

Block-Approximated Exponential Random Graphs

Florian Adriaens*

KTH Royal Institute of Technology, Sweden
first.last@gmail.com

Alexandru Mara Jeffrey Lijffijt Tijn De Bie

IDLab, Ghent University, Belgium
first.last@ugent.be

Abstract—An important challenge in the field of exponential random graphs (ERGs) is the fitting of non-trivial ERGs on large graphs. By utilizing fast matrix block-approximation techniques, we propose an approximative framework to such non-trivial ERGs that result in dyadic independence (i.e., edge independent) distributions, while being able to meaningfully model *local* information of the graph (e.g., degrees) as well as *global* information (e.g., clustering coefficient, assortativity, etc.) if desired. This allows one to efficiently generate random networks with similar properties as an observed network, and the models can be used for several downstream tasks such as link prediction. Our methods are scalable to sparse graphs consisting of millions of nodes.

Empirical evaluation demonstrates competitiveness in terms of both speed and accuracy with state-of-the-art methods—which are typically based on embedding the graph into some low-dimensional space—for link prediction, showcasing the potential of a more direct and interpretable probabilistic model for this task.

Index Terms—maximum entropy models, exponential random graphs, network modeling, link prediction

I. INTRODUCTION

Network modeling is typically concerned with the following setting: given measurements on certain structural properties of a real-life network, such as node degrees, clustering coefficient, density, and so forth, one wishes to find a model for the network where samples generated by the model have similar values of the measured properties. Prominent examples include stochastic block models [1], graphons [2] and exponential random graphs (ERGs) [3].

ERGs have received significant attention in numerous research areas [3]–[6], as they provide a well-founded probabilistic framework to network modeling; their study has been described as the “statistical mechanics of networks” [7]. They arise naturally as the solution to the problem of finding the maximum entropy distribution over a set of graphs, while respecting constraints that a probabilistic statistic must equal an observed statistic in expectation. Equivalently, the parameters of an ERG can be determined by maximum likelihood estimation of an exponential distribution for explaining the observed statistics in a principled manner. The advantage of ERGs is that they can represent a wide range of structural tendencies, such as transitivity and degree heterogeneity, by capturing complicated dependence patterns that are not easily modeled by simpler probabilistic models.

The downside of classically specified ERGs is their limited scalability [8]. Inferring the parameters of an ERG is often

intractable for networks of even moderate size, because computing the partition function may require evaluating a summation over all $2^{\Theta(n^2)}$ for graphs with n nodes. The common approach is to approximately infer the parameters by MCMC sampling, after which a goodness of fit is measured by generating random networks from the ERG and comparing statistics with the observed network. Besides limited scalability, MCMC methods often have the problem of degeneracy [9], assigning most probability mass to either near-empty or near-full graphs, for example when one is counting Markov neighborhood properties such as triangles [5].

ERGs typically lose efficiency when they aim to model global constraints that suggest dependencies between edges, such as degree assortativity or the existence of many triangles. Modeling such information is important in practice: e.g. if the number of triangles is large as compared to other graphs with the same local properties, this may indicate that edge formation is a result of a triadic closure process, which can be an important property to understand and model.

Addressing this challenge, in this paper we propose the use of fast matrix block-approximation techniques to derive intelligible and scalable approximations to non-trivial ERGs for large sparse graphs ($> 10^6$ nodes). The resulting models are dyadically independent, while still meaningfully incorporating local information (degrees) and global information (triangles, assortativity, etc.).

Contributions and roadmap.

- Discussion of a general dyadic independence model that is able to incorporate arbitrary (structural) features between pairs of nodes (Section II).
- Depending on the incorporated features, fitting the model parameters may not be scalable. We propose fast matrix block-approximation techniques to provide scalable approximations to the general model (Section III).
- An empirical comparison with several state-of-the-art (embedding-based) methods for the task of predicting missing links, and a goodness-of-fit comparison with two well-known probabilistic models (Section IV).

II. MAXENT MODELS WITH STRUCTURAL CONSTRAINTS

Notation. Let $G = (V, E)$ be a graph with $|V| = n$ nodes and $|E| = m$ edges. The neighborhood of node i is denoted as $\mathcal{N}(i)$. The adjacency matrix of an observed graph G is denoted as $\hat{\mathbf{A}} = [\hat{\mathbf{A}}_{ij}] \in \mathcal{A}$, where $\mathcal{A} = \{0, 1\}^{n \times n}$ is the set of square matrices of size n . A random matrix is denoted as $\mathbf{A} \in \mathcal{A}$. The

*This work was done while the author was with Ghent University.

transpose of a matrix \mathbf{A} is denoted as \mathbf{A}^T and the Frobenius norm is denoted as $\|\mathbf{A}\|$. The expected value operator related to a distribution P is denoted as $\mathbb{E}_P[\cdot]$. The focus in this paper is restricted to undirected graphs, but extensions to directed networks are straightforward.

ERGs as maximum entropy (MaxEnt) models. The setting of this paper is that, based on some measurements (statistics) of an observed graph $\hat{\mathbf{A}}$, we wish to find a probability distribution over \mathcal{A} , i.e. the set of all possible graphs over the same set of vertices. ERGs are a family of statistical models that are often used for this task. Let $s(\cdot)$ be the vector of statistics of a graph, e.g. one component of $s(\cdot)$ could be the count of all triangles. An ERG with associated parameter vector θ is then typically introduced by positing an exponential distribution over \mathcal{A} :

$$P_\theta(\mathbf{A} = A) = \exp(\theta^T s(A) - \psi), \quad (1)$$

where ψ denotes a normalizing constant. Given an observed graph $\hat{\mathbf{A}}$, the parameters θ are then typically derived by maximum likelihood estimation (MLE).

An equivalent, but less-known, way of introducing ERGs is by looking for the maximum entropy distribution over \mathcal{A} , subject to constraints that the expected statistics are equal to the observed statistics. The MLE of (1) is equal to this maximum entropy distribution [10]. Throughout the paper, we will take the maximum entropy point-of-view.

A. An example with degree assortativity

We start by formalizing what we *wish* to solve, and end up relaxing these equations in order to have analytical expressions containing parameters that are tractable to infer. As an example, suppose we are interested in finding the maximum entropy (MaxEnt) distribution over \mathcal{A} , subject to a constraint in expectation on each individual node degree, as well as a constraint in expectation on the total *degree assortativity* in the graph, as measured by the sum of $|\mathcal{N}(i)| \cdot |\mathcal{N}(j)|$ over all edges $(i, j) \in E$. We aim to solve¹

$$\begin{aligned} & \arg \max_{P(\mathbf{A})} -\mathbb{E}_P[\log P(\mathbf{A})], \\ \text{s.t. } & \mathbb{E}_P[\sum_{i,j} \mathbf{A}_{ij} \sum_c \mathbf{A}_{ic} \sum_r \mathbf{A}_{rj}] = c, \\ & \mathbb{E}_P[\sum_j \mathbf{A}_{ij}] = d_i \quad \forall i, \\ & \mathbb{E}_P[\sum_i \mathbf{A}_{ij}] = d_j \quad \forall j, \end{aligned} \quad (2)$$

where $d_i = \sum_j \hat{\mathbf{A}}_{ij}$ is the observed degree of node i in G , and $c = \sum_{i,j} \hat{\mathbf{A}}_{ij} \sum_c \hat{\mathbf{A}}_{ic} \sum_r \hat{\mathbf{A}}_{rj}$ is the observed assortativity measure. A typical difficulty with this classically specified ERG is that in general the normalizing constant of the distribution P is infeasible to compute, because of edge dependencies introduced by the assortativity constraint. As such, the parameters of the distribution are intractable to compute exactly. This is in contrast to a MaxEnt model subject to only degree constraints. As observed by different authors [7], [11], [12], in this case P is a dyadic independence model: it

¹We implicitly assume $\sum P(\mathbf{A}) = 1$ and $P(\mathbf{A}) \geq 0$ in all the MaxEnt problem formulations, but omit them for brevity.

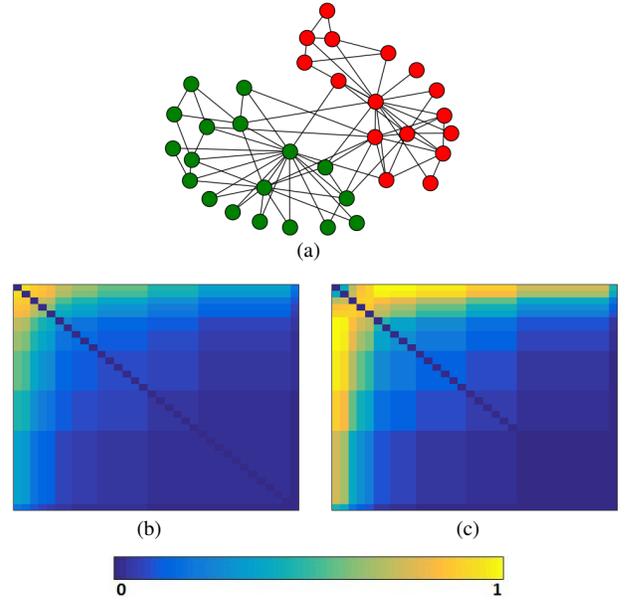


Fig. 1. (a) The *Karate* network (b) MaxEnt edge success probabilities with constraints on degrees (c) MaxEnt edge probabilities with constraints on degrees, and a global constraint on degree assortativity as given by (3). In both heatmaps, nodes are sorted in descending degree order (highest degree in the top left corner).

factorizes as a product of independent Bernoulli distributions, one for each node pair $(i, j) \in V \times V$. Moreover, in a degree-only model, the number of unique parameters to be optimized over is fully determined by the number of *unique* degrees in G . As shown by [11], for sparse G (where $m = O(n)$), the problem has in fact only $O(\sqrt{n})$ free variables instead of the $O(n)$ original variables (one original variable for each degree constraint), making inference possible on very large networks. In Section III, we show that we can do a similar reduction in variables also for more complex models.

One way to approximate (2) is to replace the row- and column sums of random variables $\sum_c \mathbf{A}_{ic}$ by their expectation:

$$\sum_c \mathbf{A}_{ic} \approx \mathbb{E}_P[\sum_c \mathbf{A}_{ic}] = d_i.$$

The approximated model of (2) then becomes

$$\begin{aligned} & \arg \max_{P(\mathbf{A})} -E_P[\log P(\mathbf{A})], \\ \text{s.t. } & \mathbb{E}_P[\sum_{i,j} d_i d_j \mathbf{A}_{ij}] = c, \\ & \mathbb{E}_P[\sum_j \mathbf{A}_{ij}] = d_i \quad \forall i, \\ & \mathbb{E}_P[\sum_i \mathbf{A}_{ij}] = d_j \quad \forall j. \end{aligned} \quad (3)$$

We omit the details, but it is easy to show that the solution to (3) is again a dyadic independence model, and the complexity depends on the number of unique degrees in G .

Let us provide some intuition on the solution of (3). First, because the distribution has maximum entropy, it is the unique distribution that injects no side information on properties that were not taken into account as constraints [13]. Secondly, if the degree assortativity cannot be explained by a model that

is inferred using *only* degree constraints, for example if c is larger than expected under a model where only degrees are constrained, then the optimum of (3) will on average assign higher probabilities between pairs of high degree nodes and between pairs of low degree nodes, at the expense of pairs of nodes where one has a large and the other a low degree. This results in a more accurate fit of the observed graph.

An example of a highly disassortative network is the *Karate* dataset [14]. It has a negative assortativity coefficient [15] of -0.48 , and thus nodes with similar degrees are less often connected. This is confirmed in a visualization of the network in Figure 1a. The dataset essentially consists of two Karate club teachers (the two highest-degree nodes), mostly connected to their own students, with few edges between the two communities. Most connections are between a high-degree node and a low-degree node, and the two teachers themselves are not connected.

Figure 1b shows the edge probabilities between all pairs of nodes, when the network is modeled by a MaxEnt model with a constraint on the expectation of each individual node degree (the degree-only model). Connections between high degree nodes are more likely, and thus it assigns most probability mass in the top left corner. On the other hand, Figure 1c shows the MaxEnt model with an additional assortativity constraint as given by (3). It builds on the degree-only model by taking into account the network’s disassortativity, and as such it lowers the edge probabilities between nodes of similar degree and increases the edge probabilities between nodes with dissimilar degrees. Note that it correctly assigns a low probability for a connection between the two Karate teachers.

B. Generalizing to arbitrary features

Building on the previous example, we can view (3) as a model that takes into account *observed* features $f_{ij} \triangleq d_i d_j$. Instead of taking the product of two node degrees, we can simply extend this to arbitrary observed features.² Denote $\mathbf{F} = [f_{ij}] \in \mathbb{R}^{n \times n}$ as an induced pairwise feature matrix. For example, the degree matrix of a node i is a matrix with ones on the i -th row and zeros elsewhere. The common neighbor matrix (triangle counting) is given by $\mathbf{F}_{\text{CN}} \triangleq \hat{\mathbf{A}}^2 = [|\mathcal{N}(i) \cap \mathcal{N}(j)|]$. Counts of different types of graphlets besides triangles can be incorporated as well, e.g., by using so-called ‘weighted motif graphs’ [16], [17] as feature matrices.

Given M of these feature matrices $\mathbf{F}_l \triangleq [f_{ij}^l]$ with $l = 1, \dots, M$, we aim to solve

$$\begin{aligned} \arg \max_{P(\mathbf{A})} & -\mathbb{E}_P[\log P(\mathbf{A})], \\ \text{s.t.} & \mathbb{E}_P[\sum_{i,j} f_{ij}^l \mathbf{A}_{ij}] = c_l \quad \forall l = 1, \dots, M. \end{aligned} \quad (4)$$

Where $c_l = \sum_{i,j} f_{ij}^l \hat{\mathbf{A}}_{ij} = \sum_{(i,j) \in E} f_{ij}^l$ are the observed statistics. Solutions to the class of models defined by (4) have the convenient property that the partition function (the normalizing constant) factorizes as a product over all possible

²Such a generalization is often less well-founded than the case of degree assortativity (where we replaced a sum of random variables with the sum of their expectations), but these models still turn out to be useful.

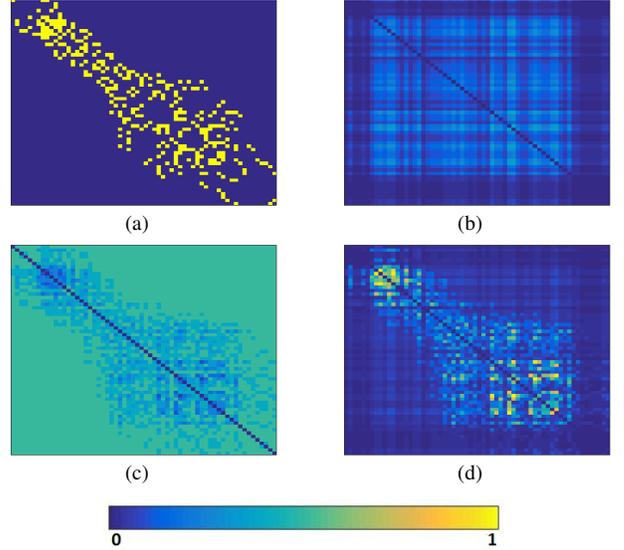


Fig. 2. MaxEnt edge success probabilities on the *Dolphins* network (a), when different kinds of constraints are taken into account; (b) MaxEnt with degrees (c) MaxEnt with only a global constraint on \mathbf{F}_{CN} (d) MaxEnt with a combined constraint on both the degrees and \mathbf{F}_{CN} , demonstrating that combining local and global information results in a superior and more realistic model.

edges. As such (4) is again a dyadic independence model, with a Bernoulli success probability for a node pair (i, j) :

$$P(\mathbf{A}_{ij} = 1) = \frac{\exp(\sum_{l=1}^M f_{ij}^l \lambda_l)}{1 + \exp(\sum_{l=1}^M f_{ij}^l \lambda_l)},$$

where λ_l denotes the Lagrange multiplier associated with the l -th constraint in (4). These multipliers are found by unconstrained minimization of the convex Lagrange dual function:

$$\begin{aligned} L(\lambda_1, \dots, \lambda_M) = & \\ & \sum_{i,j} \log(1 + \exp(\sum_{l=1}^M f_{ij}^l \lambda_l)) - \sum_{l=1}^M c_l \lambda_l. \end{aligned} \quad (5)$$

The partial derivatives $\forall l = 1, \dots, M$ are computed as

$$\frac{\partial L}{\partial \lambda_l} = \sum_{i,j} f_{ij}^l \cdot \frac{\exp(\sum_{l=1}^M f_{ij}^l \lambda_l)}{1 + \exp(\sum_{l=1}^M f_{ij}^l \lambda_l)} - c_l.$$

Instead of counting exact specifications in the original ERG, the models (4) can be seen as applying a ‘linear mask’ in order to approximately count the specifications. In fact, they are often identical to expressions found by pseudo likelihood estimation [18], [19]. Pseudo likelihood is mainly used for estimation of the original ERG parameters [19], but we argue that dyadic independence models can be valuable by themselves. For the exact derivations and a discussion regarding the pseudo likelihood, we refer to Appendix A.

In this paper, we focus on combining individual local constraints (degrees) with a limited number of global structural constraints.³ Figure 2 shows an example of the predictive power

³Two or three global constraints often suffice to obtain qualitative models.

of combining constraints on degrees with a global structural constraint \mathbf{F}_{CN} (triangle count) on the *Dolphins* dataset [20]. Figure 2a shows the connectivity of the dataset. Figure 2b shows the edge probabilities according to a degree-only MaxEnt model. It assigns higher probability to edges between two high degree nodes, but it fails to capture any form of local community structure. Figure 2c is a MaxEnt model fitted with just the structural constraint \mathbf{F}_{CN} . Although it captures community structure, it fails to make good predictions about the actual edges in the network. Indeed, most observed edges have a low probability of being present, and all node pairs with zero common neighbors are assigned a probability of 1/2 of being connected. Figure 2d shows the model when combining degree constraints and the structural constraint \mathbf{F}_{CN} . It leads to a remarkably good fit of the original network, while still leaving room for prediction and inference.

C. Scalability issues

Inference on large graphs is typically not possible when we have $M = O(n)$ constraints, as is the case when combining degree constraints with a limited number of global constraints. Minimizing (5) can be viewed as a classical learning setting with n^2 training examples and $O(n)$ weights. Standard gradient methods need $\Omega(n^2)$ computations per iteration and \mathbf{F} needs to be stored in memory. To resolve both issues, in Section III we propose to *block-approximate* the feature matrices. In particular, Theorem 1 shows that in sparse graphs ($m = O(n)$), the number of variables to be optimized over roughly reduces from $O(n)$ to $O(\sqrt{n})$. At the same time, block-approximated matrices are easily maintained in memory.

D. Label leakage & degeneracy

To avoid overfitting, it is crucial to avoid ‘label leakage’. More concretely, when predicting whether or not an edge exists, one should avoid making direct use of the actual existence of that edge while fitting the model. Indeed, if one selects $\mathbf{F} = \hat{\mathbf{A}}$, the only solution to (4) is exactly $P(\hat{\mathbf{A}}) = 1$. In other words, the model of the network is the network itself, rendering it completely useless for tasks like link prediction. In contrast, \mathbf{F}_{CN} is an excellent candidate for a link prediction feature matrix, since the number of common neighbors between two nodes is not directly related to the existence of that edge, and if it is the case, the model will actually learn the relation. Remarkably, other methods often overlook this fact. For example, the method by [21] considers embeddings defined by a truncated singular value decomposition of the adjacency matrix $\hat{\mathbf{A}}$. The Katz centrality measure [22] is another such example. Nevertheless, for other tasks that require a closer fit to the observed network (e.g., network reconstruction), it might be useful to include adjacency information, since it allows to become arbitrarily close to the observed network. For example, $\mathbf{F} = \hat{\mathbf{A}} + \beta \hat{\mathbf{A}}^2$ with small $\beta > 0$ was used in Section IV-B to obtain a close fit to the *Facebook* network.

The other side of degeneracy, i.e. assigning most probability mass to near-empty graphs, is seemingly avoided by incorpo-

rating degree constraints into the model. For example, compare Figure 2c with Figure 2d.

III. BLOCK-APPROXIMATING FEATURE MATRICES

A. Motivation

Prior work [23], [24] on improving the scalability of fitting a general MaxEnt model (4) looks for permutations of $(\lambda_1, \dots, \lambda_M)$ that leave the Lagrange dual function L (5) invariant. The convexity of L then implies that if there is a permutation that maps λ_i to λ_j , then there necessarily exists an optimum of (5) where $\lambda_i = \lambda_j$ [23, Section 4.2]. Similarly as in [23], we look for equivalent Lagrange multipliers associated with the degree constraints. Equivalent Lagrange multipliers are equated and the *reduced model* is solved by standard convex optimization methods.

However, for general feature matrices equivalences are rare. Thus, in this paper we propose to block-approximate the feature matrices. Let $\bar{\mathbf{F}} \in \mathbb{R}^{n \times n}$ be a block-approximation of \mathbf{F} , represented by a structure with $k \times k$ blocks, with each block being a submatrix with constant values. Theorem 1 states that, when fitting a MaxEnt model (4) with constraints on each node degree, as well as a global constraint as given by $\bar{\mathbf{F}}$, the number of free variables in the reduced model is at most $\sqrt{2km} + 1$ instead of the original $n + 1$ variables (n degrees and one global constraint). This implies a significant speed-up when fitting such MaxEnt models on large sparse graphs.

Lemma 1. *Let $\hat{\mathbf{A}}$ be symmetric with m non-zero entries. If the nodes are partitioned into k disjoint groups, then the sum over all groups of the number of unique degrees inside each group is at most $\sqrt{2km}$.*

Proof. The nodes are partitioned into k disjoint groups. Denote m_i as the number of ones in the rows of $\hat{\mathbf{A}}$, when the rows are restricted to the nodes in group i ($1 \leq i \leq k$). Let u_i denote the number of unique degrees in group i . Then it holds that [11, Lemma 2]:

$$u_i \leq \sqrt{2 \cdot m_i}$$

Summing over the k groups:

$$\sum_{i=1}^k u_i \leq \sqrt{2} \cdot (\sqrt{m_1} + \dots + \sqrt{m_k}).$$

Utilizing $m_1 + \dots + m_k = m$, the concavity of $\sqrt{\cdot}$ and Jensen’s inequality:

$$\sqrt{m_1} + \dots + \sqrt{m_k} \leq k \cdot \sqrt{\frac{m}{k}}.$$

Hence $\sum_{i=1}^k u_i \leq \sqrt{2km}$ for sparse matrices. \square

Theorem 1. *Let $\hat{\mathbf{A}}$ be symmetric with m non-zero entries. Let $\bar{\mathbf{F}}$ be a $k \times k$ blockmatrix, with each block being a submatrix with constant entries. A MaxEnt model (4), with constraints on each individual degree and a constraint as given by $\bar{\mathbf{F}}$, can be solved by optimizing an unconstrained convex problem with at most $\sqrt{2km} + 1$ variables.*

Proof. Let λ_i be the Lagrange multiplier associated with the degree constraint of node i . Observe that the function L (5) is invariant if one swaps λ_i with λ_j , when i and j are part of the same block *and* have the same degree $d_i = d_j$. Convexity of L then implies there exists an optimum where $\lambda_i = \lambda_j$ [23, Section 4.2]. As such, the number of free variables is equal to the number of unique degrees in each block (plus one variable for the global constraint). By Lemma 1, the total number of unique degrees summated over k blocks is bounded by $\sqrt{2km}$. \square

B. Methods

In this section, we discuss several methods for obtaining fast and qualitative block-approximations for certain classes of feature matrices. Most methods rely on a fast top d eigendecomposition of the—assumed to be sparse—adjacency matrix $\hat{\mathbf{A}}$.

1) *Spectral clustering adjacency polynomials:* A principled way of block-approximating a feature matrix \mathbf{F} is to first spectral cluster \mathbf{F} and then replacing rows with centroids. Assuming \mathbf{F} is real and symmetric, spectral clustering first calculates a truncated eigenvalue decomposition to get an optimal low-rank approximation of \mathbf{F} , after which k-means is used to cluster the nodes in the low dimensional space [25]. A block-approximation of \mathbf{F} is then found by

$$\mathbf{F} \approx \mathbf{U}\mathbf{S}\mathbf{U}^T \approx \mathbf{U}_c\mathbf{S}\mathbf{U}_c^T, \quad (6)$$

with $\mathbf{U}, \mathbf{U}_c \in \mathbb{R}^{n \times d}$ and $\mathbf{S} \in \mathbb{R}^{d \times d}$ a diagonal matrix containing the signs of the d largest (in absolute value) eigenvalues of \mathbf{F} . The matrix $\mathbf{U} = \mathbf{V}\sqrt{\Sigma}$ consists of the rescaled eigenvectors \mathbf{V} corresponding to the top d eigenvalues, where Σ denotes a diagonal matrix with the absolute values of these eigenvalues. The matrix \mathbf{U}_c is defined by replacing each row of \mathbf{U} by its respective cluster centroid. Notice that symmetry of \mathbf{F} is maintained by both of the approximations in (6). Now let $\bar{\mathbf{F}} \triangleq \mathbf{U}_c\mathbf{S}\mathbf{U}_c^T$. The following proposition gives a bound on the expected distance between $\bar{\mathbf{F}}$ and \mathbf{F} .⁴

Proposition 1. *Let $|\lambda_1| \geq |\lambda_2| \geq \dots \geq |\lambda_d|$ be the d largest eigenvalues of \mathbf{F} . Let $\phi_{OPT,k}$ denote the optimal clustering objective value with k clusters. Using k-means++ as a (randomized) clustering algorithm, the expected error $\mathbb{E}[\|\mathbf{F} - \bar{\mathbf{F}}\|]$ is at most*

$$2\sqrt{\sum_{i=1}^d |\lambda_i|} \cdot O(\log(k)) \cdot \phi_{OPT,k}$$

additively larger than any optimal rank d approximation of \mathbf{F} .

Proof. Invoking the triangle inequality, and since $\mathbf{U}\mathbf{S}\mathbf{U}^T$ is the optimal rank d approximation of \mathbf{F} , we can write

$$\begin{aligned} \|\mathbf{F} - \mathbf{U}\mathbf{S}\mathbf{U}^T\| &\leq \|\mathbf{F} - \mathbf{U}_c\mathbf{S}\mathbf{U}_c^T\| \\ &\leq \|\mathbf{F} - \mathbf{U}\mathbf{S}\mathbf{U}^T\| + \|\mathbf{U}\mathbf{S}\mathbf{U}^T - \mathbf{U}_c\mathbf{S}\mathbf{U}_c^T\|. \end{aligned}$$

⁴This should be a standard result in spectral clustering theory, but the proof is given for completeness.

The latter term can again be bounded with the triangle inequality

$$\begin{aligned} &\|\mathbf{U}\mathbf{S}\mathbf{U}^T - \mathbf{U}_c\mathbf{S}\mathbf{U}_c^T\| \\ &\leq \|\mathbf{U}\mathbf{S}\mathbf{U}^T - \mathbf{U}_c\mathbf{S}\mathbf{U}^T\| + \|\mathbf{U}_c\mathbf{S}\mathbf{U}^T - \mathbf{U}_c\mathbf{S}\mathbf{U}_c^T\| \\ &= \|(\mathbf{U} - \mathbf{U}_c)\mathbf{S}\mathbf{U}^T\| + \|\mathbf{U}_c\mathbf{S}(\mathbf{U}^T - \mathbf{U}_c^T)\|. \end{aligned} \quad (7)$$

Observe that we can write $\mathbf{U}_c = \mathbf{C}\mathbf{U}$, where $\mathbf{C} \in \mathbb{R}^{n \times n}$ denotes a matrix with in each row entries equal to $1/n_i$ for nodes that are in the same bin, and zero otherwise. Here, n_i denotes the size of a cluster i , i.e. \mathbf{C} is a matrix that replaces each row in \mathbf{U} by its cluster centroid. It's easy to see that \mathbf{C} is an orthogonal projection matrix, in the sense that $\mathbf{C}^2 = \mathbf{C}$ and $\mathbf{C} = \mathbf{C}^T$. As such, for any matrix \mathbf{B} it holds that $\|\mathbf{C}\mathbf{B}\| \leq \|\mathbf{B}\|$. Applying to the second term in (7) yields

$$\begin{aligned} &\|(\mathbf{U} - \mathbf{U}_c)\mathbf{S}\mathbf{U}^T\| + \|\mathbf{U}_c\mathbf{S}(\mathbf{U}^T - \mathbf{U}_c^T)\| \\ &\leq \|(\mathbf{U} - \mathbf{U}_c)\mathbf{S}\mathbf{U}^T\| + \|\mathbf{U}\mathbf{S}(\mathbf{U}^T - \mathbf{U}_c^T)\| \\ &\leq 2\sqrt{\sum_{i=1}^d |\lambda_i|} \cdot \|\mathbf{U} - \mathbf{U}_c\|. \end{aligned}$$

The result follows by noting that $\|\mathbf{U} - \mathbf{U}_c\|$ is the k-means objective function and applying the bound in expectation for the kmeans++ algorithm [26]. \square

The bound from Proposition 1 gives insight in how the dimension d and number of bins k affect the block-approximation (6). Increasing k will benefit the approximation since $\bar{\mathbf{F}}$ becomes closer to the optimal rank d approximation. However, for fixed k , increasing d does *not* always benefit the block-approximation. Indeed, as the clustering approximation gets worse with increasing dimension d [27], the effect on the overall block-approximation could potentially be detrimental, which is confirmed in practice. As a practical guideline, we advise to keep d small, and selecting a high k while maintaining tractable computational complexity (Section III-C).

Since \mathbf{F} is often dense, directly calculating \mathbf{F} and performing a top d eigendecomposition is not scalable both memory and timewise. Instead, for *higher-order proximity* matrices, i.e. matrices that are expressed as polynomials of $\hat{\mathbf{A}}$ with positive coefficients, one can directly use a (fast) eigendecomposition of $\hat{\mathbf{A}}$ (which is typically sparse) to get the top d eigendecomposition [21], [28]. Indeed, if \mathbf{F} is of the form

$$\mathbf{F}_p \triangleq \text{poly}(\hat{\mathbf{A}}) = q_1\hat{\mathbf{A}} + \dots + q_n\hat{\mathbf{A}}^n \quad q_i \geq 0,$$

then it's trivial to see that if α is an eigenvalue with eigenvector x_α of $\hat{\mathbf{A}}$ then $\text{poly}(\alpha)$ will be an eigenvalue of \mathbf{F}_p with the same eigenvector x_α . Hence eigenvalues are rescaled, and eigenvectors are preserved. The only difficulty is that this rescaling does not preserve the ordering. More precisely, the top d eigenvalues of \mathbf{F}_p are in general not equal to the rescaled top d eigenvalues of $\hat{\mathbf{A}}$. To get the top d eigenvalues of \mathbf{F}_p one needs to calculate $l \geq d$ eigenvalues of $\hat{\mathbf{A}}$, where l denotes the index of the d -th positive eigenvalue of $\hat{\mathbf{A}}$, when sorted according to absolute value. This is true since $q_i \geq 0$ guarantees that the ordering is preserved only for the *positive* eigenvalues of $\hat{\mathbf{A}}$. In practice l is often not significantly larger than d . For

example, for sufficiently large Erdos-Renyi graphs $l \approx 2d$ due to the semicircle law [29].

2) *Resource Allocation Index (RAI) and Adamic-Adar (AA)*: One is not limited to polynomials of $\hat{\mathbf{A}}$ for scalable block-approximation. For other practical feature matrices \mathbf{F} , we can still utilize an eigendecomposition of $\hat{\mathbf{A}}$ to get a fast block-approximation of \mathbf{F} . One such matrix often used in the complex networks community is the so-called *Resource Allocation Index (RAI)* [30]. The RAI defines a similarity score r between two nodes u and v as $r_{uv} = \sum_{k \in \mathcal{N}(u) \cap \mathcal{N}(v)} 1/d_k$. The induced matrix $\mathbf{F}_{\text{RAI}} \triangleq [r_{uv}]$ can be written in terms of the adjacency matrix $\hat{\mathbf{A}}$:

$$\mathbf{F}_{\text{RAI}} = \hat{\mathbf{A}} \mathbf{D}^{-1} \hat{\mathbf{A}}, \quad (8)$$

where \mathbf{D} is a diagonal matrix containing the degree of each node. Assume $\hat{\mathbf{A}}$ is connected, such that $d_k > 0$ and (8) is well-defined. Given a top d eigendecomposition of $\hat{\mathbf{A}} \approx \mathbf{V} \Lambda \mathbf{V}^T$, with orthonormal columns of $\mathbf{V} \in \mathbb{R}^{n \times d}$, one obtains a rank d approximation of \mathbf{F}_{RAI} as follows:

$$\begin{aligned} \mathbf{F}_{\text{RAI}} &\approx \mathbf{V} (\Lambda \mathbf{V}^T \mathbf{D}^{-1} \mathbf{V} \Lambda) \mathbf{V}^T \\ &= \tilde{\mathbf{V}} \tilde{\mathbf{V}}^T, \end{aligned}$$

where $\tilde{\mathbf{V}} \triangleq \mathbf{V} (\Lambda \mathbf{V}^T \mathbf{D}^{-1} \mathbf{V} \Lambda)^{1/2} \in \mathbb{R}^{n \times d}$. These expressions are well-defined, since the positive definiteness of $\mathbf{V}^T \mathbf{D}^{-1} \mathbf{V} \in \mathbb{R}^{d \times d}$ implies positive definiteness of $\Lambda \mathbf{V}^T \mathbf{D}^{-1} \mathbf{V} \Lambda \in \mathbb{R}^{d \times d}$, hence the principal root exists and is unique. To obtain a block-approximation $\bar{\mathbf{F}}_{\text{RAI}}$, simply cluster the rows of $\tilde{\mathbf{V}}$ into bins and replace the rows by centroids.

The *Adamic-Adar Index (AA)* [31] can be block-approximated in a very similar fashion. It is defined similarly as (8), by substituting d_k by $\log(d_k)$. Nodes with $d_k = 1$ lead to an ill-defined \mathbf{D}^{-1} matrix, but since they never occur as a common neighbor of two other nodes, these nodes are omitted in the calculations.

3) *Preferential Attachment (PA)*: The preferential attachment feature matrix [12], [32] is defined as $\mathbf{F}_{\text{PA}} \triangleq [|\mathcal{N}(i)| \cdot |\mathcal{N}(j)|]$, i.e., the matrix induced by the product of the degrees. By definition \mathbf{F}_{PA} is already rank one; it is the outer product of a vector of degrees with itself. For reasons discussed in the proof of Theorem 1, the natural way to block-approximate (in this case, exactly) $\bar{\mathbf{F}}_{\text{PA}} = \mathbf{F}_{\text{PA}}$ is by considering the unique degrees in the network.

4) *Cross/Skeleton decompositions*: Alternatively, general methods from the vast literature on scalable low-rank approximations [33]–[37] can be utilized. The most scalable methods (Cross/Skeleton decompositions) essentially sample rows and columns to obtain a low-rank decomposition. We did not utilize these methods. Instead, we restricted ourselves to the feature matrices defined above, for which a fast top d eigendecomposition of $\hat{\mathbf{A}}$ was sufficient to obtain qualitative block-approximations.

C. Overall running time

Eigendecomposition & k-means. Computing the top l eigendecomposition of $\hat{\mathbf{A}}$ is efficient for sparse symmetric matrices

TABLE I. Summary of the datasets.

Dataset	Category	$ V $	$ E $
StudentDB [44]	Relational	395	3, 423
Facebook [45]	Social	4, 039	88, 234
PPI [46]	Biological	3, 852	37, 841
Wikipedia [47]	Language	4, 777	92, 295
GR-QC [45]	Collaboration	4, 158	26, 844
BlogCatalog [48]	Social	10, 312	333, 983
YouTube [49]	Social	1, 138, 499	2, 990, 443
Flickr [45]	Social	80, 513	11, 799, 764
DBLP [45]	Collaboration	317, 080	1, 049, 866

using iterative methods [38], [39], scaling linearly with n for a fixed number of iterations. Moreover, this is only computed once. Running the k-means++ algorithm for t iterations has time complexity $O(tnk d)$ [26]. There are known instances [27] for which $t = 2^{\Theta(\sqrt{n})}$ until convergence, but the $O(\log(k))$ approximation ratio in expectation (Proposition 1) is valid even after the initialization of the clustering ($t = 1$). Hence, overall running time is linear in n for fixed t, k and d .

Optimizing the reduced Lagrange dual function. For sparse graphs ($m = O(n)$), Theorem 1 shows that the final computational step is to solve an unconstrained convex optimization problem with $O(\sqrt{kn})$ variables. The computational complexity for computing the gradient as well as the Hessian is $O(kn)$. Space complexity for storing the gradient is $O(\sqrt{kn})$ and $O(kn)$ for the Hessian, which roughly (for small k) equals the space complexity of storing the graph in sparse representation. In Section IV-C we compare three different optimization strategies, and conclude that L-BFGS [40] is particularly well-suited for this objective function. This has been observed before, as L-BFGS has been described as the “algorithm of choice” for fitting log-linear (i.e., maximum entropy) models [41], [42].

D. Combining multiple feature matrices

Theorem 1 is formulated for the case of only one block-approximated matrix $\bar{\mathbf{F}}$. Combining multiple block-approximations $\bar{\mathbf{F}}_1, \dots, \bar{\mathbf{F}}_\gamma$ can be done by considering the *greatest lower bound* of the node binning. Each $\bar{\mathbf{F}}_i$ defines a partition B_i on the set of nodes. A partition B_i is a refinement of a partition B_j if each element of B_i is a subset of some element in B_j . This relation $B_i \leq B_j$ defines a partial order [43] and the set of all partitions form a lattice. A given set of partitions $\{B_1, \dots, B_\gamma\}$ thus has a greatest lower bound $b \leq B_i$. Theorem 1 still holds for multiple $\bar{\mathbf{F}}_i$ matrices, if one replaces k with $|b|$ (and adds an additive term of γ instead of one). Assuming an equal number of bins k for each $\bar{\mathbf{F}}_i$, worst-case this amounts to a lower bound with $|b| = \min\{k^\gamma, n\}$ bins. However, for a limited number of global constraints γ and small k , we already obtain qualitative and scalable results in practice (Section IV).

IV. EVALUATION

We test the performance of our ERG models on an important downstream task: link prediction (Section IV-A). To ensure reproducibility for the link prediction experiment, we utilize the EvalNE library [50]. Model implementations, as well as

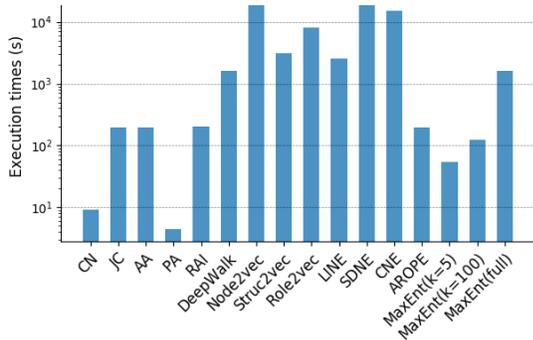


Fig. 3. Total execution times for the link prediction task on the *StudentDB*, *Facebook*, *PPI*, *Wikipedia* and *GR-QC* datasets.

customized configuration files describing the experiments, are publicly available.⁵ Section IV-B evaluates our proposed models on a social network using a goodness-of-fit approach. Detailed runtime experiments and optimization strategies are discussed in Section IV-C. Datasets are listed in Table I. Experiments were conducted on a Linux server with 256GB of RAM.

A. Link Prediction

In link prediction the aim is to identify missing links from a given network. In this task, we randomly remove 50% of the edges such that the remaining network is still connected. The reduced network is used for training, the removed edges are used for testing. We compare with eight state-of-the-art network embedding methods (Deepwalk [51], Node2vec [52], Struc2vec [53], Role2vec [54], LINE [55], SDNE [56], CNE [57] and AROPE [21]) and five common heuristics: common neighbours (CN), Adamic-Adar index (AA), Jaccard coefficient (JC), preferential attachment (PA) and Resource Allocation Index (RAI). Method descriptions, parameter tuning and further details on the experimental setup are discussed in Appendix B.

MaxEnt (full). First, we evaluate an exact model according to (4) with constraints on node degrees, and global constraints on \mathbf{F}_{CN} , \mathbf{F}_{RAI} and \mathbf{F}_{PA} . This model is denoted as MaxEnt (full) in Table II.

MaxEnt (blocked). Secondly, two block-approximated models where both $\overline{\mathbf{F}}_{\text{CN}}$ (Section III-B1) and $\overline{\mathbf{F}}_{\text{RAI}}$ (Section III-B2) are binned into $k \in \{5, 100\}$ bins and with $d = 128$. The matrix $\overline{\mathbf{F}}_{\text{PA}}$ is naturally binned by the unique degrees (Section III-B3). These models are denoted as MaxEnt ($k = 5$) and MaxEnt ($k = 100$) in Table II.

In Table II we present the average Area Under the ROC Curve (AUC) over three experiment repeats with different train/test splits for all methods. MaxEnt (full) performs well on all datasets and is never far from the optimal value achieved across all methods. As expected, the method does not scale to large networks. The block-approximated models ran on all networks and display competitive results. Figure 3 shows total execution times on the datasets where all methods successfully terminated within time.

⁵bitbucket.org/ghentdatascience/maxent-public

B. Goodness-of-fit on a social network

A standard visual approach for ERG model evaluation is to plot goodness-of-fits [8], [58]. The idea is to compare a set of higher-order network statistics, preferably statistics that are not directly modeled, with a range of the same statistics obtained by simulating random graphs from the model. We test on the *Facebook* dataset, which is essentially a combination of social circles (communities), combined with individual ego-networks. We test on three statistics that have been previously proposed to evaluate social network models [58]:

- **Minimum geodesic distance;** the shortest path distances distribution between all pairs of nodes in the network. Unreachable node pairs get assigned the value ‘Inf’.
- **Triad census;** the distribution of the number of edges between a set of three nodes.
- **Edgewise shared partners;** let CN_i denote the number of edges in the network that have exactly $0 \leq i \leq n - 2$ common neighbors. This defines a distribution $\text{CN}_0/m, \text{CN}_1/m, \dots, \text{CN}_{n-2}/m$.

We fit an approximative MaxEnt model, with constraints on node degrees, block-approximated $\overline{\mathbf{F}}_{\text{PA}}$ (unique degrees) and $\overline{\mathbf{F}}_{\text{RAI}}$ ($d = 20, k = 100$). Secondly, we fit an exact MaxEnt model with degree constraints and a polynomial constraint $\mathbf{F} = \hat{\mathbf{A}} + \beta \hat{\mathbf{A}}^2$, with $\beta = 0.025$. We compare with two other independent edge models. The Chung-Lu model (CL) [59] assigns a probability of $d_i d_j / \sum_k d_k$ to each possible node pair (i, j) , where the probabilities are clipped to one if $d_i d_j > \sum_k d_k$. Since the Chung-Lu model does not take into account other properties besides degrees, we do not expect this model to perform well on any of the aforementioned statistics.

To draw a more fair comparison, we compare with a model that takes community structure into account: the Block Two-Level Erdos-Renyi (BTER) model [60], [61]. The BTER model is a well-known scalable generative network model that can be tuned to capture and preserve two fundamental properties: degree distribution and (local) clustering coefficients.

Figure 4 shows the goodness-of-fit plots. The Chung-Lu model fails to capture any form of community structure—which the dataset most surely has—resulting in poor performance on all statistics. In contrast, both MaxEnt models with global constraints capture the triad census and edgewise shared partners distributions quite well. The geodesic distances are less well captured by the block-approximated MaxEnt model. The block-approximated MaxEnt model captures the community structure, but it still assigns a decent amount of probability mass to inter-community edges, reducing distances between nodes in different communities.

The BTER model performs well on the geodesic distances and the triad census, but interestingly, it shows adversarial oscillatory behavior on the edgewise shared partners distribution. Oscillatory behavior of statistics has been observed before in other generative models, e.g. it is known that Stochastic Kronecker Graphs have oscillating behavior of the degree distributions [62].

TABLE II. Average AUC for link prediction over three experiment repeats with different train/test splits for all methods. We use \times to indicate that a method did not finish within a pre-set time of 4 hours. Best results for each dataset are highlighted in bold.

	StudentDB	Facebook	PPI	Wiki	GR-QC	BlogCatalog	Flickr	YouTube	DBLP
CN	0.4101	0.9792	0.7737	0.8427	0.8602	0.9343	0.9379	0.5831	0.8127
JC	0.4101	0.9754	0.7613	0.4954	0.8598	0.8045	0.9316	0.5831	0.8127
AA	0.4101	0.9807	0.7764	0.8681	0.8604	0.9396	0.9383	0.5831	0.8127
PA	0.9202	0.8392	0.9022	0.9175	0.8311	0.9638	0.9676	0.9913	0.8866
RAI	0.4101	0.9813	0.776	0.8753	0.8603	0.9399	0.9376	0.5831	0.8127
DeepWalk	0.8865	0.9878	0.8867	0.8903	0.9627	0.9393	0.9772	\times	\times
Node2vec	0.9144	0.993	0.8552	0.8923	0.9593	0.922	\times	\times	\times
Struc2vec	0.92	0.8309	0.9006	0.9167	0.8215	0.96	\times	\times	\times
Role2vec	0.8653	0.9753	0.7979	0.7398	0.9386	0.8066	\times	\times	\times
LINE	0.9259	0.9875	0.8826	0.8628	0.9448	0.947	\times	\times	0.9026
SDNE	0.9695	0.9647	0.8885	0.9147	0.9066	0.9382	\times	\times	\times
CNE	0.8227	0.9082	0.8485	0.8611	0.8216	0.9193	\times	\times	\times
AROPE	0.9774	0.9863	0.899	0.9112	0.9191	0.9617	0.9825	0.9103	0.8757
MaxEnt ($k = 5$)	0.9542	0.8853	0.9018	0.9179	0.8365	0.9644	0.9699	0.9919	0.8863
MaxEnt ($k = 100$)	0.9626	0.9406	0.9026	0.9178	0.8752	0.9638	0.9699	0.9669	0.8857
MaxEnt (full)	0.9604	0.9694	0.9097	0.9182	0.9342	\times	\times	\times	\times

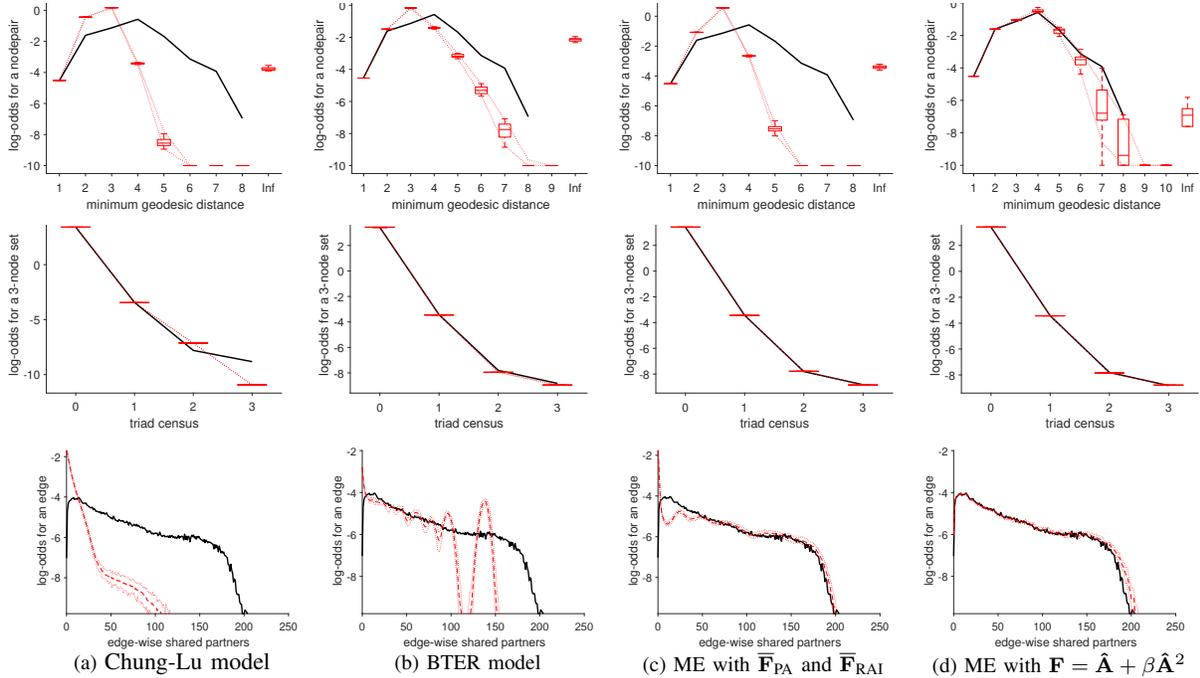


Fig. 4. Goodness of fit plots on the *Facebook* dataset for three statistics. The thick black lines are empirical distributions. The red lines are measurements from 50 randomly generated networks, according to four different independent edge models: (a) Chung-Lu (b) BTER model (c) MaxEnt with degrees and \bar{F}_{PA} and \bar{F}_{RAI} (block-approximated with $d = 20$ and $k = 100$) d) MaxEnt with degrees, and a constraint on $\mathbf{F} = \hat{\mathbf{A}} + \beta \hat{\mathbf{A}}^2$, with $\beta = 0.025$. Thin dotted red lines indicate 90% confidence intervals.

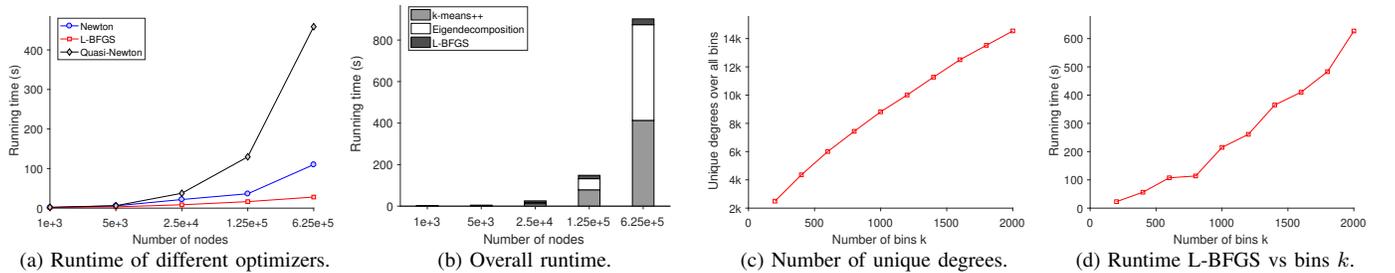


Fig. 5. Influence of graph size n and bins k on fitting times (norm gradient $< 10^{-3}$) of a MaxEnt model with constraints on degrees and block-approximated \bar{F}_{CN} . In (a) and (b) we fix $k = 100$ and let the graph size vary. In (c) and (d) we fix $n = 10^5$ and let the bin size k vary.

The MaxEnt model fitted with the polynomial constraint is a very close fit to the original network, outperforming all other models (Figure 4d). As such, it should be avoided for tasks like link prediction, but is useful to simulate networks that are close to the observed network. We refer to Section II-D for a discussion on label leakage.

C. Detailed runtime testing

Influence of the graph size. We generate synthetic Erdos-Renyi random graphs with varying number of nodes $n \in \{1, 5, 25, 125, 625\} \cdot 10^3$ with edges $m \approx 10n$. We test the influence of the graph size n on scalability, and compare three optimization strategies for optimizing (5). We solve a MaxEnt model with constraints on degrees and a block-approximated $\bar{\mathbf{F}}_{\text{CN}}$ ($d = 20$ and $k = 100$). We compare three different optimizers using the minFunc optimization package [63]: i) full Newton’s method ii) a diagonal quasi-Newton’s method iii) L-BFGS, a well-known and popular quasi-Newton method for parameter estimation in machine learning [40]. All optimizers use the same Wolfe line-search criteria to ensure global convergence. Figure 5a shows the time needed to reach a norm gradient tolerance of 10^{-3} , showing the superior performance of L-BFGS. Figure 5b shows the overall time needed to block approximate $\bar{\mathbf{F}}_{\text{CN}}$, and then fitting the reduced model with L-BFGS. K-means++ was repeated five times per instance, and each try was given a 100 iterations. The overall fitting time is mostly dominated by the eigendecomposition of $\hat{\mathbf{A}}$ and the k-means clustering.

Influence of the binning. Secondly, we generate 10 Erdos-Renyi graphs with fixed graph size $n = 10^5$ and $m \approx 10n$. The number of bins k are varied according to $200 \leq k \leq 2000$, in steps of 200. Figure 5d shows the average running time needed for L-BFGS to reach a gradient norm tolerance of 10^{-3} . Figure 5c shows the average sum of the number of unique degrees across all the bins, which is a measure of the number of variables that need to be optimized over (Section 3.3). Theorem 1 shows that this number scales as $O(\sqrt{kn})$ and Figure 5c indeed provides slight evidence for this.

V. RELATED WORK & CONCLUDING DISCUSSION

Although ERGs have a long history [4], [64] and have been successfully used as network models [8], [58], [65], they usually suffer from limited scalability and degeneracy [9]. A theoretical explanation for degeneracy is given by [66] for dense graphs. The authors of [67] try to resolve degeneracy by limiting the support of an ERG. Recently, [68] proposed slightly more scalable parameter estimates by exploiting properties of Markov chains at equilibrium, allowing for inference on a network with 10^5 nodes. Our paper circumvents both shortcomings by approximating dyadic independence models by means of block-approximated feature matrices, allowing for qualitative inference on sparse graphs with millions of nodes. With sensible feature selection, degeneracy can be avoided (Section II-D), and block-approximating typically results in additional probability smoothing.

The authors from [19] argued that the properties of the pseudo likelihood for analyzing social networks are poorly understood. In this paper, we have argued that dyadic independence models can be valuable by themselves. A possible explanation of the strength of the proposed models is the combination of both local role-based similarities (degrees) [69], with community information. Leveraging other graphlets counts, e.g., by considering weighted common graphlet counts as features, is a promising avenue for further work. Recent methods [17], [70] have been proposed to obtain approximative low-rank decompositions of such matrices.

This paper essentially provides a general framework on how to fit maximum entropy models with local and global constraints. The main idea is to smoothen the global constraints into regions of the graph where they are roughly the same, and then solving a smaller optimization problem by noting that the local constraints have limited uniqueness in those regions. Including for node- or edge attributed graphs and other types of graphs is an open avenue for further work.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC Grant Agreement no. 615517, from the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” programme, and from the FWO (project no. G091017N, G0F9816N, 3G042220).

REFERENCES

- [1] E. M. Airoldi, D. M. Blei, S. E. Fienberg, and E. P. Xing, “Mixed membership stochastic blockmodels,” in *Advances in Neural Information Processing Systems 21* (D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, eds.), pp. 33–40, Curran Associates, Inc., 2009.
- [2] C. Borgs, J. Chayes, H. Cohn, and Y. Zhao, “An lp theory of sparse graph convergence ii: Ld convergence, quotients, and right convergence,” vol. 46, August 2014.
- [3] P. Holland and S. Leinhardt, “An exponential family of probability distributions for directed graphs,” *Journal of The American Statistical Association - J AMER STATIST ASSN*, vol. 76, pp. 33–50, 03 1981.
- [4] O. Frank and D. Strauss, “Markov graphs,” *Journal of the American Statistical Association*, vol. 81, no. 395, pp. 832–842, 1986.
- [5] T. A. B. Snijders, P. E. Pattison, G. L. Robins, and M. S. Handcock, “New specifications for exponential random graph models,” *Sociological Methodology*, vol. 36, no. 1, pp. 99–153, 2006.
- [6] G. Robins, P. Pattison, Y. Kalish, and D. Lusher, “An introduction to exponential random graph (p*) models for social networks,” *Social Networks*, vol. 29, pp. 173–191, 05 2007.
- [7] J. Park and M. E. J. Newman, “Statistical mechanics of networks,” *Physical review. E, Statistical, nonlinear, and soft matter physics*, vol. 70, p. 066117, 01 2005.
- [8] S. Goodreau, “Advances in exponential random graph (p*) models applied to large social networks,” *Social networks*, vol. 29, 2007.
- [9] M. Handcock, “Assessing degeneracy in statistical models of social networks,” *Journal of the American Statistical Association*, vol. 76, pp. 33–50, 01 2003.
- [10] C. Hillar and A. Wibisono, “Maximum entropy distributions on graphs,” 2013.
- [11] T. De Bie, “Maximum entropy models and subjective interestingness: An application to tiles in binary databases,” *Data Min. Knowl. Discov.*, vol. 23, pp. 407–446, Nov. 2011.
- [12] F. Parisi, G. Caldarelli, and T. Squartini, “Entropy-based approach to missing-links prediction,” *Applied Network Science*, vol. 3, 02 2018.

- [13] T. M. Cover and J. A. Thomas, *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. New York, NY, USA: Wiley-Interscience, 2006.
- [14] W. Zachary, "An information flow model for conflict and fission in small groups1," *Journal of anthropological research*, vol. 33, 11 1976.
- [15] M. E. J. Newman, "Assortative mixing in networks," *Phys. Rev. Lett.*, vol. 89, p. 208701, Oct 2002.
- [16] N. K. Ahmed, T. L. Willke, and R. A. Rossi, "Estimation of local subgraph counts," in *2016 IEEE International Conference on Big Data (Big Data)*, pp. 586–595, Dec 2016.
- [17] R. A. Rossi, N. K. Ahmed, and E. Koh, "Higher-order network representation learning," in *WWW 2018, WWW '18*, (Republic and Canton of Geneva, CHE), 2018.
- [18] D. Strauss and M. Ikeda, "Pseudolikelihood estimation for social networks," *Journal of The American Statistical Association - J AMER STATIST ASSN*, vol. 85, pp. 204–212, 03 1990.
- [19] M. Duijn, K. Gile, and M. Handcock, "A framework for the comparison of maximum pseudo likelihood and maximum likelihood estimation of exponential family random graph models," *Social Networks*, vol. 31, pp. 52–62, 01 2009.
- [20] D. Lusseau, K. Schneider, O. Boisseau, P. Haase, E. Slooten, and S. Dawson, "The bottlenose dolphin community of doubtful sound features a large proportion of long-lasting associations," *Behavioral Ecology and Sociobiology*, vol. 54, pp. 396–405, 01 2003.
- [21] Z. Zhang, P. Cui, X. Wang, J. Pei, X. Yao, and W. Zhu, "Arbitrary-order proximity preserved network embedding," in *KDD*, 2018.
- [22] L. Katz, "A new status index derived from sociometric analysis," *Psychometrika*, vol. 18, pp. 39–43, Mar 1953.
- [23] F. Adriaens, J. Lijffijt, and T. De Bie, "Subjectively interesting connecting trees and forests," *Data Mining and Knowledge Discovery*, vol. 33, pp. 1088–1124, Jul 2019.
- [24] F. Adriaens, J. Lijffijt, and T. De Bie, "Subjectively interesting connecting trees," in *Machine Learning and Knowledge Discovery in Databases*, (Cham), Springer International Publishing, 2017.
- [25] U. Luxburg, "A tutorial on spectral clustering," *Statistics and Computing*, vol. 17, pp. 395–416, 01 2004.
- [26] D. Arthur and S. Vassilvitskii, "K-means++: The advantages of careful seeding," vol. 8, pp. 1027–1035, 01 2007.
- [27] D. Arthur and S. Vassilvitskii, "How slow is the k-means method?," *Proceedings of the Annual Symposium on Computational Geometry*, vol. 2006, pp. 144–153, 01 2006.
- [28] N. J. Higham and F. Tisseur, "Bounds for eigenvalues of matrix polynomials," *Linear Algebra and its Applications*, vol. 358, no. 1, pp. 5–22, 2003.
- [29] L. Erdos, A. Knowles, H.-T. Yau, and J. Yin, "Spectral statistics of erdos-renyi graphs i: Local semicircle law," *The Annals of Probability*, vol. 41, 05 2013.
- [30] T. Zhou, L. Lü, and Y.-C. Zhang, "Predicting missing links via local information," *The European Physical Journal B*, vol. 71, 2009.
- [31] L. A. Adamic and E. Adar, "Friends and neighbors on the web," *Social Networks*, vol. 25, no. 3, pp. 211–230, 2003.
- [32] A. Grover and J. Leskovec, "node2vec: Scalable Feature Learning for Networks," *KDD '16*, (New York, NY, USA), ACM Press, 2016.
- [33] S. Goreinov, E. Tyrtshnikov, and N. Zamarashkin, "A theory of pseudoskeleton approximations," *Linear Algebra and its Applications*, vol. 261, no. 1, pp. 1–21, 1997.
- [34] D. Achlioptas and F. Mcsherry, "Fast computation of low-rank matrix approximations," *J. ACM*, vol. 54, Apr. 2007.
- [35] I. Markovsky, *Low Rank Approximation: Algorithms, Implementation, Applications*. Springer, 2012.
- [36] N. K. Kumar and J. Schneider, "Literature survey on low rank approximation of matrices," *Linear and Multilinear Algebra*, vol. 65, no. 11, pp. 2212–2244, 2017.
- [37] P. Indyk, A. Vakilian, T. Wagner, and D. Woodruff, "Sample-optimal low-rank approximation of distance matrices," in *COLT*, 2019.
- [38] R. B. Lehoucq and D. C. Sorensen, "Deflation techniques for an implicitly restarted arnoldi iteration," *SIAM J. Matrix Analysis Applications*, vol. 17, pp. 789–821, 1996.
- [39] G. W. Stewart, "A krylov–schur algorithm for large eigenproblems," *SIAM J. Matrix Anal. Appl.*, vol. 23, pp. 601–614, Mar. 2001.
- [40] D. C. Liu and J. Nocedal, "On the limited memory bfgs method for large scale optimization," *Mathematical Programming*, vol. 45, 1989.
- [41] R. Malouf, "A comparison of algorithms for maximum entropy parameter estimation," vol. 20, pp. 49–55, 08 2002.
- [42] G. Andrew and J. Gao, "Scalable training of l_1 -regularized log-linear models," in *ICML*, January 2007.
- [43] G. Birkhoff, *Lattice Theory*. Providence: American Mathematical Society, 3rd ed., 1967.
- [44] B. Goethals, W. Le Page, and M. Mampaey, "Mining interesting sets and rules in relational databases," in *Proceedings of the 2010 ACM Symposium on Applied Computing, SAC '10*, (New York, NY, USA), pp. 997–1001, ACM, 2010.
- [45] J. Leskovec and A. Krevl, "SNAP Datasets: Stanford large network dataset collection," 2015.
- [46] B.-J. Breitkreutz, C. Stark, T. Reguly, L. Boucher, A. Breitkreutz, M. Livstone, R. Oughtred, D. H. Lackner, J. Bähler, V. Wood, *et al.*, "The biogrid interaction database: 2008 update," *Nucleic acids research*, vol. 36, pp. D637–D640, 2007.
- [47] M. Mahoney, "Large text compression benchmark," *URL*, 2011.
- [48] R. Zafarani and H. Liu, "Social computing data repository at asu," 2009.
- [49] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee, "Measurement and Analysis of Online Social Networks," in *Proceedings of the 5th ACM/Usenix Internet Measurement Conference (IMC'07)*, (San Diego, CA), October 2007.
- [50] A. Mara, J. Lijffijt, and T. D. Bie, "Evalne: A framework for evaluating network embeddings on link prediction," 2019.
- [51] B. Perozzi, R. Al-Rfou, and S. Skiena, "Deepwalk: Online learning of social representations," in *Proc. of KDD*, pp. 701–710, 2014.
- [52] A. Grover and J. Leskovec, "node2vec: Scalable feature learning for networks," in *Proc. of KDD*, pp. 855–864, 2016.
- [53] L. F. Ribeiro, P. H. Saverese, and D. R. Figueiredo, "struc2vec: Learning node representations from structural identity," in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 385–394, ACM, 2017.
- [54] N. Ahmed, R. A. Rossi, J. B. Lee, X. Kong, T. L. Willke, R. Zhou, and H. Eldardiry, "Learning role-based graph embeddings," *ArXiv*, vol. abs/1802.02896, 2018.
- [55] J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, and Q. Mei, "LINE: Large-scale information network embedding," in *Proc. of WWW*, 2015.
- [56] D. Wang, P. Cui, and W. Zhu, "Structural deep network embedding," in *Proc. of KDD*, pp. 1225–1234, 2016.
- [57] B. Kang, J. Lijffijt, and T. D. Bie, "Conditional network embeddings," in *International Conference on Learning Representations*, 2019.
- [58] D. R. Hunter, S. M. Goodreau, and M. S. Handcock, "Goodness of fit of social network models," *Journal of the American Statistical Association*, vol. 103, 2008.
- [59] F. Chung and L. Lu, "The average distance in a random graph with given expected degree," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 99, pp. 15879–82, 12 2002.
- [60] T. G. Kolda, A. Pinar, T. Plantenga, and C. Seshadhri, "A scalable generative graph model with community structure," *SIAM Journal on Scientific Computing*, vol. 36, p. C424–C452, Jan 2014.
- [61] C. Seshadhri, T. G. Kolda, and A. Pinar, "Community structure and scale-free collections of erdos-renyi graphs," *Phys. Rev. E*, May 2012.
- [62] C. Seshadhri, A. Pinar, and T. G. Kolda, "An in-depth analysis of stochastic kronecker graphs," *J. ACM*, vol. 60, May 2013.
- [63] M. Schmidt, *minFunc: unconstrained differentiable multivariate optimization in Matlab*. <http://www.cs.ubc.ca/schmidtm/Software/minFunc.html>, 2005, 2005.
- [64] M. Newman, "The structure and function of complex networks," *Computer Physics Communications*, vol. 147, pp. 40–45, 03 2003.
- [65] J. Koskinen, G. Robins, and P. Pattison, "Analysing exponential random graph (p-star) models with missing data using bayesian data augmentation," *Statistical Methodology*, vol. 7, pp. 366–384, 05 2010.
- [66] S. Chatterjee and P. Diaconis, "Estimating and understanding exponential random graph models," *Ann. Statist.*, vol. 41, pp. 2428–2461, 10 2013.
- [67] V. Karwa, S. Petrovic, and D. Bajić, "Dergms: Degeneracy-restricted exponential random graph models," 12 2016.
- [68] M. Byshkin, A. Stivala, A. Mira, G. Robins, and A. Lomi, "Fast maximum likelihood estimation via equilibrium expectation for large network data," *Scientific Reports*, vol. 8, 03 2018.
- [69] R. A. Rossi, D. Jin, S. Kim, N. Ahmed, D. Koutra, and J. B. Lee, "From community to role-based graph embeddings," *ArXiv*, vol. abs/1908.08572, 2019.
- [70] R. A. Rossi, R. Zhou, and N. K. Ahmed, "Estimation of graphlet counts in massive networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 30, pp. 44–57, Jan 2019.

