Ranking Networks

Amirali Salehi-Abari
Department of Computer Science
University of Toronto
abari@cs.toronto.edu

Craig Boutilier
Department of Computer Science
University of Toronto
cebly@cs.toronto.edu

Abstract

Latent space models for network formation assume that nodes possess latent attributes that determine their propensity to connect. We propose a new model for network formation, *ranking networks*, in which these attributes are *rankings* over some space of alternatives. Such rankings may reflect user preferences, relevance/quality judgements, etc., while ranking networks capture correlations of, say, user preferences across a social network. We present preliminary theoretical and empirical analyses of structural properties of such networks, and develop algorithmic approximations to help efficiently predict these properties. Empirical results demonstrate the quality of these approximations.

1 Introduction

Network formalisms offer a powerful tool for studying the interactions of entities in complex systems. Of special interest in network science is the development of *network formation models* that explain the emergence of common structural properties of real-world networks [2, 34]. Recent work has focused on modeling the role of both known and hidden attributes of nodes in forming networks—the intuition being that each node possesses attributes (e.g., geographical location [35] or social position [14, 7]) that determine the chance that any pair of nodes become connected. Examples of such models include spatial networks [4], random geometric graphs [31], latent space models [14, 7], and multiplicative attribute graphs [19].

Most of these models take attributes to be binary, real-valued, or integral. Surprisingly, little work has studied attribute-based network formation in which the underlying attributes correspond to *rankings* over some space of options. Such rankings may represent the preferences of agents over some set of alternatives (e.g., products, political candidates, jobs) or their subjective opinions about the quality or relevance of certain items (web pages, sports teams). Because of the increasing availability of ranked data due to recommender systems, web search, online surveying, and other applications, a deeper understanding of the possible interactions between ranked data and network formation seems imperative. In this work, we develop and begin a preliminary analysis of *ranking networks*, in which each node in the network possesses a ranking, and the similarity between two nodes' rankings determines their affinity or probability of connection. Such networks can be used for a variety of purposes, but we are especially interested in models where rankings reflect *agent preferences* over a some space of options. We discuss the extent to which we can leverage the structure of ranking networks to address algorithmic questions that arise in the analysis of networks. For instance, what topological network properties emerge under such a model? How to model parameters (including the ranking distribution and similarity metric) influence these properties? And can such structural properties be predicted efficiently?

Ranking network models can be deployed in a variety of contexts. It is widely acknowledged that the behaviours of individuals are correlated over the social networks [10, 15]; and correlated preferences offer one possible explanation of this phenomenon. Ranking networks can help predict or elicit the preferences of specific individuals to develop more efficient recommender, social choice and

advertising mechanisms. They may also have application in information retrieval and topic modeling over the web: given some common set of attributes (e.g., topics, tags), each node (e.g., web page, photo) may have a "relevance ranking" over those attributes. A ranking network can capture the correlation of such relevance rankings given the structure of underlying information network (and may help improve the efficiency of recently developed comparison-based interactive search methods [16]).The predictive power of ranking networks might be also harnessed for link prediction [24, 22]—to help predict future interactions between entities—or *link mining*—to help infer unobserved links between entities [13].

We introduce a probabilistic model for the formation of *ranking networks*, in which each entity possesses a ranking over a set of alternatives. The similarity of their rankings determines the chance of connectivity between two individuals. We first analyze some general topological properties of ranking networks (e.g., connectivity conditions, graph diameter, degree distribution, edge density, clustering coefficient). We then study distance-based ranking networks in which similarity is measured using standard ranking distance metrics, examine a few additional properties in this setting, and describe several algorithmic approximations to efficiently predict structural properties of the resulting network. Our approximation methods are also applicable for learning model parameters and inference problem. We present preliminary empirical results using the *Mallows' ϕ-model* [27], a distributional model of rankings used widely in psychometrics and machine learning, and well-suited to modeling user preferences. These results confirm the efficiency and predictive accuracy of our computational approximations.

2 Related Work

The pioneering random graph model, $G(n, p)$, of Erdős and Rènyi [11] assumes an edge between any pair of nodes occurs with probability p. Other seminal work includes the preferential attachment [2] and small-world [34] models. Since then network formation models have drawn considerable attention, with a variety of models proposed to capture structural properties of real-world networks [29, 3, 30].One can roughly divide these model into two categories. *Static models* (e.g., Erdős-Rènyi, small world) specify the network using simple static properties (e.g., probability of edge occurrence), whereas *dynamic models* (e.g., preferential attachment) involve some dynamic "growth" process by which nodes and edges are added to a network.

Our ranking network model is a random, static model. It is also falls within the class of *spatial (or latent space) networks* [4, 14] in which nodes have a set of real-valued, binary, or integer-valued *latent variables*, with the probability of an edge forming between two nodes determined by their attributes. The *Waxman random graph model* [35], often used to model the computer networks, distributes n nodes uniformly at random on the plane, with two nodes connected with probability that decreases exponentially with their Euclidean distance. Hoff *et al.* [14] develop a similar model for inferring latent variables on social networks—nodes are points in a d-dimensional Euclidean "social space." The *random dot product model* [36] is somewhat distinct in that the distance between nodes is given by the inner product of their position vectors. The *hidden variable model* [6, 7] is a generalization of the Waxman model: nodes are equipped with a hidden (real-valued or integer) random variable drawn independently from a specified distribution. Two nodes are connected according to a symmetric probability function over node attributes. Serrano et al.[33] slightly modify this approach to introduce a general class of models based on a hidden metric space where nodes—located at a specific point in this space—are connected with a probability determined by a *connection probability function* defined over node distances. The *multiplicative attribute graphs (MAG) model* [19] generalizes the Kronecker graph model [20], assigning each node a vector of categorical attributes, and using a link-affinity matrix to capture connection probabilities based on attribute interactions.

Within spatial network formation models, recent research has addressed *inference and learning* of latent attributes and social network structure. Hoff *et al.* [14] develop inference and learning techniques to infer node locations in a fixed d-dimensional Euclidean "social space." Using the MAG model, Kim and Leskovec [17] develop a scalable *variational expectation maximization* method for learning model parameters given network structure and node attributes, and inferring node's latent variables and model parameters when the network structure is observed. Kim and Leskovec [18] extend this to allow nodes to have multiple latent characteristics. Other work on learning on graphs includes *collective classification* [32] and active learning over network data [5, 26].

3 The Formation of Ranking Networks

We first provide background on latent space models for network formation. We then define *ranking networks* by building upon these models and discuss some general theoretical properties. Finally, we examine the special case of *distance-based ranking networks*, analysing some of its theoretical properties and developing computational approximations to estimate these properties.

3.1 Latent Space Network Models

Latent space or hidden variable network models [6] generally assume a set n nodes, where each node i is associated with a latent variable h_i. This represents some feature (or feature vector) of the node or individual in question. The parameters of such a model are given by a parameter vector $\boldsymbol{\theta}$ (we discuss model-specific parameterizations below). A random undirected graph is generated by assuming: (1) a distribution $\rho(h|\boldsymbol{\eta})$ specifying the probability that an arbitrary node takes on value h; and (2) a symmetric *connection probability function*, $c(h, h^{'})$, where $c(h_i, h_j)$ denotes the probability that edge e_{ij} forms between nodes i and j given their variable values. A common variant of this model defines the connection probability using the relative distance d_{ij} between the values h_i and h_j in some metric space. In this case, we assume a function $c(d) : [0, \infty) \rightarrow [0, 1]$ that maps distances into connection probabilities. We generally assume distance-based connection functions below. A canonical example of such a connection probability is Waxman's [35] function $c(d) = \beta e^{-\frac{d}{d_0}}$, widely used for modeling computer network topologies, where β controls edge density and d_0 represents a "typical" distance between nodes. Another is $c(d) = (1 + \frac{d}{\beta})^{-\alpha}$, used for modeling social networks [7]. Here β controls the average node degree while $\alpha > 1$ determines the degree of homophily. Let $\boldsymbol{\lambda}$ denote the (vector of) model parameters governing connection function $c(\cdot)$ and $\boldsymbol{\theta} = (\boldsymbol{\eta}, \boldsymbol{\lambda})$ the vector of all parameters of a latent space model.

We now describe certain structural properties that emerge under latent space models. Because we will deal with discrete ranking spaces below, we assume latent variables are discrete. First, note that *edge density*, i.e., the probability of an occurring edge between two arbitrary nodes, is given by:

$$\mathcal{E}(\boldsymbol{\theta}) = \sum_{h,h^{'}} \rho(h|\boldsymbol{\eta})\rho(h^{'}|\boldsymbol{\eta})c\left(d(h^{'}, h)|\boldsymbol{\lambda}\right). \tag{1}$$

As a result, the expected number of edges of a random node is given by $\binom{n}{2}\mathcal{E}(\boldsymbol{\theta})$. The probability of edge between a node i with observed value h_i and a random alternative node j is:

$$\mathcal{D}(h_i, \boldsymbol{\theta}) = \sum_{h^{'}} \rho(h^{'}|\boldsymbol{\eta})c\left(d(h^{'}, h_i)|\boldsymbol{\lambda}\right). \tag{2}$$

Thus, the expected degree of a node with value h_i is given by $(n-1)\mathcal{D}(h_i, \boldsymbol{\theta})$. Note that $\mathcal{E}(\boldsymbol{\theta})$ can be rewritten as:

$$\mathcal{E}(\boldsymbol{\theta}) = \sum_{h} \rho(h|\boldsymbol{\eta})\mathcal{D}(h, \boldsymbol{\theta}). \tag{3}$$

The exact *degree distribution* $P(k, \boldsymbol{\theta})$ is:

$$P(k, \boldsymbol{\theta}) = \sum_{h} \rho(h|\boldsymbol{\eta})\mathcal{G}(h, k|\boldsymbol{\theta}), \tag{4}$$

where $\mathcal{G}(h, k|\boldsymbol{\theta})$ is the probability that a node with value h has k neighbours:

$$\mathcal{G}(h, k|\boldsymbol{\theta}) = \binom{n-1}{k}\mathcal{D}(h, \boldsymbol{\theta})^k \left(1 - \mathcal{D}(h, \boldsymbol{\theta})\right)^{n-1-k}. \tag{5}$$

The *clustering coefficient* of a node i is the fraction of the *pairs* of neighboring nodes that are themselves connected. The expected clustering coefficient for a node i with value h_i can be computed by:

$$\mathcal{C}(h_i, \boldsymbol{\theta}) = \frac{1}{\mathcal{D}(h_i, \boldsymbol{\theta})^2} \sum_{h^{'}, h^{''}} \rho(h^{'})\rho(h^{''})c(d(h_i, h^{'}))c(d(h_i, h^{''}))c(d(h^{'}, h^{''})) \tag{6}$$

One can compute the average clustering by $\langle \mathcal{C}(\boldsymbol{\theta}) \rangle = \sum_h \rho(h, \boldsymbol{\eta})\mathcal{C}(h, \boldsymbol{\theta})$.

3.2 Ranking Networks: A General Model

We adapt the general latent space network model to setting where the latent attributes are *rankings* over some set of alternatives. For instance, nodes might represent individuals with attributes reflecting their preferences over some set of products, services, political candidates, etc. We assume a finite set of alternatives (or options) $\mathcal{A} = \{a_1, \ldots, a_m\}$ and a set of nodes $\mathcal{N} = \{1, \ldots, n\}$. Each node i has a ranking (or strict total order) over $\mathcal{A}$, denoted by $\succ_i$ (weaker notions, e.g., preorders, partial orders, can be accommodated, though some details of our model require modification). Let $\Omega(\mathcal{A})$ denote the set of all $m!$ rankings over $\mathcal{A}$. Our latent variables are rankings r drawn from $\Omega(\mathcal{A})$. We assume that each node i's ranking r_i is drawn independently from some (parameterized) distribution $\rho(r|\boldsymbol{\eta})$ over $\Omega(\mathcal{A})$. We also assume a *ranking distance metric* $d : \Omega(\mathcal{A})^2 \to \mathbb{R}$ which measures similarity between rankings. Finally, a *connection probability function* $c(d) : [0, \infty) \to [0, 1]$ determines the probability that two nodes i, j are connected given the distance $d(r_i, r_j)$ between their rankings. We now detail each of these three components of our model.

Distance Metric on Rankings. We use the "similarity" of two rankings to determine their distance, which will be used below to determine connection probabilities. A variety of well-known distance metrics for rankings can be used [9]. We briefly describe several common distance metrics. A natural set of distances are "d_p distances", where $d_p(r, r') = \sum_{i=1}^{m} |r(a_i) - r'(a_i)|^p$ for $p \in [1, \infty)$. The well-known *footrule* ($p = 1$) and *Spearman* ($p = 2$) distances are instances of this. *Hamming* distance is another natural model, while *Kendall's* τ distance is used, e.g., in psychometrics and social choice, where

$$d_\tau(r, r') = \sum_{k \neq l} I[r(a_k) > r(a_l) \text{ and } r'(a_k) < r'(a_l)]. \tag{7}$$

Here d_τ measures the number of pairwise swaps needed to transform r to r'. As the ranking space $\Omega(\mathcal{A})$ is discrete with finite size $m!$, there are finitely many realizable distances. It is easy to see that, for any $r, r' : 0 \leq d_\tau(r, r') \leq \binom{m}{2}$ and $d_\tau(r, r') \in \mathbb{N} \cup \{0\}$.

Ranking Distributions. The $\rho(r|\boldsymbol{\eta})$ component of our ranking network model accommodates arbitrary ranking distributions. Distributional models of rankings developed in psychometrics and statistics, and now widely used in machine learning and IR [23], include Mallows, Plackett-Luce, Bradley-Terry, and many others (see [28] for an overview). We use the Mallows' ϕ-model in our empirical experiments. It is characterized by a "modal" *reference ranking* σ and a *dispersion parameter* $\phi \in [0, 1)$, with the probability of a ranking r given by $\rho(r|\sigma, \phi) \propto \phi^{d_\tau(r,\sigma)}$.

Connection Probability Function. With ranking-based distance metrics in hand, we adopt standard connection functions for latent-space models (see above). We assume $c(d)$ is integrable and strictly decreasing. We can derive sufficient conditions for connectivity of a ranking network:

Theorem 3.1 *Assume a ranking model $(\rho(r|\boldsymbol{\eta}), c(d|\boldsymbol{\lambda}))$. The induced ranking network is connected with high probability (i.e., with probability $1 - o(1)$ where $o(1) \to 0$ and $n \to \infty$) if*

$$d_M(m) < c^{-1}\left(\frac{\log n}{n}\Big|\boldsymbol{\lambda}\right),$$

where $d_M(m)$ is the maximum possible distance under d given m alternatives.[1]

This can be used to derive suitable conditions for connectivity of specific models. For instance, using the τ distance metric and Waxman connection function $c(x|\alpha, d_0) = \alpha e^{\frac{x}{d_0}}$, the emerging ranking network will be connected with high probability if $\binom{m}{2} < -d_0 \ln \frac{\log n}{\alpha n}$.

The *small world* effect is a commonly observed property of real-world networks: the diameter (or longest shortest path between any pair of nodes) is small, as is the average shortest path length [1].

[1] We assume that m is fixed and independent of the number n of individuals. This is consistent with the applications we have in mind (e.g., social choice, recommender systems, interactive search, etc.). We plan to extend our analysis to the case where m depends on n. Omitted proofs can be found in a longer version of the paper at `http://www.cs.toronto.edu/~abari/papers/RankingNetworks.pdf`

Ranking networks exhibit these properties: the diameter $D(\boldsymbol{\theta})$ and the average shortest path $\langle l(\boldsymbol{\theta})\rangle$ can be approximated and (often) bounded by:

$$D(\boldsymbol{\theta}) \leq \left\lceil \frac{\log(n)}{\log(n-1) + \log c(d_M(m)|\boldsymbol{\lambda})} \right\rceil; \quad \langle l(\boldsymbol{\theta})\rangle \leq \frac{\log(n)}{\log(n-1) + \log c(d_M(m)|\boldsymbol{\lambda})}. \tag{8}$$

These approximations are detailed in the appendix of a longer version of this paper; and our empirical results confirm these properties and the tightness of these bounds. Ranking networks also possess the diameter-shrinking property of real-world networks [21], where diameter (and average shortest path) shrinks as the network grows:

Theorem 3.2 *Fix m and assume $\rho(r|\boldsymbol{\eta})$ distributes probability mass on more than one ranking. The asymptotic diameter of ranking networks is 2 (as $n \to \infty$).*

These and other structural properties of ranking networks can be computed readily if model parameters are given. However, due to the discrete nature of ranking space, computation can be extremely intensive, even for relatively small m, due to the combinatorial size of ranking space. This motivates the development of easy-to-compute approximations for the special class of *distance-based ranking models*, while studying properties that emerge among networks in this class. Our approximation methods also can applied for efficient learning of model parameters.

3.3 Distance-Based Ranking Models

Distance-based ranking distributions [12, 28] have ranking probabilities that decrease with increasing distance from some modal or reference ranking $\sigma \in \Omega(\mathcal{A})$:

$$\rho(r|\sigma,\omega) = \frac{1}{\psi(\omega)} \exp(-\omega d(r,\sigma)), \tag{9}$$

where $\omega \in [0,\infty)$ is a *dispersion parameter* and $\psi(\omega)$ is a normalizing constant. As $\omega \to \infty$, ρ becomes concentrated at the reference ranking σ, whereas for $\omega = 0$, ρ is the uniform distribution. The Mallows ϕ-model above is an example of such a model (with dispersion $\phi = e^{-\omega}$ and distance d_τ). While we focus on unimodal models, mixtures of such models offer additional modeling flexibility [25].

In this section, we assume that ρ (our distance-based ranking distribution) and c (our distance-based connection function) use the same distance metric. For instance, when using d_τ as our distance measure, the ranking distribution is the Mallows ϕ-model and the probability of connection between two nodes is determined based on this same distance.

We first observe that, as $\omega \to \infty$, ranking networks converge on the well-studied ER random graph model $G(n,p)$ with $p = c(0|\boldsymbol{\lambda})$. We can also bound the probability that a node with a ranking r is connected to a randomly chosen node:

Theorem 3.3 *Given reference ranking σ and a distance-based ranking model, for any fixed $\boldsymbol{\theta}$ and any $r \in \Omega(\mathcal{A})$:*

$$\mathcal{D}(\sigma_M,\boldsymbol{\theta}) \leq \mathcal{D}(r,\boldsymbol{\theta}) \leq \mathcal{D}(\sigma,\boldsymbol{\theta}),$$

where σ_M is some ranking at maximum distance from σ.[2]

This theoretical observation has a natural interpretation: individuals that possess more probable or more "popular" rankings have greater odds of connecting with others (hence have higher expected degree). This is not surprising given the nature of distance-based ranking models. In the context of social networks, this is one possible contributing factor to the intuition that people with more "popular" preferences tend to have more friends and social interactions than those with uncommon preferences, i.e., *preference popularity governs social popularity*. From a different perspective, one can use observed node degree to draw inferences about its ranking: higher degree is predictive of more common preferences. (i.e., close to the modal ranking). This can be exploited to support efficient estimation/learning of the reference ranking, the ranking distribution, and preferences of specific individuals. Using Thm. 3.3, it is straightforward to bound $\mathcal{E}(\boldsymbol{\theta})$.

