

How Many Communities Are There?

D. Franco Saldaña^a, Yi Yu^b, and Yang Feng^a

^aDepartment of Statistics, Columbia University, New York, New York; ^bStatistical Laboratory, University of Cambridge, Cambridge, UK

ABSTRACT

Stochastic blockmodels and variants thereof are among the most widely used approaches to community detection for social networks and relational data. A stochastic blockmodel partitions the nodes of a network into disjoint sets, called communities. The approach is inherently related to clustering with mixture models; and raises a similar model selection problem for the number of communities. The Bayesian information criterion (BIC) is a popular solution, however, for stochastic blockmodels, the conditional independence assumption given the communities of the endpoints among different edges is usually violated in practice. In this regard, we propose composite likelihood BIC (CL-BIC) to select the number of communities, and we show it is robust against possible misspecifications in the underlying stochastic blockmodel assumptions. We derive the requisite methodology and illustrate the approach using both simulated and real data. Supplementary materials containing the relevant computer code are available online.

ARTICLE HISTORY

Received December 2014
Revised June 2015

KEYWORDS

Community detection;
Composite likelihood;
Degree-corrected stochastic
blockmodel; Model selection;
Spectral clustering;
Stochastic blockmodel

1. Introduction

Enormous network datasets are being generated and analyzed along with an increasing interest from researchers in studying the underlying structures of a complex networked world. The potential benefits span traditional scientific fields such as epidemiology and physics, but also emerging industries, especially large-scale Internet companies. Among a variety of interesting problems arising with network data, in this article, we focus on community detection in undirected networks $G := (V, E)$, where V and E are the sets of nodes and edges, respectively. In this framework, the community detection problem can be formulated as finding the true disjoint partition of $V = V_1 \sqcup \dots \sqcup V_K$, where K is the number of communities. Although it is difficult to give a rigorous definition, communities are often regarded as tightly knit groups of nodes that are loosely connected between themselves.


The community detection problem has close connections with graph partitioning, which could be traced back to Euler, while it has its own characteristics due to the concrete physical meanings from the underlying dataset (Newman and Girvan 2004). Over the last decade, there has been a considerable amount of work on it, including minimizing ratio cut (Wei and Cheng 1989), minimizing normalized cut (Shi and Malik 2000), maximizing modularity (Newman and Girvan 2004), hierarchical clustering (Newman 2004), and edge-removal methods (Newman and Girvan 2004), to name a few. Among all the progress made by peer researchers, spectral clustering (Donath and Hoffman 1973) based on stochastic blockmodels (Holland, Laskey, and Leinhardt 1983) debuted and soon gained a majority of attention. We refer the interested readers to Spielmat and Teng (1996) and Goldenberg et al. (2010) as comprehensive reviews

on the history of spectral clustering and stochastic blockmodels, respectively.

Compared to the amount of work on spectral clustering or stochastic blockmodels, to the best of our knowledge, there is little work on the selection of the community number K . In most of the previously mentioned community detection methods, the number of communities is generally input as a pre-specified quantity. For the literature addressing the problem of selecting K , besides the block-wise edge splitting method of Chen and Lei (2014), a common practice is to use Bayesian information criterion (BIC) type criteria (Airoldi et al. 2008; Daudin, Picard, and Robin 2008) or a variational Bayes approach (Hunter, Krivitsky, and Schweinberger 2012; Latouche, Birmelé, and Ambroise 2012). An inherently related problem is that of selecting the number of components in mixture models, where the birth-and-death point process of Stephens (2000) and the allocation sampler of Nobile and Fearnside (2007) provide two fully Bayesian approaches in the case where K is finite but unknown. Based on the allocation sampler, McDaid et al. (2013) proposed an efficient Bayesian clustering algorithm that directly estimates the number of communities in stochastic blockmodels, and that exhibits similar results to the variational Bayes approach of Latouche, Birmelé, and Ambroise (2012). Nonparametric Bayesian methods based on Dirichlet process mixtures (Ferguson 1973) have also been used to estimate the number of components in this finite but unknown K setting (Fearnhead 2004), although the inconsistency of this approach has been recently shown by Miller and Harrison (2014). This community or mixture component number K , as a vital part of model selection procedures, highly depends on the model assumptions. For instance, the famous stochastic blockmodel has undesirable

CONTACT Yang Feng  yangfeng@stat.columbia.edu

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/r/JCGS.

 Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/JCGS.

© 2017 American Statistical Association, Institute of Mathematical Statistics, and Interface Foundation of North America

restrictive assumptions in the form of independent Bernoulli observations when the community assignments are known.

In this article, we study the community number selection problem with robustness consideration against model misspecification in the stochastic blockmodel and its variants. Our motivation is that, the conditional independence assumption among edges, when the communities of their endpoints are given, is usually violated in real applications. In addition, we do not restrict our interest only to exchangeable graphs. Using techniques from the composite likelihood paradigm (Lindsay 1988), we develop a composite likelihood BIC (CL-BIC) approach (Gao and Song 2010) for selecting the community number in the situation where assumed independencies in the stochastic blockmodel and other exchangeable graph models do not hold. The procedure is tested on simulated and real data, and is shown to outperform two competitors—traditional BIC and the variational Bayes criterion of Latouche, Birmelé, and Ambroise (2012), in terms of model selection consistency.

The rest of the article is organized as follows. The background for stochastic blockmodels and spectral clustering is introduced in Section 2, and the proposed CL-BIC methodology is developed in Section 3. In Section 4, several simulation examples as well as two real datasets are analyzed. The article is concluded with a short discussion in Section 5.

2. Background

First, we would like to introduce some notation. For an N -node undirected, simple, and connected network G , its symmetric adjacency matrix A is defined as $A_{ij} := 1$ if (i, j) is an element in E , and $A_{ij} := 0$ otherwise. The diagonal $\{A_{ii}\}_{i=1}^N$ is fixed to zero (i.e., self-edges are not allowed). Moreover, D and L denote the degree matrix and Laplacian matrix, respectively. Here, $D_{ii} := d_i$, and $D_{ij} := 0$ for $i \neq j$, where d_i is the degree of node i , that is, the number of edges with endpoint node i ; and $L := D^{-1/2}AD^{-1/2}$. As isolated nodes are discarded, $D^{-1/2}$ is well-defined.

2.1 Stochastic Blockmodels

2.1.1 Standard Stochastic Blockmodel

Stochastic blockmodels were first introduced in Holland, Laskey, and Leinhardt (1983). They posited independent Bernoulli random variables $\{A_{ij}\}_{1 \leq i < j \leq N}$ with success probabilities $\{P_{ij}\}$ that depend on the communities of their endpoints i and j . Consequently, all edges are conditionally independent given the corresponding communities. Moreover, each node is associated with one and only one community, with label Z_i , where $Z_i \in \{1, \dots, K\}$. Following Rohe, Chatterjee, and Yu (2011) and Choi, Wolfe, and Airoldi (2012), throughout this article we assume each Z_i is fixed and unknown, thus yielding $\mathbb{P}(A_{ij} = 1; Z_i = z_i, Z_j = z_j) = \theta_{z_i z_j}$. Treating the node assignments Z_1, \dots, Z_N as latent random variables is another popular approach in the community detection literature, and various methods including the variational Bayes criterion of Latouche, Birmelé, and Ambroise (2012) and the belief propagation algorithm of Decelle et al. (2011) efficiently approximate the corresponding *observed-data* log-likelihood of the stochastic

blockmodel, without having to add K^N multinomial terms accounting for all possible label assignments.

For $\theta := (\theta_{ab}; 1 \leq a \leq b \leq K)'$ and for any fixed community assignment $\mathbf{z} \in \{1, \dots, K\}^N$, the log-likelihood under the standard *stochastic blockmodel* (SBM) is given as

$$\ell(\theta; \mathbf{A}) := \sum_{i < j} [A_{ij} \log \theta_{z_i z_j} + (1 - A_{ij}) \log(1 - \theta_{z_i z_j})]. \quad (1)$$

For the remainder of the article, denote N_a as the size of community a , and n_{ab} as the maximum number of possible edges between communities a and b , that is, $n_{ab} := N_a N_b$ for $a \neq b$, and $n_{aa} := N_a(N_a - 1)/2$. Also, let $m_{ab} := \sum_{i < j} A_{ij} \mathbf{1}\{z_i = a, z_j = b\}$, and $\hat{\theta}_{ab} := m_{ab}/n_{ab}$ be the maximum likelihood estimator (MLE) of θ_{ab} in (1).

Under this framework, Choi, Wolfe, and Airoldi (2012) showed that the fraction of misclustered nodes converges in probability to zero under maximum likelihood fitting when K is allowed to grow no faster than \sqrt{N} . By means of a *regularized* maximum likelihood estimation approach, Rohe, Qin, and Fan (2014) further proved that this *weak* convergence can be achieved for $K = \mathcal{O}(N/\log^5 N)$.

2.1.2 Degree-Corrected Stochastic Blockmodel

Heteroscedasticity of node degrees within communities is often observed in real-world networks. To tackle this problem, Karrer and Newman (2011) proposed the *degree-corrected blockmodel* (DCBM), in which the success probabilities $\{P_{ij}\}$ are also functions of individual effects. To be more precise, the DCBM assumes that $\mathbb{P}(A_{ij} = 1; Z_i = z_i, Z_j = z_j) = \omega_i \omega_j \theta_{z_i z_j}$, where $\omega := (\omega_1, \dots, \omega_n)'$ are individual effect parameters satisfying the identifiability constraint $\sum_i \omega_i \mathbf{1}\{z_i = a\} = 1$ for each community $1 \leq a \leq K$.

To simplify technical derivations, Karrer and Newman (2011) allowed networks to contain both multi-edges and self-edges. Thus, they assumed the random variables $\{A_{ij}\}_{1 \leq i \leq j \leq N}$ to be independent Poisson, with the previously defined success probabilities $\{P_{ij}\}$ of an edge between vertices i and j replaced by the expected number of such edges. Under this framework, and for any fixed community assignment $\mathbf{z} \in \{1, \dots, K\}^N$, Karrer and Newman (2011) arrived at the log-likelihood $\ell(\theta, \omega; \mathbf{A})$ of observing the adjacency matrix $\mathbf{A} = (A_{ij})$ under the DCBM,

$$\ell(\theta, \omega; \mathbf{A}) := 2 \sum_i d_i \log \omega_i + \sum_{a,b} (m_{ab} \log \theta_{ab} - \theta_{ab}). \quad (2)$$

After allowing for the identifiability constraint on ω , the MLEs of the parameters θ_{ab} and ω_i are given by $\hat{\theta}_{ab} := m_{ab}$ and $\hat{\omega}_i := d_i / \sum_{j: z_j = z_i} d_j$, respectively.

As mentioned in Zhao, Levina, and Zhu (2012), there is no practical difference in performance between the log-likelihood (2) and its slightly more elaborate version based on the true Bernoulli observations. The reason is that the Bernoulli distribution with a small mean is well approximated by the Poisson distribution, and the sparser the network is, the better the approximation works (Perry and Wolfe 2012).

2.1.3 Mixed Membership Stochastic Blockmodel

As a methodological extension in which nodes are allowed to belong to more than one community, Airoldi et al. (2008) proposed the *mixed membership stochastic blockmodel* (MMB) for directed relational data $\{A_{ij}\}_{1 \leq i, j \leq N}$. For instance, when a social actor interacts with its different neighbors, an array of different social contexts may be taking place and thus the actor may be taking on different latent roles.

The model assumes the observed network is generated according to node-specific distributions of community membership and edge-specific indicator vectors denoting membership in one of the K communities. More specifically, each vertex i is associated with a randomly drawn vector $\vec{\pi}_i$, with π_{ia} denoting the probability of node i belonging to community a . Additionally, let the indicator vector $\vec{z}_{i \rightarrow j}$ denote the community membership of node i when he sends a message to node j , and $\vec{z}_{i \leftarrow j}$ denote the community membership of node j when he receives a message from node i . If, to account for the asymmetric interactions, we denote by $\theta := (\theta_{ab})$ the $K \times K$ matrix where θ_{ab} represents the probability of having an edge from a social actor in community a to a social actor in community b , the MMB posits that the $\{A_{ij}\}_{1 \leq i, j \leq N}$ are drawn from the following generative process.

- For each node $i \in V$:
 - Draw a K -dimensional mixed membership vector $\vec{\pi}_i \sim \text{Dirichlet}(\alpha)$, with the vector $\alpha = (\alpha_1, \dots, \alpha_K)'$ being a hyperparameter.
- For each possible edge variable A_{ij} :
 - Draw membership indicator vector for the initiator $\vec{z}_{i \rightarrow j} \sim \text{Multinomial}(\vec{\pi}_i)$.
 - Draw membership indicator vector for the receiver $\vec{z}_{i \leftarrow j} \sim \text{Multinomial}(\vec{\pi}_j)$.
 - Sample the interaction $A_{ij} \sim \text{Bernoulli}(\vec{z}_{i \rightarrow j}' \theta \vec{z}_{i \leftarrow j})$.

Upon defining the set of mixed membership vectors $\mathbf{\Pi} := \{\vec{\pi}_i : i \in V\}$ and the sets of membership indicator vectors $\mathbf{Z}_{\rightarrow} := \{\vec{z}_{i \rightarrow j} : i, j \in V\}$ and $\mathbf{Z}_{\leftarrow} := \{\vec{z}_{i \leftarrow j} : i, j \in V\}$, following Airoldi et al. (2008), we obtain the complete data log-likelihood of the *hyperparameters* $\{\theta, \alpha\}$ as

$$\begin{aligned} \ell(\theta, \alpha; \mathbf{A}, \mathbf{\Pi}, \mathbf{Z}_{\rightarrow}, \mathbf{Z}_{\leftarrow}) : \\ = \sum_{i,j} [A_{ij} \log(\vec{z}_{i \rightarrow j}' \theta \vec{z}_{i \leftarrow j}) + (1 - A_{ij}) \log(1 - \vec{z}_{i \rightarrow j}' \theta \vec{z}_{i \leftarrow j})] \\ + N \left(\log \Gamma \left(\sum_a \alpha_a \right) - \sum_a \log \Gamma(\alpha_a) \right) \\ + \sum_i \sum_a (\alpha_a - 1) \log \pi_{ia} + \text{const}, \end{aligned} \quad (3)$$

where \mathbf{A} corresponds to the observed data and $\{\mathbf{\Pi}, \mathbf{Z}_{\rightarrow}, \mathbf{Z}_{\leftarrow}\}$ are the latent variables.

To carry out posterior inference of the latent variables given the observations \mathbf{A} , Airoldi et al. (2008) proposed an efficient coordinate ascent algorithm based on a variational approximation to the true posterior. Therefore, one can compute expected posterior mixed membership vectors and posterior membership indicator vectors. We refer interested readers to Section 3 in Airoldi et al. (2008) for further details.

Consequently, following the same profile likelihood approach, for any fixed set $\{\mathbf{\Pi}, \mathbf{Z}_{\rightarrow}, \mathbf{Z}_{\leftarrow}\}$, the MLE of θ_{ab} is

given by

$$\hat{\theta}_{ab} := \sum_{i,j} A_{ij} \cdot \vec{z}_{i \rightarrow j, a} \vec{z}_{i \leftarrow j, b} / \sum_{i,j} \vec{z}_{i \rightarrow j, a} \vec{z}_{i \leftarrow j, b}. \quad (4)$$

As the MLE of α_a does not admit a closed form, Minka (2000) proposed an efficient Newton–Raphson procedure for obtaining parameter estimates in Dirichlet models, where the gradient and Hessian matrix of the complete data log-likelihood (3) with respect to α are

$$\begin{aligned} \frac{\partial \ell(\theta, \alpha; \mathbf{A})}{\partial \alpha_a} &= N \left(\Psi \left(\sum_a \alpha_a \right) - \Psi(\alpha_a) \right) + \sum_i \sum_a \log \pi_{ia} \\ \frac{\partial^2 \ell(\theta, \alpha; \mathbf{A})}{\partial \alpha_a \partial \alpha_b} &= N \left(\Psi' \left(\sum_a \alpha_a \right) - \Psi'(\alpha_a) \mathbf{1}\{a = b\} \right), \end{aligned} \quad (5)$$

and Ψ is known as the digamma function (i.e., the logarithmic derivative of the gamma function).

2.2 Spectral Clustering and SCORE

Although there is a parametric framework for the standard stochastic blockmodel, considering the computational burden, it is intractable to directly estimate both parameters θ and \mathbf{z} based on exact maximization of the log-likelihood (1). Researchers have instead resorted to spectral clustering as a computationally feasible algorithm. For comprehensive reviews, we refer interested readers to von Luxburg (2007) and Rohe, Chatterjee, and Yu (2011), in which the authors proved the consistency of spectral clustering in the standard stochastic blockmodel under proper conditions imposed on the density of the network and the eigen-structure of the Laplacian matrix. The algorithm finds the eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_K$ associated with the K eigenvalues of \mathbf{L} that are largest in magnitude, forming an $N \times K$ matrix $\mathbf{U} := (\mathbf{u}_1, \dots, \mathbf{u}_K)$, and then applies the K -means algorithm to the rows of \mathbf{U} .

