
On Mixed Memberships and Symmetric Nonnegative Matrix Factorizations

Xueyu Mao¹ Purnamrita Sarkar² Deepayan Chakrabarti³

Abstract

The problem of finding overlapping communities in networks has gained much attention recently. Optimization-based approaches use non-negative matrix factorization (NMF) or variants, but the global optimum cannot be provably attained in general. Model-based approaches, such as the popular mixed membership stochastic blockmodel or MMSB (Airoldi et al., 2008), use parameters for each node to specify the overlapping communities, but standard inference techniques cannot guarantee consistency. We link the two approaches, by (a) establishing sufficient conditions for the symmetric NMF optimization to have a unique solution under MMSB, and (b) proposing a computationally efficient algorithm called GeoNMF that is provably optimal and hence consistent for a broad parameter regime. We demonstrate its accuracy on both simulated and real-world datasets.

1. Introduction

Community detection is a fundamental problem in network analysis. It has been widely used in a diverse set of applications ranging from link prediction in social networks (Soundarajan & Hopcroft, 2012), predicting protein-protein or protein-DNA interactions in biological networks (Chen & Yuan, 2006), to network protocol design such as data forwarding in delay tolerant networks (Lu et al., 2015).

Traditional community detection assumes that every node in the network belongs to exactly one community, but many practical settings call for greater flexibility. For

¹Department of Computer Science. ²Department of Statistics and Data Sciences. ³Department of Information, Risk, and Operations Management. The University of Texas at Austin, TX, USA. Correspondence to: Xueyu Mao <xmao@cs.utexas.edu>, Purnamrita Sarkar <purna.sarkar@austin.utexas.edu>, Deepayan Chakrabarti <deepay@utexas.edu>.

instance, individuals in a social network may have multiple interests, and hence are best described as members of multiple interest-based communities. We focus on the popular mixed membership stochastic blockmodel (MMSB) (Airoldi et al., 2008) where each node i , $i \in [n]$ has a discrete probability distribution $\theta_i = (\theta_{i1}, \dots, \theta_{iK})$ over K communities. The probability of linkage between nodes i and j depends on the degree of overlap between their communities:

$$\begin{aligned}\theta_i &\sim \text{Dirichlet}(\alpha) & i \in [n] \\ \mathbf{P} &= \rho \Theta \mathbf{B} \Theta^T \\ \mathbf{A}_{ij} = \mathbf{A}_{ji} &= \text{Bernoulli}(\mathbf{P}_{ij}) & i, j \in [n]\end{aligned}$$

where θ_i is the i -th row of Θ , \mathbf{A} represents the adjacency matrix of the generated graph, and $\mathbf{B} \in \mathbb{R}^{K \times K}$ is the community-community interaction matrix. The parameter ρ controls the sparsity of the graph, so WLOG, the largest entry of \mathbf{B} can be set to 1. The parameter $\alpha_0 = \sum_i \alpha_i$ controls the amount of overlap. In particular, when $\alpha_0 \rightarrow 0$, MMSB reduces to the well known stochastic blockmodel, where every node belongs to exactly one community. Larger α_0 leads to more overlap. Since we only observe \mathbf{A} , a natural question is: how can $\{\theta_i\}$ and \mathbf{B} be recovered from \mathbf{A} in a way that is provably consistent?

1.1. Prior work

We categorize existing approaches broadly into three groups: model-based parameter inference methods, specialized algorithms that offer provable guarantees, and optimization-based methods using non-negative matrix factorization.

Model-based methods: These apply standard techniques for inference of hidden variables to the MMSB model. Examples include MCMC techniques (Chang, 2012) and variational methods (Gopalan & Blei, 2013). While these often work well in practice, there are no proofs of consistency for these methods. The MCMC methods are difficult to scale to large graphs, so we compare against the faster variational inference methods in our experiments.

Algorithms with provable guarantees: There has been work on provably consistent estimation on models similar to MMSB. Zhang et al. (2014) propose a spectral method

(OCCAM) for a model where the θ_i has unit ℓ_2 norm (unlike MMSB, where they have unit ℓ_1 norm). In addition to the standard assumptions regarding the existence of “pure” nodes¹ (which only belong to a single community) and a positive-definite \mathbf{B} , they also require \mathbf{B} to have equal diagonal entries, and assume that the ground truth communities has a unique optimum of a special loss function, and there is curvature around the optimum. Such assumptions may be hard to verify. Ray et al. (2015) and Kaufmann et al. (2016) consider models with binary community memberships. Kaufmann et al. (2016) show that the global optimum of a special loss function is consistent. However, achieving the global optimum is computationally intractable, and the scalable algorithm proposed by them (SAAC) is not provably consistent. Anandkumar et al. (2014) propose a tensor based approach for MMSB. Despite their elegant solution the computational complexity is $O(n^2K)$, which can be prohibitive for large graphs.

Optimization-based methods: If \mathbf{B} is positive-definite, the MMSB probability matrix \mathbf{P} can be written as $\mathbf{P} = \mathbf{W}\mathbf{W}^T$, where the \mathbf{W} matrix has only non-negative entries. In other words, \mathbf{W} is the solution to a Symmetric Non-negative Matrix Factorization (SNMF) problem: $\mathbf{W} = \arg \min_{\mathbf{X} \geq 0} \text{loss}(\mathbf{P}, \mathbf{X}\mathbf{X}^T)$ for some loss function that measure the “difference” between \mathbf{P} and its factorization. SNMF has been widely studied and successfully used for community detection (Kuang et al., 2015; Wang et al., 2011; 2016; Psorakis et al., 2011), but typically lacks the guarantees we desire. Our paper attempts to address these issues.

We note that Arora et al. (2012; 2013) used NMF to consistently estimate parameters of a topic model. However, their results cannot be easily applied to the MMSB inference problem. In particular, for topic models, the columns of the word-by-topic matrix specifying the probability distribution of words in a topic sum to 1. For MMSB, the rows of the node membership matrix sum to 1. The relationship of this work to the MMSB problem is unclear.

1.2. Problem Statement and Contributions

We seek to answer two problems.

Problem 1: *Given \mathbf{P} , when does the solution to the SNMF optimization yield the correct \mathbf{W} ?*

The difficulty stems from the fact that (a) the MMSB model may not always be identifiable, and (b) even if it is, the corresponding SNMF problem may not have a unique solution

¹This is a common assumption even for NMF methods for topic modeling, where each topic is assumed to have an anchor word (words belonging to only one topic). Huang et al. (2016) introduced a special optimization criterion to relax the presence of anchor words, but the optimization criterion is non-convex.

(even after allowing for permutation of communities).

