

Covariate-assisted spectral clustering

BY N. BINKIEWICZ

*Department of Statistics, University of Wisconsin, 1300 University Avenue,
Madison, Wisconsin 53706, U.S.A.*

norbertbin@gmail.com

J. T. VOGELSTEIN

*Department of Biomedical Engineering, Johns Hopkins University, 720 Rutland Avenue,
Baltimore, Maryland 21205, U.S.A.*

jovo@jhu.edu

AND K. ROHE

*Department of Statistics, University of Wisconsin, 1300 University Avenue,
Madison, Wisconsin 53706, U.S.A.*

karlrohe@stat.wisc.edu

SUMMARY

Biological and social systems consist of myriad interacting units. The interactions can be represented in the form of a graph or network. Measurements of these graphs can reveal the underlying structure of these interactions, which provides insight into the systems that generated the graphs. Moreover, in applications such as connectomics, social networks, and genomics, graph data are accompanied by contextualizing measures on each node. We utilize these node covariates to help uncover latent communities in a graph, using a modification of spectral clustering. Statistical guarantees are provided under a joint mixture model that we call the node-contextualized stochastic blockmodel, including a bound on the misclustering rate. The bound is used to derive conditions for achieving perfect clustering. For most simulated cases, covariate-assisted spectral clustering yields results superior both to regularized spectral clustering without node covariates and to an adaptation of canonical correlation analysis. We apply our clustering method to large brain graphs derived from diffusion MRI data, using the node locations or neurological region membership as covariates. In both cases, covariate-assisted spectral clustering yields clusters that are easier to interpret neurologically.

Some key words: Brain graph; Laplacian; Network; Node attribute; Stochastic blockmodel.

1. INTRODUCTION

Modern experimental techniques in areