

The Importance of Being Correlated: Implications of Dependence in Joint Spectral Inference across Multiple Networks

Konstantinos Pantazis

KPANTAZI@UMD.EDU

*Department of Mathematics
University of Maryland
College Park, MD 20742*

Avanti Athreya

DATHREY1@JHU.EDU

*Department of Applied Mathematics and Statistics
Johns Hopkins University
Baltimore, MD 21218*

Jesús Arroyo

JARROYO@TAMU.EDU

*Department of Statistics
Texas A&M University
College Station, TX 77843*

William N. Frost

WILLIAM.FROST@ROSALINDFRANKLIN.EDU

*Cell Biology and Anatomy, and Center for Brain Function and Repair
Chicago Medical School
Rosalind Franklin University of Medicine and Science
Chicago, IL 60064-3905*

Evan S. Hill

EVAN.HILL@ROSALINDFRANKLIN.EDU

*Cell Biology and Anatomy, and Center for Brain Function and Repair
Chicago Medical School
Rosalind Franklin University of Medicine and Science
Chicago, IL 60064-3905*

Vince Lyzinski

VLYZINSK@UMD.EDU

*Department of Mathematics
University of Maryland
College Park, MD 20742*

Editor: Karsten Borgwardt

Abstract

Spectral inference on multiple networks is a rapidly-developing subfield of graph statistics. Recent work has demonstrated that joint, or simultaneous, spectral embedding of multiple independent networks can deliver more accurate estimation than individual spectral decompositions of those same networks. Such inference procedures typically rely heavily on independence assumptions across the multiple network realizations, and even in this case, little attention has been paid to the induced network correlation that can be a consequence of such joint embeddings. In this paper, we present a *generalized omnibus* embedding methodology and we provide a detailed analysis of this embedding across both independent and correlated networks, the latter of which significantly extends the reach of such

procedures, and we describe how this omnibus embedding can itself induce correlation. This leads us to distinguish between *inherent* correlation—that is, the correlation that arises naturally in multisample network data—and *induced* correlation, which is an artifact of the joint embedding methodology. We show that the generalized omnibus embedding procedure is flexible and robust, and we prove both consistency and a central limit theorem for the embedded points. We examine how induced and inherent correlation can impact inference for network time series data, and we provide network analogues of classical questions such as the effective sample size for more generally correlated data. Further, we show how an appropriately calibrated generalized omnibus embedding can detect changes in real biological networks that previous embedding procedures could not discern, confirming that the effect of inherent and induced correlation can be subtle and transformative. By allowing for and deconstructing both forms of correlation, our methodology widens the scope of spectral techniques for network inference, with import in theory and practice.

Keywords: Joint graph embeddings, induced correlation, latent space network models

1. Introduction

Networks and graphs, which consist of objects of interest and a vast array of possible relationships between them, arise very naturally in fields as diverse as political science (Ward et al., 2011; Waugh et al., 2009; Conover et al., 2011; Perry et al., 2018); bioinformatics via gene and protein interaction networks (Bhowmick and Seah, 2015; Elmsallati et al., 2015; Lezon et al., 2006); physics via (for example) dimer systems and percolation (Heilmann and Lieb, 1972; Lyons, 1989; Callaway et al., 2000); and sociology/social network analysis (Castellano et al., 2009; Carrington et al., 2005; Elmer et al., 2020), to name but a few. As such, they are a useful data structure for modeling complex interactions between different experimental entities. Network data, however, is qualitatively distinct from more traditional Euclidean data, and statistical inference on networks is a comparatively new discipline, one that has seen explosive growth over the last two decades. While there is a significant literature devoted to the rigorous statistical study of single networks, multiple network inference—the analogue of the classical problem of multiple-sample Euclidean inference—is still relatively nascent.

Much recent progress in network inference has relied on extracting Euclidean representations of networks, and popular methods include spectral embeddings of network adjacency (Athreya et al., 2018a) or Laplacian (Rohe et al., 2011) matrices, representation learning (Grover and Leskovec, 2016; Ribeiro et al., 2017), or Bayesian hierarchical methods (Durante et al., 2017). Moreover, many network models (Hoff et al., 2002a) allow for important properties of network entities to be hidden, or *latent*, and posit that relationships between entities depend on these latent variables. Such models, known as *latent position networks*, have wide intuitive appeal. For instance, relationships among participants in a social network are a function of the participants’ personal interests, which are typically not directly observed. In these cases, spectral embeddings can provide useful estimates of latent variables, effectively transforming, via eigendecompositions, a non-Euclidean inference problem into a Euclidean one.

For single latent position networks, spectrally-derived estimates of important graph parameters are well-understood, and under mild assumptions, these estimates satisfy classical notions of consistency (Sussman et al., 2012; Rohe et al., 2011), asymptotic normality

(Athreya et al., 2016; Tang and Priebe, 2018), and efficiency (Tang and Priebe, 2018; Tang et al., 2017c; Xie and Xu, 2019; Athreya et al., 2018b). More recently, spectral methods have also proven useful in multi-sample network inference, including (non)parametric estimation (Durante et al., 2017; Tang et al., 2018), two-sample hypothesis testing (Tang et al., 2017a,b; Asta and Shalizi, 2014; Li and Li, 2018), and graph matching (Lyzinski et al., 2015; Zhang, 2018a,b). Typically, these methods rely upon separately embedding multiple networks into a lower-dimensional Euclidean space and then aligning the embeddings via Procrustes analysis (Gower, 1975) or point set registration methods (Myronenko and Song, 2010). An important issue in multi-sample inference, however, is the use of *multiple* networks both for improved estimation of underlying model parameters and for more streamlined testing across several populations of networks. To this end, a number of recent papers are dedicated to the development of novel techniques for simultaneously embedding several networks into a common Euclidean space, employing spectral graph techniques (Levin et al., 2017; Nielsen and Witten, 2018; Wang et al., 2019; Arroyo et al., 2020), tensor factorizations (Zhang and Xia, 2018; Zhang, 2019; Jing et al., 2020), multilayer network decompositions (Kivelä et al., 2014; Paul and Chen, 2016, 2020), and nonparametric Bayesian algorithms (Durante et al., 2017; Durante and Dunson, 2018).

While multisample joint embedding methods allow for accurate graph inference and are often superior to individual separate embeddings (Levin et al., 2017; Arroyo et al., 2020), there are a number of potential pitfalls in joint embeddings. In particular, network statisticians must confront issues of noisy vertex alignments across graphs (Lyzinski, 2018); large, high-rank matrices that arise in a joint embedding (Draves and Sussman, 2020); the relationship between individual network sparsity and the signal in a joint embedding; and the *induced* correlation across estimates that arise from the joint embedding, the last of which is inevitable in any simultaneous embedding procedure. What is more, virtually all existing procedures for multisample network inference rely, like their classical analogues, on an assumption of independence across network realizations. In this sense, existing methodology is ill-equipped to handle, at least in a principled manner, the *inherent* network correlation—for example, the natural and unavoidable correlation across edges in a network time series—to say nothing of the additional correlation induced by any dimension-reduction procedure.

This paper is the first to broaden spectral analysis to account for both types of correlation, and to examine how the correlation induced by joint spectral procedures can mask or amplify important signal. We focus on a generalization of the omnibus multiple graph embedding procedure (OMNI) of Levin et al. (2017), in which multiple networks are simultaneously embedded into a single lower-dimensional subspace, with a distinct representation for each vertex across the networks. The work of Levin et al. (2017) considers this problem in the case where the network realizations themselves are independent, though even when independent network samples are jointly embedded, correlation across the embedded point clouds is automatically induced by the OMNI procedure (this is the price we pay to circumvent the pairwise Procrustes or registration analysis necessary in separate embedding settings), to say nothing of the impact of OMNI in preserving or masking the a priori present inherent correlation across networks. It is natural, then, to seek to adapt the OMNI embedding technique in order to preserve in the embedded (independent or correlated) graphs the same correlation that would be present if edge-wise inherently correlated networks are embedded separately and then aligned. This would allow for the jointly-embedded networks

to be a more appropriate proxy in embedding space for sequences of graphs with complex dependency structures.

To understand these phenomena more rigorously, we anchor our analysis in a specific class of latent position random graphs, the *random dot product graph* (RDPG; see Young and Scheinerman (2007)). Random dot product graphs have proven to be a theoretically tractable family of latent position networks (Athreya et al., 2018a) suitable for modeling a host of complex real-data networks (Tang and Priebe, 2018; Priebe et al., 2019; Patsolic et al., 2020). In Section 2.2, we formulate several models for inherent correlation across a series of random dot product graphs, and, thereafter, examine the impact of a joint spectral embedding of such a collection of multiple networks. Given m adjacency matrices of multiple independent, n -vertex, aligned RDPGs, the OMNI embedding of Levin et al. (2017) and its more general counterpart—the genOMNI embedding we define here—provide m distinct representations for the latent attributes of each of the n vertices in the collection of graphs. This permits both consistent estimation of underlying RDPG latent positions (in which the omnibus embedding is empirically shown to be competitive with embedding the sample mean of the respective adjacencies) and inference *across* the latent positions, including testing, classification, and change-point detection. The generalized omnibus embedding jointly embeds the collection of graphs, though by construction the m distinct $n \times d$ blocks in the genOMNI embedding (there are m such blocks) are necessarily correlated. This is not a unique feature of the generalized omnibus methodology; all joint embedding procedures typically induce correlation across network pairs. What is unique, at least to our knowledge, about the genOMNI embedding is that the dual impact of induced and inherent correlation in the embedding is theoretically tractable.

In light of this, the major contributions of this paper are as follows. The first is an entirely novel treatment of method-induced correlation in the output of joint embedding procedures; indeed, the question of method-induced correlation is not considered in the original OMNI work in Levin et al. (2017) or in other existing joint graph embedding papers (e.g., Arroyo et al. (2020); Wang et al. (2019); Nielsen and Witten (2018)). The second is the development, through genOMNI, of a flexible joint embedding that can not only reproduce complex correlation in the embedded space, but can also accommodate inherently correlated network data while retaining important theoretical performance guarantees of consistency and asymptotic normality. By comparing the omnibus embedding of independent graphs to the separate embeddings of correlated (and subsequently Procrustes-aligned) graphs, we can explicitly capture the level of correlation the joint OMNI embedding induces, in the limit, across independent networks. This, in turn, motivates the creation of the generalized omnibus embedding (Definition 12), which produces more complex correlation structure in the embedded space, enabling higher-fidelity application of the omnibus methodology in real data. This replication of more complicated correlation structure renders the generalized omnibus embedding suitable for inference on network time series, because it can reproduce, via realizations of independent networks, the correlation that is an important component of a time series.

The core result underlying the above is a central limit theorem (Theorems 11 and 14) for the row-wise residuals of the estimated latent positions in both inherently correlated and/or independent RDPGs in a generalized omnibus embedding framework. In addition, we are able to precisely characterize the dual effects of inherent and induced correlation on

the limiting covariance structure across the embedded networks. In Sections 5 and 6, we show how the weights of genOMNI can be adapted for certain specific inherent correlation structures. As an illustration of the power of the more nuanced embedding correlation enabled by the genOMNI setting, we present in Section 3.2 an analysis of a motor program time series of networks in the brain of the marine mollusk *Aplysia californica*. The classical omnibus embedding on this time series homogenizes the inherent correlation across the time-series, effectively obscuring important network changes corresponding directly to transitions in animal behavior (from stimulus to gallop and crawl). Our genOMNI embedding, however, is flexible enough to permit different weightings of networks over time, and this more general joint inference procedure captures exactly the signal the earlier omnibus embedding misses. Figure 3 in Section 3.2 and Figure 7 in Section 5 demonstrate this contrast in inferential accuracy between the two.

Lastly, in Section 7, we further show, with theory, simulated and real data examples, how inherent and induced correlation across networks impact the effective sample size for subsequent inference tasks in the joint embedded space. This provides a network analogue of the classical statistical challenge of quantifying, via a comparison of sample sizes, the extent to which dependence in data can impact inference. In sum, our generalized omnibus embedding and accompanying correlation analysis form a tractable, scalable inference methodology that can be applied to independent and correlated data, carries straightforward theoretical guarantees, has demonstrable empirical utility, and correctly identifies important and subtle network changes that its predecessors miss.

Notation: For a positive integer n , we let $[n] = \{1, 2, \dots, n\}$, let 0_n be the zero $n \times n$ matrix, and let $\vec{1}_n \in \mathbb{R}^n$ (resp. $J_n \in \mathbb{R}^{n \times n}$) be the vector (resp. matrix) with all entries identically equal to one. The set of $n \times n$ real orthogonal matrices is denoted by \mathcal{O}_n . We represent a simple (no self-loops or multiple edges), un-weighted and un-directed graph as the ordered pair $G = (V, E)$, where $V = [n]$ represents the set of nodes and $E \subset \binom{[n]}{2}$ the set of edges of the graph; we denote the set of all n -vertex labeled graphs via \mathcal{G}_n . For the graph $G = (V, E)$, we will denote its adjacency matrix via $A \in \{0, 1\}^{n \times n}$; i.e., A_{ij} is equal to 1 if there exists an edge between nodes i and j in G , and 0 otherwise. Where there is no danger of confusion, we will often refer to a graph G and its adjacency matrix A interchangeably. The Kronecker product is denoted by \otimes and the direct sum by \oplus . Finally, the symbols $\|\cdot\|_F$, $\|\cdot\|$, and $\|\cdot\|_{2 \rightarrow \infty}$ correspond to the Frobenius, spectral and two-to-infinity norms respectively.

2. Background

In this section, we will introduce the modeling and spectral embedding frameworks that we build our theory and methods upon.

2.1 Random Dot Product Graphs

The theoretical developments to follow are situated in the context of the *random dot product graph* (as mentioned above, abbreviated RDPG) model of Young and Scheinerman (2007). Random dot product graphs are a special case of the more general *latent position random graphs* (abbreviated LPGs) of Hoff et al. (2002b). Every vertex in a latent position random

graph has associated to it a (typically unobserved) *latent position*, itself an object belonging to some (often Euclidean) space \mathcal{X} . Probabilities of an edge between two vertices i and j , p_{ij} , are then a function $\kappa(\cdot, \cdot) : \mathcal{X} \times \mathcal{X} \rightarrow [0, 1]$ (known as the *link function*) of their associated latent positions (x_i, x_j) . Thus $p_{ij} = \kappa(x_i, x_j)$, and edges between vertices arise independently of one another. Given these probabilities, the entries A_{ij} of the adjacency matrix A are conditionally independent Bernoulli random variables with success probabilities p_{ij} . We consolidate these probabilities into a matrix $P = (p_{ij})$, and we write $A \sim P$ to denote this relationship.

In a d -dimensional random dot product graph, the latent space is an appropriately-constrained subspace of \mathbb{R}^d , and the link function is simply the dot product of the two latent d -dimensional vectors. Random dot product graphs are often divided into two types: those in which the latent positions are fixed, and those in which the latent positions are themselves random. Specifically, we consider the case in which the latent position $X_i \in \mathbb{R}^d$ for vertex i is drawn from some distribution F on \mathbb{R}^d , and we further assume that the latent positions for each vertex are drawn independently and identically from this distribution F . Random dot product graphs have proven to be a tractable and useful model for low-rank latent position networks, and variants of the RDPG model have recently emerged that extend the framework to allow for modeling more complex network topologies (Rubin-Delanchy et al., 2017; Tang and Priebe, 2018).

Definition 1 (d -dimensional RDPG) *Let F be a distribution on a set $\mathcal{X} \in \mathbb{R}^d$ satisfying $\langle x, x' \rangle \in [0, 1]$ for all $x, x' \in \mathcal{X}$. Let $X_1, X_2, \dots, X_n \sim F$ be i.i.d. random variables distributed via F , and let $P = \mathbf{X}\mathbf{X}^T$, where $\mathbf{X} = [X_1^T | X_2^T | \dots | X_n^T]^T \in \mathbb{R}^{n \times d}$. Let A be a symmetric, hollow adjacency matrix with above diagonal entries distributed via*

$$P(A|\mathbf{X}) = \prod_{i < j} (X_i^T X_j)^{A_{ij}} (1 - X_i^T X_j)^{1 - A_{ij}}; \tag{1}$$

i.e., conditioned on \mathbf{X} the above diagonal entries are independent Bernoulli random variables with success probabilities provided by the corresponding above diagonal entries in P . The pair (A, \mathbf{X}) is then said to be an instantiation of a d -dimensional Random Dot Product Graph with distribution F , denoted $(A, \mathbf{X}) \sim \text{RDPG}(F, n)$.

Note that there is a rotational non-identifiability inherent to the RDPG model. Indeed, if $\mathbf{Y} = \mathbf{X}W$ for $W \in \mathcal{O}_d$, then the distribution over graphs induced by Eq. 1 by \mathbf{X} and \mathbf{Y} are identical; i.e., $\mathbb{P}(A|\mathbf{X}) = \mathbb{P}(A|\mathbf{Y})$ for all A . As inference in the RDPG setting often proceeds by first estimating the latent positions \mathbf{X} , which can only be done up to a rotation factor, this model is not generally suitable for inference tasks that are not rotationally invariant.

Remark 2 *We note that a generalization of the RDPG model has recently been developed, namely the Generalized Random Dot Product Graph of Rubin-Delanchy et al. (2017), which allows for modeling latent position graphs where P is not necessarily positive definite (this, for example, allows for disassortative connectivity behavior in stochastic block model networks modeled via the RDPG framework which is not possible under Definition 1). In more general latent space models (see, for example, Hoff et al. (2002a); Lei (2020); Hoff (2008)), the probability of connections in the network are computed via more general similarity functions (i.e., kernels) compared to the dot product in the RDPG. While we suspect that our*

results translate immediately to the generalized RDPG setting (and perhaps less immediately to the general latent space setting; see Tang et al. (2013b)), we do not pursue this further here.

2.2 Modeling Multiple Correlated RDPGs

Next, we present a natural extension of the RDPG model to the multiple network setting, namely the *Joint Random Dot Product Graph* from Levin et al. (2017). This model allows us to simultaneously characterize multiple graphs $(A^{(k)})_{i=1}^m$ with a common set of latent positions \mathbf{X} so that the collection of $A^{(k)}$ are conditionally independent given \mathbf{X} .

Definition 3 (Joint Random Dot Product Graph) *Let F be a distribution on a set $\mathcal{X} \in \mathbb{R}^d$ satisfying $\langle x, x' \rangle \in [0, 1]$ for all $x, x' \in \mathcal{X}$. Let $X_1, X_2, \dots, X_n \stackrel{i.i.d.}{\sim} F$, and let $P = \mathbf{X}\mathbf{X}^T$, where $\mathbf{X} = [X_1^T | X_2^T | \dots | X_n^T]^T \in \mathbb{R}^{n \times d}$. We say that the random graphs $(A^{(1)}, A^{(2)}, \dots, A^{(m)})$ are an instantiation of a Joint Random Dot Product Graph model (abbreviated JRDPG), written*

$$(A^{(1)}, A^{(2)}, \dots, A^{(m)}, \mathbf{X}) \sim \text{JRDPG}(F, n, m)$$

if marginally each $(A^{(k)}, \mathbf{X}) \sim \text{RDPG}(F, n)$ and conditioned on \mathbf{X} , the $A^{(k)}$'s are independent with distribution given by Eq. (1).

While this model allows for modeling multiple networks simultaneously, the conditional independence is ill suited for a number of inference tasks (e.g., time-series analysis in graphs as in Wang et al. (2013); Tang et al. (2013a); Bhattacharyya et al. (2020)) that necessitate more nuanced dependency structure across graphs. While correlated RDPG models exist for pairs of graphs (see, for example, Patsolic et al. (2020)), we seek a framework that allows for (pairwise) correlation across the entire collection of $A^{(i)}$. To do so, we assume, as in JRDPG, that each network is marginally distributed as an RDPG, and then we assign an edge-wise correlation among the network pairs. We will then provide a few constructive methods via which such networks can be sampled.

Definition 4 (Pairwise multiple edge-correlated RDPG) *With notation as in Definition 3, we say that the random graphs $(A^{(1)}, A^{(2)}, \dots, A^{(m)})$ are an instantiation of a R-correlated Joint Random Dot Product Graph model, written $(A^{(1)}, A^{(2)}, \dots, A^{(m)}, \mathbf{X}) \sim \text{JRDPG}(F, n, m, R)$, if*

- i. Marginally, $(A^{(k)}, \mathbf{X}) \sim \text{RDPG}(F, n)$ for every $k \in [m]$;*
- ii. The matrix $R \in [-1, 1]^{m \times m}$ is symmetric and has diagonal entries identically equal to 1. We will write the (k_1, k_2) -element of R via ρ_{k_1, k_2} .*
- iii. Conditioned on \mathbf{X} the collection*

$$\{A_{i,j}^{(k)}\}_{k \in [m], i < j}$$

is mutually independent except that for each $\{i, j\} \in \binom{V}{2}$, we have for each $k_1, k_2 \in [m]$,

$$\text{correlation}(A_{i,j}^{(k_1)}, A_{i,j}^{(k_2)}) = \rho_{k_1, k_2}.$$

Note that if all off-diagonal elements of R are identically equal to ρ , then we will often write $(A^{(1)}, A^{(2)}, \dots, A^{(m)}, \mathbf{X}) \sim \text{JRDPG}(F, n, m, \rho)$.

Below we will assume that $R \geq 0_m$ entry-wise, although negatively correlated graphs can be considered in Definition 4. The pairwise multiple edge-correlated RDPG is a better candidate for modeling time series of networks and multilayer networks than the conditionally independent JRDPG model, as it allows for generating conditionally (within graph) edge-independent networks and induces correlation across networks pairwise. A natural extension would be to allow for correlation across edges within each network, and we are actively working on this extension; see Babkin et al. (2020) for an example of how this structure could be introduced in a model related to the stochastic blockmodel.

We next illustrate two constructions that result in $\text{RDPG}(F, n, m, R)$ networks.

2.2.1 FORWARD PROPAGATION (SEQUENTIAL) MODEL

The *Forward Propagation (Sequential) model* (abbreviated $\text{JRDPG}_{\text{for}}$) fits time-varying networks with forward propagation of correlation suitable for network time-series inference tasks such as (spectral) clustering (Pensky and Zhang, 2019) and anomaly detection of vertices (or networks) (Chen et al., 2020) in a given time period. It derives the $\text{JRDPG}(F, n, m, R)$ for R equal to the symmetric matrix

$$R_{\text{for}} = \begin{bmatrix} 1 & \varrho_{1,2} & \varrho_{1,2}\varrho_{2,3} & \cdots & \prod_{k=1}^{m-1} \varrho_{k,k+1} \\ \varrho_{1,2} & 1 & \varrho_{2,3} & \cdots & \prod_{k=2}^{m-1} \varrho_{k,k+1} \\ \varrho_{1,2}\varrho_{2,3} & \varrho_{2,3} & 1 & \cdots & \prod_{k=3}^{m-1} \varrho_{k,k+1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \prod_{k=1}^{m-1} \varrho_{k,k+1} & \prod_{k=2}^{m-1} \varrho_{k,k+1} & \prod_{k=3}^{m-1} \varrho_{k,k+1} & \cdots & 1 \end{bmatrix} \in \mathbb{R}^{m \times m}.$$

The formal definition is presented as follows.

Definition 5 (Forward Propagation (Sequential) model) *With the notation as in Definition 4, we say that the random graphs $A^{(1)}, A^{(2)}, \dots, A^{(m)}$ are an instantiation of a sequential (correlated) JRDPG and we write $(A^{(1)}, A^{(2)}, \dots, A^{(m)}, \mathbf{X}) \sim \text{JRDPG}_{\text{for}}(F, n, m, \boldsymbol{\varrho})$ if*

- i. $\boldsymbol{\varrho} \in [0, 1]^{m-1}$ is a vector whose k -th element is denoted $\varrho_{k,k+1}$.
- ii. $(A^{(1)}, A^{(2)}, \dots, A^{(m)}, \mathbf{X}) \sim \text{JRDPG}(F, n, m, R)$ where for $1 \leq k_1 < k_2 \leq m$,

$$R_{k_1, k_2} := \rho_{k_1, k_2} = \prod_{k=k_1}^{k_2-1} \varrho_{k, k+1}.$$

Practically, we can sample from $(A^{(1)}, A^{(2)}, \dots, A^{(m)}, \mathbf{X}) \sim \text{JRDPG}_{\text{for}}(F, n, m, \boldsymbol{\varrho})$ by first sampling from $(A^{(1)}, \mathbf{X}) \sim \text{RDPG}(F, n)$, and then conditional on \mathbf{X} and $A^{(\ell)}$ for $\ell \geq 1$, independently sampling the edges of $A^{(\ell+1)}$ according to the following scheme

$$\forall \{i, j\} \in \binom{V}{2}, \quad A_{ij}^{(\ell+1)} \sim \begin{cases} \text{Bern}(P_{ij} + \varrho_{\ell, \ell+1}(1 - P_{ij})) & \text{if } A_{ij}^{(\ell)} = 1, \\ \text{Bern}(P_{ij}(1 - \varrho_{\ell, \ell+1})) & \text{if } A_{ij}^{(\ell)} = 0. \end{cases}$$

A straightforward induction on ℓ guarantees that marginally, $(A^{(\ell)}, \mathbf{X}) \sim \text{RDPG}(F, n)$. The form of the correlation follows from the recursion that for $k_1 < k_2$ (suppressing the conditioning on \mathbf{X} below),

$$\text{correlation}(A_{i,j}^{(k_1)}, A_{i,j}^{(k_2)}) = \varrho_{k_1, k_2} \cdot \text{correlation}(A_{i,j}^{(k_1)}, A_{i,j}^{(k_2-1)}).$$

It is also possible to model non-stationary time series of graphs (i.e., allowing for distinct (latent) distributions of $A^{(k_1)}$ and $A^{(k_2)}$ still with pairwise correlation; see Lyzinski and Sussman (2020) for details).

2.2.2 SINGLE GENERATOR MODEL

Similarly, the *Single Generator model* derives from $\text{JRDPG}(F, n, m, R)$ when the correlation matrix R is equal to

$$R_{gen} = \begin{bmatrix} 1 & \varrho_1 \varrho_2 & \cdots & \varrho_1 \varrho_m \\ \varrho_2 \varrho_1 & 1 & \cdots & \varrho_2 \varrho_m \\ \vdots & \vdots & \ddots & \vdots \\ \varrho_m \varrho_1 & \varrho_m \varrho_2 & \cdots & 1 \end{bmatrix} \in \mathbb{R}^{m \times m}.$$

Define the (generator) vector $\nu = [\varrho_1, \dots, \varrho_m]^T \in \mathbb{R}^m$, so that the correlation matrix R_{gen} can be also written as $R_{gen} = \nu \nu^T + \text{diag}(I_m - \nu \nu^T)$. It is reasonable to add the term $\text{diag}(I_m - \nu \nu^T)$; this ensures that the correlation of each graph with itself remains the same for all graphs and is equal to 1.

Definition 6 (Single Generator Model) *With notation as in Definition 4, we say that the random graphs $A^{(0)}, A^{(1)}, \dots, A^{(m)}$ are an instantiation of a multiple RDPG with generator matrix $A^{(0)}$ and we write $(A^{(0)}, A^{(1)}, A^{(2)}, \dots, A^{(m)}, \mathbf{X}) \sim \text{JRDPG}_{gen}(F, n, m, \nu)$ if*

- i. Marginally, $(A^{(0)}, \mathbf{X}) \sim \text{RDPG}(F, n)$;
- ii. $\nu \in [0, 1]^m$ is a nonnegative vector with entries in $[0, 1]$; we denote the k -th entry of ν via $\nu_k = \varrho_k$.
- iii. $(A^{(1)}, A^{(2)}, \dots, A^{(m)}, \mathbf{X}) \sim \text{JRDPG}(F, n, m, R)$ where $R = \nu \nu^T + \text{diag}(I_m - \nu \nu^T)$ so that for $1 \leq k_1 < k_2 \leq m$,

$$R_{k_1, k_2} := \rho_{k_1, k_2} = \varrho_{k_1} \varrho_{k_2}.$$

The single generator model, JRDPG_{gen} , mimics the correlated Erdős-Rényi graph pairs that are a common model in the graph matching literature (see, for example, Pedarsani and Grossglauser (2011); Cullina and Kiyavash (2016)) where the pair of graphs are noisy realizations of a background network (here $A^{(0)}$). The single generator model with m networks provides a suitable framework for studying problems of aligning multiple networks. As in the JRDPG_{for} case, we can sample from $(A^{(0)}, \dots, A^{(m)}, \mathbf{X}) \sim \text{JRDPG}_{gen}(F, n, m, \nu)$ by first sampling from $(A^{(0)}, \mathbf{X}) \sim \text{RDPG}(F, n)$, and then conditional on $A^{(0)}$ and \mathbf{X} , independently sampling the edges of each $A^{(\ell)}$, $\ell \in [m]$, according to the following scheme

$$\forall \{i, j\} \in \binom{V}{2}, \quad A_{ij}^{(\ell)} \sim \begin{cases} \text{Bern}(P_{ij} + \varrho_\ell(1 - P_{ij})) & \text{if } A_{ij}^{(0)} = 1, \\ \text{Bern}(P_{ij}(1 - \varrho_\ell)) & \text{if } A_{ij}^{(0)} = 0. \end{cases}$$

The above scheme implies that the correlation between two networks $A^{(\ell_1)}, A^{(\ell_2)}$ is given by the product $\varrho_{\ell_1} \varrho_{\ell_2}$, i.e., $\text{correlation}(A_{i,j}^{(\ell_1)}, A_{i,j}^{(\ell_2)}) = \varrho_{\ell_1} \varrho_{\ell_2}$.

2.3 Spectral Graph Embeddings

One of the key inference tasks in latent position random graphs (LPGs) is to estimate the unobserved latent positions for each of the vertices based on a single observation of the adjacency matrix of a sufficiently large graph. Since the matrix of connection probabilities for an RDPG is expressible as an outer product of the matrix of true latent positions, and since the adjacency matrix A can be regarded as a “small” perturbation of P , the inference of properties of P from an observation of A is a problem well-suited to spectral graph methods, such as singular value decompositions of adjacency or Laplacian matrices. Indeed, these spectral decompositions have been the basis for a suite of approaches to graph estimation, community detection, and hypothesis testing for random dot product graphs. For a comprehensive summary of these techniques, see Athreya et al. (2018a). Note that the popular stochastic blockmodel (SBM) with positive semidefinite block connection probabilities can be regarded as a random dot product graph. In an SBM, there are a finite number of possible latent positions for each vertex—one for each block—and the latent position exactly determines the block assignment for that vertex.

For a random dot product graph in which the latent position $X_i \in \mathbb{R}^d$ for each vertex i , $1 \leq i \leq n$, are drawn i.i.d from some distribution F on \mathbb{R}^d , a common graph inference task is to infer properties of F from an observation of the graph alone. For example, in a stochastic block model, in which the distribution F is discretely supported, we may wish to estimate the point masses in the support of F . In the graph inference setting, however, there are two sources of randomness that culminate in the generation of the actual graph: first, the randomness in the latent positions, and second, *given these latent positions*, the conditional randomness in the existence of edges between vertices.

A rank- d RDPG has a connection probability matrix P that is necessarily low rank (rank d , regardless of the number of vertices in the graph); hence random dot product graphs can be productively analyzed using low-dimensional embeddings. Under mild assumptions, the adjacency matrix A of a random dot product graph approximates the matrix $P = \mathbb{E}(A)$. To be more precise, $P = \mathbb{E}(A|\mathbf{X})$ in the sense that the spectral norm of $A - P$ can be controlled; see Oliveira (2009) and Lu and Peng (2013). It is reasonable to ask how close the spectrum and associated invariant subspaces of A are to those of P . Weyl’s Theorem (Horn and Johnson, 1985) describes how the eigenvalues of A differ from those of P . Sharp bounds on differences between the associated invariant subspaces are fewer, with the Davis-Kahan Theorem (Davis and Kahan, 1970; Bhatia, 1997; Yu et al., 2015) perhaps the best known. Because of the invariance of the inner product to orthogonal transformations, however, the RDPG exhibits a clear nonidentifiability: latent positions can be estimated only up to an orthogonal transformation.

Since P is a symmetric, positive definite, rank d -matrix of the form $P = \mathbf{X}\mathbf{X}^T$, the latent position matrix \mathbf{X} can be written as $\mathbf{X} = U_P S_P^{1/2} W$ for some orthogonal matrix W , where S_P is the diagonal matrix of the d nonzero eigenvalues of P , sorted by magnitude, and U_P the associated eigenvectors. The Davis-Kahan Theorem translates spectral norm

bounds on $A - P$ to projection operator bounds between $U_A U_A^T$ and $U_P U_P^T$, and, in turn, into Frobenius norm bounds between U_A and a rotation of U_P (Rohe et al., 2011).