Even when the conditions for Problem 1 are met, we may be unable to find a good solution in practice. This is due to two reasons. First, we only know the adjacency matrix \mathbf{A} , and not the probability matrix \mathbf{P} . Second, the general SNMF problem is non-convex, and SNMF algorithms can get stuck at local optima. Hence, it is unclear if an *algorithm* can consistently recover the MMSB parameters. This leads to our next question.

Problem 2: *Given \mathbf{A} generated from a MMSB model, can we develop a fast and provably consistent algorithm to infer the parameters?*

Our goal is to develop a fast algorithm that provably solves SNMF for an identifiable MMSB model. Note that generic SNMF algorithms typically do not have any provable guarantees.

Our contributions are as follows.

Identifiability: We show conditions that are sufficient for MMSB to be identifiable; specifically, there must be at least one “pure” exemplar of each of the K clusters (i.e., a node that belongs to that community with probability 1), and \mathbf{B} must be full rank.

Uniqueness under SNMF: We provide sufficient conditions under which an identifiable MMSB model is the unique solution for the SNMF problem; specifically, the MMSB probability matrix \mathbf{P} has a unique SNMF solution if \mathbf{B} is diagonal. It is important to note that MMSB with a diagonal \mathbf{B} still allows for interactions between different communities via members who belong to both.

Recovery algorithm: We present a new algorithm, called GeoNMF, for recovering the parameters $\{\theta_i\}$ and \mathbf{B} given only the observed adjacency matrix \mathbf{A} . The only compute-intensive part of the algorithm is the calculation of the top- K eigenvalues and eigenvectors of \mathbf{A} , for which highly optimized algorithms exist (Press et al., 1992).

Provable guarantees: Under the common assumption that θ_i are generated from a Dirichlet(α) prior, we prove the consistency of GeoNMF when \mathbf{B} is diagonal and there are “pure” nodes for each cluster (exactly the conditions needed for uniqueness of SNMF). We allow the sparsity parameter ρ to decay with the graph size n . All proofs are deferred to the appendix.

Empirical validation: On simulated networks, we compare GeoNMF against variational methods (SVI) (Gopalan & Blei, 2013). Since OCCAM, SAAC, and BSNMF (a Bayesian variant of SNMF (Psorakis et al., 2011)) are formed under different model assumptions, we exclude these for the simulation experiments for fairness. We also run experiments on Facebook and Google Plus ego

networks collected by Mcauley & Leskovec (2014); co-authorship datasets constructed by us from DBLP (Ley, 2002) and the Microsoft academic graph (MAG) (Sinha et al., 2015). These networks can have up to 150,000 nodes. On these graphs we compare GeoNMF against SVI, SAAC, OCCAM and BSNMF. We see that GeoNMF is consistently among the top, while also being one of the fastest. This establishes that GeoNMF achieves excellent accuracy and is computationally efficient in addition to being provably consistent.

2. Identifiability and Uniqueness

In order to present our results, we will now introduce some key definitions. Similar definitions appear in (Zhang et al., 2014).

Definition 2.1. A node $i \in [n]$ is called a “pure” node if $\exists j \in [K]$ such that $\theta_{ij} = 1$ and $\theta_{i\ell} = 0$ for all $\ell \in [K]$, $\ell \neq j$.

Identifiability of MMSB. MMSB is not identifiable in general. Consider the following counter example.

$$\mathbf{M}_1 = \begin{bmatrix} 0.5 & 0.5 & 0 \\ 0 & 0.5 & 0.5 \\ 0.5 & 0 & 0.5 \end{bmatrix} \quad \mathbf{M}_2 = \begin{bmatrix} 0.5 & 0.25 & 0.25 \\ 0.25 & 0.5 & 0.25 \\ 0.25 & 0.25 & 0.5 \end{bmatrix}$$

It can be easily checked that the probability matrices \mathbf{P} generated by the parameter set $(\Theta^{(1)}, \mathbf{B}^{(1)}, \rho^{(1)}) = (\mathbf{M}_1, \mathbf{I}_{3 \times 3}, 1)$ is exactly the same as that generated by $(\Theta^{(2)}, \mathbf{B}^{(2)}, \rho^{(2)}) = (\mathbf{I}_{3 \times 3}, 2\mathbf{M}_2, 0.5)$, where $\mathbf{I}_{3 \times 3}$ is the identity matrix. This example can be extended to arbitrarily large n : for every new row $\theta_i^{(2)}$ added to $\Theta^{(2)}$, add the row $\theta_i^{(1)} = \theta_i^{(2)}\mathbf{M}_1$ to $\Theta^{(1)}$. The new rows are still non-negative and sum to 1; it can be verified that $\mathbf{P}^{(1)} = \mathbf{P}^{(2)}$ even after these new node additions.

Thus, while MMSB is not identifiable in general, we can prove identifiability under certain conditions.

Theorem 2.1 (Sufficient conditions for MMSB identifiability). Suppose parameters Θ, \mathbf{B} of the MMSB model satisfy the following conditions: (a) there is at least one pure node for each community, and (b) \mathbf{B} has full rank. Then, MMSB is identifiable up to a permutation.

Since identifiability is a necessary condition for consistent recovery of parameters, we will assume these conditions from now on.

Uniqueness of SNMF for MMSB model. Even when the MMSB model is identifiable, the SNMF optimization may not have a unique solution. In other words, given an MMSB probability matrix \mathbf{P} , there might be multiple matrices \mathbf{X} such that $\mathbf{P} = \mathbf{X}\mathbf{X}^T$, even if \mathbf{P} corresponds to a unique parameter setting $(\Theta, \mathbf{B}, \rho)$ under MMSB. For SNMF to

work, $\mathbf{W} = \sqrt{\rho}\Theta\mathbf{B}^{1/2}$ must be the unique SNMF solution. When does this happen?

In general, SNMF is not unique because \mathbf{W} can be permuted, so we consider the following definition of uniqueness.

Definition 2.2. (Uniqueness of SNMF (Huang et al., 2014)) The Symmetric NMF of $\mathbf{P} = \mathbf{W}\mathbf{W}^T$ is said to be (essentially) unique if $\mathbf{P} = \tilde{\mathbf{W}}\tilde{\mathbf{W}}^T$ implies $\tilde{\mathbf{W}} = \mathbf{W}\mathbf{Z}$, where \mathbf{Z} is a permutation matrix.

Theorem 2.2 (Uniqueness of SNMF for MMSB). Consider an identifiable MMSB model where \mathbf{B} is diagonal. Then, its Symmetric NMF \mathbf{W} is unique and equals $\sqrt{\rho}\Theta\mathbf{B}^{1/2}$.

