Journal of the American Statistical Association

Taylor & Francis
Taylor & Francis Group

ISSN: 0162-1459 (Print) 1537-274X (Online) Journal homepage: www.tandfonline.com/journals/uasa20

Euclidean Mirrors and Dynamics in Network Time Series

Avanti Athreya, Zachary Lubberts, Youngser Park & Carey Priebe

To cite this article: Avanti Athreya, Zachary Lubberts, Youngser Park & Carey Priebe (2025) Euclidean Mirrors and Dynamics in Network Time Series, Journal of the American Statistical Association, 120:550, 1025-1036, DOI: 10.1080/01621459.2024.2392912

To link to this article: https://doi.org/10.1080/01621459.2024.2392912

JOURNAL OF THE AMERICAN STATISTICAL ASSOCIATION
2025, VOL. 120, NO. 550, 1025–1036: Theory and Methods
https://doi.org/10.1080/01621459.2024.2392912

Taylor & Francis
Taylor & Francis Group

OPEN ACCESS

Check for updates

Euclidean Mirrors and Dynamics in Network Time Series

Avanti Athreya[a*], Zachary Lubberts[b*], Youngser Park[c], and Carey Priebe[a]

[a]Department of Applied Mathematics and Statistics, Johns Hopkins University, Baltimore, MD; [b]Department of Statistics, University of Virginia, Charlottesville, VA; [c]Center of Imaging Science, Johns Hopkins University, Baltimore, MD

ABSTRACT

Analyzing changes in network evolution is central to statistical network inference. We consider a dynamic network model in which each node has an associated time-varying low-dimensional latent vector of feature data, and connection probabilities are functions of these vectors. Under mild assumptions, the evolution of latent vectors exhibits low-dimensional manifold structure under a suitable distance. This distance can be approximated by a measure of separation between the observed networks themselves, and there exist Euclidean representations for underlying network structure, as characterized by this distance. These Euclidean representations, called Euclidean mirrors, permit the visualization of network dynamics and lead to methods for change point and anomaly detection in networks. We illustrate our methodology with real and synthetic data, and identify change points corresponding to massive shifts in pandemic policies in a communication network of a large organization. Supplementary materials for this article are available online, including a standardized description of the materials available for reproducing the work.

ARTICLE HISTORY
Received August 2022
Accepted June 2024

KEYWORDS
Dissimilarity measure; Euclidean realizability; Network time series; Spectral decomposition

1. Introduction

The structure of many organizational and communication networks underwent a dramatic shift during the disruption of the COVID-19 pandemic in 2020 (Zuzul et al. Forthcoming). This massive shock altered network connectivity in many respects and across multiple scales, differentially impacting individual nodes, local sub-communities, and whole networks. A visualization of this can be seen in Figure 1, which illustrates the shifting structure, from the spring to the summer of 2020, in a communications network of a large corporation. Each node represents an E-mail account, and connection between nodes reflect E-mail frequency between accounts. The panel on the left of Figure 1 shows a clustering of the network into subcommunities, and the panel on the right shows how those network connections between the same individuals shifted over time.

Such transformations give rise to several important questions in statistical network inference: how to construct useful measures of dissimilarity across networks; how to estimate any such measure of dissimilarity from random network realizations; how to identify loci of change; and how to gauge differences across scales, from nodes to sub-networks to the entire network itself. Our goal in this article is to build a robust methodology to address such phenomena, and to model and infer important characteristics of network time series.

To this end, we focus on a class of time series of random networks. We define an intuitive distance between the evolution of certain random variables that govern the behavior of nodes in the networks and prove that this distance can be consistently estimated from the observed networks. When this distance is sufficiently similar to a Euclidean distance, multidimensional scaling extracts a curve in low-dimensional Euclidean space that mirrors the structure of the network dynamics. This permits a visualization of network evolution and identification of change points. Figure 2 is the result of an end-to-end case study using these techniques for a time series of communication networks in a large corporation in the months around the start of pandemic work-from-home protocols: see the dramatic change in both panels beginning in Spring 2020. See Section 3 for the methodology used to generate these figures, and Section 4 for the full details of this experiment.

Analysis of multiple networks is a key emerging subdiscipline of network inference, with approaches ranging from joint spectral embedding (Levin et al. 2017; Jones and Rubin-Delanchy 2020; Arroyo et al. 2021; Gallagher, Jones, and Rubin-Delanchy 2021; Jing et al. 2021; Pantazis et al. 2022), tensor decompositions (Zhang and Xia 2018; Lei, Chen, and Lynch 2020; Jing et al. 2021), least-squares methods (Pensky 2019; Lei and Lin 2022), maximum likelihood methods (Krivitsky and Handcock 2014) and multiscale methods via random walks on graphs (Lee and Maggioni 2011). In Padilla, Yu, and Priebe (2022) and Wang, Yu, and Rinaldo (2021), the authors consider changepoint localization for a time series of latent position random graphs (Hoff, Raftery, and Handcock 2002), a type of independent-edge network in which each node or vertex has an associated latent position that determines its probability of connection with others. The authors establish consistency for localization of a particular kind of changepoint—namely, the case in which the

CONTACT Avanti Athreya dathrey1@jhu.edu Department of Applied Mathematics and Statistics, Johns Hopkins University, Baltimore, MD 21218.
*Co-first authors.
Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/JASA.

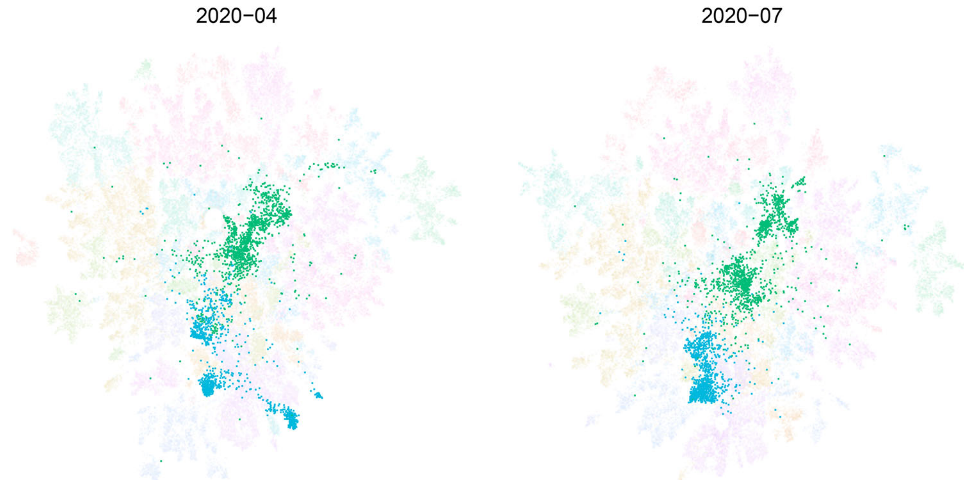

Figure 1. Evidence of structural network dissimilarity across time. Visualizations are of anonymized and aggregated Microsoft communications networks in April and July 2020, with two communities highlighted. Nodes represent E-mail accounts, with edges determined by E-mail frequency (edges not shown in these images). Left panel shows an initial Leiden clustering (Traag, Waltman, and Van Eck 2019) of the nodes into different subcommunities (colored); right panel shows the shift in connectivity structure over this time. Full visualization details are given in Section 4.1. During this time period, our analysis demonstrates that the network as a whole experiences a structural shock coincident with the pandemic work-from-home order. However, different subcommunities undergo qualitatively different changes in their structure, from a connected network that seems to diverge (green) and a less cohesive one that seems to coalesce (blue). (See Figure 4, wherein the overall network behavior, as well as the two highlighted communities with their different temporal behavior, are depicted.) The figures here are two-dimensional renderings of temporal snapshots of a large (n = 32,277) complex network; hence, conclusions based on this visualization are notional.

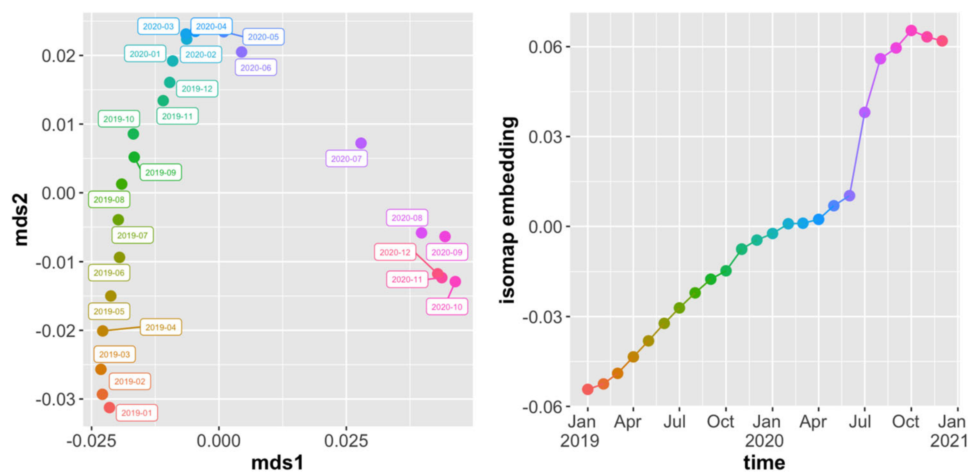

Figure 2. Classical multidimensional scaling (CMDS) and ISOMAP embeddings of estimated network dissimilarity identifies changes in network structure. Left plot shows top two dimensions CMDS of pairwise network distance matrix for anonymized, aggregated Microsoft communications networks from January 2019 to December 2020; individual networks follow a curve that progresses smoothly until spring 2020 and then exhibits a major shock. Right plot shows associated ISOMAP manifold learning representation and reveals a potential anomaly and/or changepoint. Full details for these plots are in Section 4.

latent positions are all fixed prior to some time point and after which they may be different. Asymptotic properties of these methods depend on particular model assumptions for how the networks evolve over time and relate to one another, and rigorous performance guarantees can be challenging and limited in scope. The underlying geometry of latent spaces affects network structure and evolution as well. In Smith, Asta, and Calder (2019), the authors consider the impact of different curvature and non-Euclidean properties of latent space geometry on network formation. In Wilkins-Reeves and McCormick (2022), the authors prove asymptotic results for estimators of underlying latent space curvature.

On the one hand, both single and multiple-network inference problems often have related objectives. For example, if data include multiple network realizations from the same underlying model on the same set of aligned vertices, we may wish to effectively exploit these additional realizations for more accurate estimation of common network parameters—that is, use the replications in a multiple-network setting to refine parameter estimates that govern any single network in the collection. On the other hand, multiple network inference involves statistically distinct questions, such as identifying loci of change *across* networks or detecting anomalies in a time series of networks.

Euclidean latent position networks assign to each vertex a typically unobserved vector in some low-dimensional Euclidean space $\mathbb{R}^d$; edges between vertices then arise independently. The probability of an edge between vertex i and vertex j is some fixed function κ, called the *link* function or kernel, of the two associated latent positions for the respective vertices. Latent position random graphs have the appealing characteristic of modeling network connections as functions of inherent features of the vertices themselves—these features are encoded in the latent positions—and transforming network inference into the recovery of lower-dimensional structure. More specifically, if we have a series of time-indexed latent position graphs G_t on a common aligned vertex set, then associated to each network is a matrix $\mathbf{X}_t$ whose rows are the latent vectors of the vertices. Since the edge formation probabilities are a function of pairs of rows of $\mathbf{X}_t$, the probabilistic evolution of the network time series is completely determined by the evolution of the rows of $\mathbf{X}_t$. As such, the natural object of study for inference about a time series of latent position graphs are the rows of $\mathbf{X}_t$. In particular, anomalies or changepoints in the time-series of networks correspond to changes in the $\mathbf{X}_t$ process. For example, a change in a specific network entity is associated to a change in its latent position, which can then be estimated.

The evolution of the rows of $\mathbf{X}_t$ can be deterministic, as is the case when features of the nodes in a network follow some predictable time-dependent pattern; but it can also be random, as is the case when the actors in a network have underlying preferences that are subject to random shocks. When the latent position vector $X_t(i)$ for some individual vertex i is a random variable, we have, as t varies, a stochastic process. This collection of random variables can be endowed with a metric, which under certain conditions is Euclidean realizable; that is, the random variables at each time have a representation as points in $\mathbb{R}^c$ for some dimension c, where the metric space distances between the random variables are equal to the Euclidean distances between these points (see Borg and Groenen (2005) for more on Euclidean realizability of dissimilarity matrices). This allows us to visualize the time evolution of this stochastic process as the image of a map from an interval into $\mathbb{R}^c$.

We use this idea to formulate a novel approach to network time series. We demonstrate methods for consistently estimating a Euclidean representation, or *mirror*, of the evolution of the latent position distributions from the observed networks. This mirror can reveal important underlying structure of the network dynamics, as we demonstrate in both simulated and real data, the latter of which is drawn from organizational and communication networks, revealing the changepoint corresponding to the start of pandemic work-from-home orders.

2. Model and Geometric Results

In order to model the intrinsic characteristics of the entities in our network, we consider *latent position random graphs*, which associate a vector of features in $\mathbb{R}^d$ to each vertex in the network. The connections between vertices in the network are independent given the latent positions, with connection probabilities depending on the latent position vectors of the two vertices in question. In our notation, $X \in \mathcal{X} \subset \mathbb{R}^d$ or $x \in \mathcal{X} \subset \mathbb{R}^d$, represent column vectors. If such column vectors are arranged as rows in a matrix, we specify this explicitly or we use the transpose to denote the corresponding row vector.

Definition 1 (Latent Position Graph, Random Dot Product Graph, and Generalized Random Dot Product). We say that the random graph G with adjacency matrix $\mathbf{A} \in \mathbb{R}^{n\times n}$ is a *latent position random graph (LPG)* with latent position matrix $\mathbf{X} \in \mathbb{R}^{n\times d}$, whose rows are the transpositions of the column vectors $X^1, \dots, X^n \in \mathcal{X} \subseteq \mathbb{R}^d$, and link function $\kappa : \mathcal{X} \times \mathcal{X} \to [0, 1]$, if

$$\mathbb{P}[\mathbf{A}|\mathbf{X}] = \prod_{i<j} \kappa(X^i, X^j)^{A_{i,j}}(1 - \kappa(X^i, X^j))^{1-A_{i,j}}.$$

If $\kappa(x, y) = x^\top y$, we say that G is a *random dot product graph (RDPG)* and we call $\mathbf{P} = \mathbf{X}\mathbf{X}^T$ the connection probability matrix. In this case, each A_{ij} is marginally distributed (conditionally on X^i, X^j) as Bernoulli$(\langle X^i, X^j\rangle)$.

As a generalization, suppose $x^1, y^1 \in \mathbb{R}^p$ and $x^2, y^2 \in \mathbb{R}^q$, where $p + q = d$. When $\kappa([(x^1)^\top, (x^2)^\top]^\top, [(y^1)^\top, (y^2)^\top]^\top) = (x^1)^\top y^1 - (x^2)^\top y^2$, we say that G is a *generalized random dot product graph (GRDPG)* and we call $\mathbf{P} = \mathbf{X}I_{p,q}\mathbf{X}^T$ the generalized edge connection probability matrix, where $I_{p,q} = I_p \oplus (-I_q)$.

Remark 1 (Orthogonal nonidentifiability in RDPGs). Note that if $\mathbf{X} \in \mathbb{R}^{n\times d}$ is a matrix of latent positions and $\mathbf{W} \in \mathbb{R}^{d\times d}$ is orthogonal, $\mathbf{X}$ and $\mathbf{XW}$ give rise to the same distribution over graphs. Thus, the RDPG model has a nonidentifiability up to orthogonal transformation. Analogously, the GRDPG model has a nonidentifiability up to indefinite orthogonal transformations.

Since we wish to model randomness in the underlying features of each vertex, we will consider latent positions that are themselves random variables defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. For a particular sample point $\omega \in \Omega$, let $X(t, \omega) \in \mathbb{R}^d$ be the realization of the associated latent position for this vertex at time $t \in [0, T]$. On the one hand, for fixed ω, as t varies, $X(t, \omega), 0 \le t \le T$ is the realized trajectory of a d-dimensional stochastic process. On the other hand, for a given time t, the random variable $X(t, \cdot)$ represents the constellation of possible latent positions at this time. In order for the inner product to be a well-defined link function, we require that the distribution of $X(t, \cdot)$ follow an *inner-product distribution*:

Definition 2. Let F be a probability distribution on $\mathbb{R}^d$. We say that F is a *d-dimensional inner product distribution* if $0 \le x^\top y \le 1$ for all $\mathbf{x}, \mathbf{y} \in \mathrm{supp}F$. We will suppose throughout this work that for a d-dimensional inner product distribution F and $X \sim F$, $\mathbb{E}[XX^\top]$ has rank d.

We wish to quantify the difference between the random vectors $X(t, \cdot)$ and $X(t', \cdot)$. Suppose that the graphs come from an RDPG or GRDPG model, where at each time t, the latent positions of each graph vertex are drawn independently from a common inner product latent position distribution F_t. Because $X(t, \cdot)$ is a latent position, we necessarily have $X(t, \cdot) \in L^2(\Omega)$; for notational simplicity, we will use $X(t, \cdot)$ and X_t interchangeably. We define a norm, which we call the *maximum directional variation norm*, on this space of random variables; this norm leads to a natural metric $d_{MV}(X_t, X_{t'})$, both of which are described below. In the definition below, and throughout the article, we use $\|\cdot\|$ to denote the Euclidean norm in $\mathbb{R}^d$, $\|\cdot\|_2$ to denote the spectral norm of a matrix, and $\|\cdot\|_F$ to denote the Frobenius norm of a matrix.

Definition 3 (Maximum directional variation norm and metric). For a random vector $X \in L^2(\Omega)$, we define

$$\|X\|_{MV} = \max_u \mathbb{E}[\langle X, u\rangle^2]^{1/2} = \|\mathbb{E}[XX^\top]\|_2^{1/2},$$

where the maximization is over $u \in \mathbb{R}^d$ with $\|u\| = 1$. We define an associated metric d_{MV} by minimizing the norm of the difference between the random variables $X_t, X_{t'}$ over all orthogonal transformations, which aligns these distributions.

$$d_{MV}(X_t, X_{t'}) = \min_W \|X_t - WX_{t'}\|_{MV} = \min_W \left\| \mathbb{E}[(X_t - WX_{t'})(X_t - WX_{t'})^\top] \right\|_2^{1/2}, \quad (1)$$

where the matrix norm on the right hand side is the spectral norm. Given a map $\varphi : [0, T] \to L^2(\Omega)$ that assigns time points t to random variables $\varphi(t) = X_t$, we may write $d_{MV}(\varphi(t), \varphi(t')) := d_{MV}(X_t, X_{t'})$.

The minimization in (1) is a variant of the classical Procrustes alignment problem, so we may refer to the latent positions after this rotation as "Procrustes-aligned."

Remark 2. If X has mean zero, the $\|X\|_{MV}$ considers the square of spectral norm of its covariance matrix; that is, the norm $\|\cdot\|_{MV}$ gives the maximal directional *variation* when X is centered. In the cases of interest, we wish to capture features of the variance of the *drift* in the latent position, $X_t - WX_{t'}$: this is the origin of the name for this metric and its associated norm. The metric d_{MV} is not properly a metric on $L^2(\Omega)$, since if $X = WY$ a.s. for some orthogonal matrix W, then $d_{MV}(X, Y) = 0$. However, if we consider the equivalence relation defined by $X \sim Y$ whenever $X = WY$ for some orthogonal matrix, this *is* a metric on the corresponding set of equivalence classes. This means that we are able to absorb the non-identifiability from the original parameterization, obtaining a new parameter space with a metric space structure where the underlying distribution is identifiable.

One of the central contributions of this article is that the d_{MV} metric captures important features of the time-varying distributions F_t. To describe a family of networks indexed by time, each of which is generated by a matrix of latent positions that are themselves random, we consider a *latent position stochastic process*.

Definition 4 (Latent position process). Let $\mathcal{F}_t, 0 \leq t \leq T$ be a filtration of $\mathcal{F}$. A *latent position process* $\varphi(t)$ is an $\mathcal{F}_t$-adapted map $\varphi : [0, T] \to (L^2(\Omega), d_{MV})$ such that for each $t \in [0, T]$, $\varphi(t) = X(t, \cdot)$ has an inner product distribution. We say that a latent position process is *nonbacktracking* if $\varphi(t) = \varphi(t')$ implies $\varphi(s) = \varphi(t)$ for all $s \in [t, t']$.

Once we have the latent position stochastic process, we can construct a time series of latent position random networks whose vertices have independent, identically distributed latent positions given by $\varphi(t)$.

Definition 5 (Time Series of LPGs). Let φ be a latent position process, and fix a given number of vertices n and collection of times $\mathcal{T} \subseteq [0, T]$. We draw an iid sample $\omega_j \in \Omega$ for $1 \leq j \leq n$, and obtain the latent position matrices $\mathbf{X}_t \in \mathbb{R}^{n \times d}$ for $t \in \mathcal{T}$ by appending the rows $X(t, \omega_j)$, $1 \leq j \leq n$. The *time series of LPGs (TSG)* $\{G_t : t \in \mathcal{T}\}$ are conditionally independent LPGs with latent position matrices $\mathbf{X}_t, t \in \mathcal{T}$.

We emphasize that each vertex in the TSG corresponds to a single $\omega \in \Omega$, which induces dependence between the latent positions for that vertex across time points, but the latent position trajectories of any two distinct vertices are independent of one another across all times. Since these trajectories form an iid sample from the latent position process, it is natural to measure their evolution over time using the metric on the corresponding random variables, namely $d_{MV}(X_t, X_{t'})$. In the definition of this distance, the expectation is over $\omega \in \Omega$, which means that it depends on the joint distribution of X_t and $X_{t'}$. In particular, d_{MV} depends on more than just the marginal distributions of the random vectors X_t and $X_{t'}$ individually, but takes into account their dependence inherited from the latent position process φ.

A key question is whether the image $\varphi([0, T])$ has useful geometric structure when equipped with the metric d_{MV}. It turns out that, under mild conditions, this image is a manifold. In addition, the map φ admits a Euclidean analogue, called a *mirror*, which is a finite-dimensional curve that retains important signal from the generating stochastic process for the network time series. To make this precise, we define the notions of Euclidean realizability and approximate Euclidean realizability, below, and provide several examples of latent position processes that satisfy these requirements.

Definition 6 (Notions of Euclidean realizability). Let φ be a latent position process.

We say that φ is *approximately (Lipschitz) Euclidean c-realizable* with *mirror* ψ and realizability constant $C > 0$ if there exists a Lipschitz continuous curve $\psi : [0, T] \to \mathbb{R}^c$ such that

$$\left| d_{MV}(\varphi(t), \varphi(t')) - \|\psi(t) - \psi(t')\| \right| \leq C|t - t'| \text{ for all } t, t' \in [0, T].$$

For a fixed $\alpha \in (0, 1)$, we say that φ is *approximately α-Hölder Euclidean c-realizable* if ψ is α-Hölder continuous, and there is some $C > 0$ such that

$$\left| d_{MV}(\varphi(t), \varphi(t')) - \|\psi(t) - \psi(t')\| \right| \leq C|t - t'|^\alpha \text{ for all } t, t' \in [0, T].$$

Rather than c-realizable, we may simply say *realizable* if there is some c for which this holds; we simply say that φ is Hölder Euclidean realizable if the condition holds for some $\alpha \in (0, 1]$.

Remark 3. If there exists a Lipschitz curve ψ in $\mathbb{R}^c$ for which

$$d_{MV}(\varphi(t), \varphi(t')) = \|\psi(t) - \psi(t')\|$$

we say the latent position process is *exactly Euclidean realizable*. While this is seldom the case for most interesting latent position processes, it can be instructive to consider what this implies: that pairwise d_{MV} distances between the latent position process at t and t' coincide exactly with Euclidean distances along the curve ψ at t and t'. Hence, the term *mirror*, a Euclidean-space curve that replicates (with some distortion in the approximately realizable case) the time-varying d_{MV} distance. For useful intuition, consider a one-dimensional Brownian motion B_t. While this is not a latent position process, its covariance operator $R(s, t) = \mathbb{E}[(B_t - B_s)^2]$ is exactly $\mathbb{E}[(B_t - B_s)^2] = (t - s)$, corresponding to the distance between points along the line $x(t) = t$ between t and s.

In practice, the latent position process is unobserved, so it is unclear whether the Euclidean realizability condition holds. However, we show that the d_{MV} distance can be consistently estimated, so the question of realizability may be resolved at least in part by inspection of the screenplot of the estimated distance matrix. We remark further on this point after Theorem 6.

Note that if φ is approximately c-realizable, it is c'-realizable for any $c' > c$, and a tradeoff exists between the choice of dimension c and the accuracy of the approximation, as measured by C and α. The realizability dimension c can be interpreted as a choice in a dimension reduction procedure. Namely, the dimension c corresponds to a curve ψ in $\mathbb{R}^c$, along which pairwise Euclidean distances locally approximate those of the maximum variational distances along the latent position process.

As such, none of c, C, or ψ, as defined above, need be unique. This leads naturally to the question of an "optimal" mirror—that is, one that best captures, in Euclidean space, the salient features of the d_{MV} distance. To make this precise, suppose φ is a latent position process.

For any associated mirror ψ, consider the functional $\mathcal{L}(\psi)$ given by

$$\mathcal{L}(\psi) = \int_0^T \int_0^T \left| d_{MV}^2(\varphi(s), \varphi(t)) - \|\psi(s) - \psi(t)\|^2 \right|^2 \mathrm{d}s\, \mathrm{d}t \quad (2)$$

As we show below, there exists a solution to the variational problem of minimizing this functional over the class of mirrors $\mathcal{S}(c, \alpha, C)$ that are α-Hölder with realizability constant C, and satisfy $\int_0^T \psi(t)\, \mathrm{d}t = 0$. We call this minimizer an *optimal* mirror for this α, c, and C. While any mirror satisfying the realizability constraints estimates the d_{MV} distance well locally, a minimizer of this functional also estimates the d_{MV} distance well in a global sense.

Theorem 1 (Existence of Optimal Mirrors). Let φ be a latent position process, which is approximately α-Hölder Euclidean c-realizable with realizability constant C. Let $\mathcal{S}(c, \alpha, C)$ be the class of mirrors defined above. Then there exists a solution to the variational problem

$$\inf_{\psi \in \mathcal{S}(c,\alpha,C)} \mathcal{L}(\psi) \quad (3)$$

Theorem 2 (Uniqueness of Optimal Mirrors). If φ is a latent position process which is exactly α-Hölder Euclidean c-realizable, the solution to the variational problem in (3) is unique up to orthogonal transformations.

In addition to the existence of optimal mirrors, an approximate Euclidean realizable latent position process φ has the property that its image is a manifold.

Theorem 3 (Manifold properties of a nonbacktracking latent position process). Let φ be a nonbacktracking latent position process which is approximately Euclidean realizable. Then $\mathcal{M} = \varphi([0, T])$ is homeomorphic to an interval $[0, I]$. In particular, it is a topological one-manifold with boundary. If φ is injective and approximately α-Hölder Euclidean realizable, the same conclusion holds.

If we suppose that the trajectories of φ satisfy a certain degree of smoothness, it turns out that the map φ into the space of random variables equipped with the d_{MV} metric also has this degree of smoothness.

Theorem 4. (Smooth trajectories and smooth latent position processes) Suppose $X(\cdot, \omega) : [0, T] \to \mathbb{R}^d$ is α-Hölder continuous with some $\alpha \in (0, 1]$ for almost every $\omega \in \Omega$, such that

$$\|X(t, \omega) - X(s, \omega)\| \le L(\omega)|t - s|^\alpha,$$

where the random variable $L \in L^2(\Omega)$. Let $\mathcal{M} = \varphi([0, T])$. Then $\varphi : [0, T] \to (\mathcal{M}, d_{MV})$ is Hölder continuous with this same α.

Remark 4. In the above definitions of realizability, regularity conditions are imposed on ψ, which takes values in $\mathbb{R}^c$, rather than on φ, which gives random variables as output. Moreover, ψ is the Euclidean realization of the manifold $\varphi([0, T])$ in the space of random variables; this as an approximately distance-preserving representation of those random variables, each of which captures the full state of the system with all of the given entities at any time t. As we show, estimates of this Euclidean mirror, derived from observations of graph connectivity structure at a collection of time points, can recover important features of the time-varying latent positions.

There are several natural classes of latent position processes that are approximately Lipschitz or α-Hölder Euclidean realizable. The next theorem demonstrates approximate α-Hölder Euclidean realizability for any latent position process expressible as the sum of a deterministic drift and a martingale term whose increments have well-controlled variance.

Theorem 5 (Approximate Holder realizability of variance-controlled martingale-plus-drift processes). Suppose M_t is an $\mathcal{F}_t$-martingale with respect to the filtration $\{\mathcal{F}_t : 0 \le t \le T\}$, and suppose $\gamma : [0, T] \to \mathbb{R}^d$ is Lipschitz continuous. Let $\varphi(t) = \gamma(t) + M_t$. Then

$$d_{MV}(X_t, X_s)^2 \le \|\mathrm{cov}(M_t - M_s)\|_2 + \|\gamma(t) - \gamma(s)\|^2.$$

When M_t satisfies $\|\mathrm{cov}(M_t - M_s)\|_2 \le C(t-s)$, and $\gamma(t) = a(t)v$ for some $v \in \mathbb{R}^d$ and Lipschitz continuous $a : [0, T] \to \mathbb{R}$, then φ is approximately α-Hölder Euclidean realizable with $\alpha = 1/2$ and $c = 1$.

Example 1. Consider $X_t = \gamma(t) + B_t$, where B_t is a d-dimensional Brownian motion, and $\gamma : [0, T] \to \mathbb{R}^d$ is a Lipschitz continuous function of the form $\gamma(t) = a(t)v$. The $\varphi(t)$ is approximately α-Hölder Euclidean realizable, with $\alpha = 1/2$, and the Euclidean mirror $\psi(t)$ is $\psi(t) = a(t)\|v\|$, so $c = 1$. In the supplementary material, we provide simulations of a network time series with this latent position process and show that our estimated mirror matches ψ well for a network of 2000 nodes.

Example 2. Consider $X_t = \gamma(t) + I_t$, where $I_t = \int_0^t B_s\, \mathrm{d}s$, B_s is a d-dimensional Brownian motion, and $\gamma(t) = (at + b)v$ is a function describing the mean of X_t over time, with $a, b \in \mathbb{R}, v \in \mathbb{R}^d$. Then each sample path of X_t is continuously differentiable in t, and $\mathbb{E}[X_t] = \gamma(t)$ is as well. If $\{\mathcal{F}_t : 0 \le t \le T\}$ is the canonical filtration generated by Brownian motion, then $I(t)$ is *not* an $\mathcal{F}_t$-martingale. Then

$$d_{MV}(X_t, X_{t'})^2 = a^2(t - t')^2\|v\|^2 + \sigma^2[(t - t')^2(t + 2t')/3],$$

so φ is approximately Euclidean realizable with $\psi(t) = \sqrt{a^2\|v\|^2 + \sigma^2 T}t$, so again $c = 1$.

The latent positions for the vertices in our network are not typically observed—instead, we only see the connectivity between the nodes in the network, from which a given realization of the latent positions can, under certain model assumptions, be accurately estimated. In order to compare the networks at times t and t', we can consider estimates of the networks' latent positions at these two times as noisy observations from the joint distribution of $(X_t, X_{t'})$, and deploy these estimates in an approximation of the distances $d_{MV}(X_t, X_{t'})$. Using these approximate distances, we can then estimate the curve $\psi(t)$, giving a visualization for the evolution of *all* of the latent positions in the random graphs over time.

Suppose that G is a random dot product graph with latent position matrix $\mathbf{X}$, where the rows of $\mathbf{X}$ are independent, identically distributed draws from a latent position distribution F on $\mathbb{R}^d$. Let $\mathbf{A}$ be the adjacency matrix for this graph. As shown in Sussman et al. (2012), a spectral decomposition of the adjacency matrix yields consistent estimates for the underlying matrix of latent positions. We introduce the following definition.

Definition 7 (Adjacency Spectral Embedding). Given an adjacency matrix $\mathbf{A}$, we define the *adjacency spectral embedding (ASE)* with dimension d as $\hat{\mathbf{X}} = \hat{\mathbf{U}}\hat{\mathbf{S}}^{1/2}$, where $\hat{\mathbf{U}} \in \mathbb{R}^{n\times d}$ is the matrix of d top eigenvectors of $\mathbf{A}$ and $\hat{\mathbf{S}} \in \mathbb{R}^{d\times d}$ is the diagonal matrix with the d largest eigenvalues of $\mathbf{A}$ on the diagonal.

As we show in the next section, we will use the ASE of the observed adjacency matrices in our TSG to estimate the d_{MV} distance between latent position random variables over time, and in turn, to estimate the Euclidean mirror, which records important underlying structure for the time series of networks.

3. Statistical Estimation of Euclidean Mirrors

Given a finite sample from a time series of graphs with approximately α-Hölder Euclidean realizable latent position process φ, our goal is to estimate a finite-sample analogue of an optimal Euclidean mirror ψ. The distances $d_{MV}(\varphi(t), \varphi(s))$ can be used to recover a version of the mirror at these sampled times (up to rigid transformations) from classical multidimensional scaling (CMDS). As such, the crucial estimation problem is one of accurately estimating the distances $d_{MV}(\varphi(t), \varphi(s))$. To this end, we define the *estimated pairwise distances* between any two such $n \times d$ latent position matrices $\hat{\mathbf{X}}_t$ and $\hat{\mathbf{X}}_s$ as follows:

$$\hat{d}_{MV}(\hat{\mathbf{X}}_t, \hat{\mathbf{X}}_s) := \min_{W\in\mathcal{O}^{d\times d}} \frac{1}{\sqrt{n}}\|\hat{\mathbf{X}}_t - \hat{\mathbf{X}}_s W\|_2, \tag{4}$$

where $\mathcal{O}^{d\times d}$ is the set of real orthogonal matrices of order d, and $\|\cdot\|_2$ denotes the spectral norm. Note that when $\mathbf{U}, \mathbf{V} \in M_{n,d}(\mathbb{R})$ have orthonormal columns, we have the following well-known relations between $\hat{d}_{MV}$ and the spectral norm of their $\sin\Theta$ matrix:

$$\|\sin\Theta(\mathbf{U},\mathbf{V})\|_2 \le \sqrt{n}\hat{d}_{MV}(\mathbf{U},\mathbf{V}) \le \sqrt{2}\|\sin\Theta(\mathbf{U},\mathbf{V})\|_2.$$

Our central result is that, when our networks have a sufficiently large number of vertices n, $\hat{d}_{MV}$ provides a consistent estimate of d_{MV}.

Theorem 6. Suppose φ is a LPP such that $\varphi(t)$ takes values in $\mathbb{R}^d$ for all $t \in [0,T]$. Let $\{G_t\}_{t\in\mathcal{T}}$ be a time series of LPGs whose latent positions follow this LPP. For each $t \in \mathcal{T}$, let $\hat{\mathbf{X}}_t \in \mathbb{R}^{n\times d}$ be the matrix of estimated latent positions from the ASE of each graph G_t. Then for all $s, t \in \mathcal{T}$, with overwhelming probability as $n \to \infty$,

$$\left|\hat{d}_{MV}(\hat{\mathbf{X}}_t, \hat{\mathbf{X}}_s)^2 - d_{MV}(\varphi(t),\varphi(s))^2\right| \le \frac{\log(n)}{\sqrt{n}}.$$

The functional $\mathcal{L}(\psi)$ in (2) requires information of $d_{MV}(\varphi(t), \varphi(s))$ for all $t, s \in [0,T]$. In the finite-sample case, however, we only have a fixed, finite set of time points $\mathcal{T} = \{t_i\}_{i=1}^m \subseteq [0,T]$, with $t_i < t_{i+1}$ for all i. To address finite-sample estimation of an analogue of an optimal mirror, we introduce the functional $\widehat{\mathcal{L}}$ defined on sets of size m of vectors in $\mathbb{R}^c$:

$$\widehat{\mathcal{L}}(\{v_i\}_{i=1}^m) = \sum_{i,j=1}^m \Big|d_{MV}(\varphi(t_i),\varphi(t_j))^2 - \|v_i - v_j\|^2\Big|^2 \ \Delta t_i\, \Delta t_j. \tag{5}$$

The time steps are defined as $\Delta t_i := t_i - t_{i-1}$, with $t_0 = 0$. Note that the time steps Δt_i need not be constant for (5), allowing us to consider real-data settings in which the network observations may not be equally spaced in time. Suppose we know the true matrix D_φ of pairwise distances whose i,jth entry is $d_{MV}(\varphi(t_i), \varphi(t_j))$. When the Δt_i are equal and the process is exactly Euclidean realizable, classical multidimensional scaling (CMDS) applied to this matrix yields a collection of vectors $\{\psi(t_i) : 1 \le i \le m\}$, unique up to rotation, that minimizes $\widehat{\mathcal{L}}(\{v_i\}_{i=1}^m)$ (see additional details after Corollary 1). We call this the finite-sample mirror for the latent position process φ. When the Δt_i are not all equal, we suggest obtaining an initial solution for (5) via CMDS, obtaining $V_0 \in \mathbb{R}^{m\times c}$, with rows $v_i \in \mathbb{R}^c$. Then one can employ a general-purpose nonlinear solver to (5) with the initial point V_0.

Having defined the dissimilarity matrices

$$\mathcal{D}_\varphi = [d_{MV}(\varphi(s),\varphi(t))]_{s,t\in\mathcal{T}}, \quad \mathcal{D}_\psi = [\|\psi(s) - \psi(t)\|_2]_{s,t\in\mathcal{T}},$$

where $\{\psi(t)\}_{t\in\mathcal{T}}$ achieves the minimum value of $\widehat{\mathcal{L}}$, we note that the first records the pairwise distances between the latent position process at times t and s; the second records the differences between the finite-sample optimal Euclidean mirror at these times. Of course, the true distances are not observed, and must be estimated. The estimates for these quantities are then

$$\widehat{\mathcal{D}}_\varphi = [\hat{d}_{MV}(\hat{\mathbf{X}}_s, \hat{\mathbf{X}}_t)]_{s,t\in\mathcal{T}}, \quad \mathcal{D}_{\hat{\psi}} = [\|\hat{\psi}(s) - \hat{\psi}(t)\|_2]_{s,t\in\mathcal{T}},$$

where $\{\hat{\psi}(t)\}_{t\in\mathcal{T}}$ is the output of CMDS applied to the matrix $\widehat{\mathcal{D}}_\varphi$. This means our full mirror estimation procedure is as follows:

Suppose $A^{(2)}$ is the matrix of squared entries of A. Theorem 6 then guarantees that the square of each entrywise difference between $\hat{\mathcal{D}}_\varphi^{(2)}$ and $\mathcal{D}_{\hat{\psi}}^{(2)}$ is bounded, with high probability, by $\log^2(n)/n$. Since both $\hat{\mathcal{D}}_\varphi^{(2)}$ and $\mathcal{D}_{\hat{\psi}}^{(2)}$ are $m \times m$ matrices, we immediately derive the following corollary:

Algorithm 1 Mirror estimation

1: **Input**: graph adjacency matrices $A_{t_1}, \ldots, A_{t_m}$, embedding dimensions d and c.
2: Compute Adjacency Spectral Embedding of A_{t_i} to obtain $\hat{\mathbf{X}}_{t_1}, \ldots, \hat{\mathbf{X}}_{t_m} \in M_{n,d}(\mathbb{R})$.
3: For $i,j \in \{1, \ldots, m\}$ compute matrix entry $\hat{\mathcal{D}}_\varphi(i,j) = \hat{d}_{MV}(\hat{\mathbf{X}}_{t_i}, \hat{\mathbf{X}}_{t_j})$.
4: Apply CMDS to $\hat{\mathcal{D}}_\varphi$ to yield $\hat{\psi}(t_i) \in \mathbb{R}^c, 1 \leq i \leq m$. The $\hat{\psi}(t_i)$ are defined as the rows of $\hat{U}\hat{S}^{1/2}$, where $\hat{U}\hat{S}\hat{U}^T$ is the closest rank-c matrix to $-\frac{1}{2}P\hat{\mathcal{D}}_\varphi^{(2)}P$ with respect to the Frobenius norm, and $P = I - J/m$, for J the $m \times m$ matrix of all ones.
5: **Output**: Return mirror estimates $\hat{\psi}(t_1), \ldots, \hat{\psi}(t_m)$.

Corollary 1. Suppose the setting of Theorem 6. For fixed m, with overwhelming probability,

$$\|\widehat{\mathcal{D}}_\varphi^{(2)} - \mathcal{D}_\varphi^{(2)}\|_F \leq \frac{m\log(n)}{\sqrt{n}}.$$

We recall that CMDS computes the scaled eigenvectors of the matrix $-\frac{1}{2}PA^{(2)}P$, where $P = I - J_m/m$ is a projection matrix, $A^{(2)}$ is a matrix of squared distances, and J_m is the $m \times m$ matrix of all ones. This matrix may be written as USU^T for some $m \times c$ matrix U with orthonormal columns and diagonal matrix S. This means that $\psi(t_i)$, the value at t_i of the finite-sample optimal Euclidean mirror associated to φ, is simply the ith row of the matrix $US^{1/2}$. We will analogously denote the ith row of $\hat{U}\hat{S}^{1/2}$, the output of CMDS applied to $\widehat{\mathcal{D}}_\varphi$, by $\hat{\psi}(t_i)$.

Remark 5. In practice, where c is unknown, the selection of the mirror dimension c is a model selection problem. However, in light of Corollary 1 and the Hoffman-Wielandt inequality, we see that the eigenvalues of the estimated projected distance matrix $-\frac{1}{2}P\widehat{\mathcal{D}}_\varphi^{(2)}P$ approximate those of the theoretical one, $-\frac{1}{2}P\mathcal{D}_\varphi^{(2)}P$. As such, the singular values of $-\frac{1}{2}P\widehat{\mathcal{D}}_\varphi^{(2)}P$ are consistent estimates for the true singular values, meaning that the correct choice of c will be revealed for large networks.

It turns out that for large networks, the invariant subspace associated to $\widehat{\mathcal{D}}_\varphi^{(2)}$ corresponding to its largest eigenvalues is an accurate approximation to the corresponding subspace of $\mathcal{D}_\varphi^{(2)}$, which matches that of $\mathcal{D}_\psi^{(2)}$ when we have approximate Euclidean realizability. This suggests that applying CMDS to the estimated dissimilarity matrix $\widehat{\mathcal{D}}_\varphi$ can recover the finite-sample optimal mirror $\psi(t)$ up to a rotation: in other words, $\hat{\psi}(t) \approx R\psi(t)$ for some real orthogonal matrix R and all $t \in \mathcal{T}$.

Theorem 7. Suppose φ is approximately Euclidean c-realizable and let m be a given positive integer. Let $\hat{U}, U \in \mathbb{R}^{m\times c}$ be the top c eigenvectors, and $\hat{S}, S \in \mathbb{R}^{c\times c}$ be the diagonal matrices with diagonal entries equal to the top c eigenvalues of $\hat{E}_\varphi = -\frac{1}{2}P\widehat{\mathcal{D}}_\varphi^{(2)}P^\top$ and $E_\varphi = -\frac{1}{2}P\mathcal{D}_\varphi^{(2)}P^\top$, respectively, where $P = I - J_m/m$, J_m is the all-ones matrix of order m. Suppose $S_{i,i} > 0$ for $1 \leq i \leq c$. Then with overwhelming probability, there is a real orthogonal matrix $R \in \mathcal{O}^{c\times c}$ such that

$$\|\hat{U} - UR\|_F \leq \frac{2^{3/2}}{\lambda_c(E_\varphi)}\left(\frac{m\log(n)}{\sqrt{n}} + \left(\sum_{i=c+1}^{m}\lambda_i^2\left(E_\varphi\right)\right)^{1/2}\right).$$

Call this upper bound $B = B(n, m, c)$. The spectrally-scaled CMDS output satisfies

$$\|\hat{U}\hat{S}^{1/2} - US^{1/2}R\|_F \leq B\lambda_1^{1/2}(E_\varphi) \times \left(2 + 4B\kappa^{1/2} + (1+2B)\frac{m\log(n)}{\sqrt{n}\lambda_c(E_\varphi)}\right),$$

where $\kappa = \lambda_1(E_\varphi)/\lambda_c(E_\varphi)$. In particular, we have

$$\sum_{i=1}^{m}\|\hat{\psi}(t_i) - R\psi(t_i)\|^2 \leq B^2\lambda_1(E_\varphi) \times \left(2 + 4B\kappa^{1/2} + (1+2B)\frac{m\log(n)}{\sqrt{n}\lambda_c(E_\varphi)}\right)^2.$$

If all but the top c eigenvalues of $\mathcal{D}_\varphi$ are sufficiently small—as is the case when $\mathcal{D}_\varphi$ is rank c—Theorem 7 ensures that a Euclidean mirror can be consistently estimated. As such, if the important aspects of a finitely-sampled latent position process, such as changepoints or anomalies, are reflected in low-dimensional Euclidean space, then we recover an optimal finite-sample mirror consistently through CMDS applied to the estimated distance matrix. We encapsulate our consistency results and connections between true distances, their estimates, and associated Euclidean mirrors in Figure 3.

The right-hand side of Figure 3 lists the true and typically unobserved distance measure d_{MV}, and from it, immediately below, the matrix of pairwise distances $\mathcal{D}_\varphi$. If this dissimilarity is Euclidean realizable in c dimensions, then classical multidimensional scaling will recover this mirror, denoted by ψ, up to Euclidean distance-preserving transformations.

On the left-hand side of Figure 3, we see how to compute an estimate of d_{MV} from spectral embeddings of a pair of observed network adjacencies. Theorem 6 grants that the estimated dissimilarity matrix of pairwise distance $\hat{\mathcal{D}}_\varphi$ will be close to the true dissimilarity $\mathcal{D}_\varphi$, and if the latent position process is Euclidean realizable, Theorem 7 establishes that classical multidimensional scaling applied to $\hat{\mathcal{D}}_\varphi$ serves as a consistent estimate for ψ.

This figure describes our overall approach to the problem of inference in time series of networks: first, we construct a useful dissimilarity measure that captures important features of the underlying LPP; next, show how this dissimilarity can be consistently estimated; and finally, extract an estimated mirror that provides a low-dimensional representation of network evolution. Our methodology is not restricted simply to the specific distance measure d_{MV} that we have defined here, and other notions of distance may have conceptual, theoretical, or computational benefits depending on the underlying process.

Remark 6. Using CMDS on the distance matrix as our estimate means that permuting the time points $t_1, \ldots, t_k$ to $t_{\sigma(1)}, \ldots, t_{\sigma(k)}$ where $\sigma \in S_k$ produces the same set of points in $\mathbb{R}^c$ (up to an orthogonal transformation), in a permuted order. As such, *if the original times are associated with these points*, we recover the identical smooth curve. If the times are not retained, only the original time ordering recovers this smooth trajectory.

