

Weighted Random Dot Product Graphs

Bernardo Marengo¹ , Paola Bermolen¹ , Marcelo Fiori¹ ,
Federico Larroca¹  and Gonzalo Mateos² 

¹*Facultad de Ingeniería, Universidad de la República, Uruguay, e-mail: bmarengo@fing.edu.uy; paola@fing.edu.uy; mfiori@fing.edu.uy; flarroca@fing.edu.uy*

²*Dept. of Electrical and Computer Engineering, University of Rochester, Rochester, NY, USA, e-mail: gmateosb@ur.rochester.edu*

Abstract: Modeling of intricate relational patterns has become a cornerstone of contemporary statistical research and related data science fields. Networks, represented as graphs, offer a natural framework for this analysis. This paper extends the Random Dot Product Graph (RDPG) model to accommodate weighted graphs, markedly broadening the model’s scope to scenarios where edges exhibit heterogeneous weight distributions. We propose a nonparametric weighted (W)RDPG model that assigns a sequence of latent positions to each node. Inner products of these nodal vectors specify the moments of their incident edge weights’ distribution via moment-generating functions. In this way, and unlike prior art, the WRDPG can discriminate between weight distributions that share the same mean but differ in other higher-order moments. We derive statistical guarantees for an estimator of the nodal’s latent positions adapted from the workhorse adjacency spectral embedding, establishing its consistency and asymptotic normality. We also contribute a generative framework that enables sampling of graphs that adhere to a (prescribed or data-fitted) WRDPG, facilitating, e.g., the analysis and testing of observed graph metrics using judicious reference distributions. The paper is organized to formalize the model’s definition, the estimation (or nodal embedding) process and its guarantees, as well as the methodologies for generating weighted graphs, all complemented by illustrative and reproducible examples showcasing the WRDPG’s effectiveness in various network analytic applications.

MSC2020 subject classifications: Primary 05C80, 62G05; secondary 60F05.

Keywords and phrases: Random weighted graph, latent positions model, random dot product graph, statistical network analysis.

Received May 2025.

1. Introduction

The study of network data has become a central theme in modern statistical research due to the growing demands to understand complex relational structures found in social, biological, and technological systems; see e.g., [15]. Networks, represented as graphs $G = (V, E)$ with vertex set V and edge set E , offer a natural framework to describe pairwise interactions among entities. Statistical approaches to the analysis of network data not only address issues of modeling, inference, and prediction, but also provide insights into the underlying mechanisms governing connectivity and community structure in complex systems.

A powerful approach within this domain is the use of latent variable models [9], where each vertex is endowed with an unobserved feature vector that encapsulates its propensity to

establish relational ties. This general class of models offers a probabilistic lens through which the observed network structure can be understood in terms of latent geometric or probabilistic constructs, e.g., [14, Ch. 2.2] and [12, Ch. 10.7]. Among these models, the Random Dot Product Graph (RDPG) stands out for its interpretability and theoretical tractability, providing a bridge between latent geometry and observed connectivity patterns.

In the standard RDPG model, we consider a simple undirected and unweighted random graph $G = (V, E)$ with N vertices, where each vertex $i \in V$ is associated with a latent position $\mathbf{x}_i \in \mathbb{R}^d$. Typically, the dimensionality of the embedding space $d \ll N$. The RDPG model posits that for any pair of vertices i and j , the presence of an edge is determined by a Bernoulli random variable $A_{ij} \in \{0, 1\}$ with success probability given by the dot product of the latent positions, that is,

$$A_{ij} \sim \text{Bernoulli}(\mathbf{x}_i^\top \mathbf{x}_j), \quad \text{for } i \neq j,$$

where the A_{ij} 's are mutually independent. Since self loops are excluded, $A_{ii} \equiv 0$ for all $i \in V$. Here, $\mathbf{A} \in \{0, 1\}^{N \times N}$ is the graph's symmetric adjacency matrix, and the inner product $\mathbf{x}_i^\top \mathbf{x}_j$ must lie in the interval $[0, 1]$ to yield as a valid probability. RDPGs can also incorporate random latent positions by defining a distribution F supported on $\mathcal{X} \subset \mathbb{R}^d$, such that $\mathbf{x}^\top \mathbf{y} \in [0, 1]$ for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$.

One can verify that Erdős-Rényi (ER) graphs, stochastic block models (SBMs) with a positive semidefinite (PSD) probability matrix, or other more sophisticated random graph models are particular cases of an RDPG [1]. The resulting expressiveness has been a main driver behind the popularity of the model. Moreover, the RDPG definition provides a natural geometric interpretation of connectivity: angles between latent position vectors indicate affinity between vertices, and their magnitudes indicate how well connected they are. For $d \leq 3$, visual inspection of the nodes' vector representations can reveal community structure. For higher dimensions or more complex scenarios, angle-based clustering of nodal embeddings can also be used [18, 25].

The associated RDPG inference problem of estimating latent positions from graph observations enjoys strong asymptotic properties, facilitating statistical inference tasks by bringing to bear tools of geometrical data analysis in latent space [1]. Let us briefly describe the inference method in such a setup. Consider stacking all the nodes' latent position vectors in the rows of matrix $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]^\top \in \mathbb{R}^{N \times d}$. Given an observed graph \mathbf{A} and a prescribed embedding dimension d (typically obtained using an elbow rule on \mathbf{A} 's eigenvalue scree plot, using e.g. [37]), the goal is to estimate \mathbf{X} . For the RDPG model, the edge-formation probabilities are given by the entries $P_{ij} = \mathbf{x}_i^\top \mathbf{x}_j$ of the PSD matrix $\mathbf{P} = \mathbf{X}\mathbf{X}^\top$. Latent positions are uniquely defined only up to an orthogonal transformation, since for any orthogonal matrix $\mathbf{Q} \in O(d)$, the rotated positions $\mathbf{Y} = \mathbf{X}\mathbf{Q}$ satisfy $\mathbf{P} = \mathbf{Y}\mathbf{Y}^\top$ as well. The Adjacency Spectral Embedding (ASE) estimator solves the following least-squares approximation [25] to obtain

$$\hat{\mathbf{X}} \in \underset{\mathbf{X} \in \mathbb{R}^{N \times d}}{\text{argmin}} \|\mathbf{A} - \mathbf{X}\mathbf{X}^\top\|_F^2. \tag{1}$$

In other words, $\hat{\mathbf{P}} = \hat{\mathbf{X}}\hat{\mathbf{X}}^\top$ is the best rank- d PSD approximation to the adjacency matrix \mathbf{A} , in the Frobenius-norm sense. It is straightforward to see that $\hat{\mathbf{X}} = \hat{\mathbf{U}}\hat{\mathbf{D}}^{1/2}$ verifies (1), where $\mathbf{A} = \mathbf{U}_\mathbf{A}\mathbf{D}_\mathbf{A}\mathbf{U}_\mathbf{A}^\top$ is the eigendecomposition of \mathbf{A} , $\hat{\mathbf{D}} \in \mathbb{R}^{d \times d}$ is a diagonal matrix with the d

largest-magnitude eigenvalues of \mathbf{A} , and $\hat{\mathbf{U}} \in \mathbb{R}^{N \times d}$ are the associated eigenvectors from $\mathbf{U}_{\mathbf{A}}$.¹ Consistency and asymptotic ($N \rightarrow \infty$) normality results for ASE are available; see e.g., [1].

In this work we extend the RDPG model to account for *weighted* graphs, endowed with a weight map $w : E \mapsto \mathbb{R}_+$ that assigns a nonnegative value to each edge. The adjacency matrix entries for the weighted graph are $W_{ij} = W_{ji} = w(i, j)$ for all $(i, j) \in E$, while the absence of an edge is represented as $W_{ij} = W_{ji} = 0$. This formulation naturally encompasses unweighted graphs as a special case, where edge weights are constrained to binary values (i.e., $w \equiv 1$).

Several extensions of the standard (or *vanilla*) RDPG model have been developed to incorporate edge weights. A common approach is to introduce a parametric distribution F_{θ} with parameter $\theta \in \mathbb{R}^L$ to model weighted adjacency entries [7, 30]. For instance, in the case of Poisson-distributed weights, one might set $L = 1$ and assume $W_{ij} \sim \text{Poisson}(\theta)$. In this formulation, each node $i \in V$ is assigned a collection of latent vectors $\mathbf{x}_i[l] \in \mathbb{R}^d$ for $l = 1, \dots, L$, such that the weight of an edge between nodes i and j follows the distribution

$$W_{ij} \sim F_{(\mathbf{x}_i^{\top}[1]\mathbf{x}_j[1], \dots, \mathbf{x}_i^{\top}[L]\mathbf{x}_j[L])},$$

independently across edges. To model sparse graphs, this distribution may include a point mass at zero. The classic RDPG model is recovered by setting F_{θ} as a Bernoulli distribution with success probability $\theta \in [0, 1]$. Despite its flexibility, this approach has some noteworthy limitations. A key drawback is that all edges must share the same parametric family of weight distributions, differing only in their parameters. While mixture models can introduce some heterogeneity, they still require an explicit assumption about the underlying distribution families. This requirement significantly restricts the model's applicability, particularly if edges follow heterogeneous, unknown, or multimodal weight distributions.

To address this shortcoming, recent work by Gallagher et al. [8] proposes a nonparametric alternative. Therein each node has a single associated latent position $\mathbf{z}_i \in \mathcal{Z}$, which is endowed with a probability distribution F . It is postulated that, given a family $\{H(\mathbf{z}_1, \mathbf{z}_2) : \mathbf{z}_1, \mathbf{z}_2 \in \mathcal{Z}\}$ of symmetric real-valued distributions, there exists a map $\phi : \mathcal{Z} \mapsto \mathbb{R}^d$ such that if $W_{ij} \sim H(\mathbf{z}_i, \mathbf{z}_j)$, then $\mathbb{E}[W_{ij}] = \phi^{\top}(\mathbf{z}_i)\mathbf{I}_{p,q}\phi(\mathbf{z}_j)$, where $\mathbf{I}_{p,q}$ is a diagonal matrix of p ones and q minus ones, such that $p + q = d$. This diagonal matrix facilitates modeling heterophilic (or disassortative) behavior, as in [23]. Interestingly, one can consistently recover $\mathbf{x}_i = \phi(\mathbf{z}_i)$ via the ASE of \mathbf{W} , and the estimated $\hat{\mathbf{x}}_i$'s are asymptotically normal. While this method improves flexibility and avoids restrictive parametric assumptions, it can only recover $\phi(\mathbf{z}_i)$, i.e., the latent positions for the mean adjacency matrix $\mathbb{E}[\mathbf{W}]$. Consequently, the model is unable to differentiate between pairs of edges that stem from distinct distributions sharing an identical mean. These challenges highlight the ongoing need for more expressive, discriminative, and robust latent space models for weighted graphs.

1.1. Proposed approach and contributions

Instead, we propose that the sequence of nodal vectors be related to the weight distribution's moment-generating function (MGF). Assume each node has a sequence of latent positions $\{\mathbf{x}_i[k]\}_{k \geq 0}$ such that the inner products $\{\mathbf{x}_i^{\top}[k]\mathbf{x}_j[k]\}_{k \geq 0}$ form an admissible moment sequence. Our weighted (W)RDPG model specifies the k -th order moments of the random

¹Under mild assumptions on the latent positions distribution F it is possible to show that the top d eigenvalues of \mathbf{A} are nonnegative with probability tending to 1 as $N \rightarrow \infty$, so $\hat{\mathbf{X}}$ is well defined, see [1] for details.




Fig 1. Latent position estimation. Given an adjacency matrix \mathbf{A} we compute its k -th entry-wise power $\mathbf{A}^{(k)}$. The ASE of $\mathbf{A}^{(k)}$ yields the estimates $\hat{\mathbf{X}}[k]$; see also Section 2.2.

entries in the adjacency matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$ are given by $\mathbb{E} \left[W_{ij}^k \right] = \mathbf{x}_i^T [k] \mathbf{x}_j [k]$, for all $k \geq 0$ (see Section 2.1).

The WRDPG model offers several key advantages: (i) it is nonparametric; (ii) it considers higher-order moments beyond the mean; while providing (iii) statistical guarantees for the associated estimator; and (iv) a generative mechanism that can reproduce the structure and the weights’ distribution of real networks. Next, we elaborate on these attractive features to better position our technical contributions in context. Unlike [7, 30], the nonparametric WRDPG graph modeling framework does not require prior assumptions about a specific weight distribution. Nodal vectors can be used to describe high-order moments of said distribution, thus enhancing the model’s representation and discriminative power. For example, different from [8] the WRDPG model can distinguish between distributions with the same mean but differing standard deviations; see also Section 2.4 and Figure 5 for an illustrative example.

Furthermore, framing our model in terms of (symmetric) adjacency matrices allows us to define latent position estimators based on their spectral decomposition; see Section 2.2. Statistical guarantees can be derived using tools to control the matrix spectrum under perturbations. As shown in [19], for each fixed moment index k , the latent position matrix $\mathbf{X}[k]$ can be consistently recovered (up to an orthogonal transformation) via the ASE of $\mathbf{W}^{(k)}$, the matrix of entry-wise k -th powers of \mathbf{W} ; see Figure 1 for a schematic diagram of the latent position estimation procedure. In Section 3 we generalize the result in [19] by proving that consistency holds for unbounded sub-Weibull edge weights under a stronger convergence metric that ensures uniform convergence for all $\hat{\mathbf{x}}_i$ ’s. We also establish that the estimator is asymptotically normal as $N \rightarrow \infty$, a result that is novel in our WRDPG setup.

In Section 4 we show how one can generate graphs adhering to the proposed WRDPG model, which is non-trivial in a nonparameteric setting and was not considered in [8]. We develop a methodology to sample random weighted graphs whose edge weight distribution is defined by a sequence of moments given by inner products of connected nodes’ latent positions. Depending on the characteristics of the weight distribution—whether discrete, continuous, or a mixture—we can solve for the latent position sequence to ensure the inner products match the prescribed moments. For real-valued (continuous) weights, we rely on the maximum entropy principle subject to moment constraints. We introduce a primal-dual approach to find a probability density function that maximizes entropy, offering improved numerical stability relative to previous algorithms [24]. Results for mixture distributions allow us to generate graphs that simultaneously replicate the sparsity pattern of the network and the edge weights. In Figure 2 we present a schematic diagram for the graph generation pipeline.

Moreover, given an observed weighted network (instead of the ground-truth latent position sequence as in Figure 2), we show how we can use this generative procedure to sample graphs

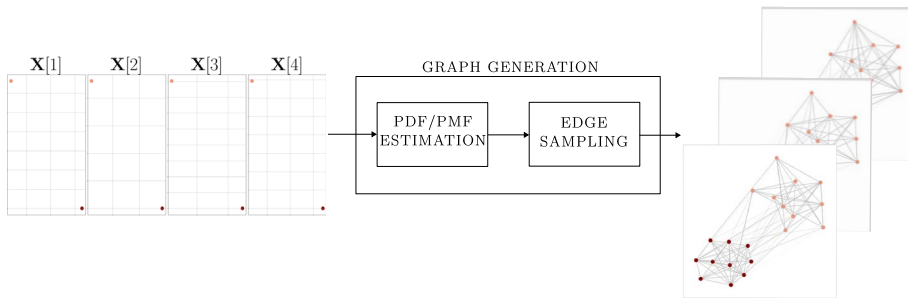


Fig 2. *Graph generation.* Given the latent positions of each vertex $\{\mathbf{X}[k]\}_{k \geq 0}$, we estimate a weight distribution whose sequence of moments is given by the corresponding dot products. Edge weights are then sampled from this estimated distribution; see also Section 4.

that are similar to the original one, in a well-defined statistical sense. This can be useful for several statistical inference tasks involving network data. If, for instance, one would like to assess the significance of some observed graph characteristic, this procedure allows to generate several ‘comparable’ graphs and construct judicious reference distributions [15, Section 6.2]. In a related use case, one could perform hypothesis testing to determine if a given graph adheres to the WRDPG model.

The rest of the paper is organized as follows. In Section 2 we formally define the WRDPG model, describe the embedding method, and present several illustrative examples. The estimation problem is addressed in Section 3, and the proofs of the asymptotic statistical guarantees are presented in detail. Finally, in Section 4, we present our methodology for generating WRDPG graphs. We include several reproducible examples to demonstrate the ability of the generative framework to produce weighted graphs with desired characteristics.

Notational conventions

Throughout the paper we use the following notation. Real numbers are denoted by plain symbols (either upper or lowercase), such as $i, j, N \in \mathbb{R}$. Vectors in a Euclidean space \mathbb{R}^d are denoted in bold, lowercase, e.g., $\mathbf{x} \in \mathbb{R}^d$, and are assumed to be column vectors. We reserve bold uppercase letters for matrices, such as $\mathbf{A} \in \mathbb{R}^{n \times m}$, with \mathbf{I} denoting the identity matrix. The symbol $\mathbf{A}^{(k)}$ stands for the k -fold Hadamard (i.e., entrywise) product of matrix \mathbf{A} with itself, that is:

$$\mathbf{A}^{(k)} := \underbrace{\mathbf{A} \circ \mathbf{A} \circ \dots \circ \mathbf{A}}_{k \text{ times}}$$

The group of orthogonal matrices of dimension d is denoted by $O(d)$, i.e., $\mathbf{Q} \in O(d)$ iff $\mathbf{Q} \in \mathbb{R}^{d \times d}$ and $\mathbf{Q}\mathbf{Q}^\top = \mathbf{I}$, where $(\cdot)^\top$ stands for transposition.

The dot product between $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ is denoted by $\mathbf{x}^\top \mathbf{y}$, whereas $\|\mathbf{x}\|_2$ and $\|\mathbf{x}\|_\infty$ denote the Euclidean (i.e., L_2) and maximum (i.e., L_∞) norm of vector \mathbf{x} . When dealing with matrices, $\|\mathbf{A}\|_F$, $\|\mathbf{A}\|_{\max}$, $\|\mathbf{A}\|_2$ and $\|\mathbf{A}\|_{2 \rightarrow \infty}$ denote the Frobenius, maximum (i.e., largest entry in magnitude), spectral (i.e., largest singular value) and $2 \rightarrow \infty$ norm of matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$, respectively. The latter is defined as

$$\|\mathbf{A}\|_{2 \rightarrow \infty} := \sup_{\|\mathbf{x}\|_2=1} \|\mathbf{A}\mathbf{x}\|_\infty,$$

see [5]. The $2 \rightarrow \infty$ norm is thus the operator norm of \mathbf{A} acting as $\mathbf{A}(\cdot) : \mathbb{R}^n \mapsto \mathbb{R}^m$, when \mathbb{R}^n is endowed with the Euclidean norm and \mathbb{R}^m with the ∞ norm.

We repeatedly make use of the following relations between the $2 \rightarrow \infty$, spectral, and Frobenius norms of any matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$:

$$\begin{aligned} \|\mathbf{A}\|_{2 \rightarrow \infty} &\leq \|\mathbf{A}\|_2 \leq \|\mathbf{A}\|_F, \\ \|\mathbf{A}\|_2 &\leq \min \{ \sqrt{n} \|\mathbf{A}\|_{2 \rightarrow \infty}, \sqrt{m} \|\mathbf{A}^\top\|_{2 \rightarrow \infty} \}. \end{aligned}$$

The $2 \rightarrow \infty$ norm is not submultiplicative; however, the following relations hold for any matrices $\mathbf{A}, \mathbf{B}, \mathbf{C}$ of suitable dimensions:

$$\begin{aligned} \|\mathbf{AB}\|_{2 \rightarrow \infty} &\leq \|\mathbf{A}\|_{2 \rightarrow \infty} \|\mathbf{B}\|_2 \leq \|\mathbf{A}\|_2 \|\mathbf{B}\|_2, \\ \|\mathbf{CA}\|_{2 \rightarrow \infty} &\leq \|\mathbf{C}\|_\infty \|\mathbf{A}\|_{2 \rightarrow \infty} \leq \sqrt{n} \|\mathbf{C}\|_2 \|\mathbf{A}\|_2, \end{aligned}$$

where $\|\mathbf{A}\|_\infty := \max_i \sum_j |A_{ij}|$ is the maximum absolute row sum of \mathbf{A} .

For a function $f : \mathbb{N} \mapsto \mathbb{R}^+$, we say that a random variable $Y \in \mathbb{R}$ is $\mathcal{O}_{\mathbb{P}}(f(N))$ if for each $\alpha > 0$ there exists $N_0 \in \mathbb{N}$ and $C > 0$ such that for all $N \geq N_0$ it holds that $\mathbb{P}[|Y| < Cf(N)] \geq 1 - N^{-\alpha}$. In that case we write $Y = \mathcal{O}_{\mathbb{P}}(f(N))$. Furthermore, we say that Y is $\Theta_{\mathbb{P}}(f(N))$ if for each $\alpha > 0$ there exists $N_0 \in \mathbb{N}$ and $c, C > 0$ such that for all $N \geq N_0$ it holds that $\mathbb{P}[cf(N) < Y < Cf(N)] \geq 1 - N^{-\alpha}$. In that case we write $Y = \Theta_{\mathbb{P}}(f(N))$.

2. Weighted RDPG model

Here we define the WRDPG model and the ASE-based latent position estimator. Illustrative examples are presented to ground these concepts. Furthermore, we show the discriminative power for community detection inherited from considering higher-order moments beyond the mean. Finally, we discuss the impact of the number of nodes on the accuracy of the moment sequence reconstruction.

2.1. Model specification

In this section, we formally define our WRDPG model. We follow the rationale of the vanilla RDPG and define the model in terms of random latent position sequences per node. As previewed in Section 1, the inner product of the latent positions will be related to the MGF of the edge weights' distribution, thus requiring some preliminary definitions.

Definition 2.1 (Admissible moment sequence). A sequence $\{m[k]\}_{k \geq 0}$ of real numbers is an admissible moment sequence if $m[0] = 1$ and for all $p \geq 0$ the matrix

$$\mathbf{M} = \begin{pmatrix} m[0] & m[1] & m[2] & \dots & m[p] \\ m[1] & m[2] & m[3] & \dots & m[p+1] \\ m[2] & m[3] & m[4] & \dots & m[p+2] \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ m[p] & m[p+1] & m[p+2] & \dots & m[2p] \end{pmatrix} \in \mathbb{R}^{(p+1) \times (p+1)}$$

whose entries are $M_{ij} = m[i + j - 2]$ is PSD.

Remark 2.1. Definition 2.1 is enough to guarantee that there exists a probability measure μ in \mathbb{R} such that $m[k]$ is the k -th moment of μ ; see e.g., [35, Ch. 3].

Definition 2.2 (Weighted inner-product distribution). Let F be a probability distribution with support $\text{supp}F = \mathcal{X} \subset (\mathbb{R}^d)^\infty$. We say that F is a weighted inner-product distribution if for all $\{\mathbf{x}[k]\}_{k \geq 0}, \{\mathbf{y}[k]\}_{k \geq 0} \in \mathcal{X}$, the sequence $\{\mathbf{x}^\top[k]\mathbf{y}[k]\}_{k \geq 0}$ is an admissible moment sequence.

Definition 2.3 (WRDPG). Let F be a weighted inner-product distribution and let $\{\mathbf{x}_1[k]\}_{k \geq 0}, \{\mathbf{x}_2[k]\}_{k \geq 0}, \dots, \{\mathbf{x}_N[k]\}_{k \geq 0}$ be i.i.d. with distribution F . Given the sequence $\mathbf{X}_k := \{\mathbf{X}[k]\}_k$, with $\mathbf{X}[k] = [\mathbf{x}_1[k], \dots, \mathbf{x}_N[k]]^\top \in \mathbb{R}^{N \times d}$, the weighted (W)RDPG model prescribes the k -th order moments of a random adjacency matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$ are given by $\mathbb{E} \left[W_{ij}^k \right] = \mathbf{x}_i^\top[k]\mathbf{x}_j[k]$, for all $k \geq 0$. Formally, the WRDPG model specifies the MGF of $\mathbf{W} \in \mathbb{R}^{N \times N}$ as

$$\mathbb{E} \left[e^{tW_{ij}} | \mathbf{X}_k \right] = \sum_{k=0}^{\infty} \frac{t^k \mathbb{E} \left[W_{ij}^k \right]}{k!} = \sum_{k=0}^{\infty} \frac{t^k \mathbf{x}_i^\top[k]\mathbf{x}_j[k]}{k!} \tag{2}$$

and the entries W_{ij} are independent, i.e., edge independent. In such a case, we write $(\mathbf{W}, \mathbf{X}_k) \sim \text{WRDPG}(F)$.

Remark 2.2 (Nonidentifiability of latent positions). As is the case for the vanilla RDPG model, latent positions are invariant to orthogonal transformations, since for any $\mathbf{Q} \in O(d)$, $\mathbf{X}[k]\mathbf{X}^\top[k] = \mathbf{X}[k]\mathbf{Q}(\mathbf{X}[k]\mathbf{Q})^\top$. This implies that given a latent position sequence \mathbf{X}_k , for each index k one may choose a matrix $\mathbf{Q}_k \in O(d)$ (which may vary with k) and construct a sequence $\mathbf{Y}_k := \{\mathbf{X}[k]\mathbf{Q}_k\}_k$, which will result in the same distribution of graphs as that given by \mathbf{X}_k .

Remark 2.3 (Model extensions). For simplicity of exposition, the WRDPG model is defined assuming a common embedding dimension d across all moment indices k . More generally, the framework can be extended to allow the latent positions $\mathbf{X}[k]$ to lie in moment-specific dimensions d_k , leading to representations of the form $\mathbb{E} \left[W_{ij}^k \right] = \mathbf{x}_i^\top[k]\mathbf{x}_j[k]$, with $\mathbf{x}_i[k] \in \mathbb{R}^{d_k}$. Such dimensions could be selected in a data-driven manner in practical applications, for instance by applying an elbow rule-based procedure to the eigenvalue scree plot of the corresponding empirical k -th order moment matrix.

In addition, the WRDPG model can be generalized to allow for indefinite inner products by replacing $\mathbf{x}_i^\top[k]\mathbf{x}_j[k]$ with $\mathbf{x}_i^\top[k]\mathbf{I}_k^{p,q}\mathbf{x}_j[k]$, where

$$\mathbf{I}_k^{p,q} := \text{diag}(\underbrace{1, \dots, 1}_{p_k}, \underbrace{-1, \dots, -1}_{q_k}), \quad p_k + q_k = d_k.$$

This extension accommodates heterophilous graph structures and aligns the WRDPG framework with the generalized RDPG model of [23]. Since our methodology builds on the spectral techniques of [8], which explicitly account for such indefinite inner product structures, analogous asymptotic guarantees to those developed in Section 3 are expected to hold in this generalized setting.

2.2. Estimation of latent positions

The vectors $\mathbf{x}_i[k]$ can be estimated via an ASE of the matrix $\mathbf{W}^{(k)}$, the entrywise k -th power of an observed symmetric weighted adjacency matrix \mathbf{W} . Indeed, for each index k , we stack the latent positions of all nodes into the matrix $\mathbf{X}[k] = [\mathbf{x}_1[k], \dots, \mathbf{x}_N[k]]^\top \in \mathbb{R}^{N \times d}$. If one had access to the moment matrix $\mathbf{M}_k := \mathbf{X}[k]\mathbf{X}^\top[k]$, it would be straightforward to recover $\mathbf{X}[k]$ (up to an orthogonal transformation) from the eigendecomposition of \mathbf{M}_k . However, since \mathbf{M}_k is typically unobserved, we instead rely on the observed weighted adjacency matrix \mathbf{W} . Because each entry of $\mathbf{W}^{(k)}$ has expectation equal to the corresponding entry of \mathbf{M}_k , we approximate $\mathbf{X}[k]$ by solving [cf. (1)]

$$\hat{\mathbf{X}}[k] \in \operatorname{argmin}_{\mathbf{X} \in \mathbb{R}^{N \times d}} \left\| \mathbf{W}^{(k)} - \mathbf{X}\mathbf{X}^\top \right\|_F^2,$$

for all $k \geq 0$. A solution is readily obtained by setting $\hat{\mathbf{X}}[k] = \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{1/2}$, where $\mathbf{W}^{(k)} = \mathbf{U}_{W_k} \mathbf{D}_{W_k} \mathbf{U}_{W_k}^\top$ is the eigendecomposition of $\mathbf{W}^{(k)}$, $\hat{\mathbf{D}}_k \in \mathbb{R}^{d \times d}$ is the diagonal matrix of the d largest-magnitude eigenvalues, and the columns of $\hat{\mathbf{U}}_k \in \mathbb{R}^{N \times d}$ contain the corresponding eigenvectors. In Appendix B, we show that, under mild assumptions on the weighted inner-product distribution F (Assumptions 1 and 2 in Section 3), the top d eigenvalues of $\mathbf{W}^{(k)}$ are nonnegative, ensuring that $\hat{\mathbf{X}}[k]$ is well-defined. We refer to this estimator as the ASE of $\mathbf{X}[k]$.

2.3. Examples

Next, we derive expressions for the latent position vectors in some simple, yet classical, models. We also simulate such WRDPG instances and show that latent position estimation via the ASE yields satisfactory results. We will derive statistical guarantees for ASE in Section 3, justifying our empirical findings.

2.3.1. Erdős-Rényi graph with Gaussian weights

Consider a graph with $N = 1000$ nodes. We first sample the presence or absence of each edge independently as Bernoulli(p). Then, for each edge we independently sample its weight from a $\mathcal{N}(\mu, \sigma^2)$ distribution. This means that W_{ij} is either 0 with probability $1 - p$ or follows a $\mathcal{N}(\mu, \sigma^2)$ distribution with probability p . One can show that the latent position sequence for every node is $\mathbf{x}[0] = 1$ and $\mathbf{x}[k] = \sqrt{pm_N[k]}$ for $k \geq 1$, where $m_N[k]$ is the k -th moment of the normal distribution with parameters μ and σ^2 .

Figure 3 shows a histogram of the recovered latent positions up to order $k = 6$, with $p = 0.5$, $\mu = 1$ and $\sigma = 0.1$. The red dashed line indicates the true latent positions in each case. Apparently, for each k the estimated positions follow a normal distribution centered around the ground-truth value $\mathbf{x}[k]$. This observation is supported by plotting the pdf of the limiting Gaussian distribution – i.e., the asymptotic distribution of the estimated embeddings as $N \rightarrow \infty$, according to Theorem 3.3, which is proven in Section 3; see also Corollary 3.2.




Fig 3. True (dashed vertical line) and estimated (histograms) latent positions for an Erdős-Rényi model with $\mathcal{N}(1, 0.01)$ weights. Pdf for limiting Gaussians, as given by Corollary 3.2, are plotted with dashed lines in each panel. The empirical distributions closely follow their limiting counterparts derived from the asymptotic normality results established in Section 3.

2.3.2. Two-block SBM with arbitrary weights’ distribution

In this case, the setup is similar as before; only now is the presence or absence of an edge given by a 2-block SBM. That is, each node belongs to exactly one of two communities, and an edge is formed between two nodes with probability p_1 if both belong to community 1, with probability p_2 if both belong to community 2, and with probability q if they belong to different communities. So, conditioned on individual node assignments to communities, the presence or absence of edges is given by the block probability matrix:

$$\mathbf{B} = \begin{pmatrix} p_1 & q \\ q & p_2 \end{pmatrix}. \tag{3}$$

After sampling edges, all weights are independently sampled from the same arbitrary edge-weight distribution.

Let $\mathbf{x}_{C_1}, \mathbf{x}_{C_2} \in (\mathbb{R}^2)^\infty$ denote the latent position sequence for each community. To compute the analytical latent positions for this model, we begin by looking at the mean, i.e., $\mathbf{x}_{C_1}[1], \mathbf{x}_{C_2}[1] \in \mathbb{R}^2$. If two nodes belong to community 1, then an edge is formed between them with probability p_1 , and its weight follows the prescribed distribution. This implies that

$$\mathbf{x}_{C_1}^\top[1]\mathbf{x}_{C_1}[1] = p_1 m_d[1] \Rightarrow \|\mathbf{x}_{C_1}[1]\| = \sqrt{p_1 m_d[1]},$$

where $m_d[1]$ is the first moment (the mean) of the chosen weights’ distribution. Due to the nonidentifiability of latent positions, we can arbitrarily place the latent positions for community 1 along the x -axis, and therefore choose:

$$\mathbf{x}_{C_1}[1] = (\sqrt{p_1 m_d[1]}, 0)^\top. \tag{4}$$

Similarly, we conclude that for community 2 we must have $\|\mathbf{x}_{C_2}[1]\| = \sqrt{p_2 m_d[1]}$. Also, an edge between two nodes in different communities is present with probability q and its weight

follows the prescribed distribution, so

$$\mathbf{x}_{C_1}^\top [1] \mathbf{x}_{C_2} [1] = q m_d [1].$$

From these two constraints, using (4) we conclude that

$$\mathbf{x}_{C_2} [1] = \left(q \sqrt{\frac{m_d [1]}{p_1}}, \sqrt{m_d [1] \left(p_2 - \frac{q^2}{p_1} \right)} \right)^\top.$$

Following the analysis above, one can derive the higher-order terms in the per-community sequences. Indeed, if $m_d [k]$ is the k -th moment for the chosen weights' distribution, imposing that the inner products between latent positions equals the moments of inter- and intra-communities connections leads to:

$$\begin{aligned} \|\mathbf{x}_{C_1} [k]\|^2 &= p_1 m_d [k] \\ \|\mathbf{x}_{C_2} [k]\|^2 &= p_2 m_d [k] \\ \mathbf{x}_{C_1}^\top [k] \mathbf{x}_{C_2} [k] &= q m_d [k]. \end{aligned}$$

Again, arbitrarily placing the latent positions for community 1 along the x -axis we can find the latent positions for each community:

$$\begin{aligned} \mathbf{x}_{C_1} [k] &= \left(\sqrt{p_1 m_d [k]}, 0 \right)^\top \\ \mathbf{x}_{C_2} [k] &= \left(q \sqrt{\frac{m_d [k]}{p_1}}, \sqrt{m_d [k] \left(p_2 - \frac{q^2}{p_1} \right)} \right)^\top. \end{aligned} \tag{5}$$

Note that (5) is valid for $k \geq 1$. For $k = 0$ we can arbitrarily set $\mathbf{x}_{C_i} [0] = (1, 0)^\top$ for $i = 1, 2$ in order to: *a*) maintain the dimensionality of the embeddings; and *b*) force the 0-th moment of each edge to be equal to 1 [cf. Definition 2.1].

We simulated the above setup for a network with $N = 1000$ nodes (700 in community 1 and 300 in community 2), block probability matrix

$$\mathbf{B} = \begin{pmatrix} 0.7 & 0.1 \\ 0.1 & 0.3 \end{pmatrix},$$

and weights sampled from a $\mathcal{N}(\mu, \sigma^2)$ distribution, with $\mu = 1$ and $\sigma = 0.1$. Results for the ASE of $\mathbf{W}^{(k)}$ up to order $k = 6$ are shown in Figure 4. As expected, the estimated embeddings for each community are centered around the analytically derived ones in (5), with an ellipse-like outline that suggests an approximately normal distribution. As with the previous ER example, this is corroborated by plotting the 95% confidence level sets of the limiting Gaussians for each community, as given by the explicit formulae in Corollary 3.2.

2.4. Discriminative power of higher-order spectral embeddings

To demonstrate the ability of our model to distinguish between communities, we simulate a two-block weighted SBM consisting of $N = 2000$ nodes. Edges are established independently with probability $p = 0.5$. Edge weights follow a Gaussian distribution with mean $\mu = 5$ and standard deviation $\sigma = 0.1$, except for those between the second block of nodes, indexed $i =$




Fig 4. Estimated (blue and red circles) and true latent positions (black crosses) for a two-block SBM with $\mathcal{N}(1, 0.01)$ weights. The 95% confidence level sets for the limiting normal distributions, as given by Corollary 3.2, are shown as dashed lines. Once more, the accuracy of the theoretical predictions is apparent.

1001, . . . , 2000, where the weights instead follow a Poisson distribution with rate parameter $\lambda = 5.1$.

For the same reasons as in the previous example, in this setting the matrix $\mathbf{X}[k]$ has at most two distinct columns for each k . Figure 5 displays the estimated embeddings $\hat{\mathbf{x}}_i[k]$ obtained via the ASE of $\mathbf{W}^{(k)}$ for $k = 1, 2, 3$ with embedding dimension $d = 2$, where nodes are color-coded by community membership. The 95% confidence level sets for the limiting normal distributions, as given by Theorem 3.3, are shown in Figure 5 as dashed lines. Notably, for each value of k , the simulated points closely follow the normal distribution predicted by the theorem in the large- N limit, providing empirical support for the asymptotic result.

Observe that for $k = 1$, the node embeddings are nearly indistinguishable across the two communities. This is expected, as the vectors $\hat{\mathbf{x}}_i[1]$ cluster around the points $(\sqrt{\mu p}, 0)^\top \approx (1.58, 0)^\top$ for the gaussian community, and $(\sqrt{\mu p}, \sqrt{p(\lambda - \mu)})^\top \approx (1.58, 0.22)^\top$ for the Poisson community, reflecting the almost identical expected weight of edges in both distributions. When $k = 2$, the embeddings begin to reveal community structure, as shown in the center panel of Figure 5. In this simplified example, closed-form expressions for higher-order embeddings can be readily derived as well. Assuming, without loss of generality, that the embeddings $\mathbf{x}_i[k]$ lie along the x -axis for $i = 1, \dots, 1000$, we obtain:

$$\mathbf{x}_i[k] = \begin{cases} \left(\left(\sqrt{pm_k^N}, 0 \right)^\top, & i \leq 1000, \\ \left(\sqrt{pm_k^N}, \sqrt{p(m_k^P - m_k^N)} \right)^\top, & i > 1000, \end{cases}$$

where m_k^N and m_k^P are the k -th moments of the univariate $\mathcal{N}(\mu, \sigma^2)$ and Poisson distributions, respectively. For $k = 2$, these correspond to the approximate coordinates $(3.54, 0)^\top$ and $(3.54, 1.75)^\top$ for the two groups. However, since the confidence sets of the limiting multivariate Gaussians still overlap, the separation between the communities is not yet clearly pronounced. At $k = 3$, the confidence sets no longer intersect, and the embeddings exhibit a clear separation between the two blocks, as shown in the right panel of Figure 5.




Fig 5. Theoretical latent positions (black crosses) and ASE embeddings of $\mathbf{W}^{(k)}$ for Gaussian ($\mu = 5$ and $\sigma = 0.1$; in red) and Poisson ($\lambda = 5.1$; in blue) distributed weights for $d = 2$ and $k = 1$ (left), $k = 2$ (center), and $k = 3$ (right). Nodes with different weight distributions are clearly revealed for $k = 3$, but they overlap for $k = 1$. The 95% confidence level sets for the limiting normal distributions, as given by Theorem 3.3, are shown as dashed lines.

Recall that the model proposed in [8] is limited to embeddings derived solely from the mean of the weight distribution, specifically, the ASE of the matrix $\mathbf{W}^{(1)} = \mathbf{W}$. In the weighted SBM described above, both blocks are constructed to have almost identical expected edge weights. As a consequence, the embeddings $\hat{\mathbf{x}}_i[1]$ obtained via an ASE of \mathbf{W} are concentrated around the same point in latent space regardless of the underlying block membership, and meaningful separation between the two blocks begins to emerge only when considering higher-order embeddings. Therefore, approaches restricted to first-order information, such as [8], are inherently incapable of discriminating between communities whose structure is encoded in the higher-order moments of edge weights.

2.5. Accuracy of moment recovery with varying number of nodes

Accurate estimation of higher-order moments is known to be highly sensitive to the amount of data available. In particular, the variance of moment estimators tends to grow rapidly with the order of the moment, requiring markedly larger sample sizes to obtain stable estimates [3]. This is because higher-order moments are dominated by extreme values in the data, making them especially prone to noise and outliers. Admittedly, this a challenge facing the proposed model.

To illustrate how this behavior affects the WRDPG model, we simulate a two-block weighted SBM with $N = 2000$ nodes, where 70% of them are assigned to community 1. The interconnection probabilities are given by the matrix

$$\mathbf{B} = \begin{pmatrix} 0.7 & 0.3 \\ 0.3 & 0.5 \end{pmatrix},$$

and edge weights are sampled from a Gaussian distribution with mean $\mu = 1$ and standard deviation $\sigma = 0.5$.

Figure 6 depicts the estimated latent position matrices $\hat{\mathbf{X}}[k]$ corresponding to moments $k = 1$ and $k = 4$. To assess the quality of the embeddings, we compare the entries of the empirical moment matrices $\hat{\mathbf{M}}[k] = \hat{\mathbf{X}}[k]\hat{\mathbf{X}}^T[k]$ with the true moments. The results, shown in the first two columns of Figure 6, indicate that the embeddings closely follow a mixture of




Fig 6. Inference results for a two-class SBM with Gaussian weights and $N = 2000$ nodes (first and second columns) and $N = 200$ nodes (third and fourth columns). The plots on the second and fourth columns show histograms of the estimated $\hat{\mathbf{M}}[k]$ and the vertical lines indicate the true moments. For $N = 2000$ embeddings and moments are accurately estimated up to order $k = 4$, while accuracy degrades in the $N = 200$ setting. Also, for fixed sample size performance degrades as the order increases from $k = 1$ (top row) to $k = 4$ (bottom row).

multivariate Gaussian distributions and that, as expected, the accuracy of $\hat{\mathbf{M}}[k]$ degrades as the moment order increases.

This observation is further reinforced in a second experiment presented in the third and fourth columns of Figure 6, where the number of nodes (i.e., the sample size) is reduced to $N = 200$. Although the embeddings corresponding to $k = 1$ remain fairly accurate, the limited sample size noticeably impairs the estimation for $k = 4$, illustrating the practical limitations imposed by high-order moment estimation. These examples underscore the importance of accounting for finite-sample effects when working with moment-based estimators in network models such as WRDPG.

3. Asymptotic results

Recall the ASE estimator introduced in Section 2.2. Here we establish the asymptotic results that, for a fixed index k , characterize the behavior of the estimated latent positions $\hat{\mathbf{X}}[k]$ when the number of nodes N goes to infinity. Given the inherent rotational ambiguity in the WRDPG model, the latent position sequence is estimable up to an unknown orthogonal transformation. Thus, our main results (consistency and asymptotic normality) are stated in terms of a sequence of orthogonal matrices $\mathbf{Q}_k \in O(d)$.

For these results to hold, we make the following two assumptions. Here, F is an inner-product distribution as in Definition 2.2, $\mathbf{X}_k := \{\mathbf{X}[k]\}_{k \geq 0}$ is a sequence of latent position matrices, and \mathbf{W} is the adjacency matrix of a WRDPG graph, i.e., $(\mathbf{W}, \mathbf{X}_k) \sim \text{WRDPG}(F)$ as in Definition 2.3.

Assumption 1. Let $\{\mathbf{x}[k]\}_{k \geq 0} \sim F$. Then for each k , the second moment matrix $\Delta_k = \mathbb{E}[\mathbf{x}[k]\mathbf{x}^\top[k]]$ has full rank d .

Assumption 2 (Sub-Weibull weights). There exists a constant $\theta > 0$ such that for each pair $1 \leq i < j \leq N$, there is a constant $C_{ij} > 0$ satisfying:

$$\mathbb{P} [|W_{ij}| \geq t | \mathbf{X}_k] \leq 2 \exp \left(- \left(\frac{t}{C_{ij}} \right)^{1/\theta} \right) \text{ for all } t \geq 0.$$

Remark 3.1. Since the class of sub-Weibull random variables (rvs) is closed under multiplication [34, Proposition 2.3], it follows that W_{ij}^k is also sub-Weibull, with parameter $k\theta$. This will allow us to establish our asymptotic results by invoking a concentration bound for sums of sub-Weibull variables, which is proven in Proposition A.2.

3.1. Consistency

We first establish the consistency of the estimated latent positions (up to an unknown orthogonal transformation).

Theorem 3.1. *Let F be an inner product distribution satisfying Assumption 1 and consider $(\mathbf{W}, \mathbf{X}_k) \sim \text{WRDPG}(F)$ satisfying Assumption 2. Then, for each index k there exists an orthogonal matrix $\mathbf{Q}_k \in O(d)$ such that*

$$\|\hat{\mathbf{X}}[k] \mathbf{Q}_k - \mathbf{X}[k]\|_{2 \rightarrow \infty} = O_{\mathbb{P}} \left(N^{-1/2} \log^{k\theta} N \right).$$

Theorem 3.1 generalizes a previous consistency result for the WRDPG model [19, Theorem 1] in two key aspects. First, it extends the result to accommodate unbounded edge weights. Second, it provides a bound on the difference between the estimated and actual latent positions (up to an unknown orthogonal transformation) in terms of the $2 \rightarrow \infty$ norm rather than the Frobenius norm. As noted in [5], the $2 \rightarrow \infty$ norm offers tighter control over the estimation error. This can be seen from the identity

$$\|\mathbf{A}\|_{2 \rightarrow \infty} = \max_{1 \leq i \leq n} \|\mathbf{a}_i^\top\|_2,$$

where \mathbf{a}_i^\top denotes the i -th row of $\mathbf{A} \in \mathbb{R}^{n \times m}$. This means that the $2 \rightarrow \infty$ norm corresponds to the maximum Euclidean norm of the rows of \mathbf{A} , so Theorem 3.1 bounds the maximum error between the estimated and true latent vectors for *any* node in a WRDPG graph. In contrast, controlling the Frobenius norm does not guarantee such a uniform bound.

Our consistency result aligns with previous findings for the vanilla RDPG model [1, 18], in that we provide a bound for the $2 \rightarrow \infty$ norm. It also generalizes the consistency result in [8], as their theorem applies only to $\hat{\mathbf{X}}[1]$, meaning they establish consistency solely for the estimated latent positions associated with the mean. Nevertheless, our proof follows their methodology, adapting the necessary arguments to accommodate our more general setup. We note that their proof scheme closely follows that of the RDPG model, as first established in [18, Theorem 15]. Additionally, it is worth mentioning that a more general result on matrix perturbations can be obtained using similar techniques; see [5].

As in the unweighted case, the proof of Theorem 3.1 relies on expressing the difference between $\hat{\mathbf{X}}[k] \mathbf{Q}_k$ and $\mathbf{X}[k]$ as the sum of a dominant term and a series of remainder terms. The following proposition shows that these remainders are of a lower order than the dominant term. Its proof is deferred to Appendix A.

Proposition 3.2. Let $(\mathbf{W}, \mathbf{X}_k) \sim \text{WRDPG}(F)$ as in Theorem 3.1. For each fixed integer $k \geq 0$, let $\mathbf{W}_k := \mathbf{W}^{(k)}$. Also let $\mathbf{M}_k := \mathbf{X}[k]\mathbf{X}^\top[k]$ and denote its spectral decomposition as $\mathbf{M}_k = \mathbf{U}_k \mathbf{D}_k \mathbf{U}_k^\top$, with $\mathbf{D}_k \in \mathbb{R}^{d \times d}$ and $\mathbf{U}_k \in \mathbb{R}^{N \times d}$. For $\mathbf{S}_k \in O(d)$, define \mathbf{R}_{k_1} , \mathbf{R}_{k_2} , \mathbf{R}_{k_3} and \mathbf{R}_{k_4} as

$$\begin{aligned}\mathbf{R}_{k_1} &= \mathbf{U}_k \left(\mathbf{U}_k^\top \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{1/2} - \mathbf{D}_k^{1/2} \mathbf{S}_k \right), \\ \mathbf{R}_{k_2} &= (\mathbf{I} - \mathbf{U}_k \mathbf{U}_k^\top) (\mathbf{W}_k - \mathbf{M}_k) \left(\hat{\mathbf{U}}_k - \mathbf{U}_k \mathbf{S}_k \right) \hat{\mathbf{D}}_k^{-1/2}, \\ \mathbf{R}_{k_3} &= -\mathbf{U}_k \mathbf{U}_k^\top (\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k \mathbf{S}_k \hat{\mathbf{D}}_k^{-1/2}, \\ \mathbf{R}_{k_4} &= (\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k \left(\mathbf{S}_k \hat{\mathbf{D}}_k^{-1/2} - \mathbf{D}_k^{-1/2} \mathbf{S}_k \right).\end{aligned}$$

Then there exists a choice of $\mathbf{S}_k \in O(d)$ such that the following holds:

$$\begin{aligned}\|\mathbf{R}_{k_1}\|_{2 \rightarrow \infty} &= O_{\mathbb{P}} \left(N^{-1} \log^{k\theta} N \right), \\ \|\mathbf{R}_{k_2}\|_{2 \rightarrow \infty} &= O_{\mathbb{P}} \left(N^{-3/2} \log^{2k\theta} N \right), \\ \|\mathbf{R}_{k_3}\|_{2 \rightarrow \infty} &= O_{\mathbb{P}} \left(N^{-1} \log^{k\theta} N \right), \\ \|\mathbf{R}_{k_4}\|_{2 \rightarrow \infty} &= O_{\mathbb{P}} \left(N^{-3/2} \log^{2k\theta} N \right).\end{aligned}$$

Using this proposition, we can prove Theorem 3.1.

Proof of Theorem 3.1. First, assume $\mathbf{X}[k] = \mathbf{U}_k \mathbf{D}_k^{1/2}$. Fix $\mathbf{S}_k \in O(d)$ such that Proposition 3.2 holds. Since $\hat{\mathbf{X}}[k] = \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{1/2}$, we have that:

$$\begin{aligned}\hat{\mathbf{X}}[k] - \mathbf{X}[k] \mathbf{S}_k &= \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{1/2} - \mathbf{U}_k \mathbf{D}_k^{1/2} \mathbf{S}_k \\ &= \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{1/2} - \mathbf{U}_k \mathbf{U}_k^\top \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{1/2} + \mathbf{U}_k \left(\mathbf{U}_k^\top \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{1/2} - \mathbf{D}_k^{1/2} \mathbf{S}_k \right) \\ &= \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{1/2} - \mathbf{U}_k \mathbf{U}_k^\top \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{1/2} + \mathbf{R}_{k_1}.\end{aligned}$$

Note that $\hat{\mathbf{U}}_k \hat{\mathbf{D}}_k = \mathbf{W}_k \hat{\mathbf{U}}_k$, which implies that $\hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{1/2} = \mathbf{W}_k \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{-1/2}$. So:

$$\begin{aligned}\hat{\mathbf{X}}[k] - \mathbf{X}[k] \mathbf{S}_k &= \mathbf{W}_k \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{-1/2} - \mathbf{U}_k \mathbf{U}_k^\top \mathbf{W}_k \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{-1/2} + \mathbf{R}_{k_1} \\ &= \mathbf{W}_k \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{-1/2} - \mathbf{M}_k \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{-1/2} + \mathbf{M}_k \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{-1/2} - \mathbf{U}_k \mathbf{U}_k^\top \mathbf{W}_k \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{-1/2} + \mathbf{R}_{k_1} \\ &= (\mathbf{W}_k - \mathbf{M}_k) \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{-1/2} + \mathbf{U}_k \mathbf{U}_k^\top \mathbf{M}_k \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{-1/2} - \mathbf{U}_k \mathbf{U}_k^\top \mathbf{W}_k \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{-1/2} + \mathbf{R}_{k_1} \tag{6} \\ &= (\mathbf{I} - \mathbf{U}_k \mathbf{U}_k^\top) (\mathbf{W}_k - \mathbf{M}_k) \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{-1/2} + \mathbf{R}_{k_1}, \tag{7}\end{aligned}$$

where in (6) we have used the fact that $\mathbf{M}_k = \mathbf{U}_k \mathbf{U}_k^\top \mathbf{M}_k$. From (7) we find that:

$$\begin{aligned}\hat{\mathbf{X}}[k] - \mathbf{X}[k] \mathbf{S}_k &= (\mathbf{I} - \mathbf{U}_k \mathbf{U}_k^\top) (\mathbf{W}_k - \mathbf{M}_k) \left(\mathbf{U}_k \mathbf{S}_k - \mathbf{U}_k \mathbf{S}_k + \hat{\mathbf{U}}_k \right) \hat{\mathbf{D}}_k^{-1/2} + \mathbf{R}_{k_1} \\ &= (\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k \mathbf{S}_k \hat{\mathbf{D}}_k^{-1/2} + \mathbf{R}_{k_1} + \mathbf{R}_{k_2} + \mathbf{R}_{k_3} \\ &= (\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k \left(\mathbf{D}_k^{-1/2} \mathbf{S}_k + \mathbf{S}_k \hat{\mathbf{D}}_k^{-1/2} - \mathbf{D}_k^{-1/2} \mathbf{S}_k \right) + \\ &\quad + \mathbf{R}_{k_1} + \mathbf{R}_{k_2} + \mathbf{R}_{k_3}\end{aligned}$$

$$= (\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k \mathbf{D}_k^{-1/2} \mathbf{S}_k + \mathbf{R}_{k_1} + \mathbf{R}_{k_2} + \mathbf{R}_{k_3} + \mathbf{R}_{k_4}. \tag{8}$$

Using Proposition 3.2, we conclude that

$$\|\hat{\mathbf{X}}[k] - \mathbf{X}[k] \mathbf{S}_k\|_{2 \rightarrow \infty} \leq \|(\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k \mathbf{D}_k^{-1/2} \mathbf{S}_k\|_{2 \rightarrow \infty} + \mathcal{O}_{\mathbb{P}}\left(N^{-1} \log^{k\theta}(N)\right).$$

Since $\|\mathbf{A}\mathbf{B}\|_{2 \rightarrow \infty} \leq \|\mathbf{A}\|_{2 \rightarrow \infty} \|\mathbf{B}\|_2$, we have:

$$\begin{aligned} \|\hat{\mathbf{X}}[k] - \mathbf{X}[k] \mathbf{S}_k\|_{2 \rightarrow \infty} &\leq \|(\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k\|_{2 \rightarrow \infty} \|\mathbf{D}_k^{-1/2}\|_2 + \mathcal{O}_{\mathbb{P}}\left(N^{-1} \log^{k\theta}(N)\right) \\ &\leq \|(\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k\|_{2 \rightarrow \infty} \mathcal{O}_{\mathbb{P}}\left(N^{-1/2}\right) + \mathcal{O}_{\mathbb{P}}\left(N^{-1} \log^{k\theta}(N)\right), \end{aligned} \tag{9}$$

where for the second inequality we have used that $\|\mathbf{D}_k^{-1/2}\|_2 = \mathcal{O}_{\mathbb{P}}(N^{-1/2})$ due to the fact that the top d eigenvalues of \mathbf{M}_k are $\Theta_{\mathbb{P}}(N)$; see Appendix B.

To bound the term $\|(\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k\|_{2 \rightarrow \infty}$, we first note that each entry of that matrix is a linear combination of centered, sub-Weibull rvs. Indeed,

$$[(\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k]_{ij} = \sum_{l=1}^N (W_{il} - M_{il}) U_{lj} = \sum_{l \neq i} (W_{il} - M_{il}) U_{lj} - M_{ii} U_{ii},$$

where we denote the (i, j) element of \mathbf{W}_k , \mathbf{M}_k , and \mathbf{U}_k by W_{ij} , M_{ij} , and U_{ij} , respectively. Since the sum in the rightmost term consists of centered, independent sub-Weibull rvs with parameter $k\theta$, we can apply a concentration result for this setting that we prove in Proposition A.2 of Appendix A. Using this result, we find that for all sufficiently large t ,

$$\mathbb{P}\left[\left|[(\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k]_{ij}\right| \geq t\right] \leq 2 \exp\left(-\frac{ct^2}{L_k^2 \sum_{l \neq i} |U_{il}|^2}\right),$$

for some constant $K_2 > 0$ depending on $k\theta$. By choosing $t = K_2 \log^{k\theta} N$ we find that $[(\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k]_{ij}$ is $\mathcal{O}_{\mathbb{P}}(\log^{k\theta} N)$, and by summing over $j = 1, \dots, d$ we have that each row of $(\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k$ is also $\mathcal{O}_{\mathbb{P}}(\log^{k\theta} N)$, so its $2 \rightarrow \infty$ norm is of the same order.

In Appendix B we show that Assumption 1 implies that $\lambda_i(\mathbf{M}_k) = \Theta_{\mathbb{P}}(N)$ for all $i = 1, \dots, d$. Equation (9) then implies that

$$\begin{aligned} \|\hat{\mathbf{X}}[k] - \mathbf{X}[k] \mathbf{S}_k\|_{2 \rightarrow \infty} &= \mathcal{O}_{\mathbb{P}}\left(N^{-1/2} \log^{k\theta} N\right) + \mathcal{O}_{\mathbb{P}}\left(N^{-1} \log^{k\theta}(N)\right) \\ &= \mathcal{O}_{\mathbb{P}}\left(N^{-1/2} \log^{k\theta}(N)\right). \end{aligned}$$

Choosing $\mathbf{Q}_k = \mathbf{S}_k^{\top}$ shows that $\|\hat{\mathbf{X}}[k] \mathbf{Q}_k - \mathbf{X}[k]\|_{2 \rightarrow \infty} = \mathcal{O}_{\mathbb{P}}(N^{-1/2} \log^{k\theta}(N))$, as desired.

If $\mathbf{X}[k] \neq \mathbf{U}_k \mathbf{D}_k^{1/2}$, we have that $\mathbf{X}[k] \mathbf{T}_k = \mathbf{U}_k \mathbf{D}_k^{1/2}$ for some $\mathbf{T}_k \in O(d)$. Note that the above calculations imply that

$$\|\hat{\mathbf{X}}[k] - \mathbf{U}_k \mathbf{D}_k^{1/2} \mathbf{S}_k\|_{2 \rightarrow \infty} = \mathcal{O}_{\mathbb{P}}\left(N^{-1/2} \log^{k\theta} N\right).$$

Therefore, in this case it is enough to choose $\mathbf{Q}_k = \mathbf{S}_k^{\top} \mathbf{T}_k^{\top} \in O(d)$, since

$$\|\hat{\mathbf{X}}[k] \mathbf{Q}_k - \mathbf{X}[k]\|_{2 \rightarrow \infty} = \|(\hat{\mathbf{X}}[k] - \mathbf{X}[k] \mathbf{T}_k \mathbf{S}_k) \mathbf{Q}_k\|_{2 \rightarrow \infty} \leq \|\hat{\mathbf{X}}[k] - \mathbf{U}_k \mathbf{D}_k^{1/2} \mathbf{S}_k\|_{2 \rightarrow \infty}.$$

□

3.2. Asymptotic normality

Next, we show that latent positions behave, asymptotically as $N \rightarrow \infty$, as multivariate normal random variables. We also explicitly calculate the covariance matrix of such a normal in terms of the second-moment matrix Δ_k of the latent positions and the latent positions themselves.

