

Clustering Network Layers with the Strata Multilayer Stochastic Block Model

Natalie Stanley, Saray Shai, Dane Taylor, and Peter J. Mucha

Abstract—Multilayer networks are a useful data structure for simultaneously capturing multiple types of relationships between a set of nodes. In such networks, each relational definition gives rise to a layer. While each layer provides its own set of information, community structure across layers can be collectively utilized to discover and quantify underlying relational patterns between nodes. To concisely extract information from a multilayer network, we propose to identify and combine sets of layers with meaningful similarities in community structure. In this paper, we describe the “strata multilayer stochastic block model” (sMLSBM), a probabilistic model for multilayer community structure. The central extension of the model is that there exist groups of layers, called “strata”, which are defined such that all layers in a given stratum have community structure described by a common stochastic block model (SBM). That is, layers in a stratum exhibit similar node-to-community assignments and SBM probability parameters. Fitting the sMLSBM to a multilayer network provides a joint clustering that yields node-to-community and layer-to-stratum assignments, which cooperatively aid one another during inference. We describe an algorithm for separating layers into their appropriate strata and an inference technique for estimating the SBM parameters for each stratum. We demonstrate our method using synthetic networks and a multilayer network inferred from data collected in the Human Microbiome Project.

Index Terms—Stochastic block models, clustering, multilayer networks, strata, probabilistic models

1 INTRODUCTION

MODELING relational information between a set of entities can often be successfully achieved through a network representation. Here, entities correspond to nodes and edges reflect some connection between them. In many applications, there are multiple ways to define an edge that can be collectively analyzed for a more thorough understanding of the data. Multilayer networks provide a framework to do this, in that each relational definition leads to a new layer in the network [1], [2], [3]. Such data and corresponding networks have shown to be useful in many contexts, such as, in the comparison of genetic and protein-protein interactions in a cell [4], in understanding underlying relationships and community structure across social networks [5], and in the analysis of temporal networks [6]. Furthermore, recent advances in the mathematical foundations for multilayer networks have made analysis of these types of data more feasible. In particular, [3] has introduced a mathematical formalism with tensors. Doing so allows for the calculation of important network quantities, such as centrality and clustering coefficients, as well as modularity [6]. Thus, given the inherent multiplexity of network data across fields as well as recent theoretical developments for

handling these types of data, there exists a need for the development of appropriate tools that can leverage information from all layers to elucidate structural patterns.

Each layer in a multilayer network provides its own information about interactions between nodes, and it is useful to ask whether sets of layers are providing redundant information. Addressing this question requires the development of an approach to compress networks into a reduced-layer representation such that it effectively retains the information from the original multilayer network. Aggregating layers can potentially result in a loss of information, but it can also successfully corroborate the existence of underlying structural patterns. Moreover, this can lead to improved identification of structural patterns, including enhanced community detection [7]. This idea of reducibility in multilayer networks was previously explored in [8]: using an information-theoretic notion of distance between pairs of network layers, the authors performed hierarchical clustering of layers and chose the partition that maximized a quality function reflecting information loss due to the aggregation of layers.

Inspired by the ideas in [8] that groups of layers often provide redundant information, we seek to further explore this idea to identify sets of layers, which we denote as “strata”, with each stratum described by a single probabilistic model based on community structure. This effectively amounts to defining *local* probabilistic network models, and is analogous to biclustering [9] or co-clustering [10] problems. Moreover, our method can be regarded as a joint clustering procedure, in which the nodes and layers of networks are clustered simultaneously. Just as in [10], where the objective is to jointly cluster words and documents such that joint word-document subgroups correspond to particular topics, our objective is to cluster network layers such that each stratum is a set of layers with a characteristic community structure. To achieve this goal, we have developed the

- N. Stanley is with the Curriculum in Bioinformatics and Computational Biology and with the Carolina Center for Interdisciplinary Applied Mathematics, Department of Mathematics, University of North Carolina, Chapel Hill. E-mail: stanleyn@email.unc.edu.
- S. Shai, D. Taylor, and P. J. Mucha are with the Carolina Center for Interdisciplinary Applied Mathematics, Department of Mathematics, University of North Carolina, Chapel Hill. E-mail: {sshai, taylordr}@live.unc.edu, mucha@unc.edu.

Manuscript received 9 Oct. 2015; revised 12 Jan. 2016; accepted 11 Feb. 2016. Date of publication 24 Mar. 2016; date of current version 10 June 2016.

Recommended for acceptance by R. D’Souza.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TNSE.2016.2537545

strata multilayer stochastic block model (sMLSBM). We additionally emphasize that by collectively utilizing similar layers in a principled way, we can achieve more robust community detection and parameter inference for the probabilistic community detection models that describe each stratum.

1.1 Network Comparison Based on Community Structure

The problem of aggregating layers in a multilayer network is closely related to the problem of clustering networks. That is, given an ensemble of networks, one aims to identify sets such that networks within a set have similar characteristics. These characteristics, or “features” in this context, can describe any of the following: micro-scale structural properties such as subgraph motifs [11], [12]; multiscale properties such as community structure [13], [14], [15], the spectra of network-related matrices [16] and by defining latent roles [17]. Although clustering layers in a multilayer network is closely related to clustering networks in an ensemble, these are distinct problems with different difficulties and nuances. We focus on the prior pursuit; however, we expect for certain network ensembles that it will be beneficial to modify and apply our methods to the clustering of networks.

In this work, we analyze and compare layers in a multilayer network based on their community structure. Community detection in single-layer networks is an essential tool for understanding the organization and functional relatedness between nodes in a network [18], [19]. Although there are many definitions for what constitutes a “community” [20], one often assumes an “assortative community” in which there is a prevalence of edges between nodes in the same community as compared to the amount of edges connecting these nodes to the remaining network. In seeking to identify such communities, numerous approaches have been proposed, including those based on maximizing a modularity measure [21] and fitting a generative probabilistic model [22]. Because each of these approaches present computational challenges for efficiently detecting communities, numerous heuristics exist for developing practical algorithms [19], [23], [24], [25], [26].

While our approach is to define a probabilistic model for multilayer community structure, we note that there have previously been approaches to understand similarities in network ensembles that are grounded in exploiting similarities in community structure between networks. In [14], the authors seek to partition a group of networks into subgroups through construction of a network of networks (NoN). Communities in the NoN are chosen such that the networks representing the nodes are sufficiently similar in their underlying community structure. In one significant application of this method, the authors clustered gene co-expression networks and found an increased number of significant functional enrichment categories for biological processes. Similarly, in [15], the authors explore mesoscopic similarity between layers using an informational theoretic approach. While they have designed their method to handle any feature of network architecture, they highlight their ability to quantify similarity between network layers based on node-to-community assignments in the layers.

In seeking a statistically-grounded approach for studying communities in multilayer networks, we consider the

stochastic block model (SBM) [27], a popular generative model for community structure in networks. The assumption of the SBM is that nodes in a particular community are related to nodes within and between communities in the same way, thus allowing SBMs to describe several types of communities (e.g., assortative, disassortative, core-periphery, etc. [20], [28]). There are many other appealing aspects of stochastic block models; for example, a model-based approach allows for the denoising of networks through the removal of false edges and the addition of missing edges [22], [29]. The inference procedure for fitting SBMs to an undirected network with N nodes and K communities involves learning the two parameters, π and \mathbf{Z} . Parameter π is a $K \times K$ symmetric matrix, where π_{mn} gives the probability of an edge existing between a given node in community m and another node in community n . Matrix \mathbf{Z} is an $N \times K$ indicator matrix, wherein each binary entry Z_{im} indicates whether or not node i is in community m . Each row of \mathbf{Z} is constrained such that $\sum_{m=1}^K Z_{im} = 1$, i.e., each node only belongs to 1 community. We also define vector z , which has entries $z_i = \operatorname{argmax}_m \{Z_{im}\}$ that indicate the community to which node i belongs. For a given network, these parameters are often inferred through a maximum likelihood approach, and once learned, they provide information about the within and between community relatedness.

1.2 Related Work on Multilayer SBMs

Due to the ubiquity of network data with multiple network layers, community detection in multilayer networks constitutes an important body of research. Important directions include generalizing the modularity measure [6] and studying dynamics [30] for this more general setting.

Given the usefulness of SBMs for the understanding of node organization in single-layer networks, it is important to extend SBMs to the multilayer framework, and indeed this direction of research is receiving growing attention [7], [31], [32], [33], [34]. In this context, the general assumption is that there are shared patterns in community structure across the layers of a multilayer network, and the goal is to define and identify a stochastic block model that captures this structure. These works have explored many types of applications that can arise involving multilayer networks, and have therefore given rise to several complementary models for multilayer stochastic block models (MLSBMs). We now briefly summarize this previous work that is very related, but notably different, from the model we study herein.

