EM-Based Smooth Graphon Estimation Using Bayesian and Spline-Based Approaches

Benjamin Sischka and Göran Kauermann

September 16, 2021

Abstract

This paper proposes the estimation of a smooth graphon model for network data analysis using principles of the EM algorithm. The approach considers both variability with respect to ordering the nodes of a network and smooth estimation of the graphon by nonparametric regression. To do so, (linear) B-splines are used, which allow for smooth estimation of the graphon, conditional on the node ordering. This provides the M-step. The true ordering of the nodes arising from the graphon model remains unobserved and Bayesian ideas are employed to obtain posterior samples given the network data. This yields the E-step. Combining both steps gives an EMbased approach for smooth graphon estimation. Unlike common other methods, this procedure does *not* require the restriction of a monotonic marginal function. The proposed graphon estimate allows to explore node-ordering strategies and therefore to compare the common degree-based node ranking with the ordering conditional on the network. Variability and uncertainty are taken into account using MCMC techniques. Examples and simulation studies support the applicability of the approach.

Keywords: Graphon model; EM algorithm; MCMC; Gibbs sampling; B-spline surface; Social network; Political network; Connectome

Benjamin Sischka is Research Assistant, Department of Statistics, Ludwig-Maximilians-Universität München, 80539 München, Germany (E-mail: benjamin.sischka@stat.uni-muenchen.de). Göran Kauermann is Professor, Department of Statistics, Ludwig-Maximilians-Universität München, 80539 München, Germany (E-mail: goeran.kauermann@stat.uni-muenchen.de).

1 Introduction

The analysis of network data has achieved increasing interest in the last years. Goldenberg et al. (2010), Hunter et al. (2012), Fienberg (2012) and Salter-Townshend et al. (2012), respectively, published survey articles demonstrating the state-of-the-art in the field. We also refer to Kolaczyk (2009), Kolaczyk and Csárdi (2014) and Lusher et al. (2013) for monographs in the field of statistical network data analysis, see also Kolaczyk (2017). The statistical workhorse models for network data are Exponential Random Graph Models (ERGM), Stochastic Block Models (SBM) and Latent Space Models. ERGMs make use of an exponential family distribution to model the network's adjacency matrix as a random matrix. This model class was proposed by Frank and Strauss (1986) and is extensively discussed in Snijders et al. (2006). SBMs as well as the latent space models are both model classes which make use of latent quantities relating to the nodes, that are employed to explore the network structure. They were proposed in their nowadays form by Holland et al. (1983) and Hoff et al. (2002), respectively.

A different modeling strategy, which is also based on latent quantities, results through comprehending the network adjacency matrix $\mathbf{Y} \in \{0,1\}^{N \times N}$ to be generated by a so called graphon. The graphon as data generating model comes into play by assuming that we draw N random variables

$$U_1, \dots, U_N \stackrel{i.i.d.}{\sim} \text{Uniform}[0, 1]$$
 (1)

and, given U_i and U_j , simulate the network entries Y_{ij} conditionally independently through

$$Y_{ij}|U_i = u_i, U_j = u_j \sim \text{Binomial}(1, w(u_i, u_j)).$$
(2)

The function w(,) is thereby called a graphon (= graph function). In case of undirected networks we additionally require symmetry so that $Y_{ij} = Y_{ji}$ (and hence in principle we assume w(u, v) = w(v, u)). We remain with the commonly used convention and exclude so called self-loops, which means $Y_{ii} = 0$ for i = 1, ..., N. Apparently, the graphon function w(,) is not unique with regard to the data generating process (2). This is because for any measure-preserving bijection $\varphi : [0, 1] \to [0, 1]$ the permuted graphon $w(\varphi(u), \varphi(v))$ yields the same model. More generally, as has been stated by Diaconis and Janson (2008), two graphons w(,) and w'(,) represent the same generating model if and only if there exist two measure preserving mappings – not necessarily bijections – $\tilde{\varphi}$ and $\tilde{\varphi}' : [0, 1] \to [0, 1]$ such that $w(\tilde{\varphi}(u), \tilde{\varphi}(v)) = w'(\tilde{\varphi}'(u), \tilde{\varphi}'(v))$ for almost all $(u, v) \in [0, 1]^2$. Some papers therefore add a further attribute to achieve uniqueness, e.g. see Bickel and Chen (2009) or Yang et al. (2014). The common setting to do so is to postulate that

$$g(u) = \int w(u, v) \,\mathrm{d}v \tag{3}$$

is strictly increasing in u, which leads to the so called canonical representation of the graphon w(,). Note that the distribution of $g(U_i)$ with U_i following (1) can be interpreted as (asymptotic) distribution of the degree proportion. However, the additional condition (3) incorporated to tackle the identifiability issue can be problematic, especially in combination with other postulations. For reconstruction reasons commonly some kind of smoothness is assumed in the sense that w(,) satisfies (at least piecewise) some Lipschitz or Hölder conditions of continuity, e.g. see Olhede and Wolfe (2014), Chan and Airoldi (2014) or Gao et al. (2015a). But then, considering an arbitrary continuous graphon, the transformation into its canonical representation neither usually retains the continuity nor are there necessarily any guarantees that such a canonical representation exists. Taking only graphon functions into account which are continuous in their canonical representation would apparently be a strong restriction of the generality of graphon models. We therefore consider the more general model class which does not include the canonical constraint. Still, we also discuss estimation under the restriction (3).

Graphon estimation for modeling network data has recently found attention in the statistical literature. A general discussion is given in Orbanz and Roy (2014) and You (2020) recently launched an R package for graphon estimation. Graphons can be related to ERGMs, at least for simple statistics like two-star or triangles, as shown in Diaconis and Chatterjee (2013). He and Zheng (2015) make use of this connection and propose to use asymptotic properties of graphons to derive estimates in high dimensional ERGMs.

Moreover, graphon models can be seen as a generalized model class which also includes SBMs and latent space models (e.g. see Bickel et al., 2011). Wolfe and Olhede (2013) and Yang et al. (2014) discuss non-parametric estimation of graphons including tests on the validity of prespecified graphon shapes, see also Chan and Airoldi (2014), Airoldi et al. (2013) or Olhede and Wolfe (2014). One of the most propagated strategies in the graphon estimation literature is to approximate the graphon through a SBM, see Choi and Wolfe (2014), Choi (2017) or Klopp et al. (2017). Gao et al. (2015a) discuss optimal graphon estimation for the SBM approximation. SBMs are, however, by definition discontinuous models, meaning that the corresponding graphon function is discontinuous and hence not smooth. In this paper, in contrast, we focus on smooth and hence continuous graphon functions. For a general discussion on graphons we refer to Borgs et al. (2008), Lovász (2012), Diaconis and Janson (2008) or Bickel and Chen (2009).

Estimating graphons can generally be reduced to probability matrix estimation, where the goal is to gain information about the specific pairwise probabilities $P(Y_{ij} = 1|U_i = u_i, U_j = u_j) = w(u_i, u_j)$. From this perspective, the U_i are not well-defined components. This reduction – considering the graphon values only at unspecific positions – allows to circumvent the identifiability issue from above. Chatterjee (2015) provides convergence results for general matrix estimation using singular value decomposition, see also Gao et al. (2015a), Zhang et al. (2015) or Xu (2017). Gao et al. (2016) extends this towards partially observed matrices/networks, see also Gao et al. (2015b) for a Bayesian approach. The cited work, however, yields estimates of specific graphon values at unspecific positions and thus does not lead to a smooth graphon estimate on the domain $[0, 1]^2$ which is focused in this work.

In this paper we propose to use penalized linear B-splines for graphon estimation. This borrows ideas suggested in Kauermann et al. (2013) for copula estimation, since B-splines easily allow to accommodate side constraints such as symmetry in form of w(u, v) = w(v, u)or, if required, condition (3) for the resulting estimate. This, in contrast, is difficult to accommodate in histogram or kernel based estimation. Penalized estimation with B-splines has thereby a long standing tradition in smooth estimation, starting with Eilers and Marx (1996) and Ruppert et al. (2003, 2009), see also Wood (2017). We extend the idea here to graphon estimation. However, before this smoothing approach can be applied, we also need to fill the lack of information about the U_i . On the other hand, conditional on the observed network data $\boldsymbol{y} = \{y_{ij}\}_{i,j=1,\dots,N}$, a presumption about the positions of the U_i can only be made in relation to w(,). Since both the implicitly supposed parameters of the B-spline approximation of w(,) and the latent U_i are unknown, this is a typical task for an EM type algorithm. Moreover, MCMC techniques can be applied in this context to approximate numerically the complex posterior distribution of the U_i , which together yields a MCEM algorithm.

The paper is organized as follows. Section 2 displays the main ideas of pursuing an EM based algorithm for smooth graphon estimation. Section 3 describes the procedure in detail. Section 4 and Section 5 showcase results for both simulations and real-world data examples. A discussion concludes the paper.

2 Graphon Representation and EM Motivation

We assume that the graphon $w : [0, 1]^2 \to [0, 1]$ is a smooth function, meaning that w(,) is sufficiently differentiable in both arguments. We call w(,) to have a canonical representation if $g(u) = \int w(u, v) \, dv$ is strictly increasing, see, among others, Chan and Airoldi (2014). We assume further that w(,) is symmetric and generates a network of size N through the following process. For N independent uniform variables

$$U_i \overset{i.i.d.}{\sim} \text{Uniform}[0,1], \quad i = 1, \dots, N$$

we obtain the symmetric network through

$$P(Y_{ij} = 1 | U_i = u_i, U_j = u_j) = w(u_i, u_j)$$
(4)

for $1 \leq i < j \leq N$, where $Y_{ji} = Y_{ij}$ and $Y_{ii} = 0$. The variables U_i remain unobservable and as data we only obtain the observed network \boldsymbol{y} . This implies in particular that no information on the values of U_i is given and hence estimation of w(,) occurs to be difficult. We propose to tackle the estimation problem using an EM algorithm. In that sense, we calculate (or rather approximate by simulating) $E(\boldsymbol{U}|\boldsymbol{y})$ from which we derive an appropriate ordering, giving the E-step, which in turn allows to estimate w(,) using smoothing techniques, providing the M-step. For the E-step we take a Bayesian view by looking at the posterior density of $\boldsymbol{U} = (U_1, \ldots, U_N)$ given \boldsymbol{y} . Note that since U_1, \ldots, U_N are *i.i.d.* uniform we have

$$f(\boldsymbol{u}|\boldsymbol{y}) \propto \prod_{\substack{i,j\\j>i}} w(u_i, u_j)^{y_{ij}} (1 - w(u_i, u_j))^{1-y_{ij}}.$$

If we look at the univariate distribution of a single variable U_k given the entire network \boldsymbol{y} , this results through

$$f_k(u_k|\boldsymbol{y}) \propto \int \dots \int \prod_{\substack{i,j \ j>i}} w(u_i, u_j)^{y_{ij}} (1 - w(u_i, u_j))^{1 - y_{ij}} \, \mathrm{d}u_1 \dots \mathrm{d}u_{k-1} \, \mathrm{d}u_{k+1} \dots \mathrm{d}u_N.$$
(5)

Apparently, both the joint and the marginal posterior distribution are too complex to be calculated analytically, in particular if N is large. We will therefore explore (5) by pursuing a Bayesian approach using MCMC techniques. Details are given in the next sections. Letting $\hat{u}_k^{(m)}$ denote the rescaled ranks of the posterior means (see equation (9) below) of the MCMC samples in the *m*th step of the EM algorithm, we then maximize the likelihood resulting from (4) with u_i replaced by $\hat{u}_i^{(m)}$. To do so, we make use of penalized B-spline smoothing, meaning that we replace the unknown graphon w(,) by an approximate spline base representation $B(,)\theta$, where B(,) is a bivariate spline and θ is the vector of unknown spline coefficients to be estimated in the M-step. Both steps the E- and the M-step will be introduced in detail below. Before doing so, however, we propose a simple first E-step, which also serves as initial step of the corresponding EM iteration.

Even though the probability model (4) used for the construction of networks is simple, it can not directly be used for estimation. The reason is that variables U_i are unobservable and hence can not directly be employed to estimate the graphon w(,). Instead, in the recent literature the graphon w(,) is usually estimated by smoothing the observed adjacency matrix \boldsymbol{y} . The preceding rearrangement of the nodes thereby should effect an appropriate reordering of the rows and columns which in turn should reflect the ordering of the U_i . Let $\psi : \{1, \ldots, N\} \rightarrow \{1, \ldots, N\}$ be a permutation such that

$$U_{\psi(i)} \le U_{\psi(j)}$$

with i < j. This means $U_{\psi(i)} = U_{(i)}$, where $U_{(1)} \leq U_{(2)} \leq \ldots \leq U_{(N)}$ define the ordered variables U_i . Note that since U_i , $i = 1, \ldots, N$ are not observable, we can also not observe $\psi()$ which therefore needs to be estimated. A common approach is to make use of the degree. Let therefore $\hat{\psi} : \{1, \ldots, N\} \rightarrow \{1, \ldots, N\}$ be a permutation such that

$$degree(\hat{\psi}(i)) \le degree(\hat{\psi}(j))$$
 (6)

for i < j. Note that $\hat{\psi}()$ can serve as an initial estimate for $\psi()$, although it requires the assumption that g() from (3) is strictly increasing. We therefore use it as initial E-step and define the corresponding resulting initial prediction for U_j based on this simple sorting through

$$\hat{u}_j^{(0)} = \frac{rank(degree(j))}{N+1},\tag{7}$$

where rank(degree(j)) is the rank from smallest to largest of the *j*th element of the tuple (degree(i) : i = 1, ..., N). This is equivalent to define $\hat{u}_{\hat{\psi}(j)}^{(0)} = j/(N+1)$, where i/(N+1), i = 1, ..., N represent the expected values of N ordered independently Uniform[0, 1] distributed variables. Chan and Airoldi (2014) prove – in case of graphons with canonical representation – asymptotic convergence rates for $\hat{\psi}(j)$ towards $\psi()$, meaning that $|\psi(j) - \hat{\psi}(j)|/N \to 0$ for all j = 1, ..., N.

We can now replace w(,) in (4) by its empirical version $\hat{w}^{(0)}(,)$ which is defined through

$$\hat{w}^{(0)}(u,v) = y_{\hat{\psi}(\lceil uN \rceil)\hat{\psi}(\lceil vN \rceil)},$$

where $\lceil uN \rceil$ defines the smallest value which is greater or equal to uN. Note that $\hat{w}^{(0)}(,)$ just mimics the ordered adjacency matrix scaled towards the unit square. These calculations provide the initial steps in the subsequently introduced EM algorithm.

3 EM Algorithm for smooth Graphons

3.1 Bayesian Approach for the E-Step

We pursue the E-step by exploiting the posterior distribution of U. This is done by constructing an appropriate MCMC Gibbs sampling scheme based on the full posterior distribution of U_k . Note that by conditioning on y and all U_j except of U_k one gets

$$f_k(u_k|u_1,\ldots,u_{k-1},u_{k+1},\ldots,u_N,\boldsymbol{y}) \propto \prod_{j \neq k} w(u_k,u_j)^{y_{kj}} (1 - w(u_k,u_j))^{1 - y_{kj}}.$$
(8)

We pretend in this section that the graphon w(,) is known, which is the general setting in the E-step. This allows to easily sample from (8) using Gibbs sampling as MCMC technique. To do so, we assume $\mathbf{u}^{<t>} = (u_1^{<t>}, \ldots, u_N^{<t>})$ to be the current state of the Markov chain. To update the kth component we then set $u_j^{<t+1>} := u_j^{<t>}$ for $j \neq k$, while component u_k is updated by drawing from (8). To pursue this, we make use of a normal proposal using a logit link. To be specific, let $z_k^{<t>} = \log(u_k^{<t>}/(1-u_k^{<t>})) = \log t (u_k^{<t>})$. We then propose to draw $z_k^* = z_k^{<t>} + \operatorname{Normal}(0, \sigma^2)$ and set $u_k^* = \log t^{-1}(z_k^*) = \exp(z_k^*)/(1+\exp(z_k^*))$. Hence, the proposal density for U_k is proportional to

$$\begin{aligned} q(u_k^*|u_k^{}) &= \frac{\partial u_k^*}{\partial z_k^*} \phi(z_k^*|z_k^{}) \\ &\propto \frac{1}{u_k^*(1-u_k^*)} \exp\left(-\frac{1}{2} \frac{(\text{logit } (u_k^*) - \text{logit } (u_k^{}))^2}{\sigma^2}\right), \end{aligned}$$

where $\phi()$ is the standard normal density. Consequently, the ratio of proposals equals

$$\frac{q_k(u_k^{}|u_k^*)}{q_k(u_k^*|u_k^{})} = \frac{u_k^*(1-u_k^*)}{u_k^{}(1-u_k^{})}.$$

The proposed value u_k^* is accepted (and hence we set $u_k^{< t+1>} := u_k^*$) with probability

$$\min\left\{1, \prod_{j \neq k} \left[\left(\frac{w(u_k^*, u_j^{< t>})}{w(u_k^{< t>}, u_j^{< t>})} \right)^{y_{kj}} \left(\frac{1 - w(u_k^*, u_j^{< t>})}{1 - w(u_k^{< t>}, u_j^{< t>})} \right)^{1 - y_{kj}} \right] \frac{u_k^*(1 - u_k^*)}{u_k^{< t>}(1 - u_k^{< t>})} \right\}$$

If we do not accept u_k^* , we set $u_k^{\langle t+1 \rangle} := u_k^{\langle t \rangle}$. Based on the resulting Markov chain we can estimate the posterior mean $E(U_k | \boldsymbol{y})$ by taking the mean of the simulated values, observing an appropriate burn-in phase of the algorithm. To be specific, we use the MCMC sequence to estimate the posterior mean in the *m*th step of the EM algorithm through

$$\bar{u}_k^{(m)} = \frac{1}{n} \sum_{s=1}^n u_k^{\langle s \cdot N \cdot r \rangle},$$

where $r \in \mathbb{N}$ describes a thinning factor and n is the number of MCMC states which are taken into account. We then use the ranks of the posterior means to reorder the network matrix accordingly. This corresponds to setting the missing values of U_j according to (7) to

$$\hat{u}_{j}^{(m)} = \frac{rank(\bar{u}_{j}^{(m)})}{N+1}.$$
(9)

The Gibbs sampling approach is straightforward and simple but requires the knowledge of the graphon w(,). Apparently, in the EM algorithm we replace w(,) in the formula above by the current estimate resulting through the M-step, which is described subsequently. We define the final estimate resulting through (9) as \hat{u}_j^{EM} .

3.2 Spline based Graphon Estimation for the M-Step

3.2.1 Linear B-Splines

For smooth estimation of the graphon w(,) we first formulate a spline based approximation through

$$w_{\boldsymbol{\theta}}^{spline}(u,v) = \boldsymbol{B}(u,v) \,\boldsymbol{\theta} = [\boldsymbol{B}(u) \otimes \boldsymbol{B}(v)] \,\boldsymbol{\theta},\tag{10}$$



Figure 1: Normalized univariate linear B-spline basis for the approximation of the graphon. The (equidistant) inner knots are denoted by τ_j with $j = 1, \ldots, K$.

where $\boldsymbol{B}(u) \in \mathbb{R}^{1 \times K}$ is a linear B-spline basis on [0, 1], normalized to have maximum value 1, see Figure 1. Parameter vector $\boldsymbol{\theta} \in \mathbb{R}^{K^2 \times 1}$ is indexed through

$$\boldsymbol{\theta} = (\theta_{11}, \dots, \theta_{1K}, \theta_{21}, \dots, \theta_{K1}, \dots, \theta_{KK})^{\top}$$

Using (10) we obtain the likelihood

$$l(\boldsymbol{\theta}) = \sum_{i} \sum_{j \neq i} \left[y_{ij} \log \left(\boldsymbol{B}_{ij} \boldsymbol{\theta} \right) + (1 - y_{ij}) \log \left(1 - \boldsymbol{B}_{ij} \boldsymbol{\theta} \right) \right],$$

where $\boldsymbol{B}_{ij} = \boldsymbol{B}(u_i) \otimes \boldsymbol{B}(u_j)$. Taking the derivative leads to the score function

$$\boldsymbol{s}(\boldsymbol{\theta}) = \sum_{i} \sum_{j \neq i} \boldsymbol{B}_{ij}^{\top} \left(\frac{y_{ij}}{w_{\boldsymbol{\theta}}^{spline}(u_i, u_j)} - \frac{1 - y_{ij}}{1 - w_{\boldsymbol{\theta}}^{spline}(u_i, u_j)} \right).$$

Moreover, taking the expected second order derivative leads to the Fisher matrix

$$\boldsymbol{F}(\boldsymbol{\theta}) = \sum_{i} \sum_{j \neq i} \boldsymbol{B}_{ij}^{\top} \boldsymbol{B}_{ij} \left[w_{\boldsymbol{\theta}}^{spline} \left(u_i, u_j \right) \cdot \left(1 - w_{\boldsymbol{\theta}}^{spline} \left(u_i, u_j \right) \right) \right]^{-1}.$$

Our intention is to maximize $l(\boldsymbol{\theta})$, which could be simply done by Fisher scoring. The resulting maximizer does, however, not by default lead to a proper estimate, meaning

to fulfill symmetry and boundedness. Furthermore, in case we aim to estimate a graphon with canonical representation, we need to incorporate the constraint from (3). We therefore impose additional (linear) side constraints on $\boldsymbol{\theta}$. Considering the identifiability issue we get the marginal function with (10) through

$$g_{\boldsymbol{\theta}}^{spline}(u) = \left[\boldsymbol{B}(u) \otimes \int_{0}^{1} \boldsymbol{B}(v) \, \mathrm{d}v \right] \boldsymbol{\theta}.$$
 (11)

For normalized B-splines we can easily calculate the integral and for a standardized basis with equidistant knots we obtain

$$\int_{0}^{1} \boldsymbol{B}(v) \, \mathrm{d}v = \left(\int_{0}^{1} B_{1}(v) \, \mathrm{d}v, \int_{0}^{1} B_{2}(v) \, \mathrm{d}v, \dots, \int_{0}^{1} B_{K}(v) \, \mathrm{d}v \right)$$
$$= \underbrace{\frac{1}{K-1} \left(\frac{1}{2}, 1, \dots, 1, \frac{1}{2} \right)}_{=:\boldsymbol{A}}.$$

This allows to rewrite (11) to $g_{\theta}^{spline}(u) = [\mathbf{B}(u) \otimes \mathbf{A}] \boldsymbol{\theta}$. Hence, the marginal function $g_{\theta}^{spline}()$ is also expressed as a linear B-spline and a monotonicity constraint is easily accommodated by postulating monotonicity at the knots τ_1, \ldots, τ_K . That is we need

$$g_{\boldsymbol{\theta}}^{spline}\left(\tau_{l}\right) - g_{\boldsymbol{\theta}}^{spline}\left(\tau_{l-1}\right) \ge 0 \quad \Leftrightarrow \quad \left[\left(\boldsymbol{B}\left(\tau_{l}\right) - \boldsymbol{B}\left(\tau_{l-1}\right)\right) \otimes \boldsymbol{A}\right] \boldsymbol{\theta} \ge 0 \tag{12}$$

for l = 2, ..., K, which is a simple linear constraint on the coefficient vector.

Imposing symmetry on the graphon can also easily be accommodated as linear constraints $\theta_{pq} = \theta_{qp}$ for $p \neq q$. Finally, we need $w_{\theta}^{spline}(,)$ being bounded to [0, 1], which is again a simple linear constraint. All in all, we can write the side constraints as $C\theta \geq 0$ and $D\theta = 0$ for matrices C and D chosen accordingly, where the constraint from (12) for the canonical representation can be added if required. With the above linear constraints and the maximization task of $l(\theta)$ we obtain an (iterated) quadratic programming problem, which can be solved using standard software (see e.g. Andersen et al., 2004 or Turlach and Weingessel, 2013).

3.2.2 Penalized Estimation

Following ideas from the penalized spline estimation (see Eilers and Marx, 1996 or Ruppert et al., 2009) we may additionally impose a penalty on the coefficients to achieve smoothness. This is necessary since we intend to choose K large and unpenalized estimation will lead to wiggled estimates. We refer to Eilers and Marx (1996) for a motivation of penalized spline estimation. To do so, we penalize the difference between "neighbouring" elements of $\boldsymbol{\theta}$ to achieve smoothness. Let therefore

$$\boldsymbol{L} = \begin{pmatrix} 1 & -1 & 0 & \dots & 0 \\ 0 & 1 & -1 & \dots & 0 \\ \vdots & \ddots & & & \vdots \\ 0 & \dots & 0 & 1 & -1 \end{pmatrix} \in \mathbb{R}^{(K-1) \times K}$$

be the first order difference matrix. We then penalize $[\mathbf{L} \otimes \mathbf{I}] \boldsymbol{\theta}$ and $[\mathbf{I} \otimes \mathbf{L}] \boldsymbol{\theta}$, where \mathbf{I} is the identity matrix of appropriate size. This is leading to the penalized likelihood

$$l_{\boldsymbol{P}}(\boldsymbol{\theta}, \lambda) = l(\boldsymbol{\theta}) - \frac{1}{2}\lambda \boldsymbol{\theta}^{\top} \boldsymbol{P} \boldsymbol{\theta},$$

where $\boldsymbol{P} = (\boldsymbol{L} \otimes \boldsymbol{I})^{\top} (\boldsymbol{L} \otimes \boldsymbol{I}) + (\boldsymbol{I} \otimes \boldsymbol{L})^{\top} (\boldsymbol{I} \otimes \boldsymbol{L})$ and λ serves as smoothing parameter. The corresponding penalization score function is given through

$$\boldsymbol{s}_{\boldsymbol{P}}(\boldsymbol{\theta},\lambda) = \boldsymbol{s}(\boldsymbol{\theta}) - \lambda \boldsymbol{P}\boldsymbol{\theta}$$

and the penalized Fisher matrix in the form

$$\boldsymbol{F}_{\boldsymbol{P}}(\boldsymbol{\theta},\lambda) = \boldsymbol{F}(\boldsymbol{\theta}) + \lambda \boldsymbol{P}.$$

The estimate apparently depends on the penalty parameter λ , which is suppressed in the notation. Setting $\lambda \to 0$ gives an unpenalized fit while setting $\lambda \to \infty$ leads to a constant graphon, i.e. an Erdős-Rényi model. The smoothing parameter λ therefore needs to be chosen data driven. We here follow Kauermann et al. (2013) and make use of the Akaike

Information Criterion (AIC) (Hurvich and Tsai, 1989, see also Burnham and Anderson, 2010). To do so, we define the corrected AIC through

$$AIC_c(\lambda) = -2l(\hat{\boldsymbol{\theta}}_{\boldsymbol{P}}) + 2df(\lambda) + \frac{2df(\lambda)(df(\lambda)+1)}{(N(N-1)) - df(\lambda) - 1},$$

where $\hat{\boldsymbol{\theta}}_{\boldsymbol{P}}$ is the penalized parameter estimate and $df(\lambda)$ is the degree of the model, which we define in the usual way as trace of the product of the inverse penalized Fisher matrix and the unpenalized Fisher matrix. To be specific

$$df(\lambda) = tr\left\{ \boldsymbol{F}_{\boldsymbol{P}}^{-1}(\hat{\boldsymbol{\theta}}_{\boldsymbol{P}}, \lambda) \boldsymbol{F}(\hat{\boldsymbol{\theta}}_{\boldsymbol{P}}) \right\}.$$

We subsequently denote with $\hat{w}^{(1)}(,)$ and $\hat{w}^{EM}(,)$ the penalized B-spline estimates of $w_{\theta}^{spline}(,)$ in the first and the final EM step, meaning that $\hat{w}^{(1)}(,)$ and $\hat{w}^{EM}(,)$ are based on $\hat{u}^{(0)}$ and \hat{u}^{EM} , respectively.