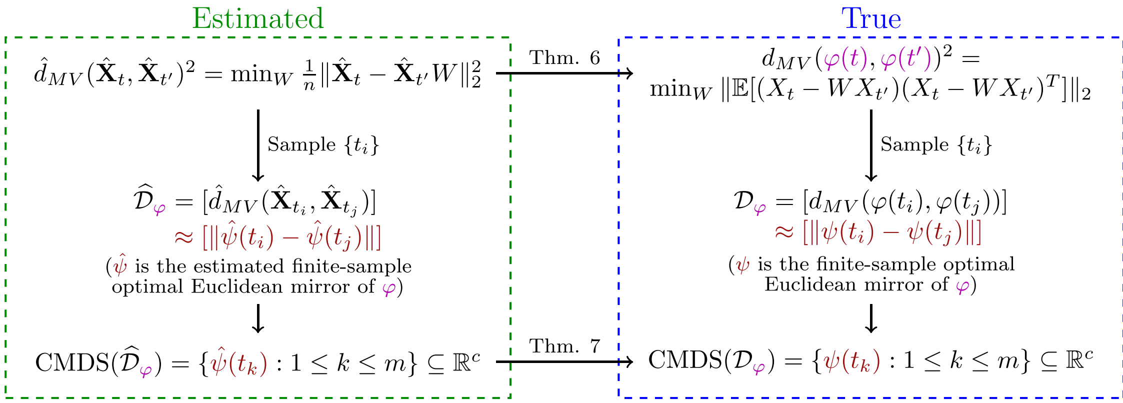

Figure 3. Consistent estimation of network dissimilarity and Euclidean mirrors.

This does not change the content of the theorem: one just replaces t_i with $t_{\sigma(i)}$ in the statement. On the other hand, if we have an *unordered* collection of networks, and not a time series of graphs ordered naturally by time, our estimation procedure will yield a collection of points in $\mathbb{R}^c$, but these will not typically fall on a one-dimensional curve.

When we have exact Euclidean realizability or when the tail eigenvalues of $\mathcal{D}_\varphi$ can be bounded directly, we obtain the following two corollaries of Theorem 7.

Corollary 2. Suppose $\mathcal{D}_\varphi$ is a Euclidean distance matrix with dimension c. Then $\mathcal{D}_\varphi = \mathcal{D}_\psi$, and for fixed m, there is a constant $C = C(m)$ such that with overwhelming probability, there is a real orthogonal matrix $R \in \mathcal{O}^{c\times c}$ such that the CMDS output satisfies

$$\|\hat{U}\hat{S}^{1/2} - US^{1/2}R\|_F \le C\frac{\log(n)}{\sqrt{n}}.$$

In particular, we have

$$\sum_{i=1}^{m} \|\hat{\psi}(t_i) - R\psi(t_i)\|^2 \le C\frac{\log^2(n)}{n}.$$

Suppose that ψ is Lipschitz continuous with constant L and φ has realizability constant B. If we further assume that there exists a constant A such that $d_{MV}(\varphi(t), \varphi(s)) \le A$ and $\|\psi(t) - \psi(s)\|_2 \le A$ for all $s, t \in [0, T]$, then we can bound the sum of the tail eigenvalues of $\mathcal{D}_\varphi^{(2)}$, turning the approximate Lipschitz Euclidean realizability assumption into an eigenvalue bound.

Note that we can always choose $A = (2L + B)T$ from the realizability assumptions, but in certain cases, A may be smaller, and in particular, not grow linearly with T. While this corollary is stated for the Lipschitz case, a version of it may be formulated for the α-Hölder case as well.

Corollary 3. Suppose $\mathcal{D}_\varphi$ is approximately Lipschitz Euclidean c-realizable with realizability constant B. Suppose $d_{MV}(\varphi(t), \varphi(s)), \|\psi(t) - \psi(s)\| \le A$ for all $s, t \in [0, T]$. Suppose that $t_i = Ti/m$ for $i = 1, \ldots, m$. Then

$$\|\mathcal{D}_\varphi^{(2)} - \mathcal{D}_\psi^{(2)}\|_F \le 2ABT\sqrt{(m^2-1)/6} \le 0.82ABTm.$$

For m fixed, there is a constant $C = C(m)$ such that with high probability, there is a rotation matrix $R \in \mathcal{O}^{c\times c}$ such that the CMDS output satisfies

$$\|\hat{U}\hat{S}^{1/2} - US^{1/2}R\|_F \le C\left(\frac{\log(n)}{\sqrt{n}} + 0.82ABT\right).$$

Example 3. In Example 2, which is approximately Lipschitz Euclidean realizable, we have

$$B \le \frac{\sigma^2 T}{a\|v\| + \sqrt{a\|v\| + \sigma^2 T}},$$

and $A \le T\sqrt{a^2\|v\|^2 + \sigma^2 T}$.

Remark 7. The relationship between the true and estimated network features from Figure 3 is equally appropriate for certain changes to the distance metric. For example, consider a latent position process with $\varphi(t) = c(t)\varphi(0)$, $c(t) \in [0, 1]$, corresponding to a global change in the density of the network, but one that leaves the community structure unchanged.

Using the adjacency spectral embedding with the eigenvectors scaled by the eigenvalues will detect these global transformations in sparsity, while using unit eigenvectors of the adjacency spectral embedding (unscaled by their respective eigenvalues) will ignore changes of this type, instead focusing only on divergences in community structure. These different computations of network dissimilarity will result in distinct mirrors, highlighting distinct changes in the networks over time.

Together, these theorems ensure that time-dependent underlying low-dimensional structure associated to network evolution can be consistently recovered. In what follows, we will see how this methodology can be employed in real and synthetic data to reveal important structural features and potential anomalies in network time series.

4. Experiments

4.1. Organizational Network Data and Pandemic-Induced Shifts

We start with a discussion of the visualization of the communication network in Figure 1 and the output of our mirror-based analysis in Figure 2. We consider a time series of weighted communication networks, arising from the E-mail communications between 32,277 entities in a large organization, with one network generated each month from January 2019 to December 2020, a period of 24 months. This data was studied through the lens of modularity in Zuzul et al. (Forthcoming).

To generate the visualization in Figure 1, we first cluster the graph G_1 using Leiden clustering and then apply Node2Vec (Grover and Leskovec 2016) to each graph G_t to obtain embeddings $\hat{\mathbf{X}}_t \in \mathbb{R}^{n\times 33}$. Applying the Uniform Manifold Approximation and Projection (UMAP) algorithm of McInnes, Healy, and Melville (2018) to the $\hat{\mathbf{X}}_t$ matrices, we obtain a layout with points $\hat{\mathbf{X}}'_t \in \mathbb{R}^{n\times 2}$, colored using the labels from the Leiden clustering.

To obtain our mirror estimates, we again apply Leiden clustering (Traag, Waltman, and Van Eck 2019) to the January 2019 network, obtaining 33 clusters that we retain throughout the two year period. We make use of this clustering to compute the Graph Encoder Embedding (GEE) of Shen, Wang, and Priebe (2022) for each time, which produces spectrally-derived estimates of invertible transformations of the original latent positions. For each time $t = 1, \dots, 24$, we obtain a matrix $\hat{\mathbf{Z}}_t \in \mathbb{R}^{32277\times 33}$, each row of which provides an estimate of these transformed latent positions. Constructing the distance matrix $\widehat{\mathcal{D}}_\varphi = [\hat{d}_{MV}(\hat{\mathbf{Z}}_t, \hat{\mathbf{Z}}_{t'})] \in \mathbb{R}^{24\times 24}$, we apply CMDS to obtain the estimated curve $\hat{\psi}$ shown in the left panel of Figure 2, where the choice of dimension $c = 2$ is based on the scree plot of $\widehat{\mathcal{D}}_\varphi$. The nonlinear dimensionality reduction technique ISOMAP (Tenenbaum, de Silva, and Langford 2000), which relies on a spectral decomposition of geodesic distances, can be applied to these points to extract an estimated one-dimensional curve, which we plot against time in the right panel of Figure 2. Since the ISOMAP embedding generates points whose Euclidean distances approximate the geodesic distances between points on the mirror, larger changes in the y-axis of this figure correspond to significant changes in the networks. This one-dimensional curve exhibits some changes from the previous trend in Spring 2020 and a much sharper qualitative transformation in July 2020. What is striking is that both these qualitative shifts correspond to policy changes: in Spring 2020, there was an initial shift in operations, widely regarded at the time as temporary. In mid-summer 2020, nearly the peak of the second wave of COVID-19, it was much clearer that these organizational shifts were likely permanent, or at least significantly longer-lived. In Figure 4, the top panel plots the result of our methods applied to the induced subgraphs corresponding to each of the 33 communities; these are represented by the grey trajectories. The trajectories of two subcommunities have been highlighted in the top plot: the green curve shows a constant rate of change throughout the two-year period, and does not exhibit a noticeable pandemic effect. The blue curve, on the other hand, shows a significant flattening in early 2020, followed by rapid changes in summer. Thus, we can see a differential effect of the pandemic on different work groups within the organization. In Figure 4, both bottom panels show methods for identifying changepoints over the 24 months, with consistent results. We start by generating the ISOMAP embedding ι of $\hat{\psi}$, yielding $\iota_t \in \mathbb{R}$ for $t = 1, \dots, 24$. In the bottom left panel, for each time starting in June 2019, we plot the *sigmage* (see Good 1992) of its ISOMAP embedding relative to the previous 5 months. That is, we measure the distance of the ISOMAP embedding to the mean of the previous 5 months' embeddings, relative to the standard deviation of those embeddings: letting $\mu_t = \frac{1}{5}\sum_{i=1}^{5} \iota_{t-i}$ for each time $t > 5$, the sigmage s_t is given by

$$s_t = \frac{|\iota_t - \mu_t|}{\sqrt{\frac{1}{4}\sum_{i=1}^{5}(\iota_t - \mu_t)^2}}.$$

Note that since the computation of the sigmages requires a window of time-points, we are only able to produce these estimates starting in June 2019. We see apparent outliers in March and April, and again in July–September 2020. The right panel shows the ISOMAP curve with a moving prediction confidence interval of width five standard deviations, generated from simple linear regression applied to the previous five time points (which is why we again only have an interval starting in June). This method indicates the same set of outliers as the previous one, but allows for some more detailed analysis: In March and April 2020, it appears that the behavior is anomalous because the network *stopped* drifting, while the behavior in the summer of that year is anomalous because it made a significant jump from its previous position.

