

SPECTRAL AND MATRIX FACTORIZATION METHODS FOR CONSISTENT COMMUNITY DETECTION IN MULTI-LAYER NETWORKS

BY SUBHADEEP PAUL¹ AND YUGUO CHEN²

¹*Department of Statistics, The Ohio State University, paul.963@osu.edu*

²*Department of Statistics, University of Illinois, yuguo@illinois.edu*

We consider the problem of estimating a consensus community structure by combining information from multiple layers of a multi-layer network using methods based on the spectral clustering or a low-rank matrix factorization. As a general theme, these “intermediate fusion” methods involve obtaining a low column rank matrix by optimizing an objective function and then using the columns of the matrix for clustering. However, the theoretical properties of these methods remain largely unexplored. In the absence of statistical guarantees on the objective functions, it is difficult to determine if the algorithms optimizing the objectives will return good community structures. We investigate the consistency properties of the global optimizer of some of these objective functions under the multi-layer stochastic blockmodel. For this purpose, we derive several new asymptotic results showing consistency of the intermediate fusion techniques along with the spectral clustering of mean adjacency matrix under a high dimensional setup, where the number of nodes, the number of layers and the number of communities of the multi-layer graph grow. Our numerical study shows that the intermediate fusion techniques outperform late fusion methods, namely spectral clustering on aggregate spectral kernel and module allegiance matrix in sparse networks, while they outperform the spectral clustering of mean adjacency matrix in multi-layer networks that contain layers with both homophilic and heterophilic communities.

1. Introduction. The study of multi-layer networks has received significant interest recently, driven by its myriad of applications in neuroscience, economics, genetics and social sciences [4, 12, 14, 19]. A multi-layer network is a powerful representation of relational data with the nodes representing the entities of interest and the network layers representing the multiple relations among those entities. While the term “multi-layer network” is often used in a more general context, we focus our attention only on a network where the nodes are connected only within a layer and there are no inter-layer edges (such networks are also called “multiplex networks” in the literature).

A dynamic or time-varying network represents different states of a single network over time. A dynamic network can also be represented as a multi-layer network, with the same node in consecutive time period usually being linked by an edge to respect the time ordering [2, 10, 19]. When appropriate for the application, for example, in the problem of consensus community detection, we can ignore the time order and consider a time-varying network as a regular multi-layer network with no inter-layer edges [5, 12].

The problem of consensus community detection in multi-layer and dynamic networks has many important applications. Often in such networks one underlying community structure is in force while the different layers of interactions are merely different manifestations of the unobserved community structure. For example, in the multi-layer twitter networks in Greene and Cunningham [11], ground truth community memberships can be attributed to the users

Received April 2017; revised December 2018.

MSC2010 subject classifications. Primary 62F12, 62H30; secondary 90B15, 15A23.

Key words and phrases. Community detection, consistency, co-regularized spectral clustering, multi-layer networks, multi-layer stochastic block model, orthogonal linked matrix factorization.

(nodes) based on attributes more fundamental and independent of the observed twitter interactions (e.g., political views, country of origin, football clubs), whereas the interactions provide multiple sources of information about the same latent community structure. Combining information from these multiple sources would then lead to enhanced performance in the learning task. Moreover, different representations of the same phenomenon often provide complimentary information, any one of which is not sufficient to describe the underlying process (see Liu et al. [17] and the examples therein).

Even in situations where the hypothesis of a single constant community structure may not be true, for example, in the analysis of dynamic brain networks, it is still often desirable to obtain a consensus partition that does not vary over time, but is a static average partition that remains in force throughout the experiment. Such an overall partition is crucial to obtain stable modules of brain regions as baseline for computing measures of local and global dynamism in the brain, for example, “flexibility” and “integration” in Bassett et al. [3] and Braun et al. [5].

The present problem is also related to a more general class of problems that generally goes under the theme of multi-view clustering and has received considerable attention over the last decade, particularly in the computer science community. Numerous methods have been proposed to combine information from multiple views of a multi-view relational data for clustering. The goal is usually to leverage the diversity and often complimentary nature of the information in different layers to outperform simply summing the layers or using any one of the layers [17]. A great many of those methods use spectral clustering or a low rank matrix factorization as a basis [9, 15, 17, 18, 22, 29, 34, 39].

The “linked matrix factorization” algorithm in Tang, Lu and Dhillon [34] and “RESCAL” algorithm in Nickel, Tresp and Kriegel [22] approximate the adjacency matrices in each layer of a multi-layer graph, or each slice of a three way tensor, with a low rank symmetric matrix factorization. While one of the factors is shared the other one varies across layers or slices. Although the algorithms employed in the two papers are quite different, the factorization in both cases is computed by minimizing an identical joint Frobenius norm objective function. Dong et al. [9] use similar common low rank matrix factorization ideas with a slightly different objective function to obtain a “joint spectrum” of a multi-layer graph which is subsequently used for clustering.

The co-regularized spectral clustering in [15] with centroid based co-regularization maximizes the combined normalized cut objective function over the Laplacian matrices from all views of the data, subject to a smoothness penalty. This idea is similar to the evolutionary spectral clustering used in Chi et al. [7] for clustering dynamic networks with a temporal smoothness penalty, and is part of a general theme of co-regularization in multi-view machine learning [38]. The co-regularization framework was extended to “joint nonnegative matrix factorization” using a Frobenius norm based objective function in Liu et al. [17]. See Sun [32] and Xu, Tao and Xu [38] for surveys of multi-view learning methods.

However, there is a lack of theoretical understanding of the objective functions in these spectral and matrix factorization based methods. Researchers often rely on simulations and applications to specific datasets to compare the methods. However, this approach fails to explore different scenarios that might arise in practice. For example, in multi-layer network applications, the component layers might have very different sparsity, signal quality and node degree distributions. Hence it is important to explore the utility of the methods under different statistical models and asymptotic settings through a principled theoretical study.

In this article we investigate the consistency properties of various methods for community detection under data generated from a multi-layer network model, the multi-layer stochastic blockmodel (MLSBM) [1, 12, 23, 24, 26, 30, 36, 37]. We derive several asymptotic results to show consistency of the global optimizers of co-regularized spectral clustering and orthogonal linked matrix factorization under a high dimensional asymptotic setup where the number

of nodes, the number of layers and the number of communities all grow. We use slight variations of the original algorithms to compute the solutions to the respective optimization problems. We note that both algorithms are not guaranteed to reach a global optimum. The present paper is an attempt to prove goodness of the objective functions rather than the algorithms, and is concerned with the following question: If it is possible to compute a global optimum or approximate one reasonably, will the global solution be consistent under a random graph model, namely the MLSBM?

In addition to the two methods mentioned above, we also consider two baseline methods previously used in literature. The first method is performing spectral clustering [21, 28] on the mean of the adjacency matrices from different layers of the multi-layer network. This method has also been considered in Tang, Lu and Dhillon [34] and Dong et al. [9] as a baseline method and is generally thought to be a simple but effective procedure [15]. In addition to including this method in our numerical comparisons, we also study its asymptotic consistency under MLSBM.

The second baseline method first computes a low dimensional spectral embedding (matrix of eigenvectors corresponding to top eigenvalues) and creates a spectral kernel for each layer, and then aggregates these kernel functions. A single layer community detection method is then applied to this aggregate spectral kernel [34]. Another variation of this idea is to compute the community assignments in each layer independently using a single layer method (e.g., modularity or spectral clustering), and then create a “module allegiance matrix” on which a single layer community detection method can be applied to compute the consensus communities [5].

The rest of the article is organized as follows. Section 2 describes the methods and algorithms considered in the article. Section 3 describes the MLSBM, defines mis-clustering rate and proves correct recovery in the noiseless case. Section 4 describes the consistency results. Section 5 contains a simulation study to numerically evaluate the methods. Section 6 gives concluding remarks. All the proofs are given in the Supplementary Material [25].

2. Methods and algorithms. We define an undirected multi-layer network with M layers as a collection of graphs $\mathcal{G} = \{G^{(1)}, \dots, G^{(M)}\}$ over a common set of n vertices. The vertices represent the entities/actors, while the layers represent different types of interactions among the entities. For the layer of the m th type, we define the adjacency matrix $A^{(m)}$ corresponding to that layer as follows: $A_{ij}^{(m)} = 1$, if there is an edge of type m between nodes i and j , and $A_{ij}^{(m)} = 0$, otherwise.

We define the vector of degrees of node i as $\mathbf{d}_i = \{d_i^{(m)}; m = 1, \dots, M\}$, where $d_i^{(m)} = \sum_j A_{ij}^{(m)}$ is its degree of the m th type. Then the normalized graph Laplacian matrix for the m th layer can be defined as $L^{(m)} = (D^{(m)})^{-1/2} A^{(m)} (D^{(m)})^{-1/2}$, where $D^{(m)}$ is a diagonal matrix with the degrees of the m th type of the nodes as elements, that is, $D_{ii}^{(m)} = d_i^{(m)}$. Together the M adjacency matrices create the three-way $n \times n \times M$ adjacency tensor of the multi-layer network $\mathbb{A} = \{A^{(1)}, \dots, A^{(M)}\}$. The corresponding Laplacian tensor is defined as $\mathbb{L} = \{L^{(1)}, \dots, L^{(M)}\}$. We denote the number of communities in the network by k . It will be assumed to be known throughout the paper. We use the notations $\|\cdot\|_2$, $\|\cdot\|_F$ and $\|\cdot\|_\Sigma$ to denote the spectral (operator) norm, the Frobenius norm and the trace norm, respectively, while $\text{tr}(\cdot)$ denotes the matrix trace. We will use $\sin \Theta(U, V)$ to denote the diagonal matrix whose elements are sines of the principle angles between the subspaces \mathcal{U} and \mathcal{V} , spanned by the columns of the matrices U and V respectively (Definition 1.5.3 in Stewart and Sun [31]).

We consider the following methods and algorithms for consensus clustering in multi-layer networks. The first two methods are so called “intermediate fusion” techniques whereby the multiple layers are integrated through a clustering objective function [17]. Such methods are often preferred over “early” and “late” fusion techniques due to superior performance [40].

2.1. *Linked matrix factorization.* The first of the intermediate fusion methods is the linked matrix factorization (LMF) for clustering multiple graphs in Tang, Lu and Dhillon [34]. Our adaptation of the method is slightly different from the one described in [34] in the sense that we enforce the columns of the shared factor to be strictly orthonormal and consequently drop the Frobenius norm regularization term (indeed this has been suggested in Tang, Lu and Dhillon [34]). In our simulations, we found the performance of both methods to be the same. To avoid confusion, we call our adaptation the orthogonal LMF (OLMF). Note that LMF has the identical objective function as the RESCAL algorithm, which is a three-way tensor factorization for learning in multi-relational data [22]. However the algorithm for RESCAL is different from that of LMF.

The OLMF solves the following optimization problem on the adjacency tensor of a multi-layer network:

$$(2.1) \quad [\hat{P}, (\hat{\Lambda}^{(1)}, \dots, \hat{\Lambda}^{(M)})] = \arg \min_{P^T P = I} \sum_{m=1}^M \|A^{(m)} - P \Lambda^{(m)} P^T\|_F^2,$$

where $P \in \mathbb{R}^{n \times k}$ is a common factor matrix and $\Lambda^{(m)} \in \mathbb{R}^{k \times k}$ are M layer specific symmetric factor matrices. This is equivalent to the following optimization problem (see the Supplementary Material [25] for a proof):

$$(2.2) \quad \begin{aligned} \hat{P} &= \arg \max_{P^T P = I} \sum_{m=1}^M \|P^T A^{(m)} P\|_F^2, \\ \hat{\Lambda}^{(m)} &= \hat{P}^T A^{(m)} \hat{P}, \quad m = \{1, \dots, M\}. \end{aligned}$$

We will refer to the objective function in (2.2) as $F(\mathbb{A}, P)$. While we require P to have orthonormal columns, we do not put any constraint on the $\Lambda^{(m)}$ matrices, and specifically we do not require them to be diagonal matrices. Note that in general $[\hat{P}, \hat{\Lambda}^{(m)}]$ is not the solution of the problem of finding the best at most rank k approximating matrix for $L^{(m)}$. Hence in general, the matrices $\hat{\Lambda}^{(m)}$ are not the diagonal matrices of singular values. Intuitively the shared factor P is expected to capture the common characteristics of the nodes in a multi-layer network including the latent community structure, while the different $\Lambda^{(m)}$ matrices capture the layer/relation specific characteristics.

We propose a BFGS algorithm to solve the OLMF optimization problem, similar to the algorithm in Tang, Lu and Dhillon [34]. The gradients are given by

$$\begin{aligned} \frac{\partial O}{\partial P} &:= - \sum_m (I - P P^T) A^{(m)} P \Lambda^{(m)}, \\ \frac{\partial O}{\partial \Lambda^{(m)}} &:= - P^T (A^{(m)} - P \Lambda^{(m)} P^T) P, \quad m = 1, \dots, M, \end{aligned}$$

where O denotes the objective function in (2.1). Once the algorithm converges, we cluster the rows of the matrix P using the k -means algorithm. Since each row in P corresponds to one of the nodes, this gives a community assignment for the nodes.

2.2. *Co-regularized spectral clustering.* The second intermediate fusion method we study is the co-regularization based approach to multi-layer spectral clustering due to Kumar, Rai and Daume [15]. The idea of co-regularization has also been previously applied to various learning problems [38]. We adopt the centroid based co-regularization method from Kumar, Rai and Daume [15] unchanged in the context of multi-layer networks. The method,

applied to the adjacency tensor, is based on solving the following optimization problem:

$$(2.3) \quad [\hat{U}^{(1)}, \dots, \hat{U}^{(M)}, \hat{U}^*] = \underset{\substack{U^{(m)T} U^{(m)} = I, \forall m, \\ U^{*T} U^* = I}}{\arg \max} \sum_{m=1}^M \{ \text{tr}(U^{(m)T} A^{(m)} U^{(m)}) \\ + \gamma_m \text{tr}(U^{*T} U^{(m)} U^{(m)T} U^*) \},$$

where $U^{(1)}, \dots, U^{(M)}$ and U^* are $n \times k$ matrices with orthonormal columns. We denote \mathbb{U} as the tensor containing the matrices $\{U^{(1)}, \dots, U^{(M)}\}$. The objective function of the optimization problem in (2.3) is denoted as $F(\mathbb{A}, \mathbb{U}, U^*)$. The optimization problem can be easily solved by alternating eigen-decomposition of the matrices $A^{(m)} - \gamma_m U^* U^{*T}$ and $\sum_m \gamma_m U^{(m)} U^{(m)T}$ [15]. After the algorithm converges, consensus community assignments for the nodes can be obtained by clustering the rows of the matrix \hat{U}^* with the k -means algorithm.

Note that the objective function contains two parts. The first part is the usual association cut spectral clustering objective function for different layers. The second part is a penalty function that seeks to maximize the cohesion between the eigenspaces obtained from different layers. To see this, we have the following proposition that characterizes the second part in terms of $\|\sin \Theta(U^{(m)}, U^*)\|_F$, which measures the distance between the column spaces spanned by $U^{(m)}$ and U^* [31]. The proof of this proposition, along with all lemmas and theorems, can be found in the Supplementary Material [25].

PROPOSITION 1. For $U^{(m)}$ and U^* as defined above, we have

$$\begin{aligned} \text{tr}(U^{*T} U^{(m)} U^{(m)T} U^*) &= k - \frac{1}{2} \|U^* U^{*T} - U^{(m)} U^{(m)T}\|_F^2 \\ &= k - \|\sin \Theta(U^{(m)}, U^*)\|_F^2. \end{aligned}$$

The penalty function alone is maximized when all the subspaces are identical, since $\|\sin \Theta(U^{(m)}, U^*)\|_F$ is 0 when the subspaces spanned by $U^{(m)}$ and U^* are identical [31]. Hence the objective function represents a trade-off between optimizing the community structure in each layer (which might be noisy) and maintaining similarity with the mean community structure. The weights γ_m 's should be chosen to reflect both the desired trade-off between these two competing goals and the relative importance of the different layers. In particular, small values of γ_m 's will prevent sharing information across layers, which will result in estimates of $U^{(m)}$ being the one that is best for its own layer and the U^* simply being the matrix of eigenvectors of $\sum_m \gamma_m U^{(m)} U^{(m)T}$. On the other hand, large values of γ_m 's will ensure the $U^{(m)}$'s try to achieve similarity with a common U^* in expense of being sub-optimal for its own layer.

2.3. *Spectral clustering on mean adjacency matrix.* The first of the two baseline procedures we consider collapses the multi-layer network into a single layer network by taking the mean of the adjacency matrices from each of the layers. The usual single layer spectral clustering algorithm [21, 28] is then applied to the resultant matrix. This procedure can be thought of as an ‘‘early integration’’ or ‘‘early fusion’’ technique, since data from multiple layers are aggregated before any processing is made [40]. Spectral clustering on some form of the aggregate matrix has appeared as a ‘‘baseline procedure’’ in Tang, Lu and Dhillon [34], Kumar, Rai and Daume [15], Dong et al. [9] and Tang, Wang and Liu [33]. In particular, consensus community detection proceeds through spectral clustering of the matrix $\bar{A} = \frac{1}{M} \sum_{m=1}^M A^{(m)}$. Consistency results for this method under the stochastic block model (SBM) were derived in

Han, Xu and Airolidi [12] in the scenario when the number of layers grows but the number of nodes does not. Chen and Hero III [6] also derived phase transition results for a weighted version of this method under a model they characterize as “multi-layer signal plus noise model.”

2.4. Aggregate spectral kernel and module allegiance matrix. The other baseline method we consider is a “late fusion” technique, where we first compute the eigenvector matrices $U^{(m)}$ ’s corresponding to the top k eigenvalues from each of the M layers of the graph and construct the aggregate spectral kernel matrix

$$K_{n \times n} = \frac{1}{M} \sum_{m=1}^M U^{(m)} U^{(m)T}.$$

However, instead of using kernel k -means to cluster the resulting matrix K as in Tang, Lu and Dhillon [34] and Dong et al. [9], we apply spectral clustering to this matrix again to obtain the community assignments. We call this method “aggregate spectral kernel.” This is in spirit of clustering the “module allegiance matrix” described in Braun et al. [5], where community assignment for each layer is first obtained using the Newman–Girvan modularity [20], and subsequently an $n \times n$ module allegiance matrix is formed, each of whose elements counts the number of times two nodes appear in the same module.

We use both the aggregate spectral kernel and the module allegiance matrix methods in our numerical study. It is worth pointing out that these methods are distinct from the majority voting method described in [12, 23]. Although, much like the majority voting, these methods process each layer separately and fuse information later, one advantage is that both the aggregate spectral kernel and module allegiance matrix methods avoid the cumbersome issue of label switching ambiguity. To see this, assume we have two community assignment matrices Z_1 and Z_2 with $Z_1 = Z_2 P$, where P is a permutation matrix, that is, Z_2 gives the same community assignments as Z_1 but with its labels switched. However when we compute the module allegiance matrix, we have $Z_1 Z_1^T = Z_2 P P^T Z_2^T = Z_2 Z_2^T$. The same is true for the aggregate spectral kernel. Intuitively, for each element they are concerned with whether two nodes belong to the same community or not, irrespective of which community that is. Hence they do not require solving a linear sum assignment problem as is required for majority voting.

3. Models and mis-clustering. The multi-layer stochastic block model (MLSBM) is a statistical model of multi-layer networks with a shared latent community structure [12, 13, 23]. We define the k block, M layer, n node MLSBM as follows. Each node of the network is assigned a community label vector of length k , which takes the value of 1 at the position corresponding to its community and 0 in all other positions. Let Z denote the $n \times k$ community assignment matrix whose i th row Z_i is the community label vector for the i th node.

Given the community labels of two nodes, the edges between them in different layers are formed independently following a Bernoulli distribution with a probability that depends only on the community assignments and the relation the edge represents (i.e., type or layer of the edge). Hence within a community the nodes have “stochastic equivalence” in the sense that the probability of an edge formation (in any layer) with another node is the same for all the nodes in a community. We further assume that there is at least one node in each community which implies that there is at least one nonzero element in each column of Z .

