



Latent structure blockmodels for Bayesian spectral graph clustering

Francesco Sanna Passino¹ · Nicholas A. Heard¹

Received: 7 July 2021 / Accepted: 20 January 2022
© The Author(s) 2022

Abstract

Spectral embedding of network adjacency matrices often produces node representations living approximately around low-dimensional submanifold structures. In particular, hidden substructure is expected to arise when the graph is generated from a latent position model. Furthermore, the presence of communities within the network might generate community-specific submanifold structures in the embedding, but this is not explicitly accounted for in most statistical models for networks. In this article, a class of models called latent structure block models (LSBM) is proposed to address such scenarios, allowing for graph clustering when community-specific one-dimensional manifold structure is present. LSBMs focus on a specific class of latent space model, the random dot product graph (RDPG), and assign a latent submanifold to the latent positions of each community. A Bayesian model for the embeddings arising from LSBMs is discussed, and shown to have a good performance on simulated and real-world network data. The model is able to correctly recover the underlying communities living in a one-dimensional manifold, even when the parametric form of the underlying curves is unknown, achieving remarkable results on a variety of real data.

Keywords Community detection · Latent structure model · Random dot product graph · Spectral clustering.

1 Introduction

Network-valued data are commonly observed in many real-world applications. They are typically represented by graph adjacency matrices, consisting of binary indicators summarising which nodes are connected. Spectral embedding (Luo et al. 2003) is often the first preprocessing step in the analysis of graph adjacency matrices: the nodes are embedded onto a low-dimensional space via eigendecompositions or singular value decompositions. This work discusses a Bayesian network model for clustering the nodes of the graph in the embedding space, when community-specific substructure is present. Therefore, the proposed methodology could be classified among spectral graph clustering methodologies,

which have been extensively studied in the literature. Recent examples include Priebe et al. (2019); Pensky and Zhang (2019); Yang et al. (2021); Sanna Passino et al. (2021). More generally, spectral clustering is an active research area, extensively studied both from a theoretical and applied perspective (for some examples, see Couillet and Benaych-Georges 2016; Hofmeyr 2019; Zhu et al. 2019; Amini and Razaei 2021; Han et al. 2021). In general, research efforts in spectral graph clustering (Rohe et al. 2011; Sussman et al. 2012; Lyzinski et al. 2014) have been primarily focused on studying its properties under simple community structures, for example the stochastic block model (SBM, Holland et al. 1983). On the other hand, in practical applications, it is common to observe more complex community-specific submanifold structure (Priebe et al. 2017; Sanna Passino et al. 2021), which is not captured by spectral graph clustering models at present. In this work, a Bayesian model for spectral graph clustering under such a scenario is proposed, and applied on a variety of real-world networks.

A popular statistical model for graph adjacency matrices is the latent position model (LPM, Hoff et al. 2002). Each node is assumed to have a low-dimensional vector representation $\mathbf{x}_i \in \mathbb{R}^d$, such that the probability of connection between two nodes is $\kappa(\mathbf{x}_i, \mathbf{x}_j)$, for a given *kernel*

FSP and NAH gratefully acknowledge funding from the Microsoft Security AI research grant “*Understanding the enterprise: Host-based event prediction for automatic defence in cyber-security*”.

✉ Francesco Sanna Passino
f.sannapassino@imperial.ac.uk

Nicholas A. Heard
n.heard@imperial.ac.uk

¹ Department of Mathematics, Imperial College London, London, UK




Fig. 1 Latent positions (black) and spectral estimates (gray) from a simulated Hardy–Weinberg LSM

function $\kappa : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, 1]$. Examples of commonly used kernel functions are the inner product $\kappa(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{x}_i^\top \mathbf{x}_j$ (Young and Scheinerman 2007; Athreya et al. 2018), the logistic distance link $\kappa(\mathbf{x}_i, \mathbf{x}_j) = [1 + \exp(-\|\mathbf{x}_i - \mathbf{x}_j\|)]^{-1}$ (Hoff et al. 2002; Salter-Townshend and McCormick 2017), where $\|\cdot\|$ is a norm, or the Bernoulli–Poisson link $\kappa(\mathbf{x}_i, \mathbf{x}_j) = 1 - \exp(-\mathbf{x}_i^\top \mathbf{x}_j)$ (for example, Todeschini et al. 2020), where the latent positions lie in $\mathbb{R}_+^d$. Rubin-Delanchy (2020) notes that spectral embeddings of adjacency matrices generated from LPMs produce node representations living near low-dimensional submanifold structures. Therefore, for subsequent inferential tasks, for example clustering, it is necessary to employ methods which take into account this latent manifold structure.

If the kernel function of the LPM is the inner product, spectral embedding provides consistent estimates of the latent positions (Athreya et al. 2016; Rubin-Delanchy et al. 2017), up to orthogonal rotations. This model is usually referred to as the random dot product graph (RDPG, Young and Scheinerman 2007; Athreya et al. 2018). In this work, RDPGs where the latent positions lie on a one-dimensional submanifold of the latent space are considered. Such a model is known in the literature as a latent structure model (LSM, Athreya et al. 2021). An example is displayed in Fig. 1, which shows the latent positions, and corresponding estimates obtained using adjacency spectral embedding (ASE; see Sect. 2 for details), of a simulated graph with 1000 nodes, where the latent positions are assumed to be drawn from the Hardy–Weinberg curve $\mathbf{f}(\theta) = \{\theta^2, (1 - \theta)^2, 2\theta(1 - \theta)\}$, with $\theta \sim \text{Uniform}[0, 1]$ (Athreya et al. 2021; Trosset et al. 2020).

In addition, nodes are commonly divided into inferred groups such that it can be plausible to assume that the latent positions live on group-specific submanifolds. The objective




Fig. 2 Fitted quadratic curves for the two-dimensional (directed) spectral embedding of a computer network graph with two communities

of community detection on graphs is to recover such latent groups. For example, in computer networks, it is common to observe community-specific curves in the latent space. Figure 2 displays an example, representing the two-dimensional latent positions, estimated using directed adjacency spectral embedding (DASE; see Sect. 2), of 278 computers in a network of HTTP/HTTPS connections from machines hosted in two computer laboratories at Imperial College London (further details are given in Sect. 5.5). The latent positions corresponding to the two communities appear to be distributed around two quadratic curves with different parameters. The best-fitting quadratic curves in $\mathbb{R}^2$ passing through the origin are also displayed, representing an estimate of the underlying community-specific submanifolds.

Motivated by the practical application to computer networks, this article develops inferential methods for LSMs allowing for latent community structure. Nodes belonging to different communities are assumed to correspond to community-specific submanifolds. This leads to the definition of *latent structure blockmodels (LSBM)*, which admit communities with manifold structure in the latent space. The main contribution is a Bayesian probability model for LSBM embeddings, which enables community detection in RDPG models with underlying community-specific manifold structure.

Naturally, for clustering purposes, procedures that exploit the structure of the submanifold are expected to perform more effectively than standard alternatives such as spectral clustering with k -means (von Luxburg 2007; Rohe et al. 2011) or Gaussian mixture models (Rubin-Delanchy et al. 2017). The problem of clustering in the presence of group-specific curves has been addressed for distance-based clustering methods (Diday 1971; Diday and Simon 1976), often using kernel-based methodologies (see Bouveyron et al. (2015), and the

review of Scholkopf and Smola (2018) or, more recently, deep learning techniques (Ye et al. 2020), but it is understudied for the purposes of graph clustering. Exploiting the underlying structure is often impractical, since the submanifold is usually unknown. Within the context of LSMs, Trosset et al. (2020) propose using an extension of Isomap (Tenenbaum et al. 2000) in the presence of noise to learn a one-dimensional unknown manifold from the estimated latent positions. In this work, the problem of submanifold estimation is addressed using flexible functional models, namely Gaussian processes, in the embedding space. Applications on simulated and real-world computer network data show that the proposed methodology is able to successfully recover the underlying communities even in complex cases with substantial overlap between the groups. Recent work from Dunson and Wu (2021) also demonstrate the potential of Gaussian processes for manifold inference from noisy observations.

Geometries arising from network models have been extensively studied in the literature (see, for example, Asta and Shalizi 2015; McCormick and Zheng 2015; Smith et al. 2019). The approach in this paper is also related to the findings of Priebe et al. (2017), where a semiparametric Gaussian mixture model with quadratic structure is used to model the embeddings arising from a subset of the neurons in the *Drosophila* connectome. In particular, the *Kenyon Cell* neurons are not well represented by a stochastic block-model, which is otherwise appropriate for other neuron types. Consequently, those neurons are modelled via a continuous curve in the latent space, estimated via semiparametric maximum likelihood (Kiefer and Wolfowitz 1956; Lindsay 1983). However, the real-world examples presented in this article in Sect. 5 require community-specific curves in the latent space for all the communities, not only for a subset of the nodes. Furthermore, estimation can be conducted within the Bayesian paradigm, which conveniently allows marginalisation of the model parameters.

This article is organised as follows: Sect. 2 formally introduces RDPGs and spectral embedding methods. Sections 3 and 4 present the main contributions of this work: the latent structure blockmodel (LSBM), and Bayesian inference for the communities of LSBM embeddings. The method is applied to simulated and real-world networks in Sect. 5, and the article is then concluded with a discussion in Sect. 6.

2 Random dot product graphs and latent structure models

Mathematically, a network $\mathbb{G} = (V, E)$ is represented by a set $V = \{1, 2, \dots, n\}$ of nodes, and an edge set E , such that $(i, j) \in E$ if i forms a link with j . If $(i, j) \in E \Rightarrow (j, i) \in E$, the network is *undirected*, otherwise it is *directed*. If the node set is partitioned into two sets V_1 and V_2 , $V_1 \cap V_2 = \emptyset$,

such that $E \subset V_1 \times V_2$, the graph is *bipartite*. A network can be summarised by the adjacency matrix $A \in \{0, 1\}^{n \times n}$, such that $A_{ij} = \mathbb{1}_E\{(i, j)\}$, where $\mathbb{1}_{\{\cdot\}}$ is an indicator function. Latent position models (LPMs, Hoff et al. 2002,) for undirected graphs postulate that the edges are sampled independently, with

$$\mathbb{P}(A_{ij} = 1 \mid \mathbf{x}_i, \mathbf{x}_j) = \kappa(\mathbf{x}_i, \mathbf{x}_j).$$

A special case of LPMs is the random dot product graph (RDPG, Young and Scheinerman 2007), defined below.

Definition 1 [Random dot product graph] For an integer $d > 0$, let F be an inner product distribution on $\mathcal{X} \subset \mathbb{R}^d$, such that for all $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, $0 \leq \mathbf{x}^\top \mathbf{x}' \leq 1$. Also, let $X = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top \in \mathcal{X}^n$ and $A \in \{0, 1\}^{n \times n}$ be symmetric. Then $(A, X) \sim \text{RDPG}_d(F^n)$ if

$$\begin{aligned} \mathbf{x}_1, \dots, \mathbf{x}_n &\stackrel{iid}{\sim} F \\ A_{ij} \mid X &\sim \text{Bernoulli}(\mathbf{x}_i^\top \mathbf{x}_j), \quad 1 \leq i < j \leq n. \end{aligned} \quad (1)$$

Given a realisation of the adjacency matrix A , the first inferential objective is to estimate the latent positions $\mathbf{x}_1, \dots, \mathbf{x}_n$. In RDPGs, the latent positions are inherently unidentifiable; in particular, multiplying the latent positions by any orthogonal matrix $Q \in \mathbb{O}(d)$, the orthogonal group with signature d , leaves the inner product (1) unchanged: $\mathbf{x}^\top \mathbf{x}' = (Q\mathbf{x})^\top Q\mathbf{x}'$. Therefore, the latent positions can only be estimated *up to orthogonal rotations*. Under such a restriction, the latent positions are consistently estimated via spectral embedding methods. In particular, the adjacency spectral embedding (ASE), defined below, has convenient asymptotic properties.

Definition 2 [ASE—Adjacency spectral embedding] For a given integer $d \in \{1, \dots, n\}$ and a binary symmetric adjacency matrix $A \in \{0, 1\}^{n \times n}$, the d -dimensional adjacency spectral embedding (ASE) $\hat{X} = [\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n]^\top$ of A is

$$\hat{X} = \Gamma \Lambda^{1/2} \in \mathbb{R}^{n \times d},$$

Where Λ is a $d \times d$ diagonal matrix containing the absolute values of the d largest eigenvalues in magnitude, in decreasing order, and Γ is a $n \times d$ matrix containing corresponding orthonormal eigenvectors.

Alternatively, the Laplacian spectral embedding is also frequently used, and considers a spectral decomposition of the modified Laplacian $D^{-1/2} A D^{-1/2}$ instead, where D is the degree matrix. In this work, the focus will be mainly on ASE. Using ASE, the latent positions are consistently estimated up to orthogonal rotations, and a central limit theorem is available (Athreya et al. 2016; Rubin-Delanchy et al. 2017; Athreya et al. 2018).

Theorem 1 [ASE-CLT—ASE central limit theorem] Let $(A^{(n)}, X^{(n)}) \sim \text{RDPG}_d(F^n)$, $n = 1, 2, \dots$, be a sequence of adjacency matrices and corresponding latent positions, and let $\hat{X}^{(n)}$ be the d -dimensional ASE of $A^{(n)}$. For an integer $m > 0$, and for the sequences of points $\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathcal{X}$ and $\mathbf{u}_1, \dots, \mathbf{u}_m \in \mathbb{R}^d$, there exists a sequence of orthogonal matrices $\mathbf{Q}_1, \mathbf{Q}_2, \dots \in \mathbb{O}(d)$ such that:

$$\lim_{n \rightarrow \infty} \mathbb{P} \left\{ \bigcap_{i=1}^m \sqrt{n} \left(\mathbf{Q}_n \hat{\mathbf{x}}_i^{(n)} - \mathbf{x}_i^{(n)} \right) \leq \mathbf{u}_i \mid \mathbf{x}_i^{(n)} = \mathbf{x}_i \right\} = \prod_{i=1}^m \Phi\{\mathbf{u}_i, \Sigma(\mathbf{x}_i)\},$$

Where $\Phi\{\cdot\}$ is the CDF of a d -dimensional normal distribution, and $\Sigma(\cdot)$ is a covariance matrix which depends on the true value of the latent position.

Therefore, informally, it could be assumed that, for n large, $\mathbf{Q} \hat{\mathbf{x}}_i \sim \mathbb{N}_d\{\mathbf{x}_i, \Sigma(\mathbf{x}_i)\}$ for some $\mathbf{Q} \in \mathbb{O}(d)$, where $\mathbb{N}_d\{\cdot\}$ is the d -dimensional multivariate normal distribution. Theorem 1 provides strong theoretical justification for the choice of normal likelihood assumed in the Bayesian model presented in the next section.

This article is mainly concerned with community detection in RDPGs when the latent positions lie on community-specific one-dimensional submanifolds $\mathcal{S}_k \subset \mathbb{R}^d$, $k = 1, \dots, K$. The proposed methodology builds upon latent structure models (LSMs, Athreya et al. 2021), a subset of RDPGs. In LSMs, it is assumed that the latent positions of the nodes are determined by draws from a univariate underlying distribution G on $[0, 1]$, inducing a distribution F on a structural support univariate submanifold $\mathcal{S} \subset \mathbb{R}^d$, such that:

$$\mathbf{x}_i \stackrel{iid}{\sim} F, \quad i = 1, \dots, n. \quad (2)$$

In particular, the distribution F on $\mathcal{S}$ is the distribution of the vector-valued transformation $\mathbf{f}(\theta)$ of a univariate random variable $\theta \sim G$, where $\mathbf{f}: [0, 1] \rightarrow \mathcal{S}$ is a function mapping draws from θ to $\mathcal{S}$. The inverse function $\mathbf{f}^{-1}: \mathcal{S} \rightarrow [0, 1]$ could be interpreted as the arc-length parametrisation of $\mathcal{S}$. In simple terms, each node is assigned a draw θ_i from the underlying distribution G , representing how far along the submanifold $\mathcal{S}$ the corresponding latent position lies, such that:

$$\mathbf{x}_i = \mathbf{f}(\theta_i).$$

Therefore, conditional on $\mathbf{f}$, the latent positions are uniquely determined by θ_i .

If the graph is directed or bipartite, each node is assigned two latent positions $\mathbf{x}_i$ and $\mathbf{x}'_i$, and the random dot product graph model takes the form $\mathbb{P}(A_{ij} = 1 \mid \mathbf{x}_i, \mathbf{x}'_j) = \mathbf{x}_i^\top \mathbf{x}'_j$. In

this case, the singular value decomposition of A can be used over the eigendecomposition.

Definition 3 [DASE—Directed adjacency spectral embedding] For an integer $d \in \{1, \dots, n\}$ and an adjacency matrix $A \in \{0, 1\}^{n \times n}$, the d -dimensional directed adjacency spectral embeddings (DASE) $\hat{X} = [\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n]^\top$ and $\hat{X}' = [\hat{\mathbf{x}}'_1, \dots, \hat{\mathbf{x}}'_n]^\top$ of A are

$$\hat{X} = US^{1/2} \in \mathbb{R}^{n \times d}, \quad \hat{X}' = VS^{1/2} \in \mathbb{R}^{n \times d},$$

Where S is a $d \times d$ diagonal matrix containing the d largest singular values, in decreasing order, and U and V are a $n \times d$ matrix containing corresponding left- and right-singular vectors.

Note that ASE and DASE could also be interpreted as instances of multidimensional scaling (MDS, Torgerson 1952; Shepard 1962). Clearly, DASE could be also used on rectangular adjacency matrices arising from bipartite graphs. Furthermore, Jones and Rubin-Delanchy (2020) proves the equivalent of Theorem 1 for DASE, demonstrating that DASE also provides a consistent procedure for estimation of the latent positions in directed or bipartite graphs. Analogous constructions to undirected LSMs could be posited for directed or bipartite models, where the $\mathbf{x}_i$'s and $\mathbf{x}'_i$'s lie on univariate structural support submanifolds.

3 Latent structure blockmodels

This section extends LSM, explicitly allowing for community structure. In particular, it is assumed that each node is assigned a latent community membership $z_i \in \{1, \dots, K\}$, $i = 1, \dots, n$, and each community is associated with a different one-dimensional structural support submanifold $\mathcal{S}_k \subset \mathbb{R}^d$, $k = 1, \dots, K$. Therefore, it is assumed that $F = \sum_{k=1}^K \eta_k F_k$ is a mixture distribution with components $F_1, \dots, F_K$ supported on the submanifolds $\mathcal{S}_1, \dots, \mathcal{S}_K$, with weights $\eta_1, \dots, \eta_K$ such that for each k , $\eta_k \geq 0$ and $\sum_{k=1}^K \eta_k = 1$. Assuming community allocations $\mathbf{z} = (z_1, \dots, z_n)$, the latent positions are obtained as

$$\mathbf{x}_i \mid z_i \sim F_{z_i}, \quad i = 1, \dots, n,$$

Where, similarly to (2), F_{z_i} is the distribution of the community-specific vector-valued transformation $\mathbf{f}_{z_i}(\theta)$, of a univariate random variable $\theta \sim G$, which is instead shared across communities. The vector-valued functions $\mathbf{f}_k = (f_{k,1}, \dots, f_{k,d}): \mathcal{G} \rightarrow \mathcal{S}_k$, $k = 1, \dots, K$, map the latent draw from the distribution G to $\mathcal{S}_k$. The resulting model will be referred to as the *latent structure blockmodel* (LSBM). Note that, for generality, the support of the underlying distribution G is assumed here to be $\mathcal{G} \subset \mathbb{R}$. Furthermore, G is

common to *all* the nodes, and the pair (θ_i, z_i) , where $\theta_i \sim G$, uniquely determines the latent position $\mathbf{x}_i$ through $\mathbf{f}_{z_i}$, such that

$$\mathbf{x}_i = \mathbf{f}_{z_i}(\theta_i).$$

Note that, in the framework described above, the submanifolds $S_1, \dots, S_K$ are one-dimensional, corresponding to curves, but the LSBM could be extended to higher-dimensional settings of the underlying subspaces, postulating draws from a multivariate distribution G supported on $\mathcal{G} \subseteq \mathbb{R}^p$ with $1 \leq p < d$.

Since the latent structure blockmodel is a special case of the random dot product graph, the LSBM latent positions are also estimated consistently via ASE, up to orthogonal rotations, conditional on knowing the functions $\mathbf{f}_k(\cdot)$. Therefore, by Theorem 1, approximately:

$$\mathbf{Q}\hat{\mathbf{x}}_i \sim \mathbb{N}\{\mathbf{f}_{z_i}(\theta_i), \mathbf{\Sigma}(\mathbf{x}_i)\}. \quad (3)$$

Special cases of the LSBM include the stochastic blockmodel (SBM, Holland et al. 1983), the degree-corrected SBM (DCSBM, Karrer and Newman 2011), and other more complex latent structure models with clustering structure, as demonstrated in the following examples. Note that, despite the similarity in name, LSBMs are *different* from latent block models (LBMs; see, for example, Keribin et al. 2015; Wyse et al. 2017), which instead generalise the SBM to bipartite graphs.

Example 1 [Stochastic blockmodel, SBM] In an SBM (Holland et al. 1983), the edge probability is determined by the community allocation of the nodes: $A_{ij} \sim \text{Bernoulli}(B_{z_i z_j})$, where $\mathbf{B} \in [0, 1]^{K \times K}$ is a matrix of probabilities for connections between communities. An SBM characterised by a nonnegative definite matrix $\mathbf{B}$ of rank d can be expressed as an LSBM, assigning community-specific latent positions $\mathbf{x}_i = \mathbf{e}_{z_i} \in \mathbb{R}^d$, such that for each (k, ℓ) , $B_{k\ell} = \mathbf{e}_k^\top \mathbf{e}_\ell$. Therefore, $\mathbf{f}_k(\theta_i) = \mathbf{e}_k$, with each $\theta_i = 1$ for identifiability. It follows that $f_{k,j}(\theta_i) = v_{k,j}$.

Example 2 [Degree-corrected SBM, DCSBM] DCSBMs (Karrer and Newman 2011) extend SBMs, allowing for heterogeneous degree distributions within communities. The edge probability depends on the community allocation of the nodes, and degree-correction parameters $\theta \in \mathbb{R}^n$ for each node, such that $A_{ij} \sim \text{Bernoulli}(\theta_i \theta_j B_{z_i z_j})$. In a latent structure blockmodel interpretation, the latent positions are $\mathbf{x}_i = \theta_i \mathbf{e}_{z_i} \in \mathbb{R}^d$. Therefore, $\mathbf{f}_k(\theta_i) = \theta_i \mathbf{e}_k$, with $f_{k,j}(\theta_i) = \theta_i v_{k,j}$. For identifiability, one could set each $v_{k,1} = 1$.

Example 3 [Quadratic latent structure blockmodel] For an LSBM with quadratic $\mathbf{f}_k(\cdot)$, it can be postulated that, conditional on a community allocation z_i , $\mathbf{x}_i = \mathbf{f}_{z_i}(\theta_i) =$

$\alpha_{z_i} \theta_i^2 + \beta_{z_i} \theta_i + \gamma_{z_i}$, with $\alpha_k, \beta_k, \gamma_k \in \mathbb{R}^d$. Under this model: $f_{k,j}(\theta_i) = \alpha_{k,j} \theta_i^2 + \beta_{k,j} \theta_i + \gamma_{k,j}$. Note that the model is not identifiable: for $v \in \mathbb{R}$, then $(\alpha_{z_i}/v^2)(v\theta_i)^2 + (\beta_{z_i}/v)(v\theta_i) + \gamma_{z_i}$ is equivalent to $\mathbf{f}_{z_i}(\theta_i)$. A possible solution is to fix each $\beta_{k,1} = 1$.

Figure 3 shows example embeddings arising from the three models described above. From these plots, it is evident that taking into account the underlying structure is essential for successful community detection.

4 Bayesian modelling of LSBM embeddings

Under the LSBM, the inferential objective is to recover the community allocations $\mathbf{z} = (z_1, \dots, z_n)$ given a realisation of the adjacency matrix $\mathbf{A}$. Assuming normality of the rows of ASE for LSBMs (3), the inferential problem consists of making joint inference about $\mathbf{z}$ and the latent functions $\mathbf{f}_k = (f_{k,1}, \dots, f_{k,d}) : \mathcal{G} \rightarrow \mathbb{R}^d$. The prior for $\mathbf{z}$ follows a Categorical-Dirichlet structure:

$$\begin{aligned} z_i &\sim \text{Categorical}(\boldsymbol{\eta}), \quad \boldsymbol{\eta} = (\eta_1, \dots, \eta_K), \quad i = 1, \dots, n, \\ \boldsymbol{\eta} &\sim \text{Dirichlet}(v/K, \dots, v/K), \end{aligned} \quad (4)$$

Where $v, \eta_k \in \mathbb{R}_+$, $k \in \{1, \dots, K\}$, and $\sum_{k=1}^K \eta_k = 1$.

Following the ASE-CLT in Theorem 1, the estimated latent positions are assumed to be drawn from Gaussian distributions centred at the underlying function value. Conditional on the pair (θ_i, z_i) , the following distribution is postulated for $\hat{\mathbf{x}}_i$:

$$\hat{\mathbf{x}}_i \mid \theta_i, \mathbf{f}_{z_i}, \sigma_{z_i}^2 \sim \mathbb{N}_d \left\{ \mathbf{f}_{z_i}(\theta_i), \sigma_{z_i}^2 \mathbf{I}_d \right\}, \quad i = 1, \dots, n, \quad (5)$$

Where $\sigma_k^2 = (\sigma_{k,1}^2, \dots, \sigma_{k,d}^2) \in \mathbb{R}_+^d$ is a community-specific vector of variances and $\mathbf{I}_d$ is the $d \times d$ identity matrix. Note that, for simplicity, the components of the estimated latent positions are assumed to be independent. This assumption loosely corresponds to the k -means clustering approach, which has been successfully deployed in spectral graph clustering under the SBM (Rohe et al. 2011). Here, the same idea is extended to a functional setting. Furthermore, for tractability (5) assumes the variance of $\hat{\mathbf{x}}_i$ does not depend on $\mathbf{x}_i$, but only on the community allocation z_i .

For a full Bayesian model specification, prior distributions are required for the latent functions and the variances. The most popular prior for functions is the Gaussian process (GP; see, for example, Rasmussen and Williams 2006). Here, for each community k , the j -th dimension of the true latent positions are assumed to lie on a one-dimensional manifold described by a function $f_{k,j}$ with a hierarchical GP-IG prior,




Fig. 3 Scatterplots of the two-dimensional ASE of simulated graphs arising from the models in Examples 1, 2, 3, and true underlying latent curves (in black). For each graph, $n = 1000$ with $K = 2$ commu-

nities of equal size. For **a** and **b**, $\mathbf{v}_1 = [3/4, 1/4]$, $\mathbf{v}_2 = [1/4, 3/4]$ and $\theta_i \sim \text{Beta}(1, 1)$, cf. Examples 1 and 2. For **c**, $\alpha_k = [-1, -4]$, $\beta_k = [1, 1]$, $\gamma_k = [0, 0]$ and $\theta_i \sim \text{Beta}(2, 1)$, cf. Example 3

with an inverse gamma (IG) prior on the variance:

$$f_{k,j} | \sigma_{k,j}^2 \sim \text{GP}(0, \sigma_{k,j}^2 \xi_{k,j}), \quad k = 1, \dots, K, \quad j = 1, \dots, d, \\ \sigma_{k,j}^2 \sim \text{IG}(a_0, b_0), \quad k = 1, \dots, K, \quad j = 1, \dots, d, \quad (6)$$

Where $\xi_{k,j}(\cdot, \cdot)$ is a positive semi-definite kernel function and $a_0, b_0 \in \mathbb{R}_+$. Note that the terminology “kernel” is used in the literature for both the GP covariance function $\xi_{k,j}(\cdot, \cdot)$ and the function $\kappa(\cdot, \cdot)$ used in LPMs (cf. Sect. 1), but their meaning is *fundamentally different*. In particular, $\kappa: \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, 1]$ is a component of the *graph generating process*, and assumed here to be the inner product, corresponding to RDPGs. On the other hand, $\xi_{k,j}: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ is the scaled covariance function of the GP prior on the unknown function $f_{k,j}$, which is used for modelling the observed graph embeddings, or equivalently the *embedding generating process*. There are no restrictions on the possible forms of $\xi_{k,j}$, except positive semi-definiteness. Overall, the approach is similar to the overlapping mixture of Gaussian processes method (Lázaro-Gredilla et al. 2012).

The class of models that can be expressed in form (6) is vast, and includes, for example, polynomial regression and splines, under a conjugate normal-inverse-gamma prior for the regression coefficients. For example, consider any function that can be expressed in the form $f_{z_i,j}(\theta_i) = \phi_{z_i,j}(\theta_i)^\top \mathbf{w}_{z_i,j}$ for some community-specific basis functions $\phi_{k,j}: \mathbb{R} \rightarrow \mathbb{R}^{q_{k,j}}$, $q_{k,j} \in \mathbb{Z}_+$, and corresponding coefficients $\mathbf{w}_{k,j} \in \mathbb{R}^{q_{k,j}}$. If the coefficients are given a normal-inverse-gamma prior

$$(\mathbf{w}_{k,j}, \sigma_{k,j}^2) \sim \text{NIG}(\mathbf{0}, \Delta_{k,j}, a_0, b_0) \\ = \mathbb{N}_{q_{k,j}}(\mathbf{0}, \sigma_{k,j}^2 \Delta_{k,j}) \text{IG}(a_0, b_0),$$

Where $\Delta_{k,j} \in \mathbb{R}^{q_{k,j} \times q_{k,j}}$ is a positive definite matrix, then $f_{k,j}$ takes form (6), with the kernel function

$$\xi_{k,j}(\theta, \theta') = \phi_{k,j}(\theta)^\top \Delta_{k,j} \phi_{k,j}(\theta'). \quad (7)$$

Considering the examples in Sect. 3, the SBM (cf. Example 1) corresponds to $\xi_{k,j}(\theta, \theta') = \Delta_{k,j}$, $\Delta_{k,j} \in \mathbb{R}_+$, whereas the DCSBM (cf. Example 2) corresponds to $\xi_{k,j}(\theta, \theta') = \theta \theta' \Delta_{k,j}$, $\Delta_{k,j} \in \mathbb{R}_+$. For the quadratic LSBM (cf. Example 3), the GP kernel takes the form $\xi_{k,j}(\theta, \theta') = (1, \theta, \theta^2) \Delta_{k,j} (1, \theta', \theta'^2)^\top$ for a positive definite scaling matrix $\Delta_{k,j} \in \mathbb{R}^{3 \times 3}$.

The LSBM specification is completed with a prior for each θ_i value, which specifies the unobserved location of the latent position $\mathbf{x}_i$ along each submanifold curve; for $\mu_\theta \in \mathbb{R}$, $\sigma_\theta^2 \in \mathbb{R}_+$,

$$\theta_i \sim \mathcal{N}(\mu_\theta, \sigma_\theta^2), \quad i = 1, \dots, n.$$

4.1 Posterior and marginal distributions

The posterior distribution for $(f_{k,j}, \sigma_{k,j}^2)$ has the same GP-IG structure as (6), with updated parameters:

$$f_{k,j} | \sigma_{k,j}^2, \mathbf{z}, \theta, \hat{\mathbf{X}} \sim \text{GP}(\mu_{k,j}^*, \sigma_{k,j}^2 \xi_{k,j}^*), \\ \sigma_{k,j}^2 | \mathbf{z}, \theta, \hat{\mathbf{X}} \sim \text{Inv-Gamma}(a_k, b_{k,j}),$$

With $k = 1, \dots, K$, $j = 1, \dots, d$. The parameters are updated as follows:

$$\mu_{k,j}^*(\theta) = \Xi_{k,j}(\theta, \theta_k^*) \{ \Xi_{k,j}(\theta_k^*, \theta_k^*) + \mathbf{I}_{n_k} \}^{-1} \hat{\mathbf{X}}_{k,j}, \\ \xi_{k,j}^*(\theta, \theta') = \xi_{k,j}(\theta, \theta') \\ - \Xi_{k,j}(\theta, \theta_k^*) \{ \Xi_{k,j}(\theta_k^*, \theta_k^*) + \mathbf{I}_{n_k} \}^{-1} \Xi_{k,j}(\theta_k^*, \theta'), \\ a_k = a_0 + n_k/2, \\ b_{k,j} = b_0 + \hat{\mathbf{X}}_{k,j}^\top \{ \Xi_{k,j}(\theta_k^*, \theta_k^*) + \mathbf{I}_{n_k} \}^{-1} \hat{\mathbf{X}}_{k,j}/2, \quad (8)$$

Where $n_k = \sum_{i=1}^n \mathbb{1}_k\{z_i\}$, $\hat{X}_{k,j} \in \mathbb{R}^{n_k}$ is the subset of values of $\hat{X}_j$ for which $z_i = k$, and $\theta_k^* \in \mathbb{R}^{n_k}$ is the vector θ , restricted to the entries such that $z_i = k$. Furthermore, $\Xi_{k,j}$ is a vector-valued and matrix-valued extension of $\xi_{k,j}$, such that $[\Xi_{k,j}(\theta, \theta')]_{\ell, \ell'} = \xi_{k,j}(\theta_\ell, \theta'_{\ell'})$. The structure of the GP-IG yields an analytic expression for the posterior predictive distribution for a new observation $\mathbf{x}^* = (x_1^*, \dots, x_d^*)$ in community z^* ,

$$\hat{x}_j^* | z^*, \mathbf{z}, \theta, \theta^*, \hat{X} \sim t_{2a_{z^*}} \left(\mu_{z^*,j}^*(\theta^*), \frac{b_{z^*,j}^*}{a_{z^*}} \left\{ 1 + \xi_{z^*,j}^*(\theta^*, \theta^*) \right\} \right), \quad (9)$$

Where $t_v(\mu, \sigma)$ denotes a Student's t distribution with v degrees of freedom, mean μ and scale parameter σ . Furthermore, the prior probabilities η for the community assignments can be integrated out, obtaining

$$p(\mathbf{z}) = \frac{\Gamma(v) \prod_{k=1}^K \Gamma(n_k + v/K)}{\Gamma(v/K)^K \Gamma(n + v)}, \quad (10)$$

Where $n_k = \sum_{i=1}^n \mathbb{1}_k\{z_i\}$. The two distributions (9) and (10) are key components for the Bayesian inference algorithm discussed in the next section.

4.2 Posterior inference

After marginalisation of the pairs $(f_{k,j}, \sigma_{k,j}^2)$ and η , inference is limited to the community allocations $\mathbf{z}$ and latent parameters θ . The marginal posterior distribution $p(\mathbf{z}, \theta | \hat{X})$ is analytically intractable; therefore, inference is performed using collapsed Metropolis-within-Gibbs Markov Chain Monte Carlo (MCMC) sampling. In this work, MCMC methods are used, but an alternative inferential algorithm often deployed in the community-detection literature is variational Bayesian inference (see, for example, Latouche et al. 2012), which is also applicable to GPs (Cheng and Boots 2017).

For the community allocations $\mathbf{z}$, the Gibbs sampling step uses the following decomposition:

$$p(z_i = k | \mathbf{z}^{-i}, \hat{X}, \theta) \propto p(z_i = k | \mathbf{z}^{-i}) p(\hat{\mathbf{x}}_i | z_i = k, \mathbf{z}^{-i}, \theta, \hat{X}^{-i}),$$

Where the superscript $-i$ denotes that the i -th row (or element) is removed from the corresponding matrix (or vector). Using (10), the first term is

$$p(z_i = k | \mathbf{z}^{-i}) = \frac{n_k^{-i} + v/K}{n - 1 + v}.$$

For the second term, using (9), the posterior predictive distribution for $\hat{\mathbf{x}}_i$ given $z_i = k$ can be written as the product of

d independent Student's t distributions, where

$$\hat{x}_{i,j} | z_i = k, \mathbf{z}^{-i}, \theta, \hat{X}^{-i} \sim t_{2a_k^{-i}} \left(\mu_{k,j}^{\star-i}(\theta_i), \frac{b_{k,j}^{-i}}{a_k^{-i}} \left\{ 1 + \xi_{k,j}^{\star-i}(\theta_i, \theta_i) \right\} \right). \quad (11)$$

Note that the quantities $\mu_{k,j}^{\star-i}$, $\xi_{k,j}^{\star-i}$, a_k^{-i} and $b_{k,j}^{-i}$ are calculated as described in (8), excluding the contribution of the i -th node.

In order to mitigate identifiability issues, it is necessary to assume that some of the parameters are known a priori. For example, assuming for each community k that $f_{k,1}(\theta_i) = \theta_i$, corresponding to a linear model in θ with no intercept and unit slope in the first dimension, gives the predictive distribution:

$$\hat{x}_{i,1} | z_i = k, \mathbf{z}^{-i}, \theta, \hat{X}^{-i} \sim t_{2a_{k,j}^{-i}} \left(\theta_i, \frac{1}{a_k^{-i}} \left\{ b_0 + \frac{1}{2} \sum_{h \neq i: z_h = k} (\hat{x}_{h,1} - \theta_h)^2 \right\} \right) \quad (12)$$

Finally, for updates to θ_i , a standard Metropolis-within-Gibbs step can be used. For a proposed value θ^* sampled from a proposal distribution $q(\cdot | \theta_i)$, the acceptance probability takes the value

$$\min \left\{ 1, \frac{p(\hat{\mathbf{x}}_i | z_i, \mathbf{z}^{-i}, \theta^*, \theta^{-i}, \hat{X}^{-i}) p(\theta^*) q(\theta_i | \theta^*)}{p(\hat{\mathbf{x}}_i | z_i, \mathbf{z}^{-i}, \theta_i, \theta^{-i}, \hat{X}^{-i}) p(\theta_i) q(\theta^* | \theta_i)} \right\}. \quad (13)$$

The proposal distribution $q(\theta^* | \theta_i)$ in this work is a normal distribution $\mathbb{N}(\theta^* | \theta_i, \sigma_*^2)$, $\sigma_*^2 \in \mathbb{R}_+$, implying that the ratio of proposal distributions in (13) cancels out by symmetry.

4.3 Inference on the number of communities K

So far, it has been assumed that the number of communities K is known. The LSBM prior specification (4) naturally admits a prior distribution on the number of communities K . Following Sanna Passino and Heard (2020), it could be assumed:

$$K \sim \text{Geometric}(\omega),$$

Where $\omega \in (0, 1)$. The MCMC algorithm is then augmented with additional moves for posterior inference on K : (i) Split or merge two communities; and (ii) Add or remove an empty community. An alternative approach when K is unknown could also be a nonparametric mixture of Gaussian processes (Ross and Dy 2013). For simplicity, in the next two sections, it will be initially assumed that all communities have the same functional form, corresponding, for example, to the same basis functions $\phi_{k,j}(\cdot)$ for dot product kernels (7). Then, in

Sect. 4.4, the algorithm will be extended to admit a prior distribution on the community-specific kernels.

4.3.1 Split or merge two communities

In this case, the proposal distribution follows Sanna Passino and Heard (2020). First, two nodes i and j are sampled randomly. For simplicity, assume $z_i \leq z_j$. If $z_i \neq z_j$, then the two corresponding communities are merged into a unique cluster: all nodes in community z_j are assigned to z_i . Otherwise, if $z_i = z_j$, the cluster is split into two different communities, proposed as follows: (i) Node i is assigned to community z_i ($z_i^* = z_i$), and node j to community $K^* = K + 1$ ($z_j^* = K^*$); (ii) The remaining nodes in community z_i are allocated in random order to clusters z_i^* or z_j^* according to their posterior predictive distribution (11) or (12), restricted to the two communities, and calculated sequentially.

It follows that the proposal distribution $q(K^*, \mathbf{z}^* | K, \mathbf{z})$ for a split move corresponds to the product of renormalised posterior predictive distributions, leading to the following acceptance probability:

$$\alpha(K^*, \mathbf{z}^* | K, \mathbf{z}) = \min \left\{ 1, \frac{p(\hat{X} | K^*, \mathbf{z}^*, \boldsymbol{\theta}) p(\mathbf{z}^* | K^*) p(K^*)}{p(\hat{X} | K, \mathbf{z}, \boldsymbol{\theta}) p(\mathbf{z} | K) p(K) q(K^*, \mathbf{z}^* | K, \mathbf{z})} \right\}, \quad (14)$$

Where $p(\hat{X} | K, \mathbf{z}, \boldsymbol{\theta})$ is the marginal likelihood. Note that the ratio of marginal likelihoods only depends on the two communities involved in the split and merge moves. Under (5) and (6), the community-specific marginal on the j -th dimension is:

$$\hat{X}_{k,j} | K, \mathbf{z}, \boldsymbol{\theta} \sim t_{2a} \left(\mathbf{0}_{n_k}, \frac{b_0}{a_0} \{ \boldsymbol{\Sigma}_{k,j}(\boldsymbol{\theta}_k^*, \boldsymbol{\theta}_k^*) + \mathbf{I}_{n_k} \} \right), \quad (15)$$

Where the notation is identical to (8), and the Student's t distribution is n_k -dimensional, with mean equal to the identically zero vector $\mathbf{0}$. The full marginal $p(\hat{X} | K, \mathbf{z}, \boldsymbol{\theta})$ is the product of marginals (15) on all dimensions and communities. If $f_{k,1}(\theta_i) = \theta_i$, cf. (12), the marginal likelihood is

$$p(\hat{X}_{k,1} | K, \mathbf{z}, \boldsymbol{\theta}) = \frac{\Gamma(a_{n_k})}{\Gamma(a_0)} \frac{b_0^{a_0}}{b_{n_k}^{a_{n_k}}} (2\pi)^{-n_k/2},$$

Where $a_{n_k} = a_0 + n_k/2$, $b_{n_k} = b_0 + \sum_{i:z_i=k} (\hat{x}_{i,1} - \theta_i)^2/2$.

4.3.2 Add or remove an empty community

When adding or removing an empty community, the acceptance probability is:

$$\alpha(K^* | K) = \min \left\{ 1, \frac{p(\mathbf{z} | K^*) p(K^*) q_{\emptyset}}{p(\mathbf{z} | K) p(K)} \right\},$$

Where $q_{\emptyset} = q(K | K^*, \mathbf{z}) / q(K^* | K, \mathbf{z})$ is the proposal ratio, equal to (i) $q_{\emptyset} = 2$ if the proposed number of clusters K^* equals the number of non-empty communities in $\mathbf{z}$; (ii) $q_{\emptyset} = 0.5$ if there are no empty clusters in $\mathbf{z}$; and (iii) $q_{\emptyset} = 1$ otherwise. Note that the acceptance probability is identical to Sect. 4.3, Sanna Passino and Heard (2020), and it does not depend on the marginal likelihoods.

4.4 Inference with different community-specific kernels

When communities are assumed to have *different* functional forms, it is required to introduce a prior distribution $p(\boldsymbol{\xi}_k)$, $\boldsymbol{\xi}_k = (\xi_{k,1}, \dots, \xi_{k,d})$, on the GP kernels, supported on one or more classes of possible kernels $\mathcal{K}$. Under this formulation, a proposal to change the community-specific kernel could be introduced. Conditional on the allocations $\mathbf{z}$, the k -th community is assigned a kernel $\boldsymbol{\xi}^* = (\xi_1^*, \dots, \xi_d^*)$ with probability:

$$p(\boldsymbol{\xi}_k = \boldsymbol{\xi}^* | \hat{X}, K, \mathbf{z}, \boldsymbol{\theta}) \propto p(\boldsymbol{\xi}^*) \prod_{j=1}^d p(\hat{X}_{k,j} | K, \mathbf{z}, \boldsymbol{\theta}, \xi_j^*),$$

normalised for $\boldsymbol{\xi}^* \in \mathcal{K}$, where $p(\hat{X}_{k,j} | K, \mathbf{z}, \boldsymbol{\theta}, \xi_j^*)$ is marginal (15) calculated under kernel ξ_j^* . The prior distribution $p(\boldsymbol{\xi}^*)$ could also be used as proposal for the kernel of an empty community (cf. Sect. 4.3.2). Similarly, in the merge move (cf. Sect. 4.3.1), the GP kernel could be sampled at random from the two kernels assigned to z_i and z_j , correcting the acceptance probability (14) accordingly.

5 Applications and results

Inference for the LSBM is tested on synthetic LSBM data and on three real-world networks. As discussed in Sect. 4.2, the first dimension $\hat{X}_1$ is assumed to be linear in θ_i , with no intercept and unit slope. It follows that θ_i is initialised to $\hat{x}_{i,1} + \varepsilon_i$, where $\varepsilon_i \sim \mathbb{N}(0, \sigma_\varepsilon^2)$, for a small σ_ε^2 , usually equal to 0.01. Note that such an assumption links the proposed Bayesian model for LSBM embeddings to Bayesian errors-in-variables models (Dellaportas and Stephens 1995). In the examples in this section, the kernel function is assumed to be of the dot product form (7), with Zellner's g-prior such that

$\Delta_{k,j} = n^2 \{\Phi_{k,j}(\theta)^\top \Phi_{k,j}(\theta)\}^{-1}$, where $\Phi_{k,j}(\theta) \in \mathbb{R}^{n \times q_{k,j}}$ such that the i -th row corresponds to $\phi_{k,j}(\theta_i)$. For the remaining parameters: $a_0 = 1$, $b_0 = 0.001$, $\mu_\theta = \sum_{i=1}^n \hat{x}_{i,1}/n$, $\sigma_\theta^2 = 10$, $v = 1$, $\omega = 0.1$.

The community allocations are initialised using k -means with K groups, unless otherwise specified. The final cluster configuration is estimated from the output of the MCMC algorithm described in Sect. 4.2, using the estimated posterior similarity between nodes i and j , $\hat{\pi}_{ij} = \hat{\mathbb{P}}(z_i = z_j \mid \hat{X}) = \sum_{s=1}^M \mathbb{1}_{z_{i,s}^* \{z_{j,s}^*\}}/M$, where M is the total number of posterior samples and $z_{i,s}^*$ is the s -th sample for z_i . The posterior similarity is not affected by the issue of label switching (Jasra et al. 2005). The clusters are subsequently estimated using hierarchical clustering with average linkage, with distance measure $1 - \hat{\pi}_{ij}$ (Medvedovic et al. 2004). The quality of the estimated clustering compared to the true partition, when available, is evaluated using the Adjusted Rand Index (ARI, Hubert and Arabie 1985). The results presented in this section are based on $M = 10,000$ posterior samples, with 1,000 burn-in period.

5.1 Simulated data: stochastic blockmodels

First, the performance of the LSBM and related inferential procedures is assessed on simulated data from two common models for graph clustering: stochastic blockmodels and degree-corrected stochastic blockmodels (cf. Examples 1 and 2). Furthermore, a graph is also simulated from a block-model with quadratic latent structure (cf. Example 3). In particular, it is evaluated whether the Bayesian inference algorithm for LSBMs recovers the latent communities from the two-dimensional embeddings plotted in Fig. 3 (Sect. 3), postulating the correct latent functions corresponding to the three generative models under consideration. The Gaussian process kernels implied by these three models are (i) $\xi_{k,j}(\theta, \theta') = \Delta_{k,j} \in \mathbb{R}_+$, $k \in \{1, 2\}$, $j \in \{1, 2\}$ for the SBM; (ii) $\xi_{k,2}(\theta, \theta') = \theta \theta' \Delta_{k,2}$, $k \in \{1, 2\}$ for the DCSBM; (iii) $\xi_{k,2}(\theta, \theta') = (\theta, \theta^2) \Delta_{k,2} (\theta', \theta'^2)^\top$ for the quadratic LSBM. For the DCSBM and quadratic LSBM, it is also assumed that $\hat{X}_1$ is linear in θ_i with no intercept and unit slope, corresponding to $f_{k,1}(\theta) = \theta$.

The performance of the inferential algorithm is also compared to alternative methods for spectral graph clustering: Gaussian mixture models (GMM; which underly the assumption of a SBM generative model) on the ASE $\hat{X}$, GMMs on the row-normalised ASE $\tilde{X}$ (see, for example, Rubin-Delanchy et al. 2017), and GMMs on a spherical coordinate transformation of the embedding (SCSC, see Sanna Passino et al. 2021). Note that the two latter methods postulate a DCSBM generative model. The ARI is averaged across 1000 initialisations of the GMM inferential algorithm. Furthermore, the proposed methodology is compared to a kernel-based clus-

tering technique using Gaussian processes: the parsimonious Gaussian process¹ model (PGP) of Bouveyron et al. (2015), fitted using the EM algorithm (PGP-EM). In all the implementations of PGP-EM in Sect. 5, the RBF kernel is used, with variance parameter chosen via grid-search, with the objective to maximise the resulting ARI. The responsibilities in the EM algorithm are initialised using the predictive probabilities obtained from a Gaussian mixture model fit on the ASE or row-normalised ASE, chosen according to the largest ARI. Also, the LSBM is compared to the hierarchical Louvain (HLouvain in Table 3) algorithm² for graphs, corresponding to hierarchical clustering by successive instances of the Louvain algorithm (Blondel et al. 2008). Finally, the LSBM is also compared to hierarchical clustering with complete linkage and Euclidean distance (HClust in Table 3), applied on $\hat{X}$.

The results are presented in Table 1. The performance of the LSBM appears to be on par with alternative methodologies for SBMs and DCSBMs. This is expected, since the LSBM does not have competitive advantage over alternative methodologies for inference under such models. In particular, inferring LSBMs with constant latent functions using the algorithm in Sect. 4.2 is equivalent to Bayesian inference in Gaussian mixture models with normal-inverse gamma priors. For DCSBMs, the LSBM is only marginally outperformed by spectral clustering on spherical coordinates (SCSC, Sanna Passino et al. 2021). On the other hand, LSBMs largely outperform competing methodologies when quadratic latent structure is observed in the embedding. It must be remarked that the coefficients of the latent functions used to generate the quadratic latent structure model in Fig. 3c were chosen to approximately reproduce the curves in the cybersecurity example in Fig. 2. Despite the apparent simplicity of the quadratic LSBM used in this simulation, its practical relevance is evident from plots of real-world network embeddings, and it is therefore important to devise methodologies to estimate its communities correctly. Furthermore, LSBMs could also be used when the number of communities K is unknown. In the examples in this section, inference on K using the MCMC algorithm overwhelmingly suggests $K = 2$, the correct value used in the simulation.

5.2 Simulated data: Hardy–Weinberg LSBM


Second, the performance of the LSBM is assessed on simulated data on a Hardy–Weinberg LSBM. In Athreya et al. (2021) and Trosset et al. (2020), a Hardy–Weinberg latent structure model is used, with $\mathcal{G} = [0, 1]$ and $\mathbf{f}(\theta) = \{\theta^2, 2\theta(1 - \theta), (1 - \theta)^2\}$. A permutation of the Hardy–

¹ PGP-EM was coded in *python* starting from the PGP-DA implementation in the *GitHub* repository mfauve1/PGPDA.

² Implemented in *python* in the library *scikit-network*.

Table 1 ARI for communities estimated using LSBM and alternative methodologies on the embeddings in Fig. 3

Method	LSBM($\hat{X}$)	GMM($\hat{X}$)	GMM($\tilde{X}$)	SCSC($\hat{X}$)	PGP($\hat{X}$)	HLouvain	HClust($\hat{X}$)
SBM	1.0	1.0	1.0	1.0	1.0	1.0	1.0
DCSBM	0.853	0.802	0.838	0.887	0.842	0.827	0.411
Quadratic LSBM	0.838	0.620	0.712	0.636	0.691	0.582	0.101

**Fig. 4** Scatterplots of $\{\hat{X}_1, \hat{X}_2, \hat{X}_3\}$, coloured by community, and true underlying latent positions (in black)

Weinberg curve is considered here for introducing community structure. A graph with $n = 1,000$ nodes is simulated, with $K = 2$ communities of equal size. Each node is assigned a latent score $\theta_i \sim \text{Uniform}(0, 1)$, which is used to obtain the latent position $\mathbf{x}_i$ through the mapping $\mathbf{f}_{z_i}(\theta_i)$. In particular:

$$\mathbf{f}_1(\theta) = \{(1 - \theta)^2, \theta^2, 2\theta(1 - \theta)\},$$

$$\mathbf{f}_2(\theta) = \{\theta^2, 2\theta(1 - \theta), (1 - \theta)^2\}.$$

Using the latent positions, the graph adjacency matrix is then generated under the random dot product graph kernel $\mathbb{P}(A_{ij} = 1 \mid \mathbf{x}_i, \mathbf{x}_j) = \mathbf{x}_i^\top \mathbf{x}_j$. The resulting scatterplot of the latent positions estimated via ASE is plotted in Fig. 4. For visualisation, the estimated latent positions $\hat{X}$ have been aligned to the true underlying latent positions X using a Procrustes transformation (see, for example, Dryden and Mardia 2016).


The inferential procedure is first run assuming that the parametric form of the underlying latent function is known. Therefore, the kernels are set to $\xi_{k,j}(\theta, \theta') = (1, \theta, \theta^2) \Delta_{k,j}(1, \theta', \theta'^2)^\top$, $k = 1, 2$, $j = 1, 2, 3$. Figure 5a shows the best-fitting curves for the two estimated communities after MCMC, which are almost indistinguishable from the true underlying latent curves. The Markov Chain was initialised setting $\theta_i = |\hat{x}_{i,1}|^{1/2}$, and obtaining initial values of the allocations $\mathbf{z}$ from k -means. The resulting ARI is 0.7918, which is not perfect since some of the nodes at the intersec-

tion between the two curves are not classified correctly, but corresponds to approximately 95% of nodes correctly classified. The number of communities is estimated to be $K = 2$ from the inferential algorithm, the correct value used in the simulation.

If the underlying functional relationship is unknown, a realistic guess could be given by examining the scatterplots of the embedding. The scatterplots in Fig. 4 show that, assuming linearity in θ_i on $\hat{X}_1$ with no intercept and unit slope, a quadratic or cubic polynomial function is required to model $\hat{X}_2$ and $\hat{X}_3$. Therefore, for the purposes of the MCMC inference algorithm in Sect. 4.2, $f_2(\cdot)$ and $f_3(\cdot)$ are assumed to be cubic functions, corresponding to the Gaussian process kernel $\xi_{k,j}(\theta, \theta') = (1, \theta, \theta^2, \theta^3) \Delta_{k,j}(1, \theta', \theta'^2, \theta'^3)^\top$, $j \in \{2, 3\}$, whereas $f_{k,1}(\theta) = \theta$, $k = 1, \dots, K$. The curves corresponding to the estimated clustering are plotted in Fig. 5b. Also in this case, the algorithm is able to approximately recover the curves that generated the graph, and the number of communities $K = 2$. The imperfect choice of the latent functions makes the ARI decrease to 0.6687, which still corresponds to over 90% of nodes correctly classified.

Table 2 presents further comparisons with the alternative techniques discussed in the previous section, showing that only LSBMs recover a good portion of the communities. This is hardly surprising, since alternative methodologies do not take into account the quadratic latent functions generating the latent positions, and therefore fail to capture the underlying community structure.

(a) Quadratic latent functions



(b) Cubic latent functions, linearity in first dimension




Fig. 5 Scatterplot of $\hat{X}$ and estimated generating curves obtained from the estimated clustering, coloured by community, and true underlying latent positions (in black)

Table 2 ARI for communities estimated using LSBM and alternative methodologies on the embeddings in Fig. 4

Method	LSBM($\hat{X}$) Quadratic	LSBM($\hat{X}$) Cubic + Linear	GMM($\hat{X}$)	GMM($\tilde{X}$)	SCSC($\hat{X}$)	PGP($\hat{X}$)	HLouvain	HClust($\hat{X}$)
ARI	0.791	0.668	−0.001	−0.001	0.000	0.063	−0.001	0.053




Fig. 6 Estimated marginal posterior for the number of non-empty clusters K on the Harry Potter network, under linear kernels

5.3 Undirected graphs: Harry Potter enmity graph

The LSBM is also applied to the Harry Potter enmity graph³, an undirected network with $n = 51$ nodes representing char-

acters of J.K. Rowling's series of fantasy novels. In the graph, $A_{ij} = A_{ji} = 1$ if the characters i and j are enemies, and 0 otherwise. A degree-corrected stochastic blockmodel represents a reasonable assumption for such a graph (see, for example, Modell and Rubin-Delanchy 2021): Harry Potter, the main character, and Lord Voldemort, his antagonist, attract many enemies, resulting in a large degree, whereas their followers are expected to have lower degree. The graph might be expected to contain $K = 2$ communities, since Harry Potter's friends are usually Lord Voldemort's enemies in the novels. Theorem 1 suggests that the embeddings for a DCSBM are expected to appear as *rays* passing through the origin. Therefore, $\mathbf{f}_k(\cdot)$ is assumed to be composed of linear functions such that $f_{k,1}(\theta) = \theta$ and $\xi_{k,2}(\theta, \theta') = \theta\theta'\Delta_{k,2}$. The estimated posterior distribution of the number of non-empty clusters is plotted in Fig. 6. It appears that $K = 2$ is the most appropriate number of clusters. The resulting 2-dimensional ASE embedding of A and the estimated clustering obtained using the linear LSBM are pictured in Fig. 7a. The two estimated communities roughly correspond to the Houses of Gryffindor and Slytherin. The inferential algorithm often admits $K = 3$, with a singleton cluster for

³ The network is publicly available in the *GitHub* repository [efekarakus/potter-network](https://github.com/efekarakus/potter-network).




Fig. 7 Scatterplots of $\{\hat{X}_2, \hat{X}_3\}$ vs. $\hat{X}_1$, coloured by community allocation, and best-fitting line and cubic truncated power spline passing through the origin, obtained from the estimated clustering

Quirinus Quirrell, which appears to be between the two main linear groups (cf. Fig. 7b).

Alternatively, if a polynomial form for $\mathbf{f}_k(\cdot)$ is unknown, a more flexible model is represented by regression splines, which can be also expressed in form (7). A common choice for $\phi_{k,j}(\cdot)$ is a cubic truncated power basis $\phi_{k,j} = (\phi_{k,j,1}, \dots, \phi_{k,j,6})$, $\ell \in \mathbb{Z}_+$, such that:

$$\begin{aligned} \phi_{k,j,1}(\theta) &= \theta, \quad \phi_{k,j,2}(\theta) = \theta^2, \quad \phi_{k,j,3}(\theta) = \theta^3, \\ \phi_{k,j,3+\ell}(\theta) &= (\theta - \tau_\ell)_+^3, \quad \ell = 1, 2, 3, \end{aligned} \quad (16)$$

Where (τ_1, τ_2, τ_3) are knots, and $(\cdot)_+ = \max\{0, \cdot\}$. In this application, the knots were selected as three equispaced points in the range of $\hat{X}_1$. The results are plotted in Fig. 7b. Using either functional form, the algorithm is clearly able to recover the two communities, meaningfully clustering Harry Potter's and Lord Voldemort's followers.

5.4 Directed graphs: *Drosophila* connectome

LSBM are also useful to cluster the larval *Drosophila* mushroom body connectome (Eichler et al. 2017), a directed graph representing connections between $n = 213$ neurons in the brain of a species of fly⁴. The right hemisphere mushroom body connectome contains $K = 4$ groups of neurons: Kenyon Cells, Input Neurons, Output Neurons and Projection Neurons. If two neurons are connected, then $A_{ij} = 1$, otherwise $A_{ij} = 0$, forming an asymmetric adjacency matrix $\mathbf{A} \in \{0, 1\}^{n \times n}$. The network has been extensively analysed in Priebe et al. (2017) and Athreya et al. (2018).

Following Priebe et al. (2017) and Athreya et al. (2018), after applying the DASE (Definition 3) for $d = 3$, a joint concatenated embedding $\hat{\mathbf{Y}} = [\hat{\mathbf{X}}, \hat{\mathbf{X}}'] \in \mathbb{R}^{n \times 2d}$ is obtained from $\mathbf{A}$. Based on the analysis of Priebe et al. (2017), it should be assumed that three of the communities (Input Neurons, Output Neurons and Projection Neurons) correspond to a stochastic blockmodel, resulting in Gaussian clusters, whereas the Kenyon Cells form a quadratic curve with respect to the first dimension in the embedding space. Therefore, the kernel functions implied in Priebe et al. (2017) are $\xi_{1,j}(\theta, \theta') = (\theta, \theta^2) \mathbf{\Delta}_{1,j}(\theta', \theta'^2) \boldsymbol{\tau}$, $j = 2, \dots, 2d$, with $f_{1,1}(\theta) = \theta$, for the first community (corresponding to the Kenyon Cells), and $\xi_{k,j}(\theta, \theta') = \Delta_{k,j}$, $k = 2, 3, 4$, $j = 1, \dots, 2d$ for the three remaining groups of neurons.

Following the discussion in Priebe et al. (2017), the LSBM inferential procedure is initialised using k -means with $K = 6$, and grouping three of the clusters to obtain $K = 4$ initial groups. Note that, since the output of k -means is invariant to permutations of the $K = 4$ labels for the community allocations, careful relabelling of the initial values is necessary to ensure that the Kenyon Cells effectively correspond to the first community which assumes a quadratic functional form. The most appropriate relabelling mapping is chosen here by repeatedly initialising the model with *all* possible permutations of the labels, and choosing the permutation that maximises the marginal likelihood. Note that the marginal likelihood under the Bayesian model (6) for LSBMs is analytically available in closed form (see, for example, Rasmussen and Williams 2006). The results obtained after MCMC sampling are plotted in Fig. 8. The estimated clustering has ARI 0.8643, corresponding to only 10 misclassified nodes out of 213.

⁴ The data are publicly available in the *GitHub* repository [youngser/mbstructure](https://github.com/youngser/mbstructure).




Fig. 8 Scatterplots of $\{\hat{Y}_2, \hat{Y}_3, \hat{Y}_4, \hat{Y}_5, \hat{Y}_6\}$ vs. $\hat{Y}_1$, coloured by neuron type, and best-fitting latent functions obtained from the estimated clustering

Table 3 ARI for communities estimated via LSBM and alternative methodologies on the *Drosophila* connectome

Method	LSBM($\hat{Y}$) Polynomial	LSBM($\hat{Y}$) Quadratic + SBM	GMM($\hat{Y}$)	GMM($\tilde{Y}$)	SCSC($\hat{Y}$)	PGP($\hat{Y}$)	HLouvain	HClust($\hat{Y}$)
ARI	0.875	0.864	0.599	0.585	0.667	0.555	0.087	0.321

From the scatterplots in Fig. 8, it appears that the embedding for the Input Neurons, Output Neurons and Projection Neurons could be also represented using linear functions. Furthermore, a quadratic curve for the Kenyon Cells might be too restrictive. The LSBM framework allows specification of all such choices. Here, assuming $f_{k,1}(\theta) = \theta$ for each k , the first community (Kenyon Cells) is given a cubic latent function, implying $\xi_{1,j}(\theta, \theta') = (\theta, \theta^2, \theta^3) \Delta_{1,j}(\theta', \theta'^2, \theta'^3)^\top$. For the second community (Input Neurons), a latent linear function is used: $\xi_{2,j}(\theta, \theta') = (1, \theta) \Delta_{2,j}(1, \theta')^\top$, $j = 2, \dots, 6$. Similarly, from observation of the scatterplots in Fig. 8, the following kernels, corresponding to linear latent functions, are assigned to the remaining communities: $\xi_{3,j}(\theta, \theta') = \theta \theta' \Delta_{3,j}$, $j = 2, \dots, 6$, $\xi_{4,2}(\theta, \theta') = \theta \theta' \Delta_{4,2}$, and $\xi_{4,j}(\theta, \theta') = (1, \theta) \Delta_{k,j}(1, \theta')^\top$, $j = 3, \dots, 6$. The results for these choices of covariance kernels are plotted in

Fig. 9, resulting in ARI 0.8754 for the estimated clustering, again corresponding to 10 misclassified nodes. Note that this representation seems to capture more closely the structure of the embeddings.

The performance of the LSBM is also compared to alternative methods for clustering in Table 3. In particular, Gaussian mixture models with $K = 4$ components were fitted on the concatenated DASE embedding $\hat{Y} = [\hat{X}, \hat{X}']$ (standard spectral clustering; see, for example Rubin-Delanchy et al. 2017), on its row-normalised version $\tilde{Y}$ (Ng et al. 2001; Qin and Rohe 2013), and on a transformation to spherical coordinates (Sanna Passino et al. 2021). The ARI is averaged over 1000 different initialisations. Furthermore, the LSBM is also compared to PGP (Bouveyron et al. 2015), hierarchical Louvain adapted to directed graphs (Dugué and Perez 2015), and




Fig. 9 Scatterplots of $\{\hat{Y}_2, \hat{Y}_3, \hat{Y}_4, \hat{Y}_5, \hat{Y}_6\}$ vs. $\hat{Y}_1$, coloured by neuron type, and best-fitting latent functions obtained from the estimated clustering, plotted only over the range of nodes assigned to each community




Fig. 10 Estimated marginal posterior for the number of non-empty clusters K on the *Drosophila* connectome, under two different mixture prior distributions on the kernels




Fig. 11 Scatterplots of $\{\hat{X}_2, \hat{X}_3, \hat{X}_4, \hat{X}_5\}$ vs. $\hat{X}_1$, coloured by department, and best-fitting quadratic curves passing through the origin

hierarchical clustering on $\hat{Y}$. The results in Table 3 show that the LSBM largely outperform the alternative clustering techniques, which are not able to account for the nonlinearity in the community corresponding to the Kenyon Cells.

Next, the MCMC algorithm is run for $M = 50,000$ iterations with 5,000 burn-in period, with the objective of inferring K , assuming a discrete uniform prior on the kernels used in Figs. 8 and 9. The estimated posterior barplots for the number of non-empty clusters K are plotted in Fig. 10. It appears that the choice of a prior on kernels admitting a quadratic curve for the Kenyon Cells and Gaussian clusters for the remaining groups is overly simplistic, since some of the communities appear to have a linear structure (cf. Fig. 8). Hence, the number of communities K is incorrectly estimated under this setting (cf. Fig. 10a). On the other hand, when a mixture distribution of cubic and linear kernels is used as prior on ξ_k , K is correctly estimated (cf. Fig. 10b),

confirming the impression from Fig. 9 of a better fit of such kernels on the *Drosophila* connectome.

5.5 Bipartite graphs: ICL computer laboratories

The LSBM methodology is finally applied to a bipartite graph obtained from computer network flow data collected at Imperial College London (ICL). The source nodes are $|V_1| = n = 439$ client machines located in four computer laboratories in different departments at Imperial College London, whereas the destination nodes are $|V_2| = 60,635$ Internet servers, connected to by HTTP and HTTPS in January 2020 from one or more of the 439 client computers. A total of 717,912 edges are observed. The inferential objective is to identify the location of the machines in the network, represented by their department, from the realisation of the rectangular adjacency matrix A , where $A_{ij} = 1$ if at least one connection is observed between client computer $i \in V_1$




Fig. 12 Estimated marginal posterior for the number of non-empty clusters K on the ICL NetFlow network, under quadratic kernels

and the server $j \in V_2$, and $A_{ij} = 0$ otherwise. It could be assumed that $K = 4$, representing the departments of Chemistry, Civil Engineering, Mathematics, and Medicine. After taking the DASE of A , the machines are clustered using the LSBM. The value $d = 5$ is selected using the criterion of Zhu and Ghodsi (2006), choosing the second elbow of the scree-plot of singular values. Computer network graphs of this kind have been seen to present quadratic curves in the embedding, as demonstrated, for example, by Fig. 2 in the introduction, which refers to a different set of machines. Therefore, it seems reasonable to assume that $\xi_{k,j}(\theta, \theta') = (\theta, \theta^2) \Delta_{k,j}(\theta', \theta'^2)^T$, $j = 2, \dots, d$, which implies quadratic functions passing through the origin, and $f_{k,1}(\theta) = \theta$. The quadratic model with $K = 4$ is fitted to the 5-dimensional embedding, obtaining ARI 0.9402, corresponding to just 9 misclassified nodes. The results are plotted in Fig. 11, with the corresponding best-fitting quadratic curves obtained from the estimated clustering. The result appears remarkable, considering that the communities are highly overlapping. Running the MCMC algorithm assuming K unknown shows that the number of clusters is correctly estimated (cf. Fig. 12), overwhelmingly suggesting $K = 4$, corresponding to the correct number of departments.

These LSBM results are also compared to alternative methodologies in Table 4. LSBM achieves the best performance in clustering the nodes, followed by Gaussian mixture modelling of the row-normalised adjacency spectral embedding. The GMM on $\tilde{X}$ sometimes converges to competitive

solutions, reaching ARI up to 0.94, but usually converges to sub-optimal solutions, as demonstrated by the average ARI of 0.7608.

As before, if a parametric form for $f_k(\cdot)$ is unknown, regression splines can be used, for example the truncated power basis (16), with three equispaced knots in the range of $\tilde{X}_1$. The results for the initial three dimensions are plotted in Fig. 13. The communities are recovered correctly, and the ARI is 0.9360, corresponding to 10 misclassified nodes.

6 Conclusion

An extension of the latent structure model (Athreya et al. 2021) for networks has been introduced, and inferential procedures based on Bayesian modelling of spectrally embedded nodes have been proposed. The model, referred to as the latent structure blockmodel (LSBM), allows for latent positions living on community-specific univariate structural support manifolds, using flexible Gaussian process priors. Under the Bayesian paradigm, most model parameters can be integrated out and inference can be performed efficiently. The proposed modelling framework could be utilised also when the number of communities is unknown and there is uncertainty around the choice the Gaussian process kernels, encoded by a prior distribution. The performance of the model inference has been evaluated on simulated and real-world networks. In particular, excellent results have been obtained on complex clustering tasks concerning the *Drosophila* connectome and the Imperial College NetFlow data, where a substantial overlap between communities is observed. Despite these challenges, the proposed methodology is still able to recover a correct clustering.

Overall, this work provides a modelling framework for graph embeddings arising from random dot product graphs where it is suspected that nodes belong to community-specific lower-dimensional subspaces. In particular, this article discusses the case of curves, which are one-dimensional structural support submanifolds. The methodology has been demonstrated to have the potential to recover the correct clustering structure even if the underlying parametric form of the underlying structure is unknown, using a flexible Gaussian process prior on the unknown functions. In particular, regression splines with a truncated power basis have been

Table 4 ARI for communities estimated via LSBM and alternative methodologies on the ICL NetFlow network

Method	LSBM($\hat{X}$) Quadratic	LSBM($\hat{X}$) Splines	GMM($\hat{X}$)	GMM($\tilde{X}$)	SCSC($\hat{X}$)	PGP($\hat{X}$)	HLouvain	HClust($\hat{X}$)
ARI	0.940	0.936	0.659	0.766	0.921	0.895	0.602	0.139




Fig. 13 Scatterplots of $\{\hat{X}_2, \hat{X}_3, \hat{X}_4, \hat{X}_5\}$ vs. $\hat{X}_1$, coloured by department, and estimated cubic truncated power splines passing through the origin

used, showing good performance in recovering the underlying curves.

In the model proposed in this work, it has been assumed that the variance only depends on the community allocation. This enables marginalising most of the parameters leading to efficient inference. On the other hand, this is potentially an oversimplification, since the ASE-CLT (Theorem 1) establishes that the covariance structure depends on the latent position. Further work should study efficient algorithms for estimating the parameters when an explicit functional form dependent on θ_i is incorporated in the covariance.

Code

A *python* library to reproduce the results in this paper is available in the *GitHub* repository `fraspas/lbm`.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Amini, A.A., Razaee, Z.S.: Concentration of kernel matrices with application to kernel spectral clustering. *Ann. Stat.* **49**(1), 531–556 (2021)
- Asta, D.M., Shalizi, C.R.: Geometric network comparisons. In: *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*. pp. 102–110. UAI'15, AUAI Press (2015)
- Athreya, A., Priebe, C.E., Tang, M., Lyzinski, V., Marchette, D.J., Sussman, D.L.: A limit theorem for scaled eigenvectors of random dot product graphs. *Sankhya A* **78**(1), 1–18 (2016)
- Athreya, A., Fishkind, D.E., Tang, M., Priebe, C.E., Park, Y., Vogelstein, J.T., Levin, K., Lyzinski, V., Qin, Y., Sussman, D.L.: Statistical inference on random dot product graphs: a survey. *J. Mach. Learn. Res.* **18**(226), 1–92 (2018)
- Athreya, A., Tang, M., Park, Y., Priebe, C.E.: On estimation and inference in latent structure random graphs. *Stat. Sci.* **36**(1), 68–88 (2021)
- Blondel, V.D., Guillaume, J.L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. *J. Stat. Mech. Theory Exp.* **10**, P10008 (2008)
- Bouveyron, C., Fauvel, M., Girard, S.: Kernel discriminant analysis and clustering with parsimonious Gaussian process models. *Stat. Comput.* **25**(6), 1143–1162 (2015)
- Cheng, C.A., Boots, B.: Variational inference for gaussian process models with linear complexity. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. pp. 5190–5200. NIPS'17, Curran Associates Inc., Red Hook, NY, USA (2017)
- Couillet, R., Benaych-Georges, F.: Kernel spectral clustering of large dimensional data. *Electron. J. Stat.* **10**(1), 1393–1454 (2016)
- Dellaportas, P., Stephens, D.: Bayesian analysis of errors-in-variables regression models. *Biometrics* **51**, 1085–1095 (1995)
- Diday, E.: Une nouvelle méthode en classification automatique et reconnaissance des formes la méthode des nuées dynamiques. *Revue de Statistique Appliquée* **19**(2), 19–33 (1971)
- Diday, E., Simon, J.C.: *Clustering analysis*, pp. 47–94. Springer, Berlin Heidelberg, Berlin, Heidelberg (1976)
- Dryden, I.L., Mardia, K.V.: *Statistical shape analysis, with applications in R*. John Wiley and Sons (2016)
- Dugué, N., Perez, A.: Directed Louvain: maximizing modularity in directed networks. Tech. Rep. hal-01231784, Université d'Orléans (2015)
- Dunson, D.B., Wu, N.: Inferring manifolds from noisy data using Gaussian processes. arXiv e-prints (2021)
- Eichler, K., Li, F., Litwin-Kumar, A., Park, Y., Andrade, I., Schneider-Mizell, C.M., Saumweber, T., Huser, A., Eschbach, C., Gerber, B., Fetter, R.D., Truman, J.W., Priebe, C.E., Abbott, L.F., Thum, A.S., Zlatić, M., Cardona, A.: The complete connectome of a learning and memory centre in an insect brain. *Nature* **548**(7666), 175–182 (2017)
- Han, X., Tong, X., Fan, Y.: Eigen selection in spectral clustering: a theory-guided practice. *J. Am. Stat. Assoc.* **0**(0), 1–13 (2021)
- Hoff, P.D., Raftery, A.E., Handcock, M.S.: Latent space approaches to social network analysis. *J. Am. Stat. Assoc.* **97**(460), 1090–1098 (2002)
- Hofmeyr, D.P.: Improving spectral clustering using the asymptotic value of the normalized cut. *J. Comput. Graph. Stat.* **28**(4), 980–992 (2019)
- Holland, P.W., Laskey, K.B., Leinhardt, S.: Stochastic blockmodels: first steps. *Soc. Netw.* **5**(2), 109–137 (1983)
- Hubert, L., Arabie, P.: Comparing partitions. *J. Classif.* **2**(1), 193–218 (1985)
- Jasra, A., Holmes, C.C., Stephens, D.A.: Markov Chain Monte Carlo methods and the label switching problem in Bayesian mixture modeling. *Stat. Sci.* **20**(1), 50–67 (2005)
- Jones, A., Rubin-Delanchy, P.: The multilayer random dot product graph. arXiv e-prints (2020)
- Karrer, B., Newman, M.E.J.: Stochastic blockmodels and community structure in networks. *Phys. Rev. E* **83**(1), 016107 (2011)
- Kerbin, C., Brault, V., Celeux, G., Govaert, G.: Estimation and selection for the latent block model on categorical data. *Stat. Comput.* **25**(6), 1201–1216 (2015)
- Kiefer, J., Wolfowitz, J.: Consistency of the maximum likelihood estimator in the presence of infinitely many incidental parameters. *Ann. Math. Stat.* **27**(4), 887–906 (1956)
- Latouche, P., Birmelé, E., Ambroise, C.: Variational Bayesian inference and complexity control for stochastic block models. *Stat. Model.* **12**(1), 93–115 (2012)
- Lázaro-Gredilla, M., Van Vaerenbergh, S., Lawrence, N.D.: Overlapping mixtures of Gaussian processes for the data association problem. *Pattern Recogn.* **45**(4), 1386–1395 (2012)
- Lindsay, B.G.: The geometry of mixture likelihoods: a general theory. *Ann. Stat.* **11**(1), 86–94 (1983)
- Luo, B., Wilson, R.C., Hancock, E.R.: Spectral embedding of graphs. *Pattern Recognit.* **36**(10), 2213–2230 (2003)
- Lyzinski, V., Sussman, D.L., Tang, M., Athreya, A., Priebe, C.E.: Perfect clustering for stochastic blockmodel graphs via adjacency spectral embedding. *Electron. J. Stat.* **8**(2), 2905–2922 (2014)
- McCormick, T.H., Zheng, T.: Latent surface models for networks using aggregated relational data. *J. Am. Stat. Assoc.* **110**(512), 1684–1695 (2015)
- Medvedovic, M., Yeung, K.Y., Bumgarner, R.E.: Bayesian mixture model based clustering of replicated microarray data. *Bioinformatics* **20**(8), 1222–1232 (2004)
- Modell, A., Rubin-Delanchy, P.: Spectral clustering under degree heterogeneity: a case for the random walk Laplacian. arXiv e-prints (2021)
- Ng, A.Y., Jordan, M.I., Weiss, Y.: On spectral clustering: Analysis and an algorithm. In: *Proceedings of the 14th International Conference on Neural Information Processing Systems*. pp. 849–856 (2001)
- Pensky, M., Zhang, T.: Spectral clustering in the dynamic stochastic block model. *Electron. J. Stat.* **13**(1), 678–709 (2019)
- Priebe, C.E., Park, Y., Tang, M., Athreya, A., Lyzinski, V., Vogelstein, J.T., Qin, Y., Cocanougher, B., Eichler, K., Zlatić, M., Cardona, A.: Semiparametric spectral modeling of the Drosophila connectome. arXiv e-prints (2017)
- Priebe, C.E., Park, Y., Vogelstein, J.T., Conroy, J.M., Lyzinski, V., Tang, M., Athreya, A., Cape, J., Bridgeford, E.: On a two-truths phenomenon in spectral graph clustering. *Proc. Natl. Acad. Sci.* **116**(13), 5995–6000 (2019)
- Qin, T., Rohe, K.: Regularized spectral clustering under the degree-corrected stochastic blockmodel. In: *Proceedings of the 26th International Conference on Neural Information Processing Systems*. vol. 2, pp. 3120–3128 (2013)
- Rasmussen, C.E., Williams, C.K.I.: *Gaussian Processes for Machine Learning*. The MIT Press, Adaptive Computation and Machine Learning (2006)
- Rohe, K., Chatterjee, S., Yu, B.: Spectral clustering and the high-dimensional stochastic blockmodel. *Ann. Stat.* **39**(4), 1878–1915 (2011)
- Ross, J.C., Dy, J.G.: Nonparametric mixture of Gaussian processes with constraints. In: *Proceedings of the 30th International Conference on Machine Learning - Volume 28. ICML'13* (2013)
- Rubin-Delanchy, P.: Manifold structure in graph embeddings. In: Laroche, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) *Advances in Neural Information Processing Systems*. vol. 33, pp. 11687–11699. Curran Associates, Inc. (2020)

- Rubin-Delanchy, P., Cape, J., Tang, M., Priebe, C.E.: A statistical interpretation of spectral embedding: the generalised random dot product graph. *arXiv e-prints* (2017)
- Salter-Townshend, M., McCormick, T.H.: Latent space models for multiview network data. *Ann. Appl. Stat.* **11**(3), 1217–1244 (2017)
- Sanna Passino, F., Heard, N.A.: Bayesian estimation of the latent dimension and communities in stochastic blockmodels. *Stat. Comput.* **30**(5), 1291–1307 (2020)
- Sanna Passino, F., Heard, N.A., Rubin-Delanchy, P.: Spectral clustering on spherical coordinates under the degree-corrected stochastic blockmodel. *Technometrics* (to appear) (2021)
- Scholkopf, B., Smola, A.J.: Learning with kernels: support vector machines, regularization, optimization, and beyond. MIT Press, Adaptive Computation and Machine Learning series (2018)
- Shepard, R.N.: The analysis of proximities: Multidimensional scaling with an unknown distance function. i. *Psychometrika* **27**(2), 125–140 (1962)
- Smith, A.L., Asta, D.M., Calder, C.A.: The geometry of continuous latent space models for network data. *Stat. Sci.* **34**(3), 428–453 (2019)
- Sussman, D.L., Tang, M., Fishkind, D.E., Priebe, C.E.: A consistent adjacency spectral embedding for stochastic blockmodel graphs. *J. Am. Stat. Assoc.* **107**(499), 1119–1128 (2012)
- Tenenbaum, J.B., de Silva, V., Langford, J.C.: A global geometric framework for nonlinear dimensionality reduction. *Science* **290**(5500), 2319–2323 (2000)
- Todeschini, A., Miskouridou, X., Caron, F.: Exchangeable random measures for sparse and modular graphs with overlapping communities. *J. Roy. Stat. Soc. B* **82**(2), 487–520 (2020)
- Torgerson, W.S.: Multidimensional scaling: I. Theory and method. *Psychometrika* **17**(4), 401–419 (1952)
- Trosset, M.W., Gao, M., Tang, M., Priebe, C.E.: Learning 1-dimensional submanifolds for subsequent inference on random dot product graphs. *arXiv e-prints* (2020)
- von Luxburg, U.: A tutorial on spectral clustering. *Stat. Comput.* **1**(4), 395–416 (2007)
- Wyse, J., Friel, N., Latouche, P.: Inferring structure in bipartite networks using the latent blockmodel and exact ICL. *Netw. Sci.* **5**(1), 45–69 (2017)
- Yang, C., Priebe, C.E., Park, Y., Marchette, D.J.: Simultaneous dimensionality and complexity model selection for spectral graph clustering. *J. Comput. Graph. Stat.* **30**(2), 422–441 (2021)
- Ye, X., Zhao, J., Chen, Y., Guo, L.J.: Bayesian adversarial spectral clustering with unknown cluster number. *IEEE Trans. Image Process.* **29**, 8506–8518 (2020)
- Young, S.J., Scheinerman, E.R.: Random dot product graph models for social networks. In: Bonato, A., Chung, F.R.K. (eds.) *Algorithms and Models for the Web-Graph*, pp. 138–149. Springer, Berlin Heidelberg, Berlin, Heidelberg (2007)
- Zhu, M., Ghodsi, A.: Automatic dimensionality selection from the scree plot via the use of profile likelihood. *Comput. Stat. Data. Anal.* **51**(2), 918–930 (2006)
- Zhu, X., Zhang, S., Li, Y., Zhang, J., Yang, L., Fang, Y.: Low-rank sparse subspace for spectral clustering. *IEEE Trans. Knowl. Data Eng.* **31**(8), 1532–1543 (2019)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.