using Cauchy-Schwarz inequality) yields the desired upper bound.

□

Proof of Lemma B.3. We have

$$\|\Theta B\| = \sup_{\mathbf{x} \in S^{k-1}} \|\Theta B \mathbf{x}\| = \sup_{\mathbf{y} \in \mathcal{C}} \|\Theta \mathbf{y}\|. \quad (40)$$

Let \mathcal{E} denote an ϵ -net of \mathcal{C} of smallest possible cardinality. Then we have

$$\|\Theta B\| \leq \sup_{\mathbf{y} \in \mathcal{E}} \|\Theta \mathbf{y}\| + \epsilon \|\Theta\|. \quad (41)$$

Indeed if the supremum defining $\|\Theta B\|$ on the RHS in (40) is attained at \mathbf{y}_s , and if \mathbf{y}_e is a point in \mathcal{E} such that $\|\mathbf{y}_s - \mathbf{y}_e\| \leq \epsilon$, then

$$\begin{aligned}
\|\Theta B\| &= \|\Theta \mathbf{y}_s\| \\
&= \|\Theta \mathbf{y}_e + \Theta(\mathbf{y}_s - \mathbf{y}_e)\| \\
&\leq \|\Theta \mathbf{y}_e\| + \|\Theta(\mathbf{y}_s - \mathbf{y}_e)\| && \text{(using triangle inequality)} \\
&\leq \sup_{\mathbf{y} \in \mathcal{E}} \|\Theta \mathbf{y}\| + \epsilon \|\Theta\|.
\end{aligned}$$

For any $\mathbf{y} \in \mathcal{E}$, we have

$$\|\Theta \mathbf{y}\|^2 = \langle \boldsymbol{\theta}^1, \mathbf{y} \rangle^2 + \cdots + \langle \boldsymbol{\theta}^n, \mathbf{y} \rangle^2.$$

Now note that $\|\Theta \mathbf{y}\|^2$ is the sum of n independent random variables. Indeed using Lemma B.5 we conclude that each of these random variables is bounded and that $\mathbb{E}[\|\Theta \mathbf{y}\|^2] \leq \frac{nu^2}{k}$. Thus, using Hoeffding's inequality, we have that for any $z > 0$,

$$\begin{aligned}
\Pr \left(\|\Theta \mathbf{y}\|^2 \geq \frac{nu^2}{k} + z \right) &\leq \Pr(\|\Theta \mathbf{y}\|^2 \geq \mathbb{E}[\|\Theta \mathbf{y}\|^2] + z) \\
&\leq \exp \left(\frac{-2z^2}{nu^4} \right).
\end{aligned}$$

Then using the union bound over the ϵ -net, we obtain that

$$\begin{aligned}
\Pr \left(\sup_{\mathbf{y} \in \mathcal{E}} \|\Theta \mathbf{y}\| \geq \sqrt{\frac{nu^2}{k} + z} \right) &\leq |\mathcal{E}| \exp \left(\frac{-2z^2}{nu^4} \right) \\
&\leq \left(\frac{2u}{\epsilon} + 1 \right)^k \exp \left(\frac{-2z^2}{nu^4} \right) && \text{(using Lemma B.4)}
\end{aligned}$$

Setting $z = nu^2/k$ in the above, we note that $\sup_{y \in \mathcal{E}} \|\Theta y\| \leq u\sqrt{\frac{2n}{k}}$ with probability at least $1 - \left(\frac{2u}{\epsilon} + 1\right)^k \exp\left(\frac{-2n}{k^2}\right)$, combining which with (41) yields the desired result. \square

Corollary B.6. $\|\Theta\| \leq 2\sqrt{\frac{2n}{k}}$ with probability at least $1 - p_2$, where $p_2 := 5^k \exp\left(\frac{-2n}{k^2}\right)$.

Proof. Set $B = I$ and $\epsilon = 1/2$ in Lemma B.3. \square

Corollary B.7. $\|\Theta B\| \leq 2u\sqrt{\frac{2n}{k}}$ with probability at least $1 - p_2$.