In Refs. [7], [31], [32], the authors studied situations in which many layers follow from a single SBM. In these instances, it is possible to obtain improved inference of the SBM parameters by incorporating multiple samples from a single model. For example, in Ref. [7] the authors considered an increasing number of layers, L , and explored asymptotic properties of the estimated SBM parameters. Specifically, they fit an SBM to each individual layer in a way that utilizes the information from all layers, and they showed convergence of these estimators to their true values as $L \rightarrow \infty$. For a network with L layers and K communities in each layer, their approach requires an estimate of the community assignment matrix \mathbf{Z}^l and probability matrix π^l for each layer l , the latter of which involves learning $K(K+1)L/2$ parameters.

To this end, the authors extended the variational approximation for approximating the maximum likelihood estimates of SBM parameters introduced in single-layer SBMs introduced in [35] to the multilayer setting.

Ref. [7] was followed up by Ref. [31], wherein the authors addressed issues that can arise for the model when K and/or L is large, or if the network is sparse. They proposed a modified model called the restricted multilayer stochastic block model (rMLSBM). In this model, instead of learning a set of L independent parameters, π_{mn}^l , for each pair, (m, n) , each entry in π is fully layer-dependent so as to produce a reduction in the number of free parameters. Specifically, to determine the probability of an edge between a node from community m and a node from community n in layer l , they use a logistic link function and model the probability as $\text{logit}(\pi_{mn}^l) = \pi_{mn} + \beta_l$. The β_l is an offset parameter representing the particular layer or type of edge. In this model, it is necessary to learn $K(K+1)/2 + L$ total parameters. Thus, the maximum likelihood estimate for an rMLSBM is a regularized estimator.

Consistent with the theme of fitting a single block model to a collection of layers, Ref. [32] is similar to Refs. [7] and [31] in that the authors seek to leverage information from all layers by considering the joint distribution of layers. Using this, they estimated quantities such as the marginal probabilities of node assignments to communities and the edge probabilities within and between groups. An interesting aspect of their approach is that they introduce a covariate capturing the coupling between pairs of nodes. For a network with K communities and L layers, this requires the estimation of $(2^L - 1)K^2 + (K - 1)$ parameters.

We summarize Refs. [33] and [34], which provide techniques to determine whether a single layer network is the result of an aggregation procedure in a multilayer network. In Ref. [33], the authors defined a version of multilayer stochastic block model and an inference procedure for assessing whether or not a single-layer network was actually obtained from an aggregation of layers in a multilayer network; they considered the aggregation of layers using boolean rules. Ref. [34] describes two possible generative processes for multilayer networks: the *edge-covariate* and *independent-layer* models. In the edge-covariate model, an aggregated network is defined in which a given edge (i, j) only appears in a single layer. Aggregating the layers in a multilayer network into a single network representation combines all of the edges from each of the layers. Thus, the translation of this idea into a generative model involves choosing a layer membership for each edge and sampling edges with a probability conditioned on adjacent nodes. In the independent-layer model, layers are generated independently from each other and the only constraint is that group membership of the nodes are the same across all layers.

While motivation to pursue this problem originated from [8], we point out that our approach does not provide a method for aggregating layers or reducing the number of layers in the network. Instead, it can in a sense compress the network in that the learned stochastic block model parameters for each stratum can be used to generate a sample network to serve as a consensus for that stratum.

1.3 Contributions

While the literature on MLSBMs has recently grown quickly, there is still a need for a probabilistic generative model that allows for the layers in a multilayer network to be described by multiple SBMs. To this end, we developed a novel multilayer stochastic block model, sMLSBM, that assigns network layers into disjoint sets that we call strata, where a collection of layers in a given stratum are assumed to be samples from the same underlying generative model. Our method can be viewed as a joint clustering procedure, where we seek to group layers into strata and nodes into communities. That is, we seek to simultaneously find layer-to-strata and node-to-community assignments.

In order to address practical applications that can involve multilayer networks with several strata, layers, communities and nodes, we introduce an algorithm that effectively partitions layers into strata and an inference procedure to learn the SBM parameters for each stratum. Importantly, these two steps—assigning nodes to communities and layers to strata—are combined in an iterative algorithm so that an improvement in community detection can lead to an improvement in the clustering of layers into strata, which can iteratively lead to further improvement in community detection, and so on.

To describe the model, the algorithm for fitting the model, and its performance, the remainder of this paper is organized as follows. In Section 2, we define the model and an algorithm for fitting it. In Section 3, we perform numerical experiments on synthetic networks. In Section 4, we test the model on correlation networks constructed from data from the Human Microbiome Project.

2 SMLSBM: STRATA MULTILAYER STOCHASTIC BLOCK MODEL

2.1 Network Definition

Let $G(N, \mathcal{E})$ define a single network with N nodes and a set of undirected edges, $\mathcal{E} = \{(i, j)\}$. We define a multiplex network, which is one kind of multilayer network [1], [2], by defining a set of network layers, $G^l(N, \mathcal{E}^l)$, where $l \in \mathcal{L}$ and the set $\mathcal{L} = \{1, 2, \dots, L\}$ indicates the layers' indices. We denote the collection of L layers as a set, \mathcal{G} , such that $\mathcal{G} = \{G^1, G^2, \dots, G^L\}$ makes up the multiplex network and each element of the set is the network representing a layer. Furthermore, we define $\mathcal{A} = \{\mathbf{A}^1, \mathbf{A}^2, \dots, \mathbf{A}^L\}$ to be the corresponding adjacency matrix representations of the network layers in \mathcal{G} .

2.2 Model Definition

Under the sMLSBM, the network layers, $G^l(N, \mathcal{E}^l)$ are assumed to be generated by a set of S stochastic block models, where the layers in stratum $s \in \{1, 2, \dots, S\}$, are parameterized by π^s and \mathbf{Z}^s (or equivalently, vector \mathbf{z}^s , which has entries $z_i^s = \arg\max_m \{Z_{im}^s\}$). Note that the parameters π^s and \mathbf{Z}^s for a single stratum are analogous in meaning to their respective parameters in the single-layer SBM case (see Section 1.1). For each stratum s , we let $\mathcal{L}^s \subseteq \mathcal{L}$ denote the set of layers corresponding to s , so that $\mathcal{L} = \bigcup_s \mathcal{L}^s$ and $\emptyset = \mathcal{L}^s \cap \mathcal{L}^t$ for all $s, t \in \{1, \dots, S\}$, $s \neq t$. We let $L^s = |\mathcal{L}^s|$ denote the number of layers in strata s so that $\sum_s L^s = L$. Finally, we allow the number of communities, K^s , to vary across the strata.




Fig. 1. Objective of strata multilayer stochastic block model (sMLSMB). Each of the $L = 9$ networks here represents a layer in a multilayer network. Every network layer has $N = 36$ nodes that are consistent across all layers. There are $S = 3$ strata as indicated by the three rows and the colors of nodes. Clearly, network layers within a stratum exhibit strong similarities in community structure. That is, although each layer follows an SBM with $K = 3$ communities, the SBM parameters are identical for layers within a strata but differ between layers in different strata. We would like to partition the layers into their appropriate strata and learn their associated SBM parameters, π^s and Z^s .

For a given multilayer network, our objective during inference is to identify the stratum assignment of each layer and to learn the collection of strata parameters, $\Pi = \{\pi^1, \pi^2, \dots, \pi^S\}$ and $Z = \{Z^1, Z^2, \dots, Z^S\}$. The learned SBM parameters for a stratum represent a consensus for the associated layers, and so in that sense can be interpreted as reducing the effective number of layers [8]. However, strata can also be interpreted as a way to simply identify layers with similarities in community structure. Fig. 1 shows a toy example of a multilayer network with $S = 3$ strata, where each layer has $N = 36$ nodes and $K = 3$ communities. Each individual network in this figure represents a layer in the network. The nodes in the layers belonging to each stratum are colored according to their stratum membership; moreover, it is easy to see that layers of a stratum exhibit high similarities in community structure.

As part of our procedure, we specify another parameter that we refer to as the adjacency probability matrix, θ^s , which can be computed from π^s and Z^s . Specifically, θ^s is an $N \times N$ matrix such that θ_{ij}^s gives the probability of an edge between nodes i and j in stratum s . That is, $\theta_{ij}^s = \pi_{z_i^s z_j^s}^s$, where z_i^s specifies the community number for node i in stratum s . Finally, we define the matrix Y of size $L \times S$, wherein an entry Y_{ls} is a binary indicator of whether or not layer l is assigned to stratum s . Note that $\sum_s Y_{ls} = 1$. We also define a vector y , which has entries $y_l = \arg\max_s \{Y_{ls}\}$ to indicate the strata to which layer l belongs.