As it is well known, the EM algorithm can be trapped at a local maximum of the likelihood (or, in this case, at a local minimum of the corrected AIC). Thus, depending on the specific data situation, it might be recommendable to repeat the algorithm several times to achieve a better fit.

3.3 Information on Ranking

Note that U_i is uniform and it is helpful to order U_i such that $U_{(1)} \leq U_{(2)} \leq \ldots \leq U_{(N)}$. Considering the degree based ordering $\hat{\psi}()$ from (6) as representative, this allows to define the (degree related) ranking density $f_{\hat{\psi}(k)}(u_{\hat{\psi}(k)}|\boldsymbol{y})$. For the sake of simplicity, we subsequently collapse $\hat{\psi}(k)$ to (k) so that henceforth $f_{(k)}(u_{(k)}|\boldsymbol{y})$ and $U_{(k)}$ refer to the node with the kth lowest degree. For other quantities the notation applies accordingly. This ranking density, however, again is difficult or even impossible to calculate analytically. Note that the full posterior density is given through

$$f_{(k)}(u_{(k)}|\boldsymbol{u}_{(-k)},\boldsymbol{y}) \propto \prod_{j \neq k} w(u_{(k)}, u_{(j)})^{y_{(k)(j)}} (1 - w(u_{(k)}, u_{(j)}))^{1 - y_{(k)(j)}},$$
(13)

| ID | Graphon |
|----|--|
| 1 | $w_1(u,v) = 0.8(1-u)(1-v) + 0.85(u \cdot v)$ |
| 2 | $w_{2}(u,v) = 0.5 \cdot \{F_{\mathcal{N}(0,0.25)}(\Phi^{-1}(u)) \cdot F_{\mathcal{N}(0,0.25)}(\Phi^{-1}(v)) + [1 - F_{\mathcal{N}(0,0.25)}(\Phi^{-1}(u))] \cdot [1 - F_{\mathcal{N}(0,0.25)}(\Phi^{-1}(v))]\}$ |

Table 1: Exemplary graphons considered for simulations. For the second graphon specification $\Phi()$ and $F_{\mathcal{N}(0,0.25)}$ denote the cdf's of the normal distributions with parametrization $(\mu, \sigma^2) = (0, 1)$ and (0, 0.25), respectively.

where $y_{(k)(j)}$ is the network entry which refers to the nodes k and j after the permutation according to the degree based ordering. In that regard, the MCMC sequence (after appropriate thinning) provides information about the posterior distribution of U given the network y. Hence, for the marginal posterior distribution of $U_{(k)}$ we can follow a Monte Carlo integration approach, see e.g. Gelfand and Smith (1990, sec. 2.2 and 2.3), and calculate

$$f_{(k)}(u_{(k)}|\boldsymbol{y}) \approx \frac{1}{n} \sum_{s=1}^{n} f_{(k)}(u_{(k)}|\boldsymbol{u}_{(-k)}^{\langle s \cdot N \cdot r \rangle}, \boldsymbol{y}),$$
(14)

where $\boldsymbol{u}_{(-k)}^{\langle t \rangle} = (u_{(1)}^{\langle t \rangle}, \dots, u_{(k-1)}^{\langle t \rangle}, u_{(k+1)}^{\langle t \rangle}, \dots, u_{(N)}^{\langle t \rangle})$ is the *t*th state of the Gibbs sampling sequence without the *k*th component after degree related permutation, $r \in \mathbb{N}$ describes a thinning factor and *n* is the number of MCMC states which are taken into account. For this purpose, the unspecified normalizing constant in (13) can be approximated through a Riemann sum, since we assume w(,) to fulfill certain continuity properties. Considering w(,) as unknown again, we employ $\hat{w}^{(1)}(,)$ or $\hat{w}^{EM}(,)$ as estimate in (13) and obtain $\hat{f}_{(k)}^{(1)}(u_{(k)}|\boldsymbol{y})$ or $\hat{f}_{(k)}^{EM}(u_{(k)}|\boldsymbol{y})$ as estimate of (14), respectively.

4 Simulation Studies

For evaluating our method we at first consider networks generated from a known ground truth. More precisely, for each of the two graphons from Table 1 we simulate networks with dimension N = 500 using the data generating process (4). The first graphon has a canonical representation, while the second is more general and does not fulfill that g() in (3) is strictly increasing. We start with Graphon 1 and demonstrate the benefits of iteratively applying the E- and M-step instead of ordering the nodes just with respect to their degree, i.e. based on $\hat{\psi}()$, and applying the M-step only once. We call the latter a "one-step" estimator. This "one-step" estimate $\hat{w}^{(1)}(,)$ is shown in the top right panel in Figure 2. The fit is apparently not convincing, which we want to make more evident by looking at the ranking information. First, we plot $\hat{u}_i^{(0)}$ as defined in (7) against the true values u_i . There is no concordance visible, which exhibits the problem with a simple ranking of the nodes based on the degree. We run an MCMC sample according to Section 3.1 to obtain information about the ranking, as discussed in Section 3.3. The MCMC sample allows to estimate the posterior distribution of $U_{(k)}$ as proposed in (14), which for three exemplary values of k is plotted in the bottom right plots of Figure 2. The corresponding initial estimates $\hat{u}_{(k)}^{(0)}$ are included as vertical dashed lines and it can be seen that they are not well represented by the posterior distributions. We can conclude that this network, which has been generated from $w_1(,)$, is not suitable to sort by degree for achieving an appropriate estimate of the true node ordering with respect to the data generating model. Hence, also the "one-step" estimation approach results in a very poor fit of the underlying true graphon. We therefore iterate between the E- and the M-step.

The converged EM estimate is shown in the top right plot in Figure 3. Regarding the final EM based estimates \hat{u}_i^{EM} , the comparison with the true values u_i reveals a convincing concordance. We also consider again the posterior distribution of $U_{(k)}$ for the same indices as above (bottom right plots), which indicates a plausible positioning. As a conclusion, the proposed EM algorithm provides appropriate results even if the initial ordering of the nodes by their degree is not adequate. Note that the canonical restriction from (12) has not been included here although $w_1(,)$ has canonical representation, since the shown unrestricted estimate turned out to be similarly good.