The k block, M layer, n node MLSBM with parameters $[Z, \mathbb{B} = \{B^{(1)}, \dots, B^{(M)}\}]$ can be written in the matrix form as

$$(3.1) \quad E(A^{(m)}) = \mathcal{A}^{(m)} = Z B^{(m)} Z^T, \quad B^{(m)} \in [0, 1]^{k \times k}, \quad Z \in \{0, 1\}^{n \times k},$$

where the matrices $B^{(m)}$ are $k \times k$ nonnegative symmetric matrices of probabilities. For our analysis we will assume varying constraints on the rank of $B^{(m)}$'s. A similar rank based constraint is a standard assumption in the analysis of spectral clustering for single layer SBM as well [16, 28]. We will refer to the matrix $\mathcal{A}^{(m)}$ as the population adjacency matrix for the m th layer and the tensor $\mathcal{A} = \{\mathcal{A}^{(1)}, \dots, \mathcal{A}^{(M)}\}$ as the population adjacency tensor.

3.1. Correct recovery in the noiseless case. Before we can tackle consistency of the methods, the first question that needs to be answered is whether a method can correctly recover the community assignments from the true population adjacency tensor when there is no sampling noise involved. The following lemma shows that OLMF, co-regularized spectral clustering, spectral clustering of mean adjacency matrix, and aggregate spectral kernel, all can correctly identify the node community labels from the population adjacency tensor of MLSBM.

LEMMA 1. *Let $\mathcal{A} = \{\mathcal{A}^{(1)}, \dots, \mathcal{A}^{(M)}\}$ be the three-way $n \times n \times M$ population adjacency tensor for MLSBM $[Z, \mathbb{B}]$ with each of the M slices $\mathcal{A}^{(m)} \in \mathbb{R}^{n \times n}$ defined as in (3.1). Then we have the following results:*

(i) *The optimization problem in (2.1) of orthogonal linked matrix factorization applied to the tensor \mathcal{A} has $\bar{P} = ZQ^{-1/2}$, $\bar{\Lambda}^{(m)} = Q^{1/2}B^{(m)}Q^{1/2}$, $m = 1, \dots, M$, as the unique solution up to an orthogonal matrix, where $Q = Z^T Z$ provided at least one of the $B^{(m)}$'s is full rank. Further $Z_i Q^{-1/2} = Z_j Q^{-1/2}$ if and only if $Z_i = Z_j$.*

(ii) *The optimization problem in (2.3) of co-regularized spectral clustering applied to the tensor \mathcal{A} has $\bar{U}^{(m)} = Z\mu^{(m)}$, $m = 1, \dots, M$, $\bar{U}^* = ZQ$ as the unique solution up to an orthogonal matrix, where $\mu^{(m)}$ and Q are invertible matrices provided each of the $B^{(m)}$'s is full rank. Further $Z_i Q = Z_j Q$ if and only if $Z_i = Z_j$.*

(iii) *The matrix containing the eigenvectors corresponding to the k largest eigenvalues of $\bar{A} = \frac{1}{M} \sum_{i=1}^M \mathcal{A}^{(m)}$ is ZQ for some invertible matrix $Q \in \mathbb{R}^{k \times k}$ provided the matrix $\frac{1}{M} \sum_{i=1}^M B^{(m)}$ is full rank. Further $Z_i Q = Z_j Q$ if and only if $Z_i = Z_j$.*

(iv) *Define $\bar{K} = \frac{1}{M} \sum_{i=1}^M \bar{U}^{(m)} \bar{U}^{(m)T}$, where $\bar{U}^{(m)}$ is the matrix of eigenvectors corresponding to the largest k eigenvalues of $\mathcal{A}^{(m)}$. The matrix containing the eigenvectors corresponding to the k largest eigenvalues of \bar{K} is ZQ for some invertible matrix $Q \in \mathbb{R}^{k \times k}$ provided each of the $B^{(m)}$'s is full rank. Further $Z_i Q = Z_j Q$ if and only if $Z_i = Z_j$.*

We make two observations on the results of this lemma. First, note that in all of the above methods, the matrix whose rows are clustered using k -means algorithm for community detection has only k distinct rows. Moreover, two rows are identical if and only if they are identical in the true community assignment matrix. This ensures that k -means algorithm in each case will correctly cluster the rows. Second, the methods require various conditions on the connectivity matrices $B^{(m)}$'s. In particular, the spectral clustering on mean adjacency matrix requires the aggregate connectivity matrix $\frac{1}{M} \sum_{i=1}^M B^{(m)}$ to be full rank, which is also related to the general issues associated with aggregating a multi-layer graph with diverse layers, explored from an information theoretical point in [23]. Third, the noiseless recovery in co-regularized spectral clustering does not depend on what we choose for γ_m 's. This quite counter-intuitive phenomenon is true because both parts of the objective function, the association cut and the penalty term, are separately maximized by the true communities, with the penalty term achieving its global maximum irrespective of γ_m .

3.2. *Characterizing mis-clustering.* Although Lemma 1 shows that the methods under consideration can perfectly recover community labels from the true population adjacency tensor, in reality we do not observe the true population tensor. Instead we observe a noisy sample version of it. Consequently, community assignment using the methods will lead to some error. For a given benchmark community assignment and an estimated community assignment, we define a mis-clustering rate as the proportion of nodes for which the assignments do not agree. Let \bar{e} denote the vector of true community labels extracted from Z and \hat{e} denote the vector of a candidate assignment. Then we define the mis-clustering rate

$$r = \frac{1}{n} \inf_{\Pi} d_H(\bar{e}, \Pi(\hat{e})),$$

where $\Pi(\cdot)$ is a permutation of the labels, $d_H(\cdot, \cdot)$ is the Hamming distance between two vectors, and \inf denotes the infimum over all permutations in Π .

Note that in each of the methods we consider, we obtain a low rank matrix with orthonormal columns whose rows are then clustered using the k -means algorithm for community detection. Hence we also need to relate this mis-clustering rate with the low rank matrices obtained from the methods. For a method under consideration, let $\hat{U}_{n \times k}$ be the low rank matrix with orthonormal columns it outputs, whose rows can subsequently be clustered to estimate community assignment \hat{e} . Then we have the following relationship,

$$(3.2) \quad r \leq \frac{8n_{\max}}{n} \|\hat{U} - Z(Z^T Z)^{-1/2} O\|_F^2,$$

where O is an arbitrary orthogonal matrix and n_{\max} is the number of nodes in the largest true community [28].

4. Consistency results. In this section we investigate the asymptotic consistency of consensus community detection using the methods outlined in Section 2. The asymptotic setup we consider is as follows. We let both n and M grow, and assume no relationship between their growth rate. However we will be most interested in the case when M grows faster than n . This framework is particularly suitable for consensus community detection in dynamic graphs, where the number of layers represents the number of temporal snapshots available to us and can potentially be exponentially larger compared to the number of nodes. We also let the number of communities k (which is assumed to be known in advance) to grow with both n and M .

Before proceeding with the main results we prove the following theorem with two results on a multi-layer graph with independent edges, the first of which extends the results contained in Chung and Radcliffe [8] to multi-layer graph settings and the second one is a new result using matrix Hoeffding bound [35].

THEOREM 1. *Let \mathcal{G} be a multi-layer graph with each edge being independent of all other edges of all types. Let $\mathbb{A} = \{A^{(1)}, \dots, A^{(M)}\}$ be its adjacency tensor and let $\mathcal{A} = \{\mathcal{A}^1, \dots, \mathcal{A}^{(M)}\}$ denote the expected adjacency tensor. Further, let Δ_m be the maximum expected degree for a node in layer m . Define $\bar{\Delta} = \frac{1}{M} \sum_{m=1}^M \Delta_m$ and $\bar{\Delta}' = \frac{1}{M} \sum_{m=1}^M \Delta_m^2$. Then we have the following results:*

- (i) *For any $\varepsilon > 0$, if $M\bar{\Delta} > \frac{4}{\varepsilon} \log(2n/\varepsilon)$, then with probability at least $1 - \varepsilon$,*

$$\left\| \frac{1}{M} \sum_m (A^{(m)} - \mathcal{A}^{(m)}) \right\|_2 \leq \sqrt{\frac{4\bar{\Delta}' \log(2n/\varepsilon)}{M}}.$$

(ii) If $\frac{1}{M} \sum_m \exp(-\frac{4\Delta_m \log(2Mn^3)}{2\Delta_m + 2\sqrt{4\Delta_m \log(2Mn^3)/3}}) \leq \exp(-\log(2Mn^3))$, then with probability at least $1 - o(1)$ we have

$$\left\| \frac{1}{M} \sum_m (A^{(m)} - \mathcal{A}^{(m)})^2 \right\|_2 \leq (\log n)^{(3+\delta)/2} \frac{\log 2M}{\sqrt{M}} \sqrt{\bar{\Delta}'} + \bar{\Delta}$$

for some $\delta > 0$.

We note that for result (i) on mean adjacency matrix to hold, we only require the average maximum expected degree per layer $\bar{\Delta} \gtrsim \log(2n/\varepsilon)/M$. In comparison, a similar result for adjacency matrix of a single graph in [8] (which (i) extends to multi-layer graph) requires the maximum expected degree $\Delta \gtrsim \log(2n/\varepsilon)$. Hence the result holds for multi-layer graphs where the individual layers are sparser on average. At first glance the density condition on maximum expected degrees for result (ii) looks complicated. However, note that the condition is satisfied, for example, by a choice of $\Delta_m > \frac{4}{9} \log(2n/\varepsilon)$ for each m with $\varepsilon = \frac{1}{Mn^2}$, which is the density condition for a similar result for the adjacency matrix of a single graph in [8]. The condition as in (ii) relaxes the requirement that each layer of the multi-layer network be denser than a threshold, and hence one can have layers which are sparser as long as the layers together satisfy the density condition.

Over the next few sections we use the results of Theorem 1 to prove consistency results for co-regularized spectral clustering, OLMF and spectral clustering in mean adjacency matrix. The common settings under which the results are proved are as follows. Let \mathcal{G} be a multi-layer network with M layers generated from the MLSBM with parameters $[Z, \mathbb{B}]$. Let \mathbb{A} be its adjacency tensor. Let $\lambda^{(m)}$ denote the minimum in absolute value nonzero eigenvalue of the m th layer population adjacency matrix and n_{\max} denote the number of nodes in the largest true community.