Theorem 3.3. *Let $(\mathbf{W}, \mathbf{X}_k) \sim \text{WRDPG}(F)$ be as in Theorem 3.1. For each index k , define the variance function $v_k : \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}$ as*

$$v_k(\mathbf{x}, \mathbf{y}) := \text{var} \left[W_{ij}^k | \mathbf{x}_i[k] = \mathbf{x}, \mathbf{x}_j[k] = \mathbf{y} \right],$$

where W_{ij} is the (i, j) entry of \mathbf{W} . Let $\Sigma_k : \mathbb{R}^d \mapsto \mathbb{R}^{d \times d}$ be the covariance function

$$\Sigma_k(\mathbf{x}) = \Delta_k^{-1} \mathbb{E} \left[v_k(\mathbf{x}, \mathbf{y}_k) \mathbf{y}_k \mathbf{y}_k^\top \right] \Delta_k^{-1},$$

where $\{\mathbf{y}_k\}_{k \geq 0} \sim F$ and Δ_k is the second-moment matrix $\Delta_k = \mathbb{E} \left[\mathbf{y}_k \mathbf{y}_k^\top \right]$.

Then for each k there exists a sequence of orthogonal matrices $\{\mathbf{Q}_{kN}\}_{N \geq 0}$ such that for all $\mathbf{z} \in \mathbb{R}^d$ and for any fixed row index i ,

$$\lim_{N \rightarrow \infty} \mathbb{P} \left[N^{1/2} (\hat{\mathbf{X}}[k] \mathbf{Q}_{kN} - \mathbf{X}[k])_i^\top \leq \mathbf{z} \mid \mathbf{x}_i[k] = \mathbf{x} \right] = \Phi(\mathbf{z}; \Sigma_k(\mathbf{x})),$$

where $\Phi(\cdot; \Sigma)$ stands for the cumulative distribution function of a $\mathcal{N}(\mathbf{0}, \Sigma)$ random vector.

Proof. From the proof of Theorem 3.1, we have that for each fixed k there exists $\mathbf{S}_k, \mathbf{T}_k \in O(d)$ such that for $\mathbf{Q}_k = \mathbf{S}_k^\top \mathbf{T}_k^\top$ it holds

$$\hat{\mathbf{X}}[k] \mathbf{Q}_k - \mathbf{X}[k] = \left(\hat{\mathbf{X}}[k] - \mathbf{U}_k \mathbf{D}_k^{1/2} \mathbf{S}_k \right) \mathbf{Q}_k.$$

Also from that proof [see (8)], we have that

$$\hat{\mathbf{X}}[k] - \mathbf{U}_k \mathbf{D}_k^{1/2} \mathbf{S}_k = (\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k \mathbf{D}_k^{-1/2} \mathbf{S}_k + \mathbf{R}_k,$$

where we have defined $\mathbf{R}_k = \mathbf{R}_{k1} + \mathbf{R}_{k2} + \mathbf{R}_{k3} + \mathbf{R}_{k4}$. Therefore:

$$N^{1/2} (\hat{\mathbf{X}}[k] \mathbf{Q}_k - \mathbf{X}[k]) = N^{1/2} (\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k \mathbf{D}_k^{-1/2} \mathbf{T}_k^\top + N^{1/2} \mathbf{R}_k \mathbf{Q}_k. \tag{10}$$

Using Proposition 3.2, we have that $\|N^{1/2} \mathbf{R}_k \mathbf{Q}_k\|_{2 \rightarrow \infty} \rightarrow 0$ as $N \rightarrow \infty$. Therefore, we can focus on the first summand in (10).

Recall from the proof of Theorem 3.1 that \mathbf{T}_k is such that $\mathbf{X}[k] \mathbf{T}_k = \mathbf{U}_k \mathbf{D}_k^{1/2}$, which implies that $\mathbf{U}_k \mathbf{D}_k^{-1/2} = \mathbf{X}[k] \mathbf{T}_k \mathbf{D}_k^{-1}$. Thus:

$$N^{1/2} (\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k \mathbf{D}_k^{-1/2} \mathbf{T}_k^\top = N^{1/2} (\mathbf{W}_k - \mathbf{M}_k) \mathbf{X}[k] \mathbf{T}_k \mathbf{D}_k^{-1} \mathbf{T}_k^\top.$$

Therefore, the i -th row of the first summand in (10) equals:

$$\begin{aligned} N^{1/2} \left[(\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k \mathbf{D}_k^{-1/2} \mathbf{T}_k^\top \right]_i^\top &= N^{1/2} \mathbf{T}_k \mathbf{D}_k^{-1} \mathbf{T}_k^\top \{ (\mathbf{W}_k - \mathbf{M}_k) \mathbf{X}[k] \}_i^\top \\ &= N \mathbf{T}_k \mathbf{D}_k^{-1} \mathbf{T}_k^\top \left[N^{-1/2} \sum_{j=1}^N (W_{ij}^k - M_{ij}) \mathbf{x}_j[k] \right] \end{aligned}$$

$$\begin{aligned}
 &= N\mathbf{T}_k\mathbf{D}_k^{-1}\mathbf{T}_k^\top \left[N^{-1/2} \sum_{j \neq i} \left(W_{ij}^k - M_{ij} \right) \mathbf{x}_j[k] \right] \\
 &\quad - N\mathbf{T}_k\mathbf{D}_k^{-1}\mathbf{T}_k \left(N^{-1/2} M_{ii} \mathbf{x}_i[k] \right),
 \end{aligned} \tag{11}$$

where as before we denote the (i, j) element of \mathbf{W} and \mathbf{M}_k by W_{ij} and M_{ij} , respectively. Conditioning on $\mathbf{x}_i[k] = \mathbf{x}$, we have that for the second term in (11)

$$\begin{aligned}
 \left\| N\mathbf{T}_k\mathbf{D}_k^{-1}\mathbf{T}_k \left(N^{-1/2} M_{ii} \mathbf{x}_i[k] \right) \right\|_2 &\leq N^{1/2} |M_{ii}| \left\| \mathbf{T}_k\mathbf{D}_k^{-1}\mathbf{T}_k \right\|_{2 \rightarrow \infty} \|\mathbf{x}\|_2 \\
 &\leq N^{1/2} L_k \left\| \mathbf{D}_k^{-1} \right\|_2 \|\mathbf{x}\|_2,
 \end{aligned}$$

where in the second inequality we have used that Assumption 2 implies that $|M_{ij}| < L_k$ for all (i, j) , and that $\|\mathbf{A}\|_{2 \rightarrow \infty} \leq \|\mathbf{A}\|_2$. Because $\left\| \mathbf{D}_k^{-1} \right\|_2 = O_{\mathbb{P}}(N^{-1})$, we conclude that the second term in (11) is $O_{\mathbb{P}}(N^{-1/2})$.

As for the first term in (11), conditional on $\mathbf{x}_i[k] = \mathbf{x}$ we have that $M_{ij} = \mathbf{x}^\top \mathbf{x}_j[k]$, so the terms in the sum

$$N^{-1/2} \sum_{j \neq i} \left(W_{ij}^k - M_{ij} \right) \mathbf{x}_j[k]$$

are centered, independent random variables, whose covariance matrix is

$$\tilde{\Sigma}_k(\mathbf{x}) = \mathbb{E} \left[v_k(\mathbf{x}, \mathbf{y}_k) \mathbf{y}_k \mathbf{y}_k^\top \right],$$

where $\{\mathbf{y}_k\}_{k \geq 0} \sim F$. Therefore, the multivariate central limit theorem implies that

$$N^{-1/2} \sum_{j \neq i} \left(W_{ij}^k - M_{ij} \right) \mathbf{x}_j[k] \xrightarrow{\mathcal{L}} \mathcal{N}(\mathbf{0}, \tilde{\Sigma}_k(\mathbf{x})).$$

To conclude the proof, recall that $\mathbf{T}_k \in O(d)$ is such that $\mathbf{X}[k]\mathbf{T}_k = \mathbf{U}_k\mathbf{D}_k^{1/2}$, so $\mathbf{X}[k] = \mathbf{U}_k\mathbf{D}_k^{1/2}\mathbf{T}_k^\top$ and therefore $\mathbf{X}^\top[k]\mathbf{X}[k] = \mathbf{T}_k\mathbf{D}_k\mathbf{T}_k^\top$. This implies that

$$\left(\mathbf{X}^\top[k]\mathbf{X}[k] \right)^{-1} = \mathbf{T}_k\mathbf{D}_k^{-1}\mathbf{T}_k^\top,$$

and therefore by the law of large numbers we have that almost surely it holds

$$N\mathbf{T}_k\mathbf{D}_k^{-1}\mathbf{T}_k^\top \rightarrow \Delta_k^{-1}.$$

This implies that the first term in (11) converges in distribution to a multivariate normal $\mathcal{N}(\mathbf{0}, \Sigma(\mathbf{x}))$, which completes the proof. \square

Remark 3.2. The covariance function in Theorem 3.3 depends on $v_k(\mathbf{x}, \mathbf{y})$, which is the conditional variance of W_{ij}^k given $\mathbf{x}_i[k] = \mathbf{x}$ and $\mathbf{x}_j[k] = \mathbf{y}$. Therefore, we can rewrite v_k as

$$\begin{aligned}
 v_k(\mathbf{x}, \mathbf{y}) &= \mathbb{E} \left[\left(W_{ij}^k \right)^2 \mid \mathbf{x}_i[k] = \mathbf{x}, \mathbf{x}_j[k] = \mathbf{y} \right] - \mathbb{E} \left[W_{ij}^k \mid \mathbf{x}_i[k] = \mathbf{x}, \mathbf{x}_j[k] = \mathbf{y} \right]^2 \\
 &= \mathbb{E} \left[W_{ij}^{2k} \mid \mathbf{x}_i[k] = \mathbf{x}, \mathbf{x}_j[k] = \mathbf{y} \right] - \left(\mathbf{x}^\top \mathbf{y} \right)^2,
 \end{aligned}$$

since the WRDPG model by definition imposes $\mathbb{E} \left[W_{ij}^k \mid \mathbf{x}_i[k] = \mathbf{x}, \mathbf{x}_j[k] = \mathbf{y} \right] = \mathbf{x}^\top \mathbf{y}$. The expectation of W_{ij}^{2k} need not have a closed-form expression given $\mathbf{x}_i[k] = \mathbf{x}$ and $\mathbf{x}_j[k] = \mathbf{y}$.

However, if we further condition on the events $\mathbf{x}_i[2k] = \mathbf{x}_2$ and $\mathbf{y}_i[2k] = \mathbf{y}_2$, we have that $\mathbb{E} \left[W_{ij}^{2k} \right] = \mathbf{x}_2^\top \mathbf{y}_2$. Therefore, by conditioning on $\mathbf{x}_i[k] = \mathbf{x}_1$ and $\mathbf{x}_i[2k] = \mathbf{x}_2$ we have the following Corollary of Theorem 3.3.

Corollary 3.1. *Let $(\mathbf{W}, \mathbf{X}_k) \sim \text{WRDPG}(F)$ be as in Theorem 3.1. Then for each k there exists a sequence of orthogonal matrices $\{\mathbf{Q}_{kN}\}_{N \geq 0}$ such that for all $\mathbf{z} \in \mathbb{R}^d$ and for any fixed row index i ,*

$$\lim_{N \rightarrow \infty} \mathbb{P} \left[N^{1/2} (\hat{\mathbf{X}}[k] \mathbf{Q}_{kN} - \mathbf{X}[k])_i^\top \leq \mathbf{z} \mid \mathbf{x}_i[k] = \mathbf{x}_1, \mathbf{x}_i[2k] = \mathbf{x}_2 \right] = \Phi(\mathbf{z}, \Sigma_k(\mathbf{x}_1, \mathbf{x}_2)),$$

where $\Sigma_k : \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}^{d \times d}$ is the covariance function

$$\Sigma_k(\mathbf{x}_1, \mathbf{x}_2) = \Delta_k^{-1} \mathbb{E} \left[\left(\mathbf{x}_2^\top \mathbf{y}_{2k} - (\mathbf{x}_1^\top \mathbf{y}_k)^2 \right) \mathbf{y}_k \mathbf{y}_k^\top \right] \Delta_k^{-1},$$

and $\{\mathbf{y}_k\}_{k \geq 0} \sim F$.

As discussed in Section 2.3.2, when the random graph follows a weighted SBM with C communities, for each moment index k there are at most C distinct latent positions that a node can assume. We denote these positions by $\mathbf{y}_m[k]$, for $m = 1, \dots, C$. Assume that each node belongs to community m with probability π_m , and define the vector $\boldsymbol{\pi} = [\pi_1, \pi_2, \dots, \pi_C]^\top$. Let $\mathbf{B} \in \mathbb{R}^{C \times C}$ be the matrix of edge-formation probabilities between communities – i.e., entry b_{lm} specifies the probability that an edge exists between a node in community l and a node in community m . Let d_{lm} denote the probability distribution governing the edge weights between nodes in communities l and m , with $m_{lm}[k]$ denoting its k -th moment.

Under these conventions, each entry of the weighted adjacency matrix \mathbf{W} is either zero with probability $1 - \boldsymbol{\pi}^\top \mathbf{B} \boldsymbol{\pi}$, or drawn from d_{lm} with probability $\pi_l \pi_m b_{lm}$. In this setup, the following Corollary of Theorem 3.3 holds.

Corollary 3.2. *With the above notation, for each moment index k there exists a sequence of orthogonal matrices $\{\mathbf{Q}_{kN}\}_{N \geq 0}$ such that for all $\mathbf{z} \in \mathbb{R}^d$ and for any fixed row index i ,*

$$\lim_{N \rightarrow \infty} \mathbb{P} \left[N^{1/2} (\hat{\mathbf{X}}[k] \mathbf{Q}_{kN} - \mathbf{X}[k])_i^\top \leq \mathbf{z} \mid \text{node } i \text{ belongs to community } l \right] = \Phi(\mathbf{z}, \Sigma_{kl}),$$

where the covariance matrix Σ_{kl} is given by $\Sigma_{kl} = \Delta_k^{-1} \tilde{\Sigma}_{kl} \Delta_k^{-1}$, with

$$\tilde{\Sigma}_{kl} = \sum_{m=1}^C \pi_m \left(b_{lm} m_{lm}[2k] - b_{lm}^2 m_{lm}^2[k] \right) \mathbf{y}_m[k] \mathbf{y}_m^\top[k]$$

and

$$\Delta_k = \sum_{m=1}^C \pi_m \mathbf{y}_m[k] \mathbf{y}_m^\top[k].$$

Remark 3.3 (Sparse WRDPG). The WRDPG model may be extended to sparse weighted networks by introducing a sparsity factor $\rho_N^{(k)} \rightarrow 0$ for each moment index k , so that

$$\mathbb{E} \left[W_{ij}^k \right] = \rho_N^{(k)} \mathbf{x}_i^\top[k] \mathbf{x}_j[k].$$

In this regime, the operator norm of the population moment matrix scales as $\|\mathbf{M}_k\|_2 = \Theta_{\mathbb{P}} \left(N \rho_N^{(k)} \right)$.

Under the sub-Weibull tail Assumption 2, the empirical k -th moment matrix satisfies

$$\left\| \mathbf{W}^{(k)} - \mathbf{M}_k \right\|_2 = O_{\mathbb{P}} \left(\rho_N^{(k)} \log^{k\theta} N \right).$$

Consequently, spectral consistency of the ASE for the k -th moment holds provided that the signal dominates the stochastic fluctuation, namely,

$$N \rho_N^{(k)} \gg \log^{k\theta} N.$$

For asymptotic normality of the estimated latent positions, a stronger sparsity condition is required. In particular, normal fluctuations are expected to hold whenever

$$N \rho_N^{(k)} \gg \log^{2k\theta} N,$$

ensuring that the central limit scaling dominates the residual spectral error. All in all, these sparsity requirements are analogous to those appearing in the sparse unweighted RDPG literature [29] and indicate that the WRDPG framework also admits sparse regimes with expected degrees growing at polylogarithmic rates.

4. Graph generation

So far, we have formally defined the WRDPG model and demonstrated its versatility and discriminative power through examples. In addition, we have provided statistical guarantees for the proposed ASE-based estimation method. In this section, we investigate how to generate graphs from the WRDPG model with latent positions

$$\mathbf{X}[0], \mathbf{X}[1], \dots, \mathbf{X}[K] \in \mathbb{R}^{N \times d},$$

where $\mathbf{X}[k] = [\mathbf{x}_1[k], \mathbf{x}_2[k], \dots, \mathbf{x}_N[k]]^T$. In other words, our aim is to sample the adjacency matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$, such that, for each pair $1 \leq i < j \leq N$, W_{ij} follows a distribution whose first $K + 1$ moments are given by $\mathbf{x}_i^T[k] \mathbf{x}_j[k]$, for all $k = 0, 1, \dots, K$. The latent positions could be defined so as to match a prescribed weight distribution, or, be estimated from real networks via the ASE.

In the sequel, we first consider a discrete weight distribution with finite support. We show that the weights' probability mass function (pmf) can be obtained in closed form, by solving a linear system of equations with Vandermonde structure. As we will see, our method is capable of reproducing the original characteristics of the network. Next, we address the problem of generating samples from WRDPG graphs with continuous weight distributions. To this end, we rely on the maximum entropy principle to recover the probability density function (pdf) from its moments. We develop a new primal-dual method, providing better and more robust solutions than prior art [24]; see also the results in Appendix D. Finally, we consider the case of a mixture distribution, simultaneously reproducing the connectivity structure of a real network and its weight distribution. Throughout, we provide supporting examples using synthetic and real data.

4.1. Discrete weights distribution

We begin by deriving a solution for the case where W_{ij} has a discrete distribution. Suppose that the latent positions for each node are given, and W_{ij} takes on $R + 1$ (known) distinct values v_0, v_1, \dots, v_R with (unknown) probabilities p_0, p_1, \dots, p_R .² To generate a WRDPG-adhering graph, we need to estimate these probabilities from the latent position sequence. In this case, the moments for such a distribution can be computed as

$$\mathbb{E} \left[W_{ij}^k \right] = \sum_{r=0}^R v_r^k p_r = v_0^k p_0 + v_1^k p_1 + \dots + v_R^k p_R = \mathbf{x}_i^\top [k] \mathbf{x}_j [k] := m_{ij} [k],$$

where the usual convention $0^0 = 1$ is used in case $v_r = 0$, for some r . We then obtain the following system of $K + 1$ linear equations on the pmf $\mathbf{p} = [p_0, p_1, \dots, p_R]^\top \in [0, 1]^{R+1}$:

$$\begin{cases} p_0 + p_1 + \dots + p_R & = m_{ij} [0] \\ v_0 p_0 + v_1 p_1 + \dots + v_R p_R & = m_{ij} [1] \\ v_0^2 p_0 + v_1^2 p_1 + \dots + v_R^2 p_R & = m_{ij} [2] \\ \vdots & \vdots \\ v_0^K p_0 + v_1^K p_1 + \dots + v_R^K p_R & = m_{ij} [K] \end{cases} \Leftrightarrow \mathbf{V} \mathbf{p} = \mathbf{m}, \tag{12}$$

where $\mathbf{m} = [m_{ij} [0], m_{ij} [1], \dots, m_{ij} [R]]^\top \in \mathbb{R}^{R+1}$ and $\mathbf{V} \in \mathbb{R}^{(K+1) \times (R+1)}$ is the Vandermonde matrix

$$\mathbf{V} = \begin{pmatrix} 1 & 1 & \dots & 1 \\ v_0 & v_1 & \dots & v_R \\ v_0^2 & v_1^2 & \dots & v_R^2 \\ \vdots & \vdots & \ddots & \vdots \\ v_0^K & v_1^K & \dots & v_R^K \end{pmatrix}.$$

Note that if $K = R$ then the system (12) has a unique solution. So, in order to estimate the pmf associated with edge (i, j) we need to prescribe as many moments as there are possible values for W_{ij} . Given the first R moments of the distribution, we readily obtain the probabilities as $\mathbf{p} = \mathbf{V}^{-1} \mathbf{m}$.