Finally, we explore the convergence behavior of the algorithm, which is illustrated in Figure 4 for some values of the graphon at exemplary specific positions. We see that after about 8 to 12 steps a reasonable convergence occurs.



Figure 2: Graphon estimation based on linear B-splines and $\hat{\boldsymbol{u}}^{(0)}$ for Graphon 1 (top left) from Table 1. The graphon estimate $\hat{\boldsymbol{w}}^{(1)}(,)$ is depicted in top right. The plot at bottom left illustrates the estimated $\hat{u}_i^{(0)}$ versus the true simulated u_i . The three lower right plots show the approximated posterior distribution of $U_{(k)}$ (based on the MCMC samples and with respect to the given graphon estimate) for some selected indices. The dashed vertical lines (see also numbers in the boxes) represent the estimates $\hat{u}_{(k)}^{(0)}$.



Figure 3: Graphon estimation for Graphon 1 (top left) from Table 1, based on the EM algorithm using $\hat{\boldsymbol{u}}^{(0)}$ from (7) as initialization and excluding the canonical restriction from (12). The graphon estimate $\hat{\boldsymbol{w}}^{EM}(,)$ is depicted in the top right. The plot at bottom left illustrates the comparison between the estimated \hat{u}_i^{EM} and the true simulated u_i . The three lower right plots show for some selected indices the approximated posterior distribution of $U_{(k)}$ with respect to the graphon estimate in the top right panel. The dashed vertical lines (see also numbers in the boxes) represent the estimates $\hat{u}_{(k)}^{EM}$.



Figure 4: Trajectory of $\hat{w}^{(m)}(u, v)$ for Graphon 1 from Table 1 at exemplary specific positions for the proceeding EM iterations $m = 1, \ldots, 30$.

For a further evaluation, we now have a look at Graphon 2, which does not provide a canonical representation and thus is completely intractable for simple degree based estimation methods. This simulation seeks to demonstrate that our proposed method can also handle more general graphons. The marginal function q() as defined in (3) is constant at 0.25, meaning that the degree is fully uninformative with regard to the node ordering. Nonetheless, we here remain with the degree based initialization as proposed above to show that the EM algorithm is capable of appropriately estimating the underlying structure even under bad initialization. We exclude the side constraint from (12), since we now explicitly want to enable flexible marginal functions. The resulting graphon estimate for a simulated network of dimension N = 500 is depicted in the top right plot in Figure 5. The structure of the true graphon (top left) is fully captured, which is accompanied by a convincing concordance between the estimated \hat{u}_i^{EM} and the true u_i . In accordance with that, the estimates $\hat{u}_{(k)}^{EM}$ are well covered by the corresponding ranking densities, which is illustrated for three selected indices at the stacked plots at the bottom right. To further showcase the performance of our method we repeat the estimation procedure several times based on the same simulated network but with differing random initialization. More precisely, instead of making use of the degree based ordering from (7), we set $\hat{\boldsymbol{u}}^{(0)}$ as a random permutation



Figure 5: Graphon estimation for the non-canonical Graphon 2 (top left) from Table 1, based on the EM algorithm using $\hat{\boldsymbol{u}}^{(0)}$ from (7) as initialization and excluding the canonical restriction from (12). The graphon estimate $\hat{\boldsymbol{w}}^{EM}(,)$ is depicted in the top right. The plot at bottom left illustrates the estimated \hat{u}_i^{EM} versus the true simulated u_i . The three lower right plots show for some selected indices the approximated posterior distribution of $U_{(k)}$ with respect to the graphon estimate in the top right panel. The dashed vertical lines (see also numbers in the boxes) represent the estimates $\hat{u}_{(k)}^{EM}$.



Figure 6: Graphon estimates $\hat{w}^{EM}(,)$ for the non-canonical Graphon 2 from Table 1, based on the EM algorithm using uninformative random initialization. For all repetitions the same simulated network has been used but with independently differing initialization. Beneath are the penalized likelihood and the corrected AIC given accordingly.

of (i/(N+1): i = 1, ..., N). This also exemplary characterizes the appearance of the EM algorithm being trapped at local minima of the AIC. Figure 6 illustrates the estimation results for six such repetitions. Although we start for each run with a completely uninformative random initialization, in most of the final estimates the structure of the original graphon can be fully recognized instantly. Yet, the estimates in the top right and the bottom left panel reveal a somehow differing structure. Having a more precise look, they exhibit the appearance of segment swaps and reversals with respect to the original graphon representation. For instance, at the graphon estimate in the bottom left panel the upper and the lower part of the domain [0, 1] have been swapped and, in addition, the lower part has been reversed. However, as mentioned before, applying permutations to the graphon

function has no effect on the model specification itself. Looking at the resulting AIC values beneath the graphon estimates, we nevertheless see that the smallest values result for smooth graphon estimates, while the occurrence of a "jump" due to only piecewise correct reordering of the nodes lead to increased AIC values. But still, the structure of the true graphon model is always well captured, albeit possibly in a different form of representation. Summarizing over all the estimation results, our algorithm yields very good estimates for Graphon 2, although an uninformative initial node ordering has been applied.

The example illustrates that a random reordering of the nodes for the initial E step can be applied in combination with the AIC in order to obtain an optimal smooth graphon estimate.

5 Real-World Data Examples

We complete the paper by analyzing real-world network data. To do so, we consider networks from three different domains, namely from sociology, political science, and neuroscience. For the EM based estimation we use degree ordering as initialization where the structure seems appropriate. Otherwise we apply random initialization and select the best outcome over several repetitions.

5.1 Facebook Network

A very common application and one of the roots of network sciences are friendship networks. We here consider a Facebook ego network which has been collected by McAuley and Leskovec (2012) and is available on the Stanford Large Network Dataset Collection (Leskovec and Krevl, 2014). This ego network with 333 actors is plotted in Figure 7 with two different colorings. In the left panel the illustrated ordering is based on the degree, while on the right it refers to the final EM result. Apparently, the ordering given by the final EM result seems much more appropriate with respect to the network structure. Thus, the iterative approach improves the initial degree ordering markedly. Moreover, the inherent structure of the network on the right in Figure 7 can to some extent be recognized instantly in the corresponding final graphon estimate $\hat{w}^{EM}(,)$, which is shown on the left in Figure 8. For example, the bundle of nodes in the center of the lower network part with roughly $\hat{u}_i^{EM} \in [0.65, 1]$ and the connectivity pattern among themselves is captured in the graphon estimate in the intense section on the bottom right. In accordance with this, also the estimates $\hat{u}_{(k)}^{EM}$ for some selected indices are adequately represented by the corresponding posterior distributions of $U_{(k)}$, which underlines the appropriateness of the graphon estimate $\hat{w}^{EM}(,)$.