4.1. Consistency result for co-regularized spectral clustering.

THEOREM 2. Let $[\hat{U}, \hat{U}^*]$ be the solution that maximizes the co-regularized spectral clustering objective function in (2.3) applied to \mathbb{A} , and r_{coreg} be the fraction of nodes mis-clustered by a k -means procedure applied to \hat{U}^* . Assume $M\bar{\Delta} > \frac{4}{9} \log(4n/\varepsilon)$ and all the $B^{(m)}$'s are of full rank. If we choose γ_m large enough such that $\gamma_m > \frac{\sqrt{M} \|A^{(m)}\|_2^2}{\sqrt{4\bar{\Delta} \log(4n/\varepsilon)}}$ for all m , then for any $\varepsilon > 0$, with probability at least $1 - \varepsilon$,

$$r_{\text{coreg}} \leq \frac{96n_{\max}k}{n \frac{1}{M} \sum_m \frac{(\lambda^{(m)})^2}{\Delta_m}} \sqrt{\frac{\bar{\Delta} \log(4n/\varepsilon)}{M}}.$$

Several discussions on the results of Theorem 2 are in order. First, in the following lemma we replace the deterministic condition on γ_m needed for consistency by a condition that holds only with high probability but involves quantities that depend purely on observed network statistics. Such a condition can then be easily verified in a given network.

LEMMA 2. Assume $M\bar{\Delta} > c \log(2n/\delta)$ where c and δ are positive constants. For each m , if we choose each $\gamma_m > \frac{\sqrt{M} \|A^{(m)}\|_2^2}{\sqrt{\frac{1}{M} \sum_m \|A^{(m)}\|_2 \log(4n/\varepsilon)}}$, then for sufficiently large c , we have with probability at least $1 - \delta$, $\gamma_m > \frac{\sqrt{M} \|A^{(m)}\|_2^2}{\sqrt{4\bar{\Delta} \log(4n/\varepsilon)}}$ for all m .

Although correct recovery under the noiseless case does not require any condition on γ_m 's, the consistency requires γ_m 's to be larger than a certain function of $\|A^{(m)}\|_2$ and M . In the typical case of sparse network layers that we will deal with, $\|A^{(m)}\|_2 \asymp \log n/M$, and then the condition in Lemma 2 reduces to $\gamma_m > O(\|A^{(m)}\|_2)$. Based on this result, in our simulations in Section 5 we choose $\gamma^{(m)}$ to be a constant times $\max \|A^{(m)}\|_2$, identically in each layer.

Second, since it is not immediately clear when the above bound will imply consistent community detection, we make some further assumptions to simplify the bound. In particular we interpret the bound under a multi-layer extension of the four parameter stochastic blockmodel introduced in Rohe, Chatterjee and Yu [28].

Co-regularized spectral clustering under four parameter MLSBM. We define a MLSBM of M layers and n nodes with four parameters $\mathbf{p} = \{p^{(1)}, \dots, p^{(M)}\}$, $\mathbf{q} = \{q^{(1)}, \dots, q^{(M)}\}$, k, s as follows. In layer m , the connection probability within a community is $p^{(m)}$ and between communities is $q^{(m)}$. We assume $p^{(m)} \neq q^{(m)}$ but are of the same asymptotic order with respect to n , for all m . The number of communities is k and all communities are of the same size $s = n/k$. Hence $n_{\max} = s = n/k$. We have the following lemma on the minimum eigenvalues of the population adjacency matrices $\lambda^{(m)}$'s.

LEMMA 3. *For the four parameter MLSBM, $\lambda^{(m)} = s(p^{(m)} - q^{(m)})$, for all $m = 1, \dots, M$.*

Let $a^{(m)} \frac{\Delta_m}{n} = p^{(m)}$ and $b^{(m)} \frac{\Delta_m}{n} = q^{(m)}$. Then $\lambda^{(m)} = \frac{\Delta_m}{k}(a^{(m)} - b^{(m)})$. Consequently, the common asymptotic order of $p^{(m)}$ and $q^{(m)}$ is captured in the $\frac{\Delta_m}{n}$ term and $a^{(m)} \asymp b^{(m)} \asymp 1$. However, note that the difference $a^{(m)} - b^{(m)}$ could still be very small. Define $f(\mathbf{a}, \mathbf{b}) = \frac{1}{M} \sum_m (a^{(m)} - b^{(m)})^2$. Then Theorem 2 implies

$$\begin{aligned} r_{\text{coreg}} &\lesssim \frac{\frac{n}{k}k}{n \frac{1}{M} \sum_m \frac{\Delta_m (a^{(m)} - b^{(m)})^2}{k^2}} \sqrt{\frac{\bar{\Delta} \log(4n/\varepsilon)}{M}} \\ &\asymp \frac{k^2}{\frac{1}{M} \sum_m \Delta_m (a^{(m)} - b^{(m)})^2} \sqrt{\frac{\bar{\Delta} \log(4n/\varepsilon)}{M}}. \end{aligned}$$

At this point we make a further assumption that $\Delta_m \asymp \bar{\Delta}$ for all m . Then we have

$$r_{\text{coreg}} \lesssim \frac{k^2 \sqrt{\log(4n/\varepsilon)}}{\sqrt{M \bar{\Delta} f(\mathbf{a}, \mathbf{b})}},$$

and community detection using this method is consistent as long as $k = o((M \bar{\Delta} / \log(4n/\varepsilon))^{1/4} \sqrt{f(\mathbf{a}, \mathbf{b})})$. We also note that the upper bound on misclustering rate becomes smaller as the number of layers M , the average density of the layers $\bar{\Delta}$ and a measure of community signal $f(\mathbf{a}, \mathbf{b})$ increase.

We consider three growth regimes on the density of the component layers of the multi-layer graph. In the first regime we assume the dense graph setting where the vectors \mathbf{p} and \mathbf{q} do not change with n . This implies that $\bar{\Delta} \asymp n$ and consequently

$$r_{\text{coreg}} \lesssim \frac{k^2}{\sqrt{nM / \log(4n/\varepsilon) f(\mathbf{a}, \mathbf{b})}}.$$

Hence as long as $k = o((nM / \log(4n/\varepsilon))^{1/4} \sqrt{f(\mathbf{a}, \mathbf{b})})$, $r_{\text{coreg}} \rightarrow 0$ with probability at least $1 - \varepsilon$, and we have consistent community detection.

In the second regime, we assume a semi-sparse setting where both $p^{(m)}$ and $q^{(m)}$ are of the order of $\log n/n$ for all m . Then $\bar{\Delta} \asymp \log n$ and we have

$$r_{\text{coreg}} \lesssim \frac{k^2}{f(\mathbf{a}, \mathbf{b})} \sqrt{\frac{\log(2n/\varepsilon)}{M \log n}} \asymp \frac{k^2}{\sqrt{M} f(\mathbf{a}, \mathbf{b})}.$$

This implies that in this setting, as long as $k = o(M^{1/4} \sqrt{f(\mathbf{a}, \mathbf{b})})$, $r_{\text{coreg}} \rightarrow 0$, and we have consistent community detection.

Finally in the sparse ‘‘constant degree’’ regime, where $p^{(m)}$ and $q^{(m)}$ are of the order of $1/n$ for all m , we have $\Delta_m \asymp 1$. Note that the density condition on the layers of the network for Theorem 2(i) to hold is $M \bar{\Delta} \gtrsim \log n$, which can be satisfied even in the constant degree regime if $M \gtrsim \log n$. If this is satisfied, then we have from Theorem 2 that

$$r_{\text{coreg}} \lesssim \frac{k^2}{f(\mathbf{a}, \mathbf{b})} \sqrt{\frac{\log(2n/\varepsilon)}{M}} \asymp \frac{k^2}{\sqrt{M/\log(2n/\varepsilon)} f(\mathbf{a}, \mathbf{b})}.$$

Hence consistent community detection is possible as long as $k = o((M/\log(2n/\varepsilon))^{1/4} \times \sqrt{f(\mathbf{a}, \mathbf{b})})$. Consequently, a large number of very sparse graphs can also lead to consistent community detection, whereas in single layer networks consistent recovery is not possible in the constant degree regime. This is also true for spectral clustering in mean adjacency matrix as we will see in Theorem 4, and is along the lines of the results obtained in Paul and Chen [23].

The next section develops similar results for the OLMF method.

4.2. Consistency result for orthogonal linked matrix factorization.

THEOREM 3. *Let $[\hat{P}, (\hat{\Lambda}^{(1)}, \dots, \hat{\Lambda}^{(M)})]$ be the solution that minimizes the OLMF objective function in (2.1) applied to \mathbb{A} , and r_{LMF} be the fraction of nodes misclustered by a k -means procedure applied to \hat{P} . If the assumption in part (ii) of Theorem 1 holds and at least one of the $B^{(m)}$ ’s is of full rank, then with probability at least $1 - o(1)$,*

$$r_{\text{OLMF}} \leq \frac{48n_{\max} k \bar{\Delta}^{1/2} (\bar{\Delta}^{1/2} + \bar{\Delta}^{1/4} (\log 2M)^{1/2} (\log n)^{2+\varepsilon} / M^{1/4})}{\frac{1}{M} \sum_m (\lambda^{(m)})^2 n}.$$

This bound can also be simplified under the four parameter MLSBM defined earlier. Under the four parameter MLSBM with Δ_m ’s all being of the same order, we have $\Delta_m \asymp \bar{\Delta}$ and $\bar{\Delta}' \asymp \bar{\Delta}^2$. Then the bound in Theorem 3 simplifies to

$$\begin{aligned} r_{\text{LMF}} &\lesssim \frac{\bar{\Delta}^{3/2} + \bar{\Delta}^{3/2} (\log 2M)^{1/2} (\log n)^{2+\varepsilon} / M^{1/4}}{\frac{1}{M} \sum_m \frac{\Delta_m^2 (a^{(m)} - b^{(m)})^2}{k^2}} \\ &\asymp \frac{k^2 (1 + (\log 2M)^{1/2} (\log n)^{2+\varepsilon} / M^{1/4})}{\bar{\Delta}^{1/2} f(\mathbf{a}, \mathbf{b})} \\ &\asymp \max \left\{ \frac{k^2}{\bar{\Delta}^{1/2} f(\mathbf{a}, \mathbf{b})}, \frac{k^2 ((\log 2M)^{1/2} (\log n)^{2+\varepsilon})}{M^{1/4} \bar{\Delta}^{1/2} f(\mathbf{a}, \mathbf{b})} \right\}. \end{aligned}$$

In the dense case where $p^{(m)}$ ’s and $q^{(m)}$ ’s remain constant with increasing n , $\bar{\Delta} \asymp n$ and

$$r_{\text{LMF}} \lesssim \frac{k^2}{\min\{n^{1/2} f(\mathbf{a}, \mathbf{b}), (M/(\log 2M)^2)^{1/4} (n/(\log n)^6)^{1/2} f(\mathbf{a}, \mathbf{b})\}}.$$

Hence consistent estimation is possible as long as k grows slower than the square root of the term in the denominator.