These bounds can be sufficiently sharpened to ensure that the rows of a partial spectral decomposition of A , known as the *adjacency spectral embedding* (ASE) are accurate estimates of the latent positions X_i for each vertex. With this in mind, we define the adjacency spectral embedding (ASE) as follows.

Definition 7 (Adjacency Spectral Embedding (ASE)) *Let $d \geq 1$ be a positive integer. The d -dimensional adjacency spectral embedding of a graph $A \in \mathcal{G}_n$ into \mathcal{R}^d , denoted by $\text{ASE}(A, d)$, is defined to be $\hat{\mathbf{X}} = U_A S_A^{1/2}$, where*

$$|A| = (A^\top A)^{1/2} = \left[U_A \mid \tilde{U}_A \right] \left[S_A \oplus \tilde{S}_A \right] \left[U_A \mid \tilde{U}_A \right]^\top,$$

is the spectral decomposition of $|A|$, $S_A \in \mathbb{R}^{d \times d}$ is the diagonal matrix with the d largest eigenvalues of $|A|$, and $U_A \in \mathbb{R}^{n \times d}$ the corresponding matrix of the d -largest eigenvectors.

Now, if we define $\hat{\mathbf{X}} = U_A S_A^{1/2}$, where S_A is the diagonal matrix of the top d eigenvalues of $(A^\top A)^{1/2}$, sorted by magnitude, and the columns of U_A are the associated unit eigenvectors, results in Sussman et al. (2012) and Lyzinski et al. (2014) establish that, under assumptions on the spectrum of P , the rows $\hat{\mathbf{X}}$ are consistent estimates of the latent positions $\{X_i\}$ (up to orthogonal rotation). Further, in Athreya et al. (2016), it is shown that under the RDPG, the (suitably-scaled) ASE of the adjacency matrix converges in distribution to a Gaussian mixture.

The utility of the ASE in single graph inference points us to a natural test statistic for determining whether two random dot product graphs have the same latent positions. Namely, we can perform Procrustes alignment of two graphs' embeddings (Tang et al., 2017a). Specifically, let $A^{(1)}$ and $A^{(2)}$ be the adjacency matrices of two random dot product graphs on the same vertex set, with vertices aligned so that vertex i in $A^{(1)}$ can be sensibly identified with vertex i in $A^{(2)}$ for all $i \in [n]$. Letting $\hat{\mathbf{X}}$ and $\hat{\mathbf{Y}}$ be the respective adjacency spectral embeddings of these two graphs, if the two graphs have the same generating P matrices, then it is reasonable to surmise that the Procrustes distance

$$\min_{W \in \mathcal{O}^{d \times d}} \|\hat{\mathbf{X}} - \hat{\mathbf{Y}}W\|_F, \quad (2)$$

will be small. In Tang et al. (2017a), the authors show that a scaled version of the Procrustes distance in (2) provides a valid and consistent test for the equality of latent positions for pairs of graphs. Unfortunately, the fact that a Procrustes minimization must be performed both complicates the test statistic and compromises its power.

The Procrustes alignment is necessary, though, because these two embeddings may well inhabit different d -dimensional subspaces of \mathbb{R}^n . An alternative approach is to consider jointly embedding a collection of random graphs into the *same* subspace, which is the topic of Levin et al. (2017) and which we extend here. Consider $(A^{(1)}, A^{(2)}, \dots, A^{(m)}, \mathbf{X}) \sim \text{JRDPG}(F, n, m, R)$. Once the latent position matrix is generated, the m correlated random graphs with respective adjacency matrices A_1, \dots, A_m all have the same connection probability matrix P . That is, $\mathbb{E}(A^{(k)} | \mathbf{X}) = P$ for all k . This is a direct graph-analogue of

correlated Euclidean data from the same generating distribution, and given that we have *multiple* adjacency matrices from the same distribution, it is plausible that a latent position inference procedure using all the $A^{(k)}$ matrices is superior to an inference procedure that depends on a single $A^{(k)}$. In addition, since individual graph embeddings cannot be compared without Procrustes alignments, a joint embedding procedure that eliminates post-hoc pairwise alignments can be particularly useful.

In Levin et al. (2017), the authors consider the setting where the $A^{(k)}$ are independent (i.e., the conditionally independent JRDPG model), and they build a spectral embedding of an $mn \times mn$ matrix M from the $A^{(k)}$ matrices by placing each $A^{(k)}$ on the main block-diagonal and, on the (k, ℓ) -th off-diagonal block, the average $\frac{A^{(k)}+A^{(\ell)}}{2}$. That is, for, say the $m = 2$ case, the matrix M is

$$M = \begin{bmatrix} A^{(1)} & \frac{A^{(1)}+A^{(2)}}{2} \\ \frac{A^{(1)}+A^{(2)}}{2} & A^{(2)} \end{bmatrix}, \quad (3)$$

Observe that the expected value of M is the matrix

$$\mathbb{E}[M] = \begin{bmatrix} P & P \\ P & P \end{bmatrix}, \quad (4)$$

which is still a rank d matrix. As a consequence, a d -dimensional embedding of M can produce an m -fold collection of *correlated* estimates for the rows of the latent position matrix \mathbf{X} . This *joint* or *omnibus* spectral embedding, denoted *OMNI*, is defined as follows.

Definition 8 (Omnibus Spectral Embedding) *Let $A^{(1)}, A^{(2)}, \dots, A^{(m)}$ be a collection of m graphs each in \mathcal{G}_n . Define the omnibus matrix of $A^{(1)}, A^{(2)}, \dots, A^{(m)}$ to be the $mn \times mn$ matrix M defined via*

$$M = \begin{bmatrix} A^{(1)} & \frac{A^{(1)}+A^{(2)}}{2} & \frac{A^{(1)}+A^{(3)}}{2} & \dots & \frac{A^{(1)}+A^{(m)}}{2} \\ \frac{A^{(2)}+A^{(1)}}{2} & A^{(2)} & \frac{A^{(2)}+A^{(3)}}{2} & \dots & \frac{A^{(2)}+A^{(m)}}{2} \\ \frac{A^{(3)}+A^{(1)}}{2} & \frac{A^{(3)}+A^{(2)}}{2} & A^{(3)} & \dots & \frac{A^{(3)}+A^{(m)}}{2} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ \frac{A^{(m)}+A^{(1)}}{2} & \frac{A^{(m)}+A^{(2)}}{2} & \frac{A^{(m)}+A^{(3)}}{2} & \dots & A^{(m)} \end{bmatrix}.$$

The d -dimensional Omnibus Spectral Embedding of $A^{(1)}, A^{(2)}, \dots, A^{(m)}$ is then given by

$$\text{OMNI}(A^{(1)}, A^{(2)}, \dots, A^{(m)}, d) = \text{ASE}(M, d) = U_M S_M^{1/2},$$

where $\text{ASE}(M, d)$ is the d -dimensional adjacency spectral embedding of M ; that is, $S_M \in \mathbb{R}^{d \times d}$ is the diagonal matrix of the top d eigenvalues of $|M|$, and $U_M \in \mathbb{R}^{mn \times d}$ the corresponding eigenvectors.

Note that the omnibus embedding is an $mn \times d$ -dimensional matrix. Each m -fold block of rows supplies an $n \times d$ matrix that can serve as a latent position estimate for the corresponding graph. That is, the s -th $n \times d$ block of the omnibus embedding, denoted $\hat{\mathbf{X}}^{(s)} = [U_M S_M^{1/2}]^s$, is an estimate for $\mathbf{X}^{(s)}$, the matrix of latent positions for the s -th graph.

In Levin et al. (2017), it is shown that—in parallel to the same result for the ASE—when $(A^{(1)}, A^{(2)}, \dots, A^{(m)}, \mathbf{X}) \sim \text{JRDPG}(F, n, m)$, the rows of the omnibus embedding provide simultaneous consistent estimation for the latent positions X_j , where $1 \leq j \leq n$. There are m such rows for each vertex j . In addition, Levin et al. (2017) demonstrates that for fixed m , as $n \rightarrow \infty$, the distribution of any fixed l sub-collection of the rows of the omnibus matrix, suitably scaled, converges in distribution to a mixture of Gaussians. It is important to note that in Levin et al. (2017), the jointly embedded adjacency matrices are all *independent*; as we shall see, in this work, which allows for explicit correlation in the multiple-graph model, we extend consistency and normality to the dependent case.

Even when the adjacency matrices are themselves independent, the simultaneous nature of the omnibus embedding forces correlation across the estimated latent positions, and in return it obviates the need for Procrustes alignments between $n \times d$ matrices $\hat{\mathbf{X}}^{(s)}$ and $\hat{\mathbf{X}}^{(t)}$, where $1 \leq s \leq t \leq m$. One consequence of this is that with the omnibus embedding, an empirically useful test statistic for assessing latent position equality is simply the Frobenius norm $\|\hat{\mathbf{X}}^{(s)} - \hat{\mathbf{X}}^{(t)}\|_F$ (as opposed to $\min_{W \in \mathcal{O}_d} \|\hat{\mathbf{X}}^{(s)} - \hat{\mathbf{X}}^{(t)}W\|_F$ in the separately embedded graph setting of Tang et al. (2017a)). Quantifying the correlation induced by this joint embedding, and determining how it relates to the choice of block-matrix on the off-diagonal of M , is an important question. The classical omnibus matrix uses a simple pairwise average, chosen to balance the requirements of estimation accuracy (when all graphs have the same latent positions) with the need to retain discriminatory power in hypothesis testing (when some of the graphs have different latent positions). By analyzing precisely how the off-diagonal blocks can impact this correlation, we can describe how the omnibus embedding can be used to perform inference on collections of graphs that are not necessarily independent, or can replicate some desired correlation structure in spectral estimates.

3. Inherent and Induced Correlation in Classical OMNI

Before delving into our theoretical results on the competing roles of induced versus inherent correlation in the omnibus framework, we first consider the effect of classical OMNI in the original setting of embedding conditionally independent networks considered in Levin et al. (2017); Draves and Sussman (2020). This highlights both the correlation induced by the OMNI method and the dual contributions of inherent and induced correlation in subsequent correlated graph results.

To this end, we start with the task of embedding a pair of n -vertex correlated d -dimensional random dot product graphs, $(B^{(1)}, B^{(2)}, \mathbf{X}) \sim \text{JRDPG}(F, n, 2, \rho)$. Prior to the development of methods to jointly embed the networks, a common approach was to separately embed the two graphs into a common Euclidean space, and then align the networks via orthogonal Procrustes analysis (Tang et al., 2017a). One motivating question for the present work is how to capture the effect of the correlation on the embedded pair. Consider a simple, motivating example with $n = 300$, $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$ and F , a mixture of point mass distributions, defined via:

$$F = \frac{1}{2}\delta_{\xi_1} + \frac{1}{2}\delta_{\xi_2}, \quad (5)$$




Figure 1: The effect of induced versus inherent correlation in the embedded space. In panels 1–5 we plot, for various levels of ρ , the difference in the aligned estimates of X_1 from separately spectrally embedding $(B^{(1)}, B^{(2)}, \mathbf{X}) \sim \text{JRDPG}(F, n, 2, \rho)$ (where F has the form in Eq. 5). In the sixth panel, we consider $(B^{(1)}, B^{(2)}, \mathbf{X}) \sim \text{JRDPG}(F, 300, 2)$, and plot the difference in the pair of estimates of $X_1 = \xi_1$ derived from the omnibus embedding. In all panels, the experiment was repeated $nMC = 500$ times.

where $\xi_1, \xi_2 \in \mathbb{R}^2$ satisfy

$$\begin{bmatrix} \xi_1 \\ \xi_2 \end{bmatrix} \begin{bmatrix} \xi_1 \\ \xi_2 \end{bmatrix}^T = \begin{bmatrix} 0.7 & 0.3 \\ 0.3 & 0.5 \end{bmatrix},$$

so that the RDPGs drawn from F are examples of correlated *stochastic blockmodel* random graphs (Holland et al., 1983). In order to better understand the role of the correlation in the embedded space, we separately spectrally embed each network, $\widehat{\mathbf{X}}_{B^{(1)}} = \text{ASE}(B^{(1)}, 2)$ and $\widehat{\mathbf{X}}_{B^{(2)}} = \text{ASE}(B^{(2)}, 2)$, and then align the networks via $\widehat{\mathbf{X}}_{B^{(1)}} W^{(1)}$ and $\widehat{\mathbf{X}}_{B^{(2)}} W^{(2)}$ where for each $k = 1, 2$,

$$W^{(k)} = \operatorname{argmin}_{W \in \mathcal{O}_2} \|\widehat{\mathbf{X}}_{B^{(k)}} W - \mathbf{X}\|_F.$$

In Figure 1, we plot $(\widehat{\mathbf{X}}_{B^{(1)}} W^{(1)} - \widehat{\mathbf{X}}_{B^{(2)}} W^{(2)})_1$ (i.e., the first row of $\widehat{\mathbf{X}}_{B^{(1)}} W^{(1)} - \widehat{\mathbf{X}}_{B^{(2)}} W^{(2)}$), the distance between the (aligned) estimates of $X_1 = \xi_1$ derived from the embeddings over a range of values of ρ ; note that in each panel the experiment is repeated $nMC = 500$ times. In the first five panels of the figure, we see the effect of increasing ρ on the difference, namely that the covariance of the difference is monotonically decreasing as ρ increases.

In the sixth panel of Figure 1 (again performing $nMC = 500$ Monte Carlo replicates), we consider the conditionally independent case (i.e., $\rho = 0$), so that $(B^{(1)}, B^{(2)}, \mathbf{X}) \sim \text{JRDPG}(F, 300, 2)$. We consider the omnibus spectral embedding of $(B^{(1)}, B^{(2)})$, denoted $\widehat{\mathbf{X}}_M = \text{OMNI}(B^{(1)}, B^{(2)}, 2)$, and plot the difference in the estimates of $X_1 = \xi_1$ derived from the omnibus embedding, namely $(\widehat{\mathbf{X}}_M^{(1)})_1 - (\widehat{\mathbf{X}}_M^{(2)})_1$. From the figure, we see that the

correlation induced between the independent graphs in the embedded space is (roughly) equivalent to the inherent correlation between $\rho = 0.75$ correlated networks that have been separately embedded and aligned. In the next section, we will formalize this notion of induced versus inherent correlation, and we will see that, as the Figure suggests, OMNI does indeed induce correlation of level $\rho = 0.75$ across independent graphs in the embedded space.

3.1 Central Limit Theorems and Correlation in the Embedded Space

Viewing $(B^{(1)}, B^{(2)}, \mathbf{X})$ as an element of the sequence

$$\left((B_n^{(1)}, B_n^{(2)}, \mathbf{X}_n) \right)_{n=1}^{\infty},$$

where for each $n \geq 1$, $(B_n^{(1)}, B_n^{(2)}, \mathbf{X}_n) \sim \text{JRDPG}(F, n, 2, \rho)$, the Central Limit Theorem established in Athreya et al. (2016) provides a framework for understanding the effect, in the embedded space, of the edge-wise correlation across networks. Letting $\widehat{\mathbf{X}}_{B_n^{(1)}} = \text{ASE}(B_n^{(1)}, d)$ and $\widehat{\mathbf{X}}_{B_n^{(2)}} = \text{ASE}(B_n^{(2)}, d)$, Theorem 3.3 of Athreya et al. (2016) (in the form presented in Theorem 9 of Athreya et al. (2018a)) implies that for each $k = 1, 2$, if $\Delta := \mathbb{E}[X_1 X_1^T]$ is rank d , then there exist sequences of orthogonal d -by- d matrices $(\tilde{W}_n^{(k)})_{n=1}^{\infty}$ such that for all $z \in \mathbb{R}^d$ and for any fixed index i ,

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[n^{1/2} \left(\widehat{\mathbf{X}}_{B_n^{(k)}} \tilde{W}_n^{(k)} - \mathbf{X}_n \right)_i \leq z \right] = \int_{\text{supp } F} \Phi(z, \Sigma(x)) dF(x),$$

where

$$\Sigma(x) := \Delta^{-1} \mathbb{E} \left[(x^T X_1 - (x^T X_1)^2) X_1 X_1^T \right] \Delta^{-1}; \quad (6)$$

$\Phi(\cdot, \Sigma)$ denotes the cdf of a (multivariate) Gaussian with mean zero and covariance matrix Σ , and $(\widehat{\mathbf{X}}_n \tilde{W}_n - \mathbf{X}_n)_i$ denotes the i -th row of $\widehat{\mathbf{X}}_n \tilde{W}_n - \mathbf{X}_n$. Combining the above central limit theorems for $\widehat{\mathbf{X}}_{B_n^{(1)}}$ and $\widehat{\mathbf{X}}_{B_n^{(2)}}$, we have the following theorem (proven in Section A.1).

Theorem 9 *Let $\rho \in (0, 1)$ be fixed. Let F be a distribution on a set $\mathcal{X} \subset \mathbb{R}^d$, where $\langle x, x' \rangle \in [0, 1]$ for all $x, x' \in \mathcal{X}$, and assume that $\Delta := \mathbb{E}[X_1 X_1^T]$ is rank d . Let $(B_n^{(1)}, B_n^{(2)}, \mathbf{X}_n) \sim \text{JRDPG}(F, n, 2, \rho)$, be a sequence of adjacency matrices and associated latent positions, where for each $n \geq 1$ the rows of \mathbf{X}_n are i.i.d. distributed according to F . Letting $\widehat{\mathbf{X}}_{B_n^{(1)}} = \text{ASE}(B_n^{(1)}, d)$ and $\widehat{\mathbf{X}}_{B_n^{(2)}} = \text{ASE}(B_n^{(2)}, d)$, there exist sequences of orthogonal d -by- d matrices $(W_n^{(1)})_{n=1}^{\infty}$, $(W_n^{(2)})_{n=1}^{\infty}$ such that for all $z \in \mathbb{R}^d$ and for any fixed index i ,*

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[n^{1/2} \left(\widehat{\mathbf{X}}_{B_n^{(1)}} W_n^{(1)} - \widehat{\mathbf{X}}_{B_n^{(2)}} W_n^{(2)} \right)_i \leq z \right] = \int_{\text{supp } F} \Phi \left(z, \tilde{\Sigma}(x, \rho) \right) dF(x), \quad (7)$$

where $\tilde{\Sigma}(x, \rho) = 2(1 - \rho)\Sigma(x)$.

From Theorem 9, we see that the effect of the correlation ρ in the embedding space is to introduce a dampening factor of $(1 - \rho)$ into the asymptotic limiting covariance (see

Figure 1). This is entirely reasonable; indeed, consider Y_1, Y_2 to be ρ correlated $\text{Norm}(\mu, \sigma^2)$ random variables, in which case $Y_1 - Y_2 \sim \text{Norm}(0, 2(1 - \rho)\sigma^2)$. In the embedded space, no extraneous correlation is introduced (in the limit) by separately embedding the networks and aligning the embeddings via $(W_n^{(k)})_{n=1}^\infty$ for $k = 1, 2$. Joint embedding procedures like the Omnibus method forgo these Procrustes rotations, but at a price: they induce correlation across even independent networks. To understand this, we consider the omnibus central limit theorem of Levin et al. (2017) and derive the following result. Its proof can be obtained from Levin et al. (2017), but is also an immediate consequence of our more general main result, Theorem 16, which is stated formally in the following section.

Theorem 10 (Induced correlation in classical OMNI) *Let F be a distribution on a set $\mathcal{X} \subset \mathbb{R}^d$, where $\langle x, x' \rangle \in [0, 1]$ for all $x, x' \in \mathcal{X}$, and assume that $\Delta := \mathbb{E}[X_1 X_1^T]$ is rank d . Let $(A_n^{(1)}, A_n^{(2)}, \dots, A_n^{(m)}, \mathbf{X}_n) \sim \text{JRDPG}(F, n, m)$ be a sequence of independent RDPG random graphs, and for each $n \geq 1$, let M_n denote the omnibus matrix as in Definition 8. Also, let $\widehat{\mathbf{X}}_{M_n} = \text{ASE}(M_n, d)$ and denote the s -th $n \times d$ block of the omnibus embedding $\widehat{\mathbf{X}}_{M_n}$ as $\widehat{\mathbf{X}}_{M_n}^{(s)}$. Consider fixed indices $i \in [n]$ and $s_1, s_2 \in [m]$. Then there exists a sequence of orthogonal matrices $(\tilde{W}_n)_{n=1}^\infty$ such that for all $z \in \mathbb{R}^d$, we have that*

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[n^{1/2} \left[\left(\widehat{\mathbf{X}}_{M_n}^{(s_1)} - \widehat{\mathbf{X}}_{M_n}^{(s_2)} \right) \tilde{W}_n \right]_i \leq z \right] = \int_{\text{supp } F} \Phi(z, \tilde{\Sigma}(x, 3/4)) dF(x). \quad (8)$$

Unpacking Theorem 10, $(\widehat{\mathbf{X}}_{M_n}^{(s_1)})_i$ and $(\widehat{\mathbf{X}}_{M_n}^{(s_2)})_i$ represent the estimates of X_i derived from $A^{(s_1)}$ and $A^{(s_2)}$ by the omnibus embedding paradigm. In light of Theorem 9, the induced correlation in OMNI can be understood in the context of the limiting covariance of the difference of a pair of estimates for the same underlying latent position. Comparing Eqs. (7) and (8), we see that OMNI effectively induces a correlation of level $\rho = 3/4$ uniformly across embedded network pairs.

While it is easy to surmise that this flat correlation would have a detrimental signal dampening effect on inference tasks in settings where there is nuanced inherent correlation (e.g., in time-series analysis), this theory is in the setting of conditionally independent networks. It is natural to ask whether in the presence of inherent correlation across the network pairs the dampening effect of the induced correlation in OMNI diminished. As we will see in our next result (a special case of Theorem 16 in Section 4.1), this is not the case.

Theorem 11 (Induced and inherent correlation in classical OMNI) *Let F be a distribution on a set $\mathcal{X} \subset \mathbb{R}^d$, where $\langle x, x' \rangle \in [0, 1]$ for all $x, x' \in \mathcal{X}$, and assume that $\Delta := \mathbb{E}[X_1 X_1^T]$ is rank d . Let*

$$(A_n^{(1)}, A_n^{(2)}, \dots, A_n^{(m)}, \mathbf{X}_n) \sim \text{JRDPG}(F, n, m, R)$$

be a sequence of correlated RDPG random graphs, and for each $n \geq 1$, let M_n denote the omnibus matrix as in Definition 8. Also, let $\widehat{\mathbf{X}}_{M_n} = \text{ASE}(M_n, d)$ and denote the s -th $n \times d$ block of the omnibus embedding $\widehat{\mathbf{X}}_{M_n}$ as $\widehat{\mathbf{X}}_{M_n}^{(s)}$. Consider fixed indices $i \in [n]$ and $s_1, s_2 \in [m]$. Then there exists a sequence of orthogonal matrices $(\tilde{W}_n)_{n=1}^\infty$ such that for all $z \in \mathbb{R}^d$, we have that

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[n^{1/2} \left[\left(\widehat{\mathbf{X}}_{M_n}^{(s_1)} - \widehat{\mathbf{X}}_{M_n}^{(s_2)} \right) \tilde{W}_n \right]_i \leq z \right] = \int_{\text{supp } F} \Phi \left(z, \tilde{\Sigma}(x, \rho(s_1, s_2)) \right) dF(x), \quad (9)$$

where $\rho(s_1, s_2) = \frac{3}{4} + \frac{1}{4}\rho_{s_1, s_2}$.

In the JRDPG(F, n, m, R) setting, the form of the limiting correlation highlights the separate contributions from the *method* (i.e., the induced correlation 3/4) and the *model* (i.e., the inherent correlation ρ_{s_1, s_2}). This correlation is independent of m , and the outsized effect of the induced correlation versus the downscaled inherent correlation further suggests that OMNI is not ideal for embedding temporal sequences of networks that exhibit complex dependency patterns (i.e., a change point or anomaly). Indeed, the induced correlation from OMNI whitens out the existing complex correlation structures amongst the $A^{(k)}$, and, as we will show in our next example, this whitening effect of OMNI can serve to mask significant data features and structures in complex data environs.

3.2 OMNI’s Whitening Correlation: Aplysia Spike Train Analysis

We consider a time series of networks derived from the Aplysia californica escape motor program of Hill et al. (2020). The motor program consists of a 20 min recording of the action potentials generated by 82 neurons in the dorsal pedal ganglion of an isolated brain preparation from the marine mollusk Aplysia californica. One minute into the recording a brief electrical stimulus was applied to pedal ganglion nerve 9 to elicit the animal’s rhythmic escape locomotion motor program. This consists of an initial rapid bursting, lasting several cycles, that drives the animal’s gallop behavior, followed by a slower rhythm persisting until the end of the recording that drives the animal’s crawling behavior. To extract a network time series from the recording, we binned the motor program into 24 bins, each approximately ≈ 50 second long; the binned motor program (with the second bin, containing the stimulus, highlighted) is pictured in Figure 2. Using the `meaRtools` package in R (Gelfman et al., 2018), we apply the STTC (spike time tiling correlation) method of Cutts and Eglen (2014) to convert each 50 second window into a weighted correlation matrix amongst the 82 neurons. Each of the 24 bins then yields one weighted graph on 82 vertices representing this correlation structure amongst the 82 neurons. The stimulus occurs then in the second of the 24 bins, followed by the gallop and crawling motor programs.

3.2.1 CLASSICAL OMNIBUS EMBEDDINGS AND CORRELATION MASKING

Our motivation in the following analyses is to determine whether, in light of the correlation structure in Theorem 11, the classical OMNI method can detect the distinct phases of behavior in the motor program. In order to explore the impact of the flat correlation induced in the classical omnibus embedding, we use the classical OMNI spectral embedding to embed the $\{A^{(k)}\}_{k=1}^{24}$ into a common \mathbb{R}^d (where $d = 4$ as chosen by locating the elbow in the scree plot as suggested by Zhu and Ghodsi (2006); Chatterjee (2015)). Further isolating the impact of the stimulus and the evolution of the galloping and crawling phases of the motor program, we plot the average distance between the estimated latent positions for each vertex (as the bar heights) in the embedding between embedded Graph 1 (res., embedded Graph 2) and embedded graph k for each $k \in [m]$ (where the location of the bars in Figure 3 correspond to $k = 1, 2, \dots, 24$). To wit, we plot (where $\widehat{\mathbf{X}}_{M_n} := [(\widehat{\mathbf{X}}_{M_n}^{(1)})^T \mid (\widehat{\mathbf{X}}_{M_n}^{(2)})^T \mid \dots \mid (\widehat{\mathbf{X}}_{M_n}^{(m)})^T]^T$)

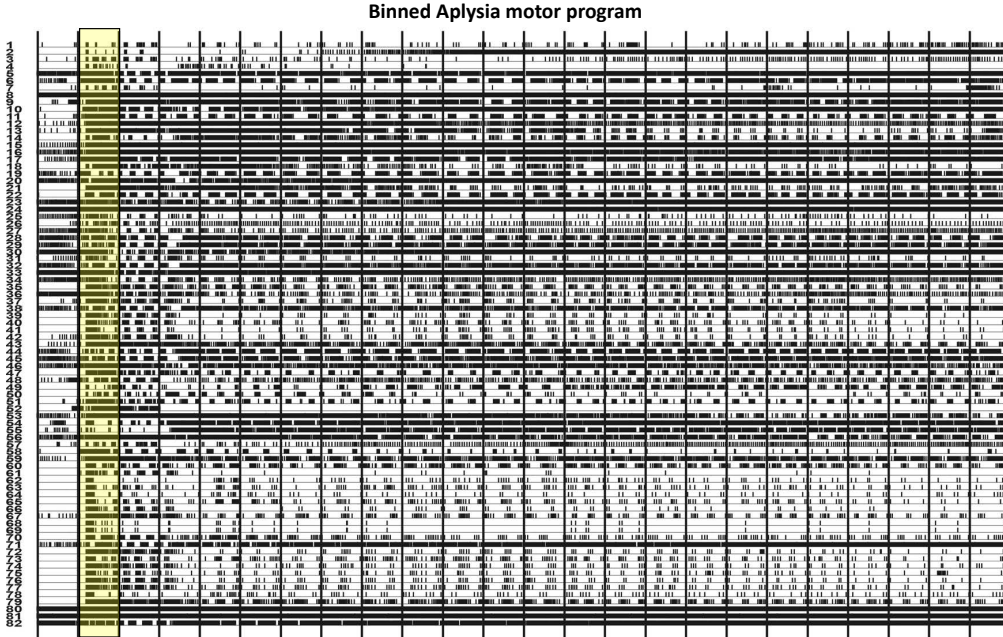


Figure 2: The 20 minute Aplysia escape motor program from Hill et al. (2020), binned into 24 windows, each approximately 50 seconds in length. The stimulus happens one minute into the motor program, in the highlighted second bin.

$$\frac{1}{82} \sum_{i=1}^{82} \left\| \left(\widehat{\mathbf{X}}_{M_n}^{(1)} \right)_i - \left(\widehat{\mathbf{X}}_{M_n}^{(k)} \right)_i \right\|_2 \text{ in the left panel in Figure 3;}$$

$$\frac{1}{82} \sum_{i=1}^{82} \left\| \left(\widehat{\mathbf{X}}_{M_n}^{(2)} \right)_i - \left(\widehat{\mathbf{X}}_{M_n}^{(k)} \right)_i \right\|_2 \text{ in the center panel in Figure 3.}$$

From Figure 2 and our knowledge of the Aplysia motor program’s evolution, we see that if the omnibus embedding is able to detect the biologically distinct phases of this network time series, then

- i. There will be a significant difference between the embeddings of graphs 1 and 2 (as demonstrated in Figure 3).
- ii. The effect of the stimulus is less apparent as the Aplysia transitions from galloping to crawling; this would be manifested as the distance between Graphs 2 and $k > 2$ increasing as the Aplysia transitions from galloping to crawling.
- iii. The distance from Graph 1 to graphs $k > 2$ should be large, as the Aplysia never returns to its spontaneous firing state in the motor program.

While we see that the omnibus methodology demonstrates the capacity to detect the anomaly (namely, the stimulus) in the second graph, the flat correlation structure induced




Figure 3: In the left (resp., center) panel, we plot the average vertex distance (as the bar heights) in the embedding between Graph 1 (resp., Graph 2) and graph k for each $k \in [m]$ in OMNI (where the bars are labeled $k = 1, 2, \dots, 24$). In the right panel, we compute the 24×24 distance matrices \mathbf{D} , and embed it into \mathbb{R}^2 . The resulting 24 data points are clustered using `Mclust` (clusters denoted by shape), and are plotted, with Graphs 1 and 2 further labeled with their corresponding number.


in the embedding space has the effect of masking the transition from galloping to crawling and creates an artificial similarity between Graphs 1 and some of the graphs $k > 2$ in the embedded space (as shown in the right panel of Figure 3).

Exploring this further, we compute the 24×24 distance matrix \mathbf{D} , where

$$\mathbf{D} = [D_{k,\ell}] \text{ where } D_{k,\ell} = \|\hat{\mathbf{X}}_{M_n}^{(k)} - \hat{\mathbf{X}}_{M_n}^{(\ell)}\|_F. \quad (10)$$

The matrix \mathbf{D} is then each embedded into \mathbb{R}^2 (where $d = 2$ is, again, chosen by the elbow in the scree plot as suggested by Zhu and Ghodsi (2006); Chatterjee (2015); see Figure 4), and the 24 data points are clustered using `Mclust` (Fraley and Raftery, 2002). The resulting clusters are plotted in the right panel of Figure 3. Points 1 and 2 (corresponding to Graphs 1 and 2 resp.) are plotted with larger symbols and labeled. In the MDS embedding, each graph is represented by a single 2-dimensional embedded data point, and the pair (Graph, Cluster) describes the cluster to which `Mclust` assigns these points. The cluster labels found by `Mclust` are given in Table 1. Ideally, the clustering would partition the motor program into the distinct phases of the biological activity. If this were the case, then Graph 1 would be a singleton cluster, Graph 2 would be a singleton cluster, there would be a cluster capturing the gallop phase (Graphs 3 and 4), and distinct clusters capturing the rhythmic crawling phase (Graphs 5–24).

This clustering reinforces the finding that while the stimulus is detected (Graph 2 is clustered apart from the others), the transition from stimulus to gallop and crawl and the distinct nature of Graph 1 (as the only spontaneous firing state measurement) are masked in this analysis. Indeed, the flat correlation structure induced by the classical OMNI embedding yields that Graph 1 is clustered together with graphs from the crawling phase, and graphs from the gallop phase are clustered together with graphs from the crawling phase. We shall see in Section 5.1 that different structures in the off-diagonal portion of the

Figure 4: Scree plot of the eigenvalues of \mathbf{D} .