2.3 Inference for sMLSMB

The procedure for fitting an sMLSMB to a given network requires finding the layer-to-strata memberships

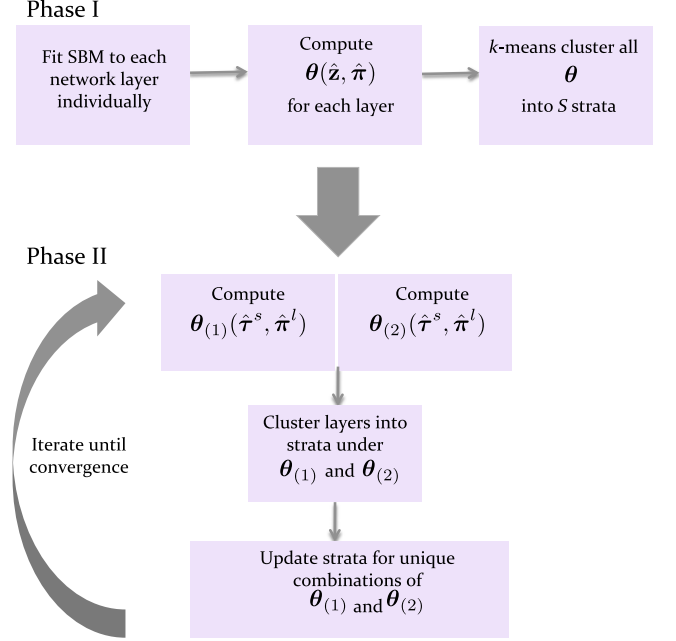


Fig. 2. Schematic illustration of our algorithm: Our algorithm for fitting an sMLSMB is broken up into two phases: an initialization phase to cluster layers into strata, and an iterative phase that allows learning of node-to-community and layer-to-strata assignments.

and node-to-community memberships that best describe the multilayer network. For notational convenience, we introduce hat notation to represent the learned parameter estimate from the inference procedure. We can write down the marginal likelihood for the collection of network layers, \mathcal{G} , as,

$$p(\mathcal{G} | \Pi) = \sum_{\mathcal{Z}} \sum_{\mathbf{Y}} p(\mathcal{G}, \mathcal{Z}, \mathbf{Y} | \Pi). \quad (1)$$

We assume the probability of an edge between two nodes in layer l belonging to stratum s can be modeled as a Bernoulli random variable, based on the community membership of the nodes. In particular, $p(A_{ij}^l = 1) \sim \text{Bernoulli}(\pi_{z_i^s z_j^s}^s)$.

Since \mathbf{Y} and \mathcal{Z} are both latent quantities, searching over all possible values quickly becomes intractable. To tackle this issue, we develop a two-phase algorithm that incorporates a clustering algorithm for choosing the best \mathbf{Y} . This greedy approach leads to a significant reduction for the size of the search space since only \mathcal{Z} must be statistically inferred. Specifically, during Phase I, we infer an SBM for each layer in isolation, and we cluster together sets of layers that have similar SBM parameters. Using these results as an initial condition in Phase II, we develop an iterative method that jointly identifies layer-to-stratum and node-to-community assignments as well as the SBM parameters for each stratum. We provide a schematic of the algorithm in Fig. 2, and below we present the two-phase algorithm in detail.

Phase I. Phase I is comprised of two parts. First, we fit an SBM to each individual layer $l \in \{1, \dots, L\}$, which yields inferred SBM parameters $\hat{\pi}^l$ and node-to-community memberships \hat{Z}^l . Then we cluster the layers based on the similarities of $\hat{\pi}^l$ and \hat{Z}^l . To infer $\hat{\pi}^l$ and \hat{Z}^l , we use the inference method described in [35]. Here, the authors used a

variational inference technique to approximate the maximum likelihood estimates for the stochastic block model parameters. For the set of L layers, this produces sets of SBM parameters for each layer, which we denote by $\hat{\Pi} = \{\hat{\pi}^1, \hat{\pi}^2, \dots, \hat{\pi}^L\}$ and $\hat{\mathbf{Z}} = \{\hat{\mathbf{Z}}^1, \hat{\mathbf{Z}}^2, \dots, \hat{\mathbf{Z}}^L\}$ (that is, at this stage of the procedure, each layer is temporarily treated as its own stratum). Note also that each $\hat{\mathbf{Z}}^l$ can be equivalently represented by vector $\hat{\mathbf{z}}^l$, as described in Section 1.1. Using the estimates $\hat{\pi}^l$ and $\hat{\mathbf{Z}}^l$ for a given layer, l , we can construct the corresponding adjacency probability matrix, $\hat{\theta}^l$, which is defined entry-wise by $\hat{\theta}_{ij}^l = \hat{\pi}_{\hat{z}_i, \hat{z}_j}^l$. Doing this for each layer results in a collection of adjacency probability matrices, $\hat{\Theta} = \{\hat{\theta}^1, \hat{\theta}^2, \dots, \hat{\theta}^L\}$.

Now, we seek an initial partition of layers into strata based on $\hat{\Theta}$. The goal is to identify S sets \mathcal{L}^s so that the matrices $\{\hat{\theta}^l\}$ with $l \in \mathcal{L}^s$ are close to one another, but they are distant from the remaining matrices, $\{\hat{\theta}^l\}$ with $l \in \mathcal{L} \setminus \mathcal{L}^s$. This is accomplished by treating each $\hat{\theta}^l$ as a feature vector and applying k -means clustering with S centers so as to identify S strata, \mathcal{L}^s . Note that S can be selected *a priori*, or approximated with a measure such as the gap statistic [36]. This gives us an initial estimate $\hat{\mathbf{Y}}$ for \mathbf{Y} . Note that this procedure initially treats each layer as a separate stratum, but provides a principled agglomeration of layers into $S \leq L$ strata.

Phase II. After a first-pass approach for assigning layers to strata, we initialize an iterative phase to more effectively estimate layer-to-strata assignments as well as the model parameters. Specifically, we would like to find the consensus SBM for each strata—that is, the $K^s \times K^s$ matrix π^s and the $N \times K^s$ matrix \mathbf{Z}^s that maximize the likelihood of the observed layers in each stratum. We let $\mathcal{A}^s = \{\mathbf{A}^l\}$ for $l \in \mathcal{L}^s$ denote the collection of adjacency matrices corresponding to the L^s layers in stratum s .

We now proceed to maximize the likelihood in each stratum, by extending the framework of Ref. [35] to a multilayer context. Note that this is similar to Ref. [7], except that we are not aiming to infer an SBM probability matrix for each layer, individually. In particular, the complete-data log-likelihood for stratum s can be written as,

$$p(\mathcal{A}^s, \mathbf{Z}^s) = p(\mathcal{A}^s | \mathbf{Z}^s) p(\mathbf{Z}^s), \quad (2)$$

where

$$p(\mathcal{A}^s | \mathbf{Z}^s) = \prod_{l \in \mathcal{L}^s} \prod_{i < j} \prod_{mn} \pi_{mn}^s A_{ij}^l (1 - \pi_{mn}^s)^{(1 - A_{ij}^l)}. \quad (3)$$

To write $p(\mathbf{Z}^s)$, it is helpful to introduce a new parameter α_m^s that represents the probability that a randomly-selected node in stratum s belongs to community m , i.e., $\alpha_m^s = p(Z_{im}^s = 1)$.