While the linear system (12) needs to be solved for every pair (i, j) , if the distributions associated with different edges share a common support v_0, v_1, \dots, v_R , then \mathbf{V} (and thus \mathbf{V}^{-1}) are the same across those edges, making computation of \mathbf{p} for the entire graph less demanding.

Example. To illustrate how the graph generative process works, we simulated a two-class weighted SBM network with $N = 500$ nodes (350 in community 1 and 150 in community 2), and block probability matrix

$$\mathbf{B} = \begin{pmatrix} 0.7 & 0.2 \\ 0.2 & 0.5 \end{pmatrix}, \tag{13}$$

where edge weights are sampled from a discrete distribution supported on the values $1, 2, \dots, 10$, with $p_i = \frac{1}{18}$ for $i \neq 5$ and $p_5 = \frac{1}{2}$. Using (5), we compute the analytical latent positions for each node, where $m_d [k]$ denotes the k -th moment of the aforementioned discrete distribution.

²Note that both v_r 's and p_r 's are dependent on i, j , but that dependence has been omitted to improve readability.




Fig 7. Comparisons between two-block SBMs generated from the base model (blue line) and from the discrete density estimated from latent positions (histograms and boxplots).

For each edge, we compute the finite moment sequence $m_{ij}[k] = \mathbf{x}_i^T[k]\mathbf{x}_j[k]$, for $k = 0, \dots, 10$. In this setup, each edge weight can take on values in $\{0, 1, 2, \dots, 10\}$, with probabilities $p_0, (1 - p_0)p_1, (1 - p_0)p_2, \dots, (1 - p_0)p_{10}$, where p_0 is the probability that nodes i and j are not connected. Specifically, $p_0 = 1 - 0.7 = 0.3$ if i and j both belong to community 1, $p_0 = 1 - 0.5 = 0.5$ if both are in community 2, and $p_0 = 1 - 0.2 = 0.8$ otherwise. Thus, edge weights are discrete random variables with $R + 1 = 11$ possible values, which requires 11 moments to uniquely solve (12).

Figure 7 shows the distributions of various metrics (degree, betweenness centrality, geodesic distance) for 100 simulated networks generated using the above procedure, along with the same metrics computed on a single reference network drawn from the base model—namely, a two-block SBM with \mathbf{B} as in (13), and edge weights sampled from the above discrete distribution. Apparently, the metrics for the network generated from the base model align closely with the distribution of corresponding metrics computed from the sampled graphs.

Remark 4.1. Vandermonde matrices—particularly those constructed from monomial bases—are prone to ill-conditioning, especially as the polynomial degree increases, or when the support of the distribution is large. This poses significant numerical challenges when recovering discrete distributions from moment sequences using direct matrix inversion methods, such as solving the system (12). To mitigate this issue, we propose an alternative formulation that leverages Chebyshev polynomials of the first kind. These polynomials are orthogonal with respect to the weight function $\frac{1}{\sqrt{1-x^2}}$ on the interval $[-1, 1]$, which endows them with superior numerical stability and conditioning within that domain [4].

Our approach is to reformulate the system (12) on the basis of Chebyshev polynomials of the first kind. Each expression on the left-hand side of the system can be interpreted as a dot product between the unknown probability vector \mathbf{p} and a polynomial expressed on the monomial basis $\{x^i\}$, evaluated at the support points $\{v_r\}$. Since any polynomial of degree K can be expressed as a linear combination of the first $K + 1$ Chebyshev polynomials, we propose to convert each monomial x^k in the system into its Chebyshev expansion. This leads to an equivalent system, with a better conditioned Vandermonde-like matrix constructed using evaluations of Chebyshev polynomials at $\{v_r\}$.

This Chebyshev-based reformulation enhances numerical stability during reconstruction; an attractive feature when dealing with empirically estimated moments, which are inevitably affected by noise due to finite sample size. Additional methodological and implementation details are provided in Appendix C.

4.2. Continuous weights distribution

We now move on to the case where the edge weights W_{ij} are continuous random variables. That is, our goal is to determine a pdf g_{ij} such that

$$\int_{\Omega} x^k g_{ij}(x) dx = \mathbf{x}_i^{\top}[k] \mathbf{x}_j[k], \quad \text{for all } k = 0, 1, \dots, K, \quad (14)$$

where $\Omega \subset \mathbb{R}$ denotes the (assumed known) support of the random variable W_{ij} .

To identify a unique pdf consistent with this partial information, we turn to the foundational work of Shore and Johnson [26], who provided an axiomatic justification for using the principle of maximum entropy in such settings. Their key result shows that when one's knowledge about a probability distribution is limited to certain expectation values or constraints, the only consistent and unbiased method for selecting a distribution is to choose the one that maximizes entropy subject to those constraints. This follows from a set of axioms that any rational updating procedure should satisfy, such as consistency, uniqueness, invariance under coordinate transformations, and system independence.

In this context, entropy refers to the differential entropy of pdf g_{ij} , namely

$$S(g_{ij}) = - \int_{\Omega} g_{ij}(x) \log g_{ij}(x) dx.$$

Maximizing $S(g_{ij})$ under the moment constraints (14) ensures that no additional structure or assumptions are imposed beyond what is strictly warranted by those constraints. Thus, following Shore and Johnson's derivation, the maximum entropy principle is not simply a heuristic, but the unique method of inference consistent with the logical requirements of rational belief updating.

Therefore, we formulate the following primal optimization problem:

$$\max_{g \in \mathcal{F}} S(g) \quad \text{s. to} \quad \int_{\Omega} x^k g(x) dx - m_k = 0, \quad \text{for all } k = 0, 1, \dots, K, \quad (15)$$

where $m_k = \mathbf{x}_i^{\top}[k] \mathbf{x}_j[k]$ and \mathcal{F} denotes the space of all probability distributions supported on Ω (as before, for clarity and to avoid notational clutter, we omit the dependence of g and m_k on indices i and j). In Appendix D, we show that the solution of (15) can be obtained by solving the corresponding dual problem

$$\min_{\lambda_0, \dots, \lambda_K \in \mathbb{R}} \left[\sum_{k=0}^K \lambda_k m_k + \int_{\Omega} \exp \left(- \sum_{k=0}^K \lambda_k x^k \right) dx - m_0 \right].$$

Since this dual formulation is convex and unconstrained, it can be tackled using standard convex optimization algorithms. In practice, we solve it using the BFGS algorithm [21] available through SciPy's `minimize` function.




Fig 8. Comparisons between two-block SBMs generated from the base model (blue line) and from the pdf estimated with our method for solving the maximum entropy problem from latent positions (histogram and boxplot). Since in this setup graphs are fully connected, we do not report results for betweenness centrality.

Example. To further illustrate the flexibility of the proposed graph generation method, we consider a two-block weighted SBM network with $N = 500$ nodes, comprising 350 nodes in community 1 and 150 nodes in community 2. The block probability matrix is given by

$$\mathbf{B} = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix},$$

which implies that the graph is fully connected. This setup is chosen because assigning probabilities strictly less than one would introduce a point mass at zero in the edge-weight distribution, thereby resulting in a mixture distribution – a setup which we address in the next section. Edge weights are sampled from distinct continuous distributions depending on the community membership of the connected nodes: for intra-community edges within community 1, weights follow a normal distribution $\mathcal{N}(6, 1)$; for intra-community edges within community 2, weights follow an exponential distribution with rate parameter $\lambda = \frac{1}{3}$; and for inter-community edges, weights follow a normal distribution $\mathcal{N}(1, 0.01)$. Using (5), we compute the analytical latent positions for each node, where $m_d[k]$ now denotes the k -th moment of the corresponding distribution associated with each block. Edge weights thus arise from a mixture of three distributions, leading to a richer moment structure.

As before, for each edge we compute the finite moment sequence $m_{ij}[k] = \mathbf{x}_i^\top[k] \mathbf{x}_j[k]$, for $k = 0, \dots, 5$. We approximate the underlying edge-weight distributions using those first six moments, estimating them via the maximum entropy principle by solving the dual problem derived in Appendix D. Figure 8 displays network-wide distributions of summary statistics (degree and geodesic distance) for 100 networks generated using this procedure, alongside the corresponding metrics of a single network sampled directly from the base model, i.e., a two-class, fully connected weighted SBM, and edge weights sampled according to the aforementioned mixture of continuous distributions. The metric distributions computed from the base model are in close agreement with those obtained from the generated graphs, thus validating the accuracy of the moment-based generation process in the presence of heterogeneous weight distributions.

4.3. Mixed weights distribution

We now focus on the case where the weights' distribution is a mixture of discrete and continuous components, with its pdf having the form

$$g_{ij}(x) = p_0^{ij} \delta(x) + (1 - p_0^{ij}) h_{ij}(x) \tag{16}$$

where $\delta(x)$ is Dirac's delta, $h_{ij}(x)$ is an unknown pdf and $p_0^{ij} \in [0, 1]$ is an unknown parameter that controls the edge-formation probability between nodes i and j . If p_0^{ij} were known, we could estimate h_{ij} from the prescribed moments $m_{ij}[k] = \mathbf{x}_i^\top[k] \mathbf{x}_j[k]$, since

$$m_{ij}[0] = p_0^{ij} + (1 - p_0^{ij}) \int_{\Omega} h_{ij}(x) dx,$$

$$m_{ij}[k] = (1 - p_0^{ij}) \int_{\Omega} x^k h_{ij}(x) dx, \quad \forall k \geq 1,$$

where Ω is the support of $h_{ij}(x)$. This implies that we are looking for a pdf h_{ij} with moments

$$\int_{\Omega} h_{ij}(x) dx = \frac{m_{ij}[0] - p_0^{ij}}{1 - p_0^{ij}},$$

$$\int_{\Omega} x^k h_{ij}(x) dx = \frac{m_{ij}[k]}{1 - p_0^{ij}}, \quad \forall k \geq 1,$$

so we can use the procedure in Section 4.2 to find h_{ij} by maximizing its entropy.

In order to estimate p_0^{ij} , we propose to simultaneously estimate it for every edge via the ASE of the binary matrix $\mathbf{A} := \mathbb{I}\{\mathbf{W} > \mathbf{0}\}$, where $\mathbb{I}\{\cdot\}$ denotes the entrywise matrix indicator function, i.e., $A_{ij} := \mathbb{I}\{W_{ij} > 0\}$. This estimator can be justified by noting that $\mathbb{E}[A_{ij}] = 1 - p_0^{ij}$. Indeed,

$$\mathbb{E}[A_{ij}] = \int_{\Omega} \mathbb{I}\{W_{ij} > 0\} g(x) dx = \int_{\Omega} (1 - p_0^{ij}) h_{ij}(x) dx = 1 - p_0^{ij},$$

where we have assumed that $\Omega \subset \mathbb{R}^+$. Let $\hat{\mathbf{X}}_A$ denote the ASE of \mathbf{A} . Then, in light of our consistency result in Section 3.1, matrix $\hat{\mathbf{P}} := \hat{\mathbf{X}}_A \hat{\mathbf{X}}_A^\top$ converges to $\mathbb{E}[\mathbf{A}]$ in the $\|\cdot\|_\infty$ sense, as $N \rightarrow \infty$.

Example (Reproducing real networks.). In this example the latent positions are not given and instead we estimate them from a real graph that is observed. As a result, now the input is a finite sequence of *approximate* moments. We study a dataset that records football matches between national teams [17]. For a given time period, we construct a graph in which nodes represent countries, and an edge exists between two countries if they have played at least one match against each other during that period. The edge weight corresponds to the number of matches played between them. Next, we compute the latent positions for the edge weight distribution by performing the ASE of $\mathbf{W}^{(k)}$, where $\mathbf{W}^{(k)}$ denotes the k -th entry-wise power of the weight matrix \mathbf{W} , with k ranging from 0 to a prespecified value K . Using these latent positions, we compute a finite moment sequence for each edge via the corresponding inner product and estimate its density using the procedure described earlier in this section.

By assuming a mixture-like density as in (16), we explicitly model the sparsity pattern of the football network. This approach enables us to generate synthetic networks that both conform




Fig 9. Graph generation metrics the football dataset [17]. Metrics for the true graph are shown with a blue solid line, while a histogram or a boxplot shows the results for the corresponding metric for 100 synthetic graphs generated using the estimated mixed densities.

to the observed sparsity pattern and exhibit edge weights that closely mimic those of the actual network. Figure 9 shows several metrics for the actual network of football matches during the period 2010–2016, alongside the corresponding metrics for 100 synthetically-generated networks using the aforementioned pdf estimation plus weighted graph sampling procedure. For this example, we used $K = 2$, meaning that the zeroth, first, and second moments were employed for density estimation.

To further demonstrate the quality of the generated graphs, we conducted a community structure analysis. Since matches between countries belonging to the same football confederation are more frequent, one would expect to observe a clear community structure in the real network, with clusters roughly corresponding to confederations. This can be verified by looking at Figure 10, where we show the results of applying the classic Louvain clustering algorithm [2] to the real graph. The algorithm detects as many clusters as there are confederations (six), with community membership reflecting, most of the times, the actual confederations teams belong to: for example, Australia is geographically in Oceania, but is a member of the Asian Football Confederation (AFC). The only major discrepancy is in North and Central America: while the countries in those subcontinents belong to the Confederation of North, Central America and Caribbean Association Football (CONCACAF), they are clustered together with southern countries, which belong to the South American Football Confederation (CONMEBOL). This might be because the former are often invited to participate in CONMEBOL’s flagship tournament, the Copa América, inflating the amount of games they play against CONMEBOL members. The fact that most Caribbean members of CONCACAF belong to a separate cluster further strengthens this hypothesis, since they rarely play in the Copa América.

We then ran the Louvain algorithm on the simulated networks and compared the results with those of the real network; see the results in Figure 11. On the left panel we show a histogram of the amount of communities in the synthetic graph’s largest connected component. As we can see, all graphs have either 5 or 6 communities in that component, which is consistent with the structure we discussed perviously. And while most graphs tend to have one community

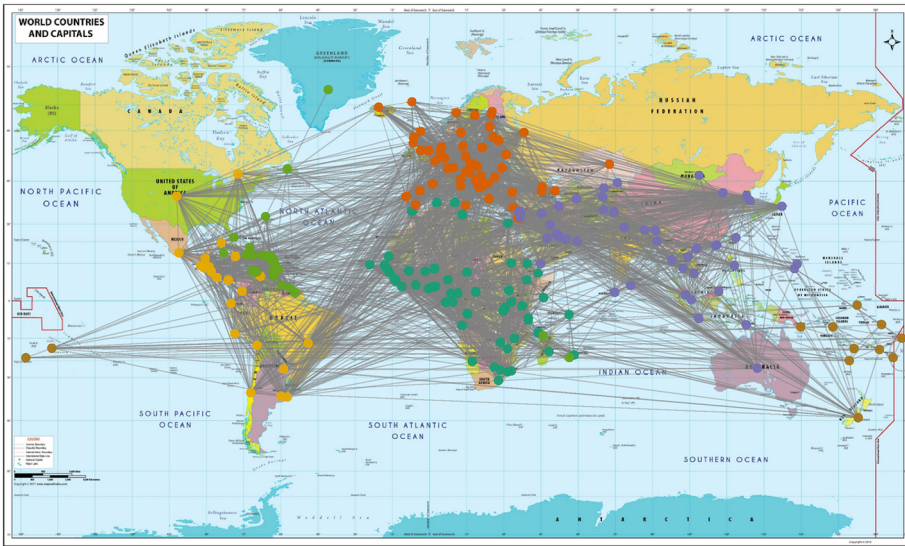


Fig 10. Result of applying the Louvain algorithm [2] to the network of international football matches. Nodes with the same color belong to the same community.

less than those in the real graph, we found that most of the times this is because countries from Oceania tend to be clustered with the AFC.

Besides simply looking at the amount of communities, we compared whether clusters in the simulated graphs matched those of the real network. To that end, we computed three different metrics of cluster agreement between each synthetic graph and the real one. Those metrics were V-measure [22], the Adjusted Rand Index [11], and the Adjusted Mutual Information [33]. On the right panel of Figure 11 we show a boxplot for each metric. Apparently, the overall community structure in the synthetic graphs follows closely that of the actual network.

Remark 4.2. Since the estimation of p_0^{ij} is based on the consistency result from Section 3.1, which guarantees convergence as the number of nodes $N \rightarrow \infty$, the finite size of the graph introduces a non-negligible finite-sample deviation. This deviation partially accounts for the (arguably quite minor) discrepancies observed in the metrics shown in Figure 9, underscoring the method's robustness.

Remark 4.3. The entropy-based graph regeneration procedure described in this section provides a way to reproduce weighted networks that match a finite set of estimated moments. This naturally suggests the possibility of using regenerated graphs to approximate the distribution of network summary statistics, such as subgraph counts, degree-based measures, or spectral functionals, in a bootstrap-like fashion. Related bootstrap methodologies have been developed recently for latent-space network models in the context of the vanilla (i.e., unweighted) RDPG model. In particular, the bootstrap framework put forth in [16] establishes consistency results for U-statistic-type graph functionals and for resampling entire networks under an RDPG generative assumption. These rely on consistent estimation of the latent positions and on representations of many network statistics as functionals of these positions.

While the elegant ideas in [16] provide a valuable conceptual foundation, extending such bootstrap guarantees to the present WRDPG framework is nontrivial. Here, regenerated networks are obtained via an entropy maximization principle using only finitely many moments,




Fig 11. Comparisons between community structure of the real football matches network and its synthetic replicates. Left: histogram of number of communities in the largest connected component (lcc) for synthetic graphs. Right: Boxplot for three metrics of clustering agreement between real and synthetic networks: V-measure, Adjusted Rand Index (ARI) and Adjusted Mutual Information (AMI).

and the latent structure itself may vary across moments. A rigorous bootstrap theory would therefore need to account for both moment truncation effects and the distributional properties induced by the entropy-based graph regeneration scheme. We regard such developments as a meaningful direction for future research.

5. Conclusions

We developed the Weighted Random Dot Product Graph (WRDPG) model, which extends the vanilla Random Dot Product Graph (RDPG) framework to effectively capture the intricacies of weighted graphs characterized by heterogeneous edge weight distributions. Our approach assigns latent positions to nodes, parameterizing edge weight distributions' moments using moment-generating functions. This innovative modeling strategy enhances the WRDPG's capability to distinguish between weight distributions that may have the same mean but differ in higher-order moments, thus providing a richer representation of network data than previous attempts to weighted RDPG modeling. We derived statistical guarantees (consistency and asymptotic normality) for our estimator of latent positions, leveraging the well-established adjacency spectral embedding method. While for simplicity we assume a common embedding dimension across all moment matrices, the WRDPG framework can be extended to allow moment-specific dimensions, and we expect robustness to dimension overspecification similar to that observed in the vanilla RDPG [28].