Similarly, Jin (2015) proposed a variant of spectral clustering for the DCBM, called Spectral Clustering On Ratios-of-Eigenvectors (SCORE). Instead of using the Laplacian matrix \mathbf{L} , SCORE collects the eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_K$ associated with the K eigenvalues of \mathbf{A} that are largest in magnitude, and then forms the $N \times K$ matrix $\mathbf{V} := (\mathbf{1}, \mathbf{v}_2/\mathbf{v}_1, \dots, \mathbf{v}_K/\mathbf{v}_1)$, where the division operator is taken entry-wise, that is, for vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, with $\prod_{\ell=1}^n b_\ell \neq 0$, $\mathbf{a}/\mathbf{b} := (a_1/b_1, \dots, a_n/b_n)'$. SCORE then applies the K -means algorithm to the rows of \mathbf{V} . The corresponding consistency results for the DCBM are also provided in Jin (2015).

3. Model Selection for the Number of Communities

3.1 Motivation

In much of the previous work, for example, Airoldi et al. (2008), Daudin, Picard, and Robin (2008), and Handcock, Raftery, and Tantrum (2007), researchers have used a BIC-penalized version of the log-likelihood (1) to choose the community number K . However, we are aware of the possible misspecifications in the underlying stochastic blockmodel assumptions and in the loss of

precision from the computational relaxation brought in by spectral clustering.

First, in network data, edges are not necessarily independent if only the communities of their endpoints are given. For instance, if two different edges A_{ij} and A_{il} have mutual endpoint i , it is highly likely that they are dependent even given the community labels of their endpoints. This misspecification problem exists in both the standard stochastic blockmodel and its variants, such as DCBM (Karrer and Newman 2011) and MMB (Airoldi et al. 2008). Second, as previously mentioned, spectral clustering is a feasible relaxation, but the loss of precision is inevitable. Several examples of this can be found in Guattery and Miller (1998). Whence, we resort to introducing CL-BIC with the concern of robustness against misspecifications in the underlying stochastic blockmodel.

We would like to emphasize that CL-BIC is not a new community detection method. Instead, under the SBM, DCBM, or MMB assumptions, it can be combined with existing community detection methods to choose the true community number.

3.2 Composite Likelihood Inference

The CL-BIC approach extends the concepts and theory of conventional BIC on likelihoods to the composite likelihood paradigm (Lindsay 1988; Varin, Reid, and Firth 2011). Composite likelihood aims at a relaxation of the computational complexity of statistical inference based on exact likelihoods. For instance, when the dependence structure for relational data is too complicated to implement, a working independence assumption can effectively recover some properties of the usual maximum likelihood estimators (Cox and Reid 2004; Varin, Reid, and Firth 2011). However, under this misspecification framework, the asymptotic variance of the resulting estimators is usually underestimated as the Fisher information. Composite marginal likelihoods (also known as *independence likelihoods*) have the same formula as conventional likelihoods in terms of being a product of marginal densities (Varin 2008), while statistical inference based on them can capture this loss of variance. Consequently, to pursue the “true” model, CL-BIC penalizes the number of parameters more than what BIC does for dependent relational data.

Before going into details, we would like to give the rationale of using stochastic blockmodels under a misspecification framework. To estimate the true joint density g of $\{A_{ij}\}_{1 \leq i < j \leq N}$, we consider the stochastic blockmodel family $\mathcal{P} = \{p_\theta : \theta \in \Theta\}$, where $\Theta = [0, 1]^{K(K+1)/2}$ for the standard stochastic blockmodel, and $\Theta = [0, 1]^{K(K+1)/2+N}$ for DCBM. The true joint density g may or may not belong to \mathcal{P} , which is a parametric family imposing independence among the $\{A_{ij}\}_{i < j}$ when only the communities of the endpoints are given.

Due to the difficulty in specifying the full, highly structured $\binom{N}{2}$ -dimensional density g , while having access to the univariate densities $p_{ij}(\cdot; \theta)$ of A_{ij} under the blockmodel family \mathcal{P} , the composite marginal likelihood paradigm compounds the first-order log-likelihood contributions to form the composite log-likelihood

$$\text{cl}(\theta; \mathbf{A}) := \sum_{i < j} \log p_{ij}(A_{ij}; \theta), \quad (6)$$

where $\text{cl}(\cdot; \mathbf{A})$ corresponds to (1) under the standard stochastic blockmodel, and corresponds to (2) in the DCBM framework. Since each component of $\text{cl}(\theta; \mathbf{A})$ in (6) is a valid log-likelihood object, the *composite score estimating equation* $\nabla_\theta \text{cl}(\theta; \mathbf{A}) = 0$ is unbiased under usual regularity conditions. The associated composite likelihood estimator (CLE) $\hat{\theta}_C$, defined as the solution to $\nabla_\theta \text{cl}(\theta; \mathbf{A}) = 0$, suggests a natural estimator of the form $\hat{p} = p_{\hat{\theta}_C}$ to minimize the expected composite Kullback–Leibler divergence (Varin and Vidoni 2005) between the assumed blockmodel p_θ and the true, but unknown, joint density g ,

$$\Delta_C(g, p; \theta) := \sum_{i < j} \mathbb{E}_g(\log g(\{A_{ij}\}_{i < j} \in \mathcal{A}_{ij}) - \log p_{ij}(A_{ij}; \theta)),$$

where $\{\mathcal{A}_{ij}\}_{i < j}$ denotes the corresponding set of marginal events.

In terms of the asymptotic properties of the CLE, following the discussion in Cox and Reid (2004), it is important to distinguish whether the available data consist of many independent replicates from a common distribution function or form a few individually large sequences. While, in the first scenario, consistency and asymptotic normality of the corresponding $\hat{\theta}_C$ hold under some regularity conditions from the classical theory of estimating equations (Varin, Reid, and Firth 2011), some difficulties arise in the second one, which includes our observations $\{A_{ij}\}_{i < j}$. Indeed, as argued in Cox and Reid (2004), if there is too much internal correlation present among the individual components of the composite score $\nabla_\theta \text{cl}(\theta; \mathbf{A})$, the estimator $\hat{\theta}_C$ will not be consistent. The CLE will retain good properties as long as the data are not too highly correlated, which is the case for spatial data with exponential correlation decay. Under this setting, Heagerty and Lele (1998) proved consistency and asymptotic normality of $\hat{\theta}_C$ in a scenario where the data are not sampled independently from a study population. Under more general settings, consistency results are expected upon using limit theorems and parametric estimation for fields (e.g., Guyon 1995); however, applying the corresponding results requires a properly defined distance on networks and α -mixing conditions based on such distance.

3.3 Composite Likelihood BIC

Taking into account the measure of model complexity in the context of composite marginal likelihoods (Varin and Vidoni 2005), we define the following criterion for selecting the community number K :

$$\text{CL-BIC}_k := -2 \text{cl}(\hat{\theta}_C; \mathbf{A}) + d_k^* \log(N(N-1)/2), \quad (7)$$

where k is the number of communities under consideration in the current model used as model index, $d_k^* := \text{trace}(\mathbf{H}_k^{-1} \mathbf{V}_k)$, $\mathbf{H}_k := E_\theta(-\nabla_\theta^2 \text{cl}(\theta; \mathbf{A}))$ and $\mathbf{V}_k := \text{Var}_\theta(\nabla_\theta \text{cl}(\theta; \mathbf{A}))$. Then the resulting estimator for the community number is

$$\hat{k}_{\text{CL-BIC}} := \arg \min_{k \in \mathcal{K}} \text{CL-BIC}_k,$$

where \mathcal{K} is the set of candidate values for k .

Note that the CLE is a function of k , since a different model index yields a different estimator $\hat{\theta}_C := \hat{\theta}_C(k)$. Assuming independent and identically distributed data replicates, which lead

to consistent and asymptotically normally distributed estimators $\hat{\theta}_C$, Gao and Song (2010) established the model selection consistency of a similar composite likelihood BIC approach for high-dimensional parametric models. While allowing for the number of potential model parameters to increase to infinity, their consistency result only holds when the true model sparsity is bounded by a universal constant.

Even though, under a misspecification framework for the blockmodel family \mathcal{P} , the observed data $\{A_{ij}\}_{i<j}$ do not form independent replicates from a common population, we anticipate the CL-BIC criterion (7) to be consistent in selecting the true community number K , at least when the correlation among the $\{A_{ij}\}_{i<j}$ is not severe and the estimators $\hat{\theta}_C$ are consistent and asymptotically normal, as in Heagerty and Lele (1998). Since all the moment conditions in the consistency results from Gao and Song (2010) hold automatically after noticing the specific forms of the blockmodel composite log-likelihoods (1)–(3), under a properly defined mixing condition on $\{A_{ij}\}_{i<j}$ (Guyon 1995), and for a bounded community number $K \leq k_0$, we conjecture that $\mathbb{P}\{\hat{k}_{\text{CL-BIC}} = K\} \rightarrow 1$ as the number of nodes N in the network grows to infinity. This theoretical study will be relegated as a future work.

