

Inference of Edge Correlations in Multilayer Networks

A. Roxana Pamfil and Sam D. Howison
Mathematical Institute, University of Oxford

Mason A. Porter
*Department of Mathematics, University of California Los Angeles and
Mathematical Institute, University of Oxford*
(Dated: September 8, 2020)

Many recent developments in network analysis have focused on multilayer networks, which one can use to encode time-dependent interactions, multiple types of interactions, and other complications that arise in complex systems. Like their monolayer counterparts, multilayer networks in applications often have mesoscale features, such as community structure. A prominent type of method for inferring such structures is the employment of multilayer stochastic block models (SBMs). A common (but potentially inadequate) assumption of these models is the sampling of edges in different layers independently, conditioned on the community labels of the nodes. In this paper, we relax this assumption of independence by incorporating edge correlations into an SBM-like model. We derive maximum-likelihood estimates of the key parameters of our model, and we propose a measure of layer correlation that reflects the similarity between connectivity patterns in different layers. Finally, we explain how to use correlated models for edge “prediction” (i.e., inference) in multilayer networks. By taking into account edge correlations, prediction accuracy improves both in synthetic networks and in a temporal network of shoppers who are connected to previously-purchased grocery products.

I. INTRODUCTION

A network is an abstract representation of a system in which entities called “nodes” interact with each other via connections called “edges” [1]. Most traditionally, in a type of network called a “graph”, each edge encodes an interaction between a pair of nodes. Networks arise in many domains and are useful for numerous practical problems, such as detecting bot accounts on Twitter [2], finding vulnerabilities in electrical grids [3], and identifying potentially harmful interactions between drugs [4]. A common feature of many networks is mesoscale (i.e., intermediate-scale) structures. Detecting such structures amounts to a type of coarse-graining, providing representations of a network that are more compact than listing all of the nodes and edges. Types of mesoscale structures include community structure [5], core-periphery structure [6, 7], role similarity [8], and others. An increasingly popular approach for modeling and detecting such structures is by using stochastic block models (SBMs) [9], a generative model that can produce networks with community structure or other mesoscale structures.

For many applications of network analysis, it is important to move beyond ordinary graphs (i.e., “monolayer networks”) to examine more complicated network structures, such as collections of interrelated networks. One can study such structures through the flexible lens of multilayer networks [10–13]. Similar to monolayer networks, a multilayer network consists of a collection of “state nodes” that are connected pairwise by edges. A state node is a manifestation of a “physical node” (which we will also sometimes call simply a “node”), which represents some entity, in a specific layer. Different layers may correspond to interactions in different time periods

(yielding a temporal network), different types of relations (yielding a multiplex network), or other possibilities. As in the setting of monolayer networks, modeling and inferring mesoscale structures in multilayer networks is a prominent research area.

A key assumption of almost all existing models of multilayer networks with mesoscale structure is that edges are generated independently, conditioned on a multilayer partition [14–21]. This independence condition applies both within each layer (which is inconsistent with the fact that real networks often include 3-cliques and other small-scale structures) and across layers (which is inconsistent with the fact that the same nodes are often adjacent to each other in multiple layers). In this paper, we focus on relaxing the edge-independence assumption that applies to edges between the same two physical nodes in different layers. We still consider each pair of nodes independently.

In Fig. 1, we show an example of a two-layer network with both strong positive and strong negative edge correlations. Incorporating such correlations into a network model is beneficial for many applications. For example, a multiplex network of air routes, where each layer corresponds to one airline, is likely to include some popular routes that appear in multiple layers (and unpopular routes may appear only in one layer)[22]. In a temporal social network, we expect people to have repeated interactions with other people [23]; this is a stronger statement than just saying that they tend to interact more within the same community over time. Such edge persistence is also common in many bipartite user-item networks: shoppers tend to buy the same grocery products over time [24], customers of a music-streaming platform listen repeatedly to their favorite songs [25], and

Wikipedia users edit specific pages several times [26].




FIG. 1: Example of a correlated multilayer network with two layers and with two blocks of nodes in each layer. Edges between two black nodes are positively correlated across layers, edges between two white nodes are negatively correlated across layers, and edges between a black node and a white node are uncorrelated across layers.

Multilayer network models that incorporate edge correlations have many important applications. One is the inference task of *edge prediction* (also called *link prediction*), where one seeks to assign probabilities of occurrence to unobserved edges. SBMs have often been used for edge prediction for both monolayer networks [4, 9, 27] and multilayer networks [15, 20, 28]. The models that we propose should yield better performance than past efforts, as they take advantage of interlayer edge correlations in data. We discuss this in more detail in Sec. III. Another application is to *graph matching* [29], where one seeks to infer a latent correspondence between nodes in two different networks when one does not know the identities of the nodes. For example, one may wish to match common users between anonymized Twitter and Facebook networks. In a series of papers [30–32], Lyzinski and collaborators established conditions under which graph matching is successful. They tested their methods on correlated Erdős–Rényi (ER) networks and correlated SBMs, which we also investigate in this paper (albeit for a different purpose). We take these models further by incorporating degree correction [33], which generates networks with heterogeneous degree distributions and is important for inference using SBM. This extension may allow one to study the graph-matching problem on more realistic network models. A third application of our work is efficient computation of correlations between pairs of layers of a multilayer network. One can use such correlation estimates to quantify the similarity between different layers and potentially to compress multilayer network data by discarding (or merging) layers that are strongly positively correlated with an existing layer [34]. Previous papers have focused primarily on node-centric notions of layer similarity [35–37], whereas our correlated models yield edge-centric measures of similarity. Benefits of our approach over related studies [34, 38, 39] include the fact

that correlation values cover an intuitive range (between -1 and 1) and that they work equally well for quantifying layer similarity and dissimilarity.

Our edge-correlated network models are also useful for community detection. Given a multilayer network, one can design an inference algorithm that determines both the parameters that describe the edge probabilities (and correlations) and a multilayer community structure that underlies these probabilities. Solving this inference problem enables the detection of “correlated communities”. Because of the additional complexity in the model, this is bound to be more difficult than standard multilayer community detection, so we leave this inference problem for future work. Instead, for the rest of the paper, we assume that we know the block structure of a network; we infer the remaining parameters, including the correlations that are the core element of our model [40]. We also assume that the block assignments \mathbf{g} are the same for all layers; the case in which communities can vary arbitrarily across layers is significantly more difficult [41], and we leave its consideration for future work. With these restrictions, one can determine \mathbf{g} using any method of choice.

Some existing models of multilayer networks incorporate interlayer dependencies by prescribing joint degree distributions [36, 42, 43], by incorporating edge overlaps [44, 45], or by modeling the appearance of new edges through preferential-attachment mechanisms [46]. The models that we describe in this paper are similar to those that were introduced in [30–32] for graph-matching purposes. Another noteworthy paper is one by Barucca et al. [47] that described a generalized version of the temporal SBM of Ghasemian et al. [17]. This generalization includes an “edge-persistence” parameter ξ , which gives the probability that an edge from one layer also occurs in the next temporal layer. For several reasons, we take a different approach. First, the model of [47] is specific to temporal networks, whereas we are also interested in other types of multilayer networks. Second, their model does not easily incorporate degree correction. Third, we want to include correlations explicitly in the model, rather than implicitly using the edge-persistence parameter ξ .

Our paper proceeds as follows. In Sec. II, we describe our models of multilayer networks with edge correlations. We start with a simple example of correlated Erdős–Rényi (ER) graphs in Sec. II A to make our exposition for more complicated models easier to follow. In Sec. II B, we integrate mesoscale structures by incorporating correlations in an SBM-like model. We then introduce degree correction in Sec. II C. For all of these models, we derive maximum-likelihood (ML) estimates both of the marginal edge-existence probabilities in each layer and of the interlayer correlations. ML estimation is common for SBMs and DCSBMs in both monolayer networks [33] and multilayer networks [16]. Although ML estimation is less powerful than performing Bayesian inference [9], the former is consistent for both SBMs and DCSBMs [48]

and it recovers many common techniques for detecting mesoscale structures in networks [49]. In Sec. III, we describe how to use our models for edge prediction, and we provide some results for synthetic networks. We then proceed with applications in Sec. IV. In Sec. IV A, we use our models to estimate pairwise layer correlations in several empirical networks. In Sec. IV B, we use our correlated models for edge prediction in a temporal network of grocery purchases. We summarize our results in Sec. V and discuss a few ideas for future work.

II. CORRELATED MODELS

In our derivations, we consider just two network layers at a time. Although this may seem limiting, we can apply our framework to generate correlated networks with more than two layers in a sequential manner (see the discussion in Sec. II A), and we can determine pairwise layer correlations for a network with arbitrarily many layers (see the applications in Sec. IV). In Sec. V, we briefly discuss the challenges that arise when modeling three or more layers simultaneously, rather than in a sequential pairwise fashion.

Consider a network with two layers and identical sets of nodes in each layer; this is known as a *node-aligned* multilayer network [10]. Let \mathbf{A}^1 and \mathbf{A}^2 denote the adjacency matrices of our two network layers. As in many generative models of networks, we assume that edges in these two layers are generated by some random process, so the entries A_{ij}^1 and A_{ij}^2 are random variables. Imposing some statistical correlation between these two sets of random variables introduces interlayer correlations in the resulting multilayer network structure.

Our goal is to propose a model of correlated networks in which each layer is, marginally, a degree-corrected stochastic block model (DCSBM) [33]. However, it is instructive to first consider the simpler cases in which each layer is marginally an Erdős–Rényi random graph (see Sec. II A) or an SBM without degree correction (see Sec. II B). Correlated ER models and correlated SBMs have been studied previously, most notably in work by Lyzinski et al. [30–32] on the graph-matching problem. However, our use of these models for estimating layer correlations is novel, as are the correlated DCSBMs that we propose in Sec. II C.

In monolayer SBMs, it is common to use either Bernoulli or Poisson random variables to generate edges between nodes. The former is generally more accurate, because it does not yield multiedges; however, the latter is more common, as it often simplifies calculations considerably [48, 50]. Nevertheless, we have found that Bernoulli models are simpler when incorporating correlations [24]. They also have several other advantages, including the fact that they work for both sparse and dense networks and that they can handle the entire correlation range between -1 and 1 . Therefore, we consider only Bernoulli models in this paper.

A. Correlated Erdős–Rényi Layers

1. Forward model

Assume that the intralayer networks that correspond to \mathbf{A}^1 and \mathbf{A}^2 are ER graphs from the $\mathcal{G}(n, p)$ ensemble [1] with edge probabilities p_1 and p_2 . For each pair of nodes (i, j) , we therefore have

$$\mathbb{P}(A_{ij}^1 = 1) = p_1, \quad (1)$$

$$\mathbb{P}(A_{ij}^2 = 1) = p_2. \quad (2)$$

To couple edges that connect the same pair of nodes in different layers, let

$$q := \mathbb{P}(A_{ij}^1 = 1, A_{ij}^2 = 1) \quad (3)$$

denote the joint probability for an edge to occur in both layers. Unless $q = p_1 p_2$, this construction implies that the random variables A_{ij}^1 and A_{ij}^2 are not independent.

The parameters p_1 , p_2 , and q (which lie in the interval $[0, 1]$) fully specify a forward model of networks with correlated ER layers. To generate a network from this model, one considers each node pair (i, j) and, independently of all other node pairs, assigns values to A_{ij}^1 and A_{ij}^2 according to the following probabilities:

$$\begin{aligned} \mathbb{P}(A_{ij}^1 = 1, A_{ij}^2 = 1) &= q, \\ \mathbb{P}(A_{ij}^1 = 1, A_{ij}^2 = 0) &= p_1 - q, \\ \mathbb{P}(A_{ij}^1 = 0, A_{ij}^2 = 1) &= p_2 - q, \\ \mathbb{P}(A_{ij}^1 = 0, A_{ij}^2 = 0) &= 1 - p_1 - p_2 + q. \end{aligned} \quad (4)$$

These expressions follow from the laws of probability and from the definitions of p_1 , p_2 , and q . For these probabilities to be well-defined, it is both necessary and sufficient that $0 \leq q \leq \min(p_1, p_2)$ and $p_1 + p_2 \leq 1 + q$. In Fig. 2, we illustrate the feasible region for p_1 and p_2 , given a value of q .

It is also possible to generate a correlated ER network in a sequential manner. First, one generates the adjacency matrix \mathbf{A}^1 by placing edges with probability p_1 . One then determines the probabilities of edges in the second layer by conditioning on the first layer:

$$\begin{aligned} \mathbb{P}(A_{ij}^2 = 1 | A_{ij}^1 = 1) &= \frac{\mathbb{P}(A_{ij}^1 = 1, A_{ij}^2 = 1)}{\mathbb{P}(A_{ij}^1 = 1)} = \frac{q}{p_1}, \\ \mathbb{P}(A_{ij}^2 = 1 | A_{ij}^1 = 0) &= \frac{\mathbb{P}(A_{ij}^1 = 0, A_{ij}^2 = 1)}{\mathbb{P}(A_{ij}^1 = 0)} = \frac{p_2 - q}{1 - p_1}, \\ \mathbb{P}(A_{ij}^2 = 0 | A_{ij}^1 = 1) &= \frac{\mathbb{P}(A_{ij}^1 = 1, A_{ij}^2 = 0)}{\mathbb{P}(A_{ij}^1 = 1)} = \frac{p_1 - q}{p_1}, \\ \mathbb{P}(A_{ij}^2 = 0 | A_{ij}^1 = 0) &= \frac{\mathbb{P}(A_{ij}^1 = 0, A_{ij}^2 = 0)}{\mathbb{P}(A_{ij}^1 = 0)} \\ &= \frac{1 - p_1 - p_2 + q}{1 - p_1}. \end{aligned} \quad (5)$$




FIG. 2: Visualization of the feasible region (gray area) for p_1 and p_2 , given a value of q . The boundaries of this region are defined by the inequalities $q \leq p_1 \leq 1$, $q \leq p_2 \leq 1$, and $p_1 + p_2 \leq 1 + q$. The hyperbola $p_1 p_2 = q$ specifies the boundary between regimes with a positive layer correlation and regimes with a negative layer correlation.