5.2 Military Alliance Network

As second real-world network example, we consider strong military alliances among the world's nations. For that purpose, we use data from the Alliance Treaty Obligations and Provisions project (Leeds et al., 2002), which provide information about all kind of military alliance agreements over an extensive period. For a suitably reduced network we here define the presence of a strong military alliance when states have entered into an offensive or defensive pact, meaning when they have signed a treaty which forces the one country to intervene by military active support if the other country comes into a conflict with offensive or defensive actions, respectively. Furthermore, we truncate the data to agreements that were in force in 2016 as the most recent available year. The best final estimation results of the EM approach over several repetitions with random initial node ordering are illustrated in Figure 9. The graphon estimate in the top right panel exhibits a very pronounced assortative structure, meaning that predominantly pairs of nodes are linked whose latent quantities are close. The final node ordering, which is visualized by the coloring in the top left network, appears reasonable with respect to the network structure. Transferring this coloring to the world map reveals a strong conformity between the geographic closeness and the closeness of the latent quantities. Together with the assortative structure, this means that countries which are geographically close form similar military alliances and further are more likely to be allied with each other.



Figure 7: Facebook ego network with node coloring referring to $\hat{u}_i^{(0)}$ (left) and \hat{u}_i^{EM} (right).



Figure 8: Graphon estimation for the Facebook ego network. The graphon estimate $\hat{w}^{EM}(,)$ (in log scale) is depicted on the left. The three right plots show the approximated posterior distribution of $U_{(k)}$ (with respect to the given graphon estimate) for some selected indices. The dashed vertical lines (see also numbers in the boxes) represent the estimates $\hat{u}_{(k)}^{EM}$.



Figure 9: Graphon estimation for the military alliance network (top left, with coloring referring to $\hat{u}_i^{EM} \in [0,1]$). The graphon estimate $\hat{w}^{EM}(,)$ (in log scale) is depicted in the top right. The lower plot shows the world map, where the colors also refer to \hat{u}_i^{EM} . As an isolated group, China, Cuba, and North Korea (colored in pink) have not been included in the estimation routine. Countries which have no strong agreement with any other country and hence do not appear in the data set are colored in gray.

5.3 Human Brain functional Coactivation Network

To conclude the real-world examples, we consider a human brain functional network which has been constructed from a meta-analysis by Crossley et al. (2013) and which is available in the Brain Connectivity Toolbox (Rubinov and Sporns, 2010a, see also Rubinov and Sporns, 2010b for detailed description). Their weighted network matrix represents the "estimated [...] similarity (Jaccard index) of the activation patterns across experimental tasks between each pair of 638 brain regions" (Crossley et al., 2013). To obtain a binary adjacency matrix out of it, we simply apply a threshold of above zero, meaning that we include a link between each pair of brain regions which is at least during one task activated simultaneously. The resulting network with a density of approximately 9.2% is depicted in the top left panel in Figure 10, where the coloring refers to \hat{u}_i^{EM} . For the evaluation, we again consider the best outcome of the EM algorithm over several repetitions with random initialization. The graphon estimate in the top right plot also reveals an assortative structure but, in addition, exhibits a conspicuous pattern of functional coactivations of some segments, which are separated with respect to the latent dimension. Regarding the spatial positions of the brain regions, the lower three plots show a strong relation between closeness in anatomical space and closeness of the \hat{u}_i^{EM} . Nevertheless, there also seem to be areas which have a similar color pattern (indicating a similar coactivation pattern) but are spatially separated. Indeed, the interaction and coactivation of distant brain areas is a familiar phenomena in neuroscience and can also be seen in Crossley et al. (2013), who pursue a different modeling strategy. Overall, we can demonstrate that the EM approach for estimating smooth graphons provides additional insight into network structures.

6 Discussion

The paper proposes a novel estimation routine for graphon estimation which explicitly takes the variability of ordering the nodes into account. The proposed semi-parametric estimation based on (linear) B-splines allows to incorporate relevant properties in the estimation and also uniqueness restrictions if required. The Bayesian approach relying on Gibbs sampling



Figure 10: Graphon estimation for the human brain functional coactivation network (top left, with coloring referring to $\hat{u}_i^{EM} \in [0, 1]$). The graphon estimate $\hat{w}^{EM}(,)$ (in log scale) is depicted in the top right. The lower three plots show the local positions of the human brain regions in anatomical space in side view (left), front view (middle), and top view (right), where the colors also refer to \hat{u}_i^{EM} .

illuminates the uncertainty about the degree and its distribution. Both steps combined give an EM type algorithm which allows for flexible graphon estimation even in large networks. The approach outperforms available routines in two aspects. First, the B-spline estimate can guarantee a smooth and continuous outcome and also a unique representation of the graphon if required. Secondly, based on the Bayesian formulation and the EM algorithm one can assess the amount of uncertainty for ordering the nodes based on their degree or based on any other strategy.

The proposed approach can also be used in other related models like stochastic block models (SBM), where we assume that nodes cluster and form within and between the clusters simple Erdős-Rényi models. Apparently, SBMs do not have a smooth underlying graphon structure so that one needs to adjust and further develop the results in this paper.

Overall, graphon estimation also provides an interesting tool for network visualizations, as demonstrated in the examples. In this respect, it is more than a modeling exercise but may also serve as tool for explanatory network data analysis.

Acknowledgments

The project was partially supported by the European Cooperation in Science and Technology [COST Action CA15109 (COSTNET)]. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. Declarations of interest: none.

References

- Airoldi, E. M., T. B. Costa, and S. H. Chan (2013). Stochastic blockmodel approximation of a graphon: Theory and consistent estimation. Advances in Neural Information Processing Systems, 692–700.
- Andersen, M. S., J. Dahl, and L. Vandenberghe (2004). CvxOpt: Open source software for convex optimization (Python). *https://cvxopt.org, version 1.2.0* [online; accessed 04-20-2018].
- Bickel, P. J. and A. Chen (2009). A nonparametric view of network models and Newman–Girvan and other modularities. *Proceedings of the National Academy of Sciences 106*(50).
- Bickel, P. J., A. Chen, and E. Levina (2011). The method of moments and degree distributions for network models. *The Annals of Statistics* 39(5), 2280–2301.
- Borgs, C., L. Chayes, L. Locász, V. T. Sós, and K. Vesztergombi (2008). Convergent sequences of dense graph I: Subgraph frequencies, metric properties and testing. Advances in Mathematics 219, 1801–1851.