In the sparser case of $\bar{\Delta} \asymp O(\log n)$ we similarly have

$$r_{\text{LMF}} \lesssim \max \left\{ \frac{k^2}{(\log n)^{1/2} f(\mathbf{a}, \mathbf{b})}, \frac{k^2 ((\log 2M)^{1/2} (\log n)^{3/2+\varepsilon})}{M^{1/4} f(\mathbf{a}, \mathbf{b})} \right\},$$

and consistency for the OLMF method provided

$$k = o \left(\min \left\{ (\log n)^{1/4}, \frac{M^{1/8}}{(\log 2M)^{1/4} (\log n)^{3/4+\varepsilon}} \right\} \sqrt{f(\mathbf{a}, \mathbf{b})} \right).$$

4.3. Consistency results for mean adjacency matrix. The final result we prove provides an upper bound on the mis-clustering rate for consensus community detection using the usual single layer spectral clustering on the mean adjacency matrix.

THEOREM 4. Define $\bar{A} = \frac{1}{M} \sum_{m=1}^M A^{(m)}$ and let $\lambda^{\bar{A}}$ denote the minimum in absolute value nonzero eigenvalue of the mean population adjacency matrix $\bar{A} = \frac{1}{M} \sum_{m=1}^M \mathcal{A}^{(m)}$. Let r_{av} be the fraction of nodes misclustered by the spectral clustering algorithm applied to \bar{A} . If $M\bar{\Delta} > \frac{4}{9} \log(2n/\varepsilon)$, and $\bar{B} = \frac{1}{M} \sum_{m=1}^M B^{(m)}$ is of full rank (i.e., rank k), then with probability at least $1 - \varepsilon$,

$$r_{\text{av}} \leq \frac{256n_{\max} k \bar{\Delta} \log(2n/\varepsilon)}{(\lambda^{\bar{A}})^2 n M}.$$

To prove this result, we employ a proof technique using Theorem 1, which is different from Han, Xu and Airolidi [12] and allows us to characterize the dependence of the misclustering rate on the growth rates of various MLSBM parameters. While the concentration result in Frobenius norm of Han, Xu and Airolidi [12] would imply consistent community detection through spectral clustering in mean adjacency matrix for fixed k as long as $n = o(M^{1/2})$, our technique yields a nonasymptotic bound on the mis-clustering rate with direct dependence on the number of communities, sparsity, signal to noise ratio along with n and M . We will next analyze the nonasymptotic bound under a simplified model and different asymptotic growth criteria on the above quantities.

Note the presence of $\lambda^{\bar{A}}$ in the denominator of the bound implies that the bound depends on the eigen-gap of the mean adjacency matrix. To interpret the bound under the four parameter MLSBM, we first prove the following lemma on the eigen-gap $\lambda^{\bar{A}}$.

LEMMA 4. For the four parameter MLSBM, $\lambda^{\bar{A}} = s \frac{1}{M} \sum_m (p^{(m)} - q^{(m)})$.

Similar to previous cases, writing the result in terms of $\bar{\Delta}$, $a^{(m)}$, and $b^{(m)}$ we have $\lambda^{\bar{A}} = \frac{1}{M} \sum_m \frac{\Delta_m}{k} (a^{(m)} - b^{(m)}) \asymp \frac{\bar{\Delta}}{k} \frac{1}{M} \sum_m (a^{(m)} - b^{(m)})$. Define $g(\mathbf{a}, \mathbf{b}) = (\frac{1}{M} \sum_m (a^{(m)} - b^{(m)}))^2$. Then from Theorem 4 we have with probability at least $1 - \varepsilon$,

$$r_{\text{av}} \lesssim \frac{\frac{n}{k} k \bar{\Delta} \log(2n/\varepsilon)}{(\frac{\bar{\Delta}}{k})^2 g(\mathbf{a}, \mathbf{b}) n M} \asymp \frac{k^2}{M \bar{\Delta} g(\mathbf{a}, \mathbf{b}) / \log(2n/\varepsilon)}.$$

This implies that $r_{\text{av}} \rightarrow 0$ as long as $k = o(\sqrt{M \bar{\Delta} g(\mathbf{a}, \mathbf{b}) / \log(2n/\varepsilon)})$, and we have consistent community detection. We note that $g(\mathbf{a}, \mathbf{b})$ and $\bar{\Delta}$ are averages over the layers of the corresponding quantities for single layer case. The above result then implies that with increasing number of layers M , the upper bound on the misclustering rate gets smaller by a factor of \sqrt{M} as compared to applying spectral clustering on any one of the layers separately as shown in Qin and Rohe [27] and Lei and Rinaldo [16] (the $\log n$ term does not

appear in Lei and Rinaldo [16] due to tighter bound on $\|A - \mathcal{A}\|_2$). We also note that the denominator in the rate for r_{av} contains the term $g(\mathbf{a}, \mathbf{b}) = (\frac{1}{M} \sum_m (a^{(m)} - b^{(m)}))^2$ instead of $f(\mathbf{a}, \mathbf{b}) = \frac{1}{M} \sum_m (a^{(m)} - b^{(m)})^2$, which appeared earlier in the rates of OLMF and co-regularized spectral clustering. From Jensen's inequality,

$$g(\mathbf{a}, \mathbf{b}) = \left(\frac{1}{M} \sum_m (a^{(m)} - b^{(m)}) \right)^2 \leq \frac{1}{M} \sum_m (a^{(m)} - b^{(m)})^2 = f(\mathbf{a}, \mathbf{b}),$$

with equality holding if and only if all the $(a^{(m)} - b^{(m)})$'s are equal. Hence equality holds if the layers are of similar signal quality, and otherwise $f(\mathbf{a}, \mathbf{b})$ is larger than $g(\mathbf{a}, \mathbf{b})$. Hence the goodness of the rate for spectral clustering in mean adjacency matrix depends on if the aggregate of the layers has good signal quality or not. In the situation where some of the layers in the multi-layer network contain heterophilic communities while the others contain homophilic communities, then $a^{(m)} - b^{(m)}$ is negative in some layers and positive in other layers. In that case $\lambda^{\bar{A}}$ could be very small and performance guarantee on spectral clustering of mean adjacency matrix become poor. These conclusions are in line with previous conclusions from minimax rates and phase transitions of consistency thresholds in [23].

In the dense regime where the vectors \mathbf{p} and \mathbf{q} do not change with n , we have the mis-clustering rate in spectral clustering in mean adjacency matrix is bounded by $r_{\text{av}} \lesssim \frac{k^2}{nMg(\mathbf{a}, \mathbf{b})/\log(2n/\varepsilon)}$. In the semi-sparse regime where both $p^{(m)}$ and $q^{(m)}$ are of the order of $\frac{\log n}{n}$ for all m , we have $\bar{\Delta} \asymp \log n$ and $r_{\text{av}} \lesssim \frac{k^2}{Mg(\mathbf{a}, \mathbf{b})}$. Finally, in the sparse constant degree regime where both $p^{(m)}$ and $q^{(m)}$ are of the order of $1/n$ for all m , we have $r_{\text{av}} \lesssim \frac{k^2}{Mg(\mathbf{a}, \mathbf{b})/\log(2n/\varepsilon)}$.

5. Simulation studies. In this section, we numerically compare the performance of the following methods through a principled simulation study: spectral clustering on mean adjacency matrix (Mean adj.), OLMF, co-regularized spectral clustering (Coreg spec), spectral clustering on aggregate spectral kernel (SpecK) and the module allegiance matrix (Module alleg.). Since the computational algorithms for both OLMF and Coreg Spec are only expected to reach a local optimum, it is important to supply good initial conditions to them and also take the best solution based on multiple initial conditions. We initialize the OLMF algorithm with P being the community assignment matrix from a randomly chosen layer and $\Lambda^{(m)}$ being the matrix containing the top k eigenvalues of $A^{(m)}$ in the diagonal. For the co-regularized spectral clustering algorithm we choose $\gamma^{(m)}$ as $4 \max \|A^{(m)}\|_2$ for all m , since the theoretical results have indicated that $\gamma^{(m)}$ should be larger than $\|A^{(m)}\|_2$ for each m .

For the first three simulations, we simulate networks from the MLSBM with the number of nodes $n = 600$ and the number of layers $M = 5$, under three different scenarios on the connection probability matrices of different layers. The performances of the methods are evaluated with increasing average degree of the multi-layer network since we would expect any reasonable method to perform better as the network gets denser. The number of communities is fixed at 3 and we assume it to be known in advance. The fourth simulation involves generating networks from MLSBM with varying number of layers and testing the performance of the methods with increasing number of layers. The fifth and final simulation considers the scenario where the multi-layer network contains layers with both heterophilic and homophilic communities.

The evaluation criterion is the normalized mutual information (NMI) with the ground truth community assignments which generate the network. The NMI is an information theoretic measure of similarity between two vectors of community assignments, with 1 indicating a perfect match and 0 indicating the vectors are random with respect to each other. The first three experiments are replicated 40 times while the last two experiments are repeated 100 times, and the average performance across the repetitions is reported.