Appendix A [Model derivation].

When considering the Lagrangian of the optimization problem (4), and setting derivatives with respect to $P(\mathbf{A})$ equal to zero, one obtains the following form for $P(\mathbf{A})$:

$$P(\mathbf{A}) = \frac{\exp(\sum_l \lambda_l \sum_{i,j} f_{ij}^l \mathbf{A}_{ij})}{Z} = \frac{\prod_{i,j} \exp(\sum_l \lambda_l f_{ij}^l \mathbf{A}_{ij})}{Z}, \quad (9)$$

where $Z = Z(\lambda_1, \dots, \lambda_M)$ is the partition function (i.e., the normalizing constant). Evaluating Z yields

$$\begin{aligned} Z &= \sum_{\mathbf{A} \in \{0,1\}^{n \times n}} \prod_{i,j} \exp(\sum_l \lambda_l f_{ij}^l \mathbf{A}_{ij}) \\ &= \prod_{i,j} \sum_{\mathbf{A}_{ij} \in \{0,1\}} \exp(\sum_l \lambda_l f_{ij}^l \mathbf{A}_{ij}) \\ &= \prod_{i,j} (1 + \exp(\sum_l \lambda_l f_{ij}^l)). \end{aligned}$$

Hence $P(\mathbf{A})$ factorizes as a product of independent Bernoulli distributions.

On the other hand, the pseudolikelihood (see e.g., [19]) of an ERG (1) with parameters θ and related statistics s is equal to

$$P_{\text{pseudo}}(\mathbf{A}) = \prod_{i,j} \frac{\exp(\sum_l \theta_l \Delta s_{ij}^l \mathbf{A}_{ij})}{1 + \exp(\sum_l \theta_l \Delta s_{ij}^l)},$$

where Δs_{ij}^l denotes the change in the statistic s^l when going from a realization of \mathbf{A} without the edge (i, j) to a realization of \mathbf{A} that includes the edge (i, j) . Hence P_{pseudo} is identical to (9), when the f_{ij}^l features are chosen to be equal to Δs_{ij}^l , as will be the case for statistics that are *counts* of certain graph related properties (degrees, triangles, etc.). For example, consider the out-degree of node i as a statistic. Then $\Delta s_{k,j}^l = 1$ when $k = i$ and $j \neq i$, and $\Delta s_{k,j}^l = 0$ elsewhere. This is exactly equal to the f_{ij}^l degree-features defined in Section II-B.

Appendix B [Experimental setup link prediction]. The sets of train and test edges are topped with equal amounts of non-edges drawn uniformly at random from the original graph. The train set is further divided in 90% train and 10% validation for hyper-parameter tuning.

We compare with the following embedding based methods:

- *Deepwalk* [51] computes node embeddings using truncated random walks and the Skipgram model approximated by hierarchical softmax.

- *Node2vec* [52] is a generalization of DeepWalk which approximates the Skipgram model by means of negative sampling. The model uses two parameters p and q that interpolate between the importance of BFS/DFS random walk strategies.
- *Struc2vec* [53] extracts structural information from the graph through node pair similarities for a range of neighbourhood sizes. This information is then summarized as a multi-layer weighted graph. A random walk on this graph is used to generate the embeddings.
- *Role2vec* [54] is a space-efficient random walk based approach which learns embeddings for different types of nodes.
- *LINE* [55] is a probabilistic approach which computes node embeddings based on first and second order similarities between network nodes.
- *SDNE* [56] uses a deep auto-encoder to generate embeddings which capture first and second order proximities.
- *CNE* [57] uses a Bayesian approach to generate embeddings using structural graph properties as priors.
- *AROPE* [21] proposes embeddings as found by the truncated singular value decompositions of higher order proximities.

The link prediction heuristics and embedding methods AROPE, CNE and MaxEnt provide node similarities which can be directly interpreted as probabilities of linking nodes. For the remaining methods we apply Logistic Regression with 5-fold cross validation on the edge embeddings to obtain link predictions.

We set $d = 128$ for all methods. The following method parameters were tuned using grid search: for DeepWalk and Struc2vec window sizes $\{5, 10, 20\}$, for Node2vec and Role2vec $p = q \in \{0.5, 1, 2\}$ and window sizes $\{5, 10, 20\}$, for LINE learning rate $\rho \in \{0.01, 0.025\}$ and number of negative samples $\{5M, 10M\}$, for SDNE $\beta \in \{2, 5, 10\}$, for AROPE higher-order proximities of orders 1 up to 4 and for CNE the learning rate $\alpha \in \{0.01, 0.05\}$.

Most methods require a binary operator that transforms node embeddings into edge embeddings. We use the same operators proposed in [52], and tuned them as additional method hyper-parameters.