3.4 Formulas

3.4.1 Standard Stochastic Blockmodel

Following our discussions in the previous section, we treat (1) as the composite marginal likelihood, under the working independence assumption that, given the community labels of their endpoints, the Bernoulli random variables $\{A_{ij}\}_{i<j}$ are independent. The first-order partial derivative of $\ell(\theta; \mathbf{A})$ with respect to θ is denoted as $\mathbf{u}(\theta) = (u(\theta_{ab}); 1 \leq a \leq b \leq k)'$, where

$$u(\theta_{ab}) = \sum_{i<j} \left[\frac{A_{ij}}{\theta_{z_i z_j}} - \frac{1 - A_{ij}}{1 - \theta_{z_i z_j}} \right] \mathbf{I}_{i,j}^{a,b},$$

and

$$\mathbf{I}_{i,j}^{a,b} = \min(\mathbb{1}\{z_i = a, z_j = b\} + \mathbb{1}\{z_i = b, z_j = a\}, 1).$$

Furthermore, the second-order partial derivative of $\ell(\theta; \mathbf{A})$ has the following components,

$$\frac{\partial^2 \ell(\theta; \mathbf{A})}{\partial \theta_{a_1 b_1} \partial \theta_{a_2 b_2}} = 0, \text{ if } (a_1, b_1) \neq (a_2, b_2)$$

and

$$\frac{\partial^2 \ell(\theta; \mathbf{A})}{\partial \theta_{ab}^2} = - \sum_{i<j} \left[\frac{A_{ij}}{\theta_{z_i z_j}^2} + \frac{1 - A_{ij}}{(1 - \theta_{z_i z_j})^2} \right] \mathbf{I}_{i,j}^{a,b}.$$

Define the Hessian matrix $\mathbf{H}_k(\theta) = E_\theta(-\partial \mathbf{u}(\theta) / \partial \theta)$, then

$$\mathbf{H}_k(\theta) = E_\theta(\text{diag}\{-\partial^2 \ell(\theta; \mathbf{A}) / \partial \theta_{ab}^2; 1 \leq a \leq b \leq k\}).$$

Define the variability matrix $\mathbf{V}_k(\theta) = \text{var}_\theta(\mathbf{u}(\theta))$ and, following Varin and Vidoni (2005), the model complexity $d_k^* = \text{trace}[\mathbf{H}_k(\theta)^{-1} \mathbf{V}_k(\theta)]$. If the underlying model is indeed a correctly specified standard stochastic blockmodel, we have $d_k^* = k(k+1)/2$ and CL-BIC reduces to the traditional BIC. Indexed by $1 \leq k \leq k_0$, the estimated criterion functions for the CL-BIC

sequence (7) are

$$\widehat{\text{CL-BIC}}_k = -2 \text{cl}(\hat{\theta}_C; \mathbf{A}) + \hat{d}_k^* \log(N(N-1)/2), \quad (8)$$

where $\hat{\theta}_C$ and \hat{d}_k^* are estimators of θ and d_k^* , respectively. For a certain k , the explicit estimator forms are given below:

$$\hat{\mathbf{H}}_k(\hat{\theta}_C) = \text{diag} \left\{ \sum_{i<j} \left[\frac{A_{ij}}{\hat{\theta}_{z_i z_j}^2} + \frac{(1 - A_{ij})}{(1 - \hat{\theta}_{z_i z_j})^2} \right] \mathbf{I}_{i,j}^{a,b} \right\}$$

and $\hat{\mathbf{V}}_k(\hat{\theta}_C) = \mathbf{u}(\hat{\theta}_C)[\mathbf{u}(\hat{\theta}_C)]^T$.

As noted by Gao and Song (2010), the above naive estimator for $\mathbf{V}_k(\theta)$ vanishes when evaluated at the CLE $\hat{\theta}_C$. An alternative proposed in Varin, Reid, and Firth (2011) is to use a jackknife covariance matrix estimator, for the asymptotic covariance matrix of $\hat{\theta}_C$, of the form

$$\text{var}_{\text{jack}}(\hat{\theta}_C) = \frac{N-1}{N} \sum_{l=1}^N (\hat{\theta}_C^{(-l)} - \hat{\theta}_C)(\hat{\theta}_C^{(-l)} - \hat{\theta}_C)^T, \quad (9)$$

where $\hat{\theta}_C^{(-l)}$ is the composite likelihood estimator of θ with the l th vertex deleted. Let $\mathbf{A}^{(-l)}$ be the $(N-1) \times (N-1)$ matrix obtained after deleting the l th row and column from the original adjacency matrix \mathbf{A} . An explicit form for $\hat{\theta}_C^{(-l)}$ is given by $\hat{\theta}_{ab}^{(-l)} = 1/n_{ab}^{(-l)} \sum_{i<j} A_{ij}^{(-l)} \mathbb{1}\{z_i = a, z_j = b\}$, with $n_{ab}^{(-l)} = N_a^{(-l)} N_b^{(-l)}$ for $a \neq b$, and $n_{aa}^{(-l)} = N_a^{(-l)}(N_a^{(-l)} - 1)/2$; naturally, $N_a^{(-l)} = N_a - 1$ if $z_l = a$ and $N_a^{(-l)} = N_a$ otherwise.

Since the asymptotic covariance matrix of $\hat{\theta}_C$ is given by the inverse Godambe information matrix, $G_k(\theta)^{-1} = \mathbf{H}_k(\theta)^{-1} \mathbf{V}_k(\theta) \mathbf{H}_k(\theta)^{-1}$, see Gao and Song (2010) and Varin, Reid, and Firth (2011), an explicit estimator for d_k^* can be obtained by right-multiplying the jackknife covariance matrix estimator (9) by $\mathbf{H}_k(\theta)$ to obtain

$$\begin{aligned} \hat{d}_k^* &= \text{trace} \left[\text{var}_{\text{jack}}(\hat{\theta}_C) \hat{\mathbf{H}}_k(\hat{\theta}_C) \right] \\ &= \sum_{1 \leq a \leq b \leq k} \left\{ \text{var}_{\text{jack}}(\hat{\theta}_{ab}) \times \sum_{i<j} \left[\frac{A_{ij}}{\hat{\theta}_{z_i z_j}^2} + \frac{1 - A_{ij}}{(1 - \hat{\theta}_{z_i z_j})^2} \right] \mathbf{I}_{i,j}^{a,b} \right\}. \end{aligned}$$

3.4.2 Degree-Corrected Stochastic Blockmodel

Similarly, we develop corresponding parallel results for DCBM. The first- and second-order partial derivatives of $\ell(\theta, \omega; \mathbf{A})$ with respect to θ are defined as follows,

$$\frac{\partial \ell(\theta, \omega; \mathbf{A})}{\partial \theta} = \mathbf{u}(\theta) = (u(\theta_{ab}); 1 \leq a \leq b \leq k)', \quad u(\theta_{ab}) = \frac{m_{ab}}{\theta_{ab}} - 1,$$

$$\frac{\partial^2 \ell(\theta, \omega; \mathbf{A})}{\partial \theta_{a_1 b_1} \partial \theta_{a_2 b_2}} = 0, \text{ if } (a_1, b_1) \neq (a_2, b_2), \quad \frac{\partial^2 \ell(\theta, \omega; \mathbf{A})}{\partial \theta_{ab}^2} = - \frac{m_{ab}}{\theta_{ab}^2},$$

$$\hat{\mathbf{H}}_k(\hat{\theta}_C) = \text{diag} \left\{ \frac{1}{\hat{\theta}_{ab}}; 1 \leq a \leq b \leq k \right\},$$

which yields

$$\hat{d}_k^* = \sum_{1 \leq a \leq b \leq k} \left\{ \text{var}_{\text{jack}}(\hat{\theta}_{ab}) / \hat{\theta}_{ab} \right\}.$$

3.4.3 Mixed Membership Stochastic Blockmodel

The estimated model complexity for MMB now involves second-order partial derivatives of $\ell(\boldsymbol{\theta}, \boldsymbol{\alpha}; \mathbf{A})$ with respect to the hyperparameters $\boldsymbol{\theta}$ and $\boldsymbol{\alpha}$. Upon noticing the form of the first term of the complete data log-likelihood (3), and recalling the Hessian matrix with respect to $\boldsymbol{\alpha}$ detailed in (5), it is easy to see that $\hat{\mathbf{H}}_k(\hat{\boldsymbol{\theta}}_C, \hat{\boldsymbol{\alpha}}_C)$ is a block matrix of the form

$$\hat{\mathbf{H}}_k(\hat{\boldsymbol{\theta}}_C, \hat{\boldsymbol{\alpha}}_C) = \begin{pmatrix} \hat{\mathbf{H}}_k(\hat{\boldsymbol{\theta}}_C) & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{H}}_k(\hat{\boldsymbol{\alpha}}_C) \end{pmatrix},$$

where $\hat{\mathbf{H}}_k(\hat{\boldsymbol{\theta}}_C)$ is a $k^2 \times k^2$ diagonal matrix given by

$$\hat{\mathbf{H}}_k(\hat{\boldsymbol{\theta}}_C) = \text{diag} \left\{ \sum_{i,j} \left[\frac{A_{ij}}{\hat{\theta}_{z_i \rightarrow j, z_i \rightarrow j}^2} + \frac{(1 - A_{ij})}{(1 - \hat{\theta}_{z_i \rightarrow j, z_i \rightarrow j})^2} \right] \mathbf{1}\{z_{i \rightarrow j} = a, z_{i \leftarrow j} = b\} \right\},$$

and $\hat{\mathbf{H}}_k(\hat{\boldsymbol{\alpha}}_C) = (\hat{\mathbf{H}}_k(\hat{\boldsymbol{\alpha}}_C)_{ab})$ is a $k \times k$ matrix with entries

$$\hat{\mathbf{H}}_k(\hat{\boldsymbol{\alpha}}_C)_{ab} = N \left(\Psi'(\hat{\alpha}_a) \mathbf{1}\{a = b\} - \Psi' \left(\sum_a \hat{\alpha}_a \right) \right).$$

