Electronic Journal of Statistics Vol. 16 (2022) 3135–3175 ISSN: 1935-7524 https://doi.org/10.1214/22-EJS2018

Spectral graph clustering via the expectation-solution algorithm^{*}

Zachary M. Pisano[†], Joshua S. Agterberg, Carey E. Priebe and Daniel Q. Naiman

Department of Applied Mathematics and Statistics, Johns Hopkins University, e-mail: zpisano1@jhu.edu; jagterb1@jhu.edu; cep@jhu.edu; daniel.naiman@jhu.edu

Abstract: The stochastic blockmodel (SBM) models the connectivity within and between disjoint subsets of nodes in networks. Prior work demonstrated that the rows of an SBM's adjacency spectral embedding (ASE) and Laplacian spectral embedding (LSE) both converge in law to Gaussian mixtures where the components are curved exponential families. Maximum likelihood estimation via the Expectation-Maximization (EM) algorithm for a full Gaussian mixture model (GMM) can then perform the task of clustering graph nodes, albeit without appealing to the components' curvature. Noting that EM is a special case of the Expectation-Solution (ES) algorithm, we propose two ES algorithms that allow us to take full advantage of these curved structures. After presenting the ES algorithm for the general curved-Gaussian mixture, we develop those corresponding to the ASE and LSE limiting distributions. Simulating from artificial SBMs and a brain connectome SBM reveals that clustering graph nodes via our ES algorithms can improve upon that of EM for a full GMM for a wide range of settings.

MSC2020 subject classifications: Primary 62H30, 60B20; secondary 62P10.

Keywords and phrases: EM algorithm, random graph, curved exponential family, estimating equations, mixture model.

Received January 2021.

1. Introduction

Statistical inference on graphs is a burgeoning field of study in statistics with applications in neuroscience [14, 18] and social networks [12], among other areas of research. Given a random graph G on n vertices generated by some model F, a typical procedure is to embed its adjacency matrix $\mathbf{A} \in \{0,1\}^{n \times n}$ into a lower dimensional space \mathbb{R}^d as a collection of n points. Commonly chosen embeddings include the *adjacency spectral embedding* (ASE) and *Laplacian spectral embedding* (LSE), obtained via the truncated eigendecomposition of \mathbf{A} and its normalized Laplacian $\mathcal{L}(\mathbf{A})$. From here one may seek to perform the task of clustering these points and — by extension — their corresponding vertices.

^{*}This work is partially supported by the D3M program of the Defense Advanced Research Projects Agency (DARPA). The authors would like to acknowledge and thank the referees, whose commentary and criticism of an earlier version of the present article improved it considerably.

[†]Corresponding author.

Z. M. Pisano et al.

The spectral graph clustering problem has been extensively studied for settings in which the graph (or graphs) on hand are posited to have been generated by a stochastic blockmodel (SBM) [10], with many results regarding consistent recovery of the block assignments being known [8]. [2] and [23] showed that the distribution of the points of the ASE and LSE both converge to curved normal mixture distributions. Moreover, they demonstrated that clustering via the Expectation-Maximization (EM) algorithm [6] for a Gaussian mixture model (GMM) performs better than doing so via the K-means algorithm. However, their implementation of the EM algorithm failed to take into account the curved structure of the mixture's component distributions, and may have therefore resulted in an increased number of clustering errors.

This paper seeks to improve upon their results by introducing an Expectation-Solution (ES) algorithm [7] that makes full use of the ASE and LSE limiting distributions' curved-normal structure. We accomplish this by noting that the complete-data likelihood equations from the EM algorithm may be reinterpreted as complete-data estimating equations used by ES. Since each component of a curved-Gaussian mixture model (C-GMM) encodes its variance as a function of the component means and potentially the mixing proportions, the estimating equations corresponding to the component variances are rendered superfluous. This leads to an iterative scheme in which the component means and mixing proportions are updated as in EM for a full GMM, but the component variances are updated by simply plugging in the new means and proportions into the variance functions.

To the best of our knowledge this paper represents the first time the ES algorithm has been utilized for a C-GMM, and owing to its relative simplicity we cannot help but recommend it as a novel tool for the practitioner's toolbox, one that is not only limited in its usefulness to the spectral clustering problem. We structure the rest of the paper as follows: In Section 2 we review the setting and background for random dot product graphs and SBMs. In Section 3 we present the ES algorithm, first for a generic C-GMM, and then for both the ASE and LSE limiting distributions. In Sections 4 and 5 we present and discuss the results of our algorithm for simulated data from artificial SBMs, a brain connectome SBM, and a rank-deficient SBM.

1.1. Notation

Except where otherwise specified we use emboldened capital letters such as $\mathbf{A}, \mathbf{B}, \mathbf{X}$, etc. to denote matrices, emboldened and italicized capital letters with one indexed subscript to denote a row of the corresponding matrix as a column vector, as in $\mathbf{A}_i, \mathbf{B}_j, \mathbf{X}_k$. Emboldened lowercase Greek letters π, τ are vectors. We take $\mathbf{\Delta}^K := \{(\pi_1, \pi_2, \ldots, \pi_K) \in \mathbb{R}_{\geq 0}^K \mid \sum_{k=1}^K \pi_k = 1\}$ to be the unit simplex in \mathbb{R}^K , and δ_c to be the probability distribution assigning point mass to its parameter c in Euclidean space. The vector of all ones and all zeros in \mathbb{R}^n are given by $\mathbf{1}_n$ and $\mathbf{0}_n$, respectively. We omit subscripts where the dimension is obvious.

Given a real symmetric $n \times n$ matrix ${\bf M}$ we write its spectral decomposition as

$$\mathbf{M} = \mathbf{U}_{\mathbf{M}} \mathbf{D}_{\mathbf{M}} \mathbf{U}_{\mathbf{M}}^{\top},$$

where $\mathbf{U}_{\mathbf{M}}$ is a unitary matrix and $\mathbf{D}_{\mathbf{M}} = \operatorname{diag}(\lambda_1^{(\mathbf{M})}, \lambda_2^{(\mathbf{M})}, \dots, \lambda_n^{(\mathbf{M})})$ such that $\lambda_1^{(\mathbf{M})} \geq \lambda_2^{(\mathbf{M})} \geq \cdots \geq \lambda_n^{(\mathbf{M})}$ are the ordered eigenvalues of \mathbf{M} . We define the normalized Laplacian of \mathbf{M} as $\mathcal{L}(\mathbf{M}) = (\operatorname{diag}(\mathbf{M}\mathbf{1}_n)^{-\frac{1}{2}})\mathbf{M}(\operatorname{diag}(\mathbf{M}\mathbf{1}_n)^{-\frac{1}{2}})$, with spectral decomposition $\tilde{\mathbf{U}}_{\mathbf{M}}\tilde{\mathbf{D}}_{\mathbf{M}}\tilde{\mathbf{U}}_{\mathbf{M}}^{\mathsf{T}}$.

In Section 4 we use the shorthand

$$\{K\text{-means}, EM, ES\} \circ \{ASE, LSE\}$$
(1)

to refer to each clustering method; e.g., $EM \circ ASE$ is to be read as "EM for the ASE." Finally, in sections where we are explicitly concerned with GMMs, any reference to the EM algorithm is understood to refer to the EM algorithm for a full GMM as outlined in [9].

2. Background and setting

The K-block SBM encodes the probabilistic connectivity within and between disjoint subsets of graph nodes. For $K \ge 2$, $\boldsymbol{\pi} \in \boldsymbol{\Delta}^{K}$, and $\mathbf{B} \in (0,1)^{K \times K}$ a rank-*d* symmetric matrix with distinct rows, we write $(\mathbf{A}, \boldsymbol{\tau}) \sim \text{SBM}(n, \mathbf{B}, \boldsymbol{\pi})$ with sparsity factor $\rho_n \in (0, 1]$ provided the following:

$$\begin{aligned} \boldsymbol{\tau} &:= (\tau_1, ..., \tau_n) \\ \tau_i \stackrel{i.i.d.}{\sim} \operatorname{Cate}(\boldsymbol{\pi}) \\ A_{ij} \mid \boldsymbol{\tau} \stackrel{ind.}{\sim} \operatorname{Bern}(\rho_n B_{\tau_i \tau_j}), \ i < j \\ A_{ij} = A_{ji} \\ \operatorname{diag}(\mathbf{A}) = \mathbf{0}_n. \end{aligned}$$

Note τ denotes the vector of block memberships. If only **A** is observed, we write $\mathbf{A} \sim \text{SBM}(n, \mathbf{B}, \pi)$.

The sparsity factor ρ_n indexes a sequence of models where the edge probabilities change with n. We mention the sparsity factor purely for the sake of completeness, but common assumptions include taking $\rho_n \equiv 1$ for all n or $\rho_n \to 0$ such that $\rho_n > \log^4 n/n$. The former assumption is equivalent to the assumption that there exists c > 0 such that $\rho_n \to c$; [23] used the latter assumption to establish concentration in spectral norm of \mathbf{A} and $\mathcal{L}(\mathbf{A})$ around $\rho_n \mathbf{P}$ and $\mathcal{L}(\mathbf{P})$ (where \mathbf{P} is discussed below). As we are primarily concerned with spectral clustering of the nodes of a single graph at a time, the sparsity factor is taken to be identically 1, and thus suppressed, throughout the remainder of the paper.

The SBM is a special case of the random dot product graph (RDPG) [22], which encodes the probability that two nodes in a random graph share an

Z. M. Pisano et al.

edge as defined by the $n \times n$ symmetric edge probability matrix \mathbf{P} such that $A_{ij} \stackrel{ind.}{\sim} \operatorname{Bern}(P_{ij})$ for i < j [27]. In lieu of an edge probability matrix, one may instead define a random graph model via $\mathbf{X} \in \mathbb{R}^{n \times d}$ such that the row magnitudes are bounded above by 1 and the dot product between any two rows falls within the unit interval. In such case we permit the mathematical convenience of allowing for self loops, leading us to formulate a random graph model with edge probability matrix $\mathbf{P} = \mathbf{X}\mathbf{X}^{\top}$. We refer to the rows of \mathbf{X} as the *latent positions* of the graph model. Note that the latent positions are inherently unidentifiable, since for any $d \times d$ orthogonal matrix \mathbf{W} we have that $(\mathbf{X}\mathbf{W})(\mathbf{X}\mathbf{W})^{\top} = \mathbf{X}\mathbf{X}^{\top}$. One arrives at an analogous means of defining a rank-d, K-block SBM by defining $\boldsymbol{\nu}_1, \ldots, \boldsymbol{\nu}_K \in \mathbb{R}^d$ such that $\|\boldsymbol{\nu}_k\| \leq 1$ for all k and $\boldsymbol{\nu}_k^{\top} \boldsymbol{\nu}_{k'} \in (0, 1)$ for each k, k' pair. If we define $\mathbf{x} \in \mathbb{R}^{K \times d}$ such that the kth row is $\boldsymbol{\nu}_k$.

For $d \leq n$, the *d*-dimensional *adjacency spectral embedding* (ASE) of \mathbf{A} is given by the $n \times d$ matrix $\mathbf{\hat{X}} := \mathbf{U}_{\mathbf{A}}^{(d)} (\mathbf{D}_{\mathbf{A}}^{(d)})^{\frac{1}{2}}$, where the columns of $\mathbf{U}_{\mathbf{A}}^{(d)}$ are the first *d* columns of $\mathbf{U}_{\mathbf{A}}$ and $\mathbf{D}_{\mathbf{A}}^{(d)}$ is the $d \times d$ principal minor of $\mathbf{D}_{\mathbf{A}}$ consisting of the *d* largest ordered eigenvalues on the diagonal [13]. The *n* rows of $\mathbf{\hat{X}}$ can thus be thought of as a collection of points in \mathbb{R}^d that estimate the true latent positions up to orthogonal transformation. We can, therefore, touch upon the notion of how a random choice of an unobservable latent position \mathbf{X}_i (say, via some distribution F on \mathbb{R}^d), informs the distribution of the corresponding estimated latent position $\mathbf{\hat{X}}_i$. For a *K*-block SBM we have that the distribution F on the latent positions is a mixture of point masses, with mixture weights π_k . [2] obtained the following result as a corollary to a more general theorem regarding RDPGs and SBMs with $\rho_n \to 0$ such that $\rho_n > \log^4 n/n$.

Theorem 2.1. Let $X_i \stackrel{i.i.d.}{\sim} F = \sum_{k=1}^K \pi_k \delta_{\nu_k}$, where $\pi \in \Delta^K$. Define $\Lambda = \mathbb{E}[X_1 X_1^\top]$. Also let $\Sigma(x) = \Lambda^{-1} \mathbb{E}[X_1 X_1^\top (x^\top X_1 - x^\top X_1 X_1^\top x)] \Lambda^{-1}$. If $\rho_n \equiv 1$, there exists a sequence of orthogonal matrices W_n such that for any fixed index i

$$\sqrt{n}(\boldsymbol{W}_{n}\hat{\boldsymbol{X}}_{i}-\boldsymbol{X}_{i})|\boldsymbol{X}_{i}=\boldsymbol{\nu}_{k}\stackrel{a}{\rightarrow}\mathcal{N}(\boldsymbol{0},\boldsymbol{\Sigma}(\boldsymbol{\nu}_{k})).$$

That is to say that the optimally rotated rows \hat{X}_i of the ASE \hat{X} for a graph generated by an SBM are approximately distributed as a mixture of curved multivariate normal distributions, each centered at the corresponding true latent positions X_i of **P**. As mentioned above, we reiterate that these true latent positions are assumed to arise from a mixture of point masses on the scaled eigenvectors of **B**. Owing to the method by which the ASE is computed, we note that the rows are identically distributed but not independent.

The Laplacian spectral embedding (LSE) is to the normalized Laplacian $\mathcal{L}(\mathbf{A})$ as the ASE is to \mathbf{A} . That is, for $d \leq n$ the *d*-dimensional LSE is given by $\check{\mathbf{X}} = \tilde{\mathbf{U}}_{\mathbf{A}}^{(d)} (\tilde{\mathbf{D}}_{\mathbf{A}}^{(d)})^{\frac{1}{2}}$, where $\tilde{\mathbf{U}}_{\mathbf{A}}^{(d)}$ and $\tilde{\mathbf{D}}_{\mathbf{A}}^{(d)}$ are analogous to $\mathbf{U}_{\mathbf{A}}^{(d)}$ and $\mathbf{D}_{\mathbf{A}}^{(d)}$. Like $\hat{\mathbf{X}}$, the *n* rows of $\check{\mathbf{X}}$ can be thought of as a collection of points in \mathbb{R}^d . The LSE is preferred in settings where the adjacency matrix is sparse or the edge probability matrix is believed to be of the form $\rho_n \mathbf{P}_n$ for some positive sequence

 $\rho_n \to 0$ as $n \to \infty$, since the normalized Laplacian is the same for **P** as it is for ρ **P**; that is,

$$\mathcal{L}(\rho \mathbf{P}) = (\operatorname{diag}((\rho \mathbf{P})\mathbf{1}_n)^{-\frac{1}{2}})(\rho \mathbf{P})(\operatorname{diag}((\rho \mathbf{P})\mathbf{1}_n)^{-\frac{1}{2}})$$
$$= \rho^{-1}\rho(\operatorname{diag}(\mathbf{P}\mathbf{1}_n)^{-\frac{1}{2}})\mathbf{P}(\operatorname{diag}(\mathbf{P}\mathbf{1}_n)^{-\frac{1}{2}})$$
$$= \mathcal{L}(\mathbf{P}).$$

[23] obtained the following result for the LSE of an SBM, analogous to the previous theorem.

Theorem 2.2. Let the setting be as in Theorem 2.1, and also let $\mu = \mathbb{E}[X_1]$ and

$$ilde{\mathbf{\Lambda}} = \mathbb{E}igg[rac{oldsymbol{X}_1 oldsymbol{X}_1^{ op}}{oldsymbol{X}_1^{ op}oldsymbol{\mu}}igg].$$

Define

$$ilde{\Sigma}(m{x}) = \mathbb{E}igg[igg(rac{ ilde{m{\Lambda}}^{-1}m{X}_1}{m{X}_1^{ op}m{\mu}} - rac{m{x}}{2m{x}^{ op}m{\mu}}igg)igg(rac{m{X}_1^{ op} ilde{m{\Lambda}}^{-1}}{m{X}_1^{ op}m{\mu}} - rac{m{x}^{ op}}{m{x}^{ op}m{\mu}}igg)igg(rac{m{x}^{ op}m{X}_1 - m{x}^{ op}m{X}_1m{X}_1^{ op}m{x}igg)igg].$$

If $\rho_n \equiv 1$ and $n_k = |\{i \leq n | X_i = \boldsymbol{\nu}_k\}|$, then there exists a sequence of orthogonal matrices \boldsymbol{W}_n such that for any fixed index i

$$n\left(\boldsymbol{W}_{n}\check{X}_{i}-\frac{\boldsymbol{\nu}_{k}}{\sqrt{\sum_{l}n_{l}\boldsymbol{\nu}_{k}^{\top}\boldsymbol{\nu}_{l}}}\right)|X_{i}=\boldsymbol{\nu}_{k}\overset{d}{\rightarrow}\mathcal{N}(0,\tilde{\boldsymbol{\Sigma}}(\boldsymbol{\nu}_{k}))$$

The limiting distributions justify the fitting of full GMMs to cluster the rows of either spectral embedding. Doing so improves upon the clustering performance of K-means. Although K-means is invariant to orthogonal transformations of the data, the procedure imposes the assumption of equal, spherical covariances on the clusters; however, the ASE and LSE clusters can be non-spherical and of varying spread as demonstrated in Figures 12–13 below, as well as in [19] and [18]. By contrast, clustering via the EM algorithm accounts for the possibility of unequal covariances; however, the curvature of the components remains unaddressed.

In practice, both rank(**B**) and K are unknown. A principled method of estimating the former is to inspect the scree-plot of the singular values of **A** and look for "elbows" defining the cut-off between singular values corresponding to signal dimensions and those corresponding to noise dimensions [17]. One can estimate the latter by means of maximizing a fitness criterion penalized by model complexity, à la the Akaike Information Criterion (AIC) [1] and the Bayesian Information Criterion (BIC) [20]; however, recent work due to [26] has called into question the validity and applicability of these classical criteria to SBM settings.

For a comprehensive look at the proofs of Theorems 2.1 and 2.2, as well as fuller details of the setting we describe, we refer the reader to [3].

3. The expectation-solution algorithms

First introduced by [7], the Expectation-Solution (ES) algorithm arises as a generalization of the EM algorithm [6], where instead of updating the parameter estimates according to a collection of complete-data likelihood equations, we update the parameters as solutions to complete-data estimating equations. When the complete-data estimating equations for a given model coincide exactly with the complete-data likelihood equations for the same model, then the ES and EM algorithms are equivalent. Thus, any comparisons drawn between EM and ES for a particular setting are actually between two different ES algorithms, a fact which allows us to use comparable convergence criteria. We will demonstrate that implementation of ES in a curved-Gaussian mixture setting allows us to make full use of the mixture components' curved structure without sacrificing the relative simplicity of EM. We review ES for the general incomplete-data setting in Appendix I.

3.1. The ES algorithm for mixtures of curved gaussians

Let us ignore, for the moment, the content of Section 2 and consider a generic K-component Gaussian mixture model (GMM) supported on \mathbb{R}^d :

$$\boldsymbol{X}_{i} \overset{i.i.d.}{\sim} \sum_{k=1}^{K} \pi_{k} \mathcal{N}(\boldsymbol{\mu}_{k}, \boldsymbol{\Sigma}_{k}), i = 1, ..., n;$$

such that $\boldsymbol{\pi} := (\pi_1, ..., \pi_K) \in \boldsymbol{\Delta}^K$ and $(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \in \mathbb{R}^d \times \mathbb{M}_d$ for all k, where \mathbb{M}_d is the space of $d \times d$ positive definite matrices. Maximum likelihood estimation of the parameters upon observing $\boldsymbol{x}_1, ..., \boldsymbol{x}_n$ can be handily accomplished by the EM algorithm. Here, the component distributions are by themselves smooth exponential family distributions, a fact that allows us to utilize various results related to the computation of maximum likelihood estimates in the M-step.

Suppose instead that our model is

$$oldsymbol{X}_i \stackrel{i.i.d.}{\sim} \sum_{k=1}^K \pi_k \mathcal{N}(oldsymbol{\mu}_k, oldsymbol{\Sigma}_k(oldsymbol{\mu}_k))$$
 $oldsymbol{\mu}_k \in \mathcal{M} \subset \mathbb{R}^d$
 $oldsymbol{\Sigma}_k : \mathcal{M} o \mathbb{M}_d, orall k$

and each Σ_k is continuous and differentiable. Here each component distribution is a curved exponential family distribution, where the component variances are functions of the means [4]; hence we refer to such a setting as a curved-Gaussian mixture model (C-GMM). If we attempt to derive complete-data likelihood equations necessary to implement an EM algorithm, we may become swiftly enmired in non-linear equations which we must solve to update the means at each iteration, depending on the structure of each Σ_k .

Here the ES algorithm circumvents the potential difficulty posed by the variance functions. We begin by considering the natural unobserved data extension of $\mathbf{X} = (\mathbf{X}_1, ..., \mathbf{X}_n)$, which should come as no surprise to those familiar with the EM algorithm:

$$\boldsymbol{Z}_{i} = (Z_{i1}, ..., Z_{iK}) \overset{i.i.d.}{\sim} \text{Multinomial}(1, K, \boldsymbol{\pi})$$
$$\boldsymbol{X}_{i} \mid Z_{ik} = 1 \sim \mathcal{N}(\boldsymbol{\mu}_{k}, \boldsymbol{\Sigma}_{k}(\boldsymbol{\mu}_{k})).$$

Let Ψ denote the parameter vector consisting of π_1, \ldots, π_{K-1} and the entries of μ_1, \ldots, μ_K . In the complete-data setting (\mathbf{X}, \mathbf{Z}) we immediately note the following:

$$\mathbb{E}_{\Psi}[Z_{ik}] = \pi_k$$
$$\mathbb{E}_{\Psi}[Z_{ik}\boldsymbol{X}_i] = \pi_k \boldsymbol{\mu}_k,$$

which give rise to the natural estimating equations

$$\frac{\sum_{i=1}^{n} Z_{ik}}{n} - \pi_k = 0$$
$$\frac{\sum_{i=1}^{n} Z_{ik} \mathbf{X}_i}{n} - \pi_k \boldsymbol{\mu}_k = 0$$

solved by

$$\hat{\pi}_k = \frac{\sum_{i=1}^n Z_{ik}}{n};$$
$$\hat{\mu}_k = \frac{\sum_{i=1}^n Z_{ik} X_i}{\sum_{i=1}^n Z_{ik}}.$$

The complete-data estimating equation $U_c((\mathbf{X}, \mathbf{Z}), \Psi) = C_{\Psi}S(\mathbf{X}, \mathbf{Z}) + b_{\Psi}(\mathbf{X}) = 0$ (Appendix I), then, is characterized by

$$C_{\Psi} = \mathbf{I}_{(d+1)K-1}$$

$$S(\mathbf{X}, \mathbf{Z}) = \frac{1}{n} \begin{bmatrix} \sum_{i=1}^{n} Z_{i1} \\ \vdots \\ \sum_{i=1}^{n} Z_{i(K-1)} \\ \sum_{i=1}^{n} Z_{i1} \mathbf{X}_{i} \\ \vdots \\ \sum_{i=1}^{n} Z_{iK} \mathbf{X}_{i} \end{bmatrix}$$

$$b_{\Psi}(\mathbf{X}) = - \begin{bmatrix} \pi_{1} \\ \vdots \\ \pi_{K-1} \\ \pi_{1} \boldsymbol{\mu}_{1} \\ \vdots \\ (1 - \sum_{k=1}^{K-1} \pi_{k}) \boldsymbol{\mu}_{K} \end{bmatrix}$$

Z. M. Pisano et al.

In computing $h(\Psi|\Psi^*) = \mathbb{E}_{\Psi^*}[S(\mathbf{X}, \mathbf{Z})|\mathbf{X} = \mathbf{x}]$, we find that we must only obtain

$$egin{aligned} &Z_{ik}^* := \mathbb{E}_{\mathbf{\Psi}^{\star}}[Z_{ik}|oldsymbol{X}_i = oldsymbol{x}_i] \ &= rac{\pi_k^* \phi(oldsymbol{x}_i|oldsymbol{\mu}_k^*, oldsymbol{\Sigma}_k(oldsymbol{\mu}_k^*))}{\sum_{j=1}^K \pi_j^* \phi(oldsymbol{x}_i|oldsymbol{\mu}_j^*, oldsymbol{\Sigma}_j(oldsymbol{\mu}_j^*))} \end{aligned}$$

Thus $h(\Psi|\Psi^*) = S(\mathbf{x}, \mathbf{Z}^*)$ by linearity of expectation. The next iterate $\hat{\Psi}$ is then obtained by solving $C_{\Psi}S(\mathbf{x}, \mathbf{Z}^*) + b_{\Psi}(\mathbf{X}) = 0$, which yields

$$\hat{\pi}_{k} = \frac{\sum_{i=1}^{n} Z_{ik}^{*}}{n};$$
$$\hat{\mu}_{k} = \frac{\sum_{i=1}^{n} Z_{ik}^{*} x_{i}}{\sum_{i=1}^{n} Z_{ik}^{*}}$$

We present this procedure written concisely as Algorithm 3.1.

Algorithm 3.1 The ES Algorithm for the Curved-Gaussian Mixture Setting
1. Initialize
$$\Psi^* = \Psi_0$$
.

2. **E-Step:** For each i and k compute

$$Z_{ik}^* = \frac{\pi_k^* \phi(\boldsymbol{x}_i | \boldsymbol{\mu}_k^*, \boldsymbol{\Sigma}_k(\boldsymbol{\mu}_k^*))}{\sum_{j=1}^K \pi_j^* \phi(\boldsymbol{x}_i | \boldsymbol{\mu}_j^*, \boldsymbol{\Sigma}_j(\boldsymbol{\mu}_j^*))}.$$

3. S-Step: Compute the entries of $\hat{\Psi}$ as

$$\hat{\pi}_k = \frac{\sum_{i=1}^n Z_{ik}^*}{n}$$
$$\hat{\mu}_k = \frac{\sum_{i=1}^n Z_{ik}^* x_i}{\sum_{i=1}^n Z_{ik}^*}$$

4. Take $\Psi^* = \hat{\Psi}$.

5. Repeat steps 2–4 until some convergence criterion is satisfied.

Note that the only difference between this ES algorithm and the usual EM algorithm for GMMs is that we "update" the component variances by plugging the usual updates for the component means into the respective $\Sigma_k(\cdot)$ instead of computing

$$\hat{\mathbf{\Sigma}}_k = rac{\sum_{i=1}^n Z_{ik}^* (oldsymbol{x}_i - \hat{oldsymbol{\mu}}_k) (oldsymbol{x}_i - \hat{oldsymbol{\mu}}_k)^ op }{\sum_{i=1}^n Z_{ik}^*}$$

for each k. As a result we estimate only (d+1)K-1 parameters, whereas in the full GMM setting we estimate another $K(d + \binom{d}{2})$ parameters comprising the component covariances. Since the model complexity penalties used to compute both AIC and BIC increase in magnitude with the number of parameters, the severe reduction in the number of parameters needed to be estimated in a C-GMM can vastly decrease these penalties. That is, if we are deciding between a GMM and C-GMM with approximately equal likelihoods for a given collection

of data, then it is clear that the C-GMM will achieve the higher (and therefore more desirable) AIC or BIC.

If we are interested in clustering the observations we perform the usual GMM clustering procedure, which is to assign each observation to the cluster with the highest posterior probability.

Remark 1. We note here the distinction between two classes of C-GMMs. We call a C-GMM **separable** if each component variance is solely a function of its respective mean and possibly its mixing proportion. Likewise, if one or more of the variances take other components' means or mixing proportions as arguments, then that C-GMM is **tied** — as in the cases of both spectral embeddings' limiting distributions.

3.2. ES algorithm for the ASE

Let us return to the content of Section 2, and assume the setting of Theorem 2.1 with $\mathbf{B} = \mathbf{x}\mathbf{x}^{\top}$, where the k^{th} row of \mathbf{x} is $\boldsymbol{\nu}_k$, and that we have observed \mathbf{A} and computed its *d*-dimensional ASE $\hat{\mathbf{X}}$. Since $F = \sum_{k=1}^{K} \pi_k \delta_{\boldsymbol{\nu}_k}$, we have that $\mathbf{\Lambda} = \sum_{k=1}^{K} \pi_k \boldsymbol{\nu}_k \boldsymbol{\nu}_k^{\top}$ and

$$\boldsymbol{\Sigma}(\boldsymbol{\nu}_k) = \boldsymbol{\Lambda}^{-1} \bigg(\sum_{j=1}^K \pi_j \boldsymbol{\nu}_j \boldsymbol{\nu}_j^\top (\boldsymbol{\nu}_k^\top \boldsymbol{\nu}_j - \boldsymbol{\nu}_k^\top \boldsymbol{\nu}_j \boldsymbol{\nu}_j^\top \boldsymbol{\nu}_k) \bigg) \boldsymbol{\Lambda}^{-1}.$$

The variance function thus takes every latent position and mixture weight as arguments, and our iterative scheme must reflect this. Therefore, let $\Sigma(\cdot|\mathbf{x}, \pi)$ denote the covariance function instead. The normal mixture model we thus consider is

$$\hat{\boldsymbol{X}}_i \sim \sum_{k=1}^K \pi_k \mathcal{N}\left(\boldsymbol{\nu}_k, \frac{\boldsymbol{\Sigma}(\boldsymbol{\nu}_k | \mathbf{x}, \boldsymbol{\pi})}{n}\right),$$

where each $\hat{\mathbf{X}}_i$ is identically distributed. Note that we have dropped the optimal orthogonal transformation \mathbf{W}_n in the statement of Theorem 2.1; hence the algorithm (presented below as Algorithm 3.2) estimates the latent positions up to this rotation, but this has no bearing on the task of clustering the graph nodes or estimating the block probability matrix and membership probabilities.

Remark 2. We can construct a separable analogue of Algorithm 3.2 by updating each variance with the newest iterate of the corresponding mean and proportion while holding all other arguments as their previous iterates, only updating them every ι iterations, or holding them constant. Such schemes (particularly the last) greatly alter the model at hand, hence we would not recommend them for use; rather, we only mention them to fill in the middle ground between Algorithms 3.1 and 3.2.

Algorithm 3.2 The ES Algorithm for the ASE

1. Initialize $\Psi^* = \Psi_0$.

2. E-Step: Compute

$$Z_{ik}^* = \frac{\pi_k^* \phi(\hat{\boldsymbol{X}}_i | \boldsymbol{\nu}_k^*, \boldsymbol{\Sigma}(\boldsymbol{\nu}_k^* | \mathbf{x}^*, \boldsymbol{\pi}^*)/n)}{\sum_{j=1}^K \pi_j^* \phi(\hat{\boldsymbol{X}}_i | \boldsymbol{\nu}_j^*, \boldsymbol{\Sigma}(\boldsymbol{\nu}_j^* | \mathbf{x}^*, \boldsymbol{\pi}^*)/n)}$$

3. S-Step: Compute

$$\begin{split} \hat{\pi}_k &= \frac{\sum_{i=1}^n Z_{ik}^*}{n} \\ \hat{\nu}_k &= \frac{\sum_{i=1}^n Z_{ik}^* \hat{X}_i}{\sum_{i=1}^n Z_{ik}^*} \end{split}$$

4. Take $\Psi^* = \hat{\Psi}$.

5. Repeat steps 2-4 until some convergence criterion is satisfied.

3.3. ES algorithm for the LSE

Let us assume the setting of Theorem 2.2. We have

$$\boldsymbol{\mu} = \sum_{k=1}^{K} \pi_k \boldsymbol{\nu}_k$$
$$\tilde{\boldsymbol{\Lambda}} = \sum_{k=1}^{K} \pi_k \frac{\boldsymbol{\nu}_k \boldsymbol{\nu}_k^{\top}}{\boldsymbol{\nu}_k^{\top} \boldsymbol{\mu}}$$
$$\tilde{\boldsymbol{\Sigma}}(\boldsymbol{\nu}_k | \mathbf{x}, \boldsymbol{\pi}) = \sum_{j=1}^{K} \pi_j \left(\frac{\tilde{\boldsymbol{\Lambda}}^{-1} \boldsymbol{\nu}_j}{\boldsymbol{\nu}_j^{\top} \boldsymbol{\mu}} - \frac{\boldsymbol{\nu}_k}{2\boldsymbol{\nu}_k^{\top} \boldsymbol{\mu}} \right) \left(\frac{\boldsymbol{\nu}_j^{\top} \tilde{\boldsymbol{\Lambda}}^{-1}}{\boldsymbol{\nu}_j^{\top} \boldsymbol{\mu}} - \frac{\boldsymbol{\nu}_k^{\top}}{\boldsymbol{\nu}_k^{\top} \boldsymbol{\mu}} \right) \left(\frac{\boldsymbol{\nu}_k^{\top} \boldsymbol{\nu}_j - \boldsymbol{\nu}_k^{\top} \boldsymbol{\nu}_j \boldsymbol{\nu}_j^{\top} \boldsymbol{\nu}_k}{\boldsymbol{\nu}_k^{\top} \boldsymbol{\mu}} \right).$$

The normal mixture model we consider here is

$$\check{\boldsymbol{X}}_i \sim \sum_{k=1}^K \pi_k \mathcal{N}\bigg(\frac{\boldsymbol{\nu}_k}{\sqrt{\sum_l n_l \boldsymbol{\nu}_l^\top \boldsymbol{\nu}_k}}, \frac{\tilde{\boldsymbol{\Sigma}}(\boldsymbol{\nu}_k | \mathbf{x}, \boldsymbol{\pi})}{n^2}\bigg),$$

where each \mathbf{X}_i is identically distributed. Owing to the presence of the unobserved n_k in the component means, the fact that the ASE and LSE are defined by the 1-1 transformation

$$\check{\mathbf{X}} = (\operatorname{diag}(\mathbf{A}\mathbf{1}_n)^{-1/2})\hat{\mathbf{X}},\tag{2}$$

and the fact that the LSE covariance function takes the ASE component means as arguments, we cannot simply invoke an analogue of the above algorithm. We propose expanding the parameter vector Ψ by treating the n_k as parameters that can be obtained by the K estimating equations

$$\mathbb{E}_{\Psi}\big[\sum_{i=1}^{n} Z_{ik}\big] - n_k = 0,$$

which would be solved by $\hat{n}_k := \sum_{i=1}^n Z_{ik} = n\hat{\pi}_k$ if we observed the true cluster labels of each row of $\check{\mathbf{X}}$.

Upon observing **A** and computing both $\hat{\mathbf{X}}$ and $\check{\mathbf{X}}$, we propose the following algorithm:

Algorithm 3.3 The ES Algorithm for the LSE

1. Initialize $\Psi^* = \Psi_0$ and take $\Sigma_k(\cdot) = \tilde{\Sigma}(\cdot | \mathbf{x}, \pi) / n^2$.

2. E-Step: Compute

$$\begin{split} \boldsymbol{\mu}_{k}^{*} &= \frac{\boldsymbol{\nu}_{k}^{*}}{\sqrt{\sum_{l} n_{l}^{*} \boldsymbol{\nu}_{l}^{*T} \boldsymbol{\nu}_{k}^{*}}} \\ Z_{ik}^{*} &= \frac{\pi_{k}^{*} \phi(\tilde{\mathbf{X}}_{i} | \boldsymbol{\mu}_{k}^{*}, \tilde{\boldsymbol{\Sigma}}(\boldsymbol{\nu}_{k}^{*} | \mathbf{x}^{*}, \boldsymbol{\pi}^{*}) / n^{2})}{\sum_{j=1}^{K} \pi_{j}^{*} \phi(\tilde{\mathbf{X}}_{i} | \boldsymbol{\mu}_{j}^{*}, \tilde{\boldsymbol{\Sigma}}(\boldsymbol{\nu}_{j}^{*} | \mathbf{x}^{*}, \boldsymbol{\pi}^{*}) / n^{2})}. \end{split}$$

3. S-Step: Compute

$$\hat{\pi}_k = \frac{\sum_{i=1}^n Z_{ik}^*}{n}$$
$$\hat{\nu}_k = \frac{\sum_{i=1}^n Z_{ik}^* \hat{X}_i}{\sum_{i=1}^n Z_{ik}^*}$$
$$\hat{n}_k = n\hat{\pi}_k.$$

4. Take $\Psi^* = \hat{\Psi}$.

5. Repeat steps 2-4 until some convergence criterion is satisfied.

Note that this algorithm makes full use of both the LSE and ASE; the \dot{X}_i are used in the E-Step, but the \hat{X}_i are used in the S-Step. Even so, clustering is to be done based on the rows of \check{X} , as they correspond exactly to the posterior probabilities computed while the algorithm runs its course. This is due to the fact that in our implementation we specify that the conditional distribution of the Z_{ik} depend purely on \check{X}_i , meanwhile the estimating equations for the ν_k remain as in the previous algorithm.

Remark 3. A major advantage of Algorithms 3.2 and 3.3 over the usual EM algorithm for a GMM is lower computational complexity to update the component variances at each iteration. Simple inspection of the M-step in the classical method reveals that the number of operations needed to update each of the K component variances is O(nd). In our ES algorithms the number of operations to perform the same task is $O(d^3K)$ for the ASE and $O(d^4K)$ for the LSE; the latter complexity may be shaved down to $O(d^3K)$ by preserving initial moment estimates of μ and $\overline{\Lambda}$ (and computing the latter's inverse) as described in Section 5. Since $d \ll n$ these ES algorithms possess far less complexity than the usual EM.

Remark 4. While Algorithms 3.2 and 3.3 do address the curvature of the spectral embeddings' limiting mixture distributions, the dependence of the estimated latent positions remains unaddressed, a shortcoming shared by the EM algorithm in these settings. Both ES and EM require the assumption of independence to carry out the E-step, since the conditional expectation of each Z_{ik} is actually

given all the observed data, not just the corresponding \hat{X}_i or \check{X}_i .

4. Simulations

For each simulation setting we generated 100 graphs and compared the performance of our ES algorithms with their respective EM analogues. The specific EM algorithm we used was the function em from the R package Mclust [21]. All four procedures were initialized at the true parameter values. Since EM is a maximum likelihood procedure, the default convergence criterion is to terminate when updates to the log-likelihood are less than 1×10^{-5} . By contrast, there is no convenient objective function associated with the ES algorithm, since the iterates arise as solutions to complete-data estimating equations; moreover, it was observed anecdotally in a few of our settings that the log-likelihood may decrease after an ES iteration. With that in mind, we noted that the EM algorithm actually is an ES algorithm where the complete-data estimating equations are the complete-data likelihood equations; and we altered the code of em to reflect this. Therefore convergence of EM and ES for the ASE was assumed when the Euclidean distance between successive estimates of the parameter vector Ψ consisting of the mixing proportions and the entries of **x** was less than 1×10^{-6} . and convergence of EM and ES for the LSE was assumed when the Euclidean distance between successive estimates of the parameter vector $\tilde{\Psi}$ consisting of the mixing proportions and the scaled latent positions was less than 1×10^{-7} for the LSE, or the number of iterations exceeded 10,000. The lower threshold for $ES \circ LSE$ was based on the fact that in computing the LSE we effectively scale the rows of the ASE towards the origin (see, e.g., Figures 12–13 below), thereby decreasing the magnitude of the latent position iterates and the extent to which they can change significantly between each pass of the algorithms.

To evaluate clustering performance we computed the adjusted Rand index (ARI) [11] for the cluster assignments from each method with the true cluster labels. For each clustering method m found in (1), we let ARI_m denote the ARI produced by that method. In the non-graph setting below (Subsection 4.1) for $\ell = A, L$ we performed Wilcoxon rank-sum tests [15], with hypotheses

 $H_o : \operatorname{med}(ARI_{EM \circ \ell SE}) = \operatorname{med}(ARI_{ES \circ \ell SE})$ $H_a : \operatorname{med}(ARI_{EM \circ \ell SE}) \neq \operatorname{med}(ARI_{ES \circ \ell SE}).$

We present these in the form of 95% confidence intervals for the median of each $\Delta_{\ell} := ARI_{EM\circ\ell SE} - ARI_{ES\circ\ell SE}$.

To evaluate accuracy of the parameter estimates, we computed the squared error of each method's terminating estimate from the true parameter vector Ψ . As will be seen in Subsection 4.2 we circumvented the issue of the latent positions' non-identifiability by rotating and centering our simulated data over the "canonical" latent postions in the first orthant by way of a Procrustes transformation. We take Ψ_{ASE} to consist of the entries of π and \mathbf{x} , as well as those of the covariance matrices $\frac{\Sigma(\nu_k | \mathbf{x}, \pi)}{n}$. As $EM \circ LSE$ does not outright produce estimates for the $\boldsymbol{\nu}_k$, we take Ψ_{LSE} to consist of the entries of π , $\boldsymbol{\mu}_k := \frac{\boldsymbol{\nu}_k}{\sqrt{\sum_l n_l \boldsymbol{\nu}_l^\top \boldsymbol{\nu}_k}}$,



FIG 1. Plots in the upper and lower rows correspond to models 1 and 2, respectively, in Section 4.1. The left-hand plots depict 95% confidence bands based on WIIcoxon rank-sum tests; median ARI differences below the dashed line (marking a difference of 0) indicate that ES tended to cluster more accurately than EM. The right hand plots depict p_A and p_L , with the dashed line representing $\alpha = 0.05$.

and the covariance matrices $\frac{\tilde{\Sigma}(\nu_k|\mathbf{x},\pi)}{n^2}$. We then performed Wilcoxon rank sum tests for the paired collection of squared errors for $\{EM, ES\} \circ ASE$, as well as $\{EM, ES\} \circ LSE$. We let p_A denote the *p*-value corresponding to the test

$$\begin{aligned} H_o: \operatorname{med}(\|\hat{\Psi}_{EM\circ ASE} - \Psi_{ASE}\|) &\leq \operatorname{med}(\|\hat{\Psi}_{ES\circ ASE} - \Psi_{ASE}\|) \\ H_a: \operatorname{med}(\|\hat{\Psi}_{EM\circ ASE} - \Psi_{ASE}\|) > \operatorname{med}(\|\hat{\Psi}_{ES\circ ASE} - \Psi_{ASE}\|), \end{aligned}$$

and p_L be that corresponding to the test

$$H_o: \operatorname{med}(\|\hat{\Psi}_{EM\circ LSE} - \Psi_{LSE}\|) \le \operatorname{med}(\|\hat{\Psi}_{ES\circ LSE} - \Psi_{LSE}\|)$$

Z. M. Pisano et al.





FIG 2. Plots in the upper and lower rows correspond to models 3 and 4, respectively, in Section 4.1. The plot may be interpreted analogously to Figure 1.

$$H_a: \operatorname{med}(\|\widehat{\Psi}_{EM \circ LSE} - \Psi_{LSE}\|) > \operatorname{med}(\|\widehat{\Psi}_{ES \circ LSE} - \Psi_{LSE}\|).$$

4.1. Non-graph setting

We first tested $ES \circ \{ASE, LSE\}$ on random data generated directly from the mixture distributions given in Theorems 2.1 and 2.2. We considered the following models:

$$\mathbf{x}_{1} = \begin{bmatrix} 0.6210 & 0.3382\\ 0.3382 & 0.6210 \end{bmatrix} \qquad \mathbf{x}_{2} = \begin{bmatrix} 0.4076 & 0.1840\\ 0.1840 & 0.4076 \end{bmatrix} \\ \mathbf{x}_{3} = \begin{bmatrix} 0.6024 & 0.3703\\ 0.3703 & 0.5319 \end{bmatrix} \qquad \mathbf{x}_{4} = \begin{bmatrix} 0.3962 & 0.2074\\ 0.2074 & 0.3721 \end{bmatrix}$$

with $\boldsymbol{\pi} = (\frac{1}{2}, \frac{1}{2})$ in all cases. These particular choices of \mathbf{x} correspond to, respectively, balanced affinity models given by (a, b) = (0.5, 0.4) and (a, b) = (0.2, 0.15), and balanced core-periphery models given by (a, b) = (0.5, 0.42) and (a, b) = (0.2, 0.15) which are included in the simulations in the following two subsections. Each randomly generated sample was taken to comprise $\hat{\mathbf{X}}$ and the simulated LSE $\check{\mathbf{X}}$ was computed via the 1–1 relationship (2), taking $\mathbf{A} = \hat{\mathbf{X}} \hat{\mathbf{X}}^{\top}$.

For models 1 and 2 (Figure 1) the EM algorithms approximately match the ES algorithms' performance for large n. Due to the sample size n featuring in the denominator of the component variances, the clusters of both mixtures shrink around the component means as n increases, pulling the points into tight, distinct clusters. EM tended to cluster more accurately than ES in models 3 and 4, and we credit this to the fact that in these settings the point clouds overlap significantly for small to moderate n, and it is not until around n = 700 that distinct clusters start to form.

In model 1 the ES algorithms tend to more accurately estimate the component means, variances, and weights than do the EM algorithms. As n increases the accuracy of the EM algorithms begins to match that of the ES algorithms; again, we suspect that this is due to the emergence of distinct clusters. In model 2, the EM algorithms perform more accurate estimation for small n, but are overtaken by the ES algorithms as n increases. Owing to the fact that \mathbf{x}_2 arises as the latent positions for a sparse SBM (relative to that corresponding to \mathbf{x}_1), we suspect that as n increases beyond 900 parameter estimation by either algorithm will be comparable as this model's clusters further separate.