[2] If more than one ranking has maximum distance, one such ranking minimizes $\mathcal{D}$ (not necessarily all).

Proposition 3.4 *Given a distance-based ranking model, $\mathcal{E}(\boldsymbol{\theta})$ is bounded by*

$$\mathcal{D}(\sigma_M, \boldsymbol{\theta}) \leq \mathcal{E}(\boldsymbol{\theta}) \leq \mathcal{D}(\sigma, \boldsymbol{\theta}),$$

where σ_M is some ranking at maximum distance from σ.

The upper bound $\mathcal{D}(\sigma, \boldsymbol{\theta})$ can be computed efficiently (in $O(m^2)$ time):

$$\mathcal{D}(\sigma, \boldsymbol{\theta}) = \hat{\mathcal{D}}(m, \omega, \boldsymbol{\eta}) = \frac{1}{\psi(\omega)} \sum_{k=0}^{d_M(m)} n_k e^{-\omega k} c(k|\boldsymbol{\lambda}), \tag{10}$$

where $d_M(m)$ is the maximum possible distance for given d when there are m alternatives (e.g., for Kendall-τ, $d_M(m) = \binom{m}{2}$) and n_i is the number of rankings of distance i from an arbitrary fixed ranking. We can efficiently compute n_i for any d either in closed form or using dynamic programming. For instance, n_i can be computed (once) for Kendall's τ in $O(m^3)$ time via:

$$T(j,k) = \begin{cases} 1, & k = 0 \text{ and } j \geq 1 \\ 0, & k > \binom{m}{2} \\ T(j, k-1) + T(j-1, k), & k \leq j-1 \\ T(j, k-1) + T(j-1, k) - T(j-1, k-j), & k > j-1 \end{cases} \tag{11}$$

If d is symmetric in the sense that $n_i = n_{d_M(m)-i}$ for all $i \leq d_M(m)$, a lower bound $\mathcal{D}(\sigma_M, \boldsymbol{\theta})$ can be computed in $O(m^2)$ time:

$$\mathcal{D}(\sigma_M, \boldsymbol{\theta}) = \check{\mathcal{D}}(m, \omega, \boldsymbol{\lambda}) = \frac{1}{\psi(\omega)} \sum_{k=0}^{d_M(m)} n_k e^{-\omega k} c(d_M(m) - k|\boldsymbol{\lambda}) \tag{12}$$

We have empirically observed that $\mathcal{D}(r, \boldsymbol{\theta})$ (usually) decreases as $d(r, \sigma)$ increases. In other words, the distance of an individual's ranking to the reference ranking is negatively correlated with its degree. Define the linear function

$$\tilde{\mathcal{D}}(d, \boldsymbol{\theta}) = \left(1 - \frac{d}{d_M(m)}\right) \mathcal{D}(\sigma, \boldsymbol{\theta}) + \left(\frac{d}{d_M(m)}\right) \mathcal{D}(\sigma_M, \boldsymbol{\theta}). \tag{13}$$

We can approximate $\mathcal{D}(r, \boldsymbol{\theta})$ by $\tilde{\mathcal{D}}(d(r, \sigma), \boldsymbol{\theta})$, which can be used to effectively approximately other structural network properties. For example, the edge density $\mathcal{E}$ and degree distribution P, respectively, be approximated by:

$$\tilde{\mathcal{E}}(\boldsymbol{\theta}) = \frac{1}{\psi(\omega)} \sum_{k=0}^{d_M(m)} n_k e^{-\omega k} \tilde{\mathcal{D}}(k, \boldsymbol{\theta}), \text{ and} \tag{14}$$

$$\tilde{P}(k, \boldsymbol{\theta}) = \frac{\binom{n-1}{k}}{\psi(\omega)} \sum_{i=0}^{d_M(m)} n_i e^{-\omega i} \tilde{\mathcal{D}}(i, \boldsymbol{\theta})^k \left(1 - \tilde{\mathcal{D}}(i, \boldsymbol{\theta})\right)^{n-1-k}. \tag{15}$$

By pre-computing n_i values, these functions can be computed in $O(m^2)$ time (cf. the $O((m!)^2)$ time required for naive exact computation). For large m, even this might problematic; but if $c(\cdot, \boldsymbol{\lambda})$ is convex, one can (loosely) approximate $\mathcal{D}$ and edge density $\mathcal{E}$ in $O(1)$ time by:

$$\tilde{\hat{\mathcal{D}}}(x, \boldsymbol{\theta}) = \left(1 - \frac{x}{d_M(m)}\right) c\left(\frac{d_M(m) e^{-\omega}}{1 + e^{-\omega}}\right) + \frac{x}{d_M(m)} c\left(\frac{d_M(m)}{1 + e^{-\omega}}\right), \text{ and} \tag{16}$$

$$\tilde{\hat{\mathcal{E}}}(\boldsymbol{\theta}) = \frac{1}{1 + e^{-\omega}} c\left(\frac{d_M(m) e^{-\omega}}{1 + e^{-\omega}}\right) + \frac{e^{-\omega}}{1 + e^{-\omega}} c\left(\frac{d_M(m)}{1 + e^{-\omega}}\right). \tag{17}$$

Our approximations $\tilde{\mathcal{E}}(\boldsymbol{\theta})$ and $\tilde{\mathcal{D}}(d(r, \sigma), \boldsymbol{\theta})$ can also be exploited for efficient model learning. Exact evaluation of the likelihood function of a ranking network when dealing with missing data (i.e., when some node rankings are unobserved) has computation of edge density $\mathcal{E}(\boldsymbol{\theta})$ and average connection probability $\mathcal{D}(r, \boldsymbol{\theta})$ as its main bottlenecks. Of course, one can learn model parameters using methods iterative methods such as EM, but direct evaluation of (an approximation of) the likelihood function using these approximations may be support more efficient learning and inference (preliminary tests seem to support this, but this is the topic of ongoing research).

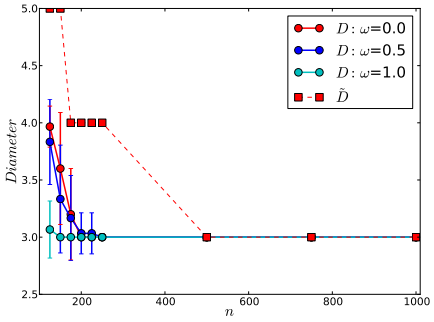

Figure 1: Mean and std. dev. of observed diameters (30 instances) for various ω and n, and its approximation $\tilde{D}$ ($m = 3$, $\alpha = 2$, $\gamma = 0.8$, and $\beta = 2$).

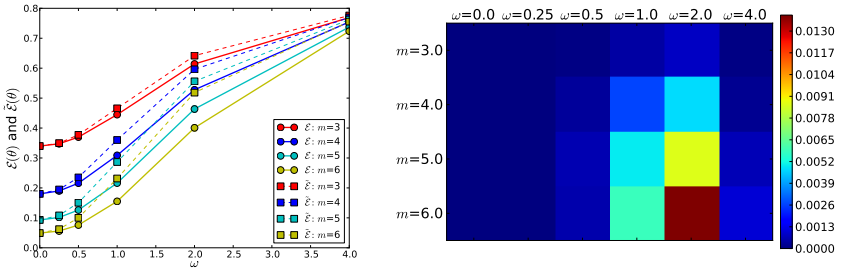

Figure 2: (a) Mean observed edge density $\mathcal{E}(\boldsymbol{\theta})$ (30 instances) and its approximation $\tilde{\mathcal{E}}(\boldsymbol{\theta})$; (b) mean squared error of approximate edge density; varying number of options m, dispersion ω ($n = 1000$, $\alpha = 2$, $\beta = 2$, $\gamma = 0.8$).

4 Numerical Case Study with Mallows Models

We describe experiments on a version of our ranking network model, using the Mallows ϕ-model under various parameter settings. We compare various structural properties (e.g., degree distribution, edge density, diameter) of the resulting networks with the predictions of our approximation methods.

Experimental Setup. We use a variant of the connection probability of [33]:

$$c(d|\boldsymbol{\lambda}) = \gamma \left(1 + \frac{d}{\beta}\right)^{-\alpha}. \tag{18}$$

Here β controls average degree and $\alpha > 1$ determines the extent of homophily. We introduce $\gamma \in (0, 1]$ to control the probability of connecting nodes with the same ranking (to account for the discrete nature of ranking space). Unless noted, all experiments are run on networks with $n = 1000$ nodes, $\alpha = 2$, $\beta = 2$ and $\gamma = 0.8$, while varying m. For each parameter setting, we report results over 30 random ranking networks.

Diameter. We first examine the effect of n and ω on the diameter of the emergent ranking networks, and compare observed diameter to our approximation $\tilde{D}$ (see Eq. 8). We fix $m = 5$. Fig. 1 confirms that diameter shrinks as n increases. Unsurprisingly, diameter decreases with increasing ω (i.e., more uniform distributions give larger diameter), largely due to the increased edge density caused by increasing ω (see below). $\tilde{D}$ provides a reasonable upper bound of diameter for any value of ω, with bounds that are very tight when n reaches 500.