In the supplementary material, we consider additional visualizations of the organizational communication networks. We plot a collection of other summary statistics, namely edge counts, maximum degree, median degree, and modularity, for each network over time. As we describe in that section, since this approach considers each network separately, these summary statistics exhibit greater variance than the ISOMAP embedding of the mirror (Figure 2 right panel, or Figure 4 bottom right panel), and they do not capture changepoints associated to pandemic policy restrictions.

4.2. Synthetic Data and Bootstrapping

In the previous section, we apply the GEE embedding to obtain estimates $\hat{\mathbf{Z}}_t$, which are then used for estimates of pairwise distances $\hat{d}_{MV}(\hat{\mathbf{Z}}_t, \hat{\mathbf{Z}}_s)$. Although the GEE differs slightly from the adjacency spectral embedding, it is computationally more tractable and yields similarly useful output. To further illustrate our underlying theory, however, we consider *synthetic data*. That is, we use real data to obtain a distribution from which we may resample. Such a network bootstrap permits us to test our asymptotic results through replicable simulations that are grounded in actual data. To this end, we consider the true latent position distribution at each time to be equally likely to be any row of the GEE-obtained estimates from the real data, $\hat{\mathbf{Z}}_t \in \mathbb{R}^{32277\times 33}$, for $t = 1, \dots, 24$. Given a sample size n_s, for each time, we sample these rows uniformly and with replacement to get a matrix of latent positions $\mathbf{X}_t \in \mathbb{R}^{n_s\times 33}$. We treat this matrix as the generating latent position matrix for independent adjacency matrices $A_t \sim \mathrm{RDPG}(\mathbf{X}_t)$. Note that if for sample i, we choose row j of $\hat{\mathbf{Z}}_1$ at time $t = 1$, then the same row j of $\hat{\mathbf{Z}}_t$ will be used for all times $t = 1, \dots, 24$ for that sample, so that the original dependence structure is preserved across time. We may now apply the methods described in our theorems, namely ASE of the adjacency matrices followed by Procrustes alignment, to obtain the estimates $\hat{\mathbf{X}}_t$, along with the associated distance

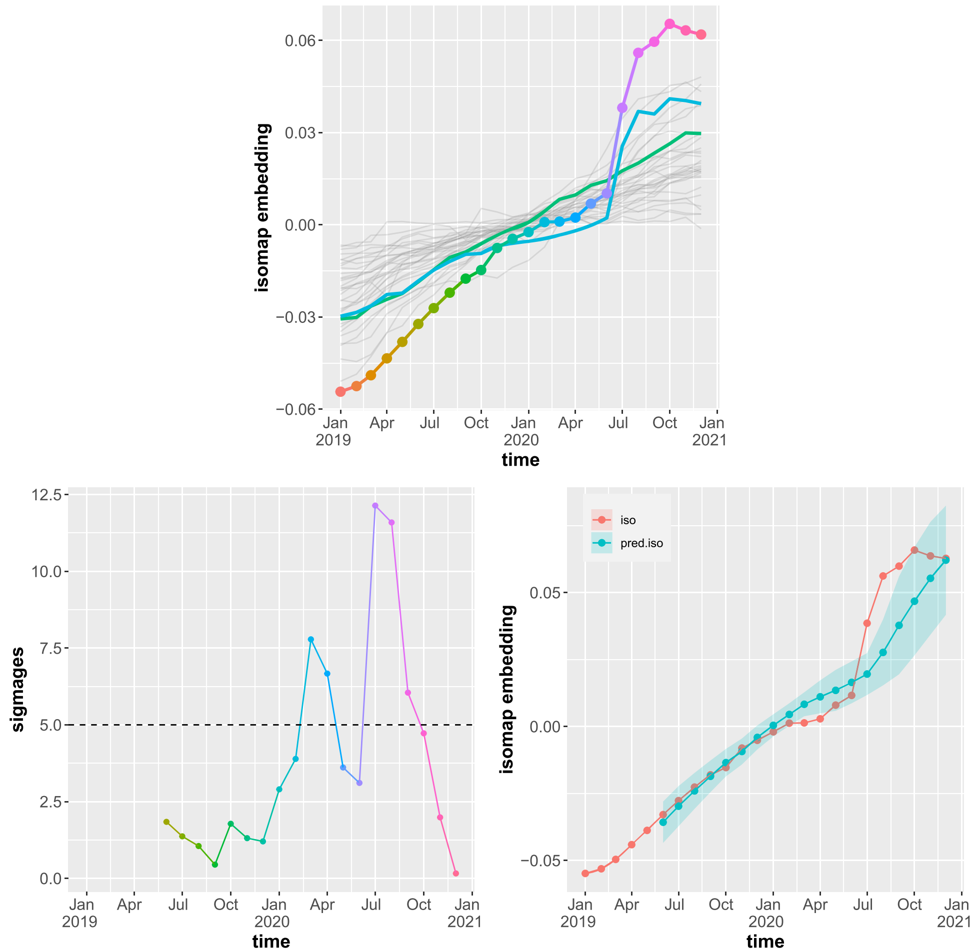

Figure 4. Differential pandemic effect on 33 subcommunities of the larger organization and change point detection. Top: For each subcommunity, we apply our methods to obtain an ISOMAP representation (single grey curve). Green and blue curves correspond to communities highlighted in Figure 1 which exhibit different changes over the time period. Overall network ISOMAP embedding from Figure 2 is overlaid for context. Bottom left panel: Comparing ISOMAP embedding at a given time to mean of previous 5 months' embeddings, measured in terms of their standard deviation. Plot of sigmages indicates clear outliers in March–April of 2020, and again in July–September. Bottom right panel: Observed ISOMAP embedding along with a running confidence interval for the predicted value of the ISOMAP embedding generated from linear regression applied to the embeddings for the previous 5 months, with width five standard deviations.

estimates. In Figure 5, we see that ISOMAP applied to the CMDS embedding of the bootstrapped data converges to the original ISOMAP curve, as predicted by our theorems.

To check whether this procedure demonstrates the pandemic effects, in Figure 6, we show the sigmages for each month, plotted over 100 replicates of this experiment, with $n_s = 30{,}000$ for each replicate. The pandemic effect in summer of 2020 is clearly visible in all but a few replicates, while the effect in March–April is still identified in the majority of replicates. We observe dramatic changes in variance for certain months, over the different replicates: this might indicate the discrepancies between the pandemic effect on different network entities, rendering the final estimate much more sensitive to the sample of rows used to generate the network.

In the appendix of supplementary material, we provide mirror estimates for an evolving stochastic blockmodel with a change in rank from 2 to 1. We demonstrate the mirror estimation procedure and model misspecification in the embedding dimension, specifically the accuracy of the first dimension of the estimated distance matrix and the noise in the second dimension at and after the collapse to a rank 1 model.

Code for these results and additional simulations can be found at *https://www.cis.jhu.edu/~parky/IsoMirror/dynamics.html*.

5. Discussion

To effectively model time series of networks, it is natural to consider network evolution governed by underlying low-dimensional dynamics. Here, we examine latent position networks in which the vertex latent positions follow a stochastic process known as a latent position process (LPP). Under mild conditions, we can associate to the LPP certain geometric structure, and understanding how that structure changes with time allows us to identify transformations in network behavior across multiple scales. To make this precise, we define the maximum directional variation norm and metric on the space of random latent positions. We describe notions of Euclidean realizability and Euclidean mirrors for this metric and process, characterizing how closely this metric can be approximated by a Euclidean distance.

Of course, the latent position process is typically unobserved; what we have instead is a time series of networks from which

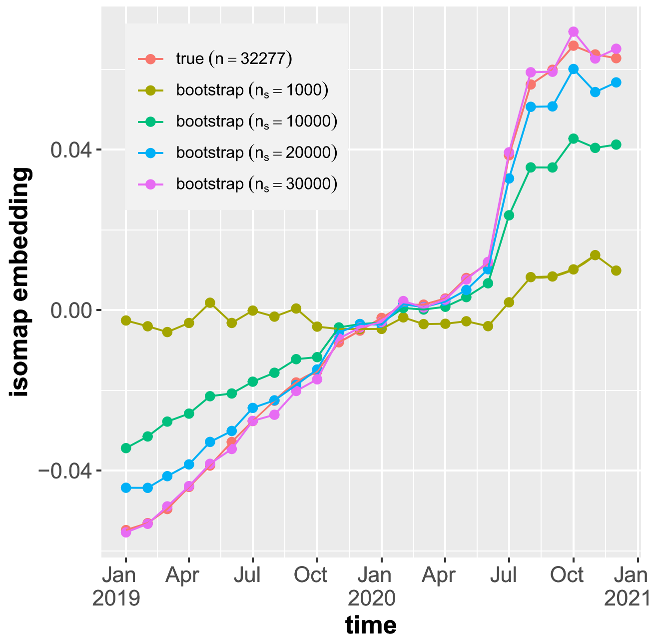

Figure 5. Convergence of bootstrapped estimates. We collect four bootstrapped time series of graphs at various network sizes n_S, sampling the latent positions with replacement from the real data of the previous section, and generating the networks as random dot product graphs with these synthetic latent position matrices. ISOMAP on the resulting network embeddings shows that the bootstrapped curve converges to the original ISOMAP curve as the number of samples n_S increases.