The data are generated according to MLSBM as defined in (3.1) in the following fashion. The community vector Z_i for each node i is generated according to a multinomial distribution with equal probability of being in any of the 3 communities. The block model matrices $B^{(m)}$'s in different layers are generated by the following scheme. Let δ be the vector of k diagonal elements and ε be the vector of $k^2 - k$ off-diagonal elements. We generate half of the elements of the ε vector from a uniform distribution $U(a, b)$ within a short range $[a, b]$ and the other half is a replication of the first half such that the matrix is symmetric. The elements of δ are generated from $U(\rho a, \rho b)$, where ρ is the parameter that controls the signal to noise ratio (SNR). We call an SNR of 2-3 as “strong” signal and an SNR which is only slightly greater than 1 as “weak” signal.

Strong community signals. In the first simulation from MLSBM, we make all the layers contain generally strong signals, but the exact SNR is randomly varied slightly so as to have some variations in signal quality across the layers. The performance of various methods under consideration is presented in Figure 1(a). Note that the layers are sparse at an average initially which is evident from the low average degree per layer: an average degree of 6 in a layer of 600 nodes, which is about 1% degree density. The layers then become denser gradually and reach about 2.5% degree density per layer. The performance of all the methods generally increases with increasing average degree. We note that spectral clustering on mean adjacency matrix, OLMF and co-regularized spectral clustering perform similarly throughout the range of the simulation. The aggregate of spectral kernel and module allegiance matrix method substantially underperform, especially in sparse multi-layer networks.

Mixed and ambiguous community signals. In this simulation, the component layers are mixed in community signal quality in the following manner. We have three layers with strong community signals and two layers where the community structure is ambiguous or almost nonexistent due to weak signal to noise ratio. This scenario is very useful to test the robustness of methods against possible variation or absence of community patterns in some of the layers. The results are presented in Figure 1(b). The OLMF method performs the best over the entire range of values of average degree, followed by co-regularized spectral clustering and spectral clustering of mean adjacency matrix. The aggregate spectral kernel and module allegiance matrix methods once again perform poorly when the average density in the layers is low, but recover subsequently as the layers become denser. The spectral kernel method performs better than the module allegiance matrix method in both the strong signals and mixed signals scenarios.

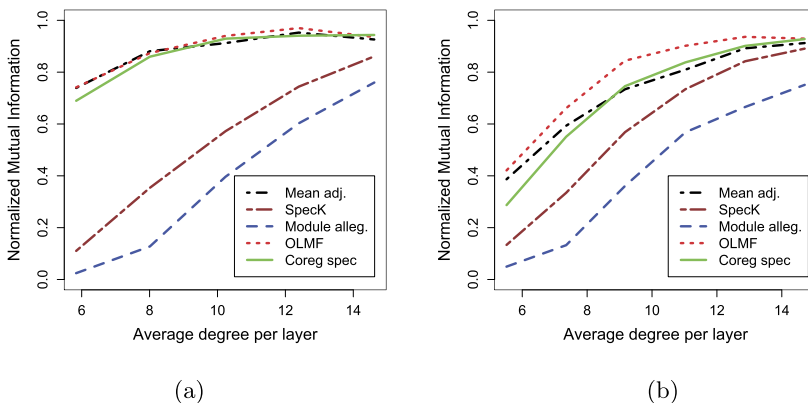


FIG. 1. Performance of various methods with increasing average degree of nodes for data generated from MLSBM with 600 nodes, 5 layers and 3 communities. (a) All layers have strong signals with some variations; (b) the layers are mixed in terms of signal quality.

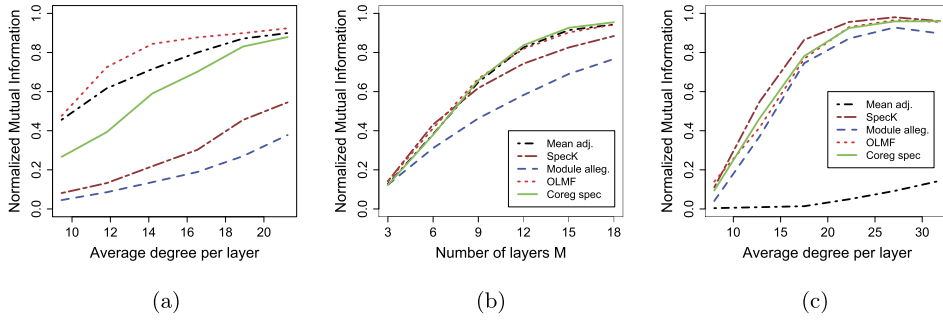


FIG. 2. Performance of various methods with (a) increasing average degree of nodes for data generated from MLSBM with 600 nodes, 5 layers and 3 communities, (b) increasing number of layers with 300 nodes and 6 communities, (c) increasing average degree of nodes with 600 nodes, 5 layers and 3 communities, where 3 layers contain homophilic communities and the other 2 contain heterophilic communities. The labels in (b) and (c) are shared for all figures.

Complementary information. The third scenario considers the so-called “complementary” principle of multiple views in multi-view learning [17]. In our case, this is equivalent to the following: none of the layers alone is sufficient to describe the community structure properly, but the layers can complement each other and together describe the community patterns. For our simulation, we generate data from MLSBM with 600 nodes, 5 layers and 3 communities with the following setting. In each of the first 3 layers, two of the communities are difficult to distinguish from noise while the third community has an SNR of 3. The fourth layer has two of the communities with high SNR and the fifth layer has the same two communities with low SNR. The performance of the competing methods are presented in Figure 2(a). We observe that both aggregate of spectral kernel and module allegiance matrix method perform poorly in this scenario as compared to the intermediate fusion methods as well as spectral clustering of mean adjacency matrix. This is expected since none of the layers alone contain complete information about the community structure and hence the eigenspaces computed separately are not very informative of the community structure. Consequently sharing information while computing individual eigenspaces as well as a consensus eigenspace is beneficial as opposed to a late fusion of individual eigenspaces. In addition, the OLMF method appears to have a clear advantage in this scenario over both co-regularized spectral clustering and spectral clustering of the mean adjacency matrix.

Increasing number of layers. This simulation setup tests the abilities of the methods to recover the community structure with a small fixed number of nodes, but increasing number of layers (and consequently more data). However, as is the case with many real world multi-layer networks, not all of the layers are strongly informative of the community structure. We fix n at 300, k at 6 and increase M from 3 to 18 in steps of 3. At every step, we add 3 layers to the multi-layer network, two of which have weak signal quality, while the third one has a strong signal. The performance of the competing methods in this simulation with 100 repetitions is depicted in Figure 2(b). We observe that the accuracy of consensus community detection in all the methods generally increases with increasing number of layers. As with the previous scenarios, we observe that OLMF, co-regularized spectral clustering, and spectral clustering of the mean adjacency matrix have more improvement in performance as compared to aggregate of spectral kernel and module allegiance matrix methods.

Layers with heterophilic communities. Finally, we consider the scenario where some layers contain homophilic (assortative) communities while others contain heterophilic (disassortative) communities. The layers with heterophilic communities have less density within the

blocks as opposed to inter-block densities. Such interactions with disassortative communities are commonly encountered in food webs and social networks. From our theoretical analysis we expect the spectral clustering of mean adjacency matrix to perform poorly in this setting. Intuitively, the mean adjacency matrix has strong inter-community edge density (in addition to strong intra-community edge density) due to the layers with heterophilic communities and consequently, the community structure is ambiguous and difficult to detect. However, the community information is separately available in all the layers irrespective of whether the communities in that layer are homophilic or heterophilic. Then one would hope, perhaps, a different way of combining information from layers will yield the community structure correctly.

Since the inter-block connection probabilities are higher than intra-block connection probabilities in the layers with heterophilic communities, the eigenvalues corresponding to the eigenvectors that contain the clustering information are all negative. Hence we need to modify some of the methods slightly for this scenario. For aggregate spectral kernel and module allegiance methods, we choose the eigenvectors corresponding to the top k eigenvalues in *absolute value* to form the $\hat{U}^{(m)}$ matrix in each layer. For co-regularized spectral clustering, we update the $\hat{U}^{(m)}$ matrix during the alternating eigen-decomposition by selecting the vectors corresponding to the top k eigenvalues in *absolute value*. The mean adjacency matrix and OLMF methods do not require any change to be made, however, we make the following optional modifications. For mean adjacency matrix during the eigen-decomposition, we choose eigenvectors corresponding to the top k eigenvalues in *absolute value* of the mean adjacency matrix, while for OLMF we only change the initialization of $\Lambda^{(m)}$ matrices to include the k largest eigenvalues in *absolute value* as its diagonal.

We fix n at 600, k at 3, M at 5, and increase the average degree per layer from 8 to 32 (from about 1% to 4.5% in degree density). We make 3 of the 5 layers contain homophilic communities by setting the ρ parameter (SNR) at 3, while we make the other two layers contain heterophilic communities by setting $\rho = 1/3$ so that the elements of δ are smaller than that of ε . The results are presented in Figure 2(c).

As expected from our theoretical results, we observe that the performance of spectral clustering in mean adjacency matrix completely breaks down and is substantially worse than the competing methods in this scenario. The other four methods behave similarly and the accuracy of community detection steadily increases with increasing degree density. This indicates that all of those four methods are capable of extracting information relevant to the community structure from layers with both homophilic and heterophilic communities and combining them without nullifying the information.

We also note that the aggregate spectral kernel method performs slightly better compared to the two intermediate fusion methods throughout the range of the simulation. We think this is because of the following reason. The relatively higher average degree per layer in the simulated networks compared to, for example, that in Figure 1(a), means the recovery of the true eigenspaces (which contain the information on community assignments) by spectral methods in each of the layers becomes increasingly accurate irrespective of whether the communities are homophilic or heterophilic [16, 27, 28]. This leads to better performance of the methods that purely rely on combinations of those independently obtained eigenspaces. Hence the aggregate spectral kernel itself becomes more effective than the intermediate fusion methods. The intermediate fusion methods on the other hand, shares information while computing the eigenspaces and the consensus eigenspace appears to underperform in the presence of layers with both homophilic and heterophilic communities.

Discussion on the simulation results. Our simulations clearly show that in sparse networks the intermediate fusion of information based methods, OLMF and co-regularized

spectral clustering, perform better than late fusion methods, aggregate of spectral kernel and module allegiance matrix method. We think sharing information across layers while computing individual layer wise spectral embeddings increases the accuracy in each of them, and hence the centroid is a more effective combination than aggregate spectral kernel or module allegiance matrix type of combination. The spectral clustering of mean adjacency matrix performs well in our simulations except the last scenario where the multi-layer network contains layers with both homophilic and heterophilic communities, in which case its performance is extremely poor. We also observe in our simulations that aggregate spectral kernel performs better than module allegiance matrix. We think the performance in module allegiance suffers because of additional noise introduced in discrete community assignments. Overall, we think the intermediate fusion methods outperform or remain competitive to the baseline methods of aggregate spectral kernel and spectral clustering in mean adjacency matrix in a wide variety of scenarios.