The above results establish that if we find a \mathbf{W} that is the symmetric NMF solution of \mathbf{P} then it is at least unique. However, two practical questions are still unanswered. First, given the non-convex nature of SNMF, how can we guarantee that we find the correct \mathbf{W} given \mathbf{P} ? Second, in practice we are given not \mathbf{P} but the noisy adjacency matrix \mathbf{A} . Typical algorithms for SNMF do not provide guarantees even for the first question.

3. Provably consistent inference for MMSB

To achieve consistent inference, we turn to the specific structure of the MMSB model. We motivate our approach in three stages. First, note that under the conditions of Theorem 2.2, the rows of \mathbf{W} form a *simplex* whose corners are formed by the pure nodes for each cluster. In addition, these corners are aligned along different axes, and hence are orthogonal to each other. Thus, if we can detect the corners of the simplex, we can recover the MMSB parameters. So the goal is to find the pure nodes from different clusters, since they define the corners.

While our goal is to get \mathbf{W} , note that it is easy to compute $\mathbf{V}\mathbf{E}^{1/2}$ where \mathbf{V}, \mathbf{E} are the eigenvectors and eigenvalues of \mathbf{P} , i.e., $\mathbf{P} = \mathbf{V}\mathbf{E}\mathbf{V}^T$. Thus, $\mathbf{W}\mathbf{W}^T = (\mathbf{V}\mathbf{E}^{1/2})(\mathbf{V}\mathbf{E}^{1/2})^T$. This implies that $\mathbf{W} = \mathbf{V}\mathbf{E}^{1/2}\mathbf{Q}$ for some orthogonal matrix \mathbf{Q} (Lemma A.1 of (Tang et al., 2013)). Essentially we should be able to identify the pure nodes by finding the corners of the simplex based on \mathbf{V} and \mathbf{E} .

Once we have found the pure nodes, it is easy to find the rotation matrix \mathbf{Q} modulo a permutation of classes, because we know that the pure nodes are on the axis for the simplex of $\Theta\mathbf{B}^{1/2}$.

Now, we note something rather striking. Let \mathcal{D} denote the diagonal matrix with expected degrees on the diagonal. Consider the population Laplacian $\mathcal{D}^{-1/2}\mathbf{P}\mathcal{D}^{-1/2}$. Its square root is given by $\mathcal{D}^{-1/2}\mathbf{V}\mathbf{E}^{1/2}$, which has the following interesting property for equal Dirichlet param-

ters $\alpha_a = \alpha_0/K$. We show in Lemma 4.1 that while the resulting rows no longer fall on a simplex, the rows with the largest norm are precisely the pure nodes, for whom the norm concentrates around $\sqrt{K/n}$. Thus, picking the rows with the largest norm of the square root gives us the pure nodes. From this, \mathbf{Q} , θ_i for other rows and the parameters ρ and \mathbf{B} can again be easily extracted.

Needless to say, this only answers the question for the expectation matrix \mathbf{P} . In reality, we have a noisy adjacency matrix. Let $\hat{\mathbf{V}}$ and $\hat{\mathbf{E}}$ denote the matrices of eigenvectors and eigenvalues of \mathbf{A} . We also establish in this paper that the rows of $\hat{\mathbf{V}}\hat{\mathbf{E}}^{1/2}$ concentrate around its population counterpart (corresponding row of $\mathbf{V}\mathbf{E}^{1/2}\mathbf{O}$ for some rotation matrix \mathbf{O}). While there are eigenvector deviation results in random matrix theory, e.g. the Davis-Kahan Theorem (Davis & Kahan, 1970), these typically provide deviation results for the whole $\hat{\mathbf{V}}$ matrix, not its rows. In a nutshell, this crucial result lets us carefully bound the errors of each step of the same basic idea executed on \mathbf{A} , the noisy proxy for \mathbf{P} .

Algorithm 1 GeoNMF

Input: Adjacency matrix \mathbf{A} ; number of communities K ; a constant ϵ_0

Output: Estimated node-community distribution matrix $\hat{\Theta}$, Community-community interaction matrix $\hat{\mathbf{B}}$, sparsity-control parameter $\hat{\rho}$;

- 1: Randomly split the set of nodes $[n]$ into two equal-sized parts \mathcal{S} and $\bar{\mathcal{S}}$.
 - 2: Obtain the top K eigen-decomposition of $\mathbf{A}(\mathcal{S}, \mathcal{S})$ as $\hat{\mathbf{V}}_1\hat{\mathbf{E}}_1\hat{\mathbf{V}}_1^T$ and of $\mathbf{A}(\bar{\mathcal{S}}, \bar{\mathcal{S}})$ as $\hat{\mathbf{V}}_2\hat{\mathbf{E}}_2\hat{\mathbf{V}}_2^T$.
 - 3: Calculate degree matrices \mathbf{D}_2 , \mathbf{D}_{12} and \mathbf{D}_{21} for the rows of $\mathbf{A}(\mathcal{S}, \bar{\mathcal{S}})$, $\mathbf{A}(\bar{\mathcal{S}}, \mathcal{S})$ and $\mathbf{A}(\bar{\mathcal{S}}, \bar{\mathcal{S}})$ respectively.
 - 4: $\hat{\mathbf{X}} = \mathbf{D}_{21}^{-1/2}\mathbf{A}_{21}\hat{\mathbf{V}}_1\hat{\mathbf{E}}_1^{-1/2}$, where $\mathbf{A}_{21} = \mathbf{A}(\bar{\mathcal{S}}, \mathcal{S})$.
 - 5: $\mathcal{F} = \left\{ i : \|\hat{\mathbf{X}}(i, :)\|_2 \geq (1 - \epsilon_0) \max_j \|\hat{\mathbf{X}}(j, :)\|_2 \right\}$
 - 6: $\mathcal{S}_p = \text{PartitionPureNodes}\left(\hat{\mathbf{X}}(\mathcal{F}, :), \sqrt{\frac{K \min_{i \in \mathcal{F}} \mathbf{D}_2(i, i)}{4n \max_{i \in \mathcal{F}} \mathbf{D}_2(i, i)}}\right)$
 - 7: $\hat{\mathbf{X}}_p = \hat{\mathbf{X}}(\mathcal{S}_p, :)$
 - 8: Get $\hat{\beta}$, where $\hat{\beta}_i = \left\| \mathbf{e}_i^T \mathbf{D}_{21}^{1/2}(\mathcal{S}_p, \mathcal{S}_p) \hat{\mathbf{X}}_p \right\|_2^2$, $i \in [K]$
 - 9: $\hat{\mathbf{B}} = \text{diag}(\hat{\beta})$
 - 10: $\hat{\rho} = \max_i \hat{\mathbf{B}}_{ii}$
 - 11: $\hat{\mathbf{B}} = \hat{\mathbf{B}}/\hat{\rho}$
 - 12: $\hat{\Theta}(\bar{\mathcal{S}}, :) = \mathbf{D}_{21}^{1/2} \hat{\mathbf{X}} \hat{\mathbf{X}}_p^{-1} \mathbf{D}_{21}^{-1/2}(\mathcal{S}_p, \mathcal{S}_p)$
 - 13: Repeat steps with \mathbf{D}_{12} , \mathbf{A}_{12} , $\hat{\mathbf{V}}_2$, and $\hat{\mathbf{E}}_2$ to obtain parameter estimates for the remaining bipartition.
-

Algorithm 1 shows our NMF algorithm based on these geometric intuitions for inference under MMSB (henceforth, GeoNMF). The complexity of GeoNMF is dominated by the one-time eigen-decomposition in step 2. Thus this algorithm is fast and scalable. The consistency of parameters

inferred under GeoNMF is shown in the next section.

Algorithm 2 PartitionPureNodes

Input: Matrix $\mathbf{M} \in \mathbb{R}^{m \times K}$, where each row represents a pure node; a constant τ

Output: A set \mathcal{S} consisting of one pure node from each cluster.

- 1: $\mathcal{S} = \{\}, \mathcal{C} = \{\}$.
 - 2: **while** $\mathcal{C} \neq [m]$ **do**
 - 3: Randomly pick one index from $[m] \setminus \mathcal{C}$, say s
 - 4: $\mathcal{S} = \mathcal{S} \cup \{s\}$
 - 5: $\mathcal{C} = \mathcal{C} \cup \{i \in [m] \setminus \mathcal{C} : \|\mathbf{M}(s, :) - \mathbf{M}(i, :)\| \leq \tau\}$
 - 6: **end while**
-

Remark 3.1. Note that Algorithm 1 produces two sets of parameters for the two partitions of the graph \mathcal{S} and $\bar{\mathcal{S}}$. In practice one may need to have parameter estimates of the entire graph. While there are many ways of doing this, the most intuitive way would be to look at the set of pure nodes in \mathcal{S} (call this \mathcal{S}_p) and those in $\bar{\mathcal{S}}$ (call this $\bar{\mathcal{S}}_p$). If one looks at the subgraph induced by the union of all these pure nodes, then with high probability, there should be K connected components, which will allow us to match the communities.

Also note that Algorithm 2 may return $k \neq K$ clusters. However, we show in Lemma 4.4 that the pure nodes extracted by our algorithm will be highly separated and with high probability we will have $k = K$ for an appropriately chosen τ .

Finally, we note that, in our implementation, we construct the candidate pure node set \mathcal{F} (step 5 of Algorithm 1) by finding all nodes with norm within ϵ_0 multiplicative error of the largest norm. We increase ϵ_0 from a small value, until $\hat{\mathbf{X}}_p$ has condition number close to one. This is helpful when n is small, where asymptotic results do not hold.

4. Analysis

We want to prove that the sample-based estimates $\hat{\Theta}$, $\hat{\mathbf{B}}$ and $\hat{\rho}$ concentrate around the corresponding population parameters Θ , \mathbf{B} , and ρ after appropriate normalization. We will show this in several steps, which follow the steps of GeoNMF.

For the following statements, denote $\beta_{\min} = \min_a \mathbf{B}_{aa}$, $\Theta_2 = \Theta(\bar{\mathcal{S}}, :)$, where $\bar{\mathcal{S}}$ is one of the random bipartitions of $[n]$. Let \mathcal{D}_{21} be the population version of \mathbf{D}_{21} defined in Algorithm 1. Also let $\hat{\mathbf{X}}_i = \mathbf{e}_i^T \mathbf{D}_{21}^{-1/2} \mathbf{A}_{21} \hat{\mathbf{V}}_1 \hat{\mathbf{E}}_1^{-1/2}$ and its population version $\mathbf{X}_i = \sqrt{\rho} \cdot \mathbf{e}_i^T \mathcal{D}_{21}^{-1/2} \Theta_2 \mathbf{B}^{1/2}$ for $i \in [\frac{n}{2}]$.

First we show the pure nodes have the largest row norm of the population version of $\hat{\mathbf{X}}$.

Lemma 4.1. Recall that $\mathbf{X} \in \mathbb{R}^{\frac{n}{2} \times K}$. If $\Theta \sim \text{Dirichlet}(\alpha)$ with $\alpha_i = \alpha_0/K$, then $\forall i \in [\frac{n}{2}]$,

$$\|\mathbf{X}_i\|_2^2 \leq \frac{2K}{n} \max_a \theta_{ia} \left(1 + O_P \left(\sqrt{\frac{K \log n}{n}} \right) \right)$$

with probability larger than $1 - O(1/n^3)$.

In particular, if node i of subgraph $\mathbf{A}(\bar{\mathcal{S}}, \bar{\mathcal{S}})$ is a pure node ($\max_a \theta_{ia} = 1$),

$$\|\mathbf{X}_i\|_2^2 \in \frac{2K}{n} \left[1 - O_P \left(\sqrt{\frac{K \log n}{n}} \right), 1 + O_P \left(\sqrt{\frac{K \log n}{n}} \right) \right].$$

Concentration of rows of $\hat{\mathbf{X}}$. We must show that the rows of the sample $\hat{\mathbf{X}}$ matrix concentrate around a suitably rotated population version. While it is known that $\hat{\mathbf{V}}$ concentrates around suitably rotated \mathbf{V} (see the variant of Davis-Kahan Theorem presented in (Yu et al., 2015)), these results are for *columns* of the \mathbf{V} matrix, not for each *row*. The trivial bound for row-wise error would be to upper bound it by the total error, which is too crude for our purposes. To get row-wise convergence, we use sample-splitting (similar ideas can be found in (McSherry, 2001; Chaudhuri et al., 2012)), as detailed in steps 1 to 4 of GeoNMF. The key idea is to split the graph in two parts and project the adjacency matrix of one part onto eigenvectors of another part. Due to independence of these two parts, one can show concentration.

Theorem 4.2. Consider an adjacency matrix \mathbf{A} generated from MMSB(Θ, \mathbf{B}, ρ), where $\Theta \sim \text{Dirichlet}(\alpha)$ with $\alpha_i = \alpha_0/K$, whose parameters satisfy the conditions of Theorem 2.2. If $\rho n = \Omega(\log n)$, then \exists orthogonal matrix $\mathbf{O} \in \mathbb{R}^{K \times K}$ that $\forall i \in [\frac{n}{2}]$,

$$\frac{\|\hat{\mathbf{X}}_i - \mathbf{X}_i \mathbf{O}\|_2}{\|\mathbf{X}_i\|_2} = O_P \left(\frac{K^2 \sqrt{\log n}}{\beta_{\min}^{5/2} \rho \sqrt{n}} \right)$$

with probability larger than $1 - O(K^2/n^2)$.

Thus, the sample-based quantity for *each row* i converges to its population variant.

Selection of pure nodes. GeoNMF selects the nodes with (almost) the highest norm. We prove that this only selects nearly pure nodes. Let $\epsilon' = O_P \left(\frac{K^2 \sqrt{\log n}}{\beta_{\min}^{5/2} \rho \sqrt{n}} \right)$ represent the row-wise error term from Theorem 4.2.

Lemma 4.3. Let \mathcal{F} be the set of nodes with $\|\hat{\mathbf{X}}_i\|_2 \geq (1 - \epsilon_0) \max_j \|\hat{\mathbf{X}}_j\|_2$. Then $\forall i \in \mathcal{F}$,

$$\max_a \theta_{ia} \geq 1 - O_P(\epsilon_0 + \epsilon')$$

with probability larger than $1 - O(K^2/n^2)$.

We choose $\epsilon_0 = O_P(\epsilon')$ and it is straightforward to show by Lemmas 4.1, 4.3, and Theorem 4.2 that if $\epsilon_0 \geq 2\epsilon'$, then \mathcal{F} includes all pure nodes from all K communities.

Clustering of pure nodes. Once the (nearly) pure nodes have been selected, we run PartitionPureNodes (Algorithm 2) on them. We show that these nodes can form exactly K well separated clusters and each cluster only contains nodes whose θ are peaked on the same element, and PartitionPureNodes can select exactly one node from each of the K communities.

Lemma 4.4. Let $\tau = \sqrt{\frac{K \min_{i \in \mathcal{F}} \mathbf{D}_2(i, i)}{4n \max_{i \in \mathcal{F}} \mathbf{D}_2(i, i)}}$, where \mathcal{F} is defined in step 5 of Algorithm 1. If all conditions in Theorem 4.2 are satisfied, then PartitionPureNodes ($\hat{\mathbf{X}}(\mathcal{F}, :), \tau$) returns one (nearly) pure node from each of the underlying K communities with probability larger than $1 - O(K^2/n^2)$.

Concentration of $(\hat{\Theta}, \hat{\mathbf{B}}, \hat{\rho})$. GeoNMF recovers Θ using $\mathbf{D}, \hat{\mathbf{X}}$, and its pure portion $\hat{\mathbf{X}}_p$ (via the inverse $\hat{\mathbf{X}}_p^{-1}$). We first prove that $\hat{\mathbf{X}}_p^{-1}$ concentrates around its expectation.

Theorem 4.5. Let \mathcal{S}_p be the set of pure nodes extracted using our algorithm. Let $\hat{\mathbf{X}}_p$ denote the rows of $\hat{\mathbf{X}}$ indexed by \mathcal{S}_p . Then, for the orthogonal matrix \mathbf{O} from Theorem 4.2,

$$\frac{\|\hat{\mathbf{X}}_p^{-1} - (\mathbf{X}_p \mathbf{O})^{-1}\|_F}{\|\mathbf{X}_p^{-1}\|_F} = O_P \left(\frac{K^{5/2} \sqrt{\log n}}{\beta_{\min}^{5/2} \rho \sqrt{n}} \right)$$

with probability larger than $1 - O(K^2/n^2)$.

Next, we shall prove consistency for $\hat{\Theta}_2 := \hat{\Theta}(\bar{\mathcal{S}}, :)$; the proof for $\hat{\Theta}(\mathcal{S}, :)$ is similar. Let $\mathbf{D}_{21p} = \mathbf{D}_{21}(\mathcal{S}_p, \mathcal{S}_p)$.

Theorem 4.6. Let $\hat{\Theta}_2 = \mathbf{D}_{21}^{1/2} \hat{\mathbf{X}}_p^{-1} \mathbf{D}_{21p}^{-1/2}$, then \exists a permutation matrix $\mathbf{\Pi} \in \mathbb{R}^{K \times K}$ such that

$$\frac{\|\hat{\Theta}_2 - \Theta_2 \mathbf{\Pi}\|_F}{\|\Theta_2\|_F} = O_P \left(\frac{K^3 \sqrt{\log n}}{\beta_{\min}^3 \rho \sqrt{n}} \right)$$

with probability larger than $1 - O(K^2/n^2)$.

Recall that \mathbf{B} and $\hat{\mathbf{B}}$ are both diagonal matrices, with diagonal components $\{\beta_a\}$ and $\{\hat{\beta}_a\}$ respectively.

Theorem 4.7. Let $\hat{\rho} \hat{\beta}_a = \|\mathbf{e}_a^T \mathbf{D}_{21}^{1/2}(\mathcal{S}_p, \mathcal{S}_p) \hat{\mathbf{X}}_p\|_2^2$. Then, \exists a permutation matrix $\mathbf{\Pi} \in \mathbb{R}^{K \times K}$ such that $\forall a \in [K]$,

$$\hat{\rho} \hat{\beta}_a \in \rho \beta_{a'} \left[1 - O_P \left(\frac{K^{5/2} \log n}{\beta_{\min}^{5/2} \rho \sqrt{n}} \right), 1 + O_P \left(\frac{K^{5/2} \log n}{\beta_{\min}^{5/2} \rho \sqrt{n}} \right) \right]$$

for some a' such that $\mathbf{\Pi}_{a'a} = 1$, with probability larger than $1 - O(K^2/n^2)$.

Remark 4.1. While the details of our algorithms were designed for obtaining rigorous theoretical guarantees, many

of these can be relaxed in practice. For instance, while we require the Dirichlet parameters to be equal, leading to balanced cluster sizes, real data experiments show that our algorithm works well for unbalanced settings as well. Similarly, the algorithm assumes a diagonal \mathbf{B} (which is sufficient for uniqueness), but empirically works well even in the presence of off-diagonal noise. Finally, splitting the nodes into \mathcal{S} and $\bar{\mathcal{S}}$ is not needed in practice.

5. Experiments

We present results on simulated and real-world datasets. Via simulations, we evaluate the sensitivity of GeoNMF to the various MMSB parameters: the skewness of the diagonal elements of \mathbf{B} and off-diagonal noise, the Dirichlet parameter α that controls the degree of overlap, the sparsity parameter ρ , and the number of communities K . Then, we evaluate GeoNMF on Facebook and Google Plus ego networks, and co-authorship networks with upto 150,000 nodes constructed from DBLP and the Microsoft Academic Network.

Baseline methods: For the real-world networks, we compare GeoNMF against the following methods²:

- Stochastic variational inference (SVI) for MMSB (Gopalan & Blei, 2013),
- a Bayesian variant of SNMF for overlapping community detection (BSNMF) (Psorakis et al., 2011),
- the OCCAM algorithm (Zhang et al., 2014) for recovering mixed memberships, and
- the SAAC algorithm (Kaufmann et al., 2016).