In a slight abuse of notation, we denote by $z_{i \rightarrow j}$ above the label assignment corresponding to node i when he sends a message to node j , and similarly for $z_{i \leftarrow j}$. The estimated model complexity is thus $\hat{d}_k^* = \text{trace}[\text{var}_{\text{jack}}(\hat{\boldsymbol{\theta}}_C, \hat{\boldsymbol{\alpha}}_C) \hat{\mathbf{H}}_k(\hat{\boldsymbol{\theta}}_C, \hat{\boldsymbol{\alpha}}_C)]$, where the jack-knife matrix $\text{var}_{\text{jack}}(\hat{\boldsymbol{\theta}}_C, \hat{\boldsymbol{\alpha}}_C)$, assuming a similar form as in (9) with $\hat{\boldsymbol{\theta}}_C^{(-l)}$ and $\hat{\boldsymbol{\alpha}}_C^{(-l)}$ estimated as explained in Section 2, provides the corresponding asymptotic covariance matrix estimator of the CLE $(\boldsymbol{\theta}_C, \boldsymbol{\alpha}_C)$.

We would like to remark that our CL-BIC approach for selecting the community number K extends beyond the realm of stochastic blockmodels. Indeed, both the latent space cluster model of Handcock, Raftery, and Tantrum (2007) and the local dependence model of Schweinberger and Handcock (2015), as well as any other (composite) likelihood-based approach that requires to select a value of K can employ our proposed CL-BIC methodology for selecting the number of communities. We leave the details of this further investigation for future research.

4. Experiments

In this section, we show the advantages of the CL-BIC approach over the traditional BIC as well as the variational Bayes approach in selecting the true number of communities via simulations and two real datasets.

4.1 Simulations

For simplicity of the presentation, we consider only the SBM and the DCBM in our simulations. For each setting, we relax the assumption that the A_{ij} 's are conditionally independent given the labels ($Z_i = z_i, Z_j = z_j$), varying both the dependence structure of the adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ and the value of the parameters $(\boldsymbol{\theta}, \boldsymbol{\omega})$. The models introduced are correlation-contaminated stochastic blockmodels, that is, we bring different types of correlation into the stochastic blockmodels, both standard and degree-corrected, mimicking real-world networks.

All of our simulated adjacency matrices have independent rows. That is, the binary variables A_{ik} and A_{jl} are independent, whenever $i \neq j$, given the corresponding community labels of their endpoints. However, for a fixed node $i \in V$, correlation does exist across different columns in the binary variables A_{ij} and A_{il} . For the standard stochastic blockmodel, correlated binary random variables are generated, following the approach in Leisch, Weingessel, and Hornik (1998), by thresholding a multivariate Gaussian vector with correlation matrix $\boldsymbol{\Sigma}$ satisfying $\Sigma_{jl} = \rho_{jl}$. Specifically, for any choice of $|\rho_{jl}| \leq 1$, we simulate correlated variables A_{ij} and A_{il} such that $\text{cov}(A_{ij}, A_{il}) = L(-\mu_j, -\mu_l, \rho_{jl}) - \theta_{z_i z_j} \theta_{z_i z_l}$. Here, following Leisch, Weingessel, and Hornik (1998), we have $L(-\mu_j, -\mu_l, \rho_{jl}) = \mathbb{P}(W_j \geq -\mu_j, W_l \geq -\mu_l)$, $\mu_j = \Phi^{-1}(\theta_{z_i z_j})$, and $\mu_l = \Phi^{-1}(\theta_{z_i z_l})$, where (W_j, W_l) is standard bivariate normal with correlation ρ_{jl} . Correlated Bernoulli variables for the degree-corrected blockmodel are generated in a similar fashion.

In each experiment, carried over 200 randomly generated adjacency matrices, we record the proportion of times the chosen number of communities for each of the different criteria for selecting K agrees with the truth. Apart from CL-BIC and BIC, we also consider the integrated likelihood variational Bayes (VB) approach of Latouche, Birmelé, and Ambroise (2012). To estimate the true community number, their method selects the candidate value k that maximizes a variational Bayes approximation to the observed data log-likelihood.

We restrict attention to candidate values for the true K in the range $k \in \{1, \dots, 18\}$, both in simulations and the real data analysis section. For Simulations 1–3, spectral clustering is used to obtain the community labels for each candidate k , whereas in the DCBM setting of Simulation 4, the SCORE algorithm is employed. Additionally, among the incorrectly selected community number trials, we calculate the median deviation between the selected community number and the true $K = 4$, as well as its robust standard deviation.

Simulation 1: Correlation among the edges within and between communities is introduced simultaneously throughout all blocks in the network, and not proceeding in a block-by-block fashion. Concretely, for each node i , all edges $\{A_{ij}\}_{i < j}$ are generated by thresholding a correlated $(N - i)$ -dimensional Gaussian random vector with correlation matrix $\boldsymbol{\Sigma} = (\rho_{jl})$. Thus, in this scenario, all edges A_{ij} and A_{il} with common endpoint i are correlated, regardless of whether j and l belong to the same community or not. Cases $\rho_{jl} = \rho$ and $\rho_{jl} = \rho^{|j-l|}$, with several choices of ρ are conducted. We consider a four-community network, $\boldsymbol{\theta} = (\theta_{ab}; 1 \leq a \leq b \leq 4)$, where $\theta_{aa} = 0.35$ for all $a = 1, \dots, 4$ and $\theta_{ab} = 0.05$ for $1 \leq a < b \leq 4$. Community sizes are 60, 90, 120, and 150, respectively. Results are collected in Table 1.

Simulation 2: Correlation among the edges within (ρ W.) and between (ρ B.) communities is introduced block-wisely. Concretely, for each node i , all edges A_{ij} and A_{il} are generated independently whenever j and l belong to different communities. If j and l belong to the same community, edges A_{ij} and A_{il} are generated by thresholding a correlated Gaussian random vector with correlation matrix $\boldsymbol{\Sigma} = (\rho_{jl})$. Parameter settings are identical to Simulation 1, with results collected in Table 2.

Simulation 3: Correlation settings are the same as in Simulation 2, but we change the value of the parameter $\boldsymbol{\theta}$ to allow for

Table 1. Comparison of CL-BIC and BIC over 200 repetitions from Simulation 1, where Eq and Dec indicate equally correlated and exponential decaying cases, respectively. Both correlation of multivariate Gaussian random variables (ρ MVN) and the corresponding *maximum* correlation between Bernoulli variables (ρ Ber.) are presented.

CORR		PROP		MEDIAN DEV		CORR		PROP		MEDIAN DEV	
ρ MVN	ρ Ber.	CL-BIC	BIC	CL-BIC	BIC	ρ MVN	ρ Ber.	CL-BIC	BIC	CL-BIC	BIC
0.10	0.06	1.00	0.40	0.0(0.0)	2.0(1.5)	0.40	0.25	1.00	0.35	0.0(0.0)	2.0(1.5)
0.15 Eq	0.09	0.92	0.14	1.0(0.0)	3.0(2.2)	0.50 Dec	0.32	1.00	0.21	0.0(0.0)	2.0(1.5)
0.20	0.12	0.81	0.03	1.0(0.4)	5.0(3.0)	0.60	0.40	0.99	0.12	1.0(0.0)	3.0(1.5)

NOTE: CORR, correlation; PROP, proportion; MEDIAN DEV, median deviation. In the MEDIAN DEV columns, results are in the form of median (robust standard deviation).

Table 2. Comparison of CL-BIC and BIC over 200 repetitions from Simulation 2, where *Ind* indicates $\rho_{jl} = 0$ for $j \neq l$. For simplicity, we omit the correlation between the corresponding Bernoulli variables.

