

LATENT HETEROGENEOUS MULTILAYER COMMUNITY DETECTION

Hafiz Tiomoko Ali¹, Sijia Liu⁵, Yasin Yilmaz², Romain Couillet⁴, Indika Rajapakse³, Alfred Hero³

¹ Huawei Noah’s Ark Lab, ² University of South Florida, ³ University of Michigan
⁴ GIPSA-lab, University Grenoble-Alpes, ⁵ MIT-IBM Watson AI Lab, IBM Research.

ABSTRACT

We propose a method for simultaneously detecting shared and unshared communities in heterogeneous multilayer weighted and undirected networks. The multilayer network is assumed to follow a generative probabilistic model that takes into account the similarities and dissimilarities between the communities. We make use of a variational Bayes approach for jointly inferring the shared and unshared hidden communities from multilayer network observations. We show that our approach outperforms state-of-the-art algorithms in detecting disparate (shared and private) communities on synthetic data as well as on real genome-wide fibroblast proliferation dataset.

Index Terms— multilayer networks, community detection, heterogeneous communities, variational Bayes.

I. INTRODUCTION

Community detection in networks is an ubiquitous area of active research as it enables the exploration of network structural properties in real-world scenarios including, but not limited to, social and biological sciences [1]. There is a vast literature on single-layer network community detection using modularity optimization [2], spectral clustering [3] and statistical inference [4]. However, single-layer networks are not well suited to real-world networks, such as the Internet of Things (IoT), transportation, social, and biological networks where multi-relational, multidimensional, multiplex or multilayer networks structures exist [5]. In social networks, for example, large company employees may be related to each other by similarity of their activities (functional relations) within the organization, and also by sharing the same office location (spatial relations) [6]. As another example, in genomics, genes might be related either by transcriptional interactions (functional relations), e.g., measured by RNA-seq profile similarity, or by chromatin interactions (spatial relations), e.g., measured by chromatin conformation capture (Hi-C) of promoter-enhancer ligations [7]. In each of these examples, there may be a community structure that is common to each layer and a structure that is distinct between layers. Such multilayer networks create a need for new community detection methods [8]. Recovering the communities from each network independently is suboptimal as this strategy does not exploit the shared information across the network layers. Thus, current research efforts aim at developing joint inference methods by multilayer aggregation [9]. The simplest aggregation approach is to collapse the multilayer network to a single-layer network on which classical community detection methods can be applied [10]–[12]. Alternatively, some researchers have suggested performing community detection separately in each layer followed by consensus aggregation of the communities across layers [6], [12], [13]. Another approach is to extend single-layer Stochastic Block Models [14] to multilayer networks [15]–[17].

In many applications, some communities might be shared between the different layers, while others not (see Figure 1). However, few

This work was partially supported by DARPA under the Deep Purple and FunCC program and by ARO under grant W911NF-15-1-0479. Couillet’s work is supported by the GSTATS UGA IDEX DataScience chair and the ANR RMT4GRAPH Project (ANR-14-CE28-0006).

methods in the literature explicitly consider this general scenario. The Multilayer Extraction algorithm proposed in [18] allows for the identification of heterogeneous multilayer communities where the communities might be shared between a subset of layers. The authors in [18] minimize a cost function and take into account similarities and dissimilarities between the layers’ communities. While the model used in [18] is realistic when considering multilayer graph connections, the method is only limited to unweighted graphs. The approach proposed in [19] extends [20] to weighted multilayer graphs and allows the extraction of coherent dense subgraphs (cliques) shared by subsets of layers. However, those methods are limited to the identification of dense communities and might fail when the communities are connected but with only a few edges.

In the present article, we propose a new model-based method to simultaneously detect shared and unshared communities between heterogeneous weighted networks. We define joint weighted stochastic block models (WSBM) that share “a part” of their community structures. We develop a mean field variational Bayes approach to infer the latent shared and private communities from the proposed multilayer WSBM. This extends the works in [21], [22] devised for WSBM in single-layer graphs. Our main contributions are

- We derive a variational Bayes algorithm for automatically inferring shared and unshared communities from multilayer weighted graphs.
- We establish that the proposed algorithm is more accurate and robust than previous approaches to community detection in multi-layer networks in extracting both shared and unshared communities from weighted graph benchmarks.
- We illustrate a real-world use of our method in multinomic molecular biology where it enables the discovery of heterogeneous multilayer communities of gene-gene interactions in human fibroblast proliferation.

Notations: Vectors are denoted with boldface lowercase letters and matrices by boldface uppercase letters. $\mathbf{I}_n \in \mathbb{R}^{n \times n}$ is the identity matrix and $\mathbf{1}_n$ is the column vector full of ones.