For the simulation experiments, we only compare GeoNMF against SVI, since these are the only two methods based specifically on the MMSB model. BSNMF has a completely different underlying model, OCCAM requires rows of Θ to have unit ℓ_2 norm and \mathbf{B} to have equal diagonal elements, and SAAC requires Θ to be a binary matrix, while MMSB requires rows of Θ to have unit ℓ_1 norm.

Since the community identities can only be recovered up-to a permutation, in both simulated and real data experiments, we figure out the order of the communities using the well known Munkres algorithm in (Munkres, 1957).

5.1. Simulated data

Our simulations with the MMSB model are shown in *Figure 1*. We use $\alpha_i = \alpha_0/K$ for $i \in [K]$. While this leads to

²We were not to run Anandkumar et al. (2014)’s main (GPU) implementation of their algorithm because a required library CULA is no longer open source, and a complementary CPU implementation did not yield good results with default settings.

balanced clusters, note that the real datasets have clusters of different sizes and we will show that GeoNMF works consistently well even for those networks (see Section 5.2). Unless otherwise stated, we set $n = 5000$, $K = 3$, and $\alpha_0 = 1$.

Evaluation Metric: Since we have ground truth Θ , we report the relative error of the inferred MMSB parameters $\hat{\Theta}$ defined as $\min_{\Pi} \frac{\|\hat{\Theta} - \Theta \Pi\|_F}{\|\Theta\|_F}$. Here the minimum is taken over all $K \times K$ permutation matrices. For each experiment, we report the average and the standard deviation over 10 random samples. Since all the baseline algorithms only return $\hat{\Theta}$, we only report relative error of that.

Sensitivity to skewness of the diagonal of \mathbf{B} : Let $\beta = \text{diag}(\mathbf{B})$. For skewed β , different communities have different strengths of connection. We use $\beta = (0.5 - \epsilon_B, 0.5, 0.5 + \epsilon_B)$ and plot the relative error against varying ϵ_B . *Figure 1(a)* shows that GeoNMF has much smaller error than SVI, and is robust to β over a wide range.

Sensitivity to off-diagonal element \mathbf{B} : While SNMF is identifiable only for diagonal \mathbf{B} , we still test GeoNMF in the setting where all off-diagonal entries of \mathbf{B} have noise ϵ . *Figure 1(b)* shows once again that GeoNMF is robust to such noise, and is much more accurate than SVI.

Sensitivity to α_0 : In *Figure 1(c)*, the relative error is plotted against increasing α_0 ; larger values corresponding to larger overlap between communities. Accuracy degrades with increasing overlap, as expected, but GeoNMF is much less affected than SVI.

Sensitivity to ρ : *Figure 1(d)* shows relative error against increasing ρ . For dense networks, both GeoNMF and SVI perform similarly, but the error of SVI increases drastically in the sparse regime (small ρ).

Scalability: *Figure 1(f)* shows the wall-clock time for networks of different sizes. Both GeoNMF and SVI scale linearly with the number of nodes, but SVI is about 100 times slower than GeoNMF.

5.2. Real-world data

Datasets: For real-data experiments, we use two kinds of networks:

- *Ego networks:* We use the Facebook and Google Plus (G-plus) ego networks, where each node can be part of multiple “circles” or “communities.”
- *Co-authorship networks*³: We construct co-authorship networks from DBLP (each community is a group

³Available at <http://www.cs.utexas.edu/~xmao/coauthorship>

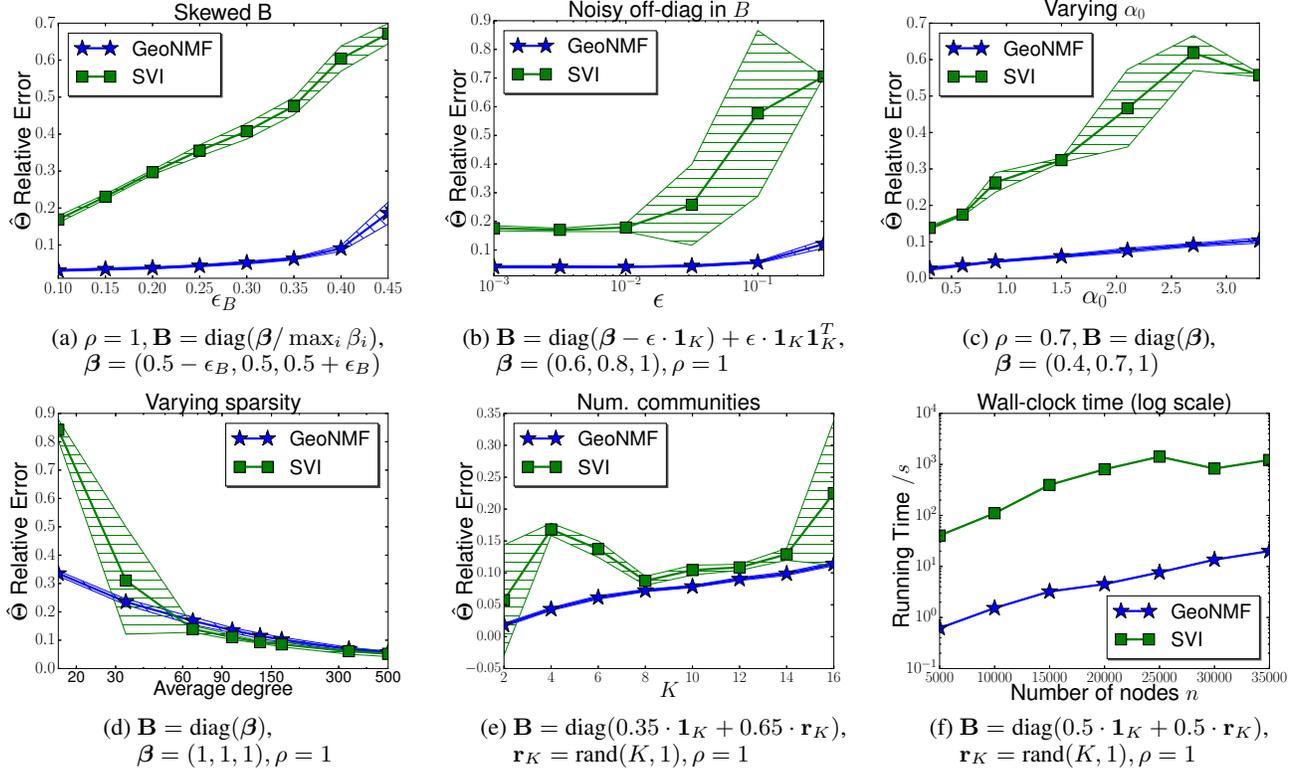


Figure 1. (a)-(e) Simulation results for varying parameters. (f) Running time.

of conferences), and from the Microsoft Academic Graph (each community is denoted by a “field of study” (FOS) tag). Each author’s θ vector is constructed by normalizing the number of papers he/she has published in conferences in a subfield (or papers that have the FOS tag).