CORR		PROP		MEDIAN DEV		CORR		PROP		MEDIAN DEV	
ρ W.	ρ B.	CL-BIC	BIC	CL-BIC	BIC	ρ W.	ρ B.	CL-BIC	BIC	CL-BIC	BIC
0.10		1.00	0.64	0.0(0.0)	2.0(0.7)		0.40	1.00	0.59	0.0(0.0)	1.0(1.5)
0.15 Eq	<i>Ind</i>	0.98	0.36	1.0(0.7)	2.0(1.5)	0.10 Eq	0.50 Dec	1.00	0.54	0.0(0.0)	2.0(0.7)
0.20		0.80	0.08	1.0(0.7)	3.0(2.2)		0.60	1.00	0.53	0.0(0.0)	2.0(1.1)
0.40		1.00	0.33	0.0(0.0)	2.0(0.7)		0.40	0.98	0.32	1.0(0.0)	2.0(1.5)
0.50 Dec	<i>Ind</i>	1.00	0.29	0.0(0.0)	2.0(1.5)	0.15 Eq	0.50 Dec	0.97	0.30	1.0(0.4)	3.0(1.5)
0.60		1.00	0.14	0.0(0.0)	2.0(1.5)		0.60	0.95	0.25	1.0(0.0)	3.0(1.5)

more general network topologies. We set $\theta = (\theta_{ab}; 1 \leq a \leq b \leq 4)'$ with $\theta_{aa} = \theta_{b4} = 0.35$ for all $a = 1, \dots, 4$ and $b = 1, 2, 3$. The remaining entries of θ are set to 0.05. Hence, following Latouche, Birmelé, and Ambroise (2012), vertices from community 4 connect with probability 0.35 to any other vertices in the network, forming a community of only hubs. Community sizes are the same as in Simulation 1, with results collected in Table 3.

Simulation 4: We follow the approach of Zhao, Levina, and Zhu (2012) in choosing the parameters (θ, ω) to generate networks from the degree-corrected blockmodel. Thus, the identifiability constraint $\sum_i \omega_i \mathbf{1}\{z_i = a\} = 1$ for each community $1 \leq a \leq K$ is replaced by the requirement that the ω_i be independently generated from a distribution with unit expectation, fixed here to be

$$\omega_i = \begin{cases} \eta_i, & \text{w.p. } 0.8, \\ 2/11, & \text{w.p. } 0.1, \\ 20/11, & \text{w.p. } 0.1, \end{cases}$$

where η_i is uniformly distributed on the interval $[0, 2]$. The vector θ , in a slight abuse of notation, is reparameterized as $\theta_n = \gamma_n \theta$, where we vary the constant γ_n to obtain different expected degrees of the network. Correlation settings and community sizes are the same as in Simulation 1, with results presented in Table 4, where choices for γ_n and θ are specified.

When the stochastic blockmodels are contaminated by the imposed correlation structure, which is expected in real-world networks, CL-BIC outperforms BIC overwhelmingly. Tables 1–2 show the improvement is more significant when the imposed correlation is larger. For instance, in the block-wise correlated case of Table 2, when we only have within-community correlation $\rho_{Dec} = 0.60$, CL-BIC does the right selection in all cases, while BIC is only successful in 14% of 200 trials.

As shown in Table 3 for the model with a community of only hubs, if the network is generated from a purely stochastic blockmodel, or if the contaminating correlation is not too strong, CL-BIC and VB have similar performance in selecting the correct $K = 4$. But again, as the imposed correlation increases, VB fails

to make the right selection more often than CL-BIC. This is particularly true in the $\rho_{Eq} = 0.20$ case, where CL-BIC makes the right selection in 85% of simulated networks, whereas VB fails in all cases, yielding models with a median of nine communities.

The same pattern translates into the DCBM setting of Table 4, where smaller values of γ_n yield sparser networks. The community number selection problem becomes more difficult as γ_n decreases, as degrees for many nodes are small, yielding noisy individual effect estimates $\hat{\omega}_i = d_i / \sum_{j:z_j=z_i} d_j$. Nevertheless, the CL-BIC approach consistently selects the correct number of communities more frequently than BIC over different correlation settings.

In addition, Figure 1 presents simulation results where the true community number K increases from 2 to 8. Following our previous examples, community sizes grow according to the sequence (60, 90, 120, 150, 60, 90, 120, 150). The selected correlation-contaminated stochastic blockmodels are $\rho_{Eq} = 0.10$ from Simulation 1, within-community correlation $\rho_{Eq} = 0.10$ from Simulation 2, and within-community correlation $\rho_{Dec} = 0.60$ from Simulation 3. As K increases and enough vertices are added into the network, CL-BIC tends to correctly estimate the true community number in all simulation settings. Even in this scenario with a growing number of communities, the proportion of times CL-BIC selects the true K is always greater than the corresponding BIC or VB estimates.

Before moving to the last simulation example, we would like to define two measures to quantify the accuracy of a given node label assignment. The first measure is a “goodness-of-fit” (GF) measure defined as

$$GF(\mathbf{z}, \hat{\mathbf{z}}_k) = \sum_{i < j} (\mathbb{1}\{z_i = z_j\} \mathbb{1}\{\hat{z}_i = \hat{z}_j\} + \mathbb{1}\{z_i \neq z_j\} \mathbb{1}\{\hat{z}_i \neq \hat{z}_j\}) / \binom{N}{2}, \quad (10)$$

where \mathbf{z} represents the true community labels and $\hat{\mathbf{z}}_k$ represents the community assignments from an estimator. Thus, the

Table 3. Comparison of CL-BIC and VB over 200 repetitions from Simulation 3. For simplicity, we omit the correlation between the corresponding Bernoulli variables.

CORR		PROP		MEDIAN DEV		CORR		PROP		MEDIAN DEV	
ρ W.	ρ B.	CL-BIC	VB	CL-BIC	VB	ρ W.	ρ B.	CL-BIC	VB	CL-BIC	VB
0.00		1.00	1.00	0.0(0.0)	0.0(0.0)	0.00		1.00	1.00	0.0(0.0)	0.0(0.0)
0.10 Eq	<i>Ind</i>	0.96	0.00	1.0(0.0)	2.0(0.0)	0.40 Dec	<i>Ind</i>	1.00	1.00	0.0(0.0)	0.0(0.0)
0.15		0.88	0.00	1.0(0.0)	4.0(2.2)	0.50		1.00	0.94	0.0(0.0)	1.0(0.0)
0.20		0.85	0.00	1.0(0.0)	5.0(1.5)	0.60		1.00	0.56	0.0(0.0)	1.0(0.0)

Table 4. Comparison of CL-BIC and BIC over 200 repetitions from Simulation 4. Before being scaled by the constant γ_n , we selected $\theta = (\theta_{ab}; 1 \leq a \leq b \leq 4)'$, where $\theta_{aa} = 7$ for all $a = 1, \dots, 4$ and $\theta_{ab} = 1$ for $1 \leq a < b \leq 4$.

CORR ρ		PROP		MEDIAN DEV		CORR ρ		PROP		MEDIAN DEV	
MVN	γ_n	CL-BIC	BIC	CL-BIC	BIC	MVN	γ_n	CL-BIC	BIC	CL-BIC	BIC
0.20 Eq	0.02	0.84	0.35	1.0(1.5)	2.0(1.5)	0.60 Dec	0.02	0.92	0.52	-1.0(1.5)	2.0(1.5)
	0.03	0.96	0.58	1.0(0.0)	3.0(1.5)		0.03	1.00	0.81	0.0(0.0)	1.0(1.5)
0.30 Eq	0.02	0.70	0.31	1.0(0.4)	2.0(0.7)	0.70 Dec	0.02	0.83	0.41	-1.0(1.5)	2.0(1.5)
	0.03	0.93	0.52	1.0(0.6)	3.0(1.5)		0.03	1.00	0.77	0.0(0.0)	1.0(0.7)
0.40 Eq	0.02	0.43	0.21	1.0(1.5)	2.0(1.5)	0.80 Dec	0.02	0.69	0.22	-1.0(1.5)	3.0(2.4)
	0.03	0.85	0.51	1.0(1.9)	3.0(1.9)		0.03	0.98	0.69	-1.0(0.4)	1.0(1.5)

measure $GF(z, \hat{z}_k)$ calculates the proportion of pairs whose estimated assignments agree with the correct labels in terms of being assigned to the same or different communities, and is commonly known as the *Rand Index* (Rand 1971) in the cluster analysis literature.

The second measure is motivated from the ‘‘assortativity’’ notion. The ratio of the median within community edge number to that of the between community edge number (MR) is defined as

$$MR(\hat{z}_k) = \text{median}_{a=1, \dots, k} (m_{aa}) / \text{median}_{a \neq b} (m_{ab}), \quad (11)$$

where k is the number of communities implied by \hat{z}_k and m_{ab} is the total number of edges between communities a and b , as given by the community assignment \hat{z}_k . It is clear that for both measures, a higher value indicates a better community detection performance.