Furthermore, we developed a generative framework that approximates desired weight distributions from a given (or estimated) latent position sequence and facilitates sampling synthetic graphs that closely mimic the structure and characteristics of real-world networks. Through a series of illustrative examples, we have demonstrated the effectiveness of the WRDPG model in practical applications. We show how our model can recover latent position structures and replicate key graph metrics, exhibiting a discriminative power essential for community

detection tasks. We have also shown that the model accommodates various edge weight distributions, ranging from discrete to continuous, or even mixed random variables. To achieve this versatility, we developed an improved computational method to estimate a continuous density by maximizing the entropy subject to moment constraints. This result is of independent interest and could permeate benefits in other, broader contexts. In light of this, a natural direction for future research is the development of bootstrap-based inference methods for the WRDPG model. Establishing rigorous guarantees in this setting, however, would require controlling both moment truncation effects and the distributional features induced by the entropy maximization scheme, extending recent bootstrap results available for the vanilla RDPG framework [16]. The findings of this study highlight the importance of incorporating higher-order moments in graph models, contributing to the methods of inference and generation in network data analysis. All techniques developed in this paper can be extended to heterophilous graphs, i.e., graphs where the moment matrices are not necessarily positive semidefinite, by bringing to bear the generalized RDPG model introduced in [23]. Asymptotic results remain tractable in this setting, since our methodology is based on that of [8], which explicitly accounts for heterophilic graphs in their model. Future work may explore further adaptations of the WRDPG model, including enhancements in (real-time) computational efficiency and applications to evolving or dynamic networks. An additional promising avenue concerns the development of clustering guarantees in settings where community structure is identifiable only through higher-order moments. Such results would require new identifiability and separation conditions across moments, as well as a careful analysis of the trade-offs inherent in leveraging higher-order information. A related open problem is the design of principled methods for aggregating information across moment matrices, such as joint embedding approaches, together with theoretical guarantees tailored to the WRDPG setting, where latent structures may vary across moments.

Appendix A: Proof of Proposition 3.2

Here we state and prove some auxiliary results that will be necessary in order to prove Proposition 3.2. We begin with a lemma that establishes a bound on the MGF of a sub-Weibull random variable near the origin.

Lemma A.1. *Let Y be a sub-Weibull random variable with parameter $\theta > 1$; that is, there exists a constant $C > 0$ such that*

$$\mathbb{P}[|Y| \geq t] \leq 2 \exp\left(-\left(\frac{t}{C}\right)^{1/\theta}\right) \quad \text{for all } t \geq 0.$$

Then there exist $T > 1$ and $C_1 > 0$, depending only on θ and C , such that:

$$\mathbb{E}\left[e^{\lambda|Y|}\right] \leq \exp\left(-C_1 \lambda^{1/(1-\theta)}\right) \quad \text{for all } \lambda \in [0, \lambda_0],$$

where

$$\lambda_0 = \frac{1}{\theta C^{1/\theta}} T^{(1-\theta)/\theta}.$$

Proof. By Fubini’s theorem, we can write:

$$\mathbb{E} [e^{\lambda|Y|}] = \int_0^\infty \lambda e^{\lambda t} \mathbb{P} [|Y| \geq t] dt.$$

Using the sub-Weibull tail bound, we obtain:

$$\mathbb{E} [e^{\lambda|Y|}] \leq 2\lambda \int_0^\infty \exp \left(\lambda t - \left(\frac{t}{C} \right)^{1/\theta} \right) dt.$$

Let $\phi_\lambda(t) = \lambda t - \left(\frac{t}{C} \right)^{1/\theta}$, and define:

$$I(\lambda) = \int_0^\infty \exp(\phi_\lambda(t)) dt.$$

The function $\phi_\lambda(t)$ attains its unique maximum at

$$t_\lambda = \left(\lambda \theta C^{1/\theta} \right)^{\theta/(1-\theta)}.$$

To apply Laplace’s method (see [6, Chapter 4]), we require $t_\lambda \geq T$ for some large enough $T > 1$, so that the integrand is sharply peaked near t_λ . Solving $t_\lambda \geq T$ yields:

$$\lambda \leq \lambda_0 = \frac{1}{\theta C^{1/\theta}} T^{(1-\theta)/\theta}.$$

Under this condition, Laplace’s method gives:

$$I(\lambda) \leq K \exp(\phi_\lambda(t_\lambda)),$$

for some constant $K > 0$ depending only on θ and C . Evaluating $\phi_\lambda(t_\lambda)$, we find:

$$\phi_\lambda(t_\lambda) = \left(\frac{1}{\theta} - 1 \right) (\theta C)^{1/(1-\theta)} \lambda^{1/(1-\theta)}.$$

Since $\theta > 1$, we define

$$C_1 = - \left(\frac{1}{\theta} - 1 \right) (\theta C)^{1/(1-\theta)} > 0,$$

so that

$$\phi_\lambda(t_\lambda) = -C_1 \lambda^{1/(1-\theta)}.$$

Putting everything together,

$$\mathbb{E} [e^{\lambda|Y|}] \leq 2\lambda I(\lambda) \leq 2\lambda K \exp \left(-C_1 \lambda^{1/(1-\theta)} \right).$$

Since $\lambda_0 \rightarrow 0$ as $T \rightarrow \infty$, we can choose a T large enough so that $2\lambda_0 K \leq 1$. Therefore, for all $\lambda \in [0, \lambda_0]$ we have:

$$\mathbb{E} [e^{\lambda|Y|}] \leq 2\lambda_0 K \exp \left(-C_1 \lambda^{1/(1-\theta)} \right) \leq \exp \left(-C_1 \lambda^{1/(1-\theta)} \right),$$

as claimed. □

We next present a concentration result for the sum of independent, centered sub-Weibull rvs with a common parameter θ . The proof follows standard arguments used in analogous concentration results for sub-Gaussian and sub-exponential rvs (see Theorems 2.6.3 and 2.8.1 in [32]).

Proposition A.2. *Let Y_1, \dots, Y_N be independent, mean-zero, sub-Weibull random variables with parameter $\theta > 1$. That is, for each $i = 1, \dots, N$, there exists a constant $C_i > 0$ such that*

$$\mathbb{P} [|Y_i| \geq t] \leq 2 \exp \left(- \left(\frac{t}{C_i} \right)^{1/\theta} \right) \quad \text{for all } t \geq 0.$$

For any $T > 0$, define $\lambda_{\min} = \frac{1}{\theta C_{\max}^{1/\theta}} T^{(1-\theta)/\theta}$, where $C_{\max} = \max_{i=1, \dots, N} C_i$.

Then there exists a large enough T such that for all $t \geq 0$, it holds that

$$\mathbb{P} \left[\left| \sum_{i=1}^N Y_i \right| \geq t \right] \leq 2 \exp \left(- \min \left\{ \lambda_{\min} t + NK_1 \lambda_{\min}^{1/(1-\theta)}, \left(\frac{t}{K_2} \right)^{1/\theta} N^{(\theta-1)/\theta} \right\} \right),$$

where $K_1, K_2 > 0$ are constants depending only on θ .

Proof. Let $S := \sum_{i=1}^N Y_i$. By Markov’s inequality and independence:

$$\mathbb{P} [S \geq t] = \mathbb{P} [e^{\lambda S} \geq e^{\lambda t}] \leq e^{-\lambda t} \prod_{i=1}^N \mathbb{E} [e^{\lambda Y_i}].$$

Now, using Lemma A.1, there exists a large enough $T > 1$ such that for every $i = 1, \dots, N$ there exists a constant $C_{1,i} > 0$ (depending only on C_i and θ) such that for all $\lambda \in [0, \lambda_{\min}]$:

$$\mathbb{E} [e^{\lambda |Y_i|}] \leq \exp \left(-C_{1,i} \lambda^{1/(1-\theta)} \right).$$

Using the fact that Y_i is mean-zero, we have:

$$\mathbb{E} [e^{\lambda Y_i}] \leq \mathbb{E} [e^{\lambda |Y_i|}] \leq \exp \left(-C_{1,i} \lambda^{1/(1-\theta)} \right).$$

Hence,

$$\mathbb{P} [S \geq t] \leq \exp \left(-\lambda t - \sum_{i=1}^N C_{1,i} \lambda^{1/(1-\theta)} \right).$$

Let $K_1 = \min_i C_{1,i}$. Then K_1 depends only on θ , since the $C_{1,i}$ ’s depend only on θ . Therefore:

$$\mathbb{P} [S \geq t] \leq \exp \left(-\lambda t - NK_1 \lambda^{1/(1-\theta)} \right). \tag{17}$$

Minimizing the exponent over $\lambda \in [0, \lambda_{\min}]$, the optimal choice is:

$$\lambda^* = \min \left\{ \lambda_{\min}, \left(\frac{t(\theta - 1)}{NK_1} \right)^{(1-\theta)/\theta} \right\}.$$

Substituting into (17), we obtain:

$$\mathbb{P} [S \geq t] \leq \exp \left(- \min \left\{ \lambda_{\min} t + NK_1 \lambda_{\min}^{1/(1-\theta)}, \left(\frac{t}{K_2} \right)^{1/\theta} N^{(\theta-1)/\theta} \right\} \right),$$

where $K_2 = \frac{\theta^\theta}{(\theta-1)^{\theta-1}}(K_1^{(\theta-1)/\theta})$ depends only on θ .

A symmetric argument for $-S$ gives the same bound for $\mathbb{P}[-S \geq t]$, and the result follows. \square

We next state a lemma that establishes a bound on the difference between the matrix of entry-wise self-products, $\mathbf{W}^{(k)}$, and the true moments matrix, $\mathbf{X}[k]\mathbf{X}^\top[k]$. The proof follows the general scheme of [8, Proposition 1], but differs in a key aspect: we rely on a different result for tail bounds of sums of independent matrices. This is due to the fact that the concentration result they invoke does not hold in our more general setup.

Lemma A.3. *Let $(\mathbf{W}, \mathbf{X}_k) \sim \text{WRDPG}(F)$ satisfy the hypotheses of Theorem 3.1. For each fixed integer $k \geq 0$, let $\mathbf{W}_k := \mathbf{W}^{(k)}$ and $\mathbf{M}_k := \mathbf{X}[k]\mathbf{X}^\top[k]$. Then*

$$\|\mathbf{W}_k - \mathbf{M}_k\|_2 = O_{\mathbb{P}}\left(\log^{k\theta} N\right).$$

Proof. Note that the definition of our model implies that $\mathbb{E}[\mathbf{W}_k] = \mathbf{M}_k$, so we are trying to control the spectral norm of the centered matrix $\mathbf{W}_k - \mathbf{M}_k$. To that end, we will use a tail bound for sums of independent matrices [31, Corollary 3.7]. That result can be stated as follows: assume we have a finite sequence $\{\mathbf{Y}_l\}_{l=1}^R$ of independent, self-adjoint, random $N \times N$ matrices that satisfy

$$\mathbb{E}\left[e^{\lambda \mathbf{Y}_l}\right] \preceq e^{g(\lambda)\mathbf{A}_l} \quad \text{for all } \lambda \in [0, \lambda_0]$$

for some function $g : [0, \lambda_0] \rightarrow \mathbb{R}$ and a finite sequence $\{\mathbf{A}_l\}_{l=1}^R$ of fixed self-adjoint matrices. Here, \preceq denotes the semidefinite ordering on Hermitian matrices, i.e., $\mathbf{A} \succeq \mathbf{B}$ if $\mathbf{A} - \mathbf{B}$ is positive semi-definite. Define the scale parameter

$$\rho = \left\| \sum_{l=1}^R \mathbf{A}_l \right\|_2.$$

Then for all $t \geq 0$ it holds:

$$\mathbb{P}\left[\left\| \sum_{l=1}^R \mathbf{Y}_l \right\|_2 \geq t\right] \leq N \inf_{\lambda \in [0, \lambda_0]} \exp(-\lambda t + g(\lambda)\rho).$$

In order to use this result, for each $1 \leq j \leq i \leq N$ let \mathbf{Y}_{ij} be the $N \times N$ matrix whose (i, j) and (j, i) entries equal $W_{ij}^k - \mathbf{x}_i^\top[k]\mathbf{x}_j[k] := Y_{ij}$, with the rest being 0. Then, $\sum_{ij} \mathbf{Y}_{ij} = \mathbf{W}_k - \mathbf{M}_k$ and $\mathbb{E}[\mathbf{Y}_{ij}] = \mathbf{0}$ since $\mathbb{E}[Y_{ij}] = 0$. Also note that if \mathbf{e}_i is the i -th vector of the canonical basis of \mathbb{R}^N , we have:

$$\mathbf{Y}_{ij} = Y_{ij}(\mathbf{e}_i \mathbf{e}_j^\top + \mathbf{e}_j \mathbf{e}_i^\top) \Rightarrow \mathbf{Y}_{ij}^p = \begin{cases} Y_{ij}^p(\mathbf{e}_i \mathbf{e}_j^\top + \mathbf{e}_j \mathbf{e}_i^\top) & \text{if } p \text{ is odd} \\ Y_{ij}^p(\mathbf{e}_i \mathbf{e}_i^\top + \mathbf{e}_j \mathbf{e}_j^\top) & \text{if } p \text{ is even, } p \geq 2 \end{cases}.$$

Therefore:

$$\begin{aligned} e^{\lambda \mathbf{Y}_{ij}} &= \mathbf{I} + \sum_{p \text{ odd}} \frac{(\lambda Y_{ij})^p}{p!} (\mathbf{e}_i \mathbf{e}_j^\top + \mathbf{e}_j \mathbf{e}_i^\top) + \sum_{p \text{ even, } p \geq 2} \frac{(\lambda Y_{ij})^p}{p!} (\mathbf{e}_i \mathbf{e}_i^\top + \mathbf{e}_j \mathbf{e}_j^\top) \\ &= \mathbf{I} + \sinh(\lambda Y_{ij})(\mathbf{e}_i \mathbf{e}_j^\top + \mathbf{e}_j \mathbf{e}_i^\top) + (\cosh(\lambda Y_{ij}) - 1)(\mathbf{e}_i \mathbf{e}_i^\top + \mathbf{e}_j \mathbf{e}_j^\top). \end{aligned}$$

This means that $e^{\lambda \mathbf{Y}_{ij}}$ has ones along the diagonal, except at the (i, i) and (j, j) entries, which are equal to $\cosh(\lambda Y_{ij})$. Its off-diagonal entries are zero, except for the (i, j) and (j, i) entries, which are equal to $\sinh(\lambda Y_{ij})$. Using the identity $e^{|x|} - \cosh(x) = |\sinh(x)|$ for all $x \in \mathbb{R}$, together with the Gershgorin circle theorem, we obtain:

$$e^{\lambda \mathbf{Y}_{ij}} \preceq \mathbf{I} + (e^{|\lambda Y_{ij}|} - 1)(\mathbf{e}_i \mathbf{e}_i^\top + \mathbf{e}_j \mathbf{e}_j^\top) = \exp\left(|\lambda Y_{ij}|(\mathbf{e}_i \mathbf{e}_i^\top + \mathbf{e}_j \mathbf{e}_j^\top)\right),$$

where the last equality follows from the fact that the matrix $\mathbf{I} + (e^{|\lambda Y_{ij}|} - 1)(\mathbf{e}_i \mathbf{e}_i^\top + \mathbf{e}_j \mathbf{e}_j^\top)$ is diagonal, with all entries equal to 1 except at the (i, i) and (j, j) positions, which are equal to $e^{|\lambda Y_{ij}|}$. Since the matrix expectation preserves the semidefinite order, this in turn implies

$$\mathbb{E}\left[e^{\lambda \mathbf{Y}_{ij}}\right] \preceq \mathbb{E}\left[\exp\left(|\lambda Y_{ij}|(\mathbf{e}_i \mathbf{e}_i^\top + \mathbf{e}_j \mathbf{e}_j^\top)\right)\right].$$

As discussed in Remark 3.1, Assumption 2 implies that Y_{ij} is a sub-Weibull random variable with parameter $k\theta$. Therefore, by Lemma A.1, there exists a constant $C_1^{i,j}$ and a threshold $\lambda_0^{i,j} > 0$ such that for all $\lambda \in [0, \lambda_0^{i,j}]$,

$$\mathbb{E}\left[e^{\lambda \mathbf{Y}_{ij}}\right] \preceq \mathbb{E}\left[\exp\left(-C_1^{i,j} \lambda^{1/(1-k\theta)}(\mathbf{e}_i \mathbf{e}_i^\top + \mathbf{e}_j \mathbf{e}_j^\top)\right)\right].$$

Letting $K_1 = \min_{i,j} C_1^{i,j}$ and $\lambda_{\min} = \min_{i,j} \lambda_0^{i,j}$, we obtain that for all $\lambda \in [0, \lambda_{\min}]$

$$\mathbb{E}\left[e^{\lambda \mathbf{Y}_{ij}}\right] \preceq \mathbb{E}\left[\exp\left(-K_1 \lambda^{1/(1-k\theta)}(\mathbf{e}_i \mathbf{e}_i^\top + \mathbf{e}_j \mathbf{e}_j^\top)\right)\right].$$

Thus, we can apply the tail bound for sums of independent random matrices stated above, with $g(\lambda) = -K_1 \lambda^{1/(1-k\theta)}$ and $\mathbf{A}_{ij} = \mathbf{e}_i \mathbf{e}_i^\top + \mathbf{e}_j \mathbf{e}_j^\top$. Therefore,

$$\mathbb{P}\left[\|\mathbf{W}_k - \mathbf{M}_k\|_2 \geq t\right] \leq N \inf_{\lambda \in [0, \lambda_{\min}]} \exp\left(-\lambda t - K_1 \lambda^{1/(1-k\theta)} \rho\right), \tag{18}$$

where

$$\rho = \left\| \sum_{1 \leq i \leq j \leq N} (\mathbf{e}_i \mathbf{e}_i^\top + \mathbf{e}_j \mathbf{e}_j^\top) \right\|_2 = \|\mathbf{N}\mathbf{I}\|_2 = N.$$

Note that the function inside the infimum in (18) is equivalent to that on the right-hand side of (17). Therefore, an argument analogous to the one in the proof of Proposition A.2 yields

$$\mathbb{P}\left[\|\mathbf{W}_k - \mathbf{M}_k\|_2 \geq t\right] \leq N \exp\left(-\min\left\{\lambda_{\min} t + N K_1 \lambda_{\min}^{1/(1-k\theta)}, \left(\frac{t}{K_2}\right)^{1/(k\theta)} N^{(k\theta-1)/(k\theta)}\right\}\right),$$

where $K_2 = \frac{(k\theta)^{k\theta}}{(k\theta-1)^{k\theta-1}} (K_1^{(k\theta-1)/(k\theta)})$. Since for large t the term involving $t^{1/(k\theta)}$ dominates the minimum, we have that, for large enough t ,

$$\mathbb{P}\left[\|\mathbf{W}_k - \mathbf{M}_k\|_2 \geq t\right] \leq N \exp\left(-\left(\frac{t}{K_2}\right)^{1/(k\theta)} N^{(k\theta-1)/(k\theta)}\right).$$

Choosing $t = K_2 \log^{k\theta} N$ then yields the desired result. □

The following propositions can be proven with the same arguments found in [18, Proposition 16, Lemma 17] and [8, Propositions 3 and 5], respectively. While we omit the detailed proofs, we provide a brief outline of the key ideas underlying each result.

Proposition A.4 follows from a Procrustes-alignment style argument, which leverages the relationship between the principal angles of the subspaces spanned by the columns of \mathbf{U}_k and $\hat{\mathbf{U}}_k$, and the eigenvalues of the difference $\hat{\mathbf{U}}_k \hat{\mathbf{U}}_k^\top - \mathbf{U}_k \mathbf{U}_k^\top$, together with an application of the Davis–Kahan theorem [36]. Proposition A.6 can be proved using similar techniques.

Proposition A.5 is derived via a Hoeffding-type concentration argument analogous to that used in the proof of Theorem 3.1. Lastly, Proposition A.7 is a direct consequence of Propositions A.4 and A.6.