6. Conclusions and discussions. In this paper we have analyzed a number of spectral and matrix factorization based techniques for multi-view clustering in terms of their asymptotic consistency properties for community detection in multi-layer networks generated from the MLSBM. We have considered a high dimensional asymptotic framework where both the number of layers (M) and the number of nodes (n) of the multi-layer graph grow. The main technical contribution of the article is to prove nonasymptotic error bounds for community detection using the global optimal solutions of both co-regularized spectral clustering and OLMF, and spectral clustering of the mean adjacency matrix in terms of model parameters of the MLSBM. As an intermediate step we have proved two concentration inequalities on two functions of adjacency matrices of a multi-layer network. We have further shown that the above-mentioned methods enjoy consistency guarantees under some conditions on the number of communities k , the maximum expected degrees of the layers Δ_m 's and the signal to noise ratios of the layers.

We have also compared five methods in terms of finite sample performance under data generated from the MLSBM through a simulation study. We found both the co-regularized spectral clustering and OLMF to be robust under varied scenarios. We also note from the simulations that widely popular methods where each layer is dealt separately and the results are fused at a later state, such as aggregating spectral kernels or module allegiance matrix, do not perform well in sparse networks when the individual layers do not contain sufficient information to recover the community structure efficiently. However, the OLMF and co-regularized spectral clustering perform well in those scenarios. We hypothesize that this is due to sharing information across layers while computing the community structure solution at each layer.

Global optimizers. Throughout the paper we have studied the properties of the global optimal solutions of the optimization problems under consideration. However, in the absence of computational methods guaranteed to achieve the global optimal solutions, it is not known whether such global optimum can ever be achieved under any circumstances. Indeed, the algorithms we have used to compute the solutions in our simulation studies are approximate algorithms that can at best reach a local optimum. To the best of our knowledge, no computationally feasible algorithm exists that can compute the global optima of the intermediate fusion objective functions with guarantees. We view the results obtained in this article as only a first step in the direction of understanding the behavior of multi-view learning methods in the context of community detection in multi-layer networks. In the future, we hope to investigate possibilities of obtaining algorithms with global optimum guarantees and extend the results obtained here to such cases.

To assure ourselves that the solutions computed by the algorithms used here are not completely away from the global solutions, we conducted a simulation study. Although we do not

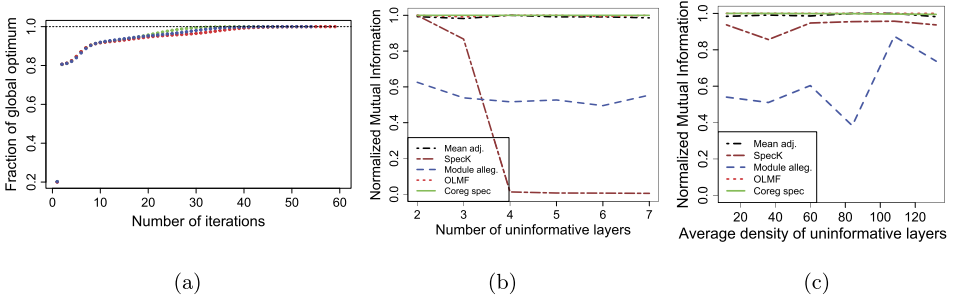


FIG. 3. (a) Objective function value as the fraction of the global optimum with increasing number of iterations for co-regularized spectral clustering. (b)–(c) Recovery of the true community structure from population adjacency tensor in the presence of uninformative layers: (b) The number of uninformative layers is increased keeping their densities fixed; (c) the density of the 3 uninformative layers are gradually increased. In both (b) and (c) there are 3 layers informative of the community structure.

know the true global optimum of the two optimization problems under study in real data situations, or even in simulated sample networks, we know the solutions for them theoretically when they are applied to the population adjacency tensor. Hence we study the objective function values at convergence for the two methods applied to population adjacency tensors with increasing degree density. Our first simulation verified that using the spectral clustering of mean adjacency matrix as initial solution, the OLMF objective function is within 10^{-18} of 0 and the co-regularized spectral clustering objective function is within 10^{-12} of the true maximum across the range of the simulation. Our second simulation whose result is presented in Figure 3(a), verified that starting from a random initial solution the objective function value for co-regularized spectral clustering goes close to the true maximum with increasing number of iterations eventually being equal to the true maximum.

Consensus community detection in the presence of noise. In this paper we have assumed the presence of an underlying consensus community assignment for the multi-layer network and focused on the problem of detecting such a structure. Indeed, Lemma 1, which shows that the methods under study can correctly recover the community structure from the population adjacency tensor (i.e., without sampling noise), is the crucial backbone of the paper on which all results are based on. Here we analyze a scenario where the community structure is truly present in some layers (perhaps a majority), while it is either absent or is different in the other layers, and the task is to detect the community structure present in the majority of the layers. Such scenarios have been previously considered in [30, 37]. Since this scenario is not the focus of the paper, we will primarily analyze whether the methods are capable of recovering the community structure from the population adjacency tensor of such a multi-layer SBM. Let M_1 layers contain the community structure of interest Z_1 and M_2 layers contain a different community structure Z_2 with $M_1 > M_2$. The community structure Z_2 could put all vertices in the same community (i.e., simply an Erdos–Renyi graph) or could be a community structure that is different from Z_1 . We concentrate on the former case, where Z_2 does not define any community structure. We will study under what conditions the methods analyzed in this paper will be able to detect the main community structure of interest Z_1 . Even for the M_2 layers with no community structure, we can write the population adjacency matrices as still being created by $Z_1 B Z_1^T$ but with B having identical values in each entry and consequently of rank 1.

The mean adjacency matrix can then be written as:

$$\bar{A} = \frac{1}{M} \sum_{m=1}^M A^{(m)} = Z_1 \left(\frac{1}{M} \left(\sum_{m=1}^{M_1} B^{(m)} + M_2 B' \right) \right) Z_1^T.$$

Nevertheless, we would require a similar condition as before, namely $\bar{B} = \frac{1}{M}(\sum_{m=1}^{M_1} B^{(m)} + M_2 B')$ will be full rank (i.e., rank k). Under this condition the spectral clustering procedure in mean adjacency matrix can extract the true community structure from the population adjacency matrix.

An (simplified) extension of the four parameter MLSBM can be defined for this case as follows: let a and b be diagonal and off-diagonal elements of $B^{(m)}$'s in the first M_1 layers with $a > b$ and c 's are the elements of B' . From Lemma 4, the smallest non-zero eigenvalue of $\bar{B} = \frac{1}{M}(M_1(a-b)I_k + (M_1b + M_2c)1_k 1_k^T)$ is $\frac{M_1}{M}(a-b)$. Since $a > b$, the $k \times k$ matrix \bar{B} has k non-zero eigenvalues and consequently is of full rank. Then a spectral clustering algorithm can recover the true community structure from the mean population adjacency matrix.

However, the noise plays a big role when we look at the sample adjacency matrices. Using the simplified result of Theorem 4 under the four parameter MLSBM, we have consistency as long as $M\bar{\Delta}g(a, b)/\log n = \omega(1)$ for a constant number of communities k . Now if we assume all layers are of similar density, then $\bar{\Delta}$ does not change by adding Erdos–Renyi graphs. However, $g(a, b) = \frac{M_1^2(a-b)^2}{(M_1+M_2)^2}$ decreases as we increase M_2 . In the case of $\bar{\Delta} \asymp \log n$, we have $M_1(a-b)^2 = \omega(M)$ as a sufficient condition for consistency. This is in contrast to the usual requirement of $(a-b)^2 = \omega(1)$. On the other hand if we assume M_1 and M_2 are fixed, but the density of the Erdos–Renyi layers gradually increases, then with the addition of such dense but uninformative layers, $\bar{\Delta}$ increases, while $g(a, b)$ remains the same. This increases the upper bound and the method does not lead to consistent community detection anymore.

Turning our attention to spectral kernel method, we will have $Z_1^T(Z_1^T Z_1)^{-1}Z_1^T$ as the spectral kernels from the M_1 layers. The matrices $\mathcal{A}^{(m)}$ are of rank 1 for the other M_2 layers and let $J^{(m)}J^{(m)T}$, where $J^{(m)}$ are matrices with orthonormal columns, be the kernels for each such $m \in \{1, \dots, M_2\}$. Then the aggregate spectral kernel is $K = \frac{1}{M}(M_1 Z_1^T(Z_1^T Z_1)^{-1}Z_1^T + \sum_{m=1}^{M_2} J^{(m)}J^{(m)T})$. Since $J^{(m)}J^{(m)T}$ is not associated with Z , the matrix Z cannot be extracted exactly from the kernel perfectly when $M_2 > 0$. For small M_2 we can still recover a subspace close to the subspace spanned by Z and the error will be governed by the Davis–Kahan theorem [31]. However, with M_2 increasing eventually we will not be able to recover the subspace at all.

For OLMF, $\bar{P} = ZQ^{-1/2}$, $\bar{\Lambda}^{(m)} = Q^{1/2}B^{(m)}Q^{1/2}$, $m = 1, \dots, M$, is still a solution of the optimization problem (2.2), and its uniqueness is ensured as long as at least one of the $B^{(m)}$'s in the first M_1 layers is of full rank.

We verify these observations on the population adjacency tensor in a simulation study, whose results are presented in Figure 3. In the first simulation (Figure 3(b)) we increase the number of layers uninformative of the community structure from 2 to 7 while keeping the number of informative layers fixed at 3. Spectral clustering on mean adjacency matrix continues to be able to recover the community structure perfectly. This behavior is replicated by both OLMF and co-regularized spectral clustering methods. However, aggregate of spectral kernel and module allegiance matrix approaches are not successful in recovering the correct community structure in the presence of uninformative layers. In the second simulation (Figure 3(c)) we keep the number of informative and uninformative layers both fixed at 3 each, and vary the density of the uninformative layers while keeping the density of the informative layers fixed. As our theoretical analysis indicates, although the aggregate of spectral kernel approaches fail to recover the community structure correctly, its performance is unaffected by increasing density of the uninformative layers. Spectral clustering on mean adjacency matrix, OLMF and co-regularized spectral clustering methods continue to be able to recover the correct community structure and are unaffected by increasing density of the uninformative layers.