In models 3 and 4 (Figure 2) the ES algorithms tend to cluster the points less accurately than their EM counterparts, particularly so when n is relatively small. Moreover, in these settings the ES algorithms largely fail to estimate the parameters more accurately than the EM algorithms, save for around n = 800 in model 3. Due to the decreasing trend in the confidence bounds as n increases, as well as the fact that \mathbf{x}_4 corresponds to a sparser SBM than does \mathbf{x}_3 , we suspect that this phenomenon is replicated at a range of n > 900 where the ES algorithms outperform the EM algorithms in model 4 as indicated by the downward trend of the p-values in the bottom-right plot.

4.2. Balanced affinity network structure

A K-block SBM is said to possess homogeneous balanced affinity structure if $B_{ii} = a$ for all $i, B_{ij} = b$ for all $i \neq j, 0 < b < a < 1$, and $\pi_k = \frac{1}{K}$ for all k [5]. We generated 100 graphs of size n = 200, 500 for each $(a, b) \in \{0.01, 0.02, \ldots, 0.99\}^2$ satisfying b < a and compared the ARI performance of each clustering method via upper-tailed classical sign tests ($\alpha = 0.025$) due to the presence of multiple ties in several settings, with the null hypotheses:

$$H_o: P(ARI_{ES\circ ASE} \ge ARI_{EM\circ ASE}) = 0.5$$
$$H_o: P(ARI_{ES\circ LSE} \ge ARI_{EM\circ LSE}) = 0.5$$
$$H_o: P(ARI_{ES\circ ASE} \ge ARI_{ES\circ LSE}) = 0.5$$

Z. M. Pisano et al.

$$H_o: P(ARI_{ES\circ LSE} \ge ARI_{ES\circ ASE}) = 0.5.$$

The use of sharp inequalities in the proportions in question was motivated by the fact that multiple choices of (a, b) result in mixtures in which perfect recovery of the cluster labels is possible by both ES and EM algorithms; moreover, if an ES algorithm manages the same ARI as its corresponding EM algorithm it does so with lower computational complexity as noted in Remark 3.

Since the latent position matrix \mathbf{x} is only identifiable up to orthogonal transformation, for each choice of (a, b) we took the canonical latent position matrix to be that centered in the first quadrant of \mathbb{R}^2 via the transformation

$$\mathbf{x} = \boldsymbol{U}_{\boldsymbol{B}} \boldsymbol{D}_{\boldsymbol{B}}^{\frac{1}{2}} \boldsymbol{U}_{\boldsymbol{B}}^{\top}.$$

For each model we used the R package *igraph* to sample from the SBM with the class assignments fixed, computed the ASE $\hat{\mathbf{X}}$ from its definition, then centered the rows over the canonical latent positions with the $d \times d$ orthogonal transformation $\hat{\mathbf{W}}$ that solved the Procrustes problem

$$\min_{\mathbf{W}} \| \hat{\mathbf{X}} \mathbf{W} - \mathbf{X} \|_F$$

subject to $\mathbf{W}^\top \mathbf{W} = \mathbf{I}$

where

$$\mathbf{X} = \begin{bmatrix} \mathbf{1}_{n_1} \otimes \boldsymbol{\nu}_1^\top \\ \mathbf{1}_{n_2} \otimes \boldsymbol{\nu}_2^\top \\ \vdots \\ \mathbf{1}_{n_K} \otimes \boldsymbol{\nu}_K^\top \end{bmatrix},$$
(3)

with $n_k := |\{i \le n | \tau_i = k\}|$. Following this the LSE **X** was computed via the 1–1 relationship (2).

The results are displayed in Figures 3–4. Both here and in subsequent simulations, EM tends to dominate when (a, b) lies close to the identity line a = b (i.e., when the simulated model is almost an Erdös-Rényi random graph [3]); but ES tends to dominate as the model moves away from the identity line until the true latent positions are so far apart that all methods tend to have equal performance. The red "strictly dominant" regions shrink as n increases from 200 to 500 due to the clusters shrinking around the component means, leading to larger black regions where the ES algorithms cluster at least as well as EM but with the benefit of lower computational complexity at each iteration.

We also repeated the simulations on \mathbf{x}_1 and \mathbf{x}_2 from the previous subsection, this time generating the data from balanced affinity SBMs defined by $(a,b) = (0.5, 0.4) \pmod{1}$ and $(a,b) = (0.2, 0.15) \pmod{2}$ on graphs of size $n = 200, 300, \ldots, 900$. Overall we observe in Figure 5 similar results to those presented in Figure 1; however, it's clear that the degree to which ES out-clusters EM has increased, particularly in model 2.



FIG 3. Black regions denote settings in which the corresponding sharp sign test in Subsection 4.2 was significant, meanwhile red regions denote significance of sign tests in which the inequality is strict. The blue regions in the upper row of plots within each figure indicate where EM strictly out-performed the corresponding ES algorithm as determined by analogous lower-tailed sign tests. have been slightly enlarged for emphasis and ease of reading. All sign tests were conducted at level $\alpha = 0.025$. Figures 4, 6–7, and 9–10 may be interpreted similarly.

4.3. Core-periphery network structure

A 2-block SBM is said to possess core-periphery structure if $\mathbf{B} = \begin{bmatrix} a \\ b \end{bmatrix}^{*}$ and $\boldsymbol{\pi} = (\pi_1, 1 - \pi_1)$ [5]. We considered the 2-block balanced core-periphery SBMs characterized by the same grid of (a, b) pairs in the balanced affinity simulations and repeated the previous experiment. The results are displayed in Figures 6–7. As in the balanced affinity results we observe that $EM \circ ASE$ tends to outperform $ES \circ ASE$ near the identity line, but the latter tends dominate as (a, b) moves further away from equal entries before the component means move



sufficiently far apart and $ES \circ ASE$ performs at least as well as $EM \circ ASE$ with lower computational complexity. In this setting, $ES \circ LSE$ tends to dominate $EM \circ LSE$ near the identity line except for exceptionally dense models. The figures also illustrate the regions in which $ES \circ ASE$ and $ES \circ LSE$ strictly dominate each other. We see that as n increases the regions of strict dominance shrink, once again, due to the clusters shrinking around the component means.

We also repeated the simulations from Section 4.1 corresponding to \mathbf{x}_3 and \mathbf{x}_4 , which respectively arise as the latent position matrices of balanced coreperiphery models with (a, b) = (0.5, 0.42) and (a, b) = (0.2, 0.15). The results displayed in Figure 8 indicate that, as in the mixture setting, ES largely fails to outperform EM, except for small to moderate n in model 3. However, the CIs in Figure 2 indicate that EM tended to strictly outperform ES for both embeddings; but in Figure 8, $ES \circ LSE$ is approximately on par with $EM \circ$



Results for SBMs 1 and 2

FIG 5. Plots in the upper and lower rows correspond to models 1 and 2, respectively, in Section 4.2, and may thus be interpreted analogously to Figures 1-2.

LSE. Both here and in the mixture setting, it was observed that ES tended to vastly overestimate the entries of the covariance matrices, hence the algorithms' seeming inability to more accurately estimate Ψ , except for $ES \circ LSE$ in model 3 when $700 \leq n \leq 900$.

4.4. Unbalanced core-periphery model

We have to this point only examined the performance of ES against that of EM for balanced 2-block SBMs, in which the blocks have equal probability of node membership. Since this case obviously fails to encapsulate a wide variety of models, we also evaluated the performance of the clustering procedures for unbalanced core-periphery model in which the more densely connected core pos-



sesses a membership probability of either 0.25 ("Small Core") and 0.75 ("Large Core") in graphs of size n = 200.

The results are presented in Figures 9-10. We see that the ES algorithms tend to dominate EM in sparse settings when the core is small, but when the core is large EM tends to dominate in a larger region near the diagonal than in balanced settings. These plots along with Figure 6 also illustrate how regions in which the ES algorithms strictly dominate each other shift as the core proportion changes.

4.5. Synthetic analysis of MRI connectome SBM

[18] investigated clustering via the EM algorithm for the ASE and LSE of a 4block SBM used to model neural connectivity in the human brain. They noted that the 2-dimensional LSE captures left-hemisphere/right-hemisphere connec-



tivity and the 2-dimensional ASE captures gray-matter/white-matter connectivity. That is, clustering according to either embedding resulted in one of "two truths."

To compare clustering performance for this setting, we generated data from the 4-block connectome estimated by [18], i.e., the model $\mathbf{A} \sim \text{SBM}(\mathbf{B}, \pi)$ where $\pi = (0.28, 0.22, 0.28, 0.22)$, **B** is the rank-4 matrix

$$\mathbf{B} = \begin{bmatrix} 0.020 & 0.044 & 0.002 & 0.009 \\ 0.044 & 0.115 & 0.010 & 0.042 \\ 0.002 & 0.010 & 0.020 & 0.045 \\ 0.009 & 0.042 & 0.045 & 0.117 \end{bmatrix}$$

and blocks 1–4 correspond to left/gray, left/white, right/gray, and right/white neurons, respectively. The eigendecomposition $\mathbf{B} = \mathbf{U}_{\mathbf{B}} \mathbf{D}_{\mathbf{B}} \mathbf{U}_{\mathbf{B}}^{\top}$ gives the latent



Results for SBMs 3 and 4

FIG 8. Plots in the upper and lower rows correspond to models 3 and 4, respectively, in Section 4.3, and may thus be interpreted analogously to Figures 1-2, 5.

position matrix (after rotating properly to the first orthant in \mathbb{R}^4) as

$$\mathbf{x} = \mathbf{B} = \mathbf{U}_{\mathbf{B}} \mathbf{D}_{\mathbf{B}}^{\frac{1}{2}} \mathbf{U}_{\mathbf{B}}^{\top} = \begin{bmatrix} \boldsymbol{\nu}_{1}^{\top} \\ \boldsymbol{\nu}_{2}^{\top} \\ \boldsymbol{\nu}_{3}^{\top} \\ \boldsymbol{\nu}_{4}^{\top} \end{bmatrix} = \begin{bmatrix} 0.0915 & 0.1076 & 0.0057 & 0.0034 \\ 0.1076 & 0.3149 & 0.0056 & 0.0649 \\ 0.0057 & 0.0056 & 0.0886 & 0.1099 \\ 0.0034 & 0.0649 & 0.1099 & 0.3173 \end{bmatrix}$$

For each $n = 500, 600, \ldots, 1200$ we generated 100 graphs, computed their ASEs and LSEs, carried out the procedures as otherwise described, and output the results to Figure 11. We opted not to compare accuracy of parameter estimation, since we observed that our algorithms tended to vastly overestimate the covariances as described in the previous subsection. Here ES tended to more successfully cluster than EM for all values of n.





4.6. Rank-deficient SBM

All of the settings in which we have conducted our simulations thus far share the property of the true rank d being equal to the number of blocks K; of course, one may quite naturally give an example of SBM in which d < K. To this end, consider the 6-block SBM given by

	0.4986	0.5982	0.5982	0.5271	0.5424	0.2109
	0.5982	0.7684	0.7154	0.7272	0.6408	0.3928
р	0.5982	0.7154	0.7178	0.6281	0.6512	0.2467
$\mathbf{D} =$	0.5271	0.7272	0.6281	0.7345	0.5548	0.4843
	0.5424	0.6408	0.6512	0.5548	0.5920	0.2020
	0.2109	0.3982	0.2467	0.4843	0.2020	0.4745





and $\boldsymbol{\pi} = (0.19, 0.18, 0.14, 0.19, 0.12, 0.18)$. The block probability matrix is rank-2, and may be exactly rewritten as $\mathbf{B} = \mathbf{x}\mathbf{x}^{\top}$, where

	0.69	0.15
	0.78	0.40
· -	0.83	0.17
$\mathbf{x} =$	0.64	0.57
	0.76	0.12
	0.16	0.67

Unlike the above connectome SBM this particular setting was not motivated by a particular application. In fact, we randomly generated this model by sampling the six rows of \mathbf{x} uniformly from the intersection of the 2-dimensional unit disc and the positive quadrant in \mathbb{R}^2 and then generating $\boldsymbol{\pi}$ from a flat Dirichlet



95% CIs of Median Difference in ARI (Brain SBM)

Fig 11.

distribution (with everything rounded to two decimal places). We have provided examples of the ASE and LSE of a graph generated from this setting (n = 600) in Figures 12–13.