Graph position in time-series	1	2	3	4	5	6	7	8	9	10	11	12
Cluster assignment	1	2	5	3	5	5	4	5	3	1	1	1
Graph position in time-series	13	14	15	16	17	18	19	20	21	22	23	24
Cluster assignment	1	5	3	1	1	4	1	3	3	1	4	3

Table 1: This table displays the cluster labels/assignments for the 24 networks when embedding the time-series using classical OMNI, then embedding the across-graph distance matrix using classical MDS scaling, and finally clustering the graphs using `Mclust`. In the MDS embedding, each graph is represented by a single 2-dimensional embedded data point, and the pair (Graph, Cluster) describes the cluster to which `Mclust` assigns these points.

Omnibus matrix can be introduced to ameliorate the issues caused by the flat correlation and in doing so recover the distinct biological phases of this motor program.

4. Generalized Omnibus Embeddings

All joint embedding procedures induce correlation across the embedded networks, and the omnibus embedding is no exception. Unfortunately, the particular structure of the correlation induced by classical OMNI (namely, the large, flat induced correlation across all graphs) may render it less effective for time-series applications or settings where the correlation varies dramatically across networks. This motivates our next contribution, in which we show that by generalizing the structure of the omnibus matrix M , more nuanced (and application appropriate) induced correlation structures are possible.

The consistency and asymptotic normality of the classical omnibus embedding rest on a few defining model assumptions: first,

$$\mathbb{E}[M] = \tilde{P} := J_m \otimes P,$$

and second, in the k -th block-row of M , the weight put on $A^{(k)}$ (which is equal to $1 + (m - 1)/2$) is strictly greater than the weight put on any $A^{(\ell)}$ for $\ell \neq k$ (as each of these is $1/2$). The low-rank RDPG structure of $\mathbb{E}[M] = \tilde{P}$ allows us to use random matrix concentration results to prove that the scaled eigenvectors of M are tightly concentrated about the scaled eigenvectors of \tilde{P} which, up to a possible rotation, are equal to $\mathbf{Z} := \tilde{\mathbf{I}}_m \otimes \mathbf{X}$. The weights in block-row k being maximized for $A^{(k)}$ ensures that the k -th block in $\text{ASE}(M, d)$ corresponds to the embedding of $A^{(k)}$, which is essential for subsequent inference (e.g., hypothesis testing (Draves and Sussman, 2020)) in the omnibus setting.

It is natural to ask if, subject to the above conditions, we can generalize the omnibus structure to permit more esoteric induced correlation in the embedded space. This motivates the following definition of the generalized omnibus matrix. Let F be a distribution on a set $\mathcal{X} \subset \mathbb{R}^d$, where $\langle x, x' \rangle \in [0, 1]$ for all $x, x' \in \mathcal{X}$. Suppose that $(A^{(1)}, A^{(2)}, \dots, A^{(m)}, \mathbf{X}) \sim \text{JRDPG}(F, n, m, R)$. The convex hull of $A^{(1)}, A^{(2)}, \dots, A^{(m)}$ is denoted via

$$\mathcal{A}_m = \left\{ \sum_{q=1}^m c_q A^{(q)} \in \mathbb{R}^{n \times n} \mid \text{with } c_q \geq 0 \text{ for all } q \in [m] \text{ and } \sum_{q=1}^m c_q = 1 \right\}.$$

Definition 12 (Generalized Omnibus Matrix) Consider $\mathfrak{M} \in \mathbb{R}^{mn \times mn}$, a generalized version of the omnibus M of Definition 8 defined as follows. \mathfrak{M} is a block matrix satisfying the following assumptions:

1. Each block entry $\mathfrak{M}^{(k, \ell)} \in \mathbb{R}^{n \times n}$, $k, \ell \in [m]$ is an element of \mathcal{A}_m ; we will write

$$\mathfrak{M}^{(k, \ell)} = \sum_{q=1}^m c_q^{(k, \ell)} A^{(q)}.$$

2. For each block row $k \in [m]$ of \mathfrak{M} , the cumulative weight of $A^{(k)}$ is greater than the cumulative weight of the rest of the $A^{(q)}$, $q \neq k$; i.e., for $q \neq k$

$$\sum_{\ell} c_q^{(k, \ell)} < \sum_{\ell} c_k^{(k, \ell)}.$$

3. \mathfrak{M} is symmetric.

Such a block matrix satisfying assumptions 1-3 above will be referred to as a Generalized Omnibus Matrix.

If \mathfrak{M} is a Generalized Omnibus Matrix, setting $\alpha(k, q) = \sum_{\ell} c_q^{(k, \ell)}$ to be the weight put on $A^{(q)}$ in the k -th block-row of \mathfrak{M} , assumption 2 above becomes, for each $k \in [m]$,

$$\alpha(k, q) < \alpha(k, k) \quad \text{for all } q \neq k. \quad (11)$$

Because \mathfrak{M} (i.e., $\mathbb{E}(\mathfrak{M}) = \tilde{P}$) is unbiased and low-rank, we can appropriately modify matrix concentration and perturbation results for the scaled eigenvectors of \mathfrak{M} , and we can establish that the rows of $\text{ASE}(\mathfrak{M}, d)$ consistently estimate the associated rows of \mathbf{Z} and the associated residual errors satisfy a distributional central limit theorem. This is the content of the following theorems, whose proofs can be found in Appendix.

Theorem 13 (Consistency of generalized omnibus embedding estimates) *Let F be a distribution on a set $\mathcal{X} \subset \mathbb{R}^d$, where $\langle x, x' \rangle \in [0, 1]$ for all $x, x' \in \mathcal{X}$, and assume that $\Delta := \mathbb{E}[X_1 X_1^T]$ is rank d . Let*

$$(A_n^{(1)}, A_n^{(2)}, \dots, A_n^{(m)}, \mathbf{X}_n) \sim \text{JRDPG}(F, n, m, R)$$

be a sequence of correlated RDPG random graphs. For each $n \geq 1$, let \mathfrak{M}_n denote the generalized omnibus matrix as in Definition 12. Let the spectral decomposition of $\tilde{P}_n = \mathbb{E}(\mathfrak{M}_n)$ be given by

$$\tilde{P}_n = \mathbb{E}(\mathfrak{M}_n) = U_{\tilde{P}_n} S_{\tilde{P}_n} U_{\tilde{P}_n}^T,$$

where $U_{\tilde{P}_n} \in \mathbb{R}^{mn \times d}$ and $S_{\tilde{P}_n} \in \mathbb{R}^{d \times d}$. Let $\hat{\mathbf{X}}_{\mathfrak{M}_n} = \text{ASE}(\mathfrak{M}_n, d) = U_{\mathfrak{M}_n} S_{\mathfrak{M}_n}^{1/2}$ denote the adjacency spectral embedding of \mathfrak{M}_n . Then there exists an orthogonal matrix $V_n \in \mathcal{O}_d$ and a constant $C > 0$ such that, with high probability for n sufficiently large,

$$\|U_{\mathfrak{M}_n} S_{\mathfrak{M}_n}^{1/2} - U_{\tilde{P}_n} S_{\tilde{P}_n}^{1/2} V_n\|_{2 \rightarrow \infty} \leq C \frac{m^{1/2} \log mn}{n^{1/2}}.$$

The above theorem is analogous to Lemma 1 in Levin et al. (2017). Interestingly, the consistency error rate for the generalized omnibus embedding estimates with pair-wise correlated adjacency matrices coincides with the consistency error rate for the classical omnibus embedding estimates with independent adjacency matrices. This is a consequence of two facts: first, the sum of the cumulative weights $\sum_{q=1}^m \alpha(k, q)$ is equal to m , and second, the extra correlation term in the matrix Bernstein concentration bound is of the same order as in the independent term.

Building upon this consistency, our principal result describes the limiting distributional behavior of low-dimensional embeddings of the generalized omnibus matrix for joint networks, stated below.

Theorem 14 (Asymptotic normality of rows of the generalized omnibus embedding)

Let F be a distribution on a set $\mathcal{X} \subset \mathbb{R}^d$, where $\langle x, x' \rangle \in [0, 1]$ for all $x, x' \in \mathcal{X}$, and assume that $\Delta := \mathbb{E}[X_1 X_1^T]$ is rank d . Let $(A_n^{(1)}, A_n^{(2)}, \dots, A_n^{(m)}, \mathbf{X}_n) \sim \text{JRDPG}(F, n, m, R)$ be a sequence of correlated RDPG random graphs, and let $\mathbf{Z}_n = [\mathbf{X}_n^T | \mathbf{X}_n^T | \dots | \mathbf{X}_n^T]^T \in \mathbb{R}^{mn \times d}$. Further, for each $n \geq 1$, let \mathfrak{M}_n denote the generalized omnibus matrix as in Definition 12 and $\hat{\mathbf{X}}_{\mathfrak{M}_n} = \text{ASE}(\mathfrak{M}_n, d)$. For fixed indices $i \in [n]$ and $s \in [m]$, let $(\hat{\mathbf{X}}_{\mathfrak{M}_n}^{(s)})_i$ denote the i -th row of the s -th $n \times d$ block of $\hat{\mathbf{X}}_{\mathfrak{M}_n}$ (i.e., the estimated latent position from graph $A_n^{(s)}$ of the hidden latent position X_i). Then, there exists a sequence of orthogonal matrices $\{\mathcal{Q}_n\}_n$ such that

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[n^{1/2} \left((\hat{\mathbf{X}}_{\mathfrak{M}_n}^{(s)} \mathcal{Q}_n - \mathbf{Z}_n \right)_i \leq z \right] = \int_{\text{supp} F} \Phi(z, \check{\Sigma}_\rho(x; s)) dF(x),$$

where

$$\check{\Sigma}_\rho(x; s) = \frac{1}{m^2} \left(\underbrace{\sum_{q=1}^m \alpha^2(s, q)}_{\text{method-induced coefficient}} + 2 \underbrace{\sum_{q < l} \alpha(s, q) \alpha(s, l) \rho_{q,l}}_{\text{model-inherent coefficient}} \right) \underbrace{\Delta^{-1} \mathbb{E}[x^T X_j (1 - x^T X_j) X_j X_j^T] \Delta^{-1}}_{:=\Sigma(x)}$$

The covariance matrix $\check{\Sigma}_\rho(x_i; s)$ corresponds to the residual of the (true) latent position of vertex i and its estimate from graph $A^{(s)}$. We note that this covariance matrix can be written as a sum of two terms, for which the first term

$$\Sigma_0(x_i; s) := \left(\frac{1}{m^2} \sum_{q=1}^m \alpha^2(s, q) \right) \Sigma(x_i)$$

corresponds to the covariance matrix for the i -th row from the s -th block of the generalized omnibus embedding under the JRDPG(F, n, m) model as in Definition 3, and the second term is accredited to the presence of the inherent correlation from the JRDPG(F, n, m, R) model. By noting that $\rho_{k,\ell} \geq 0$ and $\alpha(k, \ell) \geq 0$ for all $k, \ell \in [m]$, we deduce that $\check{\Sigma}_\rho(x_i; s) \geq \Sigma_0(x_i; s)$ entry-wise for all i , implying tighter residual errors when the estimates arise from independent graphs. This observation motivates our simulations in Section 7 where we explore the effect of the inherent correlation on effective sample size in the context of community detection in SBM's.

Remark 15 Note that for fixed n , the orthogonal matrices V_n, Q_n in Theorems 13 and 14 accordingly, can be explicitly defined. The matrix V_n is the solution of the Procrustes problem

$$\min_{V \in \mathcal{O}_d} \|U \mathfrak{M}_n - U_{\tilde{P}_n} V\|,$$

and its solution is given by $V_n := V_{1,n} V_{2,n}^T$ where the columns of the matrices $V_{1,n}, V_{2,n}$ are the left- and right-singular vectors of $U_{\tilde{P}_n}^T U \mathfrak{M}_n$ respectively. Moreover, $Q_n := V_n^T W_n$, where W_n is an orthogonal transformation such that $\mathbf{Z}_n = \hat{\mathbf{Z}}_n W_n = U_{\tilde{P}_n} S_{\tilde{P}_n}^{1/2} W_n$.

In the following example, we apply Theorem 14 with the (classical) omnibus matrix from Levin et al. (2017).

Example 1 (Classical omnibus matrix)

Suppose first that $(A^{(1)}, A^{(2)}, \dots, A^{(m)}, \mathbf{X}) \sim \text{JRDPG}(F, n, m, R)$. Let $e_k \in \mathbb{R}^m$ be the vector with the k -th entry equal to 1 and the rest equal to 0, and let $e_{k\ell} = e_k e_\ell^T + e_\ell e_k^T \in \mathbb{R}^{m \times m}$. Then, the omnibus matrix $M \in \mathbb{R}^{mn \times mn}$ is defined as

$$M = \frac{1}{2} \sum_{k \leq \ell} e_{k\ell} \otimes (A^{(k)} + A^{(\ell)})$$

As mentioned previously, M is a special case of the generalized omnibus matrix \mathfrak{M} : it satisfies all the assumptions of Definition 12. For $q \in [m]$, the coefficient matrix $C^{(q)}$ is given by

$$c_q^{(k,\ell)} = \begin{cases} 1 & \text{if } k = \ell = q \\ 1/2 & \text{if } k \neq \ell, \text{ and } q \in \{k, \ell\} \\ 0 & \text{else,} \end{cases}$$

and so

$$\alpha(s, q) = \begin{cases} 1 + (m - 1)/2 & \text{if } s = q \\ 1/2 & \text{else.} \end{cases}$$

Fix $s \in [m]$. Then, the resulting covariance matrix is

$$\check{\Sigma}_\rho(x; s) = \left(\underbrace{\frac{m+3}{4m}}_{\text{method-induced coefficient}} + \underbrace{\frac{m+1}{2m^2} \sum_{q \neq s} \rho_{q,s} + \frac{1}{2m^2} \sum_{\substack{q < l \\ q, l \neq s}} \rho_{q,l}}_{\text{model-inherent coefficient}} \right) \Sigma(x).$$

As expected, the first term coincides with the covariance structure in classical OMNI (Levin et al., 2017, Theorem 1) and further, the second term accounts for the pair-wise correlation of the adjacency matrices.

Moving beyond the classical case, the following examples of generalized omnibus structures highlight the impact the structure of \mathfrak{M} has on the limiting covariance (and hence on the limiting correlation; see Theorem 16).

Example 2 (Total Average Omnibus) In the total average case, letting

$$\bar{A} = \frac{A^{(1)} + A^{(2)} + \dots + A^{(m)}}{m},$$

we define the generalized omnibus matrix $\mathfrak{M}_{\bar{A}}$ via

$$\mathfrak{M}_{\bar{A}} = \begin{bmatrix} A^{(1)} & \bar{A} & \dots & \bar{A} \\ \bar{A} & A^{(2)} & \dots & \bar{A} \\ \vdots & \vdots & \ddots & \vdots \\ \bar{A} & \bar{A} & \dots & A^{(m)} \end{bmatrix}.$$

In this example,

$$c_q^{(k, \ell)} = \begin{cases} 1 & \text{if } k = \ell = q; \\ 0 & \text{if } k = \ell \neq q; \\ 1/m & \text{else,} \end{cases}$$

and so

$$\alpha(s, q) = \begin{cases} 1 + (m - 1)/m & \text{if } s = q \\ (m - 1)/m & \text{else.} \end{cases}$$

The associated coefficient of $\Sigma(x)$ in $\Sigma_0(x; s)$ (i.e., method coefficient) is then

$$\begin{aligned} \frac{1}{m^2} \left(\sum_{q=1}^m \alpha^2(s, q) \right) &= \frac{1}{m^2} \left((m-1)^3/m^2 + (1 + (m-1)/m)^2 \right) \\ &= \frac{m^2 + m - 1}{m^3}, \end{aligned}$$

and the model coefficient is

$$\frac{2}{m^2} \sum_{q < l} \alpha(s, q) \alpha(s, l) \rho_{q, l} = \frac{2}{m^2} \left(\frac{(m-1)(2m-1)}{m^2} \sum_{q \neq s} \rho_{q, s} + \frac{(m-1)^2}{m^2} \sum_{\substack{q < l \\ q, l \neq s}} \rho_{q, l} \right).$$

When $R \equiv 0_m$ (i.e., in the uncorrelated case), in the classical OMNI setting of Example 1 for large m the limiting covariance $\Sigma_0(x; s)$ is approximately $\Sigma(x)/4$, and is not degenerate; in the total average omnibus setting for large m , the limiting covariance is approximately $\Sigma(x)/m \approx 0$. This is sensible heuristically, as in that setting, we are effectively embedding $J_m \otimes \bar{A} \approx \tilde{P}$, and the correct scaling of the residuals would ideally be \sqrt{nm} rather than \sqrt{n} . When $R = \Theta(1)$ and m large, the model coefficient dominates the method coefficient, as even when m is large, the model coefficient is not (approximately) degenerate.

Example 3 (Weighted Pairwise Average Omnibus) In the classical omnibus setting, we have that

$$\mathfrak{M}^{(k, \ell)} = \frac{A^{(k)} + A^{(\ell)}}{2},$$

and all matrices are effectively weighted equally in the omnibus matrix. This is sensible if all $A^{(k)}$ are i.i.d., but is, perhaps, less ideal in the setting where the networks are noisily observed with the level of noise varying in $k \in [m]$. In that setting, the *Weighted Pairwise Average Omnibus* matrix defined via

$$\mathfrak{M}_W = \begin{bmatrix} A^{(1)} & \frac{w_1 A^{(1)} + w_2 A^{(2)}}{w_1 + w_2} & \frac{w_1 A^{(1)} + w_3 A^{(3)}}{w_1 + w_3} & \dots & \frac{w_1 A^{(1)} + w_m A^{(m)}}{w_1 + w_m} \\ \frac{w_2 A^{(2)} + w_1 A^{(1)}}{w_2 + w_1} & A^{(2)} & \frac{w_2 A^{(2)} + w_3 A^{(3)}}{w_2 + w_3} & \dots & \frac{w_2 A^{(2)} + w_m A^{(m)}}{w_2 + w_m} \\ \frac{w_3 A^{(3)} + w_1 A^{(1)}}{w_3 + w_1} & \frac{w_3 A^{(3)} + w_2 A^{(2)}}{w_3 + w_2} & A^{(3)} & \dots & \frac{w_3 A^{(3)} + w_m A^{(m)}}{w_3 + w_m} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ \frac{w_m A^{(m)} + w_1 A^{(1)}}{w_m + w_1} & \frac{w_m A^{(m)} + w_2 A^{(2)}}{w_m + w_2} & \frac{w_m A^{(m)} + w_3 A^{(3)}}{w_m + w_3} & \dots & A^{(m)} \end{bmatrix}$$

may be more appropriate. In this case, the block entries of the omnibus matrix are defined via

$$\mathfrak{M}_W^{(k, \ell)} = \frac{w_k A^{(k)} + w_\ell A^{(\ell)}}{w_k + w_\ell}$$

with weights $w_k > 0$ $k, \ell \in [m]$. In this example,

$$c_q^{(k, \ell)} = \begin{cases} 1 & \text{if } k = \ell = q \\ \frac{w_q}{w_q + w_\ell} & \text{if } q = k \neq \ell \\ \frac{w_q}{w_q + w_k} & \text{if } q = \ell \neq k \\ 0 & \text{else,} \end{cases}$$

and so

$$\alpha(s, q) = \begin{cases} 1 + \sum_{r \neq s} \frac{w_s}{w_s + w_r} & \text{if } s = q \\ \frac{w_q}{w_q + w_s} & \text{else.} \end{cases}$$

As (if each $w_k > 0$) $\alpha(s, q) < 1$ for $q \neq s$, we immediately have that $\alpha(s, s) > \alpha(s, q)$ for $q \neq s$. From this, it follows that in the weighted pairwise average omnibus setting, the covariance matrix is given by

$$\begin{aligned} \check{\Sigma}_\rho(x; s) &= \frac{1}{m^2} \left(\underbrace{\sum_{q \neq s} \frac{w_q^2}{(w_q + w_s)^2} + \left(1 + \sum_{r \neq s} \frac{w_s}{w_s + w_r}\right)^2}_{\text{method-induced coefficient}} \right. \\ &\quad \left. + 2 \underbrace{\sum_{q \neq s} \left(1 + \sum_{r \neq s} \frac{w_s}{w_s + w_r}\right) \frac{w_q}{w_q + w_s} \rho_{q,s} + 2 \sum_{\substack{q < l \\ q, l \neq s}} \frac{w_q w_l}{(w_q + w_s)(w_l + w_s)} \rho_{q,l}}_{\text{model-inherent coefficient}} \right) \Sigma(x) \end{aligned}$$

While the method coefficient of $\Sigma_0(x; s)$ in $\check{\Sigma}_\rho(x; s)$ is not easily expressed in general, specific examples can nonetheless be instructive. Consider the setting where $w_1 = w$ and $w_k = 1$ for all $k \neq 1$. Considering $s = 1$ in Theorem 14 provides that the method coefficient of $\check{\Sigma}_\rho(x; s)$ is

$$\begin{aligned} \frac{1}{m^2} \left(\sum_{q=1}^m \alpha^2(s, q) \right) &= \frac{1}{m^2} \left((m-1) \frac{1}{(1+w)^2} + \left(1 + \frac{(m-1)w}{1+w}\right)^2 \right) \\ &= \frac{m-1 + (1+mw)^2}{m^2(1+w)^2}. \end{aligned}$$

Similarly, the model coefficient is

$$\frac{2}{m^2} \sum_{q < l} \alpha(s, q) \alpha(s, l) \rho_{q,l} = \frac{2}{m^2} \left(\frac{1+mw}{(w+1)^2} \sum_{q \neq s} \rho_{q,s} + \frac{1}{(1+w)^2} \sum_{\substack{q < l \\ q, l \neq s}} \rho_{q,l} \right)$$

If $w \gg m$ is large, then the method coefficient is approximately equal to 1 and the model coefficient is approximately equal to 0, which is the limiting covariance achieved by embedding $A^{(1)}$ separately. This will be further explained in the context of limiting induced correlation in the next section. If $w \approx 1$, then both model and method correlation are approximately equal to their analogues in the classical OMNI setting. If $w \ll 1$ is small, then the method coefficient is approximately $1/m$, and the model correlation is effectively a function of $\rho_{k,\ell}$ for $k, \ell \neq s$ for m large (assuming, for the moment, all the inherent correlations are of the same relative order). As the number of graphs increases, the relative impact on the overall embedding of $A^{(1)}$ decreases; this is in direct contrast to the classical OMNI setting, where each graph has the same (non-trivial) relative import in the embedded space.

4.1 Limiting Inherent and Induced Correlation

In the generalized omnibus embedding, similar to the classical setting (see Theorem 11), we can precisely compute the limiting correlation between estimates of the same latent position in the embedded space. To wit, we have the following theorem; as with the other main results, its proof can be found in the Appendix, specifically Section A.4.

Theorem 16 *With notation and assumptions as in Theorem 14, consider fixed indices $i \in [n]$ and $s, s_1, s_2 \in [m]$. Let $\left(\widehat{\mathbf{X}}_{\mathfrak{M}_n}^{(s)}\right)_i$ denote the estimated latent position from graph $A_n^{(s)}$ of the hidden latent position X_i . There exists a sequence of orthogonal matrices $(\mathcal{Q}_n)_{n=1}^\infty$ such that for all $z \in \mathbb{R}^d$, we have that*

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[n^{1/2} \left[\left(\widehat{\mathbf{X}}_{\mathfrak{M}_n}^{(s_1)} - \widehat{\mathbf{X}}_{\mathfrak{M}_n}^{(s_2)} \right) \mathcal{Q}_n \right]_i \leq z \right] = \int_{\text{supp } F} \Phi \left(z, \widetilde{\Sigma}_\rho(x; s_1, s_2) \right) dF(x), \quad (12)$$

where $\widetilde{\Sigma}_\rho(x; s_1, s_2)$ is given by

$$\begin{aligned} \widetilde{\Sigma}_\rho(x; s_1, s_2) &= \frac{1}{m^2} \underbrace{\left(\sum_{q=1}^m (\alpha(s_1, q) - \alpha(s_2, q))^2 \right)}_{\text{method-induced correlation}} \\ &\quad + 2 \underbrace{\sum_{q < l} (\alpha(s_1, q) - \alpha(s_2, q)) (\alpha(s_1, l) - \alpha(s_2, l)) \rho_{q,l}}_{\text{model-inherent correlation}} \Sigma(x) \quad (13) \\ &= \frac{1}{m^2} \left(2 \sum_{q < l} (\alpha(s_1, q) - \alpha(s_2, q)) (\alpha(s_1, l) - \alpha(s_2, l)) (\rho_{q,l} - 1) \right) \Sigma(x). \end{aligned}$$

The covariance matrix $\widetilde{\Sigma}_\rho(x; s_1, s_2)$ from Theorem 16 can be expressed in terms of the limiting correlation $\rho(s_1, s_2)$, where $\rho(s_1, s_2)$ via:

$$\begin{aligned} \rho(s_1, s_2) &= 1 - \underbrace{\frac{\sum_{q=1}^m (\alpha(s_1, q) - \alpha(s_2, q))^2}{2m^2}}_{\text{method-induced correlation}} \\ &\quad + \underbrace{\frac{\sum_{q < l} (\alpha(s_1, q) - \alpha(s_2, q)) (\alpha(s_2, l) - \alpha(s_1, l)) \rho_{q,l}}{m^2}}_{\text{model-inherent correlation}}. \quad (14) \end{aligned}$$

Before considering the more exotic examples described above, let's recall the classical omnibus embedding setting.

Example 1 continued: Consider fixed indices $i \in [n]$ and $s_1, s_2 \in [m]$. Then, the limiting correlation between two estimated latent positions $\left(\widehat{\mathbf{X}}_{M_n}^{(s_1)}\right)_i, \left(\widehat{\mathbf{X}}_{M_n}^{(s_2)}\right)_i$ from graphs $A^{(s_1)}, A^{(s_2)}$ respectively is given by

$$\rho(s_1, s_2) = \frac{3}{4} + \frac{1}{4} \rho_{s_1, s_2}. \quad (15)$$

To see this, note that in this case

$$\alpha(s, q) = \begin{cases} 1 + (m-1)/2 & \text{if } s = q \\ 1/2 & \text{otherwise} \end{cases}$$

Without loss of generality, suppose that $s_1 < s_2$. Then,

$$\begin{aligned} \frac{1}{m^2} \sum_{q=1}^m (\alpha(s_1, q) - \alpha(s_2, q))^2 &= \frac{1}{m^2} \left((\alpha(s_1, s_1) - \alpha(s_2, s_1))^2 + (\alpha(s_1, s_2) - \alpha(s_2, s_2))^2 \right) \\ &= \frac{1}{m^2} \left(\frac{m^2}{4} + \frac{m^2}{4} \right) = \frac{1}{2} \end{aligned}$$

and

$$\begin{aligned} \frac{1}{m^2} \sum_{q < l} (\alpha(s_1, q) - \alpha(s_2, q)) (\alpha(s_2, l) - \alpha(s_1, l)) \rho_{q,l} \\ &= \frac{1}{m^2} \underbrace{\sum_{q=1}^{s_1-1} (\alpha(s_1, q) - \alpha(s_2, q)) (\alpha(s_2, s_1) - \alpha(s_1, s_1)) \rho_{q,s_1}}_{l=s_1} \\ &\quad + \frac{1}{m^2} \underbrace{\sum_{q=1}^{s_2-1} (\alpha(s_1, q) - \alpha(s_2, q)) (\alpha(s_2, s_2) - \alpha(s_1, s_2)) \rho_{q,s_2}}_{l=s_2} \\ &= \frac{1}{m^2} \sum_{q=1}^{s_1-1} \underbrace{(\alpha(s_1, q) - \alpha(s_2, q))}_{:=0} \left(\frac{-m}{2} \right) \rho_{q,s_1} \\ &\quad + \frac{1}{m^2} \sum_{q=1}^{s_2-1} \underbrace{(\alpha(s_1, q) - \alpha(s_2, q))}_{\neq 0 \text{ only when } q=s_1} \left(\frac{m}{2} \right) \rho_{q,s_2} \\ &= \frac{1}{m^2} \left(\frac{m}{2} \right) \left(\frac{m}{2} \right) \rho_{s_1, s_2} \\ &= \frac{1}{4} \rho_{s_1, s_2}. \end{aligned}$$

Plugging these values into Eq. 14 yields the desired result.

Example 2 continued. In the total average case, we can show (similar to the classical case) that

$$\rho(s_1, s_2) = \underbrace{1 - \frac{1}{m^2}}_{\text{method-induced correlation}} + \underbrace{\frac{1}{m^2} \rho_{s_1, s_2}}_{\text{model-inherent correlation}}.$$

This is sensible, because for large m , we have that $\mathfrak{M}_{\bar{A}} \approx J_m \otimes \bar{A}$ and the embedding of each $A^{(k)}$ is effectively equivalent to repeatedly embedding \bar{A} , yielding the large m correlation approximately equal to 1. Note also that when $m = 2$, this correlation coincides with the pairwise (unweighted) average classical OMNI setting (i.e., the matrix structure as in Definition 8). However, the covariance matrix $\tilde{\Sigma}$ depends on m in the $\mathfrak{M}_{\bar{A}}$ setting, and for $m \geq 3$ the limiting embedded correlation for $\mathfrak{M}_{\bar{A}}$ is always greater than the limiting embedded correlation for the classical M .

Example 3 continued. In the weighted average OMNI setting, the general form of the limiting correlation is computationally unwieldy. However, when $w_{s_1} = w_{s_2} = w$ and $w_k = 1$, for $k \neq s_1, s_2$, a simple computation yields that the limiting correlation for the (scaled) row difference in Theorem 16 is given by (where, wlog, $s_1 < s_2$)

$$\begin{aligned} \rho(s_1, s_2) &= 1 - \underbrace{\frac{((m-1)w+1)^2}{m^2(1+w)^2}}_{\text{method correlation}} + \underbrace{\frac{((m-1)w+1)^2}{m^2(1+w)^2} \rho_{s_1, s_2}}_{\text{model correlation}} \\ &= 1 - \frac{((m-1)w+1)^2}{m^2(1+w)^2} (1 - \rho_{s_1, s_2}). \end{aligned}$$

Figure 5 further highlights this relationship between the matrix structure of \mathfrak{M}_W and the limiting covariance for the weighted pairwise average case. In this example, we set the number of graphs to be $m = 10$. The plot illustrates the correlation $\rho(s_1, s_2)$ between estimates of the same true latent position across different values of the weights $w_{s_1} = w_{s_2} = w$, where $w_k = 1$ for all $k \neq s_1, s_2$. The different line types correspond to different values of ρ_{s_1, s_2} , with the vertical red line at $w = 1$ corresponding to the classical OMNI setting of Levin et al. (2017). We see that as the weight w increases, the correlation $\rho(s_1, s_2)$ decreases (i.e., in the limit, the embedded estimates derived from $A^{(s_1)}$ and $A^{(s_2)}$ are “less” dependent on each other). On the contrary, as the weight w decreases, the correlation $\rho(s_1, s_2)$ increases (towards 1) as the $A^{(k)}$ for $k \neq s_1, s_2$ has the same outsized influence on the embeddings of $A^{(s_1)}$ and $A^{(s_2)}$ which has the effect of making the embeddings effectively identical.

Indeed, the driving force behind the limiting correlation between the embedded $A^{(s_1)}$ and $A^{(s_2)}$ is the cumulative weights of the other $A^{(k)}$ for $k \neq s_1, s_2$. If these weights are large, then the relative contribution to block-row s_1 (resp., s_2) by $A^{(s_1)}$ (resp., $A^{(s_2)}$) is low when compared to the cumulative effect of the other networks, and this masks the signal corresponding to $A^{(s_1)}$ (resp., $A^{(s_2)}$) in the s_1 -th (resp., s_2 -th) block of the embedding. If, however, the weights of the $A^{(k)}$ for $k \neq s_1, s_2$ are small, then the relative contribution to block-row s_1 (resp., s_2) by $A^{(s_1)}$ (resp., $A^{(s_2)}$) is high when compared to the cumulative effect of the other networks, and this has the effect of reducing the correlation between the embeddings of $A^{(s_1)}$ and $A^{(s_2)}$.