II. JOINT WEIGHTED STOCHASTIC BLOCK MODELS

To start with, let us recall the definition of a single-layer Stochastic Block Model generated by a weighted distribution \mathcal{D} with sufficient statistic T and natural parameter η . Given *latent* community labels $g_i \in \{1, \dots, K\}$ (with K denoting the number of communities) of each vertex i and a community-wise connectivity matrix $\theta \in \mathbb{R}^{K \times K}$, an edge is placed between two vertices i and j with an adjacency weight A_{ij} such that

$$\mathbb{P}(A_{ij} | g_i, g_j, \theta_{g_i, g_j}) \propto \exp \{T(A_{ij})\eta(\theta_{g_i, g_j})\}.$$

Following a Bayesian approach, prior distributions are attributed to the labels g_i and the community-wise connectivity matrix θ .

We denote a multilayer graph, \mathcal{G} , defining L as the number of layers and n as the number of vertices. The graph in the l -th layer is an undirected (possibly weighted) graph $\mathcal{G}^{(l)} = (\mathcal{V}, \mathcal{E}^{(l)})$ with \mathcal{V} denoting the set of common vertices and $\mathcal{E}^{(l)}$ denoting the set of edges in graph $\mathcal{G}^{(l)}$. We denote by $\mathbf{A}^{(l)}$ the adjacency matrices containing the edge weights between each pair of vertices in graph

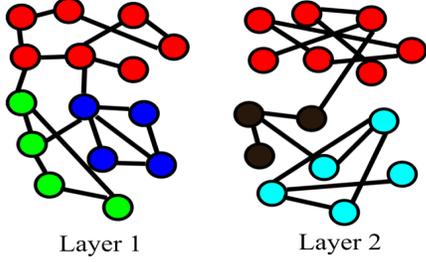


Fig. 1: Heterogeneous multilayer network. Shared communities (in red) and unshared communities in different colors for each layer.

$\mathcal{G}^{(l)}$. We propose the following generative heterogeneous community structure of the multilayer graph \mathcal{G} .

- 1) We assume that each layer is subdivided into $K^{(l)}$ non-overlapping communities among which the first $K \leq \min_l K^{(l)}$ are shared between the layers as described below.
- 2) We first generate the label $g_i^{(1)}$ of each vertex i in the first layer as $g_i^{(1)} \sim \text{Categorical}(\boldsymbol{\mu}_0^{(1)})$, where $\boldsymbol{\mu}_0^{(1)} \in \mathbb{R}^{K^{(1)}}$ contains prior probabilities that the vertices belong to one of the $K^{(1)}$ communities.
- 3) For each vertex i , if $g_i^{(1)} \in \{1, \dots, K\}$ then set $g_i^{(l)} = g_i^{(1)}$ for each layer l . Otherwise, generate for each layer l , $g_i^{(l)} \sim \text{Categorical}(\boldsymbol{\mu}_0^{(l)})$.
- 4) Given *latent* community labels $g_i^{(l)}$ (generated in steps 2 and 3) of each vertex i and community-wise connectivity matrices $\boldsymbol{\theta}^{(l)} \in \mathbb{R}^{K^{(l)} \times K^{(l)}}$ (the generation of which will be defined later), an edge is placed between two vertices i and j and is assigned an adjacency weight $A_{ij}^{(l)}$ drawn according to

$$\mathbb{P}(A_{ij}^{(l)} | g_i^{(l)}, g_j^{(l)}, \boldsymbol{\theta}_{g_i^{(l)} g_j^{(l)}}^{(l)}) \propto \exp \left\{ T^{(l)}(A_{ij}^{(l)}) \eta^{(l)}(\boldsymbol{\theta}_{g_i^{(l)} g_j^{(l)}}^{(l)}) \right\}, \quad (1)$$

where $T^{(l)}$ is the sufficient statistic and $\eta^{(l)}$ is the natural parameter of the weights distribution.

- 5) The community-wise connectivity matrices $\boldsymbol{\theta}^{(l)}$ are generated according to conjugate priors associated with the distribution characterized by $(T^{(l)}, \eta^{(l)})$, i.e.,

$$p^*(\boldsymbol{\theta}_{ab}^{(l)}) = \frac{1}{Z^{(l)}(\tau_0^{(l)})} \exp(\tau_0^{(l)} \eta^{(l)}(\boldsymbol{\theta}_{ab}^{(l)}))$$

with $\tau_0^{(l)}$ denoting the associated hyperparameters and $Z^{(l)}(\tau_0^{(l)})$ the normalization constants.

For illustration, we specialize the presentation to two communities, for which each of the matrices $\boldsymbol{\theta}^{(l)}$ are decomposed into four blocks corresponding, respectively, to the shared-shared, shared-private, private-shared, private-private interconnections (see Figure 2). As in [21], we consider each sub-matrix $\boldsymbol{\theta}_1^{(l)}, \boldsymbol{\theta}_2^{(l)}, \boldsymbol{\theta}_3^{(l)}, \boldsymbol{\theta}_4^{(l)}$ as one-dimensional vectors where the elements are stacked. Let us denote by $r_1^{(l)}, r_2^{(l)}, r_3^{(l)}, r_4^{(l)}$ the indexing variables into each of the obtained vectors, i.e., $r_1^{(l)} = 1, \dots, K^2; r_2^{(l)} = 1, \dots, K(K^{(l)} - K); r_3^{(l)} = 1, \dots, K(K^{(l)} - K); r_4^{(l)} = 1, \dots, (K^{(l)} - K)^2$. The overall prior distribution can thus be written as

$$p^*(\mathbf{g}^{(l)}, \boldsymbol{\theta}^{(l)}, l = 1, \dots, L) = \prod_{l=1}^L \prod_i (\boldsymbol{\mu}_0^{(l)})_{i, g_i^{(l)}} \times \prod_{r^{(l)}} \frac{1}{Z^{(l)}(\tau_0^{(l)})} \exp(\tau_0^{(l)} \eta^{(l)}(\boldsymbol{\theta}_{r^{(l)}}^{(l)})) \quad (2)$$



Fig. 2: Community affinity probability matrices $\boldsymbol{\theta}^{(1)}$ and $\boldsymbol{\theta}^{(2)}$ decomposed into blocks. Red blocks (common to $\mathbf{A}^{(1)}$ and $\mathbf{A}^{(2)}$) used to update shared communities while blue blocks used to update private communities of $\mathbf{A}^{(1)}$ and green blocks for the update of $\mathbf{A}^{(2)}$'s private communities.

with $r^{(l)} \equiv \{r_1^{(l)}, r_2^{(l)}, r_3^{(l)}, r_4^{(l)}\}$.

Given $\mathcal{G}^{(l)}, l = 1, \dots, L$ (equivalently their adjacency matrices $\mathbf{A}^{(l)}$), the goal is to infer the community labels $g_i^{(l)}$ for each node i in each layer l , i.e., to find the most probable clustering $\mathbf{g}^{(l)}$ of the vertices in each layer in the set of all different possible partitionings

$$[(\mathbf{g}^{(1)})^*, \dots, (\mathbf{g}^{(L)})^*] = \underset{\mathbf{g}^{(l)}, l=1, \dots, L}{\operatorname{argmax}} \mathbb{P}(\mathbf{g}^{(l)} | \mathbf{A}^{(l)}, \boldsymbol{\theta}^{(l)}, l = 1, \dots, L) \quad (3)$$

with the correlations constraints on $\mathbf{g}^{(l)}$ defined in Point 3) of Section II. The optimization problem (3) is NP-hard due to two main difficulties: the maximization is over all possible configurations of $\mathbf{g}^{(l)}$, and the calculation of the posterior distribution $\mathbb{P}(\mathbf{g}^{(l)} | \mathbf{A}^{(l)}, \boldsymbol{\theta}^{(l)})$, which is intractable due to its high dimensional integral form. Our approach to the optimization (3) is the mean field variational Bayes approximation [23], [24] that uses a factorisable distribution as an approximation to the joint posterior $p(\mathbf{g}^{(l)}, \boldsymbol{\theta}^{(l)}) \equiv \mathbb{P}(\mathbf{g}^{(l)}, \boldsymbol{\theta}^{(l)} | \mathbf{A}^{(l)})$.

III. VARIATIONAL INFERENCE

III-A. Mean field variational Bayes inference

Denote by $q(\mathbf{g}^{(l)}, \boldsymbol{\theta}^{(l)})$ an approximating (factorisable) distribution that depends on tunable shaping parameters $\boldsymbol{\mu}^{(l)}$ and $\boldsymbol{\tau}^{(l)}$. The variational Bayes algorithm fits the distribution q to the joint distribution by minimizing the KL-divergence, i.e., $q = \operatorname{argmin}_r D_{KL}(r || p)$.

Here the distribution q is taken to have the same parametric form as the prior p^*

$$q(\mathbf{g}^{(l)}, \boldsymbol{\theta}^{(l)}) = \prod_{l=1}^L \prod_i \boldsymbol{\mu}_{i, g_i^{(l)}}^{(l)} \prod_{r^{(l)}} \frac{1}{Z^{(l)}(\tau_{r^{(l)}}^{(l)})} \exp(\tau_{r^{(l)}}^{(l)} \eta^{(l)}(\boldsymbol{\theta}_{r^{(l)}}^{(l)})), \quad (4)$$

where $\tau_{r^{(l)}}^{(l)}$ and $\boldsymbol{\mu}^{(l)} \in \mathbb{R}^{n \times K^{(l)}}$ are variational parameters corresponding to the random variables $\boldsymbol{\theta}_{r^{(l)}}^{(l)}, \mathbf{g}^{(l)}$ respectively. We can rewrite the original problem (3) as follows

$$\begin{aligned} [(\mathbf{g}^{(1)})^*, \dots, (\mathbf{g}^{(L)})^*] &= \operatorname{argmax}_{\mathbf{g}^{(l)}} \int \mathbb{P}(\mathbf{g}^{(l)}, \boldsymbol{\theta}^{(l)} | \mathbf{A}^{(l)}) d\boldsymbol{\theta}^{(l)} \\ &\approx \operatorname{argmax}_{\mathbf{g}^{(l)}} \int q(\mathbf{g}^{(l)}, \boldsymbol{\theta}^{(l)}) d\boldsymbol{\theta}^{(l)} \\ &= \operatorname{argmax}_{\mathbf{g}^{(l)}} \int \prod_l \prod_i q(g_i^{(l)}) q(\boldsymbol{\theta}^{(l)}) d\boldsymbol{\theta}^{(l)} \\ &= \operatorname{argmax}_{\mathbf{g}^{(l)}} \prod_l \prod_i q(g_i^{(l)}). \end{aligned}$$

Since $q^{(l)}$ is a categorical distribution with parameter $\boldsymbol{\mu}^{(l)}$, the original problem (3) is equivalent to

$$(g_i^{(l)})^* = \operatorname{argmax}_k \boldsymbol{\mu}_{ik}^{(l)} \quad (5)$$

for each node i and layer l and thus the multilayer community detection boils down to a Maximum A Posteriori (MAP) estimator on each nodal variational parameter $\mu_i^{(l)}$ for each layer l .

III-B. Learning

As per [21], the constant model likelihood can be written as $\log \mathbb{P}(\mathbf{A}^{(l)}) = \mathcal{G}(q) + D_{KL}(q||p)$ with

$$\mathcal{G}(q) = \mathbb{E}_q \log \mathbb{P}(\mathbf{A}^{(l)} | \mathbf{g}^{(l)}, \boldsymbol{\theta}^{(l)}, l = 1, 2) + \mathbb{E}_q \frac{p^*}{q}, \quad (6)$$

where p^* is the prior distribution assigned to the parameters $\mathbf{g}^{(l)}, \boldsymbol{\theta}^{(l)}$. Since the likelihood is constant, minimizing $D_{KL}(q||p)$ (and thus making the approximation q to be the closest to the sought posterior p) is equivalent to maximizing $\mathcal{G}(q)$ over the variational parameters. In the sequel, we devise a procedure to learn the parameters for which $\mathcal{G}(q)$ is maximized.

We next address how to find the variational parameters $\tau^{(l)}, \boldsymbol{\mu}^{(l)}$ for which $\mathcal{G}(q)$ is maximized. To this end, let us first compute $\mathcal{G}(q)$ with the forms of the prior p^* and the approximation q defined in the previous section. For illustration, we specialize to $L = 2$ but the same principle applies to any number of layers. We have

$$\begin{aligned} \mathcal{G}(q) &= \mathbb{E}_q \log \mathbb{P}(\mathbf{A}^{(1)}, \mathbf{A}^{(2)} | \mathbf{g}^{(1)}, \mathbf{g}^{(2)}, \boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}) + \mathbb{E}_q \frac{p^*}{q} \\ &= \mathbb{E}_q \log \mathbb{P}(\mathbf{A}^{(1)} | \mathbf{g}^{(1)}, \mathbf{g}^{(2)}, \boldsymbol{\theta}^{(1)}) \\ &\quad + \mathbb{E}_q \log \mathbb{P}(\mathbf{A}^{(2)} | \mathbf{g}^{(1)}, \mathbf{g}^{(2)}, \boldsymbol{\theta}^{(2)}) + \mathbb{E}_q \frac{p^*}{q}, \end{aligned} \quad (7)$$

where in the last line, we use the chain rule along with the conditional independence between $\mathbf{A}^{(1)}$ and $\mathbf{A}^{(2)}$ given $\mathbf{g}^{(1)}, \mathbf{g}^{(2)}, \boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}$. The structure of the heterogeneous Joint Stochastic Block Model (Section II) couples the random variables $\mathbf{g}^{(1)}$ and $\mathbf{g}^{(2)}$ in a simple manner that can be decomposed into the following cases, which we call the *dependency cases*:

- For a vertex pair (i, j) belonging to a block with $\boldsymbol{\theta}_1^{(l)}, g_i^{(1)} = g_i^{(2)}$ and $g_j^{(1)} = g_j^{(2)}$.
- For a vertex pair (i, j) belonging to a block with $\boldsymbol{\theta}_2^{(l)}, g_i^{(1)} = g_i^{(2)}$ and $g_j^{(1)} \neq g_j^{(2)}$.
- For a vertex pair (i, j) belonging to a block with $\boldsymbol{\theta}_3^{(l)}, g_i^{(1)} \neq g_i^{(2)}$ and $g_j^{(1)} = g_j^{(2)}$.
- For a vertex pair (i, j) belonging to a block with $\boldsymbol{\theta}_4^{(l)}, g_i^{(1)} \neq g_i^{(2)}$ and $g_j^{(1)} \neq g_j^{(2)}$.

Using these *dependency cases* with (1), we obtain an expression for $\mathcal{G}(q)$ as per (7). After differentiation with respect to the sought variational parameters, we obtain updates for $\tau^{(l)}, \boldsymbol{\mu}^{(l)}$, which are stationary points of $\mathcal{G}(q)$ and correspond to local maxima. The precision of the local maxima depend on the initial values for $\boldsymbol{\mu}^{(l)}$. A single run of a single-layer clustering algorithm shall lead to the optimal solution basin of attraction. Due to the *dependency cases*, the community memberships variational parameters $\mu_{ik}^{(l)}$ depends on $g_i^{(l)}$ either belonging to the set of shared communities $\{1, \dots, K\}$, or to the set of unshared communities $(\{K+1, \dots, K^{(l)}\})$.

Algorithm 1 provides the necessary equations for the updates of the variational parameters $\tau^{(l)}$ and $\boldsymbol{\mu}^{(l)}$. Due to Equation (5), a max decision rule can then be used on $\boldsymbol{\mu}^{(l)}$ to assign labels to each node, namely $\operatorname{argmax}_k \mu_{ik}^{(l)}$ gives the label assigned to node i in graph $\mathcal{G}^{(l)}$. The label of node i is shared between different graphs $\{\mathcal{G}^{(l)}\}$ when $\operatorname{argmax}_k \mu_{ik}^{(l)} \in \{1, \dots, K\}$ and the label is unshared otherwise.

Algorithm 1 is an extension of the variational Bayes algorithm for inferring hidden communities from single-layer graphs [21], [22] to the inference of hidden *shared* and *unshared* communities from

multilayer graphs. In Algorithm 1, the updates for the parameters $\tau^{(l)}$ are done independently for each graph as in [21]. As for the community membership variational parameters $\boldsymbol{\mu}^{(l)} \in \mathbb{R}^{n \times K^{(l)}}$, the updates of the first K columns of $\boldsymbol{\mu}^{(l)}$ are identical and are computed by adding the contributions of each graph. The last $K^{(l)} - K$ columns of $\boldsymbol{\mu}^{(l)}$ are updated independently using only the information about each graph. This is quite intuitive since the first K columns of $\boldsymbol{\mu}^{(l)}$ correspond to the shared community evidences and thus they should be updated using the contributions of the graphs altogether, while the last columns correspond to unshared communities and thus the updates should be done independently for each graph.

IV. EXPERIMENTS

IV-A. Synthetic graphs

We first consider two Bernoulli SBM graphs $\mathcal{G}^{(1)}$ and $\mathcal{G}^{(2)}$ with the same intra-community probabilities and different inter-community probabilities in such a way that one graph is noisier than the other. Blindly identifying the community labels from each of the graphs would yield poor performances since we do not know in advance which graph has a clearer community structure than the other. $\mathcal{G}^{(1)}$ and $\mathcal{G}^{(2)}$ are constructed with $n = 500$ vertices, each partitioned into $K^{(l)} = 4$ ($l = 1, 2$) communities respectively among which $K = 2$ are shared between the two graphs. The community labels $g_i^{(l)}$ are assigned uniformly over the intervals $\{1, \dots, K^{(l)}\}$ in such a way that $g_i^{(1)} = g_i^{(2)}$ when $g_i^{(1)} \in \{1, \dots, K\}$. Given the community labels $g_i^{(l)}$, the entries of the adjacency matrices are generated as $A_{ij}^{(l)} \sim \text{Bernoulli}(\boldsymbol{\theta}^{(l)})$ with $\boldsymbol{\theta}^{(1)} = (p-q)\mathbf{I}_4 + q\mathbf{1}_4\mathbf{1}_4^\top$ and $\boldsymbol{\theta}^{(2)} = (p-q')\mathbf{I}_4 + q'\mathbf{1}_4\mathbf{1}_4^\top$, where we fix $p = 0.6$, $q = 0.2$ and we vary q' between 0.2 and 0.5. The larger q' is, the noisier the graph $\mathcal{G}^{(2)}$ is in comparison with $\mathcal{G}^{(1)}$ and the more difficult community recovery is when applying a community detection algorithm on $\mathcal{G}^{(2)}$ solely. The left figure in Figure 3 shows that our Joint mean field (MF) algorithm outperforms the competing method Multilayer Extraction [18] (M-E) in extracting both shared and unshared communities from the two correlated graphs. Both methods are designed to exploit the graph $\mathcal{G}^{(1)}$ to identify the communities of $\mathcal{G}^{(2)}$ (among which some are shared with $\mathcal{G}^{(1)}$). Both joint methods significantly outperform a spectral clustering (SC) algorithm and a mean field variational algorithm applied on $\mathcal{G}^{(2)}$ alone.

We next consider the same graph settings as before but with disparate distributions $A_{ij}^{(1)} \sim \text{Bernoulli}(\boldsymbol{\theta}^{(1)})$ and $A_{ij}^{(2)} \sim \text{Poisson}(\boldsymbol{\theta}^{(2)})$. Here the M-E algorithm is not exploitable since the latter is designed only for binary graphs (Bernoulli). Our joint variational algorithm is thus compared with single-graph clustering algorithms. The results are reported in the right figure of Figure 3 where our joint algorithm outperforms single-layer clustering algorithms (spectral clustering and mean field variational Bayes).

IV-B. Real world graphs

In this section, we make use of our novel approach to understand the interplay between genome structure (form) and transcription (function) based on a human fibroblast proliferation dataset [25]. This dataset consists of Hi-C contact maps [26] that capture chromatin architectures and RNA-seq data that provide gene expression levels over 8 time points. For each chromosome (1–22), we first build a correlation matrix between the RNA-seq values, where thresholding is applied to obtain a binary adjacency matrix $\mathbf{A}^{(1)}$ representing functional correspondence between different genes. The threshold was chosen as the mean of the entries of the correlation matrix. We then construct an average (over the 8 time points) Hi-C matrix $\mathbf{A}^{(2)}$ and round each entry of the average matrix

Algorithm 1 Mean field inference of heterogeneous communities in multilayer graphs.

Inputs: For $l = 1, \dots, L$, layers adjacencies $\mathbf{A}^{(l)}$, layer distributions $\mathcal{D}^{(l)} = [T^{(l)}, \eta^{(l)}, Z^{(l)}]$, number of shared communities K , total number $K^{(l)}$ of communities.

Output: $\mu^{(1)}, \dots, \mu^{(L)}$.

For $l = 1, \dots, L$, initialize $\mu^{(l)}$ and choose hyperparameters $\tau_0^{(l)}$.

repeat

for $l = 1$ to L **do**
for $r^{(l)} = 1$ to $(K^{(l)})^2$ **do**

$$\tau_{r^{(l)}}^{(l)} = \tau_0^{(l)} + \sum_{ij} \sum_{(g_i^{(l)}, g_j^{(l)})=r^{(l)}} T^{(l)}(A_{ij}^{(l)}) \mu_{i, g_i^{(l)}}^{(l)} \mu_{j, g_j^{(l)}}^{(l)}.$$

end for

end for

repeat

for $i = 1$ to n **do**
for $k = 1$ to K **do**

$$\begin{aligned} \mu_{ik}^{(1)} = & \exp \left\{ \frac{1}{L} \sum_{l=1}^L \left[\sum_{\substack{r_1^{(l)}, j \neq i \\ (k, g_j^{(l)})=r_1^{(l)}}} T^{(l)}(A_{ij}^{(l)}) \mu_{j, g_j^{(l)}}^{(l)} \bar{\eta}_{r_1^{(l)}}^{(l)} \right. \right. \\ & \left. \left. + \sum_{\substack{r_2^{(l)}, j \neq i \\ (k, g_j^{(l)})=r_2^{(l)}}} T^{(l)}(A_{ij}^{(l)}) \mu_{j, g_j^{(l)}}^{(l)} \bar{\eta}_{r_2^{(l)}}^{(l)} \right] \right\} \end{aligned}$$

$$\bar{\eta}_{r^{(l)}}^{(l)} = \frac{\partial \log Z^{(l)}}{\partial \tau^{(l)}} \Big|_{\tau^{(l)}=\tau_{r^{(l)}}^{(l)}}$$

end for

for $l = 1$ to L **do**

$$\mu_{i, 1:K}^{(l)} = \mu_{i, 1:K}^{(1)}$$

for $k = K+1$ to $K^{(l)}$ **do**

$$\begin{aligned} \mu_{ik}^{(l)} = & \exp \left\{ \sum_{\substack{r_3^{(l)}, j \neq i \\ (k, g_j^{(l)})=r_3^{(l)}}} T^{(l)}(A_{ij}^{(l)}) \mu_{j, g_j^{(l)}}^{(l)} \bar{\eta}_{r_3^{(l)}}^{(l)} \right. \\ & \left. + \sum_{\substack{r_4^{(l)}, j \neq i \\ (k, g_j^{(l)})=r_4^{(l)}}} T^{(l)}(A_{ij}^{(l)}) \mu_{j, g_j^{(l)}}^{(l)} \bar{\eta}_{r_4^{(l)}}^{(l)} \right\} \end{aligned}$$

end for

$$\text{For } k = 1, \dots, K^{(l)}, \mu_{i,k}^{(l)} = \mu_{i,k}^{(l)} / \sum_{k=1}^{K^{(l)}} \mu_{i,k}^{(l)}.$$

end for

end for

until convergence

until convergence

to the closest integer value. For the application of the variational Bayes algorithm, the entries of $\mathbf{A}^{(1)}$ are considered to be Bernoulli distributed while those of $\mathbf{A}^{(2)}$ are considered to be Poisson distributed. More sophisticated models for the sample correlation graph, e.g., Wishart distributions, could also be considered but this is left for future work.

We make use of our joint variational algorithm to automatically determine the number of shared clusters between the two graphs

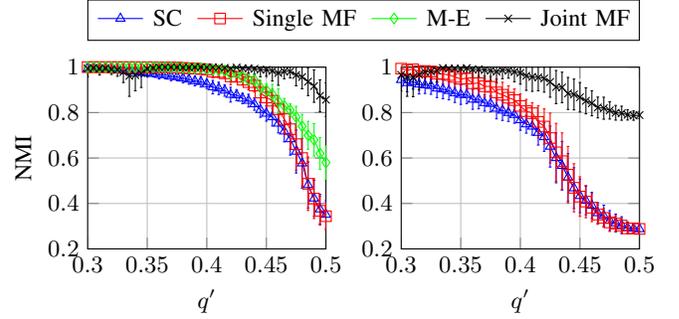


Fig. 3: Normalized Mutual Information (NMI) between communities (of noisier graph $\mathcal{G}^{(2)}$) identified by different community detection algorithms and ground truths, $n = 500$, $K = 2$ shared communities between the two graphs, $K^{(l)} = 4$. Averages over 100 randomly generated graphs.

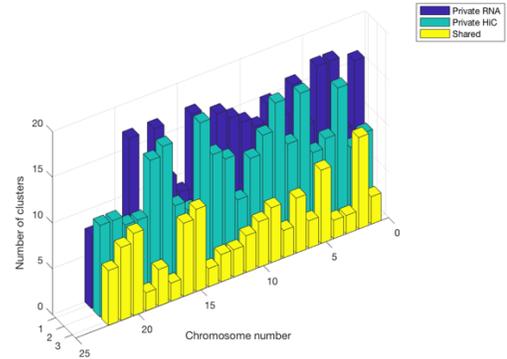


Fig. 4: Histogram of the number of clusters (private and shared) for each chromosome.

($\mathbf{A}^{(1)}$ and $\mathbf{A}^{(2)}$) as well as the number of clusters unique to each graph. We first determine the total number of clusters $K^{(1)}$, $K^{(2)}$ respectively for each graph by applying separately the single layer variational Bayes community detection algorithm [21] and choose the number of clusters maximizing the Bayesian Information Criteria (BIC) [27]. We then apply our joint method with fixed $K^{(1)}$ and $K^{(2)}$ (determined previously) and choosing the number of shared clusters K (in the range $[2, \min(K^{(1)}, K^{(2)})]$) maximizing the modularity [2] in both graphs. Figure 4 shows for each chromosome the number of shared clusters, the number of private clusters for graph $\mathbf{A}^{(1)}$ (Private RNA) and the number of private clusters for graph $\mathbf{A}^{(2)}$ (Private HiC). This analysis suggests that the genes in chromosomes 2, 5, 15, 16, 20, 21, 22 are strongly co-expressed (high connectivity in $\mathbf{A}^{(1)}$) and strongly connected (high connectivity in $\mathbf{A}^{(2)}$) since the number of shared clusters between graphs $\mathbf{A}^{(1)}$ and $\mathbf{A}^{(2)}$ is dominant compared to the number of private clusters, while the genes in the other chromosomes are either strongly co-expressed or strongly connected.

V. CONCLUSION

Our proposed joint variational algorithm is capable of extracting shared communities across all graph layers as well as identifying communities unique to each layer. The method is applicable to any multilayer network (with or without edge weights) and can provide important insights in the analysis of real-world systems as demonstrated for the human fibroblast dataset. An interesting direction of future investigation would be to consider extensions to the case where only a subset of layers share communities.

Acknowledgments: The authors thank Zeyu Sun at the University of Michigan for his help implementing the numerical experiments.

VI. REFERENCES

- [1] Santo Fortunato, "Community detection in graphs," *Physics Reports*, vol. 486, no. 3, pp. 75–174, 2010.
- [2] Mark EJ Newman, "Modularity and community structure in networks," *Proceedings of the National Academy of Sciences*, vol. 103, no. 23, pp. 8577–8582, 2006.
- [3] Mark EJ Newman, "Spectral methods for community detection and graph partitioning," *Physical Review E*, vol. 88, no. 4, pp. 042822, 2013.
- [4] Matthew B Hastings, "Community detection as an inference problem," *Physical Review E*, vol. 74, no. 3, pp. 035102, 2006.
- [5] Stefano Boccaletti, Ginestra Bianconi, Regino Criado, Charo I Del Genio, Jesús Gómez-Gardenes, Miguel Romance, Irene Sendina-Nadal, Zhen Wang, and Massimiliano Zanin, "The structure and dynamics of multilayer networks," *Physics Reports*, vol. 544, no. 1, pp. 1–122, 2014.
- [6] Brandon Oselio, Alex Kulesza, and Alfred O Hero, "Multi-layer graph analysis for dynamic social networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 4, pp. 514–523, 2014.
- [7] Jesse R Dixon, Inkyung Jung, Siddarth Selvaraj, Yin Shen, Jessica E Antosiewicz-Bourget, Ah Young Lee, Zhen Ye, Audrey Kim, Nisha Rajagopal, Wei Xie, et al., "Chromatin architecture reorganization during stem cell differentiation," *Nature*, vol. 518, no. 7539, pp. 331, 2015.
- [8] Deng Cai, Zheng Shao, Xiaofei He, Xifeng Yan, and Jiawei Han, "Community mining from multi-relational networks," in *European Conference on Principles of Data Mining and Knowledge Discovery*. Springer, 2005, pp. 445–452.
- [9] Manlio De Domenico, Andrea Lancichinetti, Alex Arenas, and Martin Rosvall, "Identifying modular flows on multilayer networks reveals highly overlapping organization in interconnected systems," *Physical Review X*, vol. 5, no. 1, pp. 011027, 2015.
- [10] Lei Tang, Xufei Wang, and Huan Liu, "Uncovering groups via heterogeneous interaction analysis," in *Data Mining, 2009. ICDM'09. Ninth IEEE International Conference on*. IEEE, 2009, pp. 503–512.
- [11] Lei Tang, Xufei Wang, and Huan Liu, "Community detection via heterogeneous interaction analysis," *Data mining and knowledge discovery*, vol. 25, no. 1, pp. 1–33, 2012.
- [12] Alessia Amelio and Clara Pizzuti, "Community detection in multidimensional networks," in *Tools with Artificial Intelligence (ICTAI), 2014 IEEE 26th International Conference on*. IEEE, 2014, pp. 352–359.
- [13] Subhadeep Paul and Yuguo Chen, "Null models and modularity based community detection in multi-layer networks," *arXiv preprint arXiv:1608.00623*, 2016.
- [14] Paul W Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt, "Stochastic blockmodels: First steps," *Social networks*, vol. 5, no. 2, pp. 109–137, 1983.
- [15] Toni Valles-Catala, Francesco A Massucci, Roger Guimera, and Marta Sales-Pardo, "Multilayer stochastic block models reveal the multilayer structure of complex networks," *Physical Review X*, vol. 6, no. 1, pp. 011036, 2016.
- [16] Perla Reyes and Abel Rodriguez, "Stochastic blockmodels for exchangeable collections of networks," *arXiv preprint arXiv:1606.05277*, 2016.
- [17] Pierre Barbillon, Sophie Donnet, Emmanuel Lazega, and Avner Bar-Hen, "Stochastic block models for multiplex networks: an application to a multilevel network of researchers," *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, vol. 180, no. 1, pp. 295–314, 2017.
- [18] Pierre Barbillon, Sophie Donnet, Emmanuel Lazega, and Avner Bar-Hen, "Stochastic block models for multiplex networks: an application to a multilevel network of researchers," *Journal of Machine Learning Research*, vol. 18, no. 1, pp. 5458–5506, 2017.
- [19] Brigitte Boden, Stephan Günnemann, Holger Hoffmann, and Thomas Seidl, "Mining coherent subgraphs in multi-layer graphs with edge labels," in *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2012, pp. 1258–1266.
- [20] Zhiping Zeng, Jianyong Wang, Lizhu Zhou, and George Karypis, "Coherent closed quasi-clique discovery from large dense graph databases," in *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2006, pp. 797–802.
- [21] Christopher Aicher, Abigail Z Jacobs, and Aaron Clauset, "Learning latent block structure in weighted networks," *Journal of Complex Networks*, p. cnu026, 2014.
- [22] Anderson Y Zhang and Harrison H Zhou, "Theoretical and computational guarantees of mean field variational inference for community detection," *arXiv preprint arXiv:1710.11268*, 2017.
- [23] Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul, "An introduction to variational methods for graphical models," *Machine learning*, vol. 37, no. 2, pp. 183–233, 1999.
- [24] David M Blei, Michael I Jordan, et al., "Variational inference for dirichlet process mixtures," *Bayesian analysis*, vol. 1, no. 1, pp. 121–143, 2006.
- [25] H. Chen, J. Chen, et al., "Functional organization of the human 4d nucleome," *Proceedings of the National Academy of Sciences*, vol. 112, no. 26, pp. 8002–8007, 2015.
- [26] Erez Lieberman-Aiden, Nynke L. van Berkum, Louise Williams, et al., "Comprehensive mapping of long-range interactions reveals folding principles of the human genome," *Science*, vol. 326, no. 5950, pp. 289–293, 2009.
- [27] Xiaoran Yan, "Bayesian model selection of stochastic block models," in *Proceedings of the 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. IEEE Press, 2016, pp. 323–328.