Proposition A.4. *For each k , let*

$$\mathbf{U}_k^\top \hat{\mathbf{U}}_k = \mathbf{S}_{k_1} \boldsymbol{\Sigma}_k \mathbf{S}_{k_2}^\top$$

be the singular value decomposition of $\mathbf{U}_k^\top \hat{\mathbf{U}}_k$. If $\mathbf{S}_k := \mathbf{S}_{k_1} \mathbf{S}_{k_2}^\top$, then it holds that

$$\|\mathbf{U}_k^\top \hat{\mathbf{U}}_k - \mathbf{S}_k\|_F = O_{\mathbb{P}}\left(N^{-1} \log^{k\theta} N\right).$$

Proposition A.5. *For each k , the following holds:*

$$\|\mathbf{U}_k^\top (\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k\|_F = O_{\mathbb{P}}\left(\log^{k\theta} N\right).$$

Proposition A.6. *For each k , the following holds:*

$$\begin{aligned} \|\hat{\mathbf{U}}_k \hat{\mathbf{U}}_k^\top - \mathbf{U}_k \mathbf{U}_k^\top\|_F &= O_{\mathbb{P}}\left(N^{-1} \log^{k\theta} N\right) \\ \|\hat{\mathbf{U}}_k - \mathbf{U}_k \mathbf{U}_k^\top \hat{\mathbf{U}}_k\|_F &= O_{\mathbb{P}}\left(N^{-1} \log^{k\theta} N\right) \end{aligned}$$

Proposition A.7. *For each k , let \mathbf{S}_k be as in Proposition A.4. Then the following holds:*

$$\begin{aligned} \|\mathbf{S}_k \hat{\mathbf{D}}_k - \mathbf{D}_k \mathbf{S}_k\|_F &= O_{\mathbb{P}}\left(\log^{k\theta} N\right), \\ \|\mathbf{S}_k \hat{\mathbf{D}}_k^{1/2} - \mathbf{D}_k^{1/2} \mathbf{S}_k\|_F &= O_{\mathbb{P}}\left(N^{-1/2} \log^{k\theta} N\right) \\ \|\mathbf{S}_k \hat{\mathbf{D}}_k^{-1/2} - \mathbf{D}_k^{-1/2} \mathbf{S}_k\|_F &= O_{\mathbb{P}}\left(N^{-3/2} \log^{k\theta} N\right). \end{aligned}$$

We are now ready to prove Proposition 3.2. We remind the reader its statement:

Proposition 3.2. For each fixed integer $k \geq 0$, let $\mathbf{W}_k := \mathbf{W}^{(k)}$. Also let $\mathbf{M}_k := \mathbf{X}[k] \mathbf{X}^\top[k]$ and denote its spectral decomposition as $\mathbf{M}_k = \mathbf{U}_k \mathbf{D}_k \mathbf{U}_k^\top$. Let $\mathbf{S}_k \in O(d)$ be as in Proposition A.7 and define \mathbf{R}_{k_1} , \mathbf{R}_{k_2} , \mathbf{R}_{k_3} and \mathbf{R}_{k_4} as

$$\begin{aligned} \mathbf{R}_{k_1} &= \mathbf{U}_k \left(\mathbf{U}_k^\top \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{1/2} - \mathbf{D}_k^{1/2} \mathbf{S}_k \right), \\ \mathbf{R}_{k_2} &= (\mathbf{I} - \mathbf{U}_k \mathbf{U}_k^\top) (\mathbf{W}_k - \mathbf{M}_k) \left(\hat{\mathbf{U}}_k - \mathbf{U}_k \mathbf{S}_k \right) \hat{\mathbf{D}}_k^{-1/2}, \\ \mathbf{R}_{k_3} &= -\mathbf{U}_k \mathbf{U}_k^\top (\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k \mathbf{S}_k \hat{\mathbf{D}}_k^{-1/2}, \\ \mathbf{R}_{k_4} &= (\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k \left(\mathbf{S}_k \hat{\mathbf{D}}_k^{-1/2} - \mathbf{D}_k^{-1/2} \mathbf{S}_k \right). \end{aligned}$$

Then the following holds:

$$\begin{aligned}\|\mathbf{R}_{k_1}\|_{2 \rightarrow \infty} &= \mathcal{O}_{\mathbb{P}}\left(N^{-1} \log^{k\theta} N\right), \\ \|\mathbf{R}_{k_2}\|_{2 \rightarrow \infty} &= \mathcal{O}_{\mathbb{P}}\left(N^{-3/2} \log^{2k\theta} N\right), \\ \|\mathbf{R}_{k_3}\|_{2 \rightarrow \infty} &= \mathcal{O}_{\mathbb{P}}\left(N^{-1} \log^{k\theta} N\right), \\ \|\mathbf{R}_{k_4}\|_{2 \rightarrow \infty} &= \mathcal{O}_{\mathbb{P}}\left(N^{-3/2} \log^{2k\theta} N\right).\end{aligned}$$

Proof. In order to bound $\mathbf{R}_{k_1} = \mathbf{U}_k \left(\mathbf{U}_k^{\top} \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{1/2} - \mathbf{D}_k^{1/2} \mathbf{S}_k \right)$, note that:

$$\begin{aligned}\|\mathbf{R}_{k_1}\|_{2 \rightarrow \infty} &\leq \|\mathbf{U}_k\|_{2 \rightarrow \infty} \left\| \mathbf{U}_k^{\top} \hat{\mathbf{U}}_k \hat{\mathbf{D}}_k^{1/2} - \mathbf{D}_k^{1/2} \mathbf{S}_k \right\|_2 \\ &\leq \|\mathbf{U}_k\|_{2 \rightarrow \infty} \left(\left\| \left(\mathbf{U}_k^{\top} \hat{\mathbf{U}}_k - \mathbf{S}_k \right) \hat{\mathbf{D}}_k^{1/2} \right\|_F + \left\| \mathbf{S}_k \hat{\mathbf{D}}_k^{1/2} - \mathbf{D}_k^{1/2} \mathbf{S}_k \right\|_F \right).\end{aligned}$$

Since $\mathbf{U}_k \mathbf{D}_k^{1/2} = \mathbf{X}[k] \mathbf{Q}_k$ with $\mathbf{Q}_k \in \mathcal{O}(d)$, we have that

$$\|\mathbf{U}_k\|_{2 \rightarrow \infty} \leq \|\mathbf{X}[k]\|_{2 \rightarrow \infty} \left\| \mathbf{D}_k^{-1/2} \right\|_2.$$

Since the inner product between the rows of $\mathbf{X}[k]$ equals the k -th moment of some random variable—which, by the definition of our model, we assume to be finite for all k —this implies that all rows of $\mathbf{X}[k]$ (which lie in \mathbb{R}^d) must have bounded norm. Therefore, $\|\mathbf{X}[k]\|_{2 \rightarrow \infty}$ is finite and does not depend on N . In Lemma B.1 of Appendix B, we show that Assumption 1 implies all nonzero eigenvalues of \mathbf{M}_k are $\Theta_{\mathbb{P}}(N)$, so $\left\| \mathbf{D}_k^{-1/2} \right\|_2 = \mathcal{O}_{\mathbb{P}}(N^{-1/2})$. Consequently, $\|\mathbf{U}_k\|_{2 \rightarrow \infty} = \mathcal{O}_{\mathbb{P}}(N^{-1/2})$. Combining this with Propositions A.4 and A.7 yields the desired bound for \mathbf{R}_{k_1} .

In a similar vein, for $\mathbf{R}_{k_3} = -\mathbf{U}_k \mathbf{U}_k^{\top} (\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k \mathbf{S}_k \hat{\mathbf{D}}_k^{-1/2}$ it holds:

$$\begin{aligned}\|\mathbf{R}_{k_3}\|_{2 \rightarrow \infty} &\leq \|\mathbf{U}_k\|_{2 \rightarrow \infty} \left\| \mathbf{U}_k^{\top} (\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k \mathbf{S}_k \hat{\mathbf{D}}_k^{-1/2} \right\|_2 \\ &\leq \|\mathbf{U}_k\|_{2 \rightarrow \infty} \left\| \mathbf{U}_k^{\top} (\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k \right\|_F \left\| \hat{\mathbf{D}}_k^{-1/2} \right\|_2.\end{aligned}$$

Using Proposition A.5 and the fact that both $\|\mathbf{U}_k\|_{2 \rightarrow \infty}$ and $\left\| \hat{\mathbf{D}}_k \right\|_2$ are $\mathcal{O}_{\mathbb{P}}(N^{-1/2})$ (the latter being a consequence of Assumptions 1 and 2, as shown in Lemma B.2 of Appendix B) then implies that $\|\mathbf{R}_{k_3}\|_{2 \rightarrow \infty}$ is $\mathcal{O}_{\mathbb{P}}(N^{-1} \log^{k\theta} N)$, as desired.

Regarding $\mathbf{R}_{k_4} = (\mathbf{W}_k - \mathbf{M}_k) \mathbf{U}_k \left(\mathbf{S}_k \hat{\mathbf{D}}_k^{-1/2} - \mathbf{D}_k^{-1/2} \mathbf{S}_k \right)$ we have that:

$$\begin{aligned}\|\mathbf{R}_{k_4}\|_{2 \rightarrow \infty} &\leq \|\mathbf{W}_k - \mathbf{M}_k\|_{\infty} \left\| \mathbf{U}_k \left(\mathbf{S}_k \hat{\mathbf{D}}_k^{-1/2} - \mathbf{D}_k^{-1/2} \mathbf{S}_k \right) \right\|_{2 \rightarrow \infty} \\ &\leq \sqrt{N} \|\mathbf{W}_k - \mathbf{M}_k\|_2 \|\mathbf{U}_k\|_{2 \rightarrow \infty} \left\| \mathbf{S}_k \hat{\mathbf{D}}_k^{-1/2} - \mathbf{D}_k^{-1/2} \mathbf{S}_k \right\|_2,\end{aligned}$$

where in the second inequality we have used that $\|\mathbf{A}\|_{\infty} \leq \sqrt{n} \|\mathbf{A}\|_2$ for all $\mathbf{A} \in \mathbb{R}^{n \times m}$. Using Lemma A.3 and Proposition A.7 together with the fact that $\|\mathbf{U}_k\|_{2 \rightarrow \infty}$ is $\mathcal{O}_{\mathbb{P}}(N^{-1/2})$ then shows the desired bound for \mathbf{R}_{k_4} .

This leaves us with the task of bounding the $2 \rightarrow \infty$ norm of \mathbf{R}_{k_2} , which we remind the reader is defined as:

$$\mathbf{R}_{k_2} = (\mathbf{I} - \mathbf{U}_k \mathbf{U}_k^{\top}) (\mathbf{W}_k - \mathbf{M}_k) \left(\hat{\mathbf{U}}_k - \mathbf{U}_k \mathbf{S}_k \right) \hat{\mathbf{D}}_k^{-1/2}.$$

Therefore:

$$\|\mathbf{R}_{k_2}\|_{2 \rightarrow \infty} \leq \|\mathbf{I} - \mathbf{U}_k \mathbf{U}_k^\top\|_{2 \rightarrow \infty} \|\mathbf{W}_k - \mathbf{M}_k\|_2 \|\hat{\mathbf{U}}_k - \mathbf{U}_k \mathbf{S}_k\|_2 \|\hat{\mathbf{D}}_k^{-1/2}\|_2. \quad (19)$$

For the first term we have that:

$$\|\mathbf{I} - \mathbf{U}_k \mathbf{U}_k^\top\|_{2 \rightarrow \infty} \leq \|\mathbf{I}\|_{2 \rightarrow \infty} + \|\mathbf{U}_k \mathbf{U}_k^\top\|_{2 \rightarrow \infty} \leq 1 + \|\mathbf{U}_k\|_{2 \rightarrow \infty} \|\mathbf{U}_k^\top\|_2.$$

Since $\|\mathbf{U}_k^\top\|_2 = \|\mathbf{U}_k\|_2 \leq \sqrt{N} \|\mathbf{U}_k\|_{2 \rightarrow \infty}$ we have that $\|\mathbf{U}_k^\top\|_2 = \mathcal{O}_{\mathbb{P}}(1)$ because $\|\mathbf{U}_k\|_{2 \rightarrow \infty} = \mathcal{O}_{\mathbb{P}}(N^{-1/2})$. Therefore

$$\|\mathbf{I} - \mathbf{U}_k \mathbf{U}_k^\top\|_{2 \rightarrow \infty} \leq 1 + \mathcal{O}_{\mathbb{P}}(N^{-1/2}) = \mathcal{O}_{\mathbb{P}}(1).$$

Also note that:

$$\begin{aligned} \|\hat{\mathbf{U}}_k - \mathbf{U}_k \mathbf{S}_k\|_2 &\leq \|\hat{\mathbf{U}}_k - \mathbf{U}_k \mathbf{U}_k^\top \hat{\mathbf{U}}_k\|_2 + \left\| \mathbf{U}_k \left(\mathbf{U}_k^\top \hat{\mathbf{U}}_k - \mathbf{S}_k \right) \right\|_2 \\ &\leq \|\hat{\mathbf{U}}_k - \mathbf{U}_k \mathbf{U}_k^\top \hat{\mathbf{U}}_k\|_2 + \|\mathbf{U}_k\|_2 \|\mathbf{U}_k^\top \hat{\mathbf{U}}_k - \mathbf{S}_k\|_2. \end{aligned}$$

By Proposition A.6 the first term is $\mathcal{O}_{\mathbb{P}}(N^{-1} \log^{k\theta} N)$, whereas the second term is $\mathcal{O}_{\mathbb{P}}(N^{-1} \log^{k\theta} N)$ because $\|\mathbf{U}_k\|_2 = \mathcal{O}_{\mathbb{P}}(1)$ and $\|\mathbf{U}_k^\top \hat{\mathbf{U}}_k - \mathbf{S}_k\|_2$ is $\mathcal{O}_{\mathbb{P}}(N^{-1} \log^{k\theta} N)$ by virtue of Proposition A.4. Therefore:

$$\|\hat{\mathbf{U}}_k - \mathbf{U}_k \mathbf{S}_k\|_2 = \mathcal{O}_{\mathbb{P}}(N^{-1} \log^{k\theta} N).$$

Combining this with Lemma A.3 and the fact that $\|\hat{\mathbf{D}}_k^{-1/2}\|_2$ is $\mathcal{O}_{\mathbb{P}}(N^{-1/2})$, from (19) we conclude $\|\mathbf{R}_{k_2}\|_{2 \rightarrow \infty} = \mathcal{O}_{\mathbb{P}}(N^{-3/2} \log^{2k\theta} N)$, which concludes the proof. \square

Appendix B: Consequences of Assumptions 1 and 2 regarding the largest eigenvalues of \mathbf{M}_k and \mathbf{W}_k

In this section, we show that Assumptions 1 and 2 imply that for all $k \geq 0$, the nonzero eigenvalues of $\mathbf{M}_k := \mathbf{X}[k] \mathbf{X}^\top[k]$ are $\Theta_{\mathbb{P}}(N)$, and that the same holds for the top d eigenvalues of $\mathbf{W}_k := \mathbf{W}^{(k)}$. We begin by proving this for \mathbf{M}_k . The proof scheme follows that of [27, Proposition 4.3]. In what follows, the notation $\lambda_i(\mathbf{A})$ denotes the i -th largest-magnitude eigenvalue of matrix \mathbf{A} .

Lemma B.1. *Let $(\mathbf{W}, \mathbf{X}_k) \sim \text{WRDPG}(F)$, where F satisfies Assumption 1 and \mathbf{W} satisfies Assumption 2. Then for all $k \geq 0$ it holds that $\lambda_i(\mathbf{M}_k) = \Theta_{\mathbb{P}}(N)$ for $i = 1, \dots, d$, while $\lambda_i(\mathbf{M}_k) = 0$ for $i = d + 1, \dots, N$.*

Proof. That $\lambda_i(\mathbf{M}_k) = 0$ for $i = d + 1, \dots, N$ is immediate since $\mathbf{M}_k = \mathbf{X}[k] \mathbf{X}^\top[k]$, so \mathbf{M}_k is a $N \times N$ matrix with rank at most d . For $i = 1, \dots, d$, note that

$$\lambda_i(\mathbf{M}_k) = \lambda_i(\mathbf{X}[k] \mathbf{X}^\top[k]) = \lambda_i(\mathbf{X}^\top[k] \mathbf{X}[k]).$$

Now, each entry of $\mathbf{X}^\top[k] \mathbf{X}[k]$ is:

$$(\mathbf{X}^\top[k] \mathbf{X}[k])_{ij} = \sum_{l=1}^N (\mathbf{x}_l[k])_i (\mathbf{x}_l[k])_j,$$

where $(\mathbf{x}_l[k])_i$ denotes the i -th entry of vector $\mathbf{x}_l[k]$. Then $(\mathbf{X}^\top[k]\mathbf{X}[k])_{ij}$ is a sum of independent random variables, each with expectation $\mathbb{E}[(\mathbf{x}_l[k])_i(\mathbf{x}_l[k])_j] = (\Delta_k)_{ij}$, where $(\Delta_k)_{ij}$ is the (i, j) entry of the second moment matrix Δ_k from Assumption 1. Since, for each k , the inner product between the rows of $\mathbf{X}[k]$ equals the k -th moment of some random variable—which, by the definition of our model, we assume to be finite for all k —it follows that all rows of $\mathbf{X}[k]$ (which lie in \mathbb{R}^d) must have bounded norm, with a bound that does not depend on N . Therefore, each term in the above sum is bounded by some constant $L_k > 0$, so by Hoeffding’s inequality [32, Theorem 2.6.2] we have that:

$$\mathbb{P}\left(|(\mathbf{X}^\top[k]\mathbf{X}[k])_{ij} - N(\Delta_k)_{ij}| \geq t\right) \leq 2 \exp\left(\frac{-2t^2}{N^2 L_k^2}\right).$$

Choosing $t = \frac{L_k}{\sqrt{2}} N^{1/2} \log^{1/2} N$ then shows that

$$|(\mathbf{X}^\top[k]\mathbf{X}[k])_{ij} - N(\Delta_k)_{ij}| = \mathcal{O}_{\mathbb{P}}\left(N^{1/2} \log^{1/2} N\right).$$

Taking a union bound then shows that $\|\mathbf{X}^\top[k]\mathbf{X}[k] - N\Delta_k\|_F$ is also $\mathcal{O}_{\mathbb{P}}\left(N^{1/2} \log^{1/2} N\right)$, and since the spectral norm is dominated by the Frobenius norm, then $\|\mathbf{X}^\top[k]\mathbf{X}[k] - N\Delta_k\|_2 = \mathcal{O}_{\mathbb{P}}\left(N^{1/2} \log^{1/2} N\right)$. A corollary of Weyl’s inequality [10, Corollary 7.3.5] then implies that:

$$|\lambda_i(\mathbf{X}^\top[k]\mathbf{X}[k]) - N\lambda_i(\Delta_k)| = \mathcal{O}_{\mathbb{P}}\left(N^{1/2} \log^{1/2} N\right).$$

Since $\lambda_i(\Delta_k) = \Theta_{\mathbb{P}}(1)$, this in turn implies that $\lambda_i(\mathbf{X}^\top[k]\mathbf{X}[k]) = \Theta_{\mathbb{P}}(N)$, which completes the proof. \square

The next lemma shows that the top d eigenvalues of \mathbf{W}_k are $\Theta_{\mathbb{P}}(N)$, while the remaining ones are within $N^{1/2} \log^{1/2} N$ of zero with high probability.