We repeated the experiment of the connectome SBM, this time electing to generate 1000 graphs for each n, and output the results to Figure 14. One can see that both ES algorithms dominate their corresponding EM for this particular setting. We include the results of additional rank-deficient simulations in Appendix III.

5. Discussion

We have described an algorithm that estimates the parameters of a curvednormal mixture model and accounts for the components' curvature. The algorithm itself is an adaptation of the usual EM algorithm for smooth-normal



ASE of Rank-2 6-Block SBM

FIG 12. The smaller points representing rows of the 2-dimensional ASE have been colored according to block membership of the corresponding nodes in the original graph. The larger orange points encircled with a black outline denote the true latent positions, i.e., the rows of X.

mixture models, where instead of updating the component variance estimates in the usual way we simply plug the usual updates to the component means into the variance function(s). Even though we developed this algorithm purely for the purpose of spectral clustering for SBMs, we suspect that we can easily adapt it to mixtures of curved exponential families in which the component distributions are not normal. We hope to explore this, as well as sufficient conditions for consistent and asymptotically normal estimates as outlined in the appendix of [7], in a future paper.

The simulation results demonstrate that taking into account the curvature of the spectral embeddings' limiting distributions allows us to improve node clustering for SBMs, and — in some settings — by a vast margin. In general



LSE of Rank-2 6-Block SBM

FIG 13. The smaller points representing rows of the 2-dimensional LSE have been colored according to block membership of the corresponding nodes in the original graph (as in the previous figure). Here the larger orange points denote the scaled Laplacian latent positions, *i.e.*, the component means of the LSE's limiting distribution.

our results indicate that EM dominates ES in settings close to Erdös-Rényi models, but ES overtakes EM in models sufficiently different from such a simple setting. In particular, the dominance our algorithms display in outperforming EM for the brain connectome model lead us to highly recommend our method for that application. Moreover, there existed at least one sample size for which at least one of our proposed algorithms clustered significantly more accurately than vanilla EM in all but one of the four SBMs which we singled out in our Subsections 4.2 and 4.3. Since we initialized all simulations at the true parameter values, we have omitted any discussion of sufficient conditions for local or global convergence of $ES \circ \{ASE, LSE\}$. However, we observed anecdotally in the brain connectome and rank-deficient settings that initializing ES far from



95% CIs of Median Difference in ARI (Low-Rank SBM)

Fig 14.

the canonical latent positions still resulted in accurate clusters. The default implementation of mclust initializes EM at the terminating values of a hierarchical agglomerative clustering procedure [21], which may also prove practically suitable for ES in future work.

Though the simulations are presented through the lens of evaluating ES against EM, we also compared both methods to K-means (Appendix II). When all three algorithms were compared to each other, we determined that their clustering performance was approximately equal for the homogeneous balanced affinity models, and that EM and ES performed vastly better than K-means in the other settings. This corroborates the prior work done by [2] and [23]; since the canonical latent positions of a 2-dimensional full-rank homogeneous balanced affinity model lie equally spaced in the first quadrant with equal covariances, the K-means assumption of spherical covariances does not particularly hinder the clustering problem. As the other settings are far more varied in the

spacing of their latent positions and the shape and orientation of their covariances, the flexibility afforded by EM and ES renders them far more effective than K-means.

The major drawback of our algorithms in the SBM setting is their seeming inability to accurately estimate the component variances. We noted that ES occasionally yielded values of Λ^{-1} and $\tilde{\Lambda}^{-1}$ with diagonal entries far exceeding 1. A possible alteration to the algorithm that may reduce or eliminate this issue entirely would be to replace Λ , μ , and $\tilde{\Lambda}$ with

$$\begin{split} \hat{\mathbf{\Lambda}} &= \frac{\sum_{i=1}^{n} \hat{\mathbf{X}}_{i} \hat{\mathbf{X}}_{i}^{\top}}{n} \\ \hat{\boldsymbol{\mu}} &= \frac{\sum_{i=1}^{n} \hat{\mathbf{X}}_{i}}{n} \\ \hat{\mathbf{\Lambda}} &= \frac{\sum_{i=1}^{n} \frac{\hat{\mathbf{X}}_{i} \hat{\mathbf{X}}_{i}^{\top}}{\hat{\mathbf{X}}_{i}^{\top} \hat{\boldsymbol{\mu}}}}{n}, \end{split}$$

respectively, since doing so holds those terms constant and would prevent slight perturbations to the component means from causing excessive inflation to the entries of the component variances. Doing so also drastically reduces the computational complexity of each iteration to those mentioned in Remark 3.

Nonetheless, the strong evidence that spectral clustering based on ES can dominate that of EM opens up avenues for future research. During the writing of this article, [25] developed one-step estimators for the latent position matrices of both the ASE and LSE, the rows of which also asymptotically converge to curved normal mixtures. These estimators make additional use of the likelihood structure of the underlying random graph model, and as a result the ensuing C-GMM possesses component variances that are locally more efficient than those of the mixtures we have considered thus far. Moreover, they find that EM clustering based on these estimators can improve upon $EM \circ \{ASE, LSE\}$. Further work due to [24] modified the limiting distributions to include the sparsity factor ρ_n in the covariance functions. Moreover, all of our simulations were conducted under the assumption that the rank of the block probability matrix d and number of blocks K are both known; hence, we have left untreated the model selection problem of deciding upon these values from a collection of candidates. [26] developed an extended ASE, which also takes a limiting C-GMM, to address this problem exactly. In future articles we hope to compare performance of $ES \circ \{ASE, LSE\}$ with EM for these additional estimators, as well as implement ES algorithms based on the latter.

Appendix I: The expectation-solution algorithm for the general incomplete data setting

We model our treatment of the ES algorithm on that found in [16].

Z. M. Pisano et al.

Suppose $\mathbf{X} \sim f_{\Psi}$ is observed with realization $\mathbf{x}, \Psi \in \mathbb{R}^p$, and there exists a natural unobserved extension of the observed data $\mathbf{Y} = (\mathbf{X}, \mathbf{Z})$. Let

$$U_c(\mathbf{Y}, \mathbf{\Psi}) = \mathbf{0}$$

be the *p*-dimensional complete-data estimating equation that could be solved if **Y** were totally observable. To apply ES, we first re-express the completedata estimating equation in terms of a linear function of a *q*-dimensional vector function $S(\mathbf{Y})$, a $(p \times q)$ -dimensional matrix function C_{Ψ} , and a *p*-dimensional vector function $b_{\Psi}(\mathbf{X})$:

$$U_c(\mathbf{Y}, \mathbf{\Psi}) = U^{(1)}(\mathbf{X}, S(\mathbf{Y}), \mathbf{\Psi})$$

= $\sum_{j=1}^q c_j(\mathbf{\Psi}) S_j(\mathbf{Y}) + b_{\mathbf{\Psi}}(\mathbf{X})$
= $C_{\mathbf{\Psi}} S(\mathbf{Y}) + b_{\mathbf{\Psi}}(\mathbf{X})$
= 0.

where $c_j(\Psi)$ is the *j*th column of C_{Ψ} . Here $S(\mathbf{Y})$ is the complete data summary statistic.

The algorithm is:

Algorithm	I.4 The ES Algorithm for the General Incomplete Data Setting
1. Initialize	$\Psi^* = \Psi_0.$
0 E Sten	Compute $h(\mathbf{u} \mathbf{u}^*) = \mathbb{E} [C(\mathbf{v}) \mathbf{v} = \mathbf{u}]$

2. E-Step: Compute $h(\Psi|\Psi^*) := \mathbb{E}_{\Psi^*}[S(\mathbf{Y})|\mathbf{X} = \mathbf{x}].$

3. S-Step: Find $\hat{\Psi}$ that solves $\mathbf{A}_{\hat{\Psi}}h(\hat{\Psi}|\Psi^*) + b_{\hat{\Psi}}(\mathbf{x}) = \mathbf{0}$.

4. Take $\Psi^* = \hat{\Psi}$.

5. Repeat steps 2–4 until some convergence criterion is satisfied.

Appendix II: Comparison of EM and ES to K-means

In our simulations we also compared the clustering performance of EM and ES against that of K-means in models 1–4 and the brain connectome. We present the results in terms of 95% confidence intervals for the median difference in ARI, but due to the presence of multiple ties in the homogeneous balanced affinity settings we elected to use the sign test with a two-sided alternative. For $\ell = A, L$ as in section 4 and Z = M, S we define

$$\Delta_{\ell}^{KZ} := ARI_{KM \circ \ell SE} - ARI_{EZ \circ \ell SE}.$$

We embodden all entries indicating strict, significant improvement of our ES algorithms over K-means.

Appendix III: Additional rank-deficient simulations

We include the results of additional simulations in two randomly generated rank-deficient SBM settings analogous to those presented in Subsection 4.6.

TABLE 1				
Model	1	Mixture		

n	95% CI for	95% CI for	95% CI for	95% CI for		
	$\operatorname{med}(\Delta_A^{KM})$	$\operatorname{med}(\Delta_L^{KM})$	$\operatorname{med}(\Delta_A^{KS})$	$\operatorname{med}(\Delta_L^{KS})$		
200	(0.0167, 0.0493)	(0, 0.0220)	(0, 0.0048)	(0, 0.015)		
300	(0, 0.0122)	(0, 0.0119)	(0, 0)	(0, 0)		
400	(0, 0)	(0, 0)	(0, 0)	(0, 0)		
500	(0, 0)	(0, 0)	(0, 0)	(0, 0)		
600	(0, 0)	(0, 0)	(0, 0)	(0, 0)		
700	(0, 0)	(0, 0)	(0, 0)	(0, 0)		
800	(0, 0)	(0, 0)	(0, 0)	(0, 0)		
900	(0, 0)	(0, 0)	(0, 0)	(0, 0)		

TABLE 2Model 2 Mixture

n	95% CI for	95% CI for	95% CI for	95% CI for
	$med(\Delta_A^{KM})$	$\operatorname{med}(\Delta_L^{KM})$	$\operatorname{med}(\Delta_A^{KS})$	$\operatorname{med}(\Delta_L^{KS})$
200	(0.0407, 0.0915)	(0.0139, 0.0439)	(0.0127, 0.0761)	(0.0140, 0.0774)
300	(0.0098, 0.0313)	(0.0061, 0.0197)	(1e-7, -0.0190)	(0, 0.0105)
400	(0.0131, 0.0168)	(-2e-7, 0.0106)	(-1e-7, 0.0023)	(0, 0.0023)
500	(-0.0067, 0.0134)	(-0.0066, 0.0068)	(0, 2e-7)	(0, 0.0066)
600	(-2e-8, 0.0059)	(0, 0.0059)	(0, 5e-9)	(0, 0.0018)
700	(0, 0.0053)	(0, 0.0051)	(0, 0)	(0, 0)
800	(0, 0.0046)	(0, 0.0046)	(0, 4e-8)	(0, 0)
900	(-2e-9, 0.0041)	(-7e-10, 0.0013)	(0, 0)	(0, 0)