Proof. This follows simply from using the inequality $\|\Theta B\| \leq \|\Theta\| \|B\|$ and the upper bound obtained in Corollary B.6. \square

Lemma B.8. $\sigma_k(\Theta B) \geq \frac{1}{4} \frac{l}{\sqrt{\alpha k + 1}} \sqrt{\frac{2n}{k}}$ with probability at least $1 - p_3$, where $p_3 := p_2 + \left(\frac{16u\sqrt{\alpha k + 1}}{l} + 1\right)^k \exp\left(\frac{-nl^4}{2k^2u^4(\alpha k + 1)^2}\right)$.

Proof. We have

$$\sigma_k(\Theta B) = \inf_{\mathbf{x} \in S^{k-1}} \|\Theta B \mathbf{x}\| = \inf_{\mathbf{y} \in \mathcal{C}} \|\Theta \mathbf{y}\|. \quad (42)$$

Let \mathcal{E} denote an ϵ -net of \mathcal{C} of smallest possible cardinality. Then we have

$$\sigma_k(\Theta B) \geq \inf_{\mathbf{y} \in \mathcal{E}} \|\Theta \mathbf{y}\| - \epsilon \|\Theta\|. \quad (43)$$

Indeed if the infimum defining $\sigma_k(\Theta B)$ on the RHS in (42) is attained at \mathbf{y}_s , and if \mathbf{y}_e is a point in \mathcal{E} such that $\|\mathbf{y}_s - \mathbf{y}_e\| \leq \epsilon$, then

$$\begin{aligned} \sigma_k(\Theta B) &= \|\Theta \mathbf{y}_s\| \\ &= \|\Theta \mathbf{y}_e + \Theta(\mathbf{y}_s - \mathbf{y}_e)\| \\ &\geq \|\Theta \mathbf{y}_e\| - \|\Theta(\mathbf{y}_s - \mathbf{y}_e)\| \quad (\text{using reverse triangle inequality}) \\ &\geq \|\Theta \mathbf{y}_e\| - \|\Theta(\mathbf{y}_s - \mathbf{y}_e)\| \\ &\geq \inf_{\mathbf{y} \in \mathcal{E}} \|\Theta \mathbf{y}\| - \epsilon \|\Theta\|. \end{aligned}$$

For any $\mathbf{y} \in \mathcal{E}$, we have

$$\|\Theta \mathbf{y}\|^2 = \langle \boldsymbol{\theta}^1, \mathbf{y} \rangle^2 + \dots + \langle \boldsymbol{\theta}^n, \mathbf{y} \rangle^2.$$

Now note that $\|\Theta \mathbf{y}\|^2$ is the sum of n independent bounded random variables. Indeed using Lemma B.5 we conclude that each of these random variables is bounded and that $\mathbb{E}[\|\Theta \mathbf{y}\|^2] \geq \frac{nl^2}{k(\alpha k + 1)}$.

Thus, using Hoeffding's inequality, we have that for any $z > 0$,

$$\begin{aligned} \Pr\left(\|\Theta \mathbf{y}\|^2 \leq \frac{nl^2}{k(\alpha k + 1)} - z\right) &\leq \Pr(\|\Theta \mathbf{y}\|^2 \leq \mathbb{E}[\|\Theta \mathbf{y}\|^2] - z) \\ &\leq \exp\left(\frac{-2z^2}{nu^4}\right). \end{aligned}$$

Then using the union bound over the ϵ -net, we obtain that

$$\begin{aligned} \Pr\left(\inf_{\mathbf{y} \in \mathcal{E}} \|\Theta \mathbf{y}\| \leq \sqrt{\frac{nl^2}{k(\alpha k + 1)} - z}\right) &\leq |\mathcal{E}| \exp\left(\frac{-2z^2}{nu^4}\right) \\ &\leq \left(\frac{2u}{\epsilon} + 1\right)^k \exp\left(\frac{-2z^2}{nu^4}\right) \quad (\text{using Lemma B.4}) \end{aligned}$$

Setting $z = \frac{1}{2} \frac{nl^2}{k(\alpha k + 1)}$, we note that $\inf_{y \in \mathcal{E}} \|\Theta y\| \geq \sqrt{\frac{1}{2} \frac{nl^2}{k(\alpha k + 1)}}$ with probability at least $1 - \left(\frac{2u}{\epsilon} + 1\right)^k \exp\left(\frac{-nl^4}{2k^2u^4(\alpha k + 1)^2}\right)$.