Edge Density. The effect of m and ω on edge density is illustrated in Fig. 2. Fig. 2(a) demonstrates that edge density increases with ω but decreases with m (compare boxes on the solid lines). Error bars are very tight (and barely observable). Our approximation $\tilde{\mathcal{E}}$ (Eq. 14) is relatively close to the observed edge density, especially when ω is relatively small ($\omega \leq 0.5$) or large ($\omega \geq 4$).

To better quantify the accuracy of our approximation, we compute the mean squared error (MSE) between the actual edge density and our approximation over 30 randomly generated ranking networks. Fig. 2(b) shows that MSE is relatively low with a maximum of roughly 0.014 for $m = 6$, $\omega = 2$. MSE first increases and then decreases with ω. Moreover, MSE increases with m especially

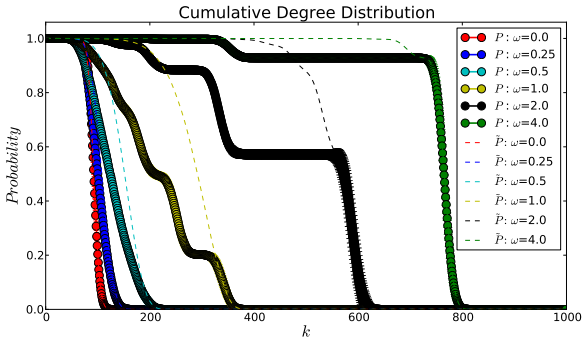

Figure 3: The mean cumulative degree distribution (30 instances) and its approximation, varying ω ($m = 5$, $n = 1000$, $\alpha = 2$, $\beta = 2$, $\gamma = 0.8$).

for $\omega = 1, 2$. This suggests that, for large m and small ω, our edge density approximation should be used cautiously. Refined approximations that account for the value of ω are being explored.

Degree Distribution. Fig. 3 shows the impact of ω on the degree distribution of ranking networks under the Mallows model, as well as the effectiveness of our approximation $\tilde{P}$ (Eq. 15). The figure shows *cumulative degree distribution* to reduce noise in the plots: given degree distribution $P(k, n)$, the cumulative degree distribution is $P_c(k, n) = \sum_{l=k}^{n} P(l, n)$ (i.e., probability that a node n has the degree $d \geq k$). The mean cumulative degree distribution is shown for various values of ω with $m = 5$ and $n = 1000$. The expected degree increases with ω; e.g., with $\omega = 0$, most of nodes have degree around 100, whereas for $\omega = 4$, most nodes have degree around 800 (since mean degrees is $n\mathcal{E}$). However, variance in the degree distribution initially increases as ω increases from 0 but then decreases. For example, variance is low at $\omega = 0$ and $\omega = 4$, but high for $\omega = 1$ and $\omega = 2$. Interestingly, with $\omega = 1$ and $\omega = 2$, the the cumulative degree distribution shows several distinct modes in the degree distribution. The relation between these modes and community structure in ranking networks is something that remains to be explored.

Fig. 3 also shows that the approximate degree distribution $\tilde{P}$ is reasonably close to actual observed degree distribution when $\omega \leq 0.5$ and $\omega \geq 4$, though it fails to account for the distinct modes, especially for $\omega = 1, 2$. Once again, more accurate approximations (sensitive to ω) may be possible.

5 Concluding Remarks and Future Work

We have introduced *ranking networks*, a class of attribute-based (or latent-space) network formation models in which node attributes are rankings over a set of options, and connections are formed between nodes based on the similarity of their underlying rankings. We studied some structural properties (e.g., diameter, connectivity, edge density) of these networks, and showed that our model possesses some characteristics of real-world networks (e.g., shrinking diameter).

This model is a starting point for the broader investigation of the impact of rankings, and preferences in particular, on network formation. Future directions include: the analysis of more realistic ranking distributions (e.g., mixture models); extensions of our model that account for heterogeneity/heterophily; and modeling the dynamics and mutability of the underlying rankings themselves in response to network connections. Of practical importance is studying the extent to which rankings, such as preferences, play a role in shaping connections in real-world networks, and how the induced correlations can be best exploited in applications such as recommender systems, advertising, social choice and voting, web search and information retrieval.

References

[1] Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1):47, 2002.

[2] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

[3] Alain Barrat, Marc Barthelemy, and Alessandro Vespignani. *Dynamical Processes on Complex Networks*. Cambridge University Press, 2008.

[4] Marc Barthélemy. Spatial networks. *Physics Reports*, 499(1):1–101, 2011.