With this approach, it is possible to generate networks with arbitrarily many layers by first sampling edges in the first layer, and then sampling edges in each subsequent layer by conditioning on the previous one. This kind of process is especially well-suited to temporal networks, in which layers have a natural ordering. For multiplex networks, it is more appropriate to extend Eqns. (4) to handle more than two layers.

It is also possible to parametrize correlated ER graphs in terms of the marginal Bernoulli probabilities, p_1 and p_2 , and the Pearson correlation

$$\begin{aligned} \rho &= \frac{\mathbb{E}[A_{ij}^1 A_{ij}^2] - \mathbb{E}[A_{ij}^1] \mathbb{E}[A_{ij}^2]}{\sigma[A_{ij}^1] \sigma[A_{ij}^2]} \\ &= \frac{q - p_1 p_2}{\sqrt{p_1(1-p_1)p_2(1-p_2)}}, \end{aligned} \quad (6)$$

where $\mathbb{E}[\cdot]$ and $\sigma[\cdot]$, respectively, denote the mean and standard deviation of a random variable. One benefit of using ρ , rather than q , as the third model parameter is that its value is easier to interpret. A value of ρ that is close to 0 indicates a weak correlation between layers, whereas values that are close to the extremes of +1 and -1 indicate a strong positive correlation and a strong negative correlation, respectively.

We can gain further intuition by considering the cases $\rho = 0$, $\rho = 1$, and $\rho = -1$. First, $\rho = 0$ if and only if $\mathbb{P}(A_{ij}^1 = 1, A_{ij}^2 = 1) = \mathbb{P}(A_{ij}^1 = 1)\mathbb{P}(A_{ij}^2 = 1)$. That is, the correlation is 0 if and only if edges are generated independently in the two layers with marginal probabilities of p_1 and p_2 . For $\rho = 1$, one can show (see [24]) that $p_1 = p_2 = q$, which corresponds to the two layers having identical network structure. Lastly, for $\rho = -1$, we have $q = 0$ and $p_1 = 1 - p_2$ (see [24]), and two nodes are adjacent in one layer if and only if they are not adjacent in the other layer.

2. Maximum-likelihood parameter estimates

We now derive ML estimates of the parameters p_1 , p_2 , and q . Let \mathcal{E} denote the set of node pairs that can form edges. For undirected networks without self-edges, there are $|\mathcal{E}| = N(N-1)/2$ such node pairs to consider, where N is the number of physical nodes. By contrast, $|\mathcal{E}| = N(N-1)$ when generating directed networks without self-edges. With this general notation, all our derivations in Sec. II are valid for both directed and undirected networks, with or without self-edges. (They are also valid for bipartite networks [24].) We consider each pair of nodes $(i, j) \in \mathcal{E}$ independently when generating edges, so the likelihood of observing adjacency matrices \mathbf{A}^1 and \mathbf{A}^2 is

$$\mathbb{P}(\mathbf{A}^1, \mathbf{A}^2 | p_1, p_2, q) = \prod_{(i,j) \in \mathcal{E}} q^{A_{ij}^1 A_{ij}^2} (p_1 - q)^{A_{ij}^1 (1-A_{ij}^2)} (p_2 - q)^{(1-A_{ij}^1) A_{ij}^2} (1 - p_1 - p_2 + q)^{(1-A_{ij}^1)(1-A_{ij}^2)}. \quad (7)$$

It is helpful to introduce the following notation:

$$\begin{aligned} e_{11} &:= |\{(i, j) \in \mathcal{E} : A_{ij}^1 = 1, A_{ij}^2 = 1\}|, \\ e_{10} &:= |\{(i, j) \in \mathcal{E} : A_{ij}^1 = 1, A_{ij}^2 = 0\}|, \\ e_{01} &:= |\{(i, j) \in \mathcal{E} : A_{ij}^1 = 0, A_{ij}^2 = 1\}|, \\ e_{00} &:= |\{(i, j) \in \mathcal{E} : A_{ij}^1 = 0, A_{ij}^2 = 0\}|. \end{aligned}$$

These quantities correspond, respectively, to the number of node pairs that are adjacent in both layers, are adjacent in the first layer but not in the second, are adjacent in the second layer but not in the first, and are not adjacent in either layer. Using this notation and taking the logarithm of (7), we arrive at the following expression for

the log-likelihood:

$$\mathcal{L} = e_{11} \log q + e_{10} \log(p_1 - q) + e_{01} \log(p_2 - q) + e_{00} \log(1 - p_1 - p_2 + q). \quad (8)$$

When fitting our model to network data, the quantities e_{11} , e_{10} , e_{01} , e_{00} are all known; and we seek to determine the values of p_1 , p_2 , and q that are best explained by the data. To do so, we maximize the log-likelihood (8) by setting its partial derivatives to 0 [51]. We obtain

$$\hat{p}_1 = \frac{e_{11} + e_{10}}{e_{11} + e_{10} + e_{01} + e_{00}}, \quad (9)$$

$$\hat{p}_2 = \frac{e_{11} + e_{01}}{e_{11} + e_{10} + e_{01} + e_{00}}, \quad (10)$$

$$\hat{q} = \frac{e_{11}}{e_{11} + e_{10} + e_{01} + e_{00}}. \quad (11)$$

In all three expressions, the denominator is equal to the number of potential edges (i.e., the cardinality of \mathcal{E}). Additionally, let $m_1 = e_{11} + e_{10}$ and $m_2 = e_{11} + e_{01}$ denote the number of observed edges in the first and the second layers, respectively. It follows that the ML estimate \hat{p}_1 is equal to the number of observed edges in layer 1 divided by the number of potential edges, and an analogous relation holds for \hat{p}_2 . The estimate \hat{q} is equal to the number of node pairs that are adjacent in both layers divided by the total number of node pairs. These results match our intuition.

We obtain an estimate of the Pearson correlation ρ between the two layers by substituting the ML estimates \hat{p}_1 , \hat{p}_2 , and \hat{q} into Eqn. (6) to obtain

$$\hat{\rho} = \frac{e_{00}e_{11} - e_{10}e_{01}}{\sqrt{(e_{11} + e_{10})(e_{11} + e_{01})(e_{10} + e_{00})(e_{01} + e_{00})}}. \quad (12)$$

One can show that maximizing the log-likelihood (8) with respect to p_1 , p_2 , and ρ (rather than with respect to p_1 , p_2 , and q) gives the same expression for $\hat{\rho}$, confirming that this is indeed an ML estimate of the correlation. Note that $\hat{\rho}$ is not defined when either layer is an empty or a complete graph, as the corresponding Bernoulli random variable has a standard deviation of 0.

In App. A, we calculate the variances that are associated with the ML estimates \hat{p}_1 , \hat{p}_2 , \hat{q} . We then show using a synthetic example that these scale as $1/N^2$, where we recall that N is the number of physical nodes. These results quantify the uncertainty around the ML estimates for correlated ER models as a function of network size.

B. Correlated SBMs

One of the ways in which real-world networks differ from ER random graphs is that the former have mesoscale structures, such as communities [5]. We use SBMs to incorporate such structures into our correlated models.

Let \mathbf{g} be a vector of block assignments, which we take to be identical for both network layers, and let K denote the number of blocks. As we explained in Sec. I, we assume throughout the present paper that we are given \mathbf{g} , and we aim to estimate the remaining model parameters. Following terminology from [27], let $\mathcal{B} = \{1, \dots, K\} \times \{1, \dots, K\}$ be the set of “edge bundles” (r, s) , each of which is described by its own set of parameters p_{rs}^1 , p_{rs}^2 , and q_{rs} . The $K \times K$ matrices \mathbf{p}^1 , \mathbf{p}^2 , and \mathbf{q} play an analogous role to p_1 , p_2 , and q in the ER layers.

Let g_i denote the block assignment of node i . A correlated two-layer SBM is described by the following set of equalities:

$$\begin{aligned} \mathbb{P}(A_{ij}^1 = 1) &= p_{g_i g_j}^1, \\ \mathbb{P}(A_{ij}^2 = 1) &= p_{g_i g_j}^2, \\ \mathbb{P}(A_{ij}^1 = 1, A_{ij}^2 = 1) &= q_{g_i g_j}. \end{aligned}$$

Lyzinski et al. proposed this forward model in [31] to study the graph-matching problem. By contrast, we focus on the inverse problem of estimating the parameters \mathbf{p}^1 , \mathbf{p}^2 , and \mathbf{q} , given some network data.

1. Maximum-likelihood parameter estimates

As in Sec. II A, suppose that we consider each node pair (i, j) independently. The likelihood of observing adjacency matrices \mathbf{A}^1 and \mathbf{A}^2 is then

$$\begin{aligned} \mathbb{P}(\mathbf{A}^1, \mathbf{A}^2 | \mathbf{g}, \mathbf{p}^1, \mathbf{p}^2, \mathbf{q}) &= \prod_{(i,j) \in \mathcal{E}} \left[q_{g_i g_j}^{A_{ij}^1 A_{ij}^2} (p_{g_i g_j}^1 - q_{g_i g_j})^{A_{ij}^1 (1 - A_{ij}^2)} (p_{g_i g_j}^2 - q_{g_i g_j})^{(1 - A_{ij}^1) A_{ij}^2} \right. \\ &\quad \left. \times (1 - p_{g_i g_j}^1 - p_{g_i g_j}^2 + q_{g_i g_j})^{(1 - A_{ij}^1)(1 - A_{ij}^2)} \right]. \end{aligned} \quad (13)$$

In this product, each factor depends on i and j only via their block memberships g_i and g_j , so we can combine

several terms. First, define

$$e_{rs}^{ab} := |\{(i, j) \in \mathcal{E} : A_{ij}^1 = a, A_{ij}^2 = b, g_i = r, g_j = s\}|$$

for $(a, b) \in \{(1, 1), (1, 0), (0, 1), (0, 0)\}$, in analogy with e_{11} , e_{10} , e_{01} , and e_{00} from Sec. II A. We can then write the log-likelihood as

$$\mathcal{L} = \sum_{(r,s) \in \mathcal{B}} \left[e_{rs}^{11} \log q_{rs} + e_{rs}^{10} \log(p_{rs}^1 - q_{rs}) + e_{rs}^{01} \log(p_{rs}^2 - q_{rs}) + e_{rs}^{00} \log(1 - p_{rs}^1 - p_{rs}^2 + q_{rs}) \right]. \quad (14)$$

The advantage of writing the log-likelihood as in (14) is that it clearly separates the contribution from different edge bundles. Using the results for ER layers from Sec. II A, we immediately obtain (without further calculations) the following ML parameter estimates:

$$\hat{p}_{rs}^1 = \frac{e_{rs}^{11} + e_{rs}^{10}}{e_{rs}^{11} + e_{rs}^{10} + e_{rs}^{01} + e_{rs}^{00}} = \frac{m_{rs}^1}{e_{rs}}, \quad (15)$$

$$\hat{p}_{rs}^2 = \frac{e_{rs}^{11} + e_{rs}^{01}}{e_{rs}^{11} + e_{rs}^{10} + e_{rs}^{01} + e_{rs}^{00}} = \frac{m_{rs}^2}{e_{rs}}, \quad (16)$$

$$\hat{q}_{rs} = \frac{e_{rs}^{11}}{e_{rs}^{11} + e_{rs}^{10} + e_{rs}^{01} + e_{rs}^{00}} = \frac{e_{rs}^{11}}{e_{rs}}, \quad (17)$$

where m_{rs}^1 and m_{rs}^2 denote the number of edges between blocks r and s in layers 1 and 2, respectively, and e_{rs} is the number of possible edges between nodes in block r and block s . When there is a single edge bundle (i.e., when we do not assume any block structure in a network), the ML estimates (15)–(17) recover those that we obtained for correlated ER networks in Sec. II A. Each edge bundle also has a corresponding Pearson correlation, whose ML estimate is

$$\hat{\rho}_{rs} = \frac{e_{rs}^{00}e_{rs}^{11} - e_{rs}^{10}e_{rs}^{01}}{\sqrt{(e_{rs}^{11} + e_{rs}^{10})(e_{rs}^{11} + e_{rs}^{01})(e_{rs}^{10} + e_{rs}^{00})(e_{rs}^{01} + e_{rs}^{00})}}. \quad (18)$$

In applications to temporal consumer–product networks, we find that different edge bundles have vastly different correlation values [24]. We anticipate that other empirical multilayer networks have similar properties.

2. Effective correlation

Although having different correlation values for different edge bundles can be useful, it is also helpful to have a single correlation measurement for a given multilayer network. For example, one may wish to use such a network diagnostic for one of the purposes that we outlined in Sec. I. One way to define an “effective correlation” is to first sample two node indices, I and J , uniformly at random and then compute the Pearson correlation of the random variables A_{IJ}^1 and A_{IJ}^2 . That is,

$$\text{corr}(A_{IJ}^1, A_{IJ}^2) = \frac{\mathbb{E}[A_{IJ}^1 A_{IJ}^2] - \mathbb{E}[A_{IJ}^1] \mathbb{E}[A_{IJ}^2]}{\sigma[A_{IJ}^1] \sigma[A_{IJ}^2]}, \quad (19)$$

where we use capital letters for the node indices I and J to emphasize that they are random variables.