Using (43), we get that

$$\sigma_k(\Theta B) \geq \sqrt{\frac{1}{2} \frac{nl^2}{k(\alpha k + 1)}} - \epsilon \|\Theta\| \quad (44)$$

with probability at least $1 - \left(\frac{2u}{\epsilon} + 1\right)^k \exp\left(\frac{-nl^4}{2k^2u^4(\alpha k + 1)^2}\right)$.

Lastly, using the upper bound on $\|\Theta\|$ derived in Corollary B.6 in (44), we get that

$$\sigma_k(\Theta B) \geq \frac{1}{2} \frac{l}{\sqrt{\alpha k + 1}} \sqrt{\frac{2n}{k}} - 2\epsilon \sqrt{\frac{2n}{k}}$$

with probability at least $1 - p_2 - \left(\frac{2u}{\epsilon} + 1\right)^k \exp\left(\frac{-nl^4}{2k^2u^4(\alpha k + 1)^2}\right)$. Setting $\epsilon = \frac{1}{8} \frac{l}{\sqrt{\alpha k + 1}}$ yields the desired result. \square

C Proof of Main Theorem

In this section, we build the proof of Theorem 3.1.

Lemma C.1. *Let $p, \gamma \in (0, 1)$. If $n > \frac{\log(p/k)}{\log I_{1-\gamma}(\alpha, (k-1)\alpha)}$, then with probability at least $1 - p$, for each $j \in [k]$, there exists a row vector \mathbf{r}^T in Θ such that*

$$\|\mathbf{r} - \mathbf{e}_j\|_\infty < \gamma. \quad (45)$$

(Here $I_x(y, z)$ denotes the regularized incomplete beta function.)

Proof. For any $j \in [k]$, define E_j as the event that there exists a row \mathbf{r}^T in Θ such that $\|\mathbf{r} - \mathbf{e}_j\|_\infty < \gamma$. Then for any $j \in [k]$, we have

$$\begin{aligned} \Pr(E_j^c) &= \prod_{i \in [n]} \Pr(\|\theta^i - \mathbf{e}_j\|_\infty \geq \gamma) \quad (\because \text{rows of } \Theta \text{ are independently sampled}) \\ &= \prod_{i \in [n]} \Pr(\theta_{ij} \leq 1 - \gamma) \quad (\because \text{rows of } \Theta \text{ belong to unit simplex}) \\ &= [I_{1-\gamma}(\alpha, (k-1)\alpha)]^n \quad (I_x(y, z) \text{ is the CDF of marginal of Dirichlet distribution}) \\ &< p/k. \quad (\text{by assumption on } n) \end{aligned} \quad (46)$$

Therefore

$$\begin{aligned} \Pr(E_1 \cap \dots \cap E_k) &= 1 - \Pr(E_1^c \cup \dots \cup E_k^c) \\ &\geq 1 - \sum_{j \in [k]} \Pr(E_j^c) \quad (\text{using the union bound}) \\ &> 1 - p. \quad (\text{using (46)}) \end{aligned}$$

\square

Proof of Theorem 3.1. Using the lower bound assumption on n and Lemma C.1, we conclude that with probability at least $1 - p$, for each $j \in [k]$, there exists a row \mathbf{r}^T in Θ such that

$$\|\mathbf{r} - \mathbf{e}_j\|_\infty < \epsilon. \quad (47)$$

Recalling the definition of Δ , we note that (47) is equivalent to

$$\|\Delta\|_{\max} < \epsilon. \quad (48)$$

Using Corollary B.7 and Lemma B.8, we conclude that

$$\kappa_0 \leq 8\kappa\sqrt{\alpha k + 1} \quad (49)$$

with probability at least $1 - p_2 - p_3$. Therefore (49) implies that

$$\begin{aligned} \min\left(\frac{1}{\sqrt{k-1}}, \frac{1}{2}\right) \frac{1}{2\sqrt{2}\kappa_0(1+80\kappa_0^2)} &\geq \epsilon_1 && \text{(using the definition of } \epsilon_1\text{)} \\ &> \epsilon && \text{(using the assumption on } \epsilon\text{)} \\ &> \|\Delta\|_{\max} && \text{(using (48))} \end{aligned} \quad (50)$$

with probability at least $1 - p_2 - p_3$.