[5] Mustafa Bilgic, Lilyana Mihalkova, and Lise Getoor. Active learning for networked data. In *Proceedings of the 27th International Conference on Machine Learning*, pages 79–86, 2010.

[6] Marián Boguñá and Romualdo Pastor-Satorras. Class of correlated random networks with hidden variables. *Physical Review E*, 68:036112, Sep 2003.

[7] Marián Boguñá, Romualdo Pastor-Satorras, Albert Díaz-Guilera, and Alex Arenas. Models of social networks based on social distance attachment. *Physical Review E*, 70:056122, Nov 2004.

[8] Béla Bollobás. *Random graphs*. Cambridge University Press, 2001.

[9] Persi Diaconis. *Group Representations in Probability and Statistics*. Institute of Mathematical Statistics, Hayward, California, 1988.

[10] David A. Easley and Jon M. Kleinberg. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, Cambridge, 2010.

[11] Paul Erdős and Alfréd Rényi. On random graphs. *Publicationes Mathematicae Debrecen*, 6:290–297, 1959.

[12] Michael A Fligner and Joseph S Verducci. Distance based ranking models. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 359–369, 1986.

[13] Lise Getoor and Christopher P. Diehl. Link mining: A survey. *ACM SIGKDD Explorations Newsletter*, 7(2):3–12, December 2005.

[14] Peter D. Hoff, Adrian E. Raftery, and Mark S. Handcock. Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97(460):1090–1098, 2002.

[15] Matthew O. Jackson. *Social and Economic Networks*. Princeton University Press, Princeton, 2008.

[16] Amin Karbasi, Stratis Ioannidis, and Laurent Massoulié. Content search through comparisons. In *Automata, Languages and Programming*, pages 601–612. Springer, 2011.

[17] Myunghwan Kim and Jure Leskovec. Modeling social networks with node attributes using the multiplicative attribute graph model. In *UAI'2011*, pages 400–409.

[18] Myunghwan Kim and Jure Leskovec. Latent multi-group membership graph model. In *Proceedings of the 29th International Conference on Machine Learning*, 2012.

[19] Myunghwan Kim and Jure Leskovec. Multiplicative attribute graph model of real-world networks. *Internet Mathematics*, 8(1-2):113–160, 2012.

[20] Jure Leskovec, Deepayan Chakrabarti, Jon Kleinberg, Christos Faloutsos, and Zoubin Ghahramani. Kronecker graphs: An approach to modeling networks. *JMLR*, 11:985–1042, 2010.

[21] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. Graphs over time: Densification laws, shrinking diameters and possible explanations. In *KDD'05*, pages 177–187, 2005.

[22] David Liben-Nowell and Jon Kleinberg. The link-prediction problem for social networks. *Journal of the American Society for Information Science and Technology*, 58(7):1019–1031, 2007.

[23] Tie-Yan Liu. Learning to rank for information retrieval. *Foundations and Trends in Information Retrieval*, 3(3):225–331, 2009.

[24] Linyuan Lü and Tao Zhou. Link prediction in complex networks: A survey. *Physica A: Statistical Mechanics and its Applications*, 390(6):1150–1170, 2011.

[25] Tyler Lu and Craig Boutilier. Learning mallows models with pairwise preferences. In *Proceedings of the 28th International Conference on Machine Learning*, pages 145–152, 2011.

[26] Sofus A. Macskassy. Using graph-based metrics with empirical risk minimization to speed up active learning on networked data. In *KDD'09*, pages 597–606.

[27] Colin L. Mallows. Non-null ranking models. In *Biometrika*, volume 44, pages 114–130, 1957.

[28] John I. Marden. *Analyzing and Modeling Rank Data*. Chapman and Hall, London, 1995.

[29] Mark Newman. *Networks: An Introduction*. Oxford University Press, 2010.

[30] Mark Newman, Albert-László Barabási, and Duncan J. Watts. *The Structure and Dynamics of Networks*. Princeton University Press, 2006.

[31] Mathew Penrose. *Random Geometric Graphs*. Oxford University Press, USA, 2003.

[32] Prithviraj Sen, Galileo Mark Namata, Mustafa Bilgic, Lise Getoor, Brian Gallagher, and Tina Eliassi-Rad. Collective classification in network data. *AI Magazine*, 29(3):93–106, 2008.

[33] M. Ángeles Serrano, Dmitri Krioukov, and Marián Boguñá. Self-similarity of complex networks and hidden metric spaces. *Physical Review Letters*, 100:078701, Feb 2008.

[34] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of small-world networks. *Nature*, 393(6684):440–442, 1998.

[35] Bernard M. Waxman. Routing of multipoint connections. *Selected Areas in Communications, IEEE Journal on*, 6(9):1617–1622, December 1988.

[36] Stephen J. Young and Edward R. Scheinerman. Random dot product graph models for social networks. In *WAW'07*, pages 138–149, 2007.