Note that $\sum_m \alpha_m^s = 1$. Using this parameter, we can write

$$p(\mathbf{Z}^s) = \prod_i \prod_m \alpha_m^s (Z_{im}^s). \quad (4)$$

It follows that the complete-data log-likelihood for the adjacency matrices representing the layers in stratum s can be expressed as,

$$\begin{aligned} \log P(\mathcal{A}^s, \mathbf{Z}^s) &= \log(P(\mathbf{Z}^s)) + \log(P(\mathcal{A}^s | \mathbf{Z}^s)) \\ &= \sum_i \sum_m Z_{im}^s \log(\alpha_m^s) \\ &\quad + \sum_{l \in \mathcal{L}^s} \sum_{i < j} \sum_{mn} A_{ij}^l \log(\pi_{mn}^s) \\ &\quad + \sum_{l \in \mathcal{L}^s} \sum_{i < j} \sum_{mn} (1 - A_{ij}^l) \log(1 - \pi_{mn}^s). \end{aligned} \quad (5)$$

Problems of this variety that involve the need to compute maximum likelihood estimates with incomplete data are typically addressed with the expectation maximization (EM) framework [37]. Doing so requires the ability to compute $P(\mathbf{Z}^s | \mathcal{A}^s)$; however, Ref. [35] showed that it is intractable to calculate the conditional distribution for the single-layer network case. To address this challenge, we use a variational approximation, analogous to approaches in [7], [32], [35]. In general, a variational approximation seeks to optimize a lower bound on the log-likelihood. To do this, we first approximate the conditional distribution, $P(\mathbf{Z}^s | \mathcal{A}^s) \approx R_{\mathcal{A}^s}$, where

$$R_{\mathcal{A}^s}(\mathbf{Z}^s) = \prod_i h(\mathbf{Z}_i^s; \boldsymbol{\tau}_i). \quad (6)$$

Here, matrix $\boldsymbol{\tau}^s$ contains entries τ_{im}^s that approximate the probability that node i belongs to community m in stratum s . Further, function $h(\cdot)$ represents the multinomial distribution, with parameters, $\{\tau_{im}^s\}$ for $m \in \{1, \dots, K^s\}$. Using this, we define the variational approximation as

$$\mathcal{J}(R_{\mathcal{A}^s}) = \ell\ell(\mathcal{A}^s) - \text{KL}(R_{\mathcal{A}^s}(\mathbf{Z}^s), P(\mathbf{Z}^s | \mathcal{A}^s)), \quad (7)$$

where $\ell\ell$ is log likelihood and KL is the Kullback-Leibler divergence.

Through maximizing $\mathcal{J}(R_{\mathcal{A}^s})$, we minimize the KL divergence between the true conditional distribution, $P(\mathbf{Z}^s | \mathcal{A}^s)$, and its approximation, $R_{\mathcal{A}^s}(\mathbf{Z}^s)$. Moreover, we follow the derivation in Ref. [35] and rewrite $\mathcal{J}(R_{\mathcal{A}^s})$ as

$$\begin{aligned} \mathcal{J}(R_{\mathcal{A}^s}) &= \sum_i \sum_m \tau_{im}^s \log(\alpha_m^s) \\ &\quad + \sum_{l \in \mathcal{L}^s} \sum_{i < j} \sum_{mn} \tau_{im}^s \tau_{jn}^s [A_{ij}^l \log(\pi_{mn}^s)] \\ &\quad + \sum_{l \in \mathcal{L}^s} \sum_{i < j} \sum_{mn} \tau_{im}^s \tau_{jn}^s [(1 - A_{ij}^l) \log(1 - \pi_{mn}^s)] \\ &\quad - \sum_i \sum_m \tau_{im}^s \log(\tau_{im}^s). \end{aligned} \quad (8)$$

We can now differentiate $\mathcal{J}(R_{\mathcal{A}^s})$ with respect to each parameter—while using Lagrange multipliers to enforce constraints (i.e., probabilities summing to 1)—to compute the updates. Doing so yields the following, where the hat notation symbolizes the current best estimate for the given parameter:

$$\hat{\alpha}_m^s = \sum_i \hat{\tau}_{im}^s / N, \quad (9)$$

$$\hat{\pi}_{mn}^s = \frac{\sum_{l \in \mathcal{L}^s} \sum_{i < j} \hat{\tau}_{im}^s \hat{\tau}_{jn}^s A_{ij}^l}{\sum_{l \in \mathcal{L}^s} \sum_{i < j} \hat{\tau}_{im}^s \hat{\tau}_{jn}^s}, \quad (10)$$

$$\hat{\tau}_{im}^s \propto \hat{\alpha}_m^s \prod_{l \in \mathcal{L}^s} \prod_{i < j} \prod_n [\hat{\pi}_{mn}^s A_{ij}^l (1 - \hat{\pi}_{mn}^s)^{1 - A_{ij}^l}]^{\tau_{jm}^s}. \quad (11)$$

To find the best estimates for $\hat{\tau}^s$ and $\hat{\pi}^s$, we alternate between updating $\hat{\tau}^s$ and $\hat{\pi}^s$ until convergence. When convergence has occurred, we refer to the resulting estimates as the consensus $\bar{\tau}^s$ and $\bar{\pi}^s$ for stratum s . Similarly, \bar{Z}^s represents the consensus indicator matrix of node-to-community assignments computed from $\bar{\tau}^s$. Note that we use the bar notation to reflect that the particular parameter estimate is for a stratum, rather than for an individual layer.

Since $\bar{\tau}^s$ and $\bar{\pi}^s$ are computed in terms of each other, we can use one of the consensus parameters to compute the other parameter in individual layers. In particular, using the fixed node-to-community assignments from $\bar{\tau}^s$, we compute the maximum-likelihood SBM parameters for a particular layer l , which we denote with a tilde and hence, $\tilde{\pi}^l$ and $\tilde{\tau}^l$. Similarly, for fixed $\bar{\pi}^s$, we compute the node-to-community assignments $\tilde{\tau}^l$. Such estimates allow us to determine whether or not the stratum consensus estimates are accurate estimates for the SBMs of individual layers of the stratum. More importantly, as we shall now describe, these layer-specific estimates allow us to design an iterative algorithm that allows for alternating between learning the node-to-community and layer-to-stratum assignments.

To this end, we represent each layer by the adjacency probability matrix, which we compute two different ways: letting $\theta(\tau, \pi)$ represent the adjacency probability matrix specified by τ and π , we define

$$\theta_{(1)}^l = \theta(\bar{\tau}^s, \tilde{\pi}^l), \quad (12)$$

$$\theta_{(2)}^l = \theta(\tilde{\tau}^l, \bar{\pi}^s). \quad (13)$$

Note that the first definition uses the strata-consensus estimate for τ^s and a layer-specific estimate for π^s , whereas the latter uses a layer-specific estimate for τ^s and the strata-consensus estimate for π^s .

During Phase I, we identified strata by clustering the adjacency probability matrices for the L layers using the k -means algorithm. We employ a similar procedure here, but instead of clustering L matrices, we now cluster $2L$ matrices, since each layer is represented in two different ways. Moreover, clustering these $2L$ matrices yields two cluster assignments for each layer. Typically, both representations of a particular layer will receive identical cluster assignments—that is, for a given l , $\theta_{(1)}^l$ and $\theta_{(2)}^l$ are assigned to the same cluster, or strata. However, an interesting case arises when the two representations induce different stratum assignments for a given layer, because this implies that there is disagreement between $\theta_{(1)}^l$ and $\theta_{(2)}^l$, which implies uncertainty in the strata assignment of that particular layer l . Because our iterative algorithm requires each layer to be assigned to a single stratum (i.e., we do not allow for mixed membership of layers into strata), layers with mixed membership according to $\theta_{(1)}^l$ and $\theta_{(2)}^l$ must be dealt with in some way. To account for these situations, we define additional strata for each combination of membership that arises. For example, if there are several layers $\{l\}$ that are clustered into stratum 1 according to $\theta_{(1)}^l$ and stratum 2 according to $\theta_{(2)}^l$, then we define a new stratum that contains only these layers. We note that there exists a variety of options for handling layers with such mixed membership after applying

k -means clustering to $\theta_{(1)}^l$ and $\theta_{(2)}^l$ (e.g., one could assign such a layer to a stratum at random); however, we leave open for future work the exploration of these other options.

After a single pass of Phase II, which requires layer-to-strata assignments (which can be encoded by vector y) as input, the algorithm yields (ideally) improved layer-to-strata assignments (as well as consensus estimates for the SBM parameters of the strata, $\bar{\tau}^s$ and $\bar{\pi}^s$). Therefore, Phase II involves iterating the above procedure until the layer-to-strata assignments do not change. We note that in principle, it is possible for new strata to arise in each iteration (i.e., because we create strata to avoid mixed membership of layers), and this can allow the number of strata to grow with each iteration; however, we did not observe this issue in any of our synthetic or real data experiments. As we will show in the following section, convergence is typically observed after just a few iterations (e.g., see, for example, the second row of Fig. 4). If such an issue arises, it may be helpful to bound the number of iterations (NOI) in Phase II.