Acknowledgments. This work was supported in part by NSF Grant DMS-1406455.

SUPPLEMENTARY MATERIAL

Supplement to “Spectral and matrix factorization methods for consistent community detection in multi-layer networks” (DOI: [10.1214/18-AOS1800SUPP](https://doi.org/10.1214/18-AOS1800SUPP); .pdf). The supplementary file contains proofs of equivalence between problems (2.1) and (2.2), Proposition 1, Lemmas 1, 2, 3 and 4, and Theorems 1, 2, 3, and 4.

REFERENCES

- [1] BARBILLON, P., DONNET, S., LAZEGA, E. and BAR-HEN, A. (2017). Stochastic block models for multiplex networks: An application to a multilevel network of researchers. *J. Roy. Statist. Soc. Ser. A* **180** 295–314. MR3600512 <https://doi.org/10.1111/rssa.12193>
- [2] BASSETT, D. S., PORTER, M. A., WYMBS, N. F., GRAFTON, S. T., CARLSON, J. M. and MUCHA, P. J. (2013). Robust detection of dynamic community structure in networks. *Chaos* **23** 013142, 16. MR3389633 <https://doi.org/10.1063/1.4790830>
- [3] BASSETT, D. S., WYMBS, N. F., PORTER, M. A., MUCHA, P. J., CARLSON, J. M. and GRAFTON, S. T. (2011). Dynamic reconfiguration of human brain networks during learning. *Proc. Natl. Acad. Sci. USA* **108** 7641–7646.
- [4] BOCCALETTI, S., BIANCONI, G., CRIADO, R., DEL GENIO, C. I., GÓMEZ-GARDEÑES, J., ROMANCE, M., SENDIÑA-NADAL, I., WANG, Z. and ZANIN, M. (2014). The structure and dynamics of multilayer networks. *Phys. Rep.* **544** 1–122. MR3270140 <https://doi.org/10.1016/j.physrep.2014.07.001>
- [5] BRAUN, U., SCHÄFER, A., WALTER, H., ERK, S., ROMANCZUK-SEIFERTH, N., HADDAD, L., SCHWEIGER, J. I., GRIMM, O., HEINZ, A. et al. (2015). Dynamic reconfiguration of frontal brain networks during executive cognition in humans. *Proc. Natl. Acad. Sci. USA* **112** 11678–11683.
- [6] CHEN, P.-Y. and HERO III, A. O. (2016). Multilayer spectral graph clustering via convex layer aggregation. Preprint. Available at [arXiv:1609.07200](https://arxiv.org/abs/1609.07200).
- [7] CHI, Y., SONG, X., ZHOU, D., HINO, K. and TSENG, B. L. (2007). Evolutionary spectral clustering by incorporating temporal smoothness. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 153–162.
- [8] CHUNG, F. and RADCLIFFE, M. (2011). On the spectra of general random graphs. *Electron. J. Combin.* **18** Paper 215, 14. MR2853072
- [9] DONG, X., FROSSARD, P., VANDERGHEYNST, P. and NEFEDOV, N. (2012). Clustering with multi-layer graphs: A spectral perspective. *IEEE Trans. Signal Process.* **60** 5820–5831. MR2990287 <https://doi.org/10.1109/TSP.2012.2212886>
- [10] GHASEMIAN, A., ZHANG, P., CLAUSET, A., MOORE, C. and PEEL, L. (2016). Detectability thresholds and optimal algorithms for community structure in dynamic networks. *Phys. Rev. X* **6** 031005.
- [11] GREENE, D. and CUNNINGHAM, P. (2013). Producing a unified graph representation from multiple social network views. *ACM Web Sci.* **2** 129–233.
- [12] HAN, Q., XU, K. S. and AIROLDI, E. M. (2015). Consistent estimation of dynamic and multi-layer block models. In *Proceedings of the 32nd International Conference on Machine Learning* 1511–1520.
- [13] HOLLAND, P. W., LASKEY, K. B. and LEINHARDT, S. (1983). Stochastic blockmodels: First steps. *Soc. Netw.* **5** 109–137. MR0718088 [https://doi.org/10.1016/0378-8733\(83\)90021-7](https://doi.org/10.1016/0378-8733(83)90021-7)
- [14] KIVELÄ, M., ARENAS, A., BARTHELEMY, M., GLEESON, J. P., MORENO, Y. and PORTER, M. A. (2014). Multilayer networks. *J. Complex Netw.* **2** 203–271.
- [15] KUMAR, A., RAI, P. and DAUME, H. (2011). Co-regularized multi-view spectral clustering. In *Advances in Neural Information Processing Systems* 1413–1421.
- [16] LEI, J. and RINALDO, A. (2015). Consistency of spectral clustering in stochastic block models. *Ann. Statist.* **43** 215–237. MR3285605 <https://doi.org/10.1214/14-AOS1274>
- [17] LIU, J., WANG, C., GAO, J. and HAN, J. (2013). Multi-view clustering via joint nonnegative matrix factorization. In *Proceedings of the 2013 SIAM International Conference on Data Mining* 252–260.
- [18] LONG, B., ZHANG, Z. M., WU, X. and YU, P. S. (2006). Spectral clustering for multi-type relational data. In *Proceedings of the 23rd International Conference on Machine Learning* 585–592.
- [19] MUCHA, P. J., RICHARDSON, T., MACON, K., PORTER, M. A. and ONNELA, J.-P. (2010). Community structure in time-dependent, multiscale, and multiplex networks. *Science* **328** 876–878. MR2662590 <https://doi.org/10.1126/science.1184819>

- [20] NEWMAN, M. E. J. and GIRVAN, M. (2004). Finding and evaluating community structure in networks. *Phys. Rev. E* **69** 026113.
- [21] NG, A. Y., JORDAN, M. I. and WEISS, Y. (2002). On spectral clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems* 849–856.
- [22] NICKEL, M., TRESP, V. and KRIEGEL, H. P. (2011). A three-way model for collective learning on multi-relational data. In *Proceedings of 28th International Conference on Machine Learning* 809–816.
- [23] PAUL, S. and CHEN, Y. (2016). Consistent community detection in multi-relational data through restricted multi-layer stochastic blockmodel. *Electron. J. Stat.* **10** 3807–3870. MR3579677 <https://doi.org/10.1214/16-EJS1211>
- [24] PAUL, S. and CHEN, Y. (2016). Null models and modularity based community detection in multi-layer networks. Preprint. Available at [arXiv:1608.00623](https://arxiv.org/abs/1608.00623).
- [25] PAUL, S. and CHEN, Y. (2020). Supplement to “Spectral and matrix factorization methods for consistent community detection in multi-layer networks.” <https://doi.org/10.1214/18-AOS1800SUPP>.
- [26] PEIXOTO, T. P. (2015). Inferring the mesoscale structure of layered, edge-valued, and time-varying networks. *Phys. Rev. E* **92** 042807.
- [27] QIN, T. and ROHE, K. (2013). Regularized spectral clustering under the degree-corrected stochastic blockmodel. In *Advances in Neural Information Processing Systems* 3120–3128.
- [28] ROHE, K., CHATTERJEE, S. and YU, B. (2011). Spectral clustering and the high-dimensional stochastic blockmodel. *Ann. Statist.* **39** 1878–1915. MR2893856 <https://doi.org/10.1214/11-AOS887>
- [29] SINGH, A. P. and GORDON, G. J. (2008). Relational learning via collective matrix factorization. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 650–658.
- [30] STANLEY, N., SHAI, S., TAYLOR, D. and MUCHA, P. J. (2016). Clustering network layers with the strata multilayer stochastic block model. *IEEE Trans. Netw. Sci. Eng.* **3** 95–105. MR3515211 <https://doi.org/10.1109/TNSE.2016.2537545>
- [31] STEWART, G. W. and SUN, J. G. (1990). *Matrix Perturbation Theory*. Computer Science and Scientific Computing. Academic Press, Boston, MA. MR1061154
- [32] SUN, S. (2013). A survey of multi-view machine learning. *Neural Comput. Appl.* **23** 2031–2038.
- [33] TANG, L., WANG, X. and LIU, H. (2012). Community detection via heterogeneous interaction analysis. *Data Min. Knowl. Discov.* **25** 1–33. MR2911723 <https://doi.org/10.1007/s10618-011-0231-0>
- [34] TANG, W., LU, Z. and DHILLON, I. S. (2009). Clustering with multiple graphs. In *Proceedings of the Ninth IEEE International Conference on Data Mining* 1016–1021.
- [35] TROPP, J. A. (2012). User-friendly tail bounds for sums of random matrices. *Found. Comput. Math.* **12** 389–434. MR2946459 <https://doi.org/10.1007/s10208-011-9099-z>
- [36] VALLES-CATALA, T., MASSUCCI, F. A., GUIMERA, R. and SALES-PARDO, M. (2016). Multilayer stochastic block models reveal the multilayer structure of complex networks. *Phys. Rev. X* **6** 011036.
- [37] WILSON, J. D., PALOWITCH, J., BHAMIDI, S. and NOBEL, A. B. (2017). Community extraction in multilayer networks with heterogeneous community structure. *J. Mach. Learn. Res.* **18** Paper No. 149, 49. MR3763783
- [38] XU, C., TAO, D. and XU, C. (2013). A survey on multi-view learning. Preprint. Available at [arXiv:1304.5634](https://arxiv.org/abs/1304.5634).
- [39] ZHOU, D. and BURGESS, C. J. (2007). Spectral clustering and transductive learning with multiple views. In *Proceedings of the 24th International Conference on Machine Learning* 1159–1166.
- [40] ŽITNIK, M. and ZUPAN, B. (2015). Data fusion by matrix factorization. *IEEE Trans. Pattern Anal. Mach. Intell.* **37** 41–53. <https://doi.org/10.1109/TPAMI.2014.2343973>