Lemma B.2. *Under the same hypotheses as Lemma B.1, for all $k \geq 0$ it holds that $\lambda_i(\mathbf{W}_k) = \Theta_{\mathbb{P}}(N)$ for $i = 1, \dots, d$, while $\lambda_i(\mathbf{W}_k) = \mathcal{O}_{\mathbb{P}}(\log^{k\theta} N)$ for $i = d + 1, \dots, N$.*

Proof. Again, by [10, Corollary 7.3.5] we have that:

$$|\lambda_i(\mathbf{W}_k) - \lambda_i(\mathbf{M}_k)| \leq \|\mathbf{W}_k - \mathbf{M}_k\|_2.$$

Therefore, using Lemma A.3 we have that $|\lambda_i(\mathbf{W}_k) - \lambda_i(\mathbf{M}_k)| = \mathcal{O}_{\mathbb{P}}(\log^{k\theta} N)$. Applying Lemma B.1 implies the desired result. \square

Appendix C: Chebyshev polynomial reformulation

We provide a detailed description of the Chebyshev-based reformulation of the moment-recovery system in (12). By expressing the system on the basis of Chebyshev polynomials of the first kind, we markedly improve the problem’s conditioning.

C.1. Transformation of monomials into Chebyshev basis

Any monomial x^k of degree $k \leq K$ can be uniquely expanded in the Chebyshev basis $\{T_j(x) : j = 0, 1, \dots, K\}$:

$$x^k = \sum_{j=0}^K c_{k,j} T_j(x),$$

where the coefficients $c_{k,j}$ are computed using recurrence relations or discrete orthogonal integrals. In practice, these coefficients can be precomputed up to degree K using the recurrence:

$$\begin{aligned} T_0(x) &= 1, \\ T_1(x) &= x, \\ T_{j+1}(x) &= 2xT_j(x) - T_{j-1}(x), \quad j \geq 1, \end{aligned}$$

and the identity relating monomials to Chebyshev polynomials (e.g., via trigonometric definitions or forward recurrence).

C.2. Construction of the Chebyshev–Vandermonde matrix

Let $\{v_0, v_1, \dots, v_R\}$ denote the support points of the discrete distribution. We define the Chebyshev–Vandermonde matrix $\mathbf{V}_C \in \mathbb{R}^{(K+1) \times (R+1)}$ with entries

$$(\mathbf{V}_C)_{k,r} = T_k(v_r^*), \quad k = 0, \dots, K, \quad r = 0, \dots, R,$$

where v_r^* are the support points mapped into the canonical interval $[-1, 1]$ via affine scaling if necessary.

The moment-recovery system then becomes an equivalent system in the Chebyshev basis. Let $\mathbf{C} \in \mathbb{R}^{(K+1) \times (K+1)}$ be the matrix with entries $C_{kj} = c_{k,j}$, so that the Chebyshev-moment vector

$$\mathbf{m}_C = \mathbf{C}\mathbf{m}$$

is the projection of the original moment vector $\mathbf{m} = [m_{ij}[0], \dots, m_{ij}[K]]^\top$ onto the Chebyshev basis. The system in the Chebyshev basis then reads

$$\mathbf{V}_C \mathbf{p} = \mathbf{m}_C.$$

C.3. Numerical solution and stability

Because Chebyshev polynomials constitute an orthogonal basis on $[-1, 1]$ with respect to the weight $(1 - x^2)^{-1/2}$, the condition number of \mathbf{V}_C grows only polynomially in K , rather than exponentially as in the monomial basis. Therefore, solving for \mathbf{p} via a standard least-squares or regularized inversion

$$\hat{\mathbf{p}} = \underset{\mathbf{p} \geq 0}{\operatorname{argmin}} \|\mathbf{V}_C \mathbf{p} - \mathbf{m}\|_2^2$$

offers improved numerical stability. Additional regularization or non-negativity constraints may be imposed to further enhance stability and ensure valid probability estimates.

Appendix D: Solution to the moments problem by entropy maximization

In Section 4.2, we formulated the problem of identifying a pdf that maximizes differential entropy subject to a set of moment constraints. Here we detail how this constrained optimization problem can be effectively tackled using a primal-dual approach. Specifically, we derive the corresponding dual formulation and present a gradient-based method for solving it, which is both computationally efficient and well-suited to our setting.

The problem is to find a pdf g_{ij} that maximizes the differential entropy

$$S(g_{ij}) = - \int_{\Omega} g_{ij}(x) \log g_{ij}(x) dx$$

subject to the moments constraints

$$\int_{\Omega} x^k g_{ij}(x) dx = \mathbf{x}_i^T [k] \mathbf{x}_j [k], \quad \text{for all } k = 0, 1, \dots, K, \tag{20}$$

where $\Omega \subset \mathbb{R}$ denotes the known support of the distribution.

From an optimization viewpoint, we wish to solve the primal problem:

$$\max_{g \in \mathcal{F}} S(g) \quad \text{s. to} \quad \int_{\Omega} x^k g(x) dx - m_k = 0, \quad \text{for all } k = 0, 1, \dots, K, \tag{21}$$

where $m_k = \mathbf{x}_i^T [k] \mathbf{x}_j [k]$ and \mathcal{F} is the space of all probability distributions that are supported in Ω (as before, dependence of g and m_k on i and j is omitted for readability).

To solve (21), we introduce Lagrange multipliers λ_k and the Lagrangian functional³

$$L(g, \lambda) = S(g) - (\lambda_0 - 1) \left(\int_{\Omega} g(x) dx - m_0 \right) - \sum_{k=1}^K \lambda_k \left(\int_{\Omega} x^k g(x) dx - m_k \right), \tag{22}$$

where $\lambda = [\lambda_0, \lambda_1, \dots, \lambda_K]^T \in \mathbb{R}^{K+1}$. We will find a solution to (21) by solving the dual problem:

$$\min_{\lambda \in \mathbb{R}^{K+1}} d(\lambda), \tag{23}$$

where $d : \mathbb{R}^{K+1} \mapsto \mathbb{R}$ is the dual function

$$d(\lambda) = \max_{g \in \mathcal{F}} L(g, \lambda).$$

Using the Euler-Lagrange equation for the calculus of variations we conclude that the function g_{λ} that extremizes L is the one that makes the functional derivative $\frac{\delta L(g, \lambda)}{\delta g(x)}$ equal to zero. By differentiation, we see that such a condition amounts to

$$\left. \frac{\delta L(g, \lambda)}{\delta g(x)} \right|_{g_{\lambda}(x)} = 0 \Leftrightarrow g_{\lambda}(x) = \exp \left(- \sum_{k=0}^K \lambda_k x^k \right). \tag{24}$$

Looking at the second variation of $L(g, \lambda)$ we conclude that this g_{λ} maximizes the Lagrangian, so the dual function can be computed as

$$d(\lambda) = L(g_{\lambda}, \lambda) = \sum_{k=0}^K \lambda_k m_k + \int_{\Omega} \exp \left(- \sum_{k=0}^K \lambda_k x^k \right) dx - m_0.$$

³Following [13] we have used $\lambda_0 - 1$ as the first Lagrange multiplier for convenience.

One can check that $d(\lambda)$ is a convex function of $\lambda_0, \lambda_1, \dots, \lambda_K$ (see, e.g., [13, Example 2.3]). Since its partial derivatives are equal to

$$\frac{\partial d(\lambda)}{\partial \lambda_k} = m_k - \int_{\Omega} x^k g_{\lambda}(x) dx$$

we have that whenever the gradient of $d(\lambda)$ is zero, then g_{λ} satisfies the moments constraints (20). In other words, if λ^* is a solution to the dual problem (23), then the function g_{λ^*} given by (24) both maximizes the Lagrangian (22) and satisfies the moments constraints (20), i.e., it is a solution to the primal problem (21). This implies that strong duality holds for the maximum entropy problem. Therefore, by solving the dual problem, we can recover the maximum entropy pdf we are seeking. Since the dual objective is convex and unconstrained, it can be solved using any standard convex optimization algorithm. In practice, we solve it using the BFGS algorithm [21] available through SciPy’s `minimize` function.

Remark D.1. If a function g_{λ} such as that in (24) satisfies the moments constraints (20), then it has maximum entropy among all pdfs g supported in Ω that satisfy such constraints. Indeed, we have that, since $\log g_{\lambda}(x) = -\sum_{k=0}^K \lambda_k x^k$,

$$\begin{aligned} S(g_{\lambda}) &= -\int_{\Omega} g_{\lambda}(x) \log g_{\lambda}(x) dx = \sum_{k=0}^K \lambda_k \int_{\Omega} g_{\lambda}(x) x^k dx = \sum_{k=0}^K \lambda_k \int_{\Omega} g(x) x^k dx \\ &= -\int_{\Omega} g(x) \log g_{\lambda}(x) dx, \end{aligned}$$

where in the third equality we have used that both g_{λ} and g satisfy the moments constraints (20). Then

$$S(g_{\lambda}) - S(g) = \int_{\Omega} g(x) \log \left(\frac{g(x)}{g_{\lambda}(x)} \right) dx = D_{KL}(g||g_{\lambda}) \geq 0$$

where D_{KL} is the Kullback-Leibler divergence.

Example. We showcase our method by testing its ability to recover an exponential distribution from its first four moments. Since the pdf of an exponential random variable with parameter α is $\alpha e^{-\alpha x}$, it can be written in the form of (24) by setting $\lambda_0 = -\log \alpha$, $\lambda_1 = \alpha$ and $\lambda_k = 0$ for all $k \geq 2$, so we can check whether the Lagrange multipliers λ_k obtained with our gradient descent solution are close to these values. The results for an exponential random variable with parameter $\alpha = 2$ are depicted in Figure 12, where we show the λ_k ’s upon convergence of our GD algorithm for 100 random initializations of λ . We also show the Lagrange multipliers obtained with the method proposed in [24]. In that work, the authors estimate λ by using Newton’s method to approximate the solutions of the system of nonlinear equations that arises from imposing that the function g_{λ} in (24) satisfies the moments constraints (20), i.e., the system

$$\int_{\Omega} x^k \exp \left(-\lambda_0 - \lambda_1 x - \lambda_2 x^2 - \dots - \lambda_R x^R \right) dx = m_k \quad \forall k = 0, \dots, K.$$

As shown in Figure 12 the Lagrange multipliers that we obtain using this approach are, most of the time, quite far from the actual values. This is because Newton’s method is sensitive to initialization and heavily relies on having an initial guess that lies on the solution’s basin




Fig 12. Box plots of Lagrange multipliers for maximum entropy estimation of an exponential rv distribution via gradient-descent (red) and the method from [24] (PyMaxEnt, green) for 100 random initializations. Gradient descent always converges to the true value, while PyMaxEnt does not.

of attraction. Our method, in contrast, always converges to the true solution. That is because our dual function is convex, so as long as the maximum entropy problem has a solution our method will converge to it no matter where it starts.

Remark D.2. This maximum entropy approach can be easily modified to accommodate discrete distributions. This is useful when we have access to fewer moments than the amount of symbols the discrete random variable takes, in which case we can no longer use the procedure described in Section 4.1.

Assuming that our random variable takes on the values v_0, \dots, v_R , we wish to find its corresponding probabilities p_0, \dots, p_R by maximizing Shannon’s entropy

$$S(p_0, \dots, p_R) = - \sum_{r=0}^R p_r \log p_r$$

subject to the moments constrains

$$\sum_{r=0}^R v_r^k p_r = m_k \quad \forall k = 0, \dots, K.$$

We define the Lagrangian in analogy to (22) and by differentiation we find that it has a maximum when $p_r = \exp \left(- \sum_{k=0}^K \lambda_k v_r^k \right)$. Thus, the dual function can be computed as:

$$d(\lambda) = \sum_{k=0}^K \lambda_k m_k + \sum_{r=0}^R \exp \left(- \sum_{k=0}^K \lambda_k v_r^k \right) - m_0,$$

which, as before, is a convex function of $\lambda_0, \dots, \lambda_K$. Therefore, we can find its minimizer using any standard convex optimization algorithm and from it compute the p_r ’s that maximize Shannon’s entropy.

Acknowledgments

Part of the results in this paper were presented at the *2024 Asilomar Conference on Signals, Systems, and Computers* [20].

The first, second, third and fourth authors are also with Centro Interdisciplinario en Ciencia de Datos y Aprendizaje Automático (CICADA), Universidad de la República, Uruguay

The authors would like to thank Matías Valdés for helpful discussions on optimization related to entropy maximization. We would also like to thank the anonymous reviewers for the valuable feedback and suggestions that led to an improved revised manuscript.

Funding

The first, second, third and fourth authors were partially supported by CSIC (I+D project 22520220100076UD) and ANII (SticAmsud LAGOON project).

The fifth author was supported in part by NSF award EECs-2231036.

Supplementary material

Accompanying code

The code to generate Figures 3 to 12 in this work is available at <https://github.com/bmarengo/wrdpg>.

References

- [1] ATHREYA, A., FISHKIND, D. E., TANG, M., PRIEBE, C. E., PARK, Y., VOGELSTEIN, J. T., LEVIN, K., LYZINSKI, V. and QIN, Y. (2017). Statistical inference on random dot product graphs: A survey. *J. Mach. Learn. Res.* **18** 8393–8484. [MR3827114](#)
- [2] BLONDEL, V. D., GUILLAUME, J.-L., LAMBIOTTE, R. and LEFEBVRE, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiments* **2008** P10008. [MR4832781](#)
- [3] BOURIN, C. and BONDON, P. (1998). Efficiency of high-order moment estimates. *IEEE Transactions on Signal Processing* **46** 255–258.
- [4] BOYD, J. P. and GALLY, D. H. (2007). Numerical experiments on the accuracy of the Chebyshev–Frobenius companion matrix method for finding the zeros of a truncated series of Chebyshev polynomials. *Journal of Computational and Applied Mathematics* **205** 281–295. [MR2324840](#)
- [5] CAPE, J., TANG, M. and PRIEBE, C. E. (2019). The two-to-infinity norm and singular subspace geometry with applications to high-dimensional statistics. *The Annals of Statistics* **47** pp. 2405–2439. [MR3988761](#)
- [6] DE BRUIJN, N. G. (2014). *Asymptotic methods in analysis*. Courier Corporation. [MR0671583](#)
- [7] DEFORD, D. R. and ROCKMORE, D. N. (2016). A random dot product model for weighted networks. *arXiv:1611.02530 [stat.AP]*.
- [8] GALLAGHER, I., JONES, A., BERTIGER, A., PRIEBE, C. E. and RUBINDELANCHY, P. (2023). Spectral embedding of weighted graphs. *J. Am. Stat. Assoc.* **0** 1–10. [MR4797912](#)

- [9] HOFF, P. D., RAFTERY, A. E. and HANDCOCK, M. S. (2002). Latent space approaches to social network analysis. *J. Am. Stat. Assoc.* **97** 1090–1098. [MR1951262](#)
- [10] HORN, R. A. and JOHNSON, C. R. (2012). *Matrix Analysis*. Cambridge University Press. [MR2978290](#)
- [11] HUBERT, L. and ARABIE, P. (1985). Comparing partitions. *Journal of Classification* **2** 193–218.
- [12] IZENMAN, A. J. (2023). *Network Models for Data Science: Theory, Algorithms, and Applications*. Cambridge University Press.
- [13] KAPUR, J. N. and KESAVAN, H. K. (1992). *Entropy Optimization Principles with Applications*. Academic Press. [MR1196459](#)
- [14] KOLACZYK, E. (2017). *Topics at the Frontier of Statistics and Network Analysis: (Re)Visiting the Foundations. SemStat Elements*. Cambridge University Press. [MR3702038](#)
- [15] KOLACZYK, E. D. (2009). *Statistical Analysis of Network Data: Methods and Models*. Springer. [MR2724362](#)
- [16] LEVIN, K. and LEVINA, E. (2025). Bootstrapping networks with latent space structure. *Electronic Journal of Statistics* **19** 745–791. [MR4870301](#)
- [17] LI, Y. and MATEOS, G. (2022). Networks of international football: communities, evolution and globalization of the game. *Applied Network Science* **7**.
- [18] LYZINSKI, V., TANG, M., ATHREYA, A., PARK, Y. and PRIEBE, C. E. (2017). Community detection and classification in hierarchical stochastic blockmodels. *IEEE Trans. Netw. Sci. Eng.* **4** 13–26. [MR3625952](#)
- [19] MARENCO, B., BERMOLLEN, P., FIORI, M., LARROCA, F. and MATEOS, G. (2022). Online Change Point Detection for Weighted and Directed Random Dot Product Graphs. *IEEE Trans. Signal Inf. Process. Netw.* **8** 144–159. [MR4397333](#)
- [20] MARENCO, B., BERMOLLEN, P., FIORI, M., LARROCA, F. and MATEOS, G. (2024). A Random Dot Product Graph Model for Weighted and Directed Networks. In *2024 58th Asilomar Conference on Signals, Systems, and Computers* 1155–1159.
- [21] NOCEDAL, J. and WRIGHT, S. J. (2006). *Numerical Optimization*, 2nd ed. Springer. [MR2244940](#)
- [22] ROSENBERG, A. and HIRSCHBERG, J. (2007). V-measure: A conditional entropy-based external cluster evaluation measure. In *Proc. of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)* 410–420.
- [23] RUBIN-DELANCHY, P., CAPE, J., TANG, M. and PRIEBE, C. E. (2022). A statistical interpretation of spectral embedding: The generalised random dot product graph. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **84** 1446–1473. [MR4494166](#)
- [24] SAAD, T. and RUAI, G. (2019). PyMaxEnt: A Python software for maximum entropy moment reconstruction. *SoftwareX* **10** 100353.
- [25] SCHEINERMAN, E. R. and TUCKER, K. (2010). Modeling graphs using dot product representations. *Comput. Stat* **25** 1–16. [MR2586721](#)
- [26] SHORE, J. and JOHNSON, R. (1980). Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on Information Theory* **26** 26–37. [MR0560389](#)

- [27] SUSSMAN, D. L., TANG, M. and PRIEBE, C. E. (2014). Consistent latent position estimation and vertex classification for random dot product graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **36** 48–57.
- [28] TAING, R. and LEVIN, K. (2026). On the Effect of Misspecifying the Embedding Dimension in Low-rank Network Models. *arXiv preprint arXiv:2601.06014*.
- [29] TANG, R., KETCHA, M., BADEA, A., CALABRESE, E. D., MARGULIES, D. S., VOGELSTEIN, J. T., PRIEBE, C. E. and SUSSMAN, D. L. (2018). Connectome smoothing via low-rank approximations. *IEEE Trans. Med. Imaging* **38** 1446–1456.
- [30] TANG, R., TANG, M., VOGELSTEIN, J. T. and PRIEBE, C. E. (2017). Robust estimation from multiple graphs under gross error contamination. *arXiv:1707.03487 [stat.ME]*.
- [31] TROPP, J. A. (2012). User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics* **12** 389–434. [MR2946459](#)
- [32] VERSHYNIN, R. (2018). *High-Dimensional Probability: An Introduction with Applications in Data Science*. *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press. [MR3837109](#)
- [33] VINH, N. X., EPPS, J. and BAILEY, J. (2009). Information theoretic measures for clusterings comparison: is a correction for chance necessary? In *Proc. International Conference on Machine Learning (ICML)* 1073–1080. [MR2738784](#)
- [34] VLADIMIROVA, M., GIRARD, S., NGUYEN, H. and ARBEL, J. (2020). Sub-Weibull distributions: Generalizing sub-Gaussian and sub-Exponential properties to heavier tailed distributions. *Stat* **9** e318. [MR4193421](#)
- [35] WIDDER, D. V. (1941). *The Laplace transform*. *Princeton Mathematical Series* **6**. Princeton University Press. [MR0005923](#)
- [36] YU, Y., WANG, T. and SAMWORTH, R. J. (2015). A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika* **102** 315–323. [MR3371006](#)
- [37] ZHU, M. and GHODSI, A. (2006). Automatic dimensionality selection from the scree plot via the use of profile likelihood. *Comput Stat Data Anal* **51** 918–930. [MR2297497](#)