TABLE 3 Model 3 Mixture

$\begin{array}{c c c c c c c c c c c c c c c c c c c $	n	95% CI for	95% CI for	95% CI for	95% CI for
$\begin{array}{c ccccccccccccccccccccccccccccccccccc$		$\operatorname{med}(\Delta_A^{KM})$	$\operatorname{med}(\Delta_L^{KM})$	$\operatorname{med}(\Delta_A^{KS})$	$\operatorname{med}(\Delta_L^{KS})$
$\begin{array}{c} 300 \\ (-0.1756, -0.1317) \\ (-0.2068, -0.1662) \\ (-0.0998, -0.0318) \\ (-0.1273, -0.0888) \\ 400 \\ (-0.2019, -0.1712) \\ (-0.2215, -0.1966) \\ (-0.1682, -0.1204) \\ (-0.1762, -0.1242) \\ 500 \\ (-0.2306, -0.2002) \\ (-0.2476, -0.2263) \\ (-0.2288, -0.1986) \\ (-0.2370, -0.2103) \\ 600 \\ (-0.2514, -0.2305) \\ (-0.2671, -0.2437) \\ (-0.2547, -0.2296) \\ (-0.2656, -0.2434) \\ 700 \\ (-0.2573, -0.2380) \\ (-0.2772, -0.2555) \\ (-0.2657, -0.2386) \\ (-0.2755, -0.2577) \\ 800 \\ (-0.2534, -0.2388) \\ (-0.2716, -0.2562) \\ (-0.2619, -0.2370) \\ (-0.2760, -0.2574) \\ \end{array}$	200	(-0.1070, -0.0537)	(-0.1450, -0.0964)	(-0.0256, -8e-5)	(-0.026, -0.0080)
$\begin{array}{c} 400 \\ (-0.2019, -0.1712) \\ (-0.2215, -0.1966) \\ (-0.1682, -0.1204) \\ (-0.1762, -0.1242) \\ 500 \\ (-0.2306, -0.2002) \\ (-0.2476, -0.2263) \\ (-0.2288, -0.1986) \\ (-0.2370, -0.2103) \\ 600 \\ (-0.2514, -0.2305) \\ (-0.2671, -0.2437) \\ (-0.2547, -0.2296) \\ (-0.2656, -0.2434) \\ 700 \\ (-0.2573, -0.2380) \\ (-0.2772, -0.2555) \\ (-0.2657, -0.2386) \\ (-0.2755, -0.2577) \\ 800 \\ (-0.2534, -0.2388) \\ (-0.2716, -0.2562) \\ (-0.2619, -0.2370) \\ (-0.2760, -0.2574) \\ \end{array}$	300	(-0.1756, -0.1317)	(-0.2068, -0.1662)	(-0.0998, -0.0318)	(-0.1273, -0.0888)
$ \begin{array}{c} 500 \\ (-0.2306, -0.2002) \\ (-0.2476, -0.2263) \\ (-0.2288, -0.1986) \\ (-0.2370, -0.2103) \\ 600 \\ (-0.2514, -0.2305) \\ (-0.2671, -0.2437) \\ (-0.2547, -0.2296) \\ (-0.2656, -0.2434) \\ 700 \\ (-0.2573, -0.2380) \\ (-0.2772, -0.2555) \\ (-0.2657, -0.2386) \\ (-0.2755, -0.2577) \\ 800 \\ (-0.2534, -0.2388) \\ (-0.2716, -0.2562) \\ (-0.2619, -0.2370) \\ (-0.2760, -0.2574) \\ \end{array} $	400	(-0.2019, -0.1712)	(-0.2215, -0.1966)	(-0.1682, -0.1204)	(-0.1762, -0.1242)
$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	500	(-0.2306, -0.2002)	(-0.2476, -0.2263)	(-0.2288, -0.1986)	(-0.2370, -0.2103)
$\begin{array}{c} 700 \\ (-0.2573, -0.2380) \\ (-0.2772, -0.2555) \\ (-0.2657, -0.2386) \\ (-0.2755, -0.2577) \\ 800 \\ (-0.2534, -0.2388) \\ (-0.2716, -0.2562) \\ (-0.2619, -0.2370) \\ (-0.2760, -0.2574) \\ \end{array}$	600	(-0.2514, -0.2305)	(-0.2671, -0.2437)	(-0.2547, -0.2296)	(-0.2656, -0.2434)
(-0.2534, -0.2388) $(-0.2716, -0.2562)$ $(-0.2619, -0.2370)$ $(-0.2760, -0.2574)$	700	(-0.2573, -0.2380)	(-0.2772, -0.2555)	(-0.2657, -0.2386)	(-0.2755, -0.2577)
	800	(-0.2534, -0.2388)	(-0.2716, -0.2562)	(-0.2619, -0.2370)	(-0.2760, -0.2574)
900(-0.2495, -0.2370)(-0.2686, -0.2586)(-0.2528, -0.2426)(-0.2722, 0.2577)	900	(-0.2495, -0.2370)	(-0.2686, -0.2586)	(-0.2528, -0.2426)	(-0.2722, 0.2577)

TABLE 4Model 4 Mixture

n	95% CI for $mod(\Delta KM)$	95% CI for $mod(\Delta^{KM})$	95% CI for $mod(\Delta KS)$	95% CI for $mod(\Delta^{KS})$
200	(-0.0543, 0.0101)	$\frac{\text{Ined}(\Delta_L)}{(-0.0777 - 0.0244)}$	(-0.0066, 0.0428)	(-0.0099 - 0.0077)
300	(-0.1207, -0.0883)	(-0.1306, -0.0908)	(-0.0316, 0.1010)	(-0.0606, -0.0346)
400	(-0.1704, -0.1367)	(-0.1772, -0.1520)	(-0.0887, -0.0466)	(-0.1018, -0.0688)
500	(-0.2116, -0.1890)	(-0.2191, -0.1912)	(-0.1310, -0.943)	(-0.1502, -0.0961)
600	(-0.2270, -0.1970)	(-0.2412, -0.2150)	(-0.1992, -0.1556)	(-0.2252, -0.1858)
700	(-0.2436, -0.2149)	(-0.2607, -0.2342)	(-0.2345, -0.2051)	(-0.2505,-0.2301)
800	(-0.2476, -0.2252)	(-0.2678, -0.2462)	(-0.2420, -0.2214)	(-0.2625, -0.2399)
900	(-0.2512, -0.2372)	(-0.2777, -0.2572)	(-0.2540, -0.2363)	(-0.2785,-0.2611)

TABLE 5					
		Model 1	1 SBM		
n	95% CI for	95% CI for	95% CI for	95% CI for	
	$\operatorname{med}(\Delta_A^{KM})$	$\operatorname{med}(\Delta_L^{KM})$	$\operatorname{med}(\Delta_A^{KS})$	$\operatorname{med}(\Delta_L^{KS})$	
200	(0.0148, 0.0530)	(7e-6, 0.0257)	(1e-5, 0.0148)	(0, 0.0026)	
300	(0, 0.0227)	(0, 0.0113)	(0, 0)	(0, 0)	
400	(0, 0.0088)	(0, 0.0086)	(0, 0)	(0, 0)	
500	(0, 0)	(0, 3e-9)	(0, 0)	(0, 0)	
600	(0, 0)	(0, 0.0020)	(0, 0)	(0, 0)	
700	(0, 0)	(0, 0)	(0, 0)	(0, 0)	
800	(0, 0)	(0, 0)	(0, 0)	(0, 0)	
900	(0, 0)	(0, 0)	(0, 0)	(0, 0)	

TABLE 6 Model 2 SBM

n	95% CI for	95% CI for	$95\%~{\rm CI}$ for	95% CI for
	$med(\Delta_A^{KM})$	$\operatorname{med}(\Delta_L^{KM})$	$\operatorname{med}(\Delta_A^{KS})$	$\operatorname{med}(\Delta_L^{KS})$
200	(0.0033, 0.0135)	(0.0010, 0.0109)	(0.0094, 0.0265)	(0, 0.0015)
300	(0.0235, 0.0510)	(0.0073, 0.0216)	(0.0131, 0.0830)	(0, 0.0008)
400	(0.0137, 0.0430)	(5e-6, 0.0186)	(-8e-7, 0.0065)	(0, 5e-6)
500	(0.0122, 0.0325)	(0.0038, 0.0115)	(0, 0.0054)	(-0.0014, 9e-7)
600	(-6e-6, 0.0098)	(1e-7, 0.0055)	(-1e-6, 8e-8)	(-0.0051, 7e-8)
700	(0, 0.0063)	(0, 0.0047)	(-4e-8, 2e-7)	(-0.0014, 2e-7)
800	(0, 0.0087)	(-1e-7, 0.0044)	(0, 0)	(-2e-9, 0)
900	(0, 0.0079)	(0, 0.0040)	(0,0)	(0, 5e-9)

TABLE 7 Model 3 SBM

n	95% CI for	95% CI for	$95\%~{\rm CI}$ for	95% CI for
	$\operatorname{med}(\Delta_A^{KM})$	$\operatorname{med}(\Delta_L^{KM})$	$\operatorname{med}(\Delta_A^{KS})$	$\operatorname{med}(\Delta_L^{KS})$
200	(-0.1061, -0.0559)	(-0.1834, -0.1487)	(-0.1552, -0.0954)	(-0.1868, -0.1220)
300	(-0.2065, -0.1188)	(-0.2767, -0.2429)	(-0.2323, -0.1825)	(-0.2796, -0.2473)
400	(-0.3298, -0.2676)	(-0.3666, -0.3380)	(-0.3401, -0.2991)	(-0.3709, -0.3437)
500	(-0.3872, -0.3664)	(-0.4409, -0.4039)	(-0.3981, -0.3631)	(-0.4373, -0.4038)
600	(-0.4284, -0.4022)	(-0.4673, -0.4512)	(-0.4409, -0.4151)	(-0.4816, -0.4551)
700	(-0.4900, -0.4617)	(-0.5192, -0.4958)	(-0.4835, -0.4501)	(-0.5214, -0.4889)
800	(-0.5236, -0.4878)	(-0.5515, -0.5234)	(-0.5133, -0.4663)	(-0.5452, -0.5111)
900	(-0.5145, -0.4678)	(-0.5365, -0.5034)	(-0.4927, -0.4568)	(-0.5307, -0.4972)

TABLE 8Model 4 SBM

n	95% CI for	95% CI for	95% CI for	95% CI for
	$med(\Delta_A^{KM})$	$\operatorname{med}(\Delta_L^{KM})$	$\operatorname{med}(\Delta_A^{KS})$	$\operatorname{med}(\Delta_L^{KS})$
200	(-0.0781, -0.0312)	(-0.1298, -0.0814)	(-0.0050, -5e-5)	(-0.1049, -0.0718)
300	(-0.1356, -0.0636)	(-0.2158, -0.1760)	(-0.0233, -0.0090)	(-0.2061, -0.1693)
400	(-0.1975, -0.1219)	(-0.2891, -0.2402)	(-0.0916, -0.0621)	(-0.2826, -0.2527)
500	(-0.2809, -0.2366)	(-0.3505, -0.3271)	(-0.2265, -0.1744)	(-0.3466, -0.3235)
600	(-0.3477, -0.3134)	(-0.4117, -0.3841)	(-0.3084, -0.2622)	(-0.4078, -0.3756)
700	(-0.4164, -0.3799)	(-0.4611, -0.4339)	(-0.3875, -0.3591)	(-0.4630, -0.4344)
800	(-0.4417, -0.4144)	(-0.4962, -0.4755)	(-0.4413, -0.4067)	(-0.5070, -0.4767)
900	(-0.4970, -0.4754)	(-0.5504, -0.5230)	(-0.4965, -0.4661)	(-0.5469, -0.5286)