3 SYNTHETIC DATA EXPERIEMENTS

3.1 Comparison of sMLSBM to Other SBM Approaches

To demonstrate a situation where the sMLSBM framework has a clear advantage over other models, we designed a synthetic experiment and compared the results to two different SBM approaches: i) fitting a single SBM to all of the layers (denoted “single SBM”), and ii). fitting a stochastic block model to each layer individually (denoted “single-layer SBM”). We generated a multilayer network, where each layer has $N = 128$ nodes, $K = 4$ communities and an expected mean degree of $c = 20$ (i.e., every network layer is expected to contain $cN/2 = 1280$ undirected edges). We specified an sMLSBM with $S = 3$ strata and 10 layers per strata, which resulted in $L = 30$ total layers. We defined π^s for each stratum s in terms of two parameters, p_{in}^s and p_{out}^s , which give the within-community edge probabilities and between-community edge probabilities, respectively. That is, we define $\pi_{mn}^s = p_{in}^s$ when $m = n$ and $\pi_{mn}^s = p_{out}^s$ when $m \neq n$. It follows that the expected mean degree is given by $c = N(p_{in}^s + (K - 1)p_{out}^s)/K$. In our experiment, we select the following SBM parameters: $(p_{in}^1, p_{out}^1) = (0.6, 0.0083)$; $(p_{in}^2, p_{out}^2) = (0.4, 0.075)$; and $(p_{in}^3, p_{out}^3) = (0.125, 0.167)$. In Fig. 3A, we show an example network layer from each strata. Nodes are colored by their community assignments in stratum 1. Note that the node-to-community assignments are different in each stratum and that the extent of block structure decreases from stratum 1 to stratum 3.

In order to compare the accuracy of fit for the three models—single-layer SBM, single SBM and sMLSBM—we quantify the inference accuracy of the SBM parameters, $\bar{\pi}^s$, and community assignments, \bar{Z}^s . First, for each layer and each model, we quantified the error (ℓ^2 norm) between $\text{vec}(\bar{\pi}^s)$ and its true value, $\text{vec}(\pi^l)$. Note that $\text{vec}(\mathbf{X})$ is the $\frac{K(K+1)}{2}$ length vector representing the lower triangle of the matrix \mathbf{X} . Moreover, to quantify error, we compute $\|\text{vec}(\pi^l) - \text{vec}(\bar{\pi}^s)\|_2$. We note that this error is well-defined because we identify $K = 4$ communities for all layers and all models. The mean error across layers under each model are

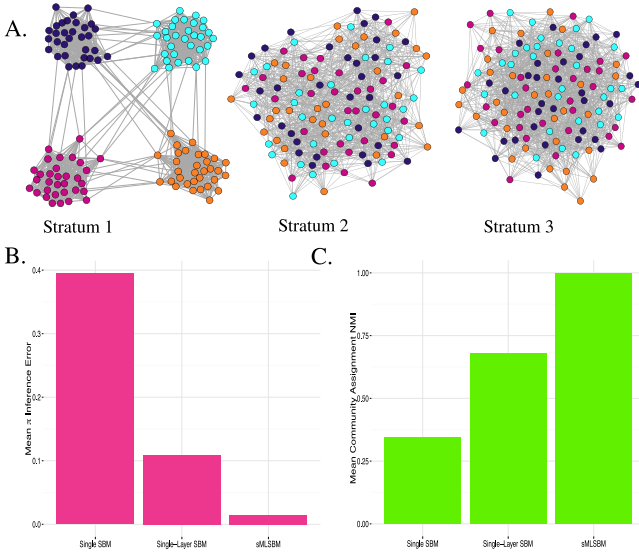


Fig. 3. *Synthetic experiment comparing sMLSBM to other SBMs.* **A.** We specified a model with $S = 3$ strata and $L = 10$ layers per stratum. A representative layer from each stratum is plotted. Note that nodes in all networks are colored according to their community membership in stratum 1. Each network has $N = 128$ nodes, $K = 4$ communities and mean degree, $c = 20$. The p_{in}^s parameters for $s = 1, 2$ and 3 are $0.6, 0.4$ and 0.25 , respectively. Corresponding values of p_{out}^s were selected to maintain the desired expected mean degree, $c=20$. **B.** We fit three types of models to the 30 network layers: i) single SBM: fitting a single SBM to all of the layers; ii) single-Layer SBM: fitting an individual SBM to each layer; and iii) sMLSBM: identifying strata and fitting an SBMs for each stratum. Each model yields an estimate $\hat{\pi}^l$ for the true SBM of each layer l , which is denoted π^l . Here s_l denotes the inferred strata for layer l . On the vertical axis we plot the mean ℓ_2 norm error $\|\text{vec}(\pi^l) - \text{vec}(\hat{\pi}^{s_l})\|_2$. **C.** For each of the three models, we computed the normalized mutual information (NMI) between the true node-to-community assignments \mathbf{z}^l and the inferred values $\hat{\mathbf{z}}^{s_l}$.

shown in Fig. 3B. In this example, sMLSBM outperforms the two other models. Second, we computed for each layer the mean normalized mutual information (NMI) [38] between the true node-to-community assignments, \mathbf{z}^l , and the inferred values, $\hat{\mathbf{z}}^{s_l}$, under each model. In other words, for each layer, we compute, $\text{NMI}(\mathbf{z}^l, \hat{\mathbf{z}}^{s_l})$. Fig. 3C shows the mean NMI for community assignments across layers. Indeed, the effects of fitting an incorrect model to a collection of layers in terms of ability to effectively estimate SBM parameters and community assignments is apparent. In particular, fitting a single SBM model results in both larger mean inference and community assignment error, compared to fitting single-layer SBMs and three strata sMLSBM. In other words, sMLSBM provides an efficient clustering into strata only when the layers are indeed related (i.e., generated from the same SBM), otherwise each layer is a stratum on its own.

3.2 Synthetic Experiment with Two Strata

Next, we further explored the performance of our algorithm (see Section 2.3) for inferring an sMLSBM under various situations: 1) in comparison to baseline clustering methods; 2) in response to an increase in the number of layers; and 3) under variations in levels of detectability. Specifically, we designed synthetic experiments in which we generated multilayer networks with either $L = 10$ or $L = 100$ layers. Every multilayer network contained $S = 2$ strata (each




Fig. 4. *Synthetic experiment with two strata.* We conducted numerical experiments with multilayer networks with $N = 128$ nodes, mean degree $c = 16$, $S = 2$ strata and $K^1 = K^2 = 4$ communities. The networks contained either $L = 10$ (left column) or $L = 100$ layers (right column), which were divided equally into the two strata. For stratum 1, we fixed the quantity $N(p_{in}^1 - p_{out}^1) = 10$, which fully specifies (p_{in}^1, p_{out}^1) since setting $c = 16$ also constrains these parameters. In contrast, we vary $N(p_{in}^2 - p_{out}^2)$. **A.** As a function of $N(p_{in}^2 - p_{out}^2)$, we plot the mean NMI to interpret the ability of sMLSBM to recover the true layer-to-strata assignments. We compare the performance of sMLSBM (purple curve) to generic k -means clustering (green symbols) of adjacency matrices. **B.** We plot the mean number of iterations required for Phase II of our algorithm (see Section 2.3) to converge. **C.** Finally, we measure the quality of node-to-community assignment results by plotting the mean NMI between the true node-to-community assignments and those inferred with sMLSBM in stratum 1 (red symbols) and stratum 2 (blue symbols).

having $K^1 = K^2 = 4$ communities), and in each layer there were $N = 128$ nodes (each having an expected mean degree of $c = 16$). Note that in this example both strata have the same node-to-community assignments. The strata were fixed to be the same size, $L^1 = L^2 = L/2$. Similar to the experiment described in Section 3.1, the SBM parameters were constructed using p_{in}^s and p_{out}^s . Since we have already specified the expected mean degree, these parameters must satisfy the constraint $c = N(p_{in}^s + p_{out}^s)/2$ for both strata. In all simulations, we fixed the SBM parameters of the first strata as $(p_{in}^1, p_{out}^1) = (.1836, .1055)$. It is also convenient to define the quantity, $N(p_{in}^1 - p_{out}^1) = 10$, which relates to the detectability of communities [39]. For example, the ability to detect community structure in a given layer and/or strata is, in general, expected to improve with increasing $N(p_{in}^s - p_{out}^s)$. For the second strata, we allow $N(p_{in}^2 - p_{out}^2)$ to vary.

We present results for this experiment in Fig. 4, wherein the left and right columns give results for $L = 10$ and $L = 100$, respectively.

Symbols in each plot represent the mean over 50 multilayer networks, and error bars show standard error. In each plot, the vertical dotted line indicates $N(p_{in}^2 - p_{out}^2) = 10$, which represents the point where the two strata are indistinguishable since $(p_{in}^1, p_{out}^1) = (p_{in}^2, p_{out}^2)$. In Fig. 4A, we show the NMI between the true layer-to-strata assignments and those inferred by sMLSBM, or $\text{NMI}(\mathbf{y}, \hat{\mathbf{y}})$. As a baseline, we compare sMLSBM results to directly clustering the layers'

adjacency matrices using the k -means algorithm with $K = 2$. We consistently observe higher NMI as a result of sMLSBM compared to k -means. More interestingly is the case with $L = 100$, where both k -means and sMLSBM perform at least moderately well at partitioning layers into strata before the point where the strata are indistinguishable. In Fig. 4B, we plot the number of iterations required for Phase II of our algorithm to converge. We observe that as the number of layers in the network increases, so does the number of required sMLSBM iterations. Moreover, the peaks in panel B. correspond to the sudden jumps in strata NMI.