We preprocessed the networks by recursively removing isolated nodes, communities without any pure nodes, and nodes with no community assignments. For the ego networks we pick networks with at least 200 nodes and the average number of nodes per community (n/K) is at least 100, giving us 3 Facebook and 40 G-plus networks. For the co-authorship networks, all communities have enough pure nodes, and after removing isolated nodes, the networks have more than 200 nodes and n/K is larger than 100. The statistics of the networks (number of nodes, average degree, number of clusters, degree of overlap etc.) are shown in Table 1. The overlap ratio is the number of overlapping nodes divided by the number of nodes. The different networks have the following subfields:

- DBLP1: Machine Learning, Theoretical Computer Science, Data Mining, Computer Vision, Artificial Intelligence, Natural Language Processing
- DBLP2: Networking and Communications, Systems,

Information Theory

- DBLP3: Databases, Data Mining, World Web Wide
- DBLP4: Programming Languages, Software Engineering, Formal Methods
- DBLP5: Computer Architecture, Computer Hardware, Real-time and Embedded Systems, Computer-aided Design
- MAG1: Computational Biology and Bioinformatics, Organic Chemistry, Genetics
- MAG2: Machine Learning, Artificial Intelligence, Mathematical Optimization

Evaluation Metric: For real data experiments, we construct Θ as follows. For the ego-networks every node has a binary vector which indicates which circle (community) each node belongs to. We normalize this to construct Θ . For the DBLP and Microsoft Academic networks we construct a row of Θ by normalizing the number of papers an author has in different conferences (ground truth communities). We present the averaged Spearman rank correlation coefficients (RC) between $\Theta(:, a), a \in [K]$ and $\hat{\Theta}(:, \sigma(a))$,

Dataset	Facebook	G-plus	DBLP1	DBLP2	DBLP3	DBLP4	DBLP5	MAG1	MAG2
# nodes n	362.0 (± 148.5)	656.2 (± 422.0)	30,566	16,817	13,315	25,481	42,351	142,788	108,064
# communities K	2.3 (± 0.58)	2.8 (± 1.74)	6	3	3	3	4	3	3
Average Degree	56.8 (± 32.3)	103.4 (± 74.9)	8.9	7.6	8.5	5.2	6.8	12.4	16.0
Overlap %	19.3 (± 29.8)	26.5 (± 32.4)	18.2	14.9	21.1	14.4	18.5	3.3	3.8

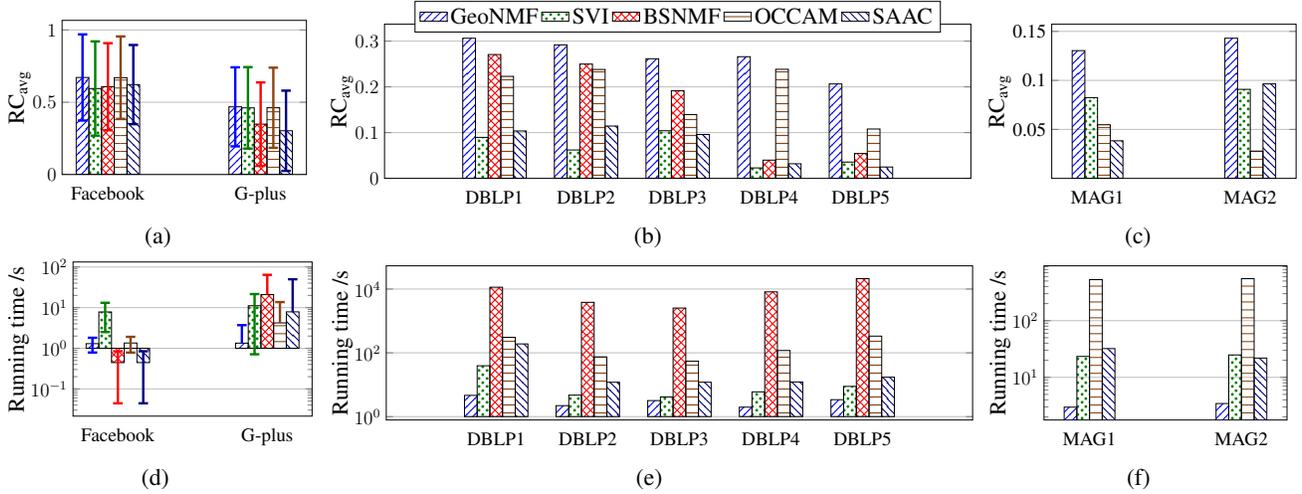


Figure 2. RC_{avg} and running time (log scale) for real datasets.

where σ is a permutation of $[K]$. The formal definition is:

$$RC_{avg}(\hat{\Theta}, \Theta) = \frac{1}{K} \max_{\sigma} \sum_{i=1}^K RC(\hat{\Theta}(:, i), \Theta(:, \sigma(i))).$$

It is easy to see that $RC_{avg}(\hat{\Theta}, \Theta)$ takes value from -1 to 1, and higher is better. Since SAAC returns binary assignment, we compute its RC_{avg} against the binary ground truth.

Performance: We report the RC_{avg} score in Figure 2(a) averaged over different Facebook and G-plus networks; in Figure 2(b) for five DBLP networks, and in Figure 2(c) for two MAG networks. We show the time in seconds (log-scale) in Figure 2(d) averaged over Facebook and G-plus networks; in Figure 2(e) for DBLP networks and in Figure 2(f) for MAG networks. We averaged over the Facebook and G-plus networks because all the performances were similar.

- For small networks like Facebook and G-plus, all algorithms perform equally well both in speed and accuracy, although GeoNMF is fast even for relatively larger G-plus networks.
- DBLP is sparser, and as a result the overall rank correlation decreases. However, GeoNMF consistently performs well. While for some networks, BSNMF and OCCAM have comparable RC_{avg} , they are much slower than GeoNMF.

- MAG is larger (hundreds of thousands of nodes) than DBLP. For these networks we could not even run BSNMF because of memory issues. Again, GeoNMF performs consistently well while outperforming others in speed.

Estimating K : While we assume that K is known apriori, K can be estimated using the USVT estimator (Chatterjee et al., 2015). For the simulated graphs, when average degree is above ten, USVT estimates K correctly. However for the real graphs, which are often sparse, it typically overestimates the true number of clusters.