- Burnham, K. P. and D. R. Anderson (2010). Model Selection and Multimodel Inference: a practical Information-theoretic Approach. Springer Science & Business Media, New York.
- Chan, S. H. and E. M. Airoldi (2014). A consistent histogram estimator for exchangeable graph models. *International Conference on Machine Learning*, 208–216.
- Chatterjee, S. (2015). Matrix estimation by universal singular value thresholding. The Annals of Statistics 43(1), 177–214.
- Choi, D. (2017). Co-clustering of nonsmooth graphons. *The Annals of Statistics* 45(4), 1488–1515.
- Choi, D. and P. J. Wolfe (2014). Co-clustering separately exchangeable network data. *The* Annals of Statistics 42(1), 29–63.
- Crossley, N. A., A. Mechelli, P. E. Vértes, T. T. Winton-Brown, A. X. Patel, C. E. Ginestet, P. McGuire, and E. T. Bullmore (2013). Cognitive relevance of the community structure of the human brain functional coactivation network. *Proceedings of the National Academy* of Sciences 110(28), 11583–11588.
- Diaconis, P. and S. Chatterjee (2013). Estimating and understanding exponential random graph models. *The Annals of Statistics* 41(5), 2428–2461.
- Diaconis, P. and S. Janson (2008). Graph limits and exchangeable random graphs. *Rendiconti di Matematica e delle sui Applicazioni 28*, 33–61.
- Eilers, P. H. C. and B. D. Marx (1996). Flexible smoothing with B-splines and penalties. Statistical Science 11(2), 89–102.
- Fienberg, S. E. (2012). A brief history of statistical models for network analysis and open challenges. Journal of Computational and Graphical Statistics 21(4), 825–839.
- Frank, O. and D. Strauss (1986). Markov graphs. Journal of the American Statistical Association 81(395), 832–842.

- Gao, C., Y. Lu, Z. Ma, and H. H. Zhou (2016). Optimal estimation and completion of matrices with biclustering structures. *The Journal of Machine Learning Research* 17(1), 5602–5630.
- Gao, C., Y. Lu, and H. H. Zhou (2015a). Rate-optimal graphon estimation. The Annals of Statistics 43(6), 2624–2652.
- Gao, C., A. W. van der Vaart, and H. H. Zhou (2015b). A general framework for bayes structured linear models. *arXiv preprint arXiv:1506.02174*.
- Gelfand, A. E. and A. F. M. Smith (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association* 85(410), 398–409.
- Goldenberg, A., A. X. Zheng, S. E. Fienberg, and E. M. Airoldi (2010). A survey of statistical network models. *Foundation and Trends in Machine Learning* 2, 129–233.
- He, R. and T. Zheng (2015). GLMLE: graph-limit enabled fast computation for fitting exponential random graph models to large social networks. Social Network Analysis and Mining 5(1), 8.
- Hoff, P. D., A. E. Raftery, and M. S. Handcock (2002). Latent space approaches to social network analysis. *Journal of the american Statistical association* 97(460), 1090–1098.
- Holland, P. W., K. B. Laskey, and S. Leinhardt (1983). Stochastic blockmodels: First steps. Social networks 5(2), 109–137.
- Hunter, D. R., P. N. Krivitsky, and M. Schweinberger (2012). Computational statistical methods for social network analysis. *Journal of Computational and Graphical Statistics* 21(4), 856–882.
- Hurvich, C. M. and C.-L. Tsai (1989). Regression and time series model selection in small samples. *Biometrika* 76(2), 297–307.
- Kauermann, G., C. Schellhase, and D. Ruppert (2013). Flexible copula density estimation with penalized hierarchical B-splines. Scandinavian Journal of Statistics 40(4), 685–705.

- Klopp, O., A. B. Tsybakov, and N. Verzelen (2017). Oracle inequalities for network models and sparse graphon estimation. *The Annals of Statistics* 45(1), 316–354.
- Kolaczyk, E. D. (2009). Statistical Analysis of Network Data. Springer, New York.
- Kolaczyk, E. D. (2017). Topics at the Frotier of Statistics and Network Analysis. Cambridge University Press.
- Kolaczyk, E. D. and G. Csárdi (2014). *Statistical Analysis of Network Data with R.* Springer, New York.
- Leeds, B. A., J. M. Ritter, S. Mc Laughlin Mitchell, and A. G. Long (2002). Alliance treaty obligations and provisions, 1815-1944. *International Interactions* 28(3), 237–260.
- Leskovec, J. and A. Krevl (2014). SNAP Datasets: Stanford Large Network Dataset Collection. *http://snap.stanford.edu/data* [online; accessed 06-14-2018].
- Lovász, L. (2012). Large Networks and Graph Limits. American Mathematical Soc.
- Lusher, D., J. Koskinen, and G. Robins (2013). Exponential Random Graph Models for social Networks. Cambridge University Press.
- McAuley, J. and J. Leskovec (2012). Learning to discover social circles in ego networks. Advances in Neural Information Processing Systems 25, 539–547.
- Olhede, S. C. and P. J. Wolfe (2014). Network histograms and universality of blockmodel approximation. *Proceedings of the National Academy of Sciences* 111(41), 14722–14727.
- Orbanz, P. and D. M. Roy (2014). Bayesian models of graphs, arrays and other exchangeable random structures. *IEEE transactions on pattern analysis and machine intelli*gence 37(2), 437–461.
- Rubinov, M. and O. Sporns (2010a). Brain Connectivity Toolbox (MATLAB). https://sites.google.com/site/bctnet [online; accessed 10-15-2019].
- Rubinov, M. and O. Sporns (2010b). Complex network measures of brain connectivity: uses and interpretations. *Neuroimage* 52(3), 1059–1069.

- Ruppert, D., M. P. Wand, and R. J. Carroll (2003). Semiparametric Regression. Cambridge University Press.
- Ruppert, D., M. P. Wand, and R. J. Carroll (2009). Semiparametric regression during 2003–2007. Electronic Journal of Statistics 3, 1193.
- Salter-Townshend, M., A. White, I. Gollini, and T. B. Murphy (2012). Review of statistical network analysis: models, algorithms and software. *Statistical Analysis and Data Mining* 5(4), 243–264.
- Snijders, T. A. B., P. E. Pattison, G. L. Robins, and M. S. Handcock (2006). New specifications for exponential random graph models. *Sociological Methodology* (1), 99–153.
- Turlach, B. A. and A. Weingessel (2013). quadprog: Functions to solve quadratic programming problems (R). https://CRAN.R-project.org/package=quadprog, version 1.5-5 [online; accessed 04-26-2018].
- Wolfe, P. J. and S. C. Olhede (2013). Nonparametric graphon estimation. arXiv preprint arXiv:1309.5936.
- Wood, S. N. (2017). P-splines with derivative based penalties and tensor product smoothing of unevenly distributed data. *Statistics and Computing* 27(4), 985–989.
- Xu, J. (2017). Rates of convergence of spectral methods for graphon estimation. arXiv preprint arXiv:1709.03183.
- Yang, J. J., Q. Han, and E. M. Airoldi (2014). Nonparametric estimation and testing of exchangeable graph models. Artificial Intelligence and Statistics, 1060–1067.
- You, K. (2020). graphon: A collection of graphon estimation methods (R). https://CRAN.R-project.org/package=graphon, version 0.3.4 [online; accessed 06-13-2020].
- Zhang, Y., E. Levina, and J. Zhu (2015). Estimating network edge probabilities by neighborhood smoothing. arXiv preprint arXiv:1509.08588.