We can calculate each term on the right-hand side of (19) by conditioning on the block assignments of the randomly chosen nodes I and J . First, for $l \in \{1, 2\}$, we have

$$\begin{aligned} \mathbb{E}[A_{IJ}^l] &= \mathbb{P}(A_{IJ}^l = 1) \\ &= \sum_{(r,s) \in \mathcal{B}} \mathbb{P}(A_{IJ}^l = 1 | g_I = r, g_J = s) \mathbb{P}(g_I = r, g_J = s) \\ &= \sum_{(r,s) \in \mathcal{B}} \hat{p}_{rs}^l \frac{e_{rs}}{|\mathcal{E}|} = \sum_{(r,s) \in \mathcal{B}} \frac{m_{rs}^l e_{rs}}{e_{rs} |\mathcal{E}|} = \frac{m_l}{|\mathcal{E}|}, \end{aligned} \quad (20)$$

where m_l denotes the number of edges in layer l . The expression (20) is the same as the probability p_l of generating an edge in layer l for the ER case. Because A_{IJ}^l is a Bernoulli random variable (in other words, it can only take the values 1 or 0), its standard deviation is

$$\sigma[A_{IJ}^l] = \sqrt{\frac{m_l}{|\mathcal{E}|} \left(1 - \frac{m_l}{|\mathcal{E}|}\right)}.$$

Lastly,

$$\begin{aligned} \mathbb{E}[A_{IJ}^1 A_{IJ}^2] &= \mathbb{P}(A_{IJ}^1 = 1, A_{IJ}^2 = 1) \\ &= \sum_{(r,s) \in \mathcal{B}} \hat{q}_{rs} \frac{e_{rs}}{|\mathcal{E}|} = \sum_{(r,s) \in \mathcal{B}} \frac{e_{rs}^{11} e_{rs}}{e_{rs} |\mathcal{E}|} = \frac{e_{11}}{|\mathcal{E}|}. \end{aligned}$$

The estimated value of the effective correlation is thus

$$\begin{aligned} \hat{\rho} &= \text{corr}(A_{IJ}^1, A_{IJ}^2) \\ &= \frac{e_{00}e_{11} - e_{10}e_{01}}{\sqrt{(e_{11} + e_{10})(e_{11} + e_{01})(e_{10} + e_{00})(e_{01} + e_{00})}}, \end{aligned} \quad (21)$$

which recovers the value in (12) for ER layers (i.e., without any block structure in the model). We stress that there is no reason a priori to expect this outcome. In fact, the analogous result does *not* hold for Poisson models [24]. In the present case, the fact that there is such a correspondence between models is convenient for practical reasons, as it implies that one can perform the simpler calculations from Sec. II A to obtain correlation estimates between network layers, even for networks with nontrivial mesoscale structure.

C. Correlated Degree-Corrected SBMs

The models that we have discussed thus far generate networks in which nodes in the same block have the same expected degree. SBMs that make this kind of assumption tend to perform poorly when they are used to infer mesoscale structure in real networks, many of which have highly heterogeneous degree distributions. This observation led to the development of degree-corrected SBMs (DCSBMs) [33]. We expect that such adjustments can also make a difference when modeling edge correlations, so we now extend the model from Sec. II B to incorporate degree correction.

We continue to work with two-layer networks, which we again specify in terms of two intralayer adjacency matrices, \mathbf{A}^1 and \mathbf{A}^2 , with a common block structure that we specify with a vector \mathbf{g} . For each node pair $(i, j) \in \mathcal{E}$, we place edges in the two layers according to the probabilities

$$\mathbb{P}(A_{ij}^1 = 1) = \theta_i^1 \theta_j^1 p_{g_i g_j}^1, \quad (22)$$

$$\mathbb{P}(A_{ij}^2 = 1) = \theta_i^2 \theta_j^2 p_{g_i g_j}^2, \quad (23)$$

$$\mathbb{P}(A_{ij}^1 = 1, A_{ij}^2 = 1) = \sqrt{\theta_i^1 \theta_j^1 \theta_i^2 \theta_j^2} q_{g_i g_j}. \quad (24)$$

We will soon justify the expression in (24). The quantities θ_i^l and θ_j^l , with $l \in \{1, 2\}$, are the degrees of nodes i and j , normalized by the mean degrees. We calculate these quantities directly from an input degree sequence, so they are not model parameters. For undirected and unipartite networks, $\theta_i^l = d_i^l / \langle d^l \rangle$, where $i \in \mathcal{N}$ and $\langle d^l \rangle$ is the mean degree in layer l . This normalization recovers the model in Sec. II B when $\theta_i^l = 1$ (i.e., when all nodes have the same degree). The model parameters p_{rs}^1 , p_{rs}^2 , and q_{rs} are now edge “propensities” that, together with the degrees, control the probabilities of edges in the layers.

The probabilities in Eqns. (22)–(23) ensure that, marginally, \mathbf{A}^1 and \mathbf{A}^2 are generated according to mono-

layer DCSBMs [33]. It is not obvious how to model the joint probability $\mathbb{P}(A_{ij}^1 = 1, A_{ij}^2 = 1)$. In particular, it is not clear how it should depend on the observed degrees of nodes i and j in layers 1 and 2 [52]. Part of the complication is that there are four such quantities for each node pair (i, j) . The choice from (24) works particularly well when $\rho = 1$ and the normalized degree sequences $\boldsymbol{\theta}^1$ and $\boldsymbol{\theta}^2$ are the same, as it reduces to a single degree-corrected SBM that generates two identical network layers. Another sensible option is to set $\mathbb{P}(A_{ij}^1 = 1, A_{ij}^2 = 1) = \theta_i^1 \theta_j^1 \theta_i^2 \theta_j^2 q_{g_i g_j}$. This choice has the nice property that edges in a particular edge bundle $(r, s) \in \mathcal{B}$ are independent if and only if $q_{rs} = p_{rs}^1 p_{rs}^2$, which matches the independence condition from Sec. II B for the setting without degree correction. However, this second model underperforms the one from (22)–(24) for edge prediction (see Sec. III and [24]). Consequently, for the rest of this paper, we use the model from Eqns. (22)–(24) as our correlated DCSBM.

1. Maximum-likelihood parameter estimates

When writing the log-likelihood for correlated DCSBMs, we can ignore any additive terms that only involve known quantities, such as the normalized degrees θ_i^l . We can thus write

$$\begin{aligned} \mathcal{L} = \sum_{(r,s) \in \mathcal{B}} \sum_{(i,j) \in \mathcal{E}} & \left[A_{ij}^1 A_{ij}^2 \log q_{rs} + A_{ij}^1 (1 - A_{ij}^2) \log \left(p_{rs}^1 - \sqrt{\frac{\theta_i^2 \theta_j^2}{\theta_i^1 \theta_j^1}} q_{rs} \right) + (1 - A_{ij}^1) A_{ij}^2 \log \left(p_{rs}^2 - \sqrt{\frac{\theta_i^1 \theta_j^1}{\theta_i^2 \theta_j^2}} q_{rs} \right) \right. \\ & \left. + (1 - A_{ij}^1)(1 - A_{ij}^2) \log \left(1 - \theta_i^1 \theta_j^1 p_{rs}^1 - \theta_i^2 \theta_j^2 p_{rs}^2 + \sqrt{\theta_i^1 \theta_j^1 \theta_i^2 \theta_j^2} q_{rs} \right) \right] \delta(g_i, r) \delta(g_j, s) + (\text{const.}). \quad (25) \end{aligned}$$

As in Sec. II B, we seek to maximize \mathcal{L} with respect to the parameters p_{rs}^1 , p_{rs}^2 , and q_{rs} by setting the corresponding derivatives to 0. However, degree-corrected models have the crucial complication that node pairs (i, j) in the same edge bundle (r, s) are no longer stochastically equivalent (i.e., the corresponding entries of the adjacency matrix are no longer sampled from independent, identically distributed random variables), so their contributions to the log-likelihood are no longer the same in general. Consequently, the ML equations for correlated DCSBMs involve $\mathcal{O}(N^2/K^2)$ terms, making them more difficult to solve efficiently.

In certain cases, we are able to make some approximations that make these ML equations easier to solve. Recall that $\theta_i^l = 1$ if the degree of node i is equal to the

mean degree in layer l . For $(i, j) \in \mathcal{E}$, we write

$$\begin{aligned} \theta_i^1 \theta_j^1 &= 1 + \varepsilon_{ij}^1, \\ \theta_i^2 \theta_j^2 &= 1 + \varepsilon_{ij}^2. \end{aligned}$$

If the degree distribution is narrow, such that all node degrees are close to the mean degree, then ε_{ij}^1 and ε_{ij}^2 are small parameters (which can be either positive or negative). In this case, a first-order Taylor expansion yields

$$\sqrt{\theta_i^1 \theta_j^1 \theta_i^2 \theta_j^2} = \sqrt{(1 + \varepsilon_{ij}^1)(1 + \varepsilon_{ij}^2)} \approx 1 + \frac{\varepsilon_{ij}^1 + \varepsilon_{ij}^2}{2}. \quad (26)$$

We also calculate

$$\sqrt{\frac{\theta_i^1 \theta_j^1}{\theta_i^2 \theta_j^2}} = \sqrt{\frac{1 + \varepsilon_{ij}^1}{1 + \varepsilon_{ij}^2}} \approx 1 + \frac{\varepsilon_{ij}^1 - \varepsilon_{ij}^2}{2} \quad (27)$$

and

$$\sqrt{\frac{\theta_i^2 \theta_j^2}{\theta_i^1 \theta_j^1}} \approx 1 + \frac{\varepsilon_{ij}^2 - \varepsilon_{ij}^1}{2}. \quad (28)$$

Using the approximations (26)–(28), we expand the first derivatives of \mathcal{L} to first order in ε_{ij}^1 and ε_{ij}^2 . (See [24] for details.) This calculation yields the following system of equations:

$$\begin{aligned} \frac{e_{rs}^{10}}{p_{rs}^1 - q_{rs}} - \frac{e_{rs}^{00}}{1 - p_{rs}^1 - p_{rs}^2 + q_{rs}} + \frac{g_{rs}^{10}}{2} \frac{q_{rs}}{(p_{rs}^1 - q_{rs})^2} - f_{rs}^1 \frac{1 - p_{rs}^2 + q_{rs}/2}{(1 - p_{rs}^1 - p_{rs}^2 + q_{rs})^2} - f_{rs}^2 \frac{p_{rs}^2 - q_{rs}/2}{(1 - p_{rs}^1 - p_{rs}^2 + q_{rs})^2} &= 0, \\ \frac{e_{rs}^{01}}{p_{rs}^2 - q_{rs}} - \frac{e_{rs}^{00}}{1 - p_{rs}^1 - p_{rs}^2 + q_{rs}} + \frac{g_{rs}^{01}}{2} \frac{q_{rs}}{(p_{rs}^2 - q_{rs})^2} - f_{rs}^1 \frac{p_{rs}^1 - q_{rs}/2}{(1 - p_{rs}^1 - p_{rs}^2 + q_{rs})^2} - f_{rs}^2 \frac{1 - p_{rs}^1 + q_{rs}/2}{(1 - p_{rs}^1 - p_{rs}^2 + q_{rs})^2} &= 0, \\ \frac{e_{rs}^{11}}{q_{rs}} - \frac{e_{rs}^{10}}{p_{rs}^1 - q_{rs}} - \frac{e_{rs}^{01}}{p_{rs}^2 - q_{rs}} + \frac{e_{rs}^{00}}{1 - p_{rs}^1 - p_{rs}^2 + q_{rs}} \\ - \frac{g_{rs}^{10}}{2} \frac{p_{rs}^1}{(p_{rs}^1 - q_{rs})^2} - \frac{g_{rs}^{01}}{2} \frac{p_{rs}^2}{(p_{rs}^2 - q_{rs})^2} + \frac{1}{2} \frac{f_{rs}^1(1 + p_{rs}^1 - p_{rs}^2) + f_{rs}^2(1 - p_{rs}^1 + p_{rs}^2)}{(1 - p_{rs}^1 - p_{rs}^2 + q_{rs})^2} &= 0. \end{aligned} \quad (29)$$

In these equations, we defined e_{rs}^{11} , e_{rs}^{10} , e_{rs}^{01} , and e_{rs}^{00} as in Sec. II B. Additionally, we set

$$\begin{aligned} g_{rs}^{10} &= \sum_{(i,j) \in \mathcal{E}} A_{ij}^1 (1 - A_{ij}^2) (\varepsilon_{ij}^2 - \varepsilon_{ij}^1) \delta(g_i, r) \delta(g_j, s), \\ g_{rs}^{01} &= \sum_{(i,j) \in \mathcal{E}} (1 - A_{ij}^1) A_{ij}^2 (\varepsilon_{ij}^1 - \varepsilon_{ij}^2) \delta(g_i, r) \delta(g_j, s), \\ f_{rs}^1 &= \sum_{(i,j) \in \mathcal{E}} (1 - A_{ij}^1) (1 - A_{ij}^2) \varepsilon_{ij}^1 \delta(g_i, r) \delta(g_j, s), \\ f_{rs}^2 &= \sum_{(i,j) \in \mathcal{E}} (1 - A_{ij}^1) (1 - A_{ij}^2) \varepsilon_{ij}^2 \delta(g_i, r) \delta(g_j, s). \end{aligned}$$

We can efficiently calculate all of these quantities from the matrices \mathbf{A}^1 and \mathbf{A}^2 .

The system of equations (29) reduces to the analogous equations for correlated SBMs if we ignore all of the terms that depend on ε_{ij}^l (i.e., the terms that correspond to perturbations of the degrees from their mean values). The zeroth-order solution that we obtain from ignoring these terms provides a good initialization of a numerical algorithm to solve (29) for the parameters p_{rs}^1 , p_{rs}^2 , and q_{rs} . In practice, when using correlated DCS-BMs for edge prediction (see Sec. III), we find that using a first-order approximation to determine p_{rs}^1 , p_{rs}^2 , and q_{rs} gives results that are almost identical to those from the zeroth-order approximation [see Eqns. (15)–(17)]. Consequently, we suggest using these zeroth-order approximation for edge-prediction applications, given that they are straightforward to calculate and have negligible impact on the quality of the results. For large networks, we also obtain a noticeable improvement in calculation speed when using these approximations.

In App. B, we compare the parameters that we estimate using the approximate system of equations (29) with those from the log-likelihood (25). As expected, the quality of the approximation depends on the shape of the

degree distribution, with larger discrepancies between the two approaches for broader degree distributions.

2. Correlation values

For the SBMs without degree correction from Sec. II B, node pairs (i, j) from a given edge bundle (r, s) have the same Pearson correlation ρ_{rs} . This no longer holds for degree-corrected models. Instead, each node pair (i, j) has its own correlation value

$$\begin{aligned} \varrho_{ij} &= \frac{\mathbb{E}[A_{ij}^1 A_{ij}^2] - \mathbb{E}[A_{ij}^1] \mathbb{E}[A_{ij}^2]}{\sigma[A_{ij}^1] \sigma[A_{ij}^2]} \\ &= \frac{\sqrt{\theta_i^1 \theta_j^1 \theta_i^2 \theta_j^2} q_{rs} - \theta_i^1 \theta_j^1 \theta_i^2 \theta_j^2 p_{rs}^1 p_{rs}^2}{\sqrt{\theta_i^1 \theta_j^1 p_{rs}^1 (1 - \theta_i^1 \theta_j^1 p_{rs}^1) \theta_i^2 \theta_j^2 p_{rs}^2 (1 - \theta_i^2 \theta_j^2 p_{rs}^2)}}. \end{aligned} \quad (30)$$

As in our earlier expansions of the ML equations, we approximate ϱ_{ij} to first order in ε_{ij}^1 and ε_{ij}^2 . We obtain

$$\begin{aligned} \varrho_{ij} \approx \rho_{rs} + \rho_{rs} \left(\frac{\varepsilon_{ij}^1}{2} \frac{p_{rs}^1}{1 - p_{rs}^1} + \frac{\varepsilon_{ij}^2}{2} \frac{p_{rs}^2}{1 - p_{rs}^2} \right. \\ \left. - \frac{\varepsilon_{ij}^1 + \varepsilon_{ij}^2}{2} \frac{p_{rs}^1 p_{rs}^2}{q_{rs} - p_{rs}^1 p_{rs}^2} \right). \end{aligned} \quad (31)$$

Ignoring terms that depend on ε_{ij}^1 and ε_{ij}^2 (i.e., terms that correspond to perturbations of the degrees from their mean values), we obtain $\varrho_{ij} \approx \rho_{rs}$. This approximation works especially well when p_{rs}^1 and p_{rs}^2 are also small, such that their respective network layers are sparse.

The case $q_{rs} = p_{rs}^1 p_{rs}^2$ requires separate consideration [53] to avoid dividing by 0. First-order approximations

in ε_{ij}^1 and ε_{ij}^2 for this case give

$$\varrho_{ij} \approx -\frac{\varepsilon_{ij}^1 + \varepsilon_{ij}^2}{2} \sqrt{\frac{p_{rs}^1}{1 - p_{rs}^1}} \sqrt{\frac{p_{rs}^2}{1 - p_{rs}^2}}. \quad (32)$$

In particular, the zeroth-order solution gives $\varrho \approx 0$, in agreement with the SBM without degree correction from Sec. II B.

III. EDGE PREDICTION

The aim of edge prediction (also called “link prediction”) in networks is to infer likely missing edges and/or spurious edges [54]. Edge prediction is useful for filling in incomplete data sets, such as protein-interaction networks (in which edges are often established as a result of costly experiments) [55] or terrorist-association networks (which are typically constructed based on partial knowledge) [56]. In the context of bipartite user-item networks, edge-prediction techniques provide candidates for personalized recommendations.

One can perform edge prediction in either a supervised or an unsupervised fashion. We briefly discuss each of these types of approaches.

Supervised methods rely on models that learn how a specified set of features relates to the presence or absence of edges. Existing methods that take advantage of multilayer structure for edge prediction typically do so through the specification of multilayer features. These include aggregations of monolayer features [57], as well as path-based [58] and neighborhood-based [59, 60] features that consider multiple layers. Although many of these features depend indirectly on the similarity between different layers, none of these methods quantify the level of correlation or use it for edge prediction.

With unsupervised methods, one obtains a ranking of node pairs such that edges are more likely among higher-ranked pairs. Common approaches include ones that are based on probabilistic models and ones that are based on similarity indices (like the Jaccard index or the Adamic-Adar index) [54]. An example of the former for multilayer networks is a method that maps each network layer independently to a hyperbolic space and then uses the hyperbolic distance between nodes in one layer to predict edges in another layer [61]. This work used node-centric notions of correlation and thereby complements the edge-centric perspective of our work. Methods that rely on similarity indices include those that first generate latent states (i.e., so-called “embeddings”) for the nodes and then rank node pairs according to the similarities of these embedding vectors [62, 63]. Tillman et al. [64] used layer-level correlations to combine monolayer similarity indices into a single score. There are also approaches that implicitly take advantage of similarities across layers, such as by extracting common higher-order structures (specifically, subgraphs with three or more nodes) and looking for patterns that differ by exactly one edge [65].

Variants of SBMs are popular choices for unsupervised edge-prediction methods [4, 9, 55, 56], including in multilayer settings [15, 20]. As an example, in a monolayer degree-corrected Bernoulli SBM, the probability that two nodes, i and j , are adjacent according to the model is $\mathbb{P}(A_{ij} = 1) = \theta_i \theta_j p_{g_i g_j}$. The pairs (i, j) for which these probabilities are relatively large but which are not adjacent (i.e., with $A_{ij} = 0$) in the actual network of interest produce a list of likely candidates for missing edges. Similarly, pairs (i, j) for which these probabilities are small but which are adjacent (i.e., with $A_{ij} = 1$) in the actual network may be spurious edges.

A. Edge Prediction Using Correlated Models

There have been several recent attempts to perform edge prediction in multilayer networks [15, 20, 28]. All of these methods use multilayer information to infer mesoscale structures in networks, but then they perform edge prediction independently in each layer, conditioned on the inferred mesoscale structure and any other model parameters. In particular, when using one of these approaches, observing that two nodes are adjacent in one layer has no bearing on their probability to be adjacent in another layer. We aim to use our correlated models to overcome this limitation.

As in our prior discussions, consider a network with two layers with intralayer adjacency matrices \mathbf{A}^1 and \mathbf{A}^2 , and let \mathbf{g} be the shared block structure of these layers. Our goal is to predict edges in the second layer, conditioned on the adjacency structure of the first layer. For each node pair $(i, j) \in \mathcal{E}$, the key quantities to calculate are the probabilities $\mathbb{P}(A_{ij}^2 = 1 | A_{ij}^1 = 1)$ and $\mathbb{P}(A_{ij}^2 = 1 | A_{ij}^1 = 0)$ for i and j to be adjacent in the second layer, conditioned on them either being adjacent or non-adjacent in the first layer. For example, using the correlated Bernoulli SBM from Sec. II B (which has no degree correction), we have

$$\mathbb{P}(A_{ij}^2 = 1 | A_{ij}^1 = 1) = \frac{q_{g_i g_j}}{p_{g_i g_j}^1}, \quad (33)$$

$$\mathbb{P}(A_{ij}^2 = 1 | A_{ij}^1 = 0) = \frac{p_{g_i g_j}^2 - q_{g_i g_j}}{1 - p_{g_i g_j}^1}. \quad (34)$$

This set of probabilities is the same across all node pairs (i, j) from the same edge bundle (r, s) . Now suppose that we have a positive correlation in this edge bundle, so $\rho_{rs} > 0$. From the definition of the Pearson correlation, it follows that $q_{rs} > p_{rs}^1 p_{rs}^2$. We then find that $\mathbb{P}(A_{ij}^2 = 1 | A_{ij}^1 = 1) > p_{rs}^2$ and $\mathbb{P}(A_{ij}^2 = 1 | A_{ij}^1 = 0) < p_{rs}^2$, whereas using a monolayer SBM would entail that $\mathbb{P}(A_{ij}^2 = 1) = p_{rs}^2$. Therefore, the effect of incorporating correlations into our edge-prediction model when these correlations are positive is (1) to increase the probability that nodes i and j are adjacent in the second layer when the corresponding edge also exists in the first layer and (2) to decrease this probability when the corresponding

edge is absent from the first layer. The effect is reversed for negative correlations.

In Table I, we summarize the two key probabilities (33)–(34) for four different correlated models, alongside the probabilities $\mathbb{P}(A_{ij}^2 = 1)$ for monolayer SBMs and DCSBMs. We include a correlated “configuration model” (CM) [66], which is a special case of the degree-corrected SBM from Sec. II C when there is only one block. (Alternatively, one can think of correlated CMs as extensions of correlated ER models that incorporate degree correction.) We use all of these models for edge prediction in synthetic networks in Sec. III B and in consumer–product networks in Sec. IV B. In particular, the two monolayer models are baselines that we hope to outperform using our correlated models.

B. Tests on Synthetic Networks

We use K -fold [67] cross-validation to assess the performance of the models from Table I on the edge-prediction task. In machine learning, this is an effective way to measure predictive performance [54]. After partitioning a given data set into K parts, one fits a model to $K - 1$ of these subsets and uses it to make predictions on the remaining (i.e., “holdout”) set. One uses each subset once as a holdout, so one does this process K times in total. For our problem, we perform 5-fold cross-validation (which is a standard choice in the machine-learning literature) by splitting the data in the second layer of a given network into 5 subsets. Effectively, this consists of hiding 20% of the entries of the adjacency matrix \mathbf{A}^2 , such that we do not know whether they are edges or not. We then train a model on 100% of the entries of \mathbf{A}^1 and 80% of the entries of \mathbf{A}^2 , and we use it to make predictions about the 20% holdout data from \mathbf{A}^2 . We do this 5 times to cover each choice of holdout data.

A common way to assess the performance of a binary classification model (i.e., a model that assigns one of two possible values to test data) is by using a receiver operating characteristic (ROC) curve. An ROC curve plots the true-positive rate (TPR) of a classifier versus the false-positive rate (FPR) for various choices of a threshold. Many models — including those that are used for edge prediction in networks — make probabilistic predictions, so specifying a threshold is necessary to convert these into binary predictions. Lowering the threshold increases both the TPR and the FPR. A model has predictive power if the former grows faster than the latter, such that the entire ROC curve lies above the diagonal line $\text{TPR} = \text{FPR}$, which gives the performance of a random classifier. As a single summary measure of a model’s predictive performance, it is common to report the area under an ROC curve (AUC). Larger AUC values are better, with a value of $\text{AUC} = 0.5$ indicating equal success as random guessing and $\text{AUC} = 1$ corresponding to perfect prediction. Even in the latter case, one still needs to determine a choice of threshold that completely separates

true positives from false positives.

The AUC is not the only possible quantity to assess edge-prediction performance, although it is very common [15, 68]. In many networks, the number of edges is much smaller than the number of non-edges (i.e., pairs of nodes that are not adjacent). In situations with such an imbalance, the area under the precision–recall (PR) curve is more sensitive than the AUC to variations in model performance. Nevertheless, we use the AUC because it has an intuitive interpretation (specifically, as the probability that the underlying model ranks a true positive above a true negative) that allows us to establish the result in App. C. Additionally, the conclusions that we draw in Sec. IV B about the performance of different models do not change significantly if we use PR curves rather than ROC curves.

We now describe how to generate synthetic network benchmarks that are suitable for testing the models in Table I. We construct these networks so that they have two tunable parameters: the Pearson correlation $\rho \in [-1, 1]$ and a community-mixing parameter $\mu \in [0, 1]$ that controls the strength of the planted mesoscale structure. See Bazzi et al. [14] for more details about the definition of μ . One can also explicitly control the degree distribution, such as by including a parameter η_k for the slope of a truncated power law (e.g., as used in [14] to sample a degree sequence in each layer). For the experiments in this section, we fix $\eta_k = -2$ and use a minimum degree of $k_{\min} = 10$ and a maximum degree of $k_{\max} = 50$. It would be interesting to explore the performance gap between degree-corrected models and models without degree correction as one varies η_k , k_{\min} , and k_{\max} , although we do not do so in the present paper. Lastly, for our numerical experiments in this section, we use networks with $N = 2000$ nodes in each layer and $n_c = 5$ communities, with community sizes sampled from a flat Dirichlet distribution (i.e., one with $\theta = 1$ in the notation of [14]).

We examine two versions, which we call CORRSBM and CORRDCSBM, of a correlated benchmark that is parametrized by the correlation ρ and the community-mixing parameter μ . For both versions of the benchmark, we generate the (undirected and unipartite) adjacency matrix \mathbf{A}^1 of the first layer in the same way. Specifically, given μ and degree-distribution parameters η_k , k_{\min} , and k_{\max} , we use the code from [69] to generate \mathbf{A}^1 and its associated block structure \mathbf{g} . We fit a monolayer model — either an SBM or a DCSBM, depending on the selected version of the benchmark — to \mathbf{A}^1 to obtain the marginal edge propensities \mathbf{p}^1 for the first layer. We then choose \mathbf{p}^2 in one of two ways. For $\rho \in [0, 1]$, we set $\mathbf{p}^2 = \mathbf{p}^1$, which ensures that we can generate networks with correlations that cover the entire range from 0 to 1. For $\rho \in [-1, 0]$, we set $\mathbf{p}^2 = \mathbf{1} - \mathbf{p}^1$, where $\mathbf{1}$ is a matrix with all entries equal to 1; this ensures that we can generate networks with correlations that cover the entire range from -1 to 0. Given \mathbf{p}^1 , \mathbf{p}^2 , and ρ , we then determine \mathbf{q} using either the correlated SBM of Sec. II B or the correlated DCSBM of Sec. II C. For the CORRDC-

TABLE I: Edge-prediction probabilities for various correlated multilayer and monolayer network models.

Model	$\mathbb{P}(A_{ij}^2 = 1 A_{ij}^1 = 1)$	$\mathbb{P}(A_{ij}^2 = 1 A_{ij}^1 = 0)$
Corr. ER	q/p_1	$(p_2 - q)/(1 - p_1)$
Corr. SBM	q_{rs}/p_{rs}^1	$(p_{rs}^2 - q_{rs})/(1 - p_{rs}^1)$
Corr. CM	$\sqrt{\theta_i^1 \theta_j^1 \theta_i^2 \theta_j^2 q} / (\theta_i^1 \theta_j^1 p_1)$	$(\theta_i^2 \theta_j^2 p_2 - \sqrt{\theta_i^1 \theta_j^1 \theta_i^2 \theta_j^2 q}) / (1 - \theta_i^1 \theta_j^1 p_1)$
Corr. DCSBM	$\sqrt{\theta_i^1 \theta_j^1 \theta_i^2 \theta_j^2 q_{rs}} / (\theta_i^1 \theta_j^1 p_{rs}^1)$	$(\theta_i^2 \theta_j^2 p_{rs}^2 - \sqrt{\theta_i^1 \theta_j^1 \theta_i^2 \theta_j^2 q_{rs}}) / (1 - \theta_i^1 \theta_j^1 p_{rs}^1)$
SBM	p_{rs}^2	
DCSBM	$\theta_i^2 \theta_j^2 p_{rs}^2$	

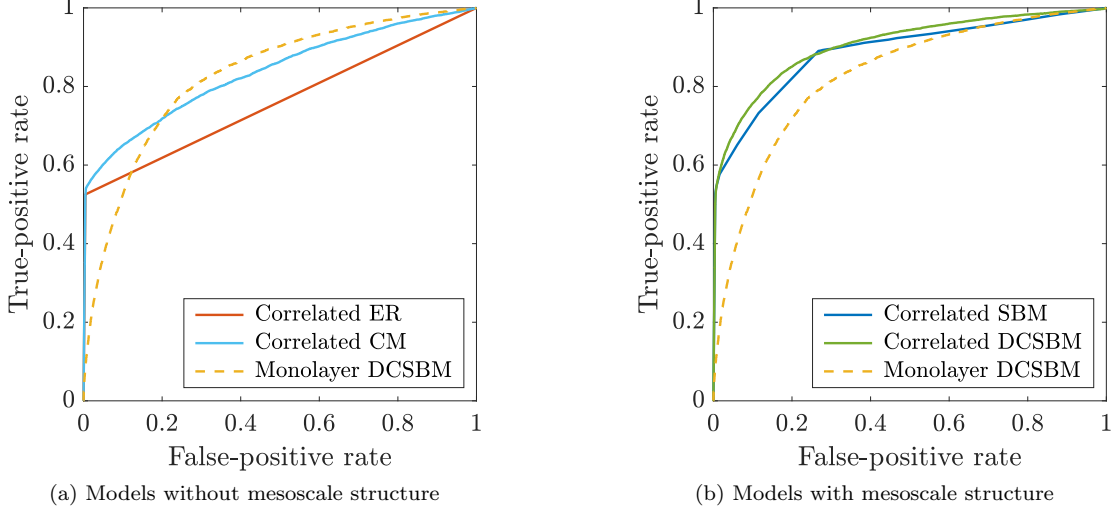


FIG. 3: ROC curves from a 5-fold cross-validation for a network that we sample from the CORRDCSBM benchmark with community-mixing parameter $\mu = 0.3$ and correlation $\rho = 0.5$. (a) Correlated models that do not incorporate any mesoscale structure compared to a monolayer DCSBM baseline. The baseline gives $\text{AUC} \approx 0.83$, whereas the AUC values for the two correlated models are approximately 0.76 (correlated ER) and 0.83 (correlated CM). (b) Correlated models that incorporate mesoscale structure compared to the same monolayer DCSBM baseline. The baseline again gives $\text{AUC} \approx 0.83$, and the AUC values for the two correlated models are approximately 0.89 (correlated SBM) and 0.91 (correlated DCSBM).

CSBM benchmark with $\rho \geq 0$, we set the normalized degrees θ_i^2 to be equal to the corresponding quantities θ_i^1 from the first layer. Again, this choice ensures that we can generate networks all the way to $\rho = 1$. (We also implemented a version of this benchmark that samples degrees independently in the second layer, and we found qualitatively similar results.) The final step consists of generating \mathbf{A}^2 given \mathbf{A}^1 , the propensities \mathbf{p}^2 and \mathbf{q} , and (for the CORRDCSBM benchmark only) the normalized degree sequences $\boldsymbol{\theta}^1$ and $\boldsymbol{\theta}^2$ for both layers. To perform this step, we first compute edge probabilities using either of the correlated models from Secs. IIB and IIC, and we then generate edges independently according to these probabilities.

We now present our results for the two variants of the benchmark. In Fig. 3, we show sample ROC curves for one network that we create using the CORRDCSBM benchmark with $\mu = 0.3$ and $\rho = 0.5$. We compare the

performance of our correlated models with a monolayer DCSBM baseline, which performs edge prediction using only information from the second network layer. Two of the correlated models outperform this baseline, and the correlated CM performs comparably well (i.e., it has a similar AUC).

In Fig. 4, we show results for the CORRSBM benchmark for two choices of the community-mixing parameter μ and several values (both positive and negative) of the Pearson correlation ρ . As expected, the AUC values for monolayer SBMs are independent of ρ , whereas the predictive performance of correlated ER models and correlated SBMs improves as we increase $|\rho|$. In particular, when $|\rho| = 1$, the two correlated models make perfect predictions. When $\rho = 0$, the performance of the correlated ER model is indistinguishable from chance (because $\text{AUC} \approx 0.5$), whereas correlated SBMs have identical performance to monolayer SBMs. The gap between

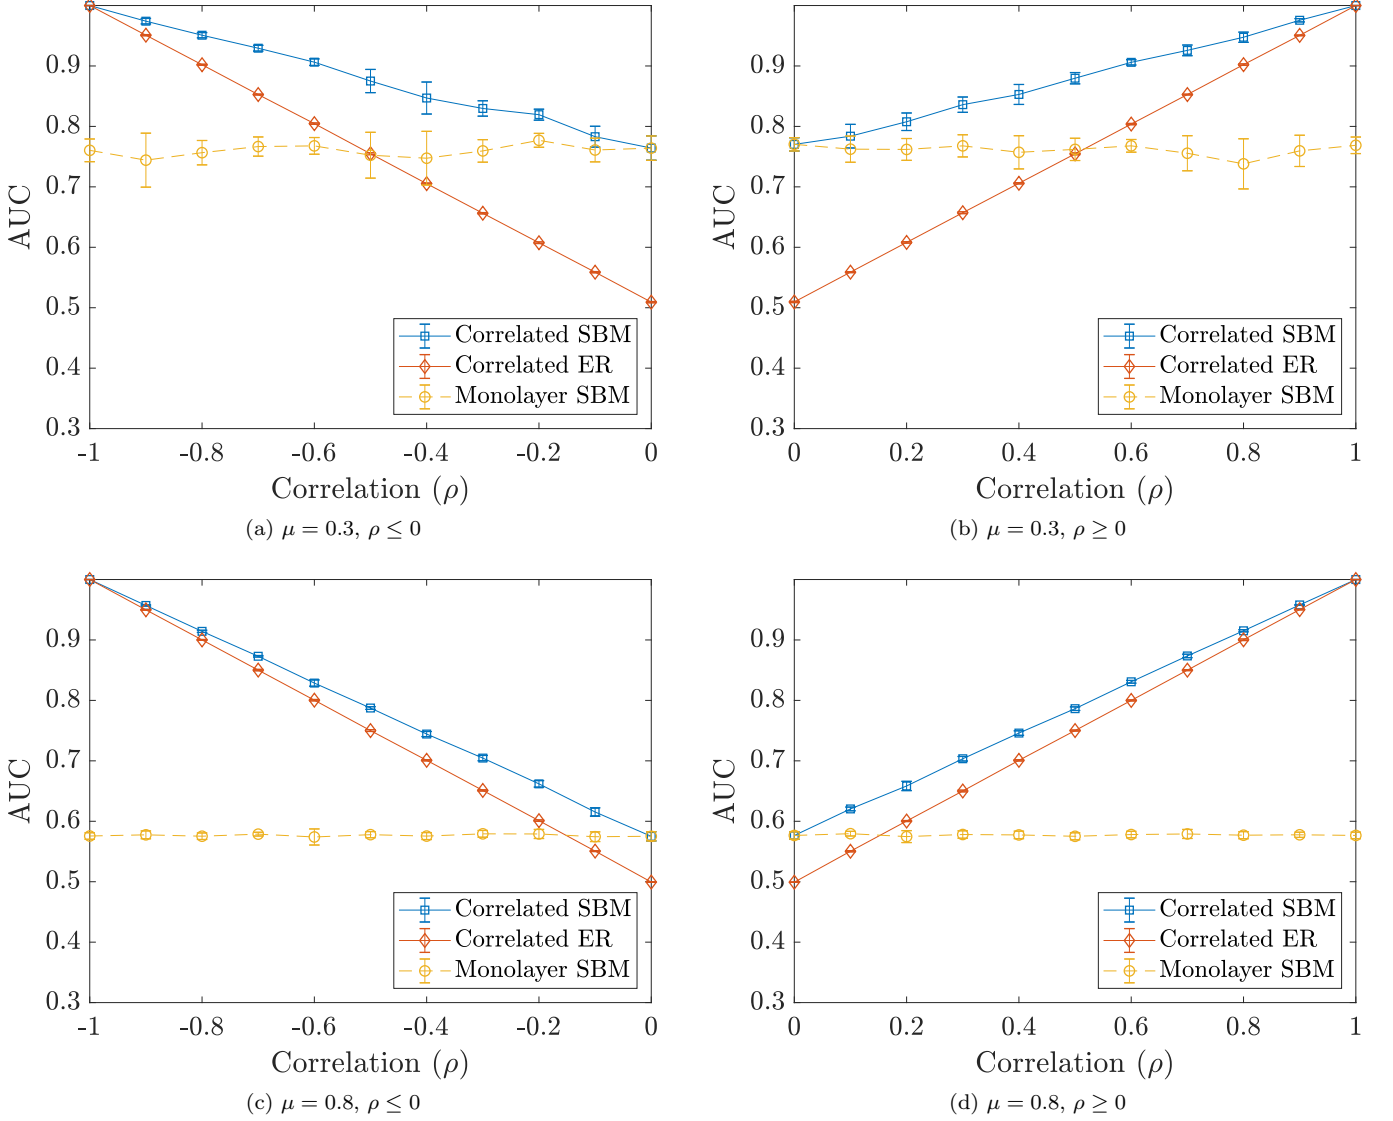


FIG. 4: Edge-prediction results on synthetic networks from the CORRSBM benchmark with (left) $\rho \leq 0$ and (right) $\rho \geq 0$ using two choices of the community-mixing parameter μ . We use $\mu = 0.3$ in the top row and $\mu = 0.8$ in the bottom row. In all plots, along the horizontal axis, we vary the correlation ρ that we use to generate network instances. On the vertical axis, we indicate the AUC for 5-fold cross-validation using a monolayer SBM (dashed curves) and correlated SBM and ER models (solid curves). Each data point is a mean across 10 trials, and the error bars correspond to one standard deviation from that mean. As expected, the AUC does not change with ρ for the monolayer model, but it increases with $|\rho|$ for the two correlated models. For progressively larger μ , for which the sampled networks have progressively weaker mesoscale structure, there is a smaller performance gap between correlated ER models and correlated SBMs.

the two correlated models is smaller for $\mu = 0.8$ than for $\mu = 0.3$, because the underlying block structure is weaker in the former case than in the latter. The AUC of the monolayer baseline is also smaller for $\mu = 0.8$ than for $\mu = 0.3$.

One striking feature in Fig. 4 is that all curves are approximately straight lines (to within sampling error). It makes sense that the performance of monolayer SBMs does not vary with ρ , as these models do not use any

information from the other layer, but the linear dependence on ρ of the other two curves is less intuitive. For the correlated ER model, we can establish rigorously (see App. C) that the AUC is approximately equal to $(1 + |\rho|)/2$ when $p_1 \approx p_2$ or when $p_1 \approx 1 - p_2$. Given that the correlated SBM curves from Fig. 4 also exhibit a linear dependence on ρ , we believe that it is possible to establish similar results for correlated models that incorporate mesoscale structure. These results have practical

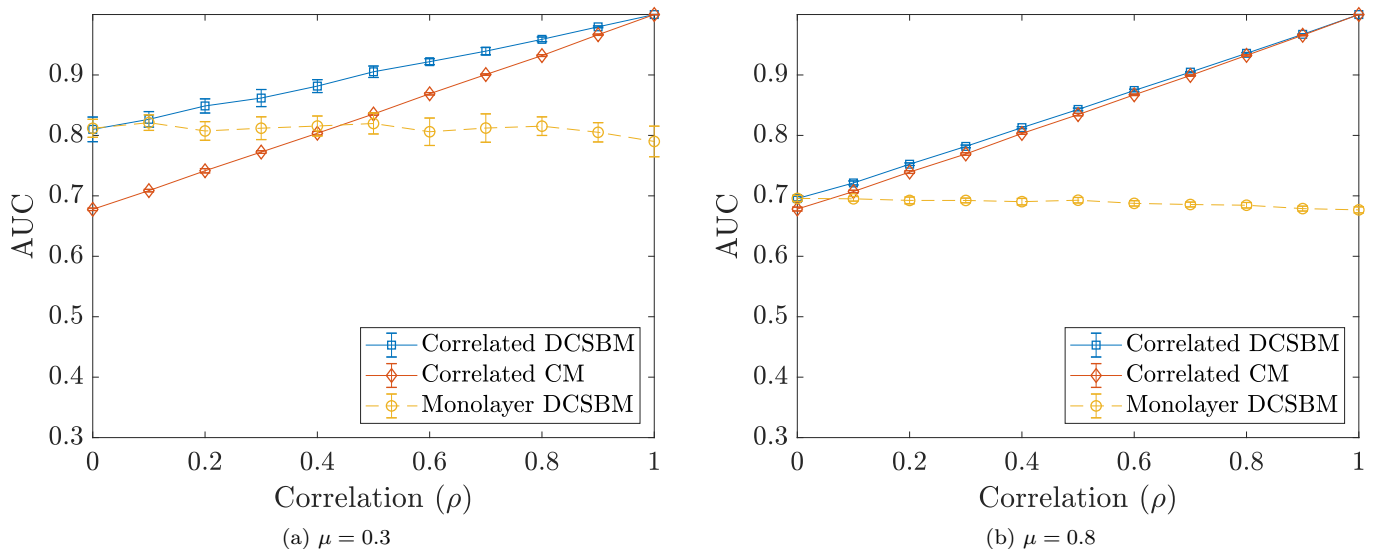


FIG. 5: Edge-prediction results on synthetic networks from the CORRDCSBM benchmark with correlation $\rho \geq 0$ and community-mixing parameters of (a) $\mu = 0.3$ and (b) $\mu = 0.8$. In both panels, along the horizontal axis, we vary the correlation ρ that we use to generate network instances. On the vertical axis, we indicate the AUC for 5-fold cross-validation using a monolayer DCSBM (dashed curves) and the correlated DCSBM and CM (solid curves). Each data point is a mean across 10 trials, and the error bars correspond to one standard deviation from that mean. As expected, the AUC is roughly independent of ρ for the monolayer model, but it increases with ρ for the two correlated models. As we increase μ , such that the sampled networks have progressively weaker mesoscale structure, we observe a substantial narrowing of the performance gap between the correlated CMs and the correlated DCSBMs.

importance, as they allow one to quickly estimate the additional benefits of using correlated models instead of monolayer SBMs for edge prediction.

In Fig. 5, we show results for the CORRDCSBM benchmark for two choices of the community-mixing parameter μ and nonnegative [70] values of the Pearson correlation ρ . As expected, when $\rho = 0$, correlated DCSBMs perform similarly to monolayer DCSBMs. As in Fig. 4, the performance of the monolayer model is roughly independent of ρ , whereas the two correlated models perform better as ρ increases. The gap between the two correlated models narrows substantially as one increases μ from 0.3 to 0.8.

IV. APPLICATIONS

We discuss two applications of correlated multilayer-network models to the analysis of empirical networks. In Sec. IV A, we report pairwise layer correlations for several multiplex networks of different sizes. In Sec. IV B, we consider a temporal bipartite network of customers and products. Using an approach similar to that from Sec. III, we demonstrate that correlated multilayer models have a better edge-prediction performance than monolayer baselines.

A. Layer Correlations in Empirical Networks

We now calculate pairwise layer correlations using the formula (21). Recall that this expression gives the effective correlation between two layers, assuming that they have identical block structures (although their edge-propensity parameters can be different). Crucially, this calculation does *not* require that one first determines the underlying block structure. In fact, as we demonstrated in Sec. II B, the effective correlation of a correlated SBM recovers the correlation of a correlated ER graph, which is straightforward to compute. Accounting for node degrees, as we did in Sec. II C for correlated DCSBMs, significantly increases the complexity of such a calculation. Additionally, as we showed in Sec. II C, correlations using a degree-corrected model are rather similar to those that one obtains without degree correction.

In Table II, we report the mean pairwise layer correlation for 9 multiplex networks. (See App. D for descriptions of these data sets.) To provide additional insight into these networks, we also report the two layers with the largest effective correlations.

We make a few observations about some of the results in Table II. For the *C. elegans* connectome, the layers that correspond to two types of chemical synapses are highly correlated with each other, and their correlation to the layer of electrical synapses is comparatively lower. For the European Union (EU) air transportation

network, the two most correlated layers are those that correspond to Scandinavian Airlines and Norwegian Air Shuttle flights; this is consistent with the findings in [38] (which were based on a different method for quantifying layer similarity). For the network of arXiv collaborations between network scientists, the two most similar categories are “physics.data-an” (which stands for “Data Analysis, Statistics and Probability”) and “cs.SI” (which stands for “Social and Information Networks”). One hypothesis is that these two labels are often used together in papers; such common usage results in a large edge overlap between the corresponding layers and hence in a large correlation value.

To quantify edge correlations at a more granular level, one can first infer block assignments \mathbf{g} and then calculate correlations ρ_{rs} between all block pairs (r, s) . One possible finding from such a calculation may be that correlations between Scandinavian Airlines and Norwegian Air Shuttle routes are significantly larger in certain geographical regions than in others.

B. Edge Prediction in Shopping Networks

The data-science company dunnhumby gave us access to “pseudonymized” transaction data from stores of a major grocery retailer in the United Kingdom. The data were pseudonymized by replacing personally identifiable information with numerical IDs, rendering it impossible to identify individual shoppers. For our analysis, we aggregate transactions over fixed time windows to construct bipartite networks of customers and products. We refer to these structures as “shopping networks”. Because some purchases occur in higher volumes than others, it is useful to incorporate edge weights. Given a customer i and a product j , the *item-penetration* weight is equal to the fraction of all of the items purchased by customer i that are product j . The *basket-penetration* weight is equal to the fraction of all baskets (i.e., distinct shopping trips) of customer i that include product j . See the doctoral dissertation [24] for more details about these weighting schemes.

We now apply the edge-prediction methodology from Sec. III to temporal shopping networks, in which edges and edge weights can change from changes in shopping behavior, with a fixed set of customers and a fixed set of products. We construct networks with two layers, which cover the three-month time periods of March–May 2013 and June–August 2013, respectively. Using the same underlying transaction data, we construct two networks for which we determine the vector \mathbf{g} of block assignments (the same one for both layers [78]) in different ways. For the first network (which we call SHOPPINGMOD), we use basket-penetration weights for the edges and apply multilayer modularity maximization [79, 80] to the weighted network to determine community assignments \mathbf{g} [81]. For the second network (which we call SHOPPINGSBM), we initially calculate item-penetration weights, and we then

apply a threshold to remove edges whose weight is below the median weight (i.e., approximately 50% of the edges). We fit a degree-corrected SBM to the resulting unweighted network using the belief-propagation algorithm from [82]. We expect better edge-prediction performance for the second network, because we detect its block structure using an SBM (as opposed to using modularity maximization, which is more restrictive).

As with our tests on synthetic networks in Sec. III B, we use 5-fold cross-validation to assess edge-prediction performance. We summarize the AUC values of our various correlated multilayer models and the monolayer baselines in Table III, and we show sample ROC curves in Fig. 6. We make a few observations about these results. First, our approximation $\text{AUC} \approx (1 + \rho)/2$ for correlated ER models is very accurate for these two networks, whose correlations are approximately 0.44 and 0.48, respectively. Second, our correlated multilayer models outperform the monolayer baselines for both networks. In particular, the very simple correlated ER model — which assigns one of two probabilities to edges, as indicated in Table I — performs about as well as the more sophisticated monolayer DCSBM for the SHOPPINGMOD network. Third, as expected, AUC values are systematically larger for SHOPPINGSBM than for SHOPPINGMOD. Finally, although incorporating mesoscale structure leads to better performance when there is no degree correction, this does not seem to be the case for degree-corrected models, as correlated DCSBMs do not perform significantly better than correlated CMs. This is also apparent in Figs. 6(b,d), where we observe almost identical ROC curves for the two models. This result suggests that, for some networks, taking into account layer correlations and degree heterogeneity alleviates the need to also consider mesoscale structure when performing edge prediction. This observation has practical implications, as a correlated CM is much easier than a correlated DCSBM to fit to data and to use for edge prediction. However, for recommendation systems, there are situations in which fitting a correlated DCSBM is beneficial, even if its edge-prediction performance is similar to that of a correlated CM. For instance, one may wish to identify relevant customers for a chosen product, irrespective of how much they buy (i.e., their degree). SBMs are able to distinguish between customers with equal degrees and identify those with the greatest predisposition to buy a particular product, whereas CMs are not.

We also note a result from [83] that a curve lies completely above (i.e., “dominates”) another ROC curve if and only if the same relationship holds for the associated PR curves. This result implies for almost all of the curves in Fig. 6 that the rankings of our models based on AUCs are almost identical to those that we would obtain if we instead base them on the areas under PR curves. The only curves whose ranking when we use PR curves is unclear from this result are the correlated ER and monolayer SBM curves in Fig. 6(c).

It is informative to consider the source of false posi-

TABLE II: Pairwise layer correlations in several multiplex networks.

Domain	Network	Number of layers	Mean correlation	Largest correlation (corresponding layers)
Social	CS Aarhus [71]	5	0.27	0.45 (“work” and “lunch” layers)
	Lazega law firm [72]	3	0.39	0.48 (“advice” and “co-work” layers)
	YouTube [73]	5	0.12	0.20 (“shared subscriptions” and “shared subscribers”)
Biological	<i>C. elegans</i> connectome [74]	3	0.47	0.85 (“MonoSyn” and “PolySyn” layers)
	<i>P. falciparum</i> genes [75]	9	0.08	0.25 (“HVR7” and “HVR9” layers)
	<i>Homo sapiens</i> proteins [76]	7	0.04	0.29 (“direct interaction” and “physical association”)
Other	FAO international trade [34]	364	0.13	0.74 (“Pastry” and “Sugar confectionery”)
	EU air transportation [22]	37	0.03	0.39 (“Scandinavian Airlines” and “Norwegian Air Shuttle”)
	ArXiv collaborations [77]	13	0.07	0.73 (“physics.data-an” and “cs.SI”)

TABLE III: Predictive performance of different models on the shopping data set, as measured by the AUC.

Model	AUC (SHOPPINGMOD)	AUC (SHOPPINGSBM)
Monolayer SBM	0.549	0.633
Correlated ER	0.724	0.743
Correlated SBM	0.742	0.793
Monolayer DCSBM	0.725	0.797
Correlated CM	0.817	0.870
Correlated DCSBM	0.818	0.875

tives in Fig. 6. Because the correlation between layers is positive, we are likely to predict an edge where none exists when all of the following conditions hold: (1) nodes (i, j) are adjacent in one layer but not the other; (2) the node pair (i, j) belongs to an edge bundle (r, s) with a large layer correlation ρ_{rs} ; and (3) nodes i and j have large degrees. Note that condition (c) applies only for correlated models with degree correction.

V. CONCLUSIONS AND DISCUSSION

We introduced models of multilayer networks in which edges that connect the same nodes in different layers are not independent. In comparison to models without edge correlations, our models offer an improved representation of many empirical networks, as interlayer correlations are a common phenomenon: flights between major airports are serviced by multiple airlines, individuals interact repeatedly with the same people, consumers often buy the same products over time, and so on. Among other potential applications, one can use our models to improve edge prediction, to study the graph-matching problem on more realistic benchmark networks, and to calculate layer correlations as insightful summary statistics for networks.

To model layer correlations, we used bivariate Bernoulli random variables to generate edges simultaneously in two network layers. (See [24] for derivations using Poisson random variables.) Correlated Bernoulli stochastic block models were proposed previously [31],

although only as forward models for generating networks, rather than for performing inference given empirical data. Another key contribution of our work is a degree-corrected variant of such a model. The maximum-likelihood equations are significantly more difficult to solve in this case, but we were able to make useful simplifications with suitable approximations. Notably, these simplified equations closely approximate those for models without degree correction for networks with almost homogeneous degree distributions.

The models in the present paper that incorporate some mesoscale structure \mathbf{g} assume that such structure is given. This setup has the benefit that one can use any desired algorithms to produce a network partition, including ones that operate on weighted or annotated networks or that use nonstandard null models in a modularity objective function. This makes our approach for analyzing correlations suitable for a wide variety of applications.

Fitting a correlated SBM to network data yields a correlation value ρ_{rs} for each edge bundle (r, s) . We have defined an effective correlation that combines all of these values into a single measure of similarity between two layers. Notably, the value of the effective correlation is independent of a network’s mesoscale structure, making it extremely easy to compute (see Eqn. (12)). We illustrated this method of assessing layer similarity for multiplex networks from social, biological, and other domains.

Another application of our work is to edge prediction in multilayer networks. Our numerical experiments revealed that simple correlated models (e.g., a correlated configuration model or a correlated SBM without degree correction) can outperform monolayer DCSBMs in terms of AUC even for moderate correlation values. We also observed such improved performance for consumer-product networks, which have significant layer correlations ($\rho \approx 0.45$). We expect that a correlated multilayer DCSBM will typically outperform a monolayer DCSBM for most empirical networks, even when there are lower levels of correlation.

There are many interesting ways to build further on our work. For example, it would be useful to be able to model all layers simultaneously, rather than in a pairwise fashion, especially for multiplex networks (in which layers do not have a natural ordering). One challenge is that




FIG. 6: ROC curves for the edge-prediction task using 5-fold cross-validation on two temporal networks, SHOPPINGMOD and SHOPPINGSBM. We consider both degree-corrected models and models without degree correction. The dotted diagonal line in each plot indicates the expected ROC curve for a random classifier. All other curves lie above this diagonal line, suggesting that they have predictive power. In all cases, correlated multilayer models outperform their monolayer counterparts. For both networks, correlated SBMs outperform correlated ER models, whereas correlated DCSBMs perform similarly to correlated CMs [as illustrated by the almost overlapping curves in panels (b) and (d)].

a multivariate Bernoulli distribution of dimension L has $2^L - 1$ parameters; this grows quickly with the number L of layers. For a temporal setting, we have proposed generating correlated networks in a sequential way by conditioning each layer on the previous one. In some cases, it will be useful to relax this “memoryless” assumption and condition a layer on all previous layers, rather than only on the most recent layer. For example, the purchases of shoppers in December one year are strongly related not

only to their purchases in November, but also to what they bought in December during the previous year.

In the present paper, we have not considered the case of nonidentical mesoscale structure across layers. (See [24] for a possible approach.) Additionally, although one can account for edge weights when fitting a block structure \mathbf{g} to use with our models, the rest of our derivations apply only to unweighted networks. Modeling correlated weighted networks entails prescribing both edge-

existence and edge-weight correlations. Yet another idea is to derive correlated models for networks with overlapping communities. Previous research [15] suggests that this can substantially improve edge-prediction performance. For correlated DCSBMs, for example, it would be useful to learn the dependence of the joint probability $\mathbb{P}(A_{ij}^1 = 1, A_{ij}^2 = 1)$ on node degrees, rather than assume the parametric form in Eqn. (24). This is a challenging problem for which maximum-likelihood estimation is unlikely to be a suitable tool, so it falls outside the scope of the present paper.

Our work is also a starting point for designing algorithms to detect correlated communities in networks by inferring \mathbf{g} alongside other model parameters. A practical outcome of such an algorithm would be a set of communities that persist across layers if and only if the edges in those communities are sufficiently highly correlated with each other. Such an approach would offer a new interpretation of what it means for a community to span multiple layers [41, 84].

Lastly, although we performed edge prediction in an unsupervised manner, it is also possible to use estimated correlations as features in supervised models [57–60] to improve their performance.

In summary, our work highlights the importance of relaxing edge-independence assumptions in statistical models of network data. Doing so provides richer insights into the structure of empirical networks, improves edge-prediction performance, and yields more realistic models on which to test community-detection, graph-matching, and other types of algorithms.

ACKNOWLEDGMENTS

ARP was funded by the EPSRC Centre for Doctoral Training in Industrially Focused Mathematical Modelling (EP/L015803/1) in partnership with dunnhumby. We are grateful to dunnhumby for providing access to grocery-shopping data and to Shreena Patel and Rosie Prior for many helpful discussions. We also thank Junbiao Lu for helpful comments.

Appendix A: Variance of ML Estimates

In Sec. II, we derived ML parameter estimates for various types of correlated network models. In this appendix, we show how to obtain the variance of these estimates and we illustrate how these variances scale with network size (i.e., the number of nodes). We focus on correlated ER models and the corresponding log-likelihood (8). The same approach also works for the other types of models that we examined.

Let $\boldsymbol{\beta} = [p_1 \ p_2 \ q]^\top$ denote the vector of parameters for a correlated ER model. Under mild conditions [85], the ML estimate $\hat{\boldsymbol{\beta}}$ converges in distribution to a multivariate normal distribution as the number N of nodes

tends to infinity. For large but finite N , the quantity $\hat{\boldsymbol{\beta}}$ is distributed approximately according to $\mathcal{N}(\boldsymbol{\beta}^*, \mathcal{I}^{-1}(\boldsymbol{\beta}^*))$, where $\boldsymbol{\beta}^*$ is the true-parameter vector and $\mathcal{I}^{-1}(\boldsymbol{\beta}^*)$ is the inverse of the Fisher information matrix evaluated at the parameter values in $\boldsymbol{\beta}^*$. The diagonal entries of $\mathcal{I}^{-1}(\boldsymbol{\beta}^*)$ provide variance estimates for p_1 , p_2 , and q .

To illustrate how these variance estimates scale with the number N of nodes, we simulate correlated ER networks with $p_1 = 0.1$, $p_2 = 0.085$, and $q = 0.05$ (which correspond to a correlation of $\rho = 0.5$) for different network sizes. In Fig. 7, we plot the 95% confidence intervals around the ML estimates for each of the three parameters in our model. We find empirically that the variance scales with $1/N^2$, so the standard deviation (and thus the width of the 95% confidence intervals) scales with $1/N$.




FIG. 7: ML estimates with 95% confidence intervals that we obtain from the inverse of the Fisher information matrix for different values of the number N of nodes. The true parameter values (dashed horizontal lines) are $p_1 = 0.1$, $p_2 = 0.085$, and $q = 0.05$. The width of the confidence intervals scales approximately as $1/N$.

Appendix B: Maximizing the Full Log-Likelihood Versus our Approximate Log-Likelihood for Correlated DCSBMs

In Sec. IIC, we derived an approximation (29) of the log-likelihood for correlated DCSBMs that enables a more efficient estimation of the parameters than maximizing the exact log-likelihood (25). We now use simulated data to illustrate how the results that one obtains using this approximation compare with maximizing the original likelihood.

We simulate networks using the CORRDCSBM benchmark (see Sec. IIIB) with $N = 1000$ nodes, $K = 5$ communities, a mixing parameter of $\mu = 0.3$, and a correlation of $\rho = 0.5$. For $N = 2000$ nodes, which we used in other experiments in the paper, we find that maximizing the full log-likelihood is prohibitively slow. Even for $N = 1000$, obtaining estimates from the full

log-likelihood takes about 30 minutes in total, whereas the calculation runs in about 5 seconds on a typical laptop when we use the approximate log-likelihood.

The quality of our approximation depends strongly on the shape of the degree distribution: accuracy degrades for distributions with larger variances and heavier tails. We illustrate this behavior in Fig. 8, where we plot the full-likelihood estimates [see Eqn. (25)] versus the approximate-likelihood estimates [see Eqn. (29)] for two networks. For both networks, we sample degrees from truncated power-law distributions using the code in [69]. In Fig. 8(a), we choose a relatively narrow degree distribution (with $k_{\text{exp}} = 0$, a minimum degree of 18, and a maximum degrees of 22). In this case, the parameter values that we estimate using the approximate log-likelihood closely match those from the full log-likelihood. For Fig. 8(b), we choose a relatively wide degree distribution (with $k_{\text{exp}} = -2$, a minimum degree of 10, and a maximum degree of 50). In this case, the approximate log-likelihood tends to overestimate parameter values, especially for larger values of these parameters. See the top-right corner of Fig. 8(b).

As this example illustrates, there is a trade-off between accuracy and speed. It may be possible to derive better approximations, such as by considering a second-order expansion instead of a first-order expansion with respect to the quantities ε_{ij}^1 and ε_{ij}^2 or by adding corrections for large-degree nodes to the approximate log-likelihood.

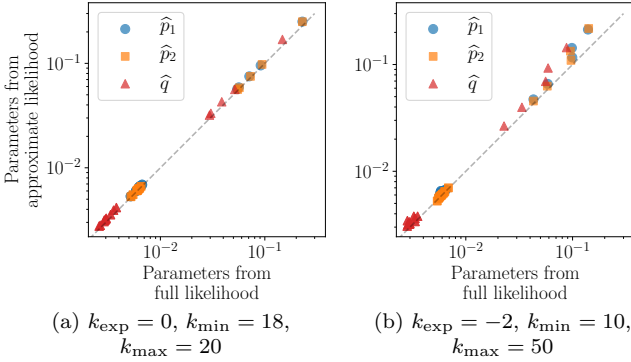


FIG. 8: ML estimates from the approximate log-likelihood (29) versus ML estimates from the full log-likelihood (25). (a) Example of a network with a relatively narrow degree distribution. For this example, the approximation works well. (b) Example of a network with a relatively wide degree distribution. For this example, we observe some discrepancies between the two sets of estimates.

Appendix C: Edge-Prediction AUC as a Function of the Pearson Correlation ρ

We establish the following result.

Proposition. *The AUC for a correlated ER model is an affine function of the Pearson correlation ρ . In particular, when $p_1 \approx p_2$ or $p_1 \approx 1 - p_2$, we have $\text{AUC}_{\text{ER}} \approx (1 + |\rho|)/2$.*

Proof. Suppose that $\rho > 0$. (The case $\rho \leq 0$ is similar.) With a correlated ER model, all unobserved interactions (i, j) in the second layer have one of two probabilities: q/p_1 if $A_{ij}^1 = 1$ and $(p_2 - q)/(1 - p_1)$ if $A_{ij}^1 = 0$. Because $\rho > 0$, we have $q > p_1 p_2$, which implies that $q/p_1 > (p_2 - q)/(1 - p_1)$. Selecting a threshold between these two probabilities amounts to predicting that everything that is an edge in the first layer is also an edge in the second layer and that everything that is not an edge in the first layer is also not an edge in the second layer. Let a and b denote the TPR and FPR, respectively, at such an intermediate threshold. In this case, a and b are the coordinates of the point at which the slope of the ROC curve changes. See the illustration in Fig. 9.




FIG. 9: Diagram of the ROC curve for an ER model, which assigns one of two possible edge probabilities to each pair of nodes.

By straightforward geometry,

$$\text{AUC}_{\text{ER}} = 1 - \frac{ab}{2} - \frac{(1-a)(1-b)}{2} - (1-a)b = \frac{1}{2} + \frac{a-b}{2}$$

is the area under this ROC curve. The next step is to estimate a and b . With a correlated ER model, the number of true positives is proportional [86] to e_{11} ; the model's prediction is correct every time that an edge that is present in the first layer is also present in the second layer. To find the TPR, one needs to divide this quantity by the number of edges (there are $e_{11} + e_{01}$ of them) in the second layer. We obtain

$$a \approx \frac{e_{11}}{e_{11} + e_{01}} \approx \frac{q}{p_2} = p_1 + \rho \sqrt{\frac{1-p_2}{p_2} p_1 (1-p_1)}. \quad (\text{C1})$$

Similarly, every time an edge that is present in the first layer is not present in the second layer counts as an incorrect prediction of the model. Therefore, the number

of false positives is proportional to e_{10} . Dividing this by the number of non-edges in the second layer yields

$$b \approx \frac{e_{10}}{e_{10} + e_{00}} \approx \frac{p_1 - q}{1 - p_2} = p_1 - \rho \sqrt{\frac{p_2}{1 - p_2}} p_1 (1 - p_1). \quad (\text{C2})$$

From (C1) and (C2), it follows that

$$\text{AUC}_{\text{ER}} \approx \frac{1}{2} + \frac{\rho}{2} \sqrt{p_1(1 - p_1)} \left(\sqrt{\frac{1 - p_2}{p_2}} + \sqrt{\frac{p_2}{1 - p_2}} \right),$$

which is an affine function of ρ . When $p_1 \approx p_2$ or $p_1 \approx 1 - p_2$, as is the case in Fig. 3, we obtain $\text{AUC}_{\text{ER}} \approx (1 + |\rho|)/2$, as desired. Using a similar argument, one can show that the same result holds (with the same assumptions on p_1 and p_2) when $\rho \leq 0$. \square

Given that the correlated SBM curves from Fig. 4 also appear to depend linearly on ρ , we believe that it is possible to establish similar results for correlated models that incorporate mesoscale structure.

Appendix D: Data Sets

We provide brief descriptions of the multiplex networks that we analyzed in Sec. IV A. For weighted networks, we disregard edge weights when calculating layer correlations. We downloaded these networks, aside from the YouTube and *P. falciparum* data sets, from <https://comunelab.fbk.eu/data.php>.

1. CS Aarhus

This is an undirected and unweighted social network of offline and online relationships between $N = 61$ members of the Department of Computer Science at Aarhus University [71]. There are $T = 5$ layers: (1) regularly eating lunch together; (2) friendships on Facebook; (3) co-authorship; (4) leisure activities; and (5) working together.

2. Lazega Law Firm

This directed, unweighted network encompasses interactions between $N = 71$ partners and associates who work at the same law firm [72]. The network has $T = 3$ layers that encode co-work, friendship, and advice relationships.

3. YouTube

This is an undirected, weighted network of interactions between $N = 15088$ YouTube users [73]. There are $T = 5$ types of interactions: (1) direct contacts (“friendships”); (2) shared contacts; (3) shared subscriptions; (4) shared subscribers; and (5) shared favorites.

4. *C. elegans* Connectome

This is a directed, unweighted network of synaptic connections between $N = 279$ neurons of the nematode *C. elegans* [74]. There are $T = 3$ layers, which correspond to electric, chemical monadic (“MonoSyn”), and chemical polyadic (“PolySyn”) junctions.

5. *P. falciparum* Genes

This is an undirected, unweighted network of $N = 307$ recombinant genes from the parasite *P. falciparum*, which causes malaria [75]. There are $T = 9$ layers that correspond to distinct highly variable regions (HVRs), in which these recombinations occur. Two genes are adjacent in a layer if they share a substring whose length is statistically significant.

6. *Homo sapiens* Proteins

This is a directed, unweighted network of interactions between $N = 18222$ proteins in *Homo sapiens* [87]. There are $T = 7$ layers, which correspond to the following types of interactions: (1) direct interactions; (2) physical associations; (3) suppressive genetic interactions; (4) association; (5) colocalization; (6) additive genetic interactions; and (7) synthetic genetic interactions. The original data is from BioGRID [76], a public database of protein interactions (for humans as well as other organisms) that is curated from different types of experiments.

7. Food and Agriculture Organization (FAO) Trade

This is a weighted, directed network of food imports and exports during the year 2010 between $N = 214$ countries [34]. There are $T = 364$ layers, which correspond to different food products.

8. European Union Air Transportation

This is an undirected, unweighted network of flights between $N = 450$ airports in Europe [22]. There are $T = 37$ layers, each of which corresponds to a different airline.

9. ArXiv Collaborations

This is an undirected, weighted coauthorship network between $N = 14489$ network scientists [77]. There are $T = 13$ layers, which correspond to different

arXiv subject areas: “physics.soc-ph”, “physics.data-an”, “physics.bio-ph”, “math-ph”, “math.OC”, “cond-

mat.dis-nn”, “cond-mat.stat-mech”, “q-bio.MN”, “q-bio”, “q-bio.BM”, “nlin.AO”, “cs.SI”, and “cs.CV”.

-
- [1] M. E. J. Newman, *Networks*, 2nd ed. (Oxford University Press, Oxford, UK, 2018).
 - [2] C. A. Davis, O. Varol, E. Ferrara, A. Flammini, and F. Menczer, in *Proceedings of the 25th International World Wide Web Conference Companion* (2016) pp. 273–274.
 - [3] G. A. Pagani and M. Aiello, *Physica A: Statistical Mechanics and its Applications* **392**, 2688 (2013).
 - [4] R. Guimerà and M. Sales-Pardo, *PLoS Computational Biology* **9**, e1003374 (2013).
 - [5] S. Fortunato and D. Hric, *Physics Reports* **659**, 1 (2016).
 - [6] P. Csermely, A. London, L.-Y. Wu, and B. Uzzi, *Journal of Complex Networks* **1**, 93 (2013).
 - [7] P. Rombach, M. A. Porter, J. H. Fowler, and P. J. Mucha, *SIAM Review* **59**, 619 (2017).
 - [8] R. A. Rossi and N. K. Ahmed, *IEEE Transactions on Knowledge and Data Engineering* **27**, 1112 (2015).
 - [9] T. P. Peixoto, in *Advances in Network Clustering and Blockmodeling*, edited by P. Doreian, V. Batagelj, and A. Ferligoj (Wiley, Hoboken, NJ, USA, 2020) pp. 289–332.
 - [10] M. Kivelä, A. Arenas, M. Barthélemy, J. P. Gleeson, Y. Moreno, and M. A. Porter, *Journal of Complex Networks* **2**, 203 (2014).
 - [11] S. Boccaletti, G. Bianconi, R. Criado, C. D. Genio, J. Gómez-Gardenes, M. Romance, I. Sendina-Nadal, Z. Wang, and M. Zanin, *Physics Reports* **544**, 1 (2014).
 - [12] M. A. Porter, *Notices of the American Mathematical Society* **65**, 1419 (2018).
 - [13] A. Aleta and Y. Moreno, *Annual Review of Condensed Matter Physics* **10**, 45 (2019).
 - [14] M. Bazzi, L. G. S. Jeub, A. Arenas, S. D. Howison, and M. A. Porter, *Physical Review Research* **2**, 023100 (2020).
 - [15] C. De Bacco, E. A. Power, D. B. Larremore, and C. Moore, *Physical Review E* **95**, 042317 (2017).
 - [16] N. Stanley, S. Shai, D. Taylor, and P. J. Mucha, *IEEE Transactions on Network Science and Engineering* **3**, 95 (2016).
 - [17] A. Ghasemian, P. Zhang, A. Clauset, C. Moore, and L. Peel, *Physical Review X* **6**, 031005 (2016).
 - [18] T. P. Peixoto and M. Rosvall, *Nature Communications* **8**, 582 (2017).
 - [19] T. P. Peixoto, *Physical Review E* **92**, 042807 (2015).
 - [20] T. Vallès-Català, F. A. Massucci, R. Guimerà, and M. Sales-Pardo, *Physical Review X* **6**, 011036 (2016).
 - [21] D. Taylor, S. Shai, N. Stanley, and P. J. Mucha, *Physical Review Letters* **116**, 228301 (2016).
 - [22] A. Cardillo, J. Gómez-Gardenes, M. Zanin, M. Romance, D. Papo, F. Del Pozo, and S. Boccaletti, *Scientific Reports* **3**, 1344 (2013).
 - [23] U. Aslak, M. Rosvall, and S. Lehmann, *Physical Review E* **97**, 062312 (2018).
 - [24] A. R. Pamfil, *Communities in Annotated, Multilayer, and Correlated Networks*, D.Phil. thesis, University of Oxford (2018).
 - [25] C. H. Park and M. Kahng, in *IEEE/ACIS 9th International Conference on Computer and Information Science* (IEEE, 2010) pp. 573–578.
 - [26] T. Yasseri, R. Sumi, A. Rung, A. Kornai, and J. Kertész, *PloS One* **7**, e38869 (2012).
 - [27] C. Aicher, A. Z. Jacobs, and A. Clauset, *Journal of Complex Networks* **3**, 221 (2014).
 - [28] M. Tarrés-Deulofeu, A. Godoy-Lorite, R. Guimerà, and M. Sales-Pardo, *Physical Review E* **99**, 032307 (2019).
 - [29] D. Conte, P. Foggia, C. Sansone, and M. Vento, *International Journal of Pattern Recognition and Artificial Intelligence* **18**, 265 (2004).
 - [30] V. Lyzinski, D. E. Fishkind, and C. E. Priebe, *Journal of Machine Learning Research* **15**, 3513 (2014).
 - [31] V. Lyzinski, D. L. Sussman, D. E. Fishkind, H. Pao, L. Chen, J. T. Vogelstein, Y. Park, and C. E. Priebe, *Parallel Computing* **47**, 70 (2015).
 - [32] V. Lyzinski and D. L. Sussman, *Information and Inference: A Journal of the IMA*, iaz031 (2020).
 - [33] B. Karrer and M. E. J. Newman, *Physical Review E* **83**, 016107 (2011).
 - [34] M. De Domenico, V. Nicosia, A. Arenas, and V. Latora, *Nature Communications* **6**, 6864 (2015).
 - [35] J. Iacovacci, Z. Wu, and G. Bianconi, *Physical Review E* **92**, 042806 (2015).
 - [36] V. Nicosia and V. Latora, *Physical Review E* **92**, 032805 (2015).
 - [37] F. Battiston, V. Nicosia, and V. Latora, *Physical Review E* **89**, 032804 (2014).
 - [38] T.-C. Kao and M. A. Porter, *Journal of Statistical Physics* **173**, 1286 (2018).
 - [39] M. De Domenico and J. Biamonte, *Physical Review X* **6**, 041062 (2016).
 - [40] In fact, previous algorithms for SBM inference [15, 16, 27, 82] have suggested that one can formulate a full inference algorithm using an iterative approach that alternates between estimating a multilayer network’s mesoscale structure and estimating its edge probabilities and correlations. Therefore, the latter task — which is the focus of the present paper — constitutes one half of an algorithm for correlated-community detection.
 - [41] A. R. Pamfil, S. D. Howison, R. Lambiotte, and M. A. Porter, *SIAM Journal on Mathematics of Data Science* **1**, 667 (2019).
 - [42] K.-M. Lee, J. Y. Kim, W.-k. Cho, K.-I. Goh, and I. Kim, *New Journal of Physics* **14**, 033027 (2012).
 - [43] B. Min, S. Do Yi, K.-M. Lee, and K.-I. Goh, *Physical Review E* **89**, 042811 (2014).
 - [44] V. Marceau, P.-A. Noël, L. Hébert-Dufresne, A. Allard, and L. J. Dubé, *Physical Review E* **84**, 026105 (2011).
 - [45] S. Funk and V. A. Jansen, *Physical Review E* **81**, 036118 (2010).
 - [46] J. Y. Kim and K.-I. Goh, *Physical review letters* **111**, 058702 (2013).
 - [47] P. Barucca, F. Lillo, P. Mazzarisi, and D. Tantari, *Journal of Statistical Mechanics: Theory and Experiment* **2018**, 123407 (2018).

- [48] Y. Zhao, E. Levina, and J. Zhu, *The Annals of Statistics* **40**, 2266 (2012).
- [49] J.-G. Young, G. St-Onge, P. Desrosiers, and L. J. Dubé, *Physical Review E* **98**, 032309 (2018).
- [50] P. O. Perry and P. J. Wolfe, arXiv preprint arXiv:1201.5871 (2012).
- [51] The Hessian is negative definite at the critical point that one obtains by setting these derivatives to 0, so this point is a local maximum (as opposed to a local minimum or saddle point) of the log-likelihood function.
- [52] One desiderata of a correlated degree-corrected model is that it satisfies the inequality $\mathbb{P}(A_{ij}^1 = 1, A_{ij}^2 = 1) \leq \mathbb{P}(A_{ij}^1 = 1)$ for all node pairs (i, j) . This requirement rules out some possibilities for $\mathbb{P}(A_{ij}^1 = 1, A_{ij}^2 = 1)$, for instance one where this quantity has no dependence on node degrees.
- [53] Situations when one or more of the denominators in Eqn. (31) are $\mathcal{O}(\varepsilon)$ also require separate derivations. However, we do not treat these exceptional cases in the present paper.
- [54] L. Lü and T. Zhou, *Physica A* **390**, 1150 (2011).
- [55] R. Guimerà and M. Sales-Pardo, *Proceedings of the National Academy of Sciences of the United States of America* **106**, 22073 (2009).
- [56] A. Clauset, C. Moore, and M. E. J. Newman, *Nature* **453**, 98 (2008).
- [57] M. Pujari and R. Kanawati, *Networks & Heterogeneous Media* **10**, 17 (2015).
- [58] M. Jalili, Y. Orouskhani, M. Asgari, N. Alipourfard, and M. Perc, *Royal Society Open Science* **4**, 160863 (2017).
- [59] D. Hristova, A. Noulas, C. Brown, M. Musolesi, and C. Mascolo, *EPJ Data Science* **5**, 24 (2016).
- [60] H. Mandal, M. Mirchev, S. Gramatikov, and I. Mishkovski, in *2018 26th Telecommunications Forum (TELFOR)* (IEEE, 2018) pp. 1–4.
- [61] K.-K. Kleineberg, M. Boguná, M. Á. Serrano, and F. Papadopoulos, *Nature Physics* **12**, 1076 (2016).
- [62] R. Matsuno and T. Murata, in *Companion Proceedings of The Web Conference* (2018) pp. 1261–1268.
- [63] L. Pio-Lopez, A. Valdeolivas, L. Tichit, É. Remy, and A. Baudot, arXiv preprint arXiv:2008.10085 (2020).
- [64] R. E. Tillman, V. K. Potluru, J. Chen, P. Reddy, and M. Veloso, arXiv preprint arXiv:2004.04704 (2020).
- [65] M. Coscia and M. Szell, arXiv preprint arXiv:2008.08351 (2020).
- [66] B. K. Fosdick, D. B. Larremore, J. Nishimura, and J. Ugander, *SIAM Review* **60**, 315 (2018).
- [67] Note that this K is different from the one that denotes the number of blocks elsewhere in this paper.
- [68] A. Ghasemian, H. Hosseinmardi, A. Galstyan, E. M. Airolidi, and A. Clauset, *Proceedings of the National Academy of Sciences of the United States of America* (2019), available at <https://doi.org/10.1073/pnas.1914950117>.
- [69] L. G. S. Jeub and M. Bazzi, A generative model for mesoscale structure in multilayer networks implemented in MATLAB, <https://github.com/MultilayerGM/MultilayerGM-MATLAB> (2016–2019), version 1.0.
- [70] The case $\rho \leq 0$ requires a different way of sampling normalized degrees θ_i^2 of nodes in the second layer to be able to take ρ all the way to $\rho = -1$. More precisely, when ρ is close to -1 , the second layer is dense if the first layer is sparse (and vice versa), so setting $\theta^1 = \theta^2$ is inadequate.
- [71] M. Magnani, B. Micenkova, and L. Rossi, arXiv preprint arXiv:1303.4986 (2013).
- [72] E. Lazega, *The Collegial Phenomenon: The Social Mechanisms of Cooperation Among Peers in a Corporate Law Partnership* (Oxford University Press, Oxford, UK, 2001).
- [73] L. Tang, X. Wang, and H. Liu, in *2009 Ninth IEEE International Conference on Data Mining* (IEEE, 2009) pp. 503–512.
- [74] B. L. Chen, D. H. Hall, and D. B. Chklovskii, *Proceedings of the National Academy of Sciences of the United States of America* **103**, 4723 (2006).
- [75] D. B. Larremore, A. Clauset, and C. O. Buckee, *PLoS Computational Biology* **9**, e1003268 (2013).
- [76] C. Stark, B.-J. Breitkreutz, T. Reguly, L. Boucher, A. Breitkreutz, and M. Tyers, *Nucleic Acids Research* **34**, D535 (2006).
- [77] M. De Domenico, A. Lancichinetti, A. Arenas, and M. Rosvall, *Physical Review X* **5**, 011027 (2015).
- [78] The assumption that communities are identical across layers is a reasonable one for this data set. When performing community detection without this restriction, we find that more than 90% of nodes stay in the same community across temporal layers. See [24] for details.
- [79] P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, and J.-P. Onnela, *Science* **328**, 876 (2010).
- [80] L. G. S. Jeub, M. Bazzi, I. S. Jutla, and P. J. Mucha, A generalized Louvain method for community detection implemented in MATLAB, <http://netwiki.amath.unc.edu/GenLouvain> (2011–2018), version 2.1.
- [81] We use GENLOUVAIN [80] to perform multilayer modularity maximization. This algorithm requires the specification of (at least) two parameters [41, 79], which we set to $\gamma = 1.2$ and $\omega = 5.0$. The large value of ω ensures that community assignments are identical across layers.
- [82] M. E. J. Newman and A. Clauset, *Nature Communications* **7** (2016).
- [83] J. Davis and M. Goadrich, in *Proceedings of the 23rd International Conference on Machine Learning* (2006) pp. 233–240.
- [84] M. Bazzi, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, and S. D. Howison, *Multiscale Modeling & Simulation: A SIAM Interdisciplinary Journal* **14**, 1 (2016).
- [85] T. J. Sweeting, *The Annals of Statistics*, 1375 (1980).
- [86] It is not exactly equal, because with K -fold cross validation, we are performing a prediction on only $1/K$ of the data at a time.
- [87] M. De Domenico, M. A. Porter, and A. Arenas, *Journal of Complex Networks* **3**, 159 (2015).