6. Conclusions

This paper explored the applicability of symmetric NMF algorithms for inference of MMSB parameters. We showed broad conditions that ensure identifiability of MMSB, and then proved sufficiency conditions for the MMSB parameters to be uniquely determined by a general symmetric NMF algorithm. Since general-purpose symmetric NMF algorithms do not have optimality guarantees, we propose a new algorithm, called GeoNMF, that adapts symmetric NMF specifically to MMSB. GeoNMF is not only provably consistent, but also shows good accuracy in simulated and real-world experiments, while also being among the fastest approaches.

References

- Airoldi, Edoardo M, Blei, David M, Fienberg, Stephen E, and Xing, Eric P. Mixed membership stochastic block-models. *Journal of Machine Learning Research*, 9: 1981–2014, 2008.
- Anandkumar, Animashree, Ge, Rong, Hsu, Daniel J, and Kakade, Sham M. A tensor approach to learning mixed membership community models. *Journal of Machine Learning Research*, 15(1):2239–2312, 2014.
- Arora, Sanjeev, Ge, Rong, and Moitra, Ankur. Learning topic models—going beyond svd. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pp. 1–10. IEEE, 2012.
- Arora, Sanjeev, Ge, Rong, Halpern, Yonatan, Mimno, David M, Moitra, Ankur, Sontag, David, Wu, Yichen, and Zhu, Michael. A practical algorithm for topic modeling with provable guarantees. In *ICML*, pp. 280–288, 2013.
- Chang, Jonathan. LDA: Collapsed gibbs sampling methods for topic models, 2012. URL <http://cran.r-project.org/web/packages/lda/index.html>.
- Chatterjee, Sourav et al. Matrix estimation by universal singular value thresholding. *The Annals of Statistics*, 43 (1):177–214, 2015.
- Chaudhuri, Kamalika, Graham, Fan Chung, and Tsias, Alexander. Spectral clustering of graphs with general degrees in the extended planted partition model. In *COLT*, volume 23, pp. 35–1, 2012.
- Chen, Jingchun and Yuan, Bo. Detecting functional modules in the yeast protein–protein interaction network. *Bioinformatics*, 22(18):2283–2290, 2006.
- Davis, Chandler and Kahan, William Morton. The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970.
- Gopalan, Prem K and Blei, David M. Efficient discovery of overlapping communities in massive networks. *Proceedings of the National Academy of Sciences*, 110(36): 14534–14539, 2013.
- Huang, Kejun, Sidiropoulos, Nicholas, and Swami, Ananthram. Non-negative matrix factorization revisited: Uniqueness and algorithm for symmetric decomposition. *Signal Processing, IEEE Transactions on*, 62(1):211–224, 2014.
- Huang, Kejun, Fu, Xiao, and Sidiropoulos, Nikolaos D. Anchor-free correlated topic modeling: Identifiability and algorithm. In *Advances in Neural Information Processing Systems*, pp. 1786–1794, 2016.
- Kaufmann, Emilie, Bonald, Thomas, and Lelarge, Marc. A spectral algorithm with additive clustering for the recovery of overlapping communities in networks. In *International Conference on Algorithmic Learning Theory*, pp. 355–370. Springer, 2016.
- Kuang, Da, Yun, Sangwoon, and Park, Haesun. Symmf: nonnegative low-rank approximation of a similarity matrix for graph clustering. *Journal of Global Optimization*, 62(3):545–574, 2015.
- Ley, Michael. The dblp computer science bibliography: Evolution, research issues, perspectives. In *International symposium on string processing and information retrieval*, pp. 1–10. Springer, 2002.
- Lu, Zongqing, Sun, Xiao, Wen, Yonggang, Cao, Guohong, and La Porta, Thomas. Algorithms and applications for community detection in weighted networks. *Parallel and Distributed Systems, IEEE Transactions on*, 26(11): 2916–2926, 2015.
- Mcauley, Julian and Leskovec, Jure. Discovering social circles in ego networks. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 8(1):4, 2014.
- McSherry, Frank. Spectral partitioning of random graphs. In *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*, pp. 529–537. IEEE, 2001.
- Munkres, James. Algorithms for the assignment and transportation problems. *Journal of the society for industrial and applied mathematics*, 5(1):32–38, 1957.
- Press, William H., Teukolsky, Saul A., Vetterling, William T., and Flannery, Brian P. *Numerical Recipes in C*. Cambridge University Press, 2nd edition, 1992.
- Psorakis, Ioannis, Roberts, Stephen, Ebden, Mark, and Sheldon, Ben. Overlapping community detection using bayesian non-negative matrix factorization. *Phys. Rev. E*, 83:066114, Jun 2011.
- Ray, A., Ghaderi, J., Sanghavi, S., and Shakkottai, S. Overlap graph clustering via successive removal. In *2014 52nd Annual Allerton Conference on Communication, Control, and Computing, Allerton 2014*, pp. 278–285, United States, 1 2015. Institute of Electrical and Electronics Engineers Inc. doi: 10.1109/ALLERTON.2014.7028467.
- Sinha, Arnab, Shen, Zhihong, Song, Yang, Ma, Hao, Eide, Darrin, Hsu, Bo-june Paul, and Wang, Kuansan. An overview of microsoft academic service (mas) and applications. In *Proceedings of the 24th international conference on world wide web*, pp. 243–246. ACM, 2015.

Soundarajan, Sucheta and Hopcroft, John. Using community information to improve the precision of link prediction methods. In *Proceedings of the 21st international conference companion on World Wide Web*, pp. 607–608. ACM, 2012.

Tang, Minh, Sussman, Daniel L, Priebe, Carey E, et al. Universally consistent vertex classification for latent positions graphs. *The Annals of Statistics*, 41(3):1406–1430, 2013.

Wang, Fei, Li, Tao, Wang, Xin, Zhu, Shenghuo, and Ding, Chris. Community discovery using nonnegative matrix factorization. *Data Mining and Knowledge Discovery*, 22(3):493–521, 2011.

Wang, Xiao, Cao, Xiaochun, Jin, Di, Cao, Yixin, and He, Dongxiao. The (un) supervised nmf methods for discovering overlapping communities as well as hubs and outliers in networks. *Physica A: Statistical Mechanics and its Applications*, 446:22–34, 2016.

Yu, Yi, Wang, Tengyao, Samworth, Richard J, et al. A useful variant of the davis–kahan theorem for statisticians. *Biometrika*, 102(2):315–323, 2015.

Zhang, Yuan, Levina, Elizaveta, and Zhu, Ji. Detecting overlapping communities in networks using spectral methods. *arXiv preprint arXiv:1412.3432*, 2014.