Finally, in Fig. 4C we show the quality of node-to-community assignments by plotting the NMI between the true and inferred node-to-community assignments as described in Section 3.1. Note that stratum 1 here represents the stratum where the majority of layers were generated from model S^1 and analogously for stratum 2. Therefore, when the strata NMI is low (panel A.), we see poorer community detection results than expected, as layers get incorrectly mixed. As the strata NMI increases, layers from the same model are assigned together and the communities NMI stabilizes. Finally, by comparing the results for $L = 100$ to those for $L = 10$, we observe an increase in number of layers, L , generally leads to an improvement in community detection and strata identification.

4 CORRELATION NETWORKS FROM THE HUMAN MICROBIOME PROJECT

As an application of sMLSBM, we consider correlation networks constructed from data from the Human Microbiome Project [40]. For various sites on the body, the human microbiome project has successfully collected multiple human samples in order to better understand interactions between bacterial species. In this context, network inference is particularly interesting, as such methods aim to capture the relationships between various organisms. Microorganisms exhibit intricate ecologies within the gut of their human host and particular body sites have been shown to possess characteristic interactions. Further, certain interactions between microbes can often be associated with particular health and disease states [41]. Microbiome data is typically collected through metagenomic sequencing and reads are further binned into groups, known as operational taxonomic units (OTUs), to represent particular organisms. The nature of this count-based sequencing data makes network inference challenging, and is thus an interesting field in itself. To demonstrate the potential use for sMLSBM in the context of the human microbiome, we applied our algorithm for learning sMLSBMs to multilayer networks constructed from the SparCC [42] network inference method.

SparCC is a correlation network inference method that aims to approximate the linear Pearson correlation between components in a system. This method performs favorably, as it accounts for the extent of diversity in the microbial community, which plays a significant role in detecting valid interactions. Furthermore, networks are constructed with the assumptions that the number of components in the system (e.g., OTUs) is large and that the correlation network should be sparse. As supplemental data in Ref. [42], the authors provided their inferred microbial interaction

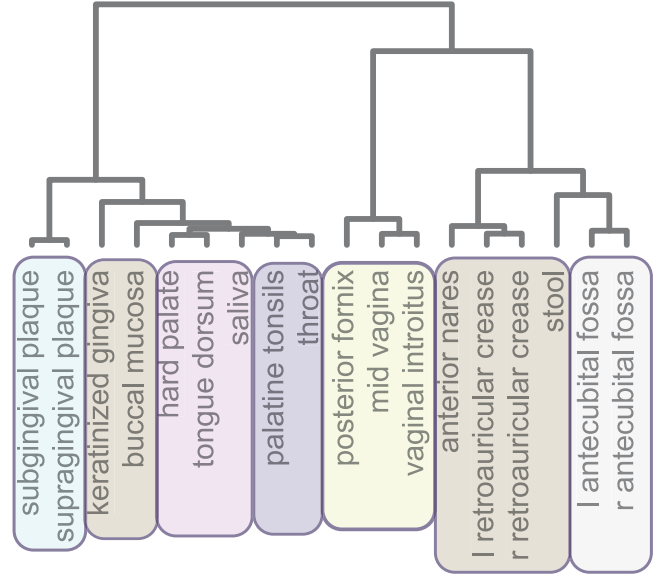


Fig. 5. Comparison of sMLSBM on the OTU interaction networks [42] for each of the body sites to a reducibility hierarchy [8]. As described in the text, we consider a multiplex network with $L = 18$ layers and $N = 213$ nodes, which we group here into $S = 6$ strata, while the dendrogram was generated by the method employed as the precursor to the reducibility framework. Colored boxes around the leaves of the dendrogram designate the body site to strata assignments obtained with sMLSBM.

networks for 18 sites in the human body, using the sparse, SparCC framework. The edges in these networks have positive and negative real-valued weights, based on the results of SparCC inference. In this analysis, we converted the SparCC networks into binary adjacency matrices by allowing a link only if the SparCC edge-weight between two OTUs was at least 0.15 (chosen as a value close to 0.2, given in Ref. [42]). To convert the 18 single-layer networks corresponding to species interactions in 18 body sites, we identified the collection of nodes (OTUs) that participated in at least two layers in terms of having at least one connecting edge weight value in the layer above the 0.15 threshold. This resulted in $N = 213$ unique OTUs (nodes) for our multilayer network analysis. We emphasize that restricting attention to nodes that participate in multiple layers was a choice we made in our focus on identifying common community structures across layers, to demonstrate the accuracy in the algorithm and inference procedures of sMLSBM. A more biologically-relevant treatment of this dataset should of course consider domain-specific expertise in formulating a network representation appropriate to the question at hand.

We inferred an sMLSBM for the multilayer network and chose to show results for $S = 6$ strata. That is, this selection leads us to find six clusters of body sites such that the microbiomes are similar between sites in the same cluster but differ from microbiomes at sites in the remaining clusters.

We indicate these six strata with colored boxes in Fig. 5. We note that due to the stochasticity of k -means in our algorithm, the communities and strata fit by sMLSBM can vary from one realization to the next. The shown strata assignments reflect those observed to yield the highest log-likelihood.

To gauge the performance of our method, we compared our strata membership results to the hierarchy obtained as

part of the reducibility method developed in [8]. To do this, we followed the following steps:

- 1) Compute the normalized Laplacian matrices for each of the 18 body site networks;
- 2) Compute the eigenvalues for each normalized Laplacian matrix;
- 3) Use these eigenvalues to compute the Von Neumann entropies for individual layers and pairs of layers;
- 4) Use the Von Neumann entropies to compute Jensen-Shannon distances between pairs of networks; and
- 5) Perform hierarchical clustering using the Jensen-Shannon distances and Ward linkage.

We show the results of this hierarchical clustering with a dendrogram in Fig. 5, which are in very good agreement with the sMLSBM results. However, as expected, we observe slight differences, since these methods cluster layers based on different criteria; in particular, sMLSBM partitioning reflects similarity only in community structure.

The results of both methods are relatively faithful to body regions in terms of groups of body sites that are spatially proximal. The only exception to this observation is the brown-colored stratum in Fig. 5, which is comprised of some seemingly unrelated body sites. While this grouping may not be intuitive, there is biological evidence to explain its plausibility. Specifically, Ref. [43] offers a state-of-the-art clustering of body sites based on biological expertise. Here, the authors have advanced understanding of microbial community composition through the application of a multinomial mixture model to define community types to characterize body sites. In particular, each sample collected through the Human Microbiome Project was assigned to 1 of 4 community types. They then quantified relationships between body sites using the p-value from a Fisher exact test on the membership of samples to community types. Similar to what we observe in the brown-colored stratum, the authors of [43] found a surprising correlation between samples from stool and oral cavity, which is reflected in our result.

In Fig. 6, we illustrate network layers for four of the six strata that we identify to highlight one advantage of having a probabilistic generative model for microbial composition shared in subsets of body sites. Specifically, each row provides information about the network layers and their fitted sMLSBM model for a particular stratum. Each grid in the figure represents the binary adjacency matrix encoding interactions between OTUs: a colored dot at position (i, j) indicates the existence of an edge (i, j) in the corresponding network layer. In the first column of each row is a sample network generated with the learned SBM parameters of that stratum, $\bar{\pi}^s$ and \bar{Z}^s . Columns 2 and 3 show two representative network layers within the stratum. Note that while some strata have more than two members, for illustrative purposes we only show two example layers. It is easy to see the very similar block structure between all networks in a given row, corroborating the usefulness of the sMLSBM approach. Finally, we highlight the usefulness of fitting sMLSBM to this multilayer network as each stratum elucidates a mechanistic understanding of the relationship between groups of OTUs, which could inspire further biological understanding or inquiry.