As a final simulation example, we analyze the performance of CL-BIC and BIC for a growing number of communities under the degree-corrected blockmodel. While the reparameterized vector $\theta_n = \gamma_n \theta$ remains as in Simulation 4, the ω_i are now independently generated from $\text{Uniform}(1/5, 9/5)$. The results are collected in Table 5, where we also record the performance of the SCORE algorithm under the true K , along with the goodness-of-fit (GF) and median ratio (MR) performance measures introduced in (10) and (11), respectively.

The true community number and community sizes grow as in the case for the standard blockmodel described in Figure 1. Although CL-BIC performs uniformly better than BIC across all validating criteria and throughout all K , the procedure does not appear to yield model selection consistent results in this example. Aside from the fact that the introduced correlation is not exponentially decaying, this poor performance as K increases




Figure 1. Comparisons between different methods for selecting the true community number K in the standard blockmodel settings of Simulations 1–3. Along the y-axis, we record the proportion of times the chosen number of communities for each of the different criteria for selecting K agrees with the truth.

Table 5. Comparison of CL-BIC and BIC over 200 repetitions from the DCBM case in Simulation 4, with $(\rho_{Eq} = 0.2, \gamma_n = 0.03)$, where the individual effect parameters ω_i are now generated from a Uniform(1/5, 9/5) distribution.

K	SCORE performance			CL-BIC					BIC				
	Misc. R.	Orac. Err.	Est. Err.	PROP	MD	RSD	GF	MR	PROP	MD	RSD	GF	MR
2	0.02	0.51	0.54	0.88	1	0	0.96	7.33	0.10	2	0.75	0.73	3.37
3	0.03	0.53	0.55	0.93	1	0.75	0.97	7.87	0.09	3	1.49	0.88	3.91
4	0.03	0.55	0.58	0.86	1	0	0.97	8.11	0.16	3	1.49	0.91	5.41
5	0.04	0.58	0.62	0.56	1	0	0.96	6.35	0.09	3	2.05	0.92	5.92
6	0.05	0.60	0.64	0.47	1	1.49	0.96	7.10	0.09	2	2.24	0.94	7.08
7	0.05	0.63	0.66	0.39	1	1.49	0.97	6.77	0.09	3	1.49	0.95	6.73
8	0.08	0.63	0.66	0.29	1	1.49	0.97	7.24	0.02	3	2.24	0.96	6.80

NOTE: PROP, proportion; MD, median deviation; RSD, robust standard deviation; GF, goodness-of-fit measure; MR: median ratio measure. Misc. R. denotes the misclustering rate of the SCORE algorithm. For $\Omega = E_{\theta}(A)$, Orac. Err. and Est. Err. are $\|\Omega_O - \Omega\| / \|\Omega\|$ and $\|\Omega_{SC} - \Omega\| / \|\Omega\|$, respectively, where $\|\cdot\|$ denotes Frobenius norm. Here, Ω_O denotes the estimate of Ω under the *oracle* scenario where we know the true community assignment $z \in \{1, \dots, K\}^N$, and Ω_{SC} is the estimate of Ω using the SCORE labeling vector.

can also be explained by the difficulty in estimating the DCBM parameters (θ, ω) in a scenario where several vertices have potentially low degrees. Indeed, even in the *oracle* scenario where we know the true community labels z_i ahead of time, and for a relatively small misclustering rate of the SCORE algorithm, Table 5 exhibits the difficulty in obtaining accurate estimates $(\hat{\theta}_C, \hat{\omega}_C)$, and in evaluating the CL-BIC criterion functions (8), under this increasing K scenario for the DCBM. Whether the increased number of parameters in the DCBM has an effect on the consistency results of CL-BIC as K increases is also an interesting line of future work.

4.2 Real Data Analysis

4.2.1 International Trade Networks

We first study an international trade dataset originally analyzed in Westveld and Hoff (2011), containing yearly international trade data between $N = 58$ countries from 1981 to 2000. For a more detailed description of this dataset, we refer the interested reader to the Appendix in Westveld and Hoff (2011). In our numerical comparisons between CL-BIC and BIC paired with the standard stochastic blockmodel log-likelihood (1), we focus on data from year 1995. For this network, an adjacency matrix A can be formed by first considering a weight matrix W with $W_{ij} = \text{Trade}_{i,j} + \text{Trade}_{j,i}$, where $\text{Trade}_{i,j}$ denotes the value of exports from country i to country j . Finally, we define $A_{ij} = 1$ if $W_{ij} \geq W_{\alpha}$, and $A_{ij} = 0$ otherwise; here W_{α} denotes the α th quantile of $\{W_{ij}\}_{1 \leq i < j \leq N}$. For the choice of $\alpha = 0.5$,

Figure 2 shows the largest connected component of the resulting network. Panel (a) shows CL-BIC selecting three communities, corresponding to countries with the highest GDPs (dark blue), industrialized European and Asian countries with medium-level GDPs (green), and developing countries in South America with the smallest GDPs (yellow). Next, in panel (b) we also show the variational Bayes solution corresponding to $k = 7$, providing finer communities for some Central and South American neighboring countries (yellow and pink, respectively) but fragmenting the high- and medium-level GDP countries into ambiguous communities. For instance, it is not clear why countries like Bolivia and Nepal belong to the same community (orange) or why the Netherlands, rather than Brazil or Italy, joined the community of countries with the highest GDPs (light blue). At last, panel (c) corresponds to the final BIC model selecting 10 communities. Under this partition, South American countries are now split into six “noisy” communities, while high GDP countries are unnecessarily fragmented into two.

We believe CL-BIC provides a better model than traditional BIC, yielding communities with countries sharing similar GDP values without dividing an entire continent into six smaller communities. On the contrary, BIC selects a model containing communities of size as small as one, which are of little, if any, practical use. The variational Bayes approach provides a meaningful solution in this example, exhibiting a similar performance as in Latouche, Birmelé, and Ambroise (2012) in terms of providing some finer community assignments.




Figure 2. Largest connected component of the international trade network for the year 1995.




Figure 3. Largest connected component of the school friendship network. Panel (c) shows the “true” grade community labels: 7th (blue), 8th (yellow), 9th (green), 10th (purple), 11th (red), and 12th (black).

4.2.2 School Friendship Networks

Now, we consider a school friendship network obtained from the National Longitudinal Study of Adolescent Health (<http://www.cpc.unc.edu/projects/addhealth>). For this network, $A_{ij} = 1$ if either student i or j reported a close friendship tie between the two, and $A_{ij} = 0$ otherwise. We focus on the network of school 7 from this dataset, and our comparisons between CL-BIC and BIC are done with respect to the degree-corrected blockmodel log-likelihood (2). With 433 vertices, Figure 3 shows the largest connected component of the resulting network. As shown in panel (a), CL-BIC selects the true community number $K = 6$, roughly agreeing with the actual grade labels, except for the black community. BIC, shown in panel (b), selects nine communities, unnecessarily splitting the 7th and 8th graders. The “true” friendship network is shown in panel (c).

We still conclude CL-BIC performs better than traditional BIC. Except for the misallocation of the black community of 12th graders, the model selected by CL-BIC correctly labels most of the remaining network. While BIC partially separates the 10th graders and the 12th graders, a substantial portion of the 10th graders are absorbed into the 9th grader community (green). In addition, BIC further fragments 7th and 8th graders into “noisy” communities. This is an extremely difficult community detection problem since, even for a “correctly” specified $k = 6$, SCORE fails to assign all 12th graders to their corresponding true grade. The black community selected by SCORE in panel (a) mainly corresponds to female students and Hispanic males, reflecting perhaps closer friendship ties among a subgroup of students recently starting junior high school.

Using the “goodness-of-fit” measure defined in (10), we found that the CL-BIC community assignment leads to $\text{GF}(\mathbf{z}, \hat{\mathbf{z}}_6) = 0.811$, which is slightly better than the $\text{GF}(\mathbf{z}, \hat{\mathbf{z}}_9) = 0.810$ obtained for BIC. For the MR measure given in (11), the results for CL-BIC and BIC are $\text{MR}(\hat{\mathbf{z}}_6) = 40.8$ and $\text{MR}(\hat{\mathbf{z}}_9) = 33.3$, respectively, again indicating the superiority of the CL-BIC solution paired with SCORE.

In both examples, BIC tends to overestimate the “true” community number K , rendering very small communities that are in turn penalized under the CL-BIC approach. This means CL-BIC successfully remedies the robustness issues brought in by spectral clustering, due to the misspecification of the underlying stochastic blockmodels, and effectively captures the loss of variance produced by using traditional BIC.

5. Discussion

There has been a tremendous amount of research in recovering the underlying structures of network data, especially on the community detection problem. Most of the existing work has focused on studying the properties of the stochastic blockmodel and its variants without looking at the possible model misspecification problem. In this article, under the standard stochastic blockmodel and its variants, we advocate the use of composite likelihood BIC for selecting the number of communities due to its simplicity in implementation and its robustness against correlated binary data.