5. Dampened Omnibus Embedding and Aplysia Spike Train Analysis

In the example in Section 3.2, we see how the uniformity of the induced correlation in the classical omnibus embedding effectively masks much of the biologically relevant signal in the Aplysia motor program. There is a dramatic spike in the data signal (the stimulus) followed by structured behavior and decay in spike intensity. In this case (and in many other time series settings where we are modeling the impact of anomalous events), it is natural to assume that the dependence between two networks at different times t_1, t_2 decays as their difference $|t_1 - t_2|$ gets bigger, and also decays over time as t_1, t_2 get bigger. This motivates




Figure 5: The plot illustrates the correlation $\rho(s_1, s_2)$ between estimates of the same true latent position across different values of the weights $w_{s_1} = w_{s_2} = w$, where $w_k = 1$ for all $k \neq s_1, s_2$. The number of graphs is $m = 10$. The different line types correspond to different values of ρ_{s_1, s_2} , with the vertical red line at $w = 1$ corresponding to the classical OMNI setting of Levin et al. (2017).

the *dampened* omnibus matrix defined via

$$\mathfrak{M}_{damp}^{(k, \ell)} = \begin{cases} \frac{w_k A^{(k)} + A^{(\ell)}}{w_k + 1} & \text{if } k > \ell, \\ A^{(k)} & \text{if } k = \ell, \\ \frac{A^{(k)} + w_\ell A^{(\ell)}}{w_\ell + 1} & \text{if } k < \ell, \end{cases}$$

where \vec{w} is a strictly increasing vector of weights. In the dampened omnibus matrix, \mathfrak{M}_{damp} , the relative contribution of each successive matrix (i.e., as $l > k$ increases in $A^{(l)}$) increases in block-row k ; this (keeping the form of $\rho(t_1, t_2)$ in mind) can have the effect of decreasing the limiting induced correlation between $A^{(k)}$ and $A^{(l)}$ as $l > k$ increases. As an example of this, consider the case where $w_\ell = \ell$ for all $\ell \in [m]$. The form of (total) correlation, though cumbersome, is nonetheless instructive. Letting $s_1, s_2 \in [m]$, $s_1 < s_2$, we have that the induced correlation is

$$\begin{aligned} \rho_{me}(s_1, s_2) = & 1 - \frac{1}{2m^2} \left[\frac{(s_1 - 1)(s_2 - s_1)^2}{(s_1 + 1)^2(s_2 + 1)^2} + \left(\frac{s_1^2 + 1}{s_1 + 1} + \sum_{l > s_1, l \neq s_2} \frac{1}{l + 1} \right)^2 \right. \\ & \left. + \sum_{s_1 < l < s_2} \left(\frac{l}{l + 1} - \frac{1}{s_2 + 1} \right)^2 + \left(\frac{(s_2 - 1)s_2 + 1}{s_2 + 1} + \sum_{l > s_2} \frac{1}{l + 1} \right)^2 \right] \end{aligned}$$




Figure 6: In a JRDPG_{for} model where the entries of $\boldsymbol{\rho}$ are identically set to 0.8, we consider the induced correlation for the classical omnibus embedding (blue), the dampened omnibus embedding (red). Considering $m = 200$ networks, the solid (resp, dashed, dotted, dot-dash) lines correspond to the limiting induced correlation between vertices of the embedded $A^{(1)}$ (resp., $A^{(100)}$, $A^{(150)}$, $A^{(175)}$) and $A^{(s_2)}$ as s_2 varies from 2 (resp., 101, 151, 176) to 200. For each curve type, the true inherent correlation across the networks is plotted in green.

while the effect of the inherent correlation is

$$\begin{aligned}
 \rho_{\text{mo}}(s_1, s_2) = & \frac{1}{m^2} \left[\sum_{l < s_1} \sum_{q < l} \frac{(s_2 - s_1)(s_1 - s_2)}{(s_1 + 1)^2 (s_2 + 1)^2} \rho_{q,l} \right. \\
 & + \sum_{q < s_1} \frac{-(s_2 - s_1)}{(s_1 + 1)(s_2 + 1)} \left(\frac{s_1^2 + 1}{s_1 + 1} + \sum_{h > s_1, h \neq s_2} \frac{1}{h + 1} \right) \rho_{q,s_1} \\
 & + \sum_{s_1 < l < s_2} \left(\sum_{q < s_1} \frac{s_2 - s_1}{(s_1 + 1)(s_2 + 1)} \left(\frac{1}{s_2 + 1} - \frac{l}{l + 1} \right) \rho_{q,l} \right. \\
 & \quad \left. + \left(\frac{s_1^2 + 1}{s_1 + 1} + \sum_{h > s_1, h \neq s_2} \frac{1}{h + 1} \right) \left(\frac{1}{s_2 + 1} - \frac{l}{l + 1} \right) \rho_{s_1,l} \right. \\
 & \quad \left. + \sum_{s_1 < q < l} \left(\frac{q}{q + 1} - \frac{1}{s_2 + 1} \right) \left(\frac{1}{s_2 + 1} - \frac{l}{l + 1} \right) \rho_{q,l} \right) \\
 & + \sum_{q < s_1} \frac{s_2 - s_1}{(s_1 + 1)(s_2 + 1)} \left(\frac{s_2^2 - s_2 + 1}{s_2 + 1} + \sum_{h > s_2} \frac{1}{h + 1} \right) \rho_{q,s_2} \\
 & + \left(\frac{s_1^2 + 1}{s_1 + 1} + \sum_{l > s_1, l \neq s_2} \frac{1}{l + 1} \right) \left(\frac{s_2^2 - s_2 + 1}{s_2 + 1} + \sum_{h > s_2} \frac{1}{h + 1} \right) \rho_{s_1,s_2} \\
 & \left. + \sum_{s_1 < q < s_2} \left(\frac{q}{q + 1} - \frac{1}{s_2 + 1} \right) \left(\frac{s_2^2 - s_2 + 1}{s_2 + 1} + \sum_{h > s_2} \frac{1}{h + 1} \right) \rho_{q,s_2} \right].
 \end{aligned}$$

With the form of ρ_{me} above, consider the situation in which $s_2 \gg s_1$. In this case, ρ_{me} is growing like $\approx 1 - s_2^2/2m^2$, which is decaying to $1/2$ as s_2 approaches m . When both $s_2, s_1 \gg 1$, ρ_{me} is growing like $\approx 1 - (s_2^2 + s_1^2)/2m^2$, which is decaying to 0 as s_2, s_1 both approach m . This corroborates the effect we see in Figure 6, where $\rho(s_1, s_2)$ is tracking the inherent (i.e., model) correlation fairly well after an initial burn in period (i.e., when s_1 is sufficiently large). For ρ_{mo} (ignoring for the moment the potential large differences in the inherent correlation), when $s_2, s_1 \gg (s_2 - s_1)$, the dominant term is $\approx \frac{1}{m^2} s_1 s_2 \rho_{s_1, s_2}$ which is tracking the true inherent correlation. When both $s_2, s_1 \gg 1$ but their difference is not necessarily relatively small, the dominant term in ρ_{mo} is of the order

$$\approx \frac{1}{m^2} s_1 s_2 \rho_{s_1, s_2} + \frac{1}{m^2} s_2 \sum_{s_1 < q < s_2} \rho_{q, s_2} .$$

We see that this only depends on ρ_{q, s_2} for $q \in \{s_1 + 1, \dots, s_2 - 1\}$ and the effect of the inherent correlation is effectively localized.

In Figure 6, we consider the JRDPG_{for} model of Section 2.2.1 where the entries of $\boldsymbol{\rho}$ are identically set to 0.8 and $m = 200$. We plot the induced correlation for the classical omnibus embedding (blue), the dampened omnibus embedding (red), where the solid (resp, dashed, dotted, dot-dash) lines correspond to the limiting induced correlation between vertices of the embedded $A^{(1)}$ (resp., $A^{(100)}$, $A^{(150)}$, $A^{(175)}$) and $A^{(s_2)}$ as s_2 varies from 2 (resp., 101 , 151 , 176) to 200 . For each curve type, the true inherent correlation across the networks is plotted in green. While the method correlation in classical OMNI prevents the induced correlation from ever tracking the true inherent correlation, we see that the more nuanced structure presented by dampened OMNI allows (after a suitable burn in period) the induced correlation to significantly better (though far from ideally) track the inherent correlation. Developing \mathfrak{M} to induce a given correlation structure is a natural next step, and we are actively pursuing this question. This decaying correlation structure appears well-suited for the Aplysia motor program, as it precisely allows for a large correlation at early time-points (corresponding to the dampening of signal after an anomalous event) that dissipates in time. We will explore this further in the next section.

5.1 Dampened Spike Train Analysis

Since the generalized omnibus embedding permits additional degrees of freedom for the off-diagonal entries, a richer spectrum of induced correlation is possible. The dampened omnibus embedding, in particular, is designed to weaken cross-graph interplay over time. Given that the correlation homogeneity of the classical omnibus makes inter-network changes less apparent, it is natural to ask if an appropriately-calibrated dampened omnibus embedding can detect what the more limited classical omnibus can not. To test this, we next apply the methodology of Section 3.2.1 to the Aplysia data of Section 3.2 using $\mathfrak{M}_{\text{damp}}$ with $w_\ell = \ell$. Results are summarized in Figure 7 where \mathbf{D}_d is the corresponding distance matrix from Eq. 10 in the dampened setting.

As in Section 3.2, we compute the matrix \mathbf{D}_d of distances between the embeddings of the 24 networks (using the dampened omnibus embedding), and next embed \mathbf{D}_d into \mathbb{R}^2 using CMDS. This yields a 2-dimensional point for each of the 24 networks in the motor program; we further cluster these points using Mclust in order to discover the distinct biological

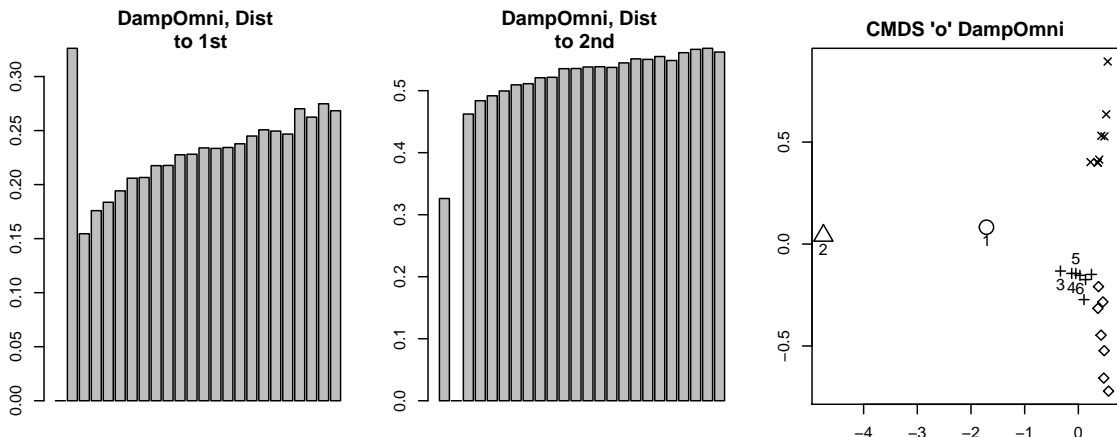


Figure 7: Using $w_\ell = \ell$ in the Dampened OMNI framework, in the left (resp., center) panel, we plot the average vertex distance (as the bar heights) in the embedding between graph 1 (resp., graph 2) and graph k for each $k \in [m]$ in Dampened OMNI (where the bars are labeled $k = 1, 2, \dots, 24$). In the right panel, we compute the 24×24 distance matrices \mathbf{D}_d , and embed it into \mathbb{R}^2 . The resulting 24 data points are clustered using `Mclust` (clusters denoted by shape), and are plotted, with graphs 1 through 6 further labeled with their corresponding number.

Graph position in time-series	1	2	3	4	5	6	7	8	9	10	11	12
Cluster assignment	1	2	3	3	3	3	3	3	3	4	4	4
Graph position in time-series	13	14	15	16	17	18	19	20	21	22	23	24
Cluster assignment	4	5	5	4	4	5	4	5	5	4	5	5

Table 2: This table displays the cluster labels for the 24 networks when embedding the time-series using dampened OMNI with weights $w_\ell = \ell$, then embedding the across-graph distance matrix using classical MDS scaling, and finally clustering the graphs using `Mclust`. In the MDS embedding, each graph is represented by a single 2-dimensional embedded data point, and the pair (Graph, Cluster) describes which cluster these points are assigned by `Mclust`.

phases of activity. The cluster labels/assignments for each network are presented in Table 2. As in the classical OMNI setting, we see from the table that dampened OMNI again isolates the stimulus in the second graph. However, unlike in the classical OMNI setting, dampened OMNI reveals additional biologically relevant structure including the uniqueness of Graph 1 (it is clustered in a singleton cluster), the transition from gallop to crawl, and an unstable dynamic in the crawling motor program.

From Table 2 and the two-dimensional embedding (Figure 7), we further see that the dampened clustering better distinguishes Graph 1 (the spontaneous firing state graph) from the rest of the motor program. Moreover, the clustering together of Graphs 3–9 and the CMDS graph embedding show that the dampened setting is able to capture (imperfectly) the transition from galloping in bins 3 and 4 to crawling in bins 5–24. The clustering on the

crawling phase further yields classes that progress in orderly fashion over the first half of the motor program, after which states irregularly alternate between 4 and 5, suggesting an unstable dynamic to the program not apparent from simple visual inspection of the firing traces.

Indeed, from Figure 3, we see that the flat correlation induced by classical OMNI masks much of this latent structure in these biological networks, while dampened omnibus induces correlation that both uncovers and clarifies this neuroscientifically relevant signal in the data. This reinforces the fact that the choice of optimal weights for the off-diagonal is *inference-task dependent* and vitally important; different weights can serve to isolate and distill different signal, both within and across graphs. The question of optimizing the weight selection for a given inference task—whether identification of specific vertices of interest, temporal trends across graphs, or change points in a network time series—is a rich object of study in its own right.

6. Induced Correlation in JRDPG_{gen} Model

In this section, we present a series of simulations that are designed further illuminate the flexibility of the generalized omnibus framework in the JRDPG_{gen} model (having explored dampened OMNI in the JRDPG_{for} model in Section 5). If we consider a collection of graphs as noisy realizations of some true, underlying, network process (i.e., the JRDPG_{gen} model), then it is reasonable to expect some of the realizations to be noisier, or of lower fidelity, than others. As such, it is desirable to have an embedding method that can account for the difference in sampling fidelity across networks. To explore this further, we consider the JRDPG_{gen} model with $\nu_i = 0.8$ for $i = 1, 2, \dots, 50$, and $\nu_i = 0.3$ for $i = 51, 52, \dots, 100$ (so that $m = 100$). In this case, the first 50 $A^{(i)}$ are higher fidelity (i.e., less noisy) copies of the generator $A^{(0)}$. It is natural then to seek to down-weight $A^{(i)}$ for $i = 51, 52, \dots, 100$ in the embedding. We can achieve this via the weighted average omnibus embedding with

$$\mathfrak{M}_W^{(i,j)} = \frac{w_i A^{(i)} + w_j A^{(j)}}{w_i + w_j}$$

where \vec{w} is a vector of positive graph weights. In Table 3, we consider the limiting embedding correlation across the networks within this model for a variety of different weight vectors. We see that upweighting the higher-fidelity networks in \mathfrak{M}_W has the effect of more closely modeling the true inherent correlation among the high fidelity networks in the embedded space versus classical OMNI (the $w_i = 1$ for all i setting). This comes at the expense of inducing more correlation between the lower fidelity samples (and across the lower fidelity and higher fidelity samples), as the higher fidelity samples are having more influence on the embeddings of the lower fidelity samples. We note also that the effect is reversed if the lower fidelity samples are upweighted.

In the model above, suppose that we want the embedding to preserve the correlation between one particular pair of graphs (wlog, between $A^{(1)}$ and $A^{(2)}$). This can be achieved via a special OMNI construction as follows. Letting m be even, for i odd define

$$\mathfrak{M}_{12}^{(i,j)} = \begin{cases} A^{(i)} & \text{if } i = j \text{ or } i \leq j - 2 \\ \frac{A^{(i)} + A^{(j)}}{2} & \text{if } i = j - 1 \\ A^{(j)} & \text{if } i > j, \end{cases}$$

	inherent corr.	$w_i \stackrel{iid}{\sim} \text{Unif}(0,1)$	$w = \vec{1}_{100}$	$w = [\vec{1}_{50}, \vec{10}_{50}]$	$w = [\vec{10}_{50}, \vec{1}_{50}]$
$\rho(1, 2)$	0.64	0.933	0.91	0.969	0.821
$\rho(1, 51)$	0.24	0.823	0.81	0.729	0.839
$\rho(51, 52)$	0.09	0.684	0.773	0.548	0.921

Table 3: Considering the $\text{JRDPG}_{\text{gen}}$ model with $\nu_i = 0.8$ for $i = 1, 2, \dots, 50$, and $\nu_i = 0.3$ for $i = 51, 52, \dots, 100$ (so that $m = 100$), we provide the correlations induced by embedding \mathfrak{M}_W for a variety of weight vectors \vec{w} . Values are rounded to three digits for ease of display.




Figure 8: Considering the $\text{JRDPG}_{\text{gen}}$ model with $\nu_i = 0.8$ for $i = 1, 2, \dots, 50$, and $\nu_i = 0.3$ for $i = 51, 52, \dots, 100$ (so that $m = 100$), we plot the difference between induced and inherent correlations when embedding via \mathfrak{M}_{12} . Darker blue indicates less accurate preservation of the correlation structure.

for i even define

$$\mathfrak{M}_{12}^{(i,j)} = \begin{cases} A^{(i)} & \text{if } i = j \text{ or } i \leq j - 1 \\ \frac{A^{(i)} + A^{(j)}}{2} & \text{if } i = j + 1 \\ A^{(j)} & \text{if } i - 1 > j. \end{cases}$$

Then it is not difficult to compute

$$\rho(1, 2) = 1 - (1 - R_{1,2}) \frac{(m-1)^2}{m^2},$$

so that $\rho(1, 2) \approx R_{1,2}$ for large m . While this choice of \mathfrak{M}_{12} may not globally preserve R in the embedded space, this demonstrates that local (i.e., between pairs or a few pairs) can be well-preserved. In Figure 8, we see that (considering again the $\text{JRDPG}_{\text{gen}}$ model with $\nu_i = 0.8$ for $i = 1, 2, \dots, 50$, and $\nu_i = 0.3$ for $i = 51, 52, \dots, 100$), the pairwise correlation is well preserved by embedding via \mathfrak{M}_{12} , not just between $A^{(1)}$ and $A^{(2)}$ but between many of the network pairs $A^{(i)}$ and $A^{(j)}$ for i, j small.

7. Effective Sample Size in Correlated Graphs

Statistical inference for multiple networks often faces questions that require to aggregate the information about the latent positions across samples of graphs. When the graphs are correlated, correlation itself adds a dampening factor that can reduce the effective sample size. In this subsection, we investigate the effect of correlation (induced and inherent) on effective sample size in subsequent inference about the latent position.

The following theorem describes the limiting behavior of the average latent positions estimated from a sample of correlated graphs. The proof is given in the Appendix.

Theorem 17 *Let $\rho \in (0, 1)$ be fixed. Let F be a distribution on a set $\mathcal{X} \subset \mathbb{R}^d$, where $\langle x, x' \rangle \in [0, 1]$ for all $x, x' \in \mathcal{X}$, and assume that $\Delta := \mathbb{E}[X_1 X_1^T]$ is rank d . Let $(A_n^{(1)}, \dots, A_n^{(m)}, \mathbf{X}_n) \sim \text{JRDPG}(F, n, m, \rho)$, be a sequence of ρ -correlated RDPG random graphs and associated latent positions.*

- a) *Letting $\widehat{\mathbf{X}}_{A_n}^{(s)} = \text{ASE}(A_n^{(s)}, d)$ for each $s = 1, \dots, m$, there exist sequences of orthogonal d -by- d matrices $(W_n^{(s)})_{n=1}^\infty$, $s = 1, \dots, m$ such that for all $z \in \mathbb{R}^d$ and for any fixed index i ,*

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[n^{1/2} \left(\frac{1}{m} \sum_{s=1}^m \widehat{\mathbf{X}}_{A_n}^{(s)} W_n^{(s)} - \mathbf{X}_n \right)_i \leq z \right] = \int_{\text{supp } F} \Phi \left(z, \widetilde{\Sigma} \left(x, \frac{1-\rho}{m} + \rho \right) \right) dF(x). \quad (16)$$

- b) *Letting $\widehat{\mathbf{X}}_{M_n} = \text{ASE}(M_n, d)$, and denoting the s -th $n \times d$ block of $\widehat{\mathbf{X}}_{M_n}$ as $\widehat{\mathbf{X}}_{M_n}^{(s)}$, there exist a sequence of orthogonal d -by- d matrices $(\widetilde{W}_n)_{n=1}^\infty$, such that for all $z \in \mathbb{R}^d$ and for any fixed index i ,*

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[n^{1/2} \left(\frac{1}{m} \sum_{s=1}^m \widehat{\mathbf{X}}_{M_n}^{(s)} \widetilde{W}_n - \mathbf{X}_n \right)_i \leq z \right] = \int_{\text{supp } F} \Phi \left(z, \widetilde{\Sigma} \left(x, \frac{1-\rho}{m} + \rho \right) \right) dF(x). \quad (17)$$

Stated simply, the previous theorem recovers the classical result about the covariance of the sample mean for correlated data. When the graphs are independent, the average of the estimated latent positions after a proper orthogonal alignment adds a factor of $1/m$ to the limiting covariance matrix, suggesting that this average concentrates for large sample size around the true mean. This is shown both for a separate embedding of each graph (part a)) or a joint embedding with classical OMNI (part b)). When the edges of the graphs are correlated by some positive constant ρ , the efficiency of this estimator is reduced, and the limiting effective sample size is $m_{\text{eff}} := \frac{m}{1+\rho(m-1)}$.

7.1 Recovery of Community Labels in Stochastic Blockmodels

We start by investigating the recovery of community labels from estimated latent positions of stochastic block model graphs. We note that the literature abounds with clustering methods for networks, including several that employ joint network embedding (see, for

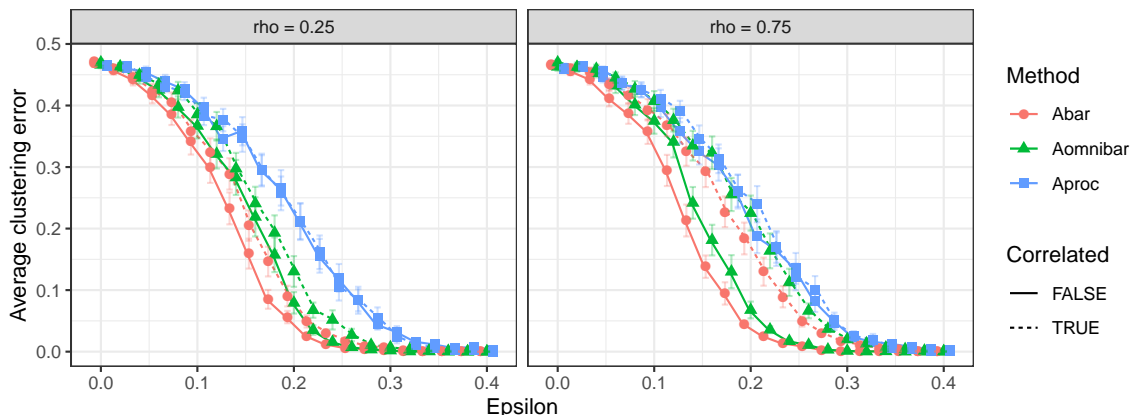


Figure 9: Community detection error in a ρ -SBM model with $K = 2$ communities, $N = 100$ number of vertices and block probability matrix P_ϵ as a function of ϵ . The latent positions are estimated by ASE embedding of the mean graph \bar{A} (red lines), OMNI embedding (green lines) and Procrustes-based pairwise embedding (blue lines). Each point is the mean of 100 MC replicates. Observe that in both panels the ASE(\bar{A} , 2) and OMNI(M , 2) embeddings in the i.i.d case perform better than their embeddings in the correlated case, while the Procrustes-based pairwise embedding remains unaffected by the effect of correlation. Also note that the performance of the Procrustes-based pairwise embedding is inferior to the other two embedding methods.

example, the methods compared in Figure 14); for a comparison of the OMNI method to existing methods in the literature, see for example Gallagher et al. (2021); Amini et al. (2019). Our focus here is primarily on understanding the effect of the induced correlation on effective sample size for this clustering inference task. In this light, the OMNI methodology is, to our knowledge, *unique* in that the explicit level of embedding correlation can be computed and the latent versus induced portions of this correlation cleanly delineated. Other joint embedding methods (see Figure 14) also induce correlation in the embedded space, and it is an open problem to derive analogues to Theorem 16 in these settings.

In addition to exploring the effect of the induced and latent correlation on effective sample size for this clustering inference task, the results of this section further illustrate the flexibility and applicability of the OMNI framework. When conducting joint estimation for latent positions across a collection of networks *with the same underlying latent positions* (as below), it is natural to consider a spectral embedding of \bar{A} , the sample mean adjacency matrix. It turns out that the regularization induced by this embedding produces a markedly accurate estimator of the latent positions and, in turn, of block model communities (Tang et al., 2018). However, because this estimator is a function of \bar{A} , which itself homogenizes differences across the networks, it is not especially robust to model misspecification, and its accuracy degrades if the underlying latent positions differ across networks. This renders it less suitable for embedding non-stationary time-series of networks, for example. At the other extreme, rather than embedding the averaged graph \bar{A} , we can consider individual embeddings of the individual adjacency matrices, which we then subject to a Procrustes alignment and posthoc averaging. This can effectively distinguish certain differences across

network latent positions, but when the latent positions are identical, the Procrustes estimate exhibits poorer performance relative to a spectral embedding of the sample mean \bar{A} (see Figures 9–11). This section demonstrates that the omnibus embedding balances the inferential strengths of these two extremes: it maintains the performance of the embedding of \bar{A} when the latent positions are identical (see Figures 9–11), while retaining (and improving upon!) the ability of the Procrustes posthoc aligned embeddings to distinguish different structure across networks (see Section 5.1).

We now consider a correlated pair of positive semidefinite K -block stochastic block models (Holland et al., 1983) as in Athreya et al. (2018a), Definition 8. In particular, letting $K = 2$, we generate a pair of adjacency matrices $(A, B, \mathbf{X}) \sim \text{JRDPG}(F, n, 2, \rho)$, where F is a mixture of point mass distributions defined by

$$F = \frac{1}{2}\delta_{\zeta_1} + \frac{1}{2}\delta_{\zeta_2}, \tag{18}$$

in which $\zeta_1, \zeta_2 \in \mathbb{R}^2$ are the latent positions and satisfy

$$\begin{bmatrix} \zeta_1 \\ \zeta_2 \end{bmatrix} \begin{bmatrix} \zeta_1 \\ \zeta_2 \end{bmatrix}^T = \begin{bmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 + \epsilon \end{bmatrix} =: P_\epsilon$$

where $\epsilon > 0$.

Observe that the stochastic block models all have the same underlying latent positions. We let $K = d = 2$ and consider the following three graph embedding techniques to help us understand the role of effective sample size on community detection in correlated random graphs:

1. **Classical OMNI embedding:** We apply the traditional omnibus embedding using M as in Equation (3) to obtain estimated latent positions $\text{ASE}(M; 2) \in \mathbb{R}^{2n \times 2}$ and we average the n rows across the two graphs to obtain $\hat{\mathbf{X}}^1 \in \mathbb{R}^{n \times 2}$.
2. **Mean graph embedding:** We use ASE to embed the mean of the two graphs, $\bar{A} = \frac{A+B}{2}$ (Tang et al., 2018) to obtain estimated latent positions $\hat{\mathbf{X}}^2 = \text{ASE}(\bar{A}; 2) \in \mathbb{R}^{n \times 2}$.
3. **Procrustes-based embedding:** We separately embed the graphs A, B , obtaining two separate matrices of estimated latent positions $\text{ASE}(A, 2), \text{ASE}(B, 2) \in \mathbb{R}^{n \times 2}$. We then align these two matrices via orthogonal Procrustes alignment (Gower, 1975), and average the aligned embeddings to obtain $\hat{\mathbf{X}}^3 \in \mathbb{R}^{n \times 2}$.

For each embedding strategy, we further consider $\rho = 0$ (i.e., conditionally independent graphs) and $\rho > 0$ (i.e., correlated graphs) in our generative model resulting in a total of six estimated latent position matrices, $\hat{\mathbf{X}}_{\text{ind}}^1, \hat{\mathbf{X}}_{\text{corr}}^1, \hat{\mathbf{X}}_{\text{ind}}^2, \hat{\mathbf{X}}_{\text{corr}}^2, \hat{\mathbf{X}}_{\text{ind}}^3, \hat{\mathbf{X}}_{\text{corr}}^3$.

In Figures 9-11, we cluster the rows of estimated latent positions $\hat{\mathbf{X}}_\ell^\kappa$, for all $\kappa \in \{1, 2, 3\}$, $\ell \in \{\text{ind}, \text{corr}\}$ into $K = 2$ communities via the model based clustering procedure in the `Mclust` R package (Fraley and Raftery, 2002). It is shown in Tang and Priebe (2016), Section 4, that Gaussian mixture model-based clustering yields substantial improvement over K -means clustering in recovering the latent communities in ASE of stochastic blockmodels. This is due to the associated limiting covariance matrices often being elliptical rather than spherical. In each figure, we plot the clustering error in recovering the latent community




Figure 10: Clustering error in recovery of community labels of estimated latent positions in a ρ -SBM model with $K = 2$ communities, $N = 100$ number of vertices and block probability matrix P_ϵ as a function of the correlation ρ . The latent positions are estimated by ASE embedding of the mean graph \bar{A} (red lines), OMNI embedding (green lines) and Procrustes-based pairwise embedding (blue lines). Each point is the mean of 100 MC replicates. In panel (a), all the methods effectively achieve the same clustering error. In panel (b), the performance of the ASE($\bar{A}, 2$) and OMNI($M, 2$) embeddings in the correlated case decrease as the correlation increases.

labels for all six estimated latent position matrices where the solid lines correspond to the independent case and the dashed lines correspond to the correlated case. The latent positions estimated by ASE($\bar{A}, 2$) are displayed in incarnadine/red lines, by OMNI embedding in green lines and by Procrustes-based pairwise embedding in blue lines. Each point is the average of 100 MC replicates.

In Figure 9, we set the number of vertices for each graph equal to $N = 100$ and we plot the clustering error versus ϵ . In panel (a) (resp. (b)) we set the correlation for the estimated latent positions in the correlated case equal to $\rho = 0.25$ (resp. $\rho = 0.75$). Observe that in both panels the ASE($\bar{A}, 2$) and OMNI($M, 2$) embeddings in the i.i.d case perform better than their corresponding embeddings in the correlated case, while the Procrustes-based pairwise embedding remains unaffected by the inherent correlation; we suspect this is due to the Procrustes step introducing additional signal across the networks which can mask the present inherent correlation. Also note that the performance of the Procrustes-based pairwise embedding is inferior to the other two embedding methods, while the omnibus and mean embeddings have comparable performance. In Figure 10, we again set $N = 100$ and we plot the clustering error versus the correlation. In panel (a) (resp. (b)) we set $\epsilon = 0.1$ (resp. $\epsilon = 0.2$) across all estimated latent positions in both cases. In panel (a), the signal in distinguishing the two communities is weak and all embedding methods achieve similar performance across ρ . In panel (b) however, the performance of clustering the estimated latent positions $\hat{\mathbf{X}}_{\text{corr}}^1$ and $\hat{\mathbf{X}}_{\text{corr}}^2$ decreases as the correlation increases, as the effective sample size of the correlated networks is decreasing.

In Figure 11, we plot the clustering error versus the number of vertices of the graphs. The rows of the figure correspond to different values of $\epsilon \in \{0.1, 0.2\}$ and the columns




Figure 11: Clustering error in recovery of community labels of estimated latent positions in a ρ -SBM model with $K = 2$ communities and block probability matrix P_ϵ as a function of the number of vertices N . The latent positions are estimated by ASE embedding of the mean graph \bar{A} (red lines), OMNI embedding (green lines) and Procrustes-based pairwise embedding (blue lines). Each point is the mean of 100 MC replicates.

correspond to different values of $\rho \in \{0.25, 0.75\}$. Comparing panels (a) and (b) we observe how the effective sample size is diminished for the estimated latent positions $\hat{\mathbf{X}}_{\text{corr}}^1, \hat{\mathbf{X}}_{\text{corr}}^2$ as the correlation increases (i.e., more vertices are needed to achieve the same error in the highly correlated case). For example, in panel (a) ($\rho = 0.25$), the clustering error for both $\hat{\mathbf{X}}_{\text{corr}}^1$ and $\hat{\mathbf{X}}_{\text{corr}}^2$ is approximately 0.2 when the number of vertices is $N \approx 250$, whereas in panel (b) ($\rho = 0.75$) the same error is achieved when $N \approx 350$. This is mirrored in panels (c) and (d) as well.

7.2 Vertex Classification in Brain Networks

Vertex classification is a problem that arises in applications where the goal is to predict vertex labels using a subset of the vertices for which this information is known a priori (Sussman et al., 2014; Chen et al., 2016). In brain networks, for instance, where the vertices correspond to brain regions or neurons, information about vertex attributes is sometimes known in more detail for a portion of the graph, and this can be used to infer these attributes in the remaining vertices (Chen et al., 2016). Here, we show how leveraging the information from a collection of networks yields improvements in vertex classification accuracy, but the presence of edge correlation between the networks can reduce the effective sample size and result in smaller prediction improvements.

The HNU1 study (Zuo et al., 2014) includes brain diffusion magnetic resonance images (dMRI) from 30 healthy subjects that were scanned 10 times each over a period of one month. These scans were used to construct a collection of brain networks with the CPAC200 atlas (Kiar et al., 2018), resulting in a sample of 300 graphs (one per each subject and each scan) with 200 aligned vertices, and binary edges denoting the existence of nerve tracts between each pair of brain regions. The vertices of the networks are labeled according to the brain hemisphere, with 94 vertices corresponding to the left hemisphere, 98 to the right hemisphere, and 8 unlabeled vertices. The post-processed brain networks were downloaded from <https://neurodata.io/mri/>.

To understand the effect of the correlation in subsequent inference, we start by measuring the correlation between pairs of graphs. Recall that under the ρ -correlated heterogeneous Erdős-Rényi (or Bernoulli) model, the correlation between two graphs $A^{(k)} \sim \text{Bernoulli}(P^{(k)})$ and $A^{(\ell)} \sim \text{Bernoulli}(P^{(\ell)})$ satisfies

$$\rho^{(k,\ell)} = \frac{\mathbb{E}[(A_{ij}^{(k)} - P_{ij}^{(k)})(A_{ij}^{(\ell)} - P_{ij}^{(\ell)})]}{\left[\text{Var}(A_{ij}^{(k)}) \text{Var}(A_{ij}^{(\ell)})\right]^{1/2}}, \quad \forall i, j \in [n], i \neq j.$$

Based on estimates of the edge probability matrices $\hat{P}^{(k)}$ and $\hat{P}^{(\ell)}$, we construct a plug-in estimator of the correlation given by

$$\hat{\rho}^{(k,\ell)} = \frac{2}{n(n-1)} \sum_{i < j} \frac{(A_{ij}^{(k)} - \hat{P}_{ij}^{(k)})(A_{ij}^{(\ell)} - \hat{P}_{ij}^{(\ell)})}{\left[\hat{P}_{ij}^{(k)}(1 - \hat{P}_{ij}^{(k)})\hat{P}_{ij}^{(\ell)}(1 - \hat{P}_{ij}^{(\ell)})\right]^{1/2}}.$$

The edge probability matrix is estimated as $\tilde{P}^{(k)} = \hat{\mathbf{X}}^{(k)} \hat{D}^{(k)} \hat{\mathbf{X}}^{(k)T}$, where $\hat{D}^{(k)} \in \mathbb{R}^{d \times d}$ is a diagonal matrix such that $\hat{D}_{uu}^{(k)}$ is the sign of the u -th eigenvalue of $A^{(k)}$ ordered by magnitude, $\hat{\mathbf{X}}^{(k)} = \text{ASE}(A^{(k)}, d)$, and $d = 15$ as suggested by an analysis to the same data set in Arroyo et al. (2020). After that, $\hat{P}^{(k)}$ is formed by trimming the entries to ensure that the values of the matrix are inside the interval $(0, 1)$. Given a value $\epsilon \in (0, 1/2)$, this estimator is defined as

$$\hat{P}_{ij}^{(k)} = \begin{cases} \epsilon & \text{if } \tilde{P}_{ij}^{(k)} \leq \epsilon, \\ \tilde{P}_{ij}^{(k)} & \text{if } \tilde{P}_{ij}^{(k)} \in (\epsilon, 1 - \epsilon), \\ 1 - \epsilon & \text{otherwise.} \end{cases}$$

In our experiments, we choose $\epsilon = 10^{-4}$, but we observe that the results remain quantitatively similar for a wide range of values. We also estimate the Pearson correlation between the graphs, which corresponds to the edge correlation under the correlated homogeneous Erdős-Rényi model, and it is given by

$$\hat{\varrho}^{(k,\ell)} = \frac{\sum_{i > j} (A_{ij}^{(k)} - \bar{A}^{(k)})(A_{ij}^{(\ell)} - \bar{A}^{(\ell)})}{\left[\sum_{i > j} (A_{ij}^{(k)} - \bar{A}^{(k)})^2 \sum_{i > j} (A_{ij}^{(\ell)} - \bar{A}^{(\ell)})^2\right]^{1/2}}.$$

Here, $\bar{A}^{(k)} = \frac{1}{n(n-1)} \sum_{i > j} A_{ij}^{(k)}$ is the edge sample mean.

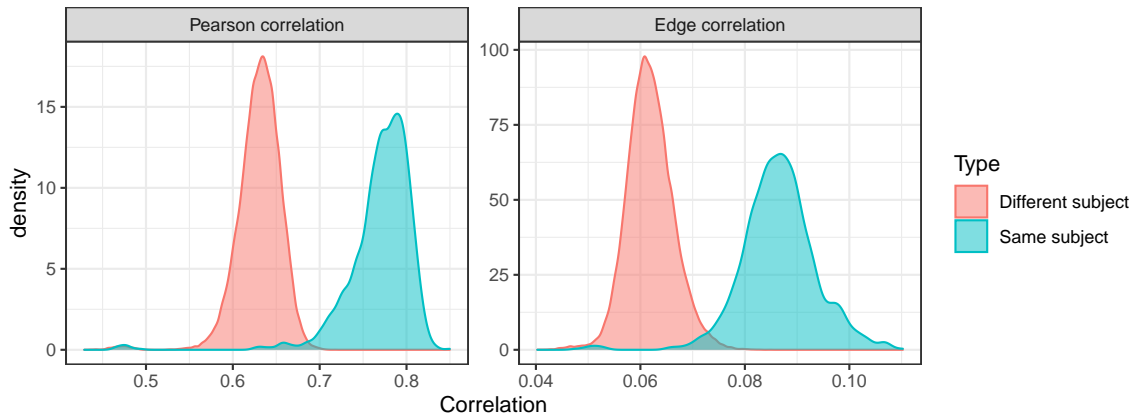


Figure 12: Empirical distribution of the correlation between pairs of graphs. The left panel shows the Pearson correlation between pairs of vectorized adjacency matrices, and the right panel shows the estimated edge correlation in the ρ -RDPG model. In both cases, the correlation between graphs corresponding to different subjects is larger than the same-subject correlation.

Figure 12 shows the empirical distribution of the estimated correlation values between pairs of networks, divided according to correlation estimates for pairs of networks corresponding to the same subject (blue curves) and pairs of networks belonging to different subjects (red curves). For both edge and Pearson correlation distributions, the red curves appear to be stochastically smaller than the blue curves, suggesting that the correlation for pairs of networks within-subject is higher than the correlation between networks of different subjects, which is expected. Thus, according to our theory, we would expect that the corresponding embeddings will show a stronger correlation, which can have consequences in subsequent inference.

We now evaluate the performance of vertex classification when a sample of m networks is observed. We focus on predicting the hemisphere labels using the information of a subset of the vertices to predict the label of the remaining ones. The classification strategy is based on constructing an unsupervised classical omnibus embedding on the m observed graphs. Note that the theory states that the correlation across the embedding will be higher for the collection of graphs with the higher inherent correlation (i.e., same subject), and this would have the effect of reducing the effective sample size. After embedding, the average of the m sets of estimated latent positions is calculated following a similar procedure to Section 7.1. The resulting embedding is used to perform 1-nearest neighbor classification for the unlabeled vertices.

We evaluate the classification performance of the omnibus embedding by randomly selecting a sample of 10% of the vertices as the training set, for which the hemisphere labels are known, and the classification accuracy is measured on the remaining 90% of the vertices with occluded hemisphere label. The accuracy is computed as the average of 250 replications over randomly selected training vertices and subjects. We compare the performance of two different choices of the training set of networks: first, a random sample of m networks from the same subject, and second, a random sample of m networks from different subjects.




Figure 13: Average vertex classification accuracy for predicting brain hemisphere. For each run, 20 random vertices were used to train the model, and the hemisphere class was predicted on the rest based on a 1-nearest neighbor classifier using the estimated latent positions of the omnibus embedding of m different graphs. The embeddings constructed from graphs corresponding to different subjects show higher gains in accuracy as the sample size increases.

Figure 13 shows the average classification accuracy (with bars corresponding to two standard errors) as a function of the number of networks used to estimate the latent positions. We observe that in general the accuracy improves as the sample size m increases, which is expected. However, the figure also shows that the gains in accuracy are much better for the classifier that uses networks from different subjects. Our theory suggests that the presence of correlation inflates the estimation error of the corresponding latent positions of the RDPG model, and thus, the degraded performance of the method is both expected and in line with our theoretical analysis.

8. Conclusion and Future Work

The generalized omnibus embedding methodology is both flexible, allowing for an array of off-diagonal weightings, and tractable, in that its component pieces can be rigorously analyzed. One such component piece, in particular, is the effect of induced and inherent correlations across the embedded networks in the omnibus framework, and in this paper we provide a first step toward understanding these correlations. The implications of this are many, both in the analysis of real data (Section 3.2), and in the setting where we desire the omnibus method to imbue identically distributed embedded data with a specific correlation structure or preserve the signal across inherently correlated networks. Furthermore, the analysis of correlated graphs in the generalized OMNI framework, enables principled, high fidelity applications in time-series (and other naturally correlated) network settings. Understanding the dual contributions of inherent and induced correlation also allows us to formulate and experiment with the network analogue of classical effective sample size analysis.

Of course, there are many significant follow-on questions and research directions. For instance, given a feasible correlation structure for a collection of networks, can we choose weights in the generalized omnibus setting that would (from i.i.d. or correlated identically distributed networks) reproduce this structure in the embedded space? In the i.i.d. setting, representing the desired correlation matrix via $\boldsymbol{\rho} = [\rho(s_1, s_2)] \in \mathbb{R}^{m \times m}$ and the matrix of generalized OMNI weights $\boldsymbol{\alpha} = [\alpha(s_1, s_2)] \in \mathbb{R}^{m \times m}$, inducing $\boldsymbol{\rho}$ in the embedding space amounts to finding feasible $\boldsymbol{\alpha}$ solving

$$\boldsymbol{\rho} = \mathbf{1} - \frac{1}{2m^2} \mathbf{D}_\alpha, \tag{19}$$

where

$$\mathbf{D}_\alpha = [\mathbf{D}_\alpha(k, \ell)] = [\|\alpha(k, \cdot) - \alpha(\ell, \cdot)\|_2^2].$$

If a solution exists, this might be easily solved via multidimensional scaling (Borg and Groenen, 2005), if not for the condition that each $\alpha(k, \cdot)$ must be nonnegative, sum to m , and be derived from symmetric $C^{(\ell)}$ matrices. While this is further complicated in the presence of inherent correlation, we are nonetheless exploring possible approximation approaches for this problem, including alternately projecting onto the cone of distance matrices (Zhang et al., 2014) and the polytope defining the constraints on the α 's. An approximate solution for this is essential, since it would allow us to use the generalized omnibus framework to produce a given correlation structure in the embedded space.

In the absence of a general solution, there are a number of \mathfrak{M} we could consider that would produce differing limiting levels of correlation amongst the embedded network pairs. For example, we could consider the forward omnibus matrix \mathfrak{M}_{for} , where

$$\mathfrak{M}_{for}^{(i,j)} = \begin{cases} \frac{(i-1)A^{(j)} + A^{(i)}}{i} & \text{if } i > j \\ A^{(i)} & \text{if } i = j \\ \frac{(j-1)A^{(i)} + A^{(j)}}{j} & \text{if } i < j \end{cases}$$

In contrast to the dampened OMNI setting, which has decaying correlation across graph pairs, in the forward model, the correlation is increasing as the graph indices are increasing, and can be used to model time series where the correlation is growing as the series progresses. While our analysis of the limiting induced correlation structure given in Eq. 14 permits for these specific, example-based constructions, a more automated approach is needed for broader applications.

As we demonstrate in this manuscript, the role of induced correlation is critical in subsequent inference. The discovery of the phenomenon of induced correlation in classical OMNI motivated the present work, and calls for further investigation into this phenomena in similar joint graph embedding methods. To shed some light on this question, we consider the same experiment as in Figure 1 with two other established joint graph embedding methods which produce estimates for each graph (along with classical OMNI); namely the Multiple Adjacency Spectral Embedding of (MASE) (Arroyo et al., 2020) and the Multiple Random Dot Product Graph (multi-RDPG) work of (Nielsen and Witten, 2018). Under the same setup as in Figure 1, we first plot the distance between the (aligned) estimates of the latent position $X_1 = \xi_1$ derived from the separate graph embeddings over a range of

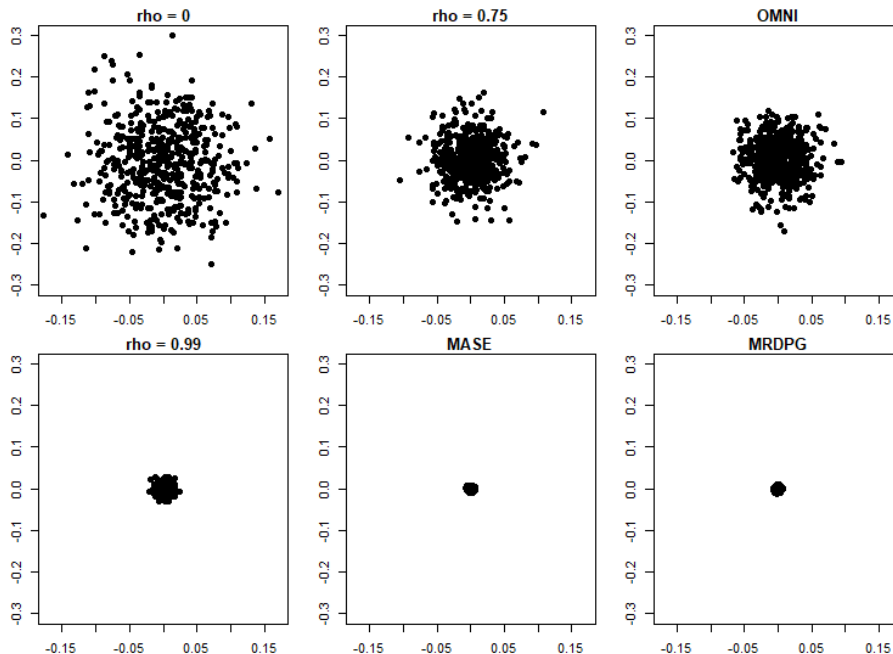


Figure 14: The effect of the induced correlation in the embedded space by the classical OMNI, MASE and multi-RDPG joint embedding methods in contrast to the inherent correlation (correlation from the graph model). The graph model ($\text{JRDPG}(F, n, m, \rho)$), the set of parameters and the number of MC iterations are the exact same as in Figure 1.

values of $\rho \in \{0, 0.75, 0.99\}$, and then, we plot the difference in the estimates of $X_1 = \xi_1$ derived from classical OMNI, MASE and multi-RDPG. The plots are depicted in Figure 14.

The top row (resp. bottom row) of panels in Figure 14 illustrates the level of correlation induced by classical OMNI (resp. MASE & multi-RDPG) in terms of the inherent correlation. In the first (resp., second) row, the first two (resp., one) figures separately embed the $\text{JRDPG}(F, n, 2, \rho)$ model followed by Procrustes alignment. The pair of graphs in the OMNI, MASE and multi-RDPG procedures was generated by $\text{JRDPG}(F, n, 2)$, i.e., when $\rho = 0$. While we are able to show explicitly that the level of correlation induced by OMNI is 0.75, the figure suggests that the correlation induced by MASE and multi-RDPG is significantly greater; indeed, the estimates across graphs by these methods are almost identical to each other (up to orthogonal transformation in case of MASE). Depending on the inference task, this additional correlation may be a blessing or a curse, and further work is needed to understand or compute this artifact of the embedding in these other methods.

A natural follow-up question is how to understand the results on induced correlation in the context of modern graph neural network based methods. As an initial experiment, we consider the GATE method of Liu et al. (2021), which uses graph auto-encoding to estimate latent structure across a collection of networks. We repeat the experiment shown in Figures 1 and 14 above using the same correlated SBM model considered therein; however, due to the more expensive computational cost of our off-the-shelf implementation in terms of time and memory, we reduce the number of vertices to 120 and the number of simulations




Figure 15: Effect of model and induced correlation in the joint graph neural network embedding of Liu et al. (2021). Considering the same framework as Figure 14 in the manuscript, we jointly embed the two correlated SBM’s with $\rho = 0.99$ (L) and $\rho = 0$ (R). In each panel, we plot the difference in the latent position estimates across two networks for one of the vertices across numerous Monte Carlo repeats.

to 50. Figure 15 shows the difference across the pair of networks of the latent position estimates for one of the vertices. Compared to Figure 14, the magnitude of this difference is much smaller, suggesting that the network embeddings obtained from the GATE method are much more similar and/or correlated regardless of the correlation value. We further explore these embeddings, and observe that the estimated edge probability matrices obtained from the GATE method are very different from the true model parameters. This is not surprising, given that the method is not primarily designed for estimating the parameters in our SBM/RDPG model; as such, it is unclear how to interpret this induced correlation in context as the aims and estimation goals of this method are potentially different from those of OMNI, MASE and multi-RDPG.

Note that the structure of the generalized omnibus mean matrix $\tilde{P} = J_m \otimes P$ is designed to maintain the low-rank property of P in \tilde{P} and allows for the leading d eigenvectors of P to be related to those of \tilde{P} (indeed $\vec{1} \otimes U_P$ provides a basis for the leading d -dimensional eigenspace of \tilde{P}). If we consider other structures in the Kronecker product that have the form $E \otimes P$, with the restriction that $\vec{1}$ is a leading eigenvector of E , can we replicate the omnibus analysis? We are exploring this at present with examples of interest, including when E has ring-graph structure. While more complex E will necessarily violate the low-rank assumption on \tilde{P} , it remains an open question whether there is still sufficient concentration (about $\vec{1} \otimes \mathbf{X}$) of the leading d -dimensional eigenspace of \mathfrak{M} to permit an accurate a low-rank approximation of \tilde{P} . Ongoing work also includes extensions of the out-of-sample framework

of Levin et al. (2018) to the generalized OMNI framework. Using the OMNI framework to embed the core vertices that are aligned across networks and then out-of-sample embedding any remaining vertices renders possible joint embedding of networks even when the vertices across the networks are only partially (or errorfully) aligned. This, in turn, enables us to use the OMNI framework towards alignment-based tasks such as seeded graph matching (as in Patsolic et al. (2014)).

Finally, given recent extensions for matrix concentrations in sparser regimes (Lei et al., 2020; Bandeira et al., 2016), we suspect that our consistency results can be extended to sparser graphs, along with distributional results and asymptotics for omnibus embeddings in which the latent positions differ and in which the OMNI matrix is asymmetric (while still maintaining $\mathbb{E}(P) = J_m \otimes P$). These extensions can be very challenging, so preliminary results will likely be constrained to only a few special cases. The range of possibilities for follow-on work highlights the important role of joint embeddings in graph inference; as such, the novel intricacies of embedded-space correlation that we examine here are a key component of multiple network inference.

Acknowledgments

This material is based on research sponsored by the Air Force Research Laboratory and Defense Advanced Research Projects Agency (DARPA) under agreement number FA8750-20-2-1001. This work is also supported in part by the D3M program of DARPA. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Research Laboratory and DARPA or the U.S. Government. The authors also gratefully acknowledge the support of NIH grant BRAIN U01-NS108637. We also gratefully acknowledge illustrative conversations with Profs. Minh Tang, Keith Levin, Daniel Sussman and Carey Priebe that helped shaped this work.

Appendix A. Proofs of Main Results

Herein we collect the proofs of the main theoretical results of the paper. Before proving the theorems, we will first establish some convenient asymptotic notation for our theory moving forward.

Definition 18 *Given a sequence of events $\{E_n\} \in \mathcal{F}$, where $n = 1, 2, \dots$, we say that E_n occurs*

- i. asymptotically almost surely and write E_n a.a.s. if $\mathbb{P}(E_n) \rightarrow 1$ as $n \rightarrow \infty$.*
- ii. with high probability, and write E_n w.h.p. , if for some $a_0 \geq 2$, there exists finite positive constant A_0 depending only on a_0 such that*

$$\mathbb{P}[E_n^c] \leq A_0 n^{-a_0}$$

for all n .

We note that E_n occurring w.h.p. is stronger than E_n occurring a.a.s., as w.h.p. implies, by the Borel-Cantelli Lemma (Chung, 2001), that $\mathbb{P}(\limsup E_n) = 1$ and with probability 1 all but finitely many E_n occur.

For matrices $A \in \mathbb{R}^{p_1 \times p_2}$ and $B \in \mathbb{R}^{p_2 \times p_3}$, we will make use of the following common matrix norm identities in the proofs below (where A_i denotes the i -th row of A , and $\|A\| = \|A\|_2$ denotes the spectral norm of A ; the symbol “:=” is used below to denote a definition).

$$\begin{aligned} \|A\|_p &:= \sup_{x \neq 0} \frac{\|Ax\|_p}{\|x\|_p} \text{ for } 1 \leq p \leq \infty; \\ \|A\|_{2 \rightarrow \infty} &:= \sup_{x \neq 0} \frac{\|Ax\|_\infty}{\|x\|_2}; \end{aligned}$$

The following lemma concerns relationships between the $2 \rightarrow \infty$ bound and other classical matrix norms. For a proof, see Cape et al. (2018).

Lemma 19 *For $A \in \mathbb{R}^{p_1 \times p_2}$ and $B \in \mathbb{R}^{p_2 \times p_3}$, we have the following:*

$$\begin{aligned} \|A\|_{2 \rightarrow \infty} &= \max_{i \in [p_1]} \|A_i\|_2; \\ \|A\|_{2 \rightarrow \infty} &\leq \|A\| \leq \sqrt{p_1} \|A\|_{2 \rightarrow \infty}; \\ \|AB\|_{2 \rightarrow \infty} &\leq \|A\|_{2 \rightarrow \infty} \|B\|; \\ \|A\| &\leq \sqrt{\|A\|_1 \|A\|_\infty} \end{aligned}$$

where A_i denotes the i -th row of A .

A.1 Correlation in the Classical Omnibus Embedding and its Relationship to Correlated Stochastic Blockmodel Graphs

We begin with the proof of Theorem 9 from Section 3.

Proof (Induced correlation in classical OMNI and its relationship with correlated block models)

To ease notation, we will frequently suppress the explicit dependence on n in the subscript of $\tilde{P} = \tilde{P}_n$, $B_n^{(k)}$, etc., noting that this dependence is to be implicitly understood throughout. For each $n \geq 1$, we write the spectral decomposition of the positive semidefinite P_n via

$$P_n = \mathbf{X}_n \mathbf{X}_n^T = U_{P_n} S_{P_n} U_{P_n}^T.$$

Let the singular value decomposition of $U_P^T U_{B^{(1)}}$ be denoted $W_{11} D_1 W_{12}^T$ and define $W^{*,1} = W_{11} W_{12}^T$. Similarly, let the singular value decomposition of $U_P^T U_{B^{(2)}}$ be denoted $W_{21} D_2 W_{22}^T$ and define $W^{*,2} = W_{21} W_{22}^T$. As in the proof of Theorem 9 (p 76) in Athreya et al. (2018a), we have that for each $k = 1, 2$ and fixed index i ,

$$\sqrt{n} \left(U_{B^{(k)}} S_{B^{(k)}}^{1/2} - U_P S_P^{1/2} W^{*,k} \right)_i = \sqrt{n} \left((B^{(k)} - P) U_P S_P^{-1/2} W^{*,k} \right)_i + O(n^{-1/2} \log n) \quad (20)$$

with high probability. Letting W_n be a sequence of orthogonal matrices such that

$$U_{P_n} S_{P_n}^{1/2} W_n = \mathbf{X}_n$$

for all $n \geq 1$, we have then that, again with high probability,

$$\sqrt{n} \left(U_{B^{(k)}} S_{B^{(k)}}^{1/2} (W^{*,k})^T W - U_P S_P^{1/2} W \right)_i = \sqrt{n} \left((B^{(k)} - P) U_P S_P^{-1/2} W \right)_i + O(n^{-1/2} \log n) \quad (21)$$

It follows then that for fixed i , we have with high probability

$$\begin{aligned} & \sqrt{n} \left(U_{B^{(1)}} S_{B^{(1)}}^{1/2} (W^{*,1})^T W - U_{B^{(2)}} S_{B^{(2)}}^{1/2} (W^{*,2})^T W \right)_i \\ &= \sqrt{n} \left(U_{B^{(1)}} S_{B^{(1)}}^{1/2} (W^{*,1})^T W - U_P S_P^{1/2} W \right)_i + \sqrt{n} \left(U_P S_P^{1/2} W - U_{B^{(2)}} S_{B^{(2)}}^{1/2} (W^{*,2})^T W \right)_i \\ &= \sqrt{n} \left((B^{(1)} - P) U_P S_P^{-1/2} W \right)_i - \sqrt{n} \left((B^{(2)} - P) U_P S_P^{-1/2} W \right)_i + O\left(\frac{\log n}{\sqrt{n}}\right) \\ &= \sqrt{n} \left((B^{(1)} - B^{(2)}) U_P S_P^{-1/2} W \right)_i + O\left(\frac{\log n}{\sqrt{n}}\right). \end{aligned} \quad (22)$$

Next, note that

$$\begin{aligned} \sqrt{n} \left((B^{(1)} - B^{(2)}) U_P S_P^{-1/2} W \right)_i &= \sqrt{n} \left((B^{(1)} - B^{(2)}) \mathbf{X} \right)_i W^T S_P^{-1} W \\ &= \sqrt{n} \left(\sum_j (B_{i,j}^{(1)} - B_{i,j}^{(2)}) X_j \right) W^T S_P^{-1} W \\ &= \left(n^{-1/2} \sum_j (B_{i,j}^{(1)} - B_{i,j}^{(2)}) X_j \right) (n W^T S_P^{-1} W). \end{aligned}$$

Conditioning on $X_i = x_i$,

$$n^{-1/2} \sum_j (B_{i,j}^{(1)} - B_{i,j}^{(2)}) X_j$$

is a scaled sum of $n - 1$ i.i.d random variables (the $(B_{i,j}^{(1)} - B_{i,j}^{(2)})X_j$), each with mean

$$\begin{aligned}\mathbb{E}((B_{i,j}^{(1)} - B_{i,j}^{(2)})X_j) &= \mathbb{E}(\mathbb{E}((B_{i,j}^{(1)} - B_{i,j}^{(2)})X_j|X_j)) \\ &= \mathbb{E}(X_j\mathbb{E}(B_{i,j}^{(1)} - B_{i,j}^{(2)}|X_j)) = 0\end{aligned}$$

and covariance matrix

$$\tilde{\Sigma}(x_i) = 2(1 - \rho) \mathbb{E}[X_j X_j^T (x_i^T X_j - (x_i^T X_j)^2)],$$

as

$$\begin{aligned}\text{Cov}((B_{i,j}^{(1)} - B_{i,j}^{(2)})X_j) &= \mathbb{E}[(B_{i,j}^{(1)} - B_{i,j}^{(2)})^2 X_j X_j^T] \\ &= \mathbb{E}(\mathbb{E}[(B_{i,j}^{(1)} - B_{i,j}^{(2)})^2 X_j X_j^T | X_j]) \\ &= \mathbb{E}(X_j X_j^T \mathbb{E}[(B_{i,j}^{(1)} - B_{i,j}^{(2)})^2 | X_j]) \\ &= \mathbb{E}(X_j X_j^T \mathbb{E}[(B_{i,j}^{(1)})^2 + (B_{i,j}^{(2)})^2 - 2B_{i,j}^{(1)} B_{i,j}^{(2)} | X_j]) \\ &= \mathbb{E}(X_j X_j^T [2x_i^T X_j - 2x_i^T X_j (x_i^T X_j + \rho(1 - x_i^T X_j))]) \\ &= \mathbb{E}(X_j X_j^T [2x_i^T X_j - 2(x_i^T X_j)^2 - \rho 2(1 - x_i^T X_j)x_i^T X_j]) \\ &= 2(1 - \rho) \mathbb{E}[X_j X_j^T (x_i^T X_j - (x_i^T X_j)^2)].\end{aligned}$$

The classical multivariate central limit theorem then yields

$$n^{-1/2} \sum_j (B_{i,j}^{(1)} - B_{i,j}^{(2)})X_j \xrightarrow{D} \mathcal{N}(0, \tilde{\Sigma}(x_i)). \quad (23)$$

The strong law of large numbers (SLLN) ensures that $\frac{1}{n} \mathbf{X}^T \mathbf{X} = \frac{1}{n} W^T S_P W \xrightarrow{a.s.} \Delta$ and therefore, $nW^T S_P^{-1} W \xrightarrow{a.s.} \Delta^{-1}$. Therefore, by the multivariate Slutsky's Theorem, conditional on $X_i = x_i$, we have that

$$\sqrt{n} \left((B^{(1)} - B^{(2)}) U_P S_P^{-1/2} W \right)_i \xrightarrow{D} \mathcal{N}(0, \tilde{\Sigma}(x_i, \rho)). \quad (24)$$

By setting $W^{(1)} = (W^{*,1})^T W$ and $W^{(2)} = (W^{*,2})^T W$, the requisite result then follows from multivariate Slutsky's applied to Eq. 22 and 24, and integrating the above display over the possible values of x_i with respect to distribution F . ■

A.2 Consistency of Generalized Omnibus Embeddings

Here, we prove our consistency result, Theorem 13 from Section 4, which is, in part, what makes the generalized omnibus embedding useful for estimation. To prove Theorem 13, we will need a Bernstein concentration inequality bound given immediately below, in A.2.1; and number of intermediate supporting lemmas (Lemmas 21–25) all of which are suitably adapted from Levin et al. (2017).

Again, to ease notation, we will on occasion suppress the explicit dependence on n in the subscript of $\mathbf{X} = \mathbf{X}_n$, $\mathfrak{M} = \mathfrak{M}_n$, $\tilde{P} = \tilde{P}_n$, etc., noting that this dependence is implicitly understood throughout. Recall that the spectral decomposition of \tilde{P} is given by

$$\tilde{P} = \mathbb{E}(\mathfrak{M}) = U_{\tilde{P}} S_{\tilde{P}} U_{\tilde{P}}^T,$$

where $U_{\tilde{P}} \in \mathbb{R}^{mn \times d}$ and $S_{\tilde{P}} \in \mathbb{R}^{d \times d}$. Also recall that the adjacency spectral embedding of \mathfrak{M} is given by $\text{ASE}(\mathfrak{M}, d) = U_{\mathfrak{M}} S_{\mathfrak{M}}^{1/2}$.

A.2.1 CONCENTRATION INEQUALITY VIA MATRIX BERNSTEIN

Lemma 20 *Let F be a distribution on a set $\mathcal{X} \in \mathbb{R}^d$ satisfying $\langle x, x' \rangle \in [0, 1]$ for all $x, x' \in \mathcal{X}$. Let $X_1, X_2, \dots, X_n, Y \stackrel{i.i.d.}{\sim} F$, and let $P = \mathbf{X}\mathbf{X}^T$ where $\mathbf{X} = [X_1^T | X_2^T | \dots | X_n^T]^T$. Let*

$$(A_n^{(1)}, A_n^{(2)}, \dots, A_n^{(m)}, \mathbf{X}_n) \sim \text{JRDPG}(F, n, m, R)$$

be a sequence of correlated RDPG random graphs, and let \mathfrak{M}_n denote the generalized omnibus matrix as in Definition 12. Then for n sufficiently large, with high probability,

$$\|\mathfrak{M} - \mathbb{E}\mathfrak{M}\| \leq 4m\sqrt{(n-1)\log mn}.$$

Proof Condition on \mathbf{X} and let $P = \mathbf{X}\mathbf{X}^T$, so that

$$\mathbb{E}\mathfrak{M} = \tilde{P} = \begin{bmatrix} P & P & \dots & P \\ P & P & \dots & P \\ \vdots & \vdots & \ddots & \vdots \\ P & P & \dots & P \end{bmatrix}.$$

For all $l \in [m]$ and $i, j \in [n]$, we define an auxiliary block matrix $E_{i,j}^{(l)} \in \mathbb{R}^{mn \times mn}$ which will help us express the difference $\mathfrak{M} - \mathbb{E}\mathfrak{M}$ as a sum of independent Hermitian matrices, which will allow us to apply Bernstein's matrix bound (see, for example, Theorem 5.4.1 in Vershynin (2018)). For $i \neq j$, let $e^{ij} = e_i e_j^T + e_j e_i^T$ where e_i is a vector in \mathbb{R}^n with all its entries equal to 0 except the i -th entry which is equal to 1, so that e^{ij} is a matrix whose entries are all equal to 0 except the (i, j) th and (j, i) th entries which are equal to 1. For $l \in [m]$, we define the $m \times m$ matrix $C^{(l)}$ (symmetric because \mathfrak{M} is symmetric) as follows,

$$C^{(l)} = \begin{bmatrix} c_\ell^{(1,1)} & c_\ell^{(1,2)} & \dots & c_\ell^{(1,m)} \\ c_\ell^{(2,1)} & c_\ell^{(2,2)} & \dots & c_\ell^{(2,m)} \\ \vdots & \vdots & \ddots & \vdots \\ c_\ell^{(m,1)} & c_\ell^{(m,2)} & \dots & c_\ell^{(m,m)} \end{bmatrix}.$$

We then define the $mn \times mn$ block (symmetric) matrix $E_{i,j}^{(l)}$ as the Kronecker product of $C^{(l)}$ and e^{ij} , i.e., $E_{i,j}^{(l)} = C^{(l)} \otimes e^{ij}$. Now, we can write $\mathfrak{M} - \mathbb{E}\mathfrak{M}$ as

$$\mathfrak{M} - \mathbb{E}\mathfrak{M} = \sum_{1 \leq \ell \leq m} \sum_{1 \leq i < j \leq n} (A_{ij}^{(\ell)} - P_{ij}) E_{i,j}^{(\ell)} = \sum_{1 \leq i < j \leq n} \underbrace{\left(\sum_{1 \leq \ell \leq m} (A_{ij}^{(\ell)} - P_{ij}) E_{i,j}^{(\ell)} \right)}_{:= B_{i,j,m}}.$$

These $B_{i,j,m}$ are then symmetric, mean-zero, independent matrices, as we have grouped edges indexed by the same (i, j) pair across graphs to account for the across graph correlation and achieve the desired independence.

As we have written $\mathfrak{M} - \mathbb{E}\mathfrak{M}$ as the sum of $\binom{n}{2}$ symmetric, mean-zero, independent matrices $B_{i,j,m}$ where for all $i, j \in [n]$, we have that

$$\|B_{i,j,m}\| \leq m \max_{\ell \in [m]} \|C^{(\ell)}\| \leq m \left(\max_{\ell \in [m]} \max_{q \in [m]} \sum_{k=1}^m c_\ell^{(q,k)} \right).$$

We can now apply the matrix Bernstein inequality to derive the desired concentration of $\mathfrak{M} - \mathbb{E}\mathfrak{M}$. To this end, let

$$L := \max_{\ell \in [m]} \max_{q \in [m]} \sum_{k=1}^m c_\ell^{(q,k)} \leq m;$$

To apply the matrix Bernstein inequality, it remains to compute the variance term

$$\begin{aligned} \nu(\mathfrak{M} - \mathbb{E}\mathfrak{M}) &= \left\| \sum_{1 \leq i < j \leq n} \mathbb{E}[B_{i,j,m}^2] \right\| \\ &\leq \left\| \sum_{1 \leq i < j \leq n} \sum_{1 \leq \ell \leq m} \mathbb{E} \left[(A_{ij}^{(\ell)} - P_{ij})^2 (E_{i,j}^{(\ell)})^2 \right] \right\| \\ &\quad + 2 \left\| \sum_{1 \leq i < j \leq n} \sum_{\ell_1 < \ell_2} \mathbb{E} \left[(A_{ij}^{(\ell_1)} - P_{ij})(A_{ij}^{(\ell_2)} - P_{ij}) E_{i,j}^{(\ell_1)} E_{i,j}^{(\ell_2)} \right] \right\|. \end{aligned} \quad (25)$$

To bound the variance term $\nu(\mathfrak{M} - \mathbb{E}\mathfrak{M})$ we will bound the two terms in Eq. 25 independently. Let $D_{ij} = e^{ij} e^{ij} = e_i e_i^T + e_j e_j^T \in \mathbb{R}^{n \times n}$. The mixed-product property of Kronecker products implies that

$$E_{i,j}^{(\ell_1)} E_{i,j}^{(\ell_2)} = (C^{(\ell_1)} \otimes e^{ij})(C^{(\ell_2)} \otimes e^{ij}) = C^{(\ell_1)} C^{(\ell_2)} \otimes D_{ij}.$$

Note that for any symmetric matrix A , we have that $\|A\| \leq \sqrt{\|A\|_1 \|A\|_\infty} = \|A\|_\infty$ (see, for example, Horn and Johnson (1985) Ex. 5.6.21), where $\|\cdot\|_1$ is the maximum column sum matrix norm and $\|\cdot\|_\infty$ the maximum row sum matrix norm. The first term is then

bounded above via

$$\begin{aligned}
 \left\| \sum_{1 \leq i < j \leq n} \sum_{1 \leq \ell \leq m} \mathbb{E} \left[(A_{ij}^{(\ell)} - P_{ij})^2 (E_{i,j}^{(\ell)})^2 \right] \right\| &= \left\| \underbrace{\sum_{1 \leq \ell \leq m} \sum_{1 \leq i < j \leq n} P_{ij} (1 - P_{ij}) [(C^{(\ell)})^2 \otimes D_{ij}]}_{:= \Delta_1} \right\| \\
 &\leq \|\Delta_1\|_\infty \\
 &\leq \frac{1}{4} (n-1) \max_{q \in [m]} \sum_{\ell=1}^m \sum_{r=1}^m \sum_{k=1}^m c_\ell^{(q,k)} c_\ell^{(k,r)} \\
 &= \frac{1}{4} (n-1) \max_{q \in [m]} \sum_{\ell=1}^m \sum_{k=1}^m c_\ell^{(q,k)} \alpha(k, \ell) \\
 &\leq \frac{1}{4} (n-1) mL.
 \end{aligned}$$

The second term is then bounded above via

$$\begin{aligned}
 &\left\| \sum_{1 \leq i < j \leq n} \sum_{\ell_1 < \ell_2} \mathbb{E} \left[(A_{ij}^{(\ell_1)} - P_{ij})(A_{ij}^{(\ell_2)} - P_{ij}) E_{i,j}^{(\ell_1)} E_{i,j}^{(\ell_2)} \right] \right\| \\
 &= \left\| \underbrace{\sum_{1 \leq i < j \leq n} P_{ij} (1 - P_{ij}) \sum_{\ell_1 < \ell_2} \rho_{\ell_1, \ell_2} C^{(\ell_1)} C^{(\ell_2)} \otimes D_{ij}}_{:= \Delta_2} \right\| \\
 &\leq \sqrt{\|\Delta_2\|_\infty \|\Delta_2\|_1} \\
 &\leq \frac{1}{4} (n-1) \sqrt{\left(\max_{q \in [m]} \sum_{\ell_1 < \ell_2} \sum_{r=1}^m \sum_{k=1}^m c_{\ell_1}^{(q,k)} c_{\ell_2}^{(k,r)} \right) \left(\max_{r \in [m]} \sum_{\ell_1 < \ell_2} \sum_{q=1}^m \sum_{k=1}^m c_{\ell_1}^{(q,k)} c_{\ell_2}^{(k,r)} \right)}.
 \end{aligned}$$

Note that

$$\begin{aligned}
 \max_{q \in [m]} \sum_{\ell_1 < \ell_2} \sum_{r=1}^m \sum_{k=1}^m c_{\ell_1}^{(q,k)} c_{\ell_2}^{(k,r)} &= \max_{q \in [m]} \sum_{\ell_1 < \ell_2} \sum_{k=1}^m c_{\ell_1}^{(q,k)} \alpha(k, \ell_2) \\
 &\leq m \cdot \max_{q \in [m]} \sum_{\ell_1=1}^m \sum_{k=1}^m c_{\ell_1}^{(q,k)} \\
 &= m \cdot \max_{q \in [m]} \sum_{\ell_1=1}^m \alpha(q, \ell_1) \\
 &= m^2,
 \end{aligned}$$

and that a similar bound holds for $\max_{r \in [m]} \sum_{\ell_1 < \ell_2} \sum_{q=1}^m \sum_{k=1}^m c_{\ell_1}^{(q,k)} c_{\ell_2}^{(k,r)}$. Therefore, we have that the second term in the variance is bounded by $(n-1)m^2/4$ yielding a total variance bound of

$$\nu(\mathfrak{M} - \mathbb{E} \mathfrak{M}) \leq (n-1)m^2.$$

To apply the matrix Bernstein bound, let $t = 4m\sqrt{(n-1)\log mn}$, and then

$$\begin{aligned} \mathbb{P}\left[\|\mathfrak{M} - \mathbb{E}\mathfrak{M}\| \geq 4m\sqrt{(n-1)\log mn}\right] &\leq 2nm \cdot \exp\left\{\frac{-8m^2(n-1)\log mn}{\nu(\mathfrak{M} - \mathbb{E}\mathfrak{M}) + \frac{4m^2L\sqrt{(n-1)\log mn}}{3}}\right\} \\ &\leq 2nm \cdot \exp\left\{\frac{-8m^2(n-1)\log mn}{(n-1)m^2 + \frac{4m^3\sqrt{(n-1)\log mn}}{3}}\right\} \\ &\leq 2nm \cdot \exp\left\{\frac{-8m^2(n-1)\log mn}{2(n-1)m^2}\right\} \\ &= 2m^{-3}n^{-3}, \end{aligned}$$

where the third inequality holds for n sufficiently large ($\sqrt{n} \geq \frac{4}{3}m\log mn$) as we assume m is not growing in n . Integrating over the \mathbf{X} then yields the desired result. \blacksquare

A.2.2 SUPPORTING LEMMAS

Lemma 21 (Observation 2 in Levin et al. (2017)) *Let F be a distribution on a set $\mathcal{X} \in \mathbb{R}^d$ satisfying $\langle x, x' \rangle \in [0, 1]$ for all $x, x' \in \mathcal{X}$. Let $X_1, X_2, \dots, X_n, Y \stackrel{i.i.d.}{\sim} F$, and let $P = \mathbf{X}\mathbf{X}^T$ where $\mathbf{X} = [X_1^T | X_2^T | \dots | X_n^T]^T$. Then with probability at least $1 - \frac{d^2}{n^2}$, we have that*

$$|\lambda_i(P) - n\lambda_i\mathbb{E}(YY^T)| \leq 2d\sqrt{n\log n}.$$

Furthermore, with high probability there exists a constant $C > 0$ such that for all $i \in [d]$, $\lambda_i(P) \geq Cn\delta$ and $\lambda_i(\tilde{P}) \geq Cnm\delta$, where $\delta = \lambda_d(\mathbb{E}(YY^T))$.

Assuming $\delta > 0$, combining Lemma 21 with the fact that the rows of $U_{\tilde{P}}S_{\tilde{P}}^{1/2}$ are bounded in Euclidean norm by 1, we obtain

$$\|U_{\tilde{P}}\|_{2 \rightarrow \infty} \leq C(mn)^{-1/2} \quad \text{w.h.p.} \quad (26)$$

Lemma 22 *With notation as in Lemma 21, assume that $\delta > 0$ so that $\mathbb{E}(YY^T)$ is full rank. Let the singular value decomposition of $U_{\tilde{P}}^T U_{\mathfrak{M}} \in \mathbb{R}^{d \times d}$ be given by $V_1 \Sigma V_2^T$. Then, there exists a constant $C > 0$ such that*

$$\|U_{\tilde{P}}^T U_{\mathfrak{M}} - V_1 V_2^T\|_F \leq C \frac{\log mn}{n} \quad \text{w.h.p.}$$

Proof [Adapted from the proof of Proposition 16 in Lyzinski et al. (2017)] Working on the intersection of the sets where both Lemma 20 and 21 hold (noting this set has high probability), note that Weyl's theorem gives that $\lambda_d(\mathfrak{M}) \geq Cnm$ for some constant C . Let $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_d$ be the singular values of $U_{\tilde{P}}^T U_{\mathfrak{M}}$, so that $\sigma_i = \cos(\theta_i)$ where the θ_i 's are the principal angles between the subspaces spanned by $U_{\tilde{P}}$ and $U_{\mathfrak{M}}$. The Davis-Kahan

theorem (see, for example, Theorem 3.6 in Bhatia (1997) Theorem VII.3.1) then implies that with high probability (where C is a constant that can change line-to-line)

$$\begin{aligned}
 \|U_{\tilde{P}}^T U_{\mathfrak{M}} - V_1 V_2^T\|_F &= \sqrt{\sum_i (1 - \sigma_i)^2} \\
 &\leq \sum_i (1 - \sigma_i^2) \\
 &\leq d \max_i |\sin(\theta_i)|^2 \\
 &\leq C \frac{d^2 \|\mathfrak{M} - \tilde{P}\|^2}{\lambda_d(\mathfrak{M})^2} \\
 &\leq C \frac{m^2 (n-1) \log mn}{n^2 m^2} \\
 &\leq C \frac{\log mn}{n},
 \end{aligned}$$

where the bounds in the fourth line follows from Davis-Kahan, and those in the second-to-last line follow from Lemma 20 (numerator) and Lemma 21 (denominator). \blacksquare

Lemma 23 *With the assumptions and notation of Lemma 22, let $V = V_1 V_2^T$. Then we have that*

$$\|V S_{\mathfrak{M}} - S_{\tilde{P}} V\|_F \leq C m^2 \log mn \quad w.h.p., \quad (27)$$

$$\|V S_{\mathfrak{M}}^{1/2} - S_{\tilde{P}}^{1/2} V\|_F \leq C \frac{m^{3/2} \log mn}{n^{1/2}} \quad w.h.p., \quad (28)$$

$$\|V S_{\mathfrak{M}}^{-1/2} - S_{\tilde{P}}^{-1/2} V\|_F \leq C \frac{m^{1/2} \log mn}{n^{3/2}} \quad w.h.p., \quad (29)$$

where $S_{\mathfrak{M}}^{-1}$, and $S_{\tilde{P}}^{-1}$ are understood to be the pseudoinverses of $S_{\mathfrak{M}}$, and $S_{\tilde{P}}$ in the unlikely event these matrices are singular.

Proof [Adapted from the proof of Proposition 17 in Lyzinski et al. (2017)] Let

$$R := U_{\mathfrak{M}} - U_{\tilde{P}} U_{\tilde{P}}^T U_{\mathfrak{M}},$$

so that

$$\begin{aligned}
 V S_{\mathfrak{M}} &= (V - U_{\tilde{P}}^T U_{\mathfrak{M}}) S_{\mathfrak{M}} + U_{\tilde{P}}^T U_{\mathfrak{M}} S_{\mathfrak{M}} \\
 &= (V - U_{\tilde{P}}^T U_{\mathfrak{M}}) S_{\mathfrak{M}} + U_{\tilde{P}}^T \mathfrak{M} U_{\mathfrak{M}} I_{\pm} \\
 &= (V - U_{\tilde{P}}^T U_{\mathfrak{M}}) S_{\mathfrak{M}} + U_{\tilde{P}}^T (\mathfrak{M} - \tilde{P}) U_{\mathfrak{M}} I_{\pm} + U_{\tilde{P}}^T \tilde{P} U_{\mathfrak{M}} I_{\pm} \\
 &= (V - U_{\tilde{P}}^T U_{\mathfrak{M}}) S_{\mathfrak{M}} + U_{\tilde{P}}^T (\mathfrak{M} - \tilde{P}) R I_{\pm} + U_{\tilde{P}}^T (\mathfrak{M} - \tilde{P}) U_{\tilde{P}} U_{\tilde{P}}^T U_{\mathfrak{M}} I_{\pm} + U_{\tilde{P}}^T \tilde{P} U_{\mathfrak{M}} I_{\pm} \\
 &= (V - U_{\tilde{P}}^T U_{\mathfrak{M}}) S_{\mathfrak{M}} + U_{\tilde{P}}^T (\mathfrak{M} - \tilde{P}) R I_{\pm} + U_{\tilde{P}}^T (\mathfrak{M} - \tilde{P}) U_{\tilde{P}} U_{\tilde{P}}^T U_{\mathfrak{M}} I_{\pm} + S_{\tilde{P}} U_{\tilde{P}}^T U_{\mathfrak{M}} I_{\pm},
 \end{aligned}$$

where I_{\pm} is a random diagonal matrix with ± 1 entries on its diagonal indicating whether the signs of the eigenvalues associated with $U_{\mathfrak{M}}$ agree for \mathfrak{M} and for $|\mathfrak{M}|$. We will first bound the Frobenius norm of $U_{\tilde{P}}^T(\mathfrak{M} - \tilde{P})U_{\tilde{P}}$. Specifically, we will prove that there exists a constant C such that w.h.p.

$$\|U_{\tilde{P}}^T(\mathfrak{M} - \tilde{P})U_{\tilde{P}}\|_F \leq d\|U_{\tilde{P}}^T(\mathfrak{M} - \tilde{P})U_{\tilde{P}}\|_{\max} \leq Cdm^2\sqrt{\log mn}. \quad (30)$$

The proof proceeds as follows. Note that $U_{\tilde{P}}^T(\mathfrak{M} - \tilde{P})U_{\tilde{P}} \in \mathbb{R}^{d \times d}$ and $\|U_{\tilde{P}}^T(\mathfrak{M} - \tilde{P})U_{\tilde{P}}\|_{\max} = \max_{i,j \in [d]} \left| \langle (\mathfrak{M} - \tilde{P})U_{\cdot j}, U_{\cdot i} \rangle \right|$. For any $1 \leq i, j \leq d$,

$$\begin{aligned} \langle (\mathfrak{M} - \tilde{P})U_{\cdot j}, U_{\cdot i} \rangle &= U_{\tilde{P}}^T(\mathfrak{M} - \tilde{P})U_{\cdot j} \\ &= 2 \sum_{1 \leq \tilde{k} < \tilde{l} \leq mn} (\mathfrak{M}_{\tilde{k}, \tilde{l}} - \tilde{P}_{\tilde{k}, \tilde{l}}) U_{\tilde{k}i} U_{\tilde{l}j} - \sum_{\tilde{k}} \tilde{P}_{\tilde{k}, \tilde{k}} U_{\tilde{k}i} U_{\tilde{k}j} \\ &= 2 \sum_{1 \leq k < l \leq n} \left(\sum_{s=1}^m \sum_{t=1}^m \mathfrak{M}_{k,l}^{(s,t)} - m^2 P_{k,l} \right) U_{ki} U_{lj} - m \sum_{k=1}^n P_{k,k} U_{ki} U_{kj} \\ &= 2 \sum_{1 \leq k < l \leq n} \left(\sum_{s=1}^m \sum_{t=1}^m \sum_{\ell=1}^m c_{\ell}^{(s,t)} A_{k,l}^{(\ell)} - m^2 P_{k,l} \right) U_{ki} U_{lj} - m \sum_{k=1}^n P_{k,k} U_{ki} U_{kj} \\ &= 2 \sum_{1 \leq k < l \leq n} \underbrace{\left(\sum_{\ell=1}^m \left(\sum_{s=1}^m \alpha(s, \ell) \right) A_{k,l}^{(\ell)} - m^2 P_{k,l} \right)}_{\text{mean 0 r.v.}} U_{ki} U_{lj} - m \sum_{k=1}^n P_{k,k} U_{ki} U_{kj}. \end{aligned}$$

Conditioned on \mathbf{X} (and hence on P), $U_{\tilde{P}}^T(\mathfrak{M} - \tilde{P})U_{\cdot j}$ consists of two terms. The first term is a sum of $\binom{n}{2}$ independent, mean zero, bounded random variables taking values in $[-2m^2 U_{ki} U_{lj}, 2m^2 U_{ki} U_{lj}]$ and the second term is of order $O(m)$, and thus, dominated by the first term. By Hoeffding's inequality,

$$\begin{aligned} \mathbb{P} \left[\left| 2 \sum_{1 \leq \tilde{k} < \tilde{l} \leq mn} (\mathfrak{M}_{\tilde{k}, \tilde{l}} - \tilde{P}_{\tilde{k}, \tilde{l}}) U_{\tilde{k}i} U_{\tilde{l}j} \right| \geq 8m^2 \sqrt{\log mn} \right] &\leq 2 \exp \left(\frac{-32m^4 \log mn}{\sum_{1 \leq k < l \leq n} (4m^2 U_{ki} U_{lj})^2} \right) \\ &\leq 2(mn)^{-2}. \end{aligned}$$

Integrating over \mathbf{X} yields the desired result that there exists a constant $C > 0$ such that

$$\|U_{\tilde{P}}^T(\mathfrak{M} - \tilde{P})U_{\tilde{P}}\|_F \leq Cm^2\sqrt{\log mn} \quad \text{w.h.p.}$$

Consider the events (where C is an appropriately chosen constant)

$$\begin{aligned} \mathcal{E}_1 &:= \left\{ \|U_{\tilde{P}}^T(\mathfrak{M} - \tilde{P})U_{\tilde{P}}\|_F \leq Cm^2 \log mn \right\} \\ \mathcal{E}_2 &:= \{ \text{the statement of Lemma 21 holds} \} \\ \mathcal{E}_3 &:= \{ \text{the statement of Lemma 22 holds} \} \\ \mathcal{E}_4 &:= \{ \text{the statement of Lemma 20 holds} \} \end{aligned}$$

As each \mathcal{E}_i is a high probability event, $\mathcal{E} = \cap_{i=1}^4 \mathcal{E}_i$ is also a high probability event. In what follows, we will condition on the events in \mathcal{E} occurring.

Conditioning on \mathcal{E} , from the Davis-Kahan theorem (shown in detail in Lemma 25), we have that

$$\|U_{\mathfrak{M}} - U_{\tilde{P}} U_{\tilde{P}}^T U_{\mathfrak{M}}\| \leq \sqrt{\frac{C \log mn}{n}} \quad (31)$$

holds for some constant C . Also, Weyl's theorem (Horn and Johnson, 1985, Section 6.3) with Lemmas 21 and 20 imply that

$$|\lambda_j(\mathfrak{M})| \leq Cm \sqrt{n \log mn}$$

for $j > d$, and that

$$|\lambda_j(\mathfrak{M}) - \lambda_j(\tilde{P})| \leq Cm \sqrt{n \log mn}$$

for $j \leq d$. As for $j \leq d$, we have that there is a constant $C > 0$ such that $\lambda_j(\tilde{P}) \geq Cnm\delta$, and so $\lambda_i(|\mathfrak{M}|) = \lambda_i(\mathfrak{M}) \geq Cnm\delta$ for all $i \in [d]$ and an appropriately chosen constant C (abusing notation, the two C 's need not be equal). Therefore, we have that $I_{\pm} = I_d$ so that

$$S_{\tilde{P}} U_{\tilde{P}}^T U_{\mathfrak{M}} I_{\pm} = S_{\tilde{P}} U_{\tilde{P}}^T U_{\mathfrak{M}} = S_{\tilde{P}} (U_{\tilde{P}}^T U_{\mathfrak{M}} - V) + S_{\tilde{P}} V.$$

We have then that

$$\begin{aligned} \|V S_{\mathfrak{M}} - S_{\tilde{P}} V\|_F &\leq \|V - U_{\tilde{P}}^T U_{\mathfrak{M}}\|_F (\|S_{\mathfrak{M}}\| + \|S_{\tilde{P}}\|) \\ &\quad + \|U_{\tilde{P}}^T (\mathfrak{M} - \tilde{P}) R\|_F + \|U_{\tilde{P}}^T (\mathfrak{M} - \tilde{P}) U_{\tilde{P}} U_{\tilde{P}}^T U_{\mathfrak{M}}\|_F. \end{aligned}$$

Now, we have that $\|\mathfrak{M}\|, \|\tilde{P}\| \leq mn$. It follows then that there exists a constant C such that (where the first bound follows from Lemma 22, and the second by combining Lemma 20 and Eq. 31)

$$\begin{aligned} \|V - U_{\tilde{P}}^T U_{\mathfrak{M}}\|_F (\|S_{\mathfrak{M}}\| + \|S_{\tilde{P}}\|) &\leq Cm \log mn \\ \|U_{\tilde{P}}^T (\mathfrak{M} - \tilde{P}) R\|_F &\leq \sqrt{d} \|U_{\tilde{P}}^T (\mathfrak{M} - \tilde{P}) R\| \leq Cm \log mn, \end{aligned}$$

so that

$$\|V S_{\mathfrak{M}} - S_{\tilde{P}} V\|_F \leq Cm \log mn + \|U_{\tilde{P}}^T (\mathfrak{M} - \tilde{P}) U_{\tilde{P}}\|_F \leq Cm^2 \log mn.$$

as desired, thus proving Eq. (27).

To prove Eq. (28) and (29), note that we have shown that with high probability, both of the following events hold:

- i. $\|V S_{\mathfrak{M}} - S_{\tilde{P}} V\|_F \leq Cm^2 \log mn$ (by the first part of the Lemma).
- ii. For all $i, j \in [d]$, $\lambda_i(\mathfrak{M}), \lambda_j(\tilde{P})$ are of order $\Omega(nm\delta)$ (by Lemmas 21 and 20).

Given these events, as we have that

$$(S_{\mathfrak{M}}^{1/2} - S_{\tilde{P}}^{1/2})(S_{\mathfrak{M}}^{1/2} + S_{\tilde{P}}^{1/2}) = (S_{\mathfrak{M}} - S_{\tilde{P}}) \Rightarrow (S_{\mathfrak{M}}^{1/2} - S_{\tilde{P}}^{1/2}) = (S_{\mathfrak{M}} - S_{\tilde{P}})(S_{\mathfrak{M}}^{1/2} + S_{\tilde{P}}^{1/2})^{-1},$$

the i, j -th entry of $VS_{\mathfrak{M}}^{1/2} - S_{\tilde{P}}^{1/2}V$ is equal to

$$V_{i,j}(\lambda_i^{1/2}(\mathfrak{M}) - \lambda_j^{1/2}(\tilde{P})) = \frac{V_{i,j}(\lambda_i(\mathfrak{M}) - \lambda_j(\tilde{P}))}{\lambda_i^{1/2}(\mathfrak{M}) + \lambda_j^{1/2}(\tilde{P})}$$

and Eq. 28 follows. To prove Eq. 29, note that the i, j -th entry of $VS_{\mathfrak{M}}^{-1/2} - S_{\tilde{P}}^{-1/2}V$ is equal to

$$V_{i,j}(\lambda_i^{-1/2}(\mathfrak{M}) - \lambda_j^{-1/2}(\tilde{P})) = \frac{V_{i,j}(\lambda_j(\tilde{P})^{1/2} - \lambda_i(\mathfrak{M})^{1/2})}{\sqrt{\lambda_i(\mathfrak{M})\lambda_j(\tilde{P})}}.$$

The proof then follows immediately. ■

Lemma 24 *With the assumptions in Lemma 22, denote the SVD of $U_{\tilde{P}}^T U_{\mathfrak{M}}$ as $V_1 \Sigma V_2^T$ and set $V := V_1 V_2^T$. There exists a constant C such that w.h.p.*

$$\|(\mathfrak{M} - \tilde{P})U_{\tilde{P}}\|_{2 \rightarrow \infty} \leq C\sqrt{dm}\sqrt{\log mn}. \quad (32)$$

Proof We note first that

$$\frac{1}{\sqrt{d}}\|(\mathfrak{M} - \tilde{P})U_{\tilde{P}}\|_{2 \rightarrow \infty} \leq \|(\mathfrak{M} - \tilde{P})U_{\tilde{P}}\|_{\max} = \max_{i \in [mn], j \in [d]} |\langle (\mathfrak{M} - \tilde{P})U_{\tilde{P}}, e_i \rangle|,$$

where $e_i \in \mathbb{R}^{mn}$ is the unit vector with all of its entries equal to 0, except for the i -th entry.

Let $s \in [m]$ arbitrary. There exists a matrix W such that

$$U_{\tilde{P}} S_{\tilde{P}}^{1/2} W = \mathbf{Z} \Rightarrow U_{\tilde{P}} = \mathbf{Z} W^T (S_{\tilde{P}})^{-1/2} = \begin{bmatrix} \mathbf{X} W^T (S_{\tilde{P}})^{-1/2} \\ \mathbf{X} W^T (S_{\tilde{P}})^{-1/2} \\ \vdots \\ \mathbf{X} W^T (S_{\tilde{P}})^{-1/2} \end{bmatrix}.$$

Therefore, for all $1 \leq k \leq n$, $U_{k,j} = U_{k+n,j} = \dots = U_{k+(m-1)n,j}$. For each $(s-1)n+1 \leq i \leq sn$ and $1 \leq j \leq d$,

$$\begin{aligned} \langle (\mathfrak{M} - \tilde{P})U_{\tilde{P}}, e_i \rangle &= e_i^T (\mathfrak{M} - \tilde{P})U_{\tilde{P}} \\ &= \sum_{\tilde{k}=1}^{mn} (\mathfrak{M}_{i,\tilde{k}} - \tilde{P}_{i,\tilde{k}}) U_{\tilde{k},j} \\ &= \sum_{k=1}^n \left[\sum_{t=1}^m \mathfrak{M}_{i,k}^{(s,t)} - mP_{i,k} \right] U_{k,j} \\ &= \sum_{k=1}^n \left[\sum_{t=1}^m \sum_{\ell=1}^m c_{\ell}^{(s,t)} A_{i,k}^{(\ell)} - mP_{i,k} \right] U_{k,j} \\ &= \sum_{k=1}^n \left[\sum_{\ell=1}^m \alpha(s, \ell) A_{i,k}^{(\ell)} - mP_{i,k} \right] U_{k,j}, \end{aligned} \quad (33)$$

where $\alpha(s, \ell) := \sum_{t=1}^m c_\ell^{(s,t)} \geq 0$ and $\sum_{\ell=1}^m \alpha(s, \ell) = m$ for all $s \in [m]$. Conditioning on \mathbf{X} (and hence on \mathbf{P}) for all $s \in [m]$, for any $(s-1)n+1 \leq i \leq sn$ and $1 \leq j \leq d$, the above expansion is a sum of n independent (in k), bounded, mean zero random variables taking values in $[-mU_{k,j}, mU_{k,j}]$. Hence, by Hoeffding's inequality,

$$\mathbb{P} \left(\left| \sum_{\tilde{k}=1}^{mn} (\mathfrak{M}_{i,\tilde{k}} - \tilde{P}_{i,\tilde{k}}) U_{\tilde{k},j} \right| \geq 2m\sqrt{\log mn} \right) \leq 2 \exp \left(- \frac{8m^2 \log mn}{\sum_{k=1}^n (2mU_{k,j})^2} \right) \leq 2(mn)^{-2},$$

where we used the fact that $\sum_{k=1}^n (2mU_{k,j})^2 \leq 4m^2$ as the columns of $U_{\tilde{P}}$ are norm 1. Therefore, by integrating over \mathbf{X} we have that w.h.p. there exists a constant $C > 0$ such that $\|(\mathfrak{M} - \tilde{P})U_{\tilde{P}}\|_{2 \rightarrow \infty} \leq C\sqrt{dm}\sqrt{\log mn}$ as desired. \blacksquare

Lemma 25 *With the assumptions in Lemma 22, denote the SVD of $U_{\tilde{P}}^T U_{\mathfrak{M}}$ as $V_1 \Sigma V_2^T$ and set $V := V_1 V_2^T$. Set also $Q^{(1)} = U_{\mathfrak{M}} - U_{\tilde{P}} V$, there exists a constant C such that w.h.p.,*

$$\|Q^{(1)}\| \leq C \frac{\log^{1/2} mn}{n^{1/2}}.$$

Proof Following the reasoning from Lemma 6.8 in Cape et al. (2018), we add and subtract $U_{\tilde{P}} U_{\tilde{P}}^T U_{\mathfrak{M}}$ and by triangle inequality,

$$\|Q^{(1)}\| = \|U_{\mathfrak{M}} - U_{\tilde{P}} V\| \leq \|U_{\mathfrak{M}} - U_{\tilde{P}} U_{\tilde{P}}^T U_{\mathfrak{M}}\| + \|U_{\tilde{P}} (U_{\tilde{P}}^T U_{\mathfrak{M}} - V)\|.$$

The first term can be rewritten as follows,

$$\|U_{\mathfrak{M}} - U_{\tilde{P}} U_{\tilde{P}}^T U_{\mathfrak{M}}\| = \|U_{\mathfrak{M}} U_{\mathfrak{M}}^T - U_{\tilde{P}} U_{\tilde{P}}^T U_{\mathfrak{M}} U_{\mathfrak{M}}^T\| = \|(I - U_{\tilde{P}} U_{\tilde{P}}^T) U_{\mathfrak{M}} U_{\mathfrak{M}}^T\| = \|\sin \Theta(U_{\mathfrak{M}}, U_{\tilde{P}})\|.$$

Given the intersection of the events in the statements of Lemmas 21 and 20, the Davis-Kahan theorem gives that,

$$\|U_{\mathfrak{M}} - U_{\tilde{P}} U_{\tilde{P}}^T U_{\mathfrak{M}}\| \leq \frac{\|\mathfrak{M} - \tilde{P}\|}{\lambda_d(\mathfrak{M})} \leq C \frac{\log^{1/2} mn}{n^{1/2}}.$$

Further, we can bound $\|U_{\tilde{P}} (U_{\tilde{P}}^T U_{\mathfrak{M}} - V)\|$ using Lemma 22; as the intersection of the events in Lemma 21, 22 and 20 has high probability, this leads us to the desired result. \blacksquare

We now address the proof of Theorem 13.