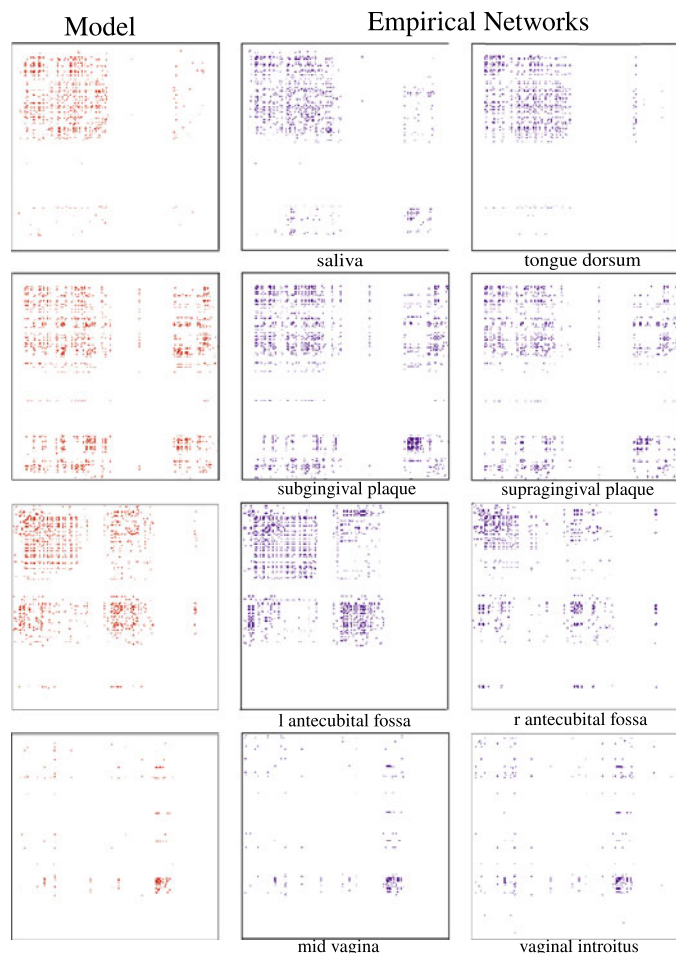


Fig. 6. *Visualization of Strata in SparCC Networks.* We visualize the adjacency matrices for SparCC networks that encode microbiome interactions at body sites. In each panel, a colored dot at position (i, j) indicates the existence of an edge (i, j) in the corresponding network layer. The four rows correspond to four different strata. In column 1, we show a sample network generated from the SBM parameters, $\bar{\pi}^s$ and \bar{Z}^s , that we inferred for that stratum. In Columns 2 and 3, we show SparCC networks from that particular stratum. Note the strong similarity across each row.

5 CONCLUSION AND FUTURE WORK

We developed a novel model for multilayer stochastic block models (MLSBMs) and an associated algorithm to jointly partition layers into strata and nodes into communities. Our model assumes that layers belonging to a stratum have community structure following the same underlying SBM. To fit sMLSBM to a multilayer network, and more specifically, a multiplex network, we iteratively alternate between rearranging layer-to-strata assignments and updating the model parameters for each stratum. Having multiple networks within a stratum—hence multiple realizations from some underlying model—helps to make inference more accurate. Particularly, more accurate assignments of nodes-to-communities within a stratum leads to improved estimation of SBM probability parameters, and vice versa. We have shown for multiplex networks with several strata (e.g., see Fig. 3) that inaccuracies can arise if one attempts to fit a single SBM to the network or study the network layers in isolation. In contrast, our model allows for an understanding of the similarities between layers in a network, in terms of their community structure.

The ability to identify strata within collections of network layers holds promise in numerous applications. One motivating application is network reducibility, whereby one compresses a multilayer network by aggregating similar layers [8]. We stress that although reducibility is a closely related pursuit, it is fundamentally different from our co-clustering pursuit of simultaneously identifying communities and strata. In particular, our approach does not provide a method for aggregating layers. Instead, sMLSBM compresses the network information in the sense that the learned SBM parameters represent a consensus for each stratum, and those consensus parameters can be used to generate a representative sample network for that stratum. For applications in which layer aggregation is sought, there are a variety of ways to aggregate layers in a strata. See, for example, Ref. [44], where the authors explore the effects on community structure for different aggregation methods. We highlight that the sMLSBM modeling approach is appropriate in situations where one seeks a generative model for community structure, and it may be particularly appropriate when application-specific evidence suggests that subsets of networks have characteristic differences in community structure.

Our comparison of sMLSBM to the reducibility method of Ref. [8] (see Fig. 5) for the application of studying microbial interaction networks reveals several extensions to sMLSBM that could make the approach more accurate and applicable to a wider range of applications. First, the reducibility method [8] does not require networks to be undirected and unweighted, and it could be quite useful to extend the sMLSBM framework to weighted and directed networks following the extensions for single-layer SBMs, as developed in [45] and [46], respectively. It would also be useful to extend to degree-corrected and overlapping (i.e., mixed-membership) communities [47], as well as mixed membership of layers into strata. Additionally, the Human Microbiome example reveals some interesting biological questions that could facilitate the development of more advanced network tools. To construct the multilayer network, negative edges were thresholded away; however, antagonistic relationships between microbes are known to be important [48]. Thus, it would be useful to develop a signed version of sMLSBM that allows edges to be either positive or negative.

The rise of a greater number of multilayer network datasets is providing the need for additional tools for the construction and analysis of such networks. The sMLSBM provides a new method to find signal in inherently noisy and complex network data.

ACKNOWLEDGMENTS

The authors thank James D. Wilson for helpful discussions about related work in multilayer networks, and in particular, multilayer stochastic block models. Research reported in this publication was supported by the Eunice Kennedy Shriver National Institute of Child Health & Human Development of the National Institutes of Health under Award Number R01HD075712, the James S. McDonnell Foundation 21st Century Science Initiative Complex Systems Scholar Award grant # 220020315, training grants T32GM067553 and T32CA201159 from the National Institutes of Health,

and the UNC Lineberger Comprehensive Cancer Center with funding provided by the University Cancer Research Fund (UCRF) via the State of North Carolina. The content is solely the responsibility of the authors and does not necessarily represent the official views of the funding agencies.

REFERENCES

- [1] M. Kivelä, A. Arenas, M. Barthélemy, J. P. Gleeson, Y. Moreno, and M. A. Porter, "Multilayer networks," *J. Complex Netw.*, vol. 2, no. 3, pp. 203–271, 2014.
- [2] S. Boccaletti, G. Bianconi, R. Criado, C. I. Del Genio, J. Gómez-Gardeñes, M. Romance, I. Sendina-Nadal, Z. Wang, and M. Zanin, "The structure and dynamics of multilayer networks," *Phys. Rep.*, vol. 544, no. 1, pp. 1–122, 2014.
- [3] M. De Domenico, A. Solé-Ribalta, E. Cozzo, M. Kivelä, Y. Moreno, M. A. Porter, S. Gómez, and A. Arenas, "Mathematical formulation of multilayer networks," *Phys. Rev. X*, vol. 3, no. 4, p. 041022, 2013.
- [4] M. Costanzo, A. Baryshnikova, J. Bellay, Y. Kim, E. D. Spear, C. S. Sevier, H. Ding, J. L. Koh, K. Toufighi, S. Mostafavi, et al., "The genetic landscape of a cell," *Science*, vol. 327, no. 5964, pp. 425–431, 2010.
- [5] D. Greene and P. Cunningham, "Producing a unified graph representation from multiple social network views," in *Proc. 5th Annu. ACM Web Sci. Conf.*, 2013, pp. 118–121.
- [6] P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, and J.-P. Onnela, "Community structure in time-dependent, multiscale, and multiplex networks," *Science*, vol. 328, no. 5980, pp. 876–878, 2010.
- [7] Q. Han, K. Xu, and E. Airoldi, "Consistent estimation of dynamic and multi-layer block models," in *Proc. 32nd Int. Conf. Mach. Learn.*, 2015, pp. 1511–1520.
- [8] M. De Domenico, V. Nicosia, A. Arenas, and V. Latora, "Structural reducibility of multilayer networks," *Nature Commun.*, vol. 6, pp. 1–9, 2015.
- [9] S. C. Madeira and A. L. Oliveira, "Biclustering algorithms for biological data analysis: A survey," *IEEE/ACM Trans. Comput. Biol. Bioinform.*, vol. 1, no. 1, pp. 24–45, Jan. 2004.
- [10] I. S. Dhillon, "Co-clustering documents and words using bipartite spectral graph partitioning," in *Proc. 7th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2001, pp. 269–274.
- [11] J. Ugander, L. Backstrom, and J. Kleinberg, "Subgraph frequencies: Mapping the empirical and extremal geography of large graph collections," in *Proc. 22nd Int. Conf. World Wide Web*, 2013, pp. 1307–1318.
- [12] K. Tsuda and T. Kudo, "Clustering graphs by weighted substructure mining," in *Proc. 23rd Int. Conf. Mach. Learn.*, 2006, pp. 953–960.
- [13] J.-P. Onnela, D. J. Fenn, S. Reid, M. A. Porter, P. J. Mucha, M. D. Fricker, and N. S. Jones, "Taxonomies of networks from community structure," *Phys. Rev. E*, vol. 86, no. 3, p. 036104, 2012.
- [14] J. Ni, H. Tong, W. Fan, and X. Zhang, "Flexible and robust multi-network clustering," in *Proc. 21th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2015, pp. 835–844.
- [15] J. Iacovacci, Z. Wu, and G. Bianconi, "Mesoscopic structures reveal the network between the layers of multiplex datasets," *arXiv preprint arXiv:1505.03824*, 2015.
- [16] U. Brandes, J. Lerner, U. Nagel, and B. Nick, "Structural trends in network ensembles," in *Complex Networks*. New York, NY, USA: Springer, 2009, pp. 83–97.
- [17] U. Brandes, J. Lerner, and U. Nagel, "Network ensemble clustering using latent roles," *Adv. Data Anal. Classification*, vol. 5, no. 2, pp. 81–94, 2011.
- [18] M. A. Porter, J.-P. Onnela, and P. J. Mucha, "Communities in networks," *Notices AMS*, vol. 56, no. 9, pp. 1082–1097, 2009.
- [19] S. Fortunato, "Community detection in graphs," *Phys. Rep.*, vol. 486, no. 3, pp. 75–174, 2010.
- [20] M. P. Rombach, M. A. Porter, J. H. Fowler, and P. J. Mucha, "Core-periphery structure in networks," *SIAM J. Appl. Math.*, vol. 74, no. 1, pp. 167–190, 2014.
- [21] M. E. Newman, "Modularity and community structure in networks," *Proc. Nat. Acad. Sci.*, vol. 103, no. 23, pp. 8577–8582, 2006.
- [22] A. Z. Jacobs and A. Clauset, "A unified view of generative models for networks: Models, methods, opportunities, and challenges," *arXiv preprint arXiv:1411.4070*, 2014.