Some extensions are possible. For instance, the proposed methodology in this work is based on the spectral clustering and SCORE algorithms, and it would be interesting to explore the combination of the CL-BIC with other community detection methods. In addition, most examples considered here are dense graphs, which are common but cannot exhaust all scenarios in real applications. Another open problem is to study whether the CL-BIC approach is consistent for the degree-corrected stochastic blockmodel, which is not necessarily true from our numerical studies.

Supplementary Materials

R Code and Trade Dataset: The R codes can be used to replicate the simulation studies and the real data analysis. The international trade network dataset is also included. More details can be found in the file README contained in the zip file. (CLBIC.zip)

Acknowledgments

The authors thank the editor, the associate editor, and two anonymous referees for their constructive comments that have greatly improved the article. Yu is supported by Richard Samworth’s Engineering and Physical Sciences Research Council Early Career Fellowship EP/J017213/1. Feng is partially supported by NSF grant DMS-1308566.

References

- Airoldi, E. M., Blei, D. M., Fienberg, S. E., and Xing, E. P. (2008), “Mixed Membership Stochastic Blockmodels,” *Journal of Machine Learning Research*, 9, 1981–2014. [171,173,174]
- Chen, K., and Lei, J. (2014), “Network Cross-Validation for Determining the Number of Communities in Network Data,” *arXiv:1411.1715v1*. [171]
- Choi, D. S., Wolfe, P. J., and Airoldi, E. M. (2012), “Stochastic Blockmodels With a Growing Number of Classes,” *Biometrika*, 99, 273–284. [172]

- Cox, D. R., and Reid, N. (2004), "A Note on Pseudolikelihood Constructed From Marginal Densities," *Biometrika*, 91, 729–737. [174]
- Daudin, J.-J., Picard, F., and Robin, S. (2008), "A Mixture Model for Random Graphs," *Statistics and Computing*, 18, 173–183. [171,173]
- Decelle, A., Krzakala, F., Moore, C., and Zdeborová, L. (2011), "Asymptotic Analysis of the Stochastic Block Model for Modular Networks and Its Algorithmic Applications," *Physical Review E*, 84, 066106–19. [172]
- Donath, W. E., and Hoffman, A. J. (1973), "Lower Bounds for the Partitioning of Graphs," *IBM Journal of Research and Development*, 17, 420–425. [171]
- Fearnhead, P. (2004), "Particle Filters for Mixture Models With an Unknown Number of Components," *Statistics and Computing*, 14, 11–21. [171]
- Ferguson, T. S. (1973), "A Bayesian Analysis of Some Nonparametric Problems," *The Annals of Statistics*, 1, 209–230. [171]
- Gao, X., and Song, P. X.-K. (2010), "Composite Likelihood Bayesian Information Criteria for Model Selection in High-Dimensional Data," *Journal of The American Statistical Association*, 105, 1531–1540. [172,175]
- Goldenberg, A., Zheng, A. X., Fienberg, S. E., and Airolidi, E. M. (2010), "A Survey of Statistical Network Models," *Foundations and Trends in Machine Learning*, 2, 129–233. [171]
- Guattery, S., and Miller, G. L. (1998), "On the Quality of Spectral Separators," *SIAM Journal on Matrix Analysis and Applications*, 19, 701–719. [174]
- Guyon, X. (1995), *Random Fields on a Network: Modeling, Statistics, and Applications*, New York: Springer-Verlag. [174,175]
- Handcock, M. S., Raftery, A. E., and Tantrum, J. M. (2007), "Model-Based Clustering for Social Networks," *Journal of the Royal Statistical Society, Series A*, 170, 301–354. [173,176]
- Heagerty, P. J., and Lele, S. R. (1998), "A Composite Likelihood Approach to Binary Spatial Data," *Journal of the American Statistical Association*, 93, 1099–1111. [174,175]
- Holland, P. W., Laskey, K. B., and Leinhardt, S. (1983), "Stochastic Blockmodels: First Steps," *Social Networks*, 5, 109–137. [171,172]
- Hunter, D. R., Krivitsky, P. N., and Schweinberger, M. (2012), "Computational Statistical Methods for Social Network Models," *Journal of Computational and Graphical Statistics*, 21, 856–882. [171]
- Jin, J. (2015), "Fast Community Detection by SCORE," *The Annals of Statistics*, 43, 57–89. [173]
- Karrer, B., and Newman, M. E. J. (2011), "Stochastic Blockmodels and Community Structure in Networks," *Physical Review E*, 83, 016107–10. [172,174]
- Latouche, P., Birmelé, E., and Ambroise, C. (2012), "Variational Bayesian Inference and Complexity Control for Stochastic Block Models," *Statistical Modelling*, 12, 93–115. [171,172,176,179]
- Leisch, F., Weingessel, A., and Hornik, K. (1998), "On the Generation of Correlated Artificial Binary Data," Working Paper No. 13, SFB, Adaptive Information Systems and Modelling in Economics and Management Science. [176]
- Lindsay, B. G. (1988), "Composite Likelihood Methods," *Contemporary Mathematics*, 80, 221–239. [172,174]
- McDaid, A. F., Murphy, T. B., Friel, N., and Hurley, N. J. (2013), "Improved Bayesian Inference for the Stochastic Blockmodel With Application to Large Networks," *Computational Statistics and Data Analysis*, 60, 12–31. [171]
- Miller, J. W., and Harrison, M. T. (2014), "Inconsistency of Pitman-Yor Process Mixtures for the Number of Components," *Journal of Machine Learning Research*, 15, 3333–3370. [171]
- Minka, T. P. (2000), "Estimating a Dirichlet Distribution," Technical Report, Microsoft Research. [173]
- Newman, M. E. J. (2004), "Detecting Community Structure in Networks," *The European Physical Journal B*, 38, 321–330. [171]
- Newman, M. E. J., and Girvan, M. (2004), "Finding and Evaluating Community Structure in Networks," *Physical Review E*, 69, 026113–15. [171]
- Nobile, A., and Fearnside, A. (2007), "Bayesian Finite Mixtures With an Unknown Number of Components: The Allocation Sampler," *Statistics and Computing*, 17, 147–162. [171]
- Perry, P. O., and Wolfe, P. J. (2012), "Null Models for Network Data," *arXiv:1201.5871v1*. [172]
- Rand, W. M. (1971), "Objective Criteria for the Evaluation of Clustering Methods," *Journal of the American Statistical Association*, 66, 846–850. [178]
- Rohe, K., Chatterjee, S., and Yu, B. (2011), "Spectral Clustering and the High-Dimensional Stochastic Blockmodel," *The Annals of Statistics*, 39, 1878–1915. [172,173]
- Rohe, K., Qin, T., and Fan, H. (2014), "The Highest Dimensional Stochastic Blockmodel With a Regularized Estimator," *Statistica Sinica*, 24, 1771–1786. [172]
- Schweinberger, M., and Handcock, M. S. (2015), "Local Dependence in Random Graph Models: Characterization, Properties and Statistical Inference," *Journal of the Royal Statistical Society, Series B*, 77, 647–676. [176]
- Shi, J., and Malik, J. (2000), "Normalized Cuts and Image Segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22, 888–905. [171]
- Spielmat, D. A., and Teng, S.-H. (1996), "Planar Graphs and Finite Element Meshes," in *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, IEEE, pp. 96–105. [171]
- Stephens, M. (2000), "Bayesian Analysis of Mixture Models With an Unknown Number of Components—An Alternative to Reversible Jump Methods," *The Annals of Statistics*, 28, 40–74. [171]
- Varin, C. (2008), "On Composite Marginal Likelihoods," *ASA Advances in Statistical Analysis*, 92, 1–28. [174]
- Varin, C., Reid, N., and Firth, D. (2011), "An Overview of Composite Likelihood Methods," *Statistica Sinica*, 21, 5–42. [174,175]
- Varin, C., and Vidoni, P. (2005), "A Note on Composite Likelihood Inference and Model Selection," *Biometrika*, 92, 519–528. [174,175]
- Luxburg, U. (2007), "A Tutorial on Spectral Clustering," *Statistics and Computing*, 17, 395–416. [173]
- Wei, Y.-C., and Cheng, C.-K. (1989), "Towards Efficient Hierarchical Designs by Ratio Cut Partitioning," in *IEEE International Conference on Computer-Aided Design, 1989. ICCAD-89. Digest of Technical Papers*, pp. 298–301. [171]
- Westveld, A. H., and Hoff, P. D. (2011), "A Mixed Effects Model for Longitudinal Relational and Network Data, With Applications to International Trade and Conflict," *The Annals of Applied Statistics*, 5, 843–872. [179]
- Zhao, Y., Levina, E., and Zhu, J. (2012), "Consistency of Community Detection in Networks Under Degree-Corrected Stochastic Block Models," *The Annals of Statistics*, 40, 2266–2292. [172,177]