Using (50), we note that the assumption of Theorem 3.3 is satisfied with probability at least $1 - p - p_2 - p_3$. Therefore the set \mathcal{J} returned by Algorithm 2 satisfies

$$\begin{aligned} \|\Pi\Theta(\mathcal{J}, \cdot) - I\|_{\max} &\leq 40\sqrt{2}\kappa_0^2\|\Delta\|_{\max} \\ &< 40\sqrt{2}\kappa_0^2\epsilon \end{aligned} \quad (51)$$

with probability at least $1 - p - p_2 - p_3$ for some $k \times k$ permutation matrix Π .

Now from Corollary B.2, we know that

$$\frac{1}{4k} \left(\frac{c_{\min}}{c_{\max}} - \frac{1}{2} \right) \geq \frac{7}{88k} \quad (52)$$

with probability at least $1 - p_1$.

Thus we have

$$\begin{aligned} 40\sqrt{2}\kappa_0^2\epsilon &< 40\sqrt{2}\kappa_0^2\epsilon_2 && \text{(using the assumption on } \epsilon\text{)} \\ &\leq 40\sqrt{2} \cdot 64\kappa^2(\alpha k + 1)\epsilon_2 && \text{(using (49))} \\ &= \frac{7}{88k} && \text{(using the definition of } \epsilon_2\text{)} \\ &\leq \frac{1}{4k} \left(\frac{c_{\min}}{c_{\max}} - \frac{1}{2} \right) && \text{(using (52))} \end{aligned} \quad (53)$$

with probability at least $1 - p - p_1 - p_2 - p_3$. Combining (51) and (53), we conclude that the assumption of Theorem 3.4 is satisfied with probability at least $1 - p - p_1 - p_2 - p_3$. Therefore for any $j \in [k]$, the vector $\hat{\theta}_j$ returned by SP+LP satisfies

$$\begin{aligned} \|\hat{\theta}_j - \theta_j\|_{\infty} &\leq 4 \cdot 40\sqrt{2}\kappa_0^2\epsilon \cdot (2\sqrt{2}k + 1) \\ &\leq 10240\sqrt{2}\kappa^2(\alpha k + 1)(2\sqrt{2}k + 1)\epsilon && \text{(using (49))} \\ &= \mathcal{O}(\alpha k^2 \kappa^2 \epsilon) \end{aligned}$$

with probability at least $1 - p - p_1 - p_2 - p_3$. Substituting the expressions for p_1, p_2 and p_3 , the probability $1 - p - p_1 - p_2 - p_3$ can be expressed as $1 - p - c_1 e^{-c_2 n}$ such that c_1, c_2 are constants that depend on α, k, κ . \square

Proof of Corollary 3.2. From Theorem 3.1, we know that the maximum distance between vectors $\hat{\theta}_1, \dots, \hat{\theta}_k$ and the columns of Θ , up to a permutation, is $\mathcal{O}(\alpha k^2 \kappa^2 \epsilon)$ with probability at least $1 - p - c_1 e^{-c_2 n}$ where c_1, c_2 are constants that depend on α, k, κ .

Similarly, the maximum distance between vectors $\hat{\theta}_1, \dots, \hat{\theta}_k$ and the columns of $\bar{\Theta}$, up to a permutation, is $\mathcal{O}(\bar{\alpha} k^2 \bar{\kappa}^2 \bar{\epsilon})$ with probability at least $1 - \bar{p} - \bar{c}_1 e^{-\bar{c}_2 n}$ where \bar{c}_1, \bar{c}_2 are constants that depend on $\bar{\alpha}, k, \bar{\kappa}$.

Combining the above two observations with the triangle inequality and the union bound yields the desired result. \square