- [23] M. A. Porter, J.-P. Onnela, and P. J. Mucha, "Communities in networks," *Notices AMS*, vol. 56, no. 9, pp. 1082–1097, 2009.
- [24] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney, "Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters," *Internet Math.*, vol. 6, no. 1, pp. 29–123, 2009.
- [25] A. Clauset, C. Moore, and M. E. Newman, "Structural inference of hierarchies in networks," in *Statistical Network Analysis: Models, Issues, and New Directions*. New York, NY, USA: Springer, 2007, pp. 1–13.
- [26] M. E. J. Newman. (2006, Sep.). Finding community structure in networks using the eigenvectors of matrices. *Phys. Rev. E* [Online]. 74, p. 036104. Available: <http://link.aps.org/doi/10.1103/PhysRevE.74.036104>
- [27] T. A. Snijders and K. Nowicki, "Estimation and prediction for stochastic blockmodels for graphs with latent block structure," *J. Classification*, vol. 14, no. 1, pp. 75–100, 1997.
- [28] C. Aicher, A. Z. Jacobs, and A. Clauset, "Learning latent block structure in weighted networks," *J. Complex Netw.*, vol. 3, no. 2, pp. 221–248, 2015.
- [29] R. Guimerà and M. Sales-Pardo, "Missing and spurious interactions and the reconstruction of complex networks," *Proc. Nat. Acad. Sci.*, vol. 106, no. 52, pp. 22 073–22 078, 2009.
- [30] M. De Domenico, A. Lancichinetti, A. Arenas, and M. Rosvall, "Identifying modular flows on multilayer networks reveals highly overlapping organization in interconnected systems," *Phys. Rev. X*, vol. 5, no. 1, p. 011027, 2015.
- [31] S. Paul and Y. Chen, "Community detection in multi-relational data with restricted multi-layer stochastic blockmodel," *arXiv preprint arXiv:1506.02699*, 2015.
- [32] P. Barbillon, S. Donnet, E. Lazega, and A. Bar-Hen, "Stochastic block models for multiplex networks: An application to networks of researchers," *arXiv preprint arXiv:1501.06444*, 2015.
- [33] T. Valles-Catala, F. A. Massucci, R. Guimerà, and M. Sales-Pardo, "Stochastic block models reveal the multilayer structure of complex networks," *arXiv preprint arXiv:1411.1098*, 2014.
- [34] T. P. Peixoto. (2015, Oct.). Inferring the mesoscale structure of layered, edge-valued, and time-varying networks. *Phys. Rev. E* [Online]. 92, p. 042807. Available: <http://link.aps.org/doi/10.1103/PhysRevE.92.042807>
- [35] J.-J. Daudin, F. Picard, and S. Robin, "A mixture model for random graphs," *Statist. Comput.*, vol. 18, no. 2, pp. 173–183, 2008.
- [36] R. Tibshirani, G. Walther, and T. Hastie, "Estimating the number of clusters in a data set via the gap statistic," *J. Roy. Statist. Soc.: Series B (Statist. Methodol.)*, vol. 63, no. 2, pp. 411–423, 2001.
- [37] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Roy. Statist. Soc. Series B (Methodological)*, vol. 39, no. 1, pp. 1–38, 1977.
- [38] L. Danon, A. Diaz-Guilera, J. Duch, and A. Arenas, "Comparing community structure identification," *J. Statist. Mech.: Theory Exp.*, vol. 2005, no. 09, p. P09008, 2005.
- [39] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová, "Inference and phase transitions in the detection of modules in sparse networks," *Phys. Rev. Lett.*, vol. 107, no. 6, p. 065701, 2011.
- [40] P. J. Turnbaugh, R. E. Ley, M. Hamady, C. M. Fraser-Liggett, R. Knight, and J. I. Gordon, "The human microbiome project," *Nature*, vol. 449, no. 7164, pp. 804–810, 2007.
- [41] K. Faust, J. F. Sathirapongsasuti, J. Izard, N. Segata, D. Gevers, J. Raes, and C. Huttenhower, "Microbial co-occurrence relationships in the human microbiome," *PLoS Comput. Biol.*, vol. 8, no. 7, p. e1002606, 2012.
- [42] J. Friedman and E. J. Alm, "Inferring correlation networks from genomic survey data," *PLoS Comput. Biol.*, vol. 8, no. 9, p. e1002687, 2012.
- [43] T. Ding and P. D. Schloss, "Dynamics and associations of microbial community types across the human body," *Nature*, vol. 509, no. 7500, pp. 357–360, 2014.
- [44] D. Taylor, S. Shai, N. Stanley, and P. J. Mucha, "Enhanced detectability of community structure in multilayer networks through layer aggregation," *arXiv preprint arXiv:1511.05271*, 2015.
- [45] C. Aicher, A. Z. Jacobs, and A. Clauset, "Learning latent block structure in weighted networks," *J. Complex Netw.*, vol. 3, no. 2, pp. 221–248, 2015.
- [46] Y. J. Wang and G. Y. Wong, "Stochastic blockmodels for directed graphs," *J. Am. Statist. Assoc.*, vol. 82, no. 397, pp. 8–19, 1987.
- [47] B. Karrer and M. E. Newman, "Stochastic blockmodels and community structure in networks," *Phys. Rev. E*, vol. 83, no. 1, p. 016107, 2011.
- [48] R. Zapién-Campos, G. Olmedo-Álvarez, and M. Santillán, "Antagonistic interactions are sufficient to explain self-assembly of bacterial communities in a homogeneous environment: A computational modeling approach," *Frontiers Microbiol.*, vol. 6, p. 489, 2015.

Natalie Stanley is a PhD student in the bioinformatics and computational biology program at the University of North Carolina at Chapel Hill. She is working under the supervision of Professor Peter Mucha. Her research interests are in the application of networks, community detection, and machine learning to computational biology problems.

Sarai Shai received a dual BS degree in mathematics and computer science from the Israel Institute of Technology (Technion) in 2008 and earned a PhD degree in computer science from the University of St Andrews in 2014. Her background includes experience as a software developer at an Israeli start-up company. She is currently a postdoctoral research associate in the Department of Mathematics at the University of North Carolina at Chapel Hill. Her research interests focus on developing mathematical and computational tools for analyzing complex networks.

Dane Taylor obtained his PhD in Applied Mathematics from the University of Colorado-Boulder in 2013, where he studied nonlinear dynamics on networks. His research ranges from the optimization of oscillatory systems such as the power grid to have enhanced synchronization properties to the modeling of spreading processes on social networks and transit systems. As a postdoctoral scholar in the UNC Department of Mathematics, Dane is currently working under Peter J. Mucha to develop networks-based methodology to analyze datasets arising from social and biological systems.

Peter J. Mucha received the PhD degree in applied and computational mathematics in 1998 from Princeton University. He was an applied mathematics instructor at the Massachusetts Institute of Technology (MIT) for three years, followed by four years as an assistant professor in mathematics at the Georgia Institute of Technology. He is currently a Bowman and Gordon Gray Distinguished Term Professor in the Departments of Mathematics and Applied Physical Sciences at the University of North Carolina at Chapel Hill. His research interests include networks, complex systems, and data science.

► **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**