Spectral graph clustering via ES

TABLE 9 Brain Connectome SBM 95% CI for 95% CI for 95% CI for 95% CI for n $\underline{\mathrm{med}}(\Delta_{I}^{KM})$ $med(\Delta_L^{KS})$ $\operatorname{med}(\Delta_A^{KM})$ $\begin{array}{c|c} \operatorname{med}(\Delta_A^{KM}) & \operatorname{med}(\Delta_L^{KM}) & \operatorname{med}(\Delta_A^{KS}) & \operatorname{med}(\Delta_L^{KS}) \\ \hline (-0.5238, -0.5044) & (-0.5679, -0.5493) & (-0.5135, -0.4820) & (-0.5643, -0.5336) \end{array}$ 500 $600 \quad (-0.5514, -0.5200) \quad (-0.5909, -0.5672) \quad (-0.5322, -0.5055) \quad (-0.5811, -0.5472)$ 700 (-0.5676, -0.5432) (-0.6148, -0.5929) (-0.5453, -0.5199) (-0.5953, -0.5665) $\begin{array}{l} 800 \\ \hline (-0.5430, -0.5032) \\ \hline (-0.5907, -0.5580) \\ \hline (-0.5229, -0.4929) \\ \hline (-0.5722, -0.5447) \\ 900 \\ \hline (-0.5442, -0.4975) \\ \hline (-0.5824, -0.5468) \\ \hline (-0.5303, -0.4956) \\ \hline (-0.5744, -0.5439) \\ \hline \end{array}$ 1000(-0.5119, -0.4769)(-0.5498, -0.5247)(-0.5029, -0.4747)(-0.5461, -0.5212)1100(-0.4952, -0.4709)(-0.5363, -0.5170)(-0.4938, -0.4682)(-0.5362, -0.5093)1200(-0.4701, -0.4462)(-0.5155, -0.4904)(-0.4719, -0.4489)(-0.5123, -0.4914)

First consider the rank-2 g-block SBM with proportion vector

$$\boldsymbol{\pi} = (.10, .27, .12, .22, .10, .19)$$

and canonical latent position matrix

$\mathbf{x} =$	0.85	0.47
	0.72	0.66
	0.12	0.26
	0.32	0.12
	0.49	0.08
	0.21	0.64

which corresponds to the block probability matrix

$$\mathbf{B} = \begin{bmatrix} 0.9434 & 0.9222 & 0.2242 & 0.3284 & 0.4541 & 0.4793 \\ 0.9222 & 0.9540 & 0.2580 & 0.3096 & 0.4056 & 0.5736 \\ 0.2242 & 0.2580 & 0.0820 & 0.0696 & 0.0796 & 0.1916 \\ 0.3284 & 0.3096 & 0.0696 & 0.1168 & 0.1664 & 0.1440 \\ 0.4541 & 0.4056 & 0.0796 & 0.1664 & 0.2465 & 0.1541 \\ 0.4793 & 0.5736 & 0.1916 & 0.1440 & 0.1541 & 0.4537 \end{bmatrix}.$$

The ASE and LSE of one particular instance (n = 600) are presented in Figures 15–16. We present the results in Figure 17; $ES \circ ASE$ maintains its dominance over $EM \circ ASE$ here, but $ES \circ LSE$ appears to falter against $EM \circ LSE$ as n increases.

Now consider another rank-2 g-block SBM with proportion vector

$$\boldsymbol{\pi} = (0.09, 0.08, 0.23, 0.20, 0.28, 0.12)$$

and canonical latent position matrix

$$\mathbf{x} = \begin{bmatrix} 0.90 & 0.29 \\ 0.69 & 0.47 \\ 0.45 & 0.06 \\ 0.48 & 0.11 \\ 0.16 & 0.48 \\ 0.04 & 0.65 \end{bmatrix}$$



ASE of Rank-2 6-Block SBM (App. III)

FIG 15. Example two-dimensional ASE of a random graph generated from the first rankdeficient SBM in Appendix III. The smaller points have been colored according to block membership of the corresponding nodes in the original graph. The larger orange points encircled with a black outline denote the true latent positions, i.e., the rows of X.

which corresponds to the block probability matrix

	0.8941	0.7573	0.4224	0.4639	0.2832	0.2245
$\mathbf{B} =$	0.7573	0.6970	0.3387	0.3829	0.3360	0.3331
	0.4224	0.3387	0.2061	0.2226	0.1008	0.0570
	0.4639	0.3829	0.2226	0.2425	0.1296	0.0907
	0.2832	0.3360	0.1008	0.1296	0.2560	0.3184
	0.2245	0.3331	0.0570	0.0907	0.3184	0.4241

The ASE and LSE of one particular instance (n = 600) are presented in Figures 18–19. We present the results in Figure 20, which are similar to those obtained in the previous setting.



LSE of Rank-2 6-Block SBM (App. III)

FIG 16. Example two-dimensional LSE of a random graph generated from the first rankdeficient SBM in Appendix III. The smaller points have been colored according to block membership of the corresponding nodes in the original graph. The larger orange points encircled with a black outline denote the scaled latent positions, i.e., the centers of the limiting C-GMM.



95% CIs of Median Difference in ARI (Low-Rank SBM, App. III)

FIG 17. Simulation results for the first rank-deficient SBM in Appendix III.



ASE of Rank-2 6-Block SBM (App. III)

FIG 18. Example two-dimensional ASE of a random graph generated from the second rankdeficient SBM in Appendix III. The smaller points have been colored according to block membership of the corresponding nodes in the original graph. The larger orange points encircled with a black outline denote the true latent positions, i.e., the rows of X.



LSE of Rank-2 6-Block SBM (App. III)

FIG 19. Example two-dimensional LSE of a random graph generated from the second rankdeficient SBM in Appendix III. The smaller points have been colored according to block membership of the corresponding nodes in the original graph. The larger orange points encircled with a black outline denote the scaled latent positions, i.e., the centers of the limiting C-GMM.



95% CIs of Median Difference in ARI (Low-Rank SBM, App. III)

FIG 20. Simulation results for the second rank-deficient SBM in Appendix III.

Z. M. Pisano et al.

References

- AKAIKE, H. (1974). A new look at the statistical model identification. IEEE Transactions on Automatic Control 16 716–723. MR0423716
- [2] ATHREYA, A., PRIEBE, C. E., TANG, M., MARCHETTE, D. J. and SUSS-MAN, D. L. (2016). A limit theorem for scaled eigenvectors of random dot product graphs. *Sankhya A* 78 1–18. MR3494576
- [3] ATHREYA, A., FISHKIND, D. E., TANG, M., PARK, Y., VOGEL-STEIN, J. T., LEVIN, K., LYZINSKI, V., QIN, Y. and SUSSMAN, D. L. (2018). Statistical inference on random dot product graphs: A survey. *Jour*nal of Machine Learning Research 18 1–92. MR3827114
- [4] BICKEL, P. J. and DOKSUM, K. A. (2015). Mathematical Statistics: Basic Ideas and Selected Topics 1, 2nd ed. CRC Press. MR3287337
- [5] CAPE, J., TANG, M. and PRIEBE, C. E. (2019). On spectral embedding performance and elucidating network structure in stochastic block model graphs. *Network Science* 7 269–291.
- [6] DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* **39** 1–38. MR0501537
- [7] ELASHOFF, M. and RYAN, L. (2004). An EM algorithm for estimating equations. Journal of Computational and Graphical Statistics 13 48–65. MR2044870
- [8] FISHKIND, D. E., SUSSMAN, D. L., VOGELSTEIN, J. T. and PRIEBE, C. E. (2013). Consistent adjacency-spectral partitioning for the stochastic block model when the model paramters are unknown. *SIAM Journal on Matrix Analysis and Applications* **34** 23–39. MR3032990
- [9] FRALEY, C. and RAFTERY, A. E. (2002). Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical* Association 97 611–631. MR1951635
- [10] HOLLAND, P. W., ASKEY, K. B. and LEINHARDT, S. (1983). Stochastic blockmodels: First steps. *Social Networks* 5 109–137. MR0718088
- [11] HUBERT, L. and ARABIE, P. (1985). Comparing partitions. Journal of Classification 2 193–218.
- [12] KARRER, B. and NEWMAN, M. E. J. (2011). Stochastic blockmodels and community structure in networks. *Physical Review E* 83. MR2788206
- [13] LYZINSKI, V., SUSSMAN, D. L., TANG, M., ATHREYA, A. and PRIEBE, C. E. (2014). Perfect clustering for stochastic blockmodel graphs via adjacency spectral embedding. *Electronic Journal of Statistics* 8 2905– 2922. MR3299126
- [14] LYZINSKI, V., TANG, M., ATHREYA, A., PARK, Y. and PRIEBE, C. E. (2017). Community Detection and Classification in Hierarchical Stochastic Blockmodels. *IEEE Transactions on Network Science and Engineering* 4 13–26. MR3625952
- [15] MANN, H. B. and WHITNEY, D. R. (1947). On a test of whether one of two random variables is stochastically larger than the other. Annals of Mathematical Statistics 18 50–60. MR0022058

- [16] MCLACHLAN, G. J. and KRISHNAN, T. (2008). The EM Algorithm and Extensions, 2nd ed. Wiley-Interscience. MR2392878
- [17] PRIEBE, C. E., PARK, Y., TANG, M., ATHREYA, A., LYZINSKI, V. and VOGELSTEIN, J. T. (2017). Semiparametric Spectral Modeling of the Drosophila Connectome. arXiv.
- [18] PRIEBE, C. E., PARK, Y., VOGELSTEIN, J. T., CONROY, J. M., LYZIN-SKI, V., TANG, M., ATHREYA, A., CAPE, J. and BRIDGEFORD, E. (2019). On a two-truths phenomenon in spectral graph clustering. In *Proceedings* of the National Academy of Sciences of the United States of America. 116 13 5995–6000. MR3939781
- [19] RUBIN-DELANCHY, P., CAPE, J., TANG, M. and PRIEBE, C. E. (2020). A statistical interpretation of spectral embedding: the generalised random dot product graph. arXiv.
- [20] SCHWARZ, G. E. (1978). Estimating the dimension of a model. The Annals of Statistics 6 464–464. MR0468014
- [21] SCRUCCA, L., FOP, M., MURPHY, T. B. and RAFTERY, A. E. (2016). mclust 5: Clustering, classification, and density estimation using Gaussian finite mixture models. *The R Journal* 8 289–317.
- [22] SUSSMAN, D. L., TANG, M., FISHKIND, D. E. and PRIEBE, C. E. (2012). A consistent adjacency spectral embedding for stochastic blockmodel graphs. *Journal of the American Statistical Association* **107** 1119– 1128. MR3010899
- [23] TANG, M. and PRIEBE, C. E. (2018). Limit theorems for eigenvectors of the normalized Laplacian for random graphs. *The Annals of Statistics* 46 2360–2415. MR3845021
- [24] XIE, F. (2021). Entrywise limit theorems of eigenvectors and their one-step refinement for sparse random graphs. *arXiv*.
- [25] XIE, F. and XU, Y. (2019). Efficient estimation for Random Dot Product Graphs via a one-step procedure. arXiv. MR4186493
- [26] YANG, C., PRIEBE, C. E., PARK, Y. and MARCHETTE, D. J. (2021). Simultaneous dimensionality and complexity model selection for spectral graph clustering. *Journal of Computational and Graphical Statistics* **30** 422–441. MR4270514
- [27] YOUNG, S. J. and SCHEINERMAN, E. (2007). Random dot product graph models for social networks. In *Proceedings of the 5th international conference on algorithms and models for the web-graph* 138–149. MR2504912