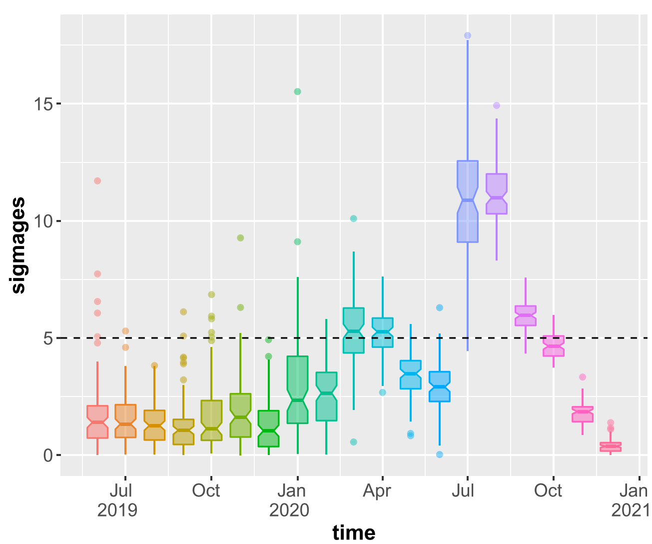

Figure 6. Pandemic effect recovered from synthetic data. For each of 100 replicates of bootstrapped data, with $n_S = 30{,}000$ for each replicate, we repeat procedure in bottom left panel of Figure 4. Sigmages plotted in a box-and-whisker plot. Pandemic effect in summer of 2020 is visible in all but a few replicates; effect in March–April is still identified in the majority of replicates.

these latent positions must be estimated. One of our key results is that the pairwise dissimilarity matrix of maximum directional variation distances between latent positions X_t and $X_{t'}$ at pairs of time points can be consistently estimated by spectrally embedding the network adjacencies at these different pairs of times and computing spectral norm distances between these embeddings. When the latent position process is such that the maximum directional variation metric between any pair of latent positions $X_t, X_{t'}$ is approximately Euclidean realizable, we find that classical multidimensional scaling applied to the estimated distances gives us an inferentially valuable low-dimensional representation of network dissimilarities across time. Further dimension reduction techniques, such as ISOMAP, can further clarify changes in network dynamics. To this last point, ISOMAP is a manifold-learning algorithm; detailed analysis of its effect on embeddings of estimated pairwise network distances can bring us closer to provable guarantees for changepoint detection. More broadly, the interplay between the probabilistic structure of the underlying latent position process and the geometric structure of the Euclidean mirror is a key component of the estimated Euclidean representation of relationships between networks across time.

We consider two estimates for the maximum directional variation distance d_{MV}, namely the spectral norm applied to the GEE estimate, and the $\hat{d}_{MV}$ estimated distance between the adjacency spectral embeddings. However, these are far from the only options, and it is an open question whether d_{MV} or another metric on the space of random variables is best for downstream inference tasks under certain model assumptions. It is also an open question whether there is a better estimate for the distance d_{MV} itself, either in terms of computational complexity or statistical properties. Of particular interest is the spectral norm distance applied to the omnibus embeddings (Levin et al. 2017) for the adjacency matrices: this likely converges to another distance on the space of random variables, potentially highlighting different features in the final CMDS embedding. Results quantifying the distribution of the errors in the CMDS embedding are key to formulating hypothesis tests for changepoint detection. The perspective described in Figure 3, which connects distance metrics for generative processes of networks to their estimates, translating manifold geometry into Euclidean geometry, is a useful contribution to time series analysis for networks. It provides mathematical formalism for network dynamics; asymptotic properties of estimates of manifold structure; and conditions for the representation of time-varying networks in low-dimensional space. Latent position networks are interpretable, estimable, and flexible enough to capture important features of real-world network time series. As such, this canonical framework invites and accommodates future approaches to joint network inference.

Supplementary Materials

The supplementary material includes proofs and additional simulations and mirror computations for latent position process networks and evolving stochastic blockmodel.

Disclosure Statement

No potential conflict of interest was reported by the author(s).

Funding

All authors gratefully acknowledge funding from Microsoft Research, the Office of Naval Research via the JHU Applied Physics Laboratory (N00024-22-D-6404), the Naval Engineering Education Consortium, the United States National Science Foundation (SES-1951005 and DMS-2113099), the Amazon-JHU AI2AI Initiative, and the Acheson Duncan Fund.

References

Arroyo, J., Athreya, A., Cape, J., Chen, G., Priebe, C. E., and Vogelstein, J. T. (2021), "Inference for Multiple Heterogeneous Networks with a Common Invariant Subspace," *Journal of Machine Learning Research*, 22, 1–49. [1025]

Borg, I., and Groenen, P. J. F. (2005), *Modern Multidimensional Scaling: Theory and Applications*, New York: Springer. [1027]

Gallagher, I., Jones, A., and Rubin-Delanchy, P. (2021), "Spectral Embedding for Dynamic Networks with Stability Guarantees," in *Advances in Neural Information Processing Systems* (Vol. 34), eds. M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, pp. 10158–10170, Curran Associates, Inc. [1025]

Good, I. J. (1992), "The Bayes/Non-Bayes Compromise: A Brief Review," *Journal of the American Statistical Association*, 87, 597–606. [1033]

Grover, A., and Leskovec, J. (2016), "node2vec: Scalable Feature Learning for Networks," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 855–864. [1033]

Hoff, P. D., Raftery, A. E., and Handcock, M. S. (2002), "Latent Space Approaches to Social Network Analysis," *Journal of the American Statistical Association*, 97, 1090–1098. [1025]

Jing, B.-Y., Li, T., Lyu, Z., and Xia, D. (2021), "Community Detection on Mixture Multilayer Networks via Regularized Tensor Decomposition," *The Annals of Statistics*, 49, 3181–3205. [1025]

Jones, A., and Rubin-Delanchy, P. (2020), "The Multilayer Random Dot Product Graph," arXiv preprint arXiv:2007.10455. [1025]

Krivitsky, P. N., and Handcock, M. S. (2014), "A Separable Model for Dynamic Networks," *Journal of the Royal Statistical Society*, Series B, 76, 29–46. [1025]

Lee, J. D., and Maggioni, M. (2011), "Multiscale Analysis of Time Series of Graphs," in *International Conference on Sampling Theory and Applications (SampTA)*, Citeseer. [1025]

Lei, J., Chen, K., and Lynch, B. (2020), "Consistent Community Detection in Multi-Layer Network Data," *Biometrika*, 107, 61–73. [1025]

Lei, J., and Lin, K. Z. (2022), "Bias-Ajusted Spectral Clustering in Multi-Layer Stochastic Block Models," *Journal of the American Statistical Association*, 118, 2433–2445. [1025]

Levin, K., Athreya, A., Tang, M., Lyzinski, V., and Priebe, C. E. (2017), "A Central Limit Theorem for an Omnibus Embedding of Random Dot Product Graphs," arXiv preprint arXiv:1705.09355. [1025,1035]

McInnes, L., Healy, J., and Melville, J. (2018), "Umap: Uniform Manifold Approximation and Projection for Dimension Reduction," arXiv preprint arXiv:1802.03426. [1033]

Padilla, O. H. M., Yu, Y., and Priebe, C. E. (2022), "Change Point Localization in Dependent Dynamic Nonparametric Random Dot Product Graphs," *The Journal of Machine Learning Research*, 23, 10661–10719. [1025]

Pantazis, K., Athreya, A., Arroyo, J., Frost, W. N., Hill, E. S., and Lyzinski, V. (2022), "The Importance of Being Correlated: Implications of Dependence in Joint Spectral Inference Across Multiple Networks," *Journal of Machine Learning Research*, 23, 1–77. [1025]

Pensky, M. (2019), "Dynamic Network Models and Graphon Estimation," *The Annals of Statistics*, 47, 2378–2403. [1025]

Shen, C., Wang, Q., and Priebe, C. E. (2022), "One-Hot Graph Encoder Embedding," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45, 7933–7938. [1033]

Smith, A. L., Asta, D. M., and Calder, C. A. (2019), "The Geometry of Continuous Latent Space Models for Network Data," *Statistical Science: A Review Journal of the Institute of Mathematical Statistics*, 34, 428–453. [1026]

Sussman, D. L., Tang, M., Fishkind, D. E., and Priebe, C. E. (2012), "A Consistent Adjacency Spectral Embedding for Stochastic Blockmodel Graphs," *Journal of the American Statistical Association*, 107, 1119–1128. [1030]

Tenenbaum, J. B., de Silva, V., and Langford, J. (2000), "A Global Geometric Framework for Nonlinear Dimensionality Reduction," *Science*, 290, 2319–2323. [1033]

Traag, V. A., Waltman, L., and Van Eck, N. J. (2019), "From Louvain to Leiden: Guaranteeing Well-Connected Communities," *Scientific Reports*, 9, 5233. [1026,1033]

Wang, D., Yu, Y., and Rinaldo, A. (2021), "Optimal Change Point Detection and Localization in Sparse Dynamic Networks," *The Annals of Statistics*, 49, 203–232. [1025]

Wilkins-Reeves, S., and McCormick, T. (2022), "Asymptotically Normal Estimation of Local Latent Network Curvature," arXiv preprint arXiv:2211.11673. [1026]

Zhang, A., and Xia, D. (2018), "Tensor svd: Statistical and Computational Limits," *IEEE Transactions on Information Theory*, 64, 7311–7338. [1025]

Zuzul, T., Pahnke, E. C., Larson, J., Bourke, P., Caurvina, N., Shah, N. P., Amini, F., Park, Y., Vogelstein, J., Weston, J., et al. (Forthcoming), "Dynamic Silos: Increased Modularity in Intra-Organizational Communication Networks During the Covid-19 Pandemic," *Management Science*, 1–21. [1025,1033]