Proof (Consistency of generalized omnibus embeddings) Let W_n be a sequence of matrices such that $\mathbf{Z} = \mathbf{Z}^* W_n$, and let V be as in Lemma 23. Define the matrices $Q^{(1)}, Q^{(2)}$ as follows

$$\begin{aligned} Q^{(1)} &= U_{\mathfrak{M}} - U_{\tilde{P}} V; \\ Q^{(2)} &= U_{\tilde{P}} U_{\tilde{P}}^T U_{\mathfrak{M}} - U_{\tilde{P}} V. \end{aligned}$$

Note that, as defined, $\mathfrak{M}U_{\mathfrak{M}} = U_{\mathfrak{M}}S_{\mathfrak{M}}I_{\pm}$ where I_{\pm} is a random sign matrix indicating whether the signs of the eigenvalues associated with $U_{\mathfrak{M}}$ agree for \mathfrak{M} and $|\mathfrak{M}|$. Given the events of Lemmas 21 and 20, we have that (by Weyl's Theorem) $I_{\pm} = I_d$; as this term appears only in the residual terms (H_2 - H_5) of the below decomposition, and as, when bounding the residual terms we assume the events of Lemmas 21 and 20, we write, with a slight abuse of notation, $\mathfrak{M}U_{\mathfrak{M}} = U_{\mathfrak{M}}S_{\mathfrak{M}}$ in the below decomposition. We decompose the term $U_{\mathfrak{M}}S_{\mathfrak{M}}^{1/2} - U_{\tilde{P}}S_{\tilde{P}}^{1/2}V$ as follows,

$$\begin{aligned}
 U_{\mathfrak{M}}S_{\mathfrak{M}}^{1/2} - U_{\tilde{P}}S_{\tilde{P}}^{1/2}V &= \underbrace{(\mathfrak{M} - \tilde{P})U_{\tilde{P}}S_{\tilde{P}}^{-1/2}V}_{:=H_1} + \underbrace{(\mathfrak{M} - \tilde{P})U_{\tilde{P}}(VS_{\mathfrak{M}}^{-1/2} - S_{\tilde{P}}^{-1/2}V)}_{:=H_2} \\
 &\quad - \underbrace{U_{\tilde{P}}U_{\tilde{P}}^T(\mathfrak{M} - \tilde{P})U_{\tilde{P}}VS_{\mathfrak{M}}^{-1/2}}_{:=H_3} + \underbrace{(I - U_{\tilde{P}}U_{\tilde{P}}^T)(\mathfrak{M} - \tilde{P})Q^{(1)}S_{\mathfrak{M}}^{-1/2}}_{:=H_4} \\
 &\quad + \underbrace{Q^{(2)}S_{\mathfrak{M}}^{1/2} + U_{\tilde{P}}(VS_{\mathfrak{M}}^{1/2} - S_{\tilde{P}}^{1/2}V)}_{:=H_5}. \tag{34}
 \end{aligned}$$

For the H_i 's, $i = 1, 2, 3, 4$, there exists a constant $C > 0$ such that the following bounds hold w.h.p.:

$$\begin{aligned}
 \|H_1\|_{2 \rightarrow \infty} &\leq \|S_{\tilde{P}}^{-1/2}\| \cdot \|(\mathfrak{M} - \tilde{P})U_{\tilde{P}}\|_{2 \rightarrow \infty} \\
 &\leq C \frac{m^{1/2} \log^{1/2} mn}{n^{1/2}}, \quad [\text{Lemmas 21, 24}] \\
 \|H_2\|_{2 \rightarrow \infty} &\leq \|VS_{\mathfrak{M}}^{-1/2} - S_{\tilde{P}}^{-1/2}V\| \cdot \|(\mathfrak{M} - \tilde{P})U_{\tilde{P}}\|_{2 \rightarrow \infty} \\
 &\leq C \frac{m^{3/2} \log^{3/2} mn}{n^{3/2}}, \quad [\text{Lemmas 23, 24}] \\
 \|H_3\|_{2 \rightarrow \infty} &\leq \|U_{\tilde{P}}\|_{2 \rightarrow \infty} \|U_{\tilde{P}}^T(\mathfrak{M} - \tilde{P})U_{\tilde{P}}\| \cdot \|S_{\mathfrak{M}}^{-1/2}\| \\
 &\leq C \frac{m \log mn}{n}, \quad [\text{Lemma 21; Eq. 26,30}] \\
 \|H_4\|_{2 \rightarrow \infty} &\leq \|(I - U_{\tilde{P}}U_{\tilde{P}}^T)\|_{2 \rightarrow \infty} \|(\mathfrak{M} - \tilde{P})\| \|Q^{(1)}\| \|S_{\mathfrak{M}}^{-1/2}\| \\
 &\leq C \frac{m^{1/2} \log mn}{n^{1/2}}. \quad [\text{Lemmas 21, 20, 25 }].
 \end{aligned}$$

Next, by Lemmas 21, 22 and Eq. 26 we get

$$\|Q^{(2)}S_{\mathfrak{M}}^{1/2}\|_{2 \rightarrow \infty} \leq \|U_{\tilde{P}}\|_{2 \rightarrow \infty} \|U_{\tilde{P}}^T U_{\mathfrak{M}} - V\| \cdot \|S_{\mathfrak{M}}^{1/2}\| \leq \frac{C \log mn}{n},$$

and by Lemma 23 and Eq (26) we get

$$\|U_{\tilde{P}}(VS_{\mathfrak{M}}^{1/2} - S_{\tilde{P}}^{1/2}V)\|_{2 \rightarrow \infty} \leq \|U_{\tilde{P}}\|_{2 \rightarrow \infty} \|(VS_{\mathfrak{M}}^{1/2} - S_{\tilde{P}}^{1/2}V)\| \leq \frac{Cm \log mn}{n}.$$

Hence, by triangle inequality,

$$\|H_5\|_{2 \rightarrow \infty} \leq C \frac{m \log mn}{n}.$$

Plugging into Eq. 34 the bounds above, for sufficiently large n , applying triangle inequality yields the desired result. \blacksquare

A.3 Central Limit Theorem for the Rows of the Generalized Omnibus Embedding

To prove Theorem 14 from Section 4, we adapt the proof of Theorem 1 in Levin et al. (2017) and further we extend it to correlated networks. The main difficulty in this adaptation is the more complex structure of the general \mathfrak{M} , which requires a number of modifications that we have already completed.

- Account for the model correlation in Lemmas 20, 23 and 24.
- The proofs of the Bernstein matrix concentration result and Lemma 23 necessitate a more delicate decomposition of \mathfrak{M} in order to leverage classical concentration inequality results. Moreover, the proof of Lemma 20 highlights how the coefficient matrices $C^{(l)}$ of the adjacency matrices $A^{(l)}$ in \mathfrak{M} fully characterize the omnibus matrix \mathfrak{M} .
- We adapt the general exchangeability result in the proof of Lemma 5 of Levin et al. (2017) (used there to bound the term ‘ B_1 ’) to our current setting, and the general block-form of \mathfrak{M} still allows for a weaker (within block) exchangeability argument to be employed, which is sufficient for our purposes.
- The model correlation and the weights in the \mathfrak{M} matrix necessitate novel decompositions to compute the relevant covariance structures.
- Considering row-wise differences of \mathfrak{M} is a key contribution to the literature.

The overall layout of the proof is as follows. Let $\widehat{\mathbf{X}}_{\mathfrak{M}_n} = \text{ASE}(\mathfrak{M}_n, d) = U_{\mathfrak{M}_n} S_{\mathfrak{M}_n}^{1/2} \in \mathbb{R}^{mn \times d}$ and let $\mathbf{Z}_n := [\mathbf{X}_n^T | \mathbf{X}_n^T | \dots | \mathbf{X}_n^T]^T \in \mathbb{R}^{mn \times d}$. Fix indices $i \in [n]$ and $s \in [m]$, and let $h = n(s-1) + i$. The quantity of interest is the h -th row (i.e., i -th row from s -th block) of

$$n^{1/2} \left(\widehat{\mathbf{X}}_{\mathfrak{M}_n} V_n^T W_n - \mathbf{Z}_n \right),$$

where $V_n, W_n \in \mathbb{R}^{d \times d}$ are suitable orthogonal transformations and W_n is such that $\mathbf{Z}_n = U_{\tilde{P}_n} S_{\tilde{P}_n}^{1/2} W_n$. This quantity can be decomposed into a sum of matrices (Eq. 34) as (dropping the subscripted dependence on n)

$$\begin{aligned} n^{1/2} \left(U_{\mathfrak{M}} S_{\mathfrak{M}}^{1/2} - U_{\tilde{P}} S_{\tilde{P}}^{1/2} V \right)_h V^T W &= n^{1/2} \left((\mathfrak{M} - \tilde{P}) U_{\tilde{P}} S_{\tilde{P}}^{-1/2} V \right)_h V^T W + (n^{1/2} R_h V) V^T W \\ &= n^{1/2} \left((\mathfrak{M} - \tilde{P}) U_{\tilde{P}} S_{\tilde{P}}^{-1/2} W \right)_h + n^{1/2} R_h W, \end{aligned}$$

where $R_h \in \mathbb{R}^{mn \times d}$ is the residual matrix. Next, Theorem 26 establishes that

$$n^{1/2} \left((\mathfrak{M} - \tilde{P}) U_{\tilde{P}} S_{\tilde{P}}^{-1/2} W \right)_h$$

converges in distribution to a mixture of normals.

We begin with a more limited conditional central limit theorem. Recall that, by the definition of the JRDPG, the latent positions of the expected omnibus matrix $\mathbb{E}\mathfrak{M} = \tilde{P} = U_{\tilde{P}}S_{\tilde{P}}U_{\tilde{P}}^T$ are given by

$$Z^* = \begin{bmatrix} X^* \\ X^* \\ \vdots \\ X^* \end{bmatrix} = U_{\tilde{P}}S_{\tilde{P}}^{1/2} \in \mathbb{R}^{mn \times d}.$$

Recall that the matrix of the true latent positions is denoted by $Z = [X^T X^T \cdots X^T]^T \in \mathbb{R}^{mn \times d}$, so that $Z = Z^*W$ for some suitable-chosen orthogonal matrix W .

Theorem 26 *With notation and assumptions as in Theorem 14, fix some $i \in [n]$ and some $s \in [m]$ and let $h = n(s-1) + i$. Conditional on $X_i = x_i \in \mathbb{R}^d$, there exists a sequence of d -by- d orthogonal matrices $\{W_n\}_n$ such that*

$$n^{1/2} \left[(\mathfrak{M}_n - \tilde{P}_n) U_{\tilde{P}_n} S_{\tilde{P}_n}^{-1/2} \right]_h W_n \xrightarrow{\mathcal{L}} \mathcal{N}(0, \check{\Sigma}(x_i)),$$

where

$$\check{\Sigma}_\rho(x_i; s) = \frac{1}{m^2} \left(\underbrace{\sum_{q=1}^m \alpha^2(s, q)}_{\text{method coefficient}} + 2 \underbrace{\sum_{q < l} \alpha(s, q) \alpha(s, l) \rho_{q,l}}_{\text{model coefficient}} \right) \underbrace{\Delta^{-1} \mathbb{E}[x_i^T X_j (1 - x_i^T X_j) X_j X_j^T] \Delta^{-1}}_{:= \Sigma(x_i)}$$

is a covariance matrix that depends on x_i .

Proof Following the proof of Lemma 6 in Levin et al. (2017), for each $n = 1, 2, \dots$, choose orthogonal $W_n \in \mathbb{R}^{d \times d}$ so that $\mathbf{Z} = Z^*W_n = U_{\tilde{P}}S_{\tilde{P}}^{1/2}W_n$. Dropping the explicit dependence on n , we rewrite the term $n^{1/2} \left[(\mathfrak{M} - \tilde{P}) U_{\tilde{P}} S_{\tilde{P}}^{-1/2} \right]_h W$ as follows

$$\begin{aligned} n^{1/2} \left[(\mathfrak{M} - \tilde{P}) U_{\tilde{P}} S_{\tilde{P}}^{-1/2} \right]_h W &= n^{1/2} \left[(\mathfrak{M} - \tilde{P}) U_{\tilde{P}} S_{\tilde{P}}^{-1/2} W \right]_h \\ &= n^{1/2} \left[(\mathfrak{M} - \tilde{P}) U_{\tilde{P}} S_{\tilde{P}}^{1/2} W W^T S_{\tilde{P}}^{-1} W \right]_h \\ &= n^{1/2} \left[(\mathfrak{M} - \tilde{P}) \mathbf{Z} \right]_h W^T S_{\tilde{P}}^{-1} W \\ &= \frac{n^{-1/2}}{m} \left[\mathfrak{M} \mathbf{Z} - \tilde{P} \mathbf{Z} \right]_h [n W^T S_{\tilde{P}}^{-1} W]. \end{aligned} \quad (35)$$

Consider next the scaled rows of $\mathfrak{M} \mathbf{Z} - \tilde{P} \mathbf{Z}$. We write (where for a matrix A , either $(A)_k$ or $[A]_k$ are used to denote the k -th row of A and $(A)_{hk}$ or $[A]_{hk}$ denoting A_{hk})

$$\begin{aligned} \frac{n^{-1/2}}{m} [\mathfrak{M} \mathbf{Z} - \tilde{P} \mathbf{Z}]_h &= \frac{n^{-1/2}}{m} \sum_{k=1}^{mn} (\mathfrak{M} - \tilde{P})_{hk} (\mathbf{Z})_k \\ &= \frac{n^{-1/2}}{m} \sum_{\ell=1}^m \sum_{j=1}^n \left(\mathfrak{M}_{i,j}^{(s,\ell)} - P_{ij} \right) X_j \\ &= \frac{n^{-1/2}}{m} \sum_{j \neq i} \left(\sum_{\ell=1}^m \mathfrak{M}_{i,j}^{(s,\ell)} - m P_{ij} \right) X_j - n^{-1/2} P_{ii} X_i. \end{aligned}$$

Conditioning on $X_i = x_i \in \mathbb{R}^d$, we first observe that

$$\frac{P_{ii}}{n^{1/2}} X_i = \frac{x_i^T x_i}{n^{1/2}} x_i \xrightarrow{a.s.} 0.$$

Moreover, the remaining portion of the scaled sum becomes

$$\begin{aligned} & n^{-1/2} \sum_{j \neq i} \frac{1}{m} \left(\sum_{\ell=1}^m \mathfrak{M}_{i,j}^{(s,\ell)} - m(x_i^T X_j) \right) X_j \\ &= n^{-1/2} \sum_{j \neq i} \frac{1}{m} \left(\sum_{\ell=1}^m \left(\sum_{q=1}^m c_q^{(s,\ell)} A_{i,j}^{(q)} \right) - m(x_i^T X_j) \right) X_j \\ &= n^{-1/2} \sum_{j \neq i} \frac{1}{m} \left(\sum_{q=1}^m \underbrace{\left(\sum_{\ell=1}^m c_q^{(s,\ell)} \right)}_{:=\alpha(s,q)} A_{i,j}^{(q)} - m(x_i^T X_j) \right) X_j, \end{aligned}$$

where $\sum_{q=1}^m \alpha(s,q) = m$ for all $s \in [m]$. The above expression is a sum of $n-1$ independent 0-mean random variables, the

$$\frac{1}{m} \left(\sum_{q=1}^m \alpha(s,q) A_{i,j}^{(q)} - m(x_i^T X_j) \right) X_j,$$

each with covariance matrix denoted by $\dot{\Sigma}(x_i)$, computed as follows (suppressing the conditioning on $X_i = x_i$)

$$\begin{aligned} \dot{\Sigma}(x_i) &= \frac{1}{m^2} \mathbb{E} \left[\left(\sum_{q=1}^m \alpha(s,q) A_{i,j}^{(q)} - m(x_i^T X_j) \right)^2 X_j X_j^T \right] \\ &= \frac{1}{m^2} \mathbb{E} \left[\mathbb{E} \left[\left(\sum_{q=1}^m \alpha(s,q) A_{i,j}^{(q)} - m(x_i^T X_j) \right)^2 X_j X_j^T \mid X_j \right] \right] \\ &= \frac{1}{m^2} \mathbb{E} \left[\left(\mathbb{E} \left[\left(\sum_{q=1}^m \alpha(s,q) A_{i,j}^{(q)} \right)^2 \mid X_j \right] - m^2 (x_i^T X_j)^2 \right) X_j X_j^T \right] \\ &= \frac{1}{m^2} \mathbb{E} \left[\left(\left(\sum_{q=1}^m \alpha^2(s,q) \right) x_i^T X_j + \left(2 \sum_{q < l} \alpha(s,q) \alpha(s,l) - m^2 \right) (x_i^T X_j)^2 \right. \right. \\ &\quad \left. \left. + \left(2 \sum_{q < l} \alpha(s,q) \alpha(s,l) \rho_{q,l} \right) x_i^T X_j (1 - x_i^T X_j) \right) X_j X_j^T \right] \\ &= \frac{1}{m^2} \left(\sum_{q=1}^m \alpha^2(s,q) + 2 \sum_{q < l} \alpha(s,q) \alpha(s,l) \rho_{q,l} \right) \mathbb{E} \left[(x_i^T X_j - (x_i^T X_j)^2) X_j X_j^T \right], \end{aligned}$$

where the third equality holds because conditioning on $P = X X^T$,

$$\mathbb{E} \left[\sum_{q=1}^m \alpha(s,q) A^{(q)} \mid P \right] = m X X^T.$$

The fourth equality holds from the fact that conditioning on X_i and X_j ,

$$\mathbb{E}[(A_{i,j}^{(q)})^2 | X_i, X_j] = X_i^T X_j$$

for all $q \in [m]$, and from the fact that conditioning on X_i and X_j , $(A_{i,j}^{(q)}, A_{i,j}^{(l)})$ are $\rho_{q,l}$ -correlated Bernoulli($X_i^T X_j$) random variables,

$$\mathbb{E}[A_{i,j}^{(q)} A_{i,j}^{(l)} | X_i, X_j] = (X_i^T X_j)^2 + \rho_{q,l} X_i^T X_j (1 - X_i^T X_j)$$

for all $q \neq l$. Lastly, the final equality holds by observing

$$m^2 = \left(\sum_{q=1}^m \alpha(s, q) \right)^2 = \sum_{q=1}^m \alpha(s, q)^2 + 2 \sum_{q < l} \alpha(s, q) \alpha(s, l).$$

Thus, by the multivariate central limit theorem we have that

$$n^{-1/2} \sum_{j \neq i} \frac{1}{m} \left(\sum_{q=1}^m \mathfrak{M}_{i,j}^{(s,q)} - m(x_i^T X_j) \right) X_j \xrightarrow{\mathcal{L}} \mathcal{N}(0, \check{\Sigma}(x_i)). \quad (36)$$

Next, recall (as in the proof of Theorem 9) $nW_n^T S_P^{-1} W_n \xrightarrow{a.s.} \Delta^{-1}$, so that by the multivariate version of Slutsky's theorem in Eq. 35 we get

$$n^{1/2} \left[(\mathfrak{M} - \tilde{P}) U_{\tilde{P}} S_{\tilde{P}}^{-1/2} \right]_h W_n \xrightarrow{\mathcal{L}} \mathcal{N}(0, \check{\Sigma}_\rho(x_i; s)),$$

where $\check{\Sigma}_\rho(x_i; s) = \Delta^{-1} \check{\Sigma}(x_i) \Delta^{-1}$. ■

This equips us for the proof of Theorem 14.

Proof (Central Limit Theorem for the rows of the generalized omnibus matrix) Fix the index $h \in [mn]$ as $h = n(s-1) + i$, where $s \in [m]$, $i \in [n]$. Recall that for each $n = 1, 2, \dots$, we choose orthogonal $W_n \in \mathbb{R}^{d \times d}$ so that $\mathbf{Z} = \mathbf{Z}^* W_n = U_{\tilde{P}} S_{\tilde{P}}^{1/2} W_n$. With V_n defined as in Lemma 22, the matrix difference

$$n^{1/2} \left(U_{\mathfrak{M}_n} S_{\mathfrak{M}_n}^{1/2} - U_{\tilde{P}_n} S_{\tilde{P}_n}^{1/2} V_n \right)_h V_n^T W_n$$

can be decomposed into a sum of matrices (as shown in Eq. 34 of the proof of Theorem 13) as follows

$$n^{1/2} \left(U_{\mathfrak{M}_n} S_{\mathfrak{M}_n}^{1/2} - U_{\tilde{P}_n} S_{\tilde{P}_n}^{1/2} V_n \right)_h V_n^T W_n = n^{1/2} \left((\mathfrak{M}_n - \tilde{P}_n) U_{\tilde{P}_n} S_{\tilde{P}_n}^{-1/2} W_n \right)_h + n^{1/2} R_{h,n} V_n W_n,$$

where $R_{h,n} \in \mathbb{R}^{mn \times d}$ is the h -th row of the residual matrix defined as $R = H_2 - H_3 + H_4 + H_5$. Further, in the proof of Theorem 13 it is shown that for any $i = 2, 3, 5$ we have

$$n^{1/2} \|H_i\|_{2 \rightarrow \infty} \leq \frac{Cm \log mn}{n^{1/2}} \quad \text{w.h.p.}$$

It remains to provide an appropriate bound for H_4 .

Rather than showing $\sqrt{n}\|H_4\|_{2 \rightarrow \infty}$ converges to 0 in probability, we will prove directly that $\sqrt{n}\|(H_4)_h\|_2$ converges to 0 in probability, which is sufficient. Adapting the exchangeability bound on the analogous term from Levin et al. (2017), we recall that (recalling that I_{\pm} is the random diagonal sign matrix designed to give $\mathfrak{M}U_{\mathfrak{M}} = U_{\mathfrak{M}}S_{\mathfrak{M}}I_{\pm}$)

$$\begin{aligned} H_4 &:= (I - U_{\tilde{P}}U_{\tilde{P}}^T)(\mathfrak{M} - \tilde{P})Q^{(1)}S_{\mathfrak{M}}^{-1/2}I_{\pm} \\ &= (I - U_{\tilde{P}}U_{\tilde{P}}^T)(\mathfrak{M} - \tilde{P})(U_{\mathfrak{M}} - U_{\tilde{P}}U_{\tilde{P}}^T U_{\mathfrak{M}})S_{\mathfrak{M}}^{-1/2}I_{\pm} \\ &\quad + (I - U_{\tilde{P}}U_{\tilde{P}}^T)(\mathfrak{M} - \tilde{P})(U_{\tilde{P}}U_{\tilde{P}}^T U_{\mathfrak{M}} - U_{\tilde{P}}V)S_{\mathfrak{M}}^{-1/2}I_{\pm}. \end{aligned}$$

Denote the first term in this expression via H_{41} (the second via H_{42}), and note that

$$\begin{aligned} H_{41} &:= (I - U_{\tilde{P}}U_{\tilde{P}}^T)(\mathfrak{M} - \tilde{P})(U_{\mathfrak{M}} - U_{\tilde{P}}U_{\tilde{P}}^T U_{\mathfrak{M}})S_{\mathfrak{M}}^{-1/2}I_{\pm} \\ &= \underbrace{(I - U_{\tilde{P}}U_{\tilde{P}}^T)(\mathfrak{M} - \tilde{P})(I - U_{\tilde{P}}U_{\tilde{P}}^T)U_{\mathfrak{M}}U_{\mathfrak{M}}^T}_{:=E_1}(U_{\mathfrak{M}}S_{\mathfrak{M}}^{-1/2}I_{\pm}) \\ H_{42} &:= (I - U_{\tilde{P}}U_{\tilde{P}}^T)(\mathfrak{M} - \tilde{P})(U_{\tilde{P}}U_{\tilde{P}}^T U_{\mathfrak{M}} - U_{\tilde{P}}V)S_{\mathfrak{M}}^{-1/2}I_{\pm} \\ &= (I - U_{\tilde{P}}U_{\tilde{P}}^T)(\mathfrak{M} - \tilde{P})U_{\tilde{P}}(U_{\tilde{P}}^T U_{\mathfrak{M}} - U_{\tilde{P}})V S_{\mathfrak{M}}^{-1/2}I_{\pm}. \end{aligned}$$

Considering first H_{41} , we have that

$$\|(H_{41})_h\|_2 \leq \|(E_1)_h\|_2 \|U_{\mathfrak{M}}S_{\mathfrak{M}}^{-1/2}\|.$$

Consider now the term $\|(E_1)_h\|_2$. For any symmetric matrix $B \in \mathbb{R}^{k \times k}$, we define $\Pi_d(B)$ to be the orthogonal projection onto the eigenspace corresponding to the eigenvectors of B with the d largest (in magnitude) eigenvalues. Similarly, let $\Pi_d^{\perp}(B)$ denote the orthogonal projection onto $(\Pi_d(B))^{\perp}$.

Note that for any permutation matrix Q , we have that

$$\begin{aligned} \Pi_d(QBQ^T) &= Q\Pi_d(B)Q^T \\ \Pi_d^{\perp}(QBQ^T) &= Q\Pi_d^{\perp}(B)Q^T. \end{aligned}$$

As in Levin et al. (2017), for any $(B, H) \in \mathbb{R}^{k \times k} \times \mathbb{R}^{k \times k}$, define the operator $\mathcal{L}(B, H)$ via:

$$\mathcal{L}(B, H) = \Pi_d^{\perp}(H)(B - H)\Pi_d^{\perp}(H)\Pi_d(B).$$

Let us consider permutations $\tilde{Q} \in \mathbb{R}^{mn \times mn}$ of the form

$$\tilde{Q} = I_m \otimes Q$$

for permutations $Q \in \mathbb{R}^{n \times n}$. Note that, for such \tilde{Q} , we have that

$$\tilde{Q}\mathfrak{M}\tilde{Q}^T = \begin{pmatrix} Q\mathfrak{M}^{(1,1)}Q^T & Q\mathfrak{M}^{(1,2)}Q^T & \dots & Q\mathfrak{M}^{(1,m)}Q^T \\ Q\mathfrak{M}^{(2,1)}Q^T & Q\mathfrak{M}^{(2,2)}Q^T & \dots & Q\mathfrak{M}^{(2,m)}Q^T \\ \vdots & \vdots & \ddots & \vdots \\ Q\mathfrak{M}^{(m,1)}Q^T & Q\mathfrak{M}^{(m,2)}Q^T & \dots & Q\mathfrak{M}^{(m,m)}Q^T \end{pmatrix},$$

and that for each (i, j) pair,

$$Q\mathfrak{M}^{(i,j)}Q^T = \sum_{\ell} c_{\ell}^{(i,j)} QA^{(\ell)}Q^T.$$

Similarly,

$$\tilde{Q}\tilde{P}\tilde{Q}^T = J_m \otimes (QPQ^T).$$

Note that $\mathcal{L}(\mathfrak{M}, \tilde{P}) = E_1$, and as these orthogonal projections are unique, we have

$$\begin{aligned} & \mathcal{L}(\tilde{Q}\mathfrak{M}\tilde{Q}^T, \tilde{Q}\tilde{P}\tilde{Q}^T) \\ &= \tilde{Q}(I - U_{\tilde{P}}U_{\tilde{P}}^T)\tilde{Q}^T\tilde{Q}(\mathfrak{M} - \tilde{P})\tilde{Q}^T\tilde{Q}(I - U_{\tilde{P}}U_{\tilde{P}}^T)\tilde{Q}^T\tilde{Q}(U_{\mathfrak{M}}U_{\mathfrak{M}}^T)\tilde{Q}^T \\ &= \tilde{Q}E_1\tilde{Q}^T. \end{aligned} \quad (37)$$

Since the rows of \mathbf{X} are i.i.d. and the correlation across a given pair of graphs is the same for all edge pairs (i.e., $\text{corr}(A_{i,j}^{(k)}, A_{i,j}^{(\ell)}) = R_{k,\ell}$ independent of i and j), together this implies that the matrix-pair entries of $(\mathfrak{M}, \tilde{P})$ are equal in distribution to those of $(\tilde{Q}\mathfrak{M}\tilde{Q}^T, \tilde{Q}\tilde{P}\tilde{Q}^T)$. Therefore, the entries of $\mathcal{L}(\tilde{Q}\mathfrak{M}\tilde{Q}^T, \tilde{Q}\tilde{P}\tilde{Q}^T) = \tilde{Q}E_1\tilde{Q}^T$ are equal in law to those of $\mathcal{L}(\mathfrak{M}, \tilde{P}) = E_1$. Therefore, for each row i we have that

$$\|(\tilde{Q}E_1)_i\|^2 = \|(\tilde{Q}E_1\tilde{Q}^T)_i\|^2 \stackrel{\mathcal{L}}{=} \|(E_1)_i\|^2.$$

This implies that if $Q_{i,j} = 1$ then for any $k \in [m]$,

$$\mathbb{E}(\|(E_1)_{(k-1)n+i}\|^2) = \mathbb{E}(\|(\tilde{Q}E_1)_{(k-1)n+i}\|^2) = \mathbb{E}(\|(E_1)_{(k-1)n+j}\|^2). \quad (38)$$

This guarantees that

$$\mathbb{E}(\|(E_1)_{(k-1)n+i}\|^2)$$

depends only on k and not on i , and note that the analogous result follows immediately for

$$\mathbb{E}(\|(\tilde{Q}E_1)_{(k-1)n+i}\|^2).$$

We can then define for $i, j \in [n]$,

$$r_k := \mathbb{E}(\|(\tilde{Q}E_1)_{(k-1)n+i}\|^2) = \mathbb{E}(\|(E_1)_{(k-1)n+j}\|^2).$$

Observe that

$$\mathbb{E}(\|E_1\|_F^2) = \sum_{k=1}^m nr_k \geq n \max_k r_k.$$

Because $h = (s-1)n + i$, for $i \in [n]$, Eq. 38 and an application of Markov's inequality yield

$$\begin{aligned} \mathbb{P}(\sqrt{n}\|(E_1)_h\| > t) &\leq \frac{n\mathbb{E}(\|(E_1)_h\|^2)}{t^2} = \frac{nr_s}{t^2} \\ &\leq \frac{\mathbb{E}(\|E_1\|_F^2)}{t^2}. \end{aligned}$$

The entries of $\mathfrak{M} - \tilde{P}$ are bounded between $[-1, 1]$, and $(I - U_{\tilde{P}}U_{\tilde{P}}^T)U_{\mathfrak{M}}U_{\mathfrak{M}}^T$ is rank d (with spectral norm bounded by 1). Hence, globally

$$\|E_1\|_F^2 \leq \|I - U_{\tilde{P}}U_{\tilde{P}}^T\|^2 \|\mathfrak{M} - \tilde{P}\|_F^2 \|(I - U_{\tilde{P}}U_{\tilde{P}}^T)U_{\mathfrak{M}}U_{\mathfrak{M}}^T\|_F^2 \leq m^2 n^2 d^2.$$

Now, we also have that, with high probability (so that the bad set has probability bounded above by Cn^{-2}), there exists a constant $C > 0$ such that

$$\begin{aligned} \|E_1\|_F^2 &\leq d^2 \|I - U_{\tilde{P}}U_{\tilde{P}}^T\|^2 \|\mathfrak{M} - \tilde{P}\|^2 \|(I - U_{\tilde{P}}U_{\tilde{P}}^T)U_{\mathfrak{M}}\|^2 \|U_{\mathfrak{M}}^T\|_F^2 \\ &\leq C \frac{m^2 n \log^2 mn}{n} = Cm^2 \log^2 mn \quad \text{by Lemma 20, and Eq. 31.} \end{aligned}$$

We then have that there exists a constant $C > 0$ such that

$$\mathbb{E}(\|E_1\|_F^2) \leq C(m^2 \log^2 mn + n^{-2} m^2 n^2 d^2) = C(m^2 \log^2 mn + m^2 d^2).$$

Letting $t = n^{1/4}$ in our Markov bound, we see that $\mathbb{P}(\sqrt{n}\|(E_1)_h\| > n^{1/4}) \rightarrow 0$. As, w.h.p., we have that $\|U_{\mathfrak{M}}S_{\mathfrak{M}}^{-1/2}\| \leq C/\sqrt{mn}$, we have that

$$n^{1/2}\|(H_{41})_h\|_2 \leq n^{1/3}\|(E_1)_h\| n^{1/6}\|U_{\mathfrak{M}}S_{\mathfrak{M}}^{-1/2}\| \xrightarrow{P} 0.$$

Turning our attention to H_{42} , we have that (w.h.p.)

$$\begin{aligned} \|H_{42}\|_F &\leq \|(I - U_{\tilde{P}}U_{\tilde{P}}^T)(\mathfrak{M} - \tilde{P})(U_{\tilde{P}}U_{\tilde{P}}^T U_{\mathfrak{M}} - U_{\tilde{P}}V)S_{\mathfrak{M}}^{-1/2}I_{\pm}\|_F \\ &\leq \|(I - U_{\tilde{P}}U_{\tilde{P}}^T)\| \cdot \|\mathfrak{M} - \tilde{P}\| \cdot \|U_{\tilde{P}}\| \cdot \|U_{\tilde{P}}^T U_{\mathfrak{M}} - V\|_F \|S_{\mathfrak{M}}^{-1/2}\| \\ &\leq C(m\sqrt{n \log mn}) \left(\frac{\log mn}{n}\right) \left(\frac{1}{(mn)^{1/2}}\right) \leq C \frac{m^{1/2} \log^{3/2} mn}{n}, \end{aligned}$$

where the last line follows from Lemmas 21, 20, and 22. Therefore, $n^{1/2}H_{42}$ converges to 0 in probability, and combined this yields that the h -th row of $n^{1/2}H_4$ converges to 0 as desired.

Since the matrix W_n is unitary, the bounds above imply that the term $n^{1/2}R_{h,n}W_n$ converges to 0 in probability. Moreover, by Lemma 26 we have

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[n^{1/2} \left((\mathfrak{M} - \tilde{P})U_{\tilde{P}}S_{\tilde{P}}^{-1/2} \right)_h W_n \leq x \right] \xrightarrow{\mathcal{D}} \int_{\text{supp}F} \Phi(x, \check{\Sigma}_{\rho}(y; s)) dF(y),$$

by integrating over the latent positions X_i . Finally, an application of Slutsky's theorem completes the proof. \blacksquare

A.4 Limiting Correlation for the General Omnibus Embedding

Here, we supply details for the computation of the limiting correlation across rows of the general omnibus embedding; this is the content of Theorem 16 from Section 4.1.

Proof We mimic the proof of Theorem 9 here, and so omit some detail. Fix some $i \in [n]$ and some $s_1, s_2 \in [m]$ and for $j = 1, 2$, let $h_j = n(s_j - 1) + i \in [mn]$. Conditioning on

$X_i = x_i$, analogous to Eq. 22, we write (where $\mathcal{Q}_n = V_n^T W_n$ as defined in the proof of Theorem 14)

$$\begin{aligned}
 & n^{1/2} \left(\left(\widehat{\mathbf{X}}_{\mathfrak{M}} V_n^T W_n \right)_{h_1} - \left(\widehat{\mathbf{X}}_{\mathfrak{M}} V_n^T W_n \right)_{h_2} \right) \\
 &= n^{1/2} \left(\left(\mathfrak{M}_n - \widetilde{P}_n \right) \mathbf{Z} \right)_{h_1} - n^{1/2} \left(\left(\mathfrak{M}_n - \widetilde{P}_n \right) \mathbf{Z} \right)_{h_2} + o_P(1) \\
 &= \frac{n^{-1/2}}{m} \left(\sum_{j \neq i} \left(\sum_{q=1}^m \mathfrak{M}_{i,j}^{(s_1,q)} - \mathfrak{M}_{i,j}^{(s_2,q)} \right) X_j \right) [n W_n^T S_P^{-1} W_n] + o_P(1) \\
 &= n^{-1/2} \left(\sum_{j \neq i} \frac{1}{m} \left(\sum_{q=1}^m \alpha(s_1, q) A_{i,j}^{(q)} - \alpha(s_2, q) A_{i,j}^{(q)} \right) X_j \right) [n W_n^T S_P^{-1} W_n] + o_P(1).
 \end{aligned}$$

Each of the $n - 1$ terms, $\frac{1}{m} \left(\sum_{q=1}^m \alpha(s_1, q) A_{i,j}^{(q)} - \alpha(s_2, q) A_{i,j}^{(q)} \right) X_j$, is an independent, mean zero, random variable, with common covariance matrix $\Phi_\rho(x_i, s_1, s_2)$. The desired result will then follow from an application of the multivariate central limit theorem and multivariate Slutsky theorems (as in the proof of Theorem 9), provided we can show the right form for $\Phi_\rho(x_i, s_1, s_2)$. To this end, we consider (suppressing the conditioning on $X_i = x_i$)

$$\begin{aligned}
 \Phi_\rho(x_i, s_1, s_2) &:= \frac{1}{m^2} \mathbb{E} \left(\left(\sum_{q=1}^m \alpha(s_1, q) A_{i,j}^{(q)} - \sum_{q=1}^m \alpha(s_2, q) A_{i,j}^{(q)} \right)^2 X_j X_j^T \right) \\
 &= \frac{1}{m^2} \left[\mathbb{E}[(x_i^T X_j) X_j X_j^T] \sum_{q=1}^m (\alpha(s_1, q) - \alpha(s_2, q))^2 \right. \\
 &\quad \left. + 2 \mathbb{E}[(x_i^T X_j)^2 X_j X_j^T] \sum_{q < l} (\alpha(s_1, q) - \alpha(s_2, q)) (\alpha(s_1, l) - \alpha(s_2, l)) \right] \\
 &\quad \left. + 2 \mathbb{E}[x_i^T X_j (1 - x_i^T X_j) X_j X_j^T] \sum_{q < l} (\alpha(s_1, q) - \alpha(s_2, q)) (\alpha(s_1, l) - \alpha(s_2, l)) \rho_{q,l} \right] \\
 &= \frac{1}{m^2} \left[\sum_{q=1}^m (\alpha(s_1, q) - \alpha(s_2, q))^2 \right. \\
 &\quad \left. + 2 \sum_{q < l} (\alpha(s_1, q) - \alpha(s_2, q)) (\alpha(s_1, l) - \alpha(s_2, l)) \rho_{q,l} \right] \Sigma(x_i) \\
 &= \frac{1}{m^2} \left(2 \sum_{q < l} (\alpha(s_1, q) - \alpha(s_2, q)) (\alpha(s_1, l) - \alpha(s_2, l)) (\rho_{q,l} - 1) \right) \Sigma(x_i), \tag{39}
 \end{aligned}$$

where $\Sigma(x_i) = \mathbb{E}[(x_i^T X_j - (x_i^T X_j)^2) X_j X_j^T]$, and the third equality follows from

$$\begin{aligned}
 0 &= \left(\sum_q \alpha(s_1, q) - \alpha(s_2, q) \right)^2 \\
 &= \sum_q (\alpha(s_1, q) - \alpha(s_2, q))^2 + 2 \sum_{q < l} (\alpha(s_1, q) - \alpha(s_2, q)) (\alpha(s_1, l) - \alpha(s_2, l)).
 \end{aligned}$$

The proof then follows mutatis mutandis as that of Theorem 9. ■

A.5 Central Limit Theorem for the Rows of the Average Embeddings

We supply details for the computation of the limiting covariance matrix of the rows of the sums of the aligned embeddings across the graphs, as described in Theorem 17 from Section 7. As in the previous theorems, the proof mimics the one of Theorem 9, and so we omit some details.

Proof To prove part a), fix some $i \in [n]$. By the same arguments than Equations (20) and (21), observe that

$$\sqrt{n} \left(\frac{1}{m} \sum_{s=1}^m \widehat{\mathbf{X}}_{A_n}^{(k)} W_n^{(s)} - \mathbf{X}_n \right)_i = \sqrt{n} \left(\frac{1}{m} \sum_{s=1}^m (A_n^{(s)} - P_n) U_P S_P^{-1/2} W_n \right)_i + O(n^{-1/2} \log n),$$

where $W^{(s)} = W^{(s,1)}(W^{(s,2)})^T$, with $W^{(s,1)}\Lambda^{(s)}(W^{(s,2)})^T = U_P^T U_{A^{(s)}}$ the singular value decomposition of the matrix on the right hand side, and W_n is a sequence of orthogonal matrices such that $U_{P_n} S_{P_n}^{1/2} W_n = \mathbf{X}_n$,

Now, observe that the first term can be written as

$$\begin{aligned} \sqrt{n} \left(\frac{1}{m} \sum_{s=1}^m (A_n^{(s)} - P_n) U_P S_P^{-1/2} W_n \right)_i &= \sqrt{n} \left(\frac{1}{m} \sum_{s=1}^m (A_n^{(s)} - P_n) \mathbf{X}_n W_n^T S_P^{-1} W_n \right)_i \\ &= \frac{1}{\sqrt{n}} \sum_{j \neq i} \left(\frac{1}{m} \sum_{s=1}^m (A_{ij}^{(s)} - P_{ij}) X_j \right) (n W_n^T S_P^{-1} W_n). \end{aligned}$$

Conditioning on $X_i = x_i$, the expression is a sum of $n - 1$ independent terms with mean zero. The covariance matrix of each term can be calculated as

$$\begin{aligned} \mathbb{E} \left(\left(\frac{1}{m} \sum_{s=1}^m (A_{ij}^{(s)} - P_{ij}) \right)^2 X_j X_j^T \right) &= \mathbb{E} \left(\frac{1}{m^2} \mathbb{E} \left(\left(\sum_{s=1}^m (A_{ij}^{(s)} - P_{ij}) \right)^2 \middle| \mathbf{X} \right) X_j X_j^T \right) \\ &= \mathbb{E} \left(\frac{1}{m^2} (m P_{ij} - m^2 P^2 + m(m-1) P_{ij} (P_{ij} + \rho(1 - P_{ij}))) X_j X_j^T \right) \\ &= \left(\frac{1 - \rho}{m} + \rho \right) \mathbb{E} (P_{ij} (1 - P_{ij}) X_j X_j^T) \\ &= \tilde{\Sigma} \left(x_i, \frac{1 - \rho}{m} + \rho \right). \end{aligned}$$

As in Theorem 9, the classical multivariate central limit theorem gives the desired result.

To prove part b), we follow similar arguments to the proofs of Theorems 9 and 16. Fix some $i \in [n]$, and observe that

$$\sqrt{n} \left(\frac{1}{m} \sum_{s=1}^m \widehat{\mathbf{X}}_{M_n}^{(s)} \tilde{W}_n - \mathbf{X}_n \right)_i = \frac{\sqrt{n}}{m} \sum_{s=1}^m \left((M_n - \tilde{P}_n) \mathbf{Z}_n (\tilde{W}_n^T S_P^{-1} \tilde{W}_n) \right)_{n(s-1)+i} + o_P(1),$$

where $\tilde{W}_n = V_n^T W_n$ as defined in the proof of Theorem 10. The first term can be written as

$$\begin{aligned}
 & \frac{n^{1/2}}{m} \sum_{s=1}^m \left((M_n - \tilde{P}_n) \mathbf{Z}_n \tilde{W}_n^T S_P^{-1} \tilde{W}_n \right)_{n(s-1)+i} \\
 &= \frac{\sqrt{n}}{m} \sum_{s=1}^m \sum_{t=1}^m \frac{1}{m} \left((M_n^{(s,t)} - P_n) \mathbf{X}_n (\tilde{W}_n^T S_P^{-1} \tilde{W}_n) \right)_i \\
 &= \frac{\sqrt{n}}{m} \sum_{j \neq i} \sum_{s=1}^m \sum_{t=1}^m \frac{1}{m} \left((M_{ij}^{(s,t)} - P_{ij}) X_j \right) (\tilde{W}_n^T S_P^{-1} \tilde{W}_n) \\
 &= \frac{\sqrt{n}}{m} \sum_{j \neq i} \sum_{s,t} \frac{1}{m} \left(\left(\frac{1}{2} (A_{ij}^{(s)} + A_{ij}^{(t)}) - P_{ij} \right) X_j \right) (\tilde{W}_n^T S_P^{-1} \tilde{W}_n) \\
 &= \frac{\sqrt{n}}{m} \sum_{j \neq i} \sum_{s,t} \left(\frac{1}{2m} (A_{ij}^{(s)} - P_{ij}) + \frac{1}{2m} (A_{ij}^{(t)} - P_{ij}) \right) X_j (\tilde{W}_n^T S_P^{-1} \tilde{W}_n) \\
 &= \frac{\sqrt{n}}{m} \sum_{j \neq i} \left(\frac{1}{2m} \sum_t \sum_s (A_{ij}^{(s)} - P_{ij}) + \frac{1}{2m} \sum_s \sum_t (A_{ij}^{(t)} - P_{ij}) \right) X_j (\tilde{W}_n^T S_P^{-1} \tilde{W}_n) \\
 &= \frac{\sqrt{n}}{m} \sum_{j \neq i} \left(\frac{1}{2} \sum_s (A_{ij}^{(s)} - P_{ij}) + \frac{1}{2} \sum_t (A_{ij}^{(t)} - P_{ij}) \right) X_j (\tilde{W}_n^T S_P^{-1} \tilde{W}_n) \\
 &= \frac{1}{\sqrt{n}} \sum_{j \neq i} \left(\frac{1}{m} \sum_{s=1}^m (A_{ij}^{(s)} - P_{ij}) \right) X_j (n \tilde{W}_n^T S_P^{-1} \tilde{W}_n).
 \end{aligned}$$

The covariance of each of the terms is calculated in the same way as in part a), and the result follows by similar arguments. \blacksquare

References

- Arash A Amini, Marina S Paez, and Lizhen Lin. Hierarchical stochastic block model for community detection in multiplex networks. *arXiv preprint arXiv:1904.05330*, 2019.
- J. Arroyo, A. Athreya, J. Cape, G. Chen, C. E. Priebe, and J. T. Vogelstein. Inference for multiple heterogeneous networks with a common invariant subspace. *Journal of Machine Learning Research*, accepted for publication, 2020.
- D. Asta and C. Shalizi. Geometric network comparison. Arxiv preprint at <http://arxiv.org/abs/1411.1350>, 2014.
- A. Athreya, V. Lyzinski, D. J. Marchette, C. E. Priebe, D. L. Sussman, and M. Tang. A limit theorem for scaled eigenvectors of random dot product graphs. *Sankhya A*, 78:1–18, 2016.

- A. Athreya, D. E. Fishkind, K. Levin, , V. Lyzinski, Y. Park, Y. Qin, D. L. Sussman, M. Tang, J. T. Vogelstein, and C. E. Priebe. Statistical inference on random dot product graphs: a survey. *Journal of Machine Learning Research*, 18, 2018a.
- A. Athreya, M. Tang, Y. Park, and C. E. Priebe. On estimation and inference in latent structure random graphs. To appear in *Statistical Science*; Arxiv preprint at <http://arxiv.org/abs/1806.01401>, 2018b.
- S. Babkin, J. R. Stewart, X. Long, and M. Schweinberger. Large-scale estimation of random graph models with local dependence. *Computational statistics & data analysis*, 152: 107029, 2020.
- Afonso S Bandeira, Ramon Van Handel, et al. Sharp nonasymptotic bounds on the norm of random matrices with independent entries. *The Annals of Probability*, 44(4):2479–2506, 2016.
- R. Bhatia. *Matrix Analysis*. Springer, 1997.
- S. Bhattacharyya, S. Chatterjee, and S. S. Mukherjee. Consistent detection and optimal localization of all detectable change points in piecewise stationary arbitrarily sparse network-sequences. *arXiv preprint arXiv:2009.02112*, 2020.
- S. S. Bhowmick and B. S. Seah. Clustering and summarizing protein-protein interaction networks: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 28(3): 638–658, 2015.
- I. Borg and P. J. F. Groenen. *Modern multidimensional scaling: Theory and applications*. Springer Science & Business Media, 2005.
- D. S. Callaway, M. E. J. Newman, S. H. Strogatz, and D. J. Watts. Network robustness and fragility: Percolation on random graphs. *Physical review letters*, 85(25):5468, 2000.
- J. Cape, M. Tang, and C. E. Priebe. The two-to-infinity norm and singular subspace geometry with applications to high-dimensional statistics. *Annals of Statistics*, 2018. Arxiv preprint at <http://arxiv.org/abs/1705.10735>.
- P. J. Carrington, J. Scott, and S. Wasserman. *Models and methods in social network analysis*, volume 28. Cambridge university press, 2005.
- C. Castellano, S. Fortunato, and V. Loreto. Statistical physics of social dynamics. *Reviews of modern physics*, 81(2):591, 2009.
- S. Chatterjee. Matrix estimation by universal singular value thresholding. *Annals of Statistics*, 43:177–214, 2015.
- G. Chen, J. Arroyo, A. Athreya, J. Cape, J. Vogelstein, Y. Park, C. White, J. Larson, W. Yang, and C. Priebe. Multiple network embedding for anomaly detection in time series of graphs. *arXiv preprint arXiv:2008.10055*, 2020.
- L. Chen, J.T. Vogelstein, V. L. Lyzinski, and C. E. Priebe. A joint graph inference case study: the *C. Elegans* chemical and electrical connectomes. *Worm*, 5:e1142041, 2016.

- K. L. Chung. *A course in probability theory*. Academic Press, 3 edition, 2001.
- M. D. Conover, B. Gonçalves, J. Ratkiewicz, A. Flammini, and F. Menczer. Predicting the political alignment of twitter users. In *2011 IEEE third international conference on privacy, security, risk and trust and 2011 IEEE third international conference on social computing*, pages 192–199. IEEE, 2011.
- D. Cullina and N. Kiyavash. Improved achievability and converse bounds for erdos-rényi graph matching. *ACM SIGMETRICS Performance Evaluation Review*, 44(1):63–72, 2016.
- C. S. Cutts and S. J. Eglén. Detecting pairwise correlations in spike trains: an objective comparison of methods and application to the study of retinal waves. *Journal of Neuroscience*, 34(43):14288–14303, 2014.
- C. Davis and W. Kahan. The rotation of eigenvectors by a perturbation. III. *SIAM Journal on Numerical Analysis*, 7:1–46, 1970.
- B. Draves and D. L. Sussman. Bias-variance tradeoffs in joint spectral embeddings. *arXiv preprint arXiv:2005.02511*, 2020.
- D. Durante and D. B. Dunson. Bayesian inference on group differences in brain networks. *Bayesian Analysis*, 13(1), 2018.
- D. Durante, D. B. Dunson, and J. T. Vogelstein. Nonparametric bayes modeling of populations of networks. *Journal of the American Statistical Association*, 112(520):1516–1530, 2017.
- T. Elmer, K. Mephram, and C. Stadtfeld. Students under lockdown: Comparisons of students’ social networks and mental health before and during the covid-19 crisis in switzerland. *Plos one*, 15(7):e0236337, 2020.
- A. Elmsallati, C. Clark, and J. Kalita. Global alignment of protein-protein interaction networks: A survey. *IEEE/ACM transactions on computational biology and bioinformatics*, 13(4):689–705, 2015.
- C. Fraley and A. E. Raftery. Model-based clustering, discriminant analysis and density estimation. *Journal of the American Statistical Association*, 97:611–631, 2002.
- I. Gallagher, A. Jones, and P. Rubin-Delanchy. Spectral embedding for dynamic networks with stability guarantees. *Advances in Neural Information Processing Systems*, 34, 2021.
- S. Gelfman, Q. Wang, Y. Lu, D. Hall, C. D. Bostick, R. Dhindsa, M. Halvorsen, K. M. McSweeney, E. Cotterill, and T. Edinburgh. meartools: An r package for the analysis of neuronal networks recorded on microelectrode arrays. *PLoS computational biology*, 14(10):e1006506, 2018.
- J. C. Gower. Generalized procrustes analysis. *Psychometrika*, 40:33–51, 1975.

- A. Grover and J. Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 855–864, 2016.
- O. J. Heilmann and E. H. Lieb. Theory of monomer-dimer systems. In *Statistical Mechanics*, pages 45–87. Springer, 1972.
- E. S. Hill, J. W. Brown, and W. N. Frost. Photodiode-based optical imaging for recording network dynamics with single-neuron resolution in non-transgenic invertebrates. *J. Vis. Exp.*, 161, e61623, 2020. doi: 10.3791/61623.
- P. D. Hoff, A. E. Raftery, and M. S. Handcock. Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97(460):1090–1098, 2002a.
- P. D. Hoff, A. E. Raftery, and M. S. Handcock. Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97:1090–1098, 2002b.
- Peter Hoff. Modeling homophily and stochastic equivalence in symmetric relational data. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems*, volume 20, pages 657–664. Curran Associates, Inc., 2008. URL <https://proceedings.neurips.cc/paper/2007/file/766ebcd59621e305170616ba3d3dac32-Paper.pdf>.
- P. W. Holland, K. Laskey, and S. Leinhardt. Stochastic blockmodels: First steps. *Social Networks*, 5:109–137, 1983.
- R. Horn and C. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.
- Bing-Yi Jing, Ting Li, Zhongyuan Lyu, and Dong Xia. Community detection on mixture multi-layer networks via regularized tensor decomposition. *arXiv preprint arXiv:2002.04457*, 2020.
- Gregory Kiar, Eric Bridgeford, Will Gray Roncal, Consortium for Reliability (CoRR), Reproducibility, Vikram Chandrashekar, Disa Mhembere, Sephira Ryman, Xi-Nian Zuo, Daniel S Marguiles, R Cameron Craddock, Carey E Priebe, Rex Jung, Vince Calhoun, Brian Caffo, Randal Burns, Michael P Milham, and Joshua Vogelstein. A high-throughput pipeline identifies robust connectomes but troublesome variability. *bioRxiv*, 2018. doi: 10.1101/188706. URL <http://m2g.io>. <https://www.biorxiv.org/content/early/2018/02/24/188706>.
- M. Kivelä, A. Arenas, M. Barthelemy, J. P. Gleeson, Y. Moreno, and M. A. Porter. Multi-layer networks. *Journal of complex networks*, 2(3):203–271, 2014.
- Jing Lei. Network representation using graph root distributions, 2020.
- Lihua Lei, Xiaodong Li, and Xingmei Lou. Consistency of spectral clustering on hierarchical stochastic block models. *arXiv preprint arXiv:2004.14531*, 2020.
- K. Levin, A. Athreya, M. Tang, V. Lyzinski, and C. E. Priebe. A central limit theorem for an omnibus embedding of random dot product graphs. *arXiv preprint arXiv:1705.09355*, 2017.

- K. Levin, F. Roosta, M. Mahoney, and C. Priebe. Out-of-sample extension of graph adjacency spectral embedding. In *International Conference on Machine Learning*, pages 2975–2984. PMLR, 2018.
- T. R. Lezon, J. R. Banavar, M. Cieplak, A. Maritan, and N. V. Fedoroff. Using the principle of entropy maximization to infer genetic interaction networks from gene expression patterns. *Proceedings of the National Academy of Sciences*, 103(50):19033–19038, 2006.
- Y. Li and H. Li. Two-sample test of community memberships of weighted stochastic block models. *arXiv preprint arXiv:1811.12593*, 2018.
- Meimei Liu, Zhengwu Zhang, and David B Dunson. Graph auto-encoding brain networks with applications to analyzing large-scale brain imaging datasets. *Neuroimage*, 245:118750, 2021.
- L. Lu and X. Peng. Spectra of edge-independent random graphs. *Electronic Journal of Combinatorics*, 20, 2013.
- R. Lyons. The ising model and percolation on trees and tree-like graphs. *Communications in Mathematical Physics*, 125(2):337–353, 1989.
- V. Lyzinski. Information recovery in shuffled graphs via graph matching. *IEEE Transactions on Information Theory*, 64(5):3254–3273, 2018.
- V. Lyzinski and D. L. Sussman. Matchability of heterogeneous networks pairs. *Information and Inference: A Journal of the IMA*, 01 2020. URL <https://doi.org/10.1093/imaiai/iaz031>. iaz031.
- V. Lyzinski, D. L. Sussman, M. Tang, A. Athreya, and C. E. Priebe. Perfect clustering for stochastic blockmodel graphs via adjacency spectral embedding. *Electronic Journal of Statistics*, 8:2905–2922, 2014.
- V. Lyzinski, D. L. Sussman, D. E. Fishkind, H. Pao, L. Chen, J. T. Vogelstein, Y. Park, and C. E. Priebe. Spectral clustering for divide-and-conquer graph matching. *Parallel Computing*, 47:70–87, 2015.
- V. Lyzinski, M. Tang, A. Athreya, Y. Park, and C. E. Priebe. Community detection and classification in hierarchical stochastic blockmodels. *IEEE Transactions in Network Science and Engineering*, 4:13–26, 2017.
- A. Myronenko and X. Song. Point set registration: Coherent point drift. *IEEE transactions on pattern analysis and machine intelligence*, 32(12):2262–2275, 2010.
- A. M. Nielsen and D. Witten. The multiple random dot product graph model. *arXiv preprint arXiv:1811.12172*, 2018.
- R. I. Oliveira. Concentration of the adjacency matrix and of the Laplacian in random graphs with independent edges. <http://arxiv.org/abs/0911.0600>, 2009.

- H. Patsolic, S. Adali, J. T. Vogelstein, Y. Park, C. E. Friebe, G. Li, and V. Lyzinski. Seeded graph matching via joint optimization of fidelity and commensurability. *arXiv preprint arXiv:1401.3813*, 2014.
- H. G. Patsolic, Y. Park, V. Lyzinski, and C. E. Priebe. Vertex nomination via seeded graph matching. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 13(3): 229–244, 2020.
- S. Paul and Y. Chen. Consistent community detection in multi-relational data through restricted multi-layer stochastic blockmodel. *Electronic Journal of Statistics*, 10(2):3807–3870, 2016.
- S. Paul and Y. Chen. Spectral and matrix factorization methods for consistent community detection in multi-layer networks. *The Annals of Statistics*, 48(1):230–250, 2020.
- P. Pedarsani and M. Grossglauser. On the privacy of anonymized networks. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '11, page 1235–1243, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450308137. doi: 10.1145/2020408.2020596. URL <https://doi.org/10.1145/2020408.2020596>.
- Marianna Pensky and Teng Zhang. Spectral clustering in the dynamic stochastic block model. *Electron. J. Statist.*, 13(1):678–709, 2019. doi: 10.1214/19-EJS1533. URL <https://doi.org/10.1214/19-EJS1533>.
- B. L. Perry, B. A. Pescosolido, and S. P. Borgatti. *Egocentric network analysis: Foundations, methods, and models*, volume 44. Cambridge university press, 2018.
- C. E. Priebe, Y. Park, J. T. Vogelstein, J. M. Conroy, V. Lyzinski, M. Tang, A. Athreya, J. Cape, and E. Bridgeford. On a two-truths phenomenon in spectral graph clustering. *Proceedings of the National Academy of Sciences*, 116(13):5995–6000, 2019.
- L. F.R. Ribeiro, P. H.P. Saverese, and D. R. Figueiredo. struc2vec: Learning node representations from structural identity. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 385–394, 2017.
- K. Rohe, S. Chatterjee, and B. Yu. Spectral clustering and the high-dimensional stochastic blockmodel. *Annals of Statistics*, 39:1878–1915, 2011.
- P. Rubin-Delanchy, C. E. Priebe, and M. Tang. The generalised random dot product graph. Arxiv preprint available at <http://arxiv.org/abs/1709.05506>, 2017.
- D. L. Sussman, M. Tang, D. E. Fishkind, and C. E. Priebe. A consistent adjacency spectral embedding for stochastic blockmodel graphs. *Journal of the American Statistical Association*, 107:1119–1128, 2012.
- D. L. Sussman, M. Tang, and C. E. Priebe. Consistent latent position estimation and vertex classification for random dot product graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36:48–57, 2014.

- M. Tang and C. E. Priebe. Limit theorems for eigenvectors of the normalized laplacian for random graphs. *Annals of Statistics*, 2018. In press.
- M. Tang and Carey E. Priebe. Limit theorems for eigenvectors of the normalized Laplacian for random graphs. Arxiv preprint., 2016.
- M. Tang, Y. Park, N. H. Lee, and C. E. Priebe. Attribute fusion in a latent process model for time series of graphs. *IEEE Transactions on Signal Processing*, 61(7):1721–1732, 2013a.
- M. Tang, D. L. Sussman, and C. E. Priebe. Universally consistent vertex classification for latent position graphs. *Annals of Statistics*, 41:1406 – 1430, 2013b.
- M. Tang, A. Athreya, D. L. Sussman, V. Lyzinski, Y. Park, and C. E. Priebe. A semiparametric two-sample hypothesis testing problem for random dot product graphs. *Journal of Computational and Graphical Statistics*, 26:344–354, 2017a.
- M. Tang, A. Athreya, D. L. Sussman, V. Lyzinski, and C. E. Priebe. A nonparametric two-sample hypothesis testing problem for random dot product graphs. *Bernoulli*, 23: 1599–1630, 2017b.
- M. Tang, J. Cape, and C. E. Priebe. Asymptotically efficient estimators for stochastic blockmodels: The naive mle, the rank-constrained mle, and the spectral. *arXiv preprint arXiv:1710.10936*, 2017c.
- R. Tang, M. Ketcha, A. Badea, E. D. Calabrese, D. S. Margulies, J. T. Vogelstein, C. E. Priebe, and D. L. Sussman. Connectome smoothing via low-rank approximations. *IEEE transactions on medical imaging*, 38(6):1446–1456, 2018.
- R. Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- H. Wang, M. Tang, Y. Park, and C. E. Priebe. Locality statistics for anomaly detection in time series of graphs. *IEEE Transactions on Signal Processing*, 62(3):703–717, 2013.
- S. Wang, J. Arroyo, J. T. Vogelstein, and C. E. Priebe. Joint embedding of graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- M. D. Ward, K. Stovel, and A. Sacks. Network analysis and political science. *Annual Review of Political Science*, 14:245–264, 2011.
- A. S. Waugh, L. Pei, J. H. Fowler, P. J. Mucha, and M. A. Porter. Party polarization in congress: A network science approach. *arXiv preprint arXiv:0907.3509*, 3(4):69, 2009.
- F. Xie and Y. Xu. Efficient estimation for random dot product graphs via a one-step procedure. *arXiv preprint arXiv:1910.04333*, 2019.
- S. Young and E. Scheinerman. Random dot product graph models for social networks. In *Proceedings of the 5th international conference on algorithms and models for the web-graph*, pages 138–149, 2007.

- Y. Yu, T. Wang, and R. J. Samworth. A useful variant of the Davis-Kahan theorem for statisticians. *Biometrika*, 102:315–323, 2015.
- A. Zhang. Cross: Efficient low-rank tensor completion. *The Annals of Statistics*, 47(2): 936–964, 2019.
- A. Zhang and D. Xia. Tensor svd: Statistical and computational limits. *IEEE Transactions on Information Theory*, 64(11):7311–7338, 2018.
- L. Zhang, G. Wahba, and M. Yuan. Distance shrinkage and euclidean embedding via regularized kernel estimation. *arXiv preprint arXiv:1409.5009*, 2014.
- Y. Zhang. Consistent polynomial-time unseeded graph matching for lipschitz graphons. *arXiv preprint arXiv:1807.11027*, 2018a.
- Y. Zhang. Unseeded low-rank graph matching by transform-based unsupervised point registration. *arXiv preprint arXiv:1807.04680*, 2018b.
- M. Zhu and A. Ghodsi. Automatic dimensionality selection from the scree plot via the use of profile likelihood. *Computational Statistics and Data Analysis*, 51:918–930, 2006.
- Xi-Nian Zuo, Jeffrey S Anderson, Pierre Bellec, Rasmus M Birn, Bharat B Biswal, Janusch Blautzik, John C.S Breitner, Randy L Buckner, Vince D Calhoun, F. Xavier Castellanos, Antao Chen, Bing Chen, Jiangtao Chen, Xu Chen, Stanley J Colcombe, William Courtney, R Cameron Craddock, Adriana Di Martino, Hao-Ming Dong, Xiaolan Fu, Qiyong Gong, Krzysztof J Gorgolewski, Ying Han, Ye He, Yong He, Erica Ho, Avram Holmes, Xiao-Hui Hou, Jeremy Huckins, Tianzi Jiang, Yi Jiang, William Kelley, Clare Kelly, Margaret King, Stephen M LaConte, Janet E Lainhart, Xu Lei, Hui-Jie Li, Kaiming Li, Kuncheng Li, Qixiang Lin, Dongqiang Liu, Jia Liu, Xun Liu, Yijun Liu, Guangming Lu, Jie Lu, Beatriz Luna, Jing Luo, Daniel Lurie, Ying Mao, Daniel S Margulies, Andrew R Mayer, Thomas Meindl, Mary E Meyerand, Weizhi Nan, Jared A Nielsen, David O’Connor, David Paulsen, Vivek Prabhakaran, Zhigang Qi, Jiang Qiu, Chunhong Shao, Zarrar Shehzad, Weijun Tang, Arno Villringer, Huiling Wang, Kai Wang, Dongtao Wei, Gao-Xia Wei, Xu-Chu Weng, Xuehai Wu, Ting Xu, Ning Yang, Zhi Yang, Yu-Feng Zang, Lei Zhang, Qinglin Zhang, Zhe Zhang, Zhiqiang Zhang, Ke Zhao, Zonglei Zhen, Yuan Zhou, Xing-Ting Zhu, and Michael P Milham. An open science resource for establishing reliability and reproducibility in functional connectomics. *Scientific Data*, 1:140049, dec 2014. ISSN 2052-4463. doi: 10.1038/sdata.2014.49. URL <http://www.nature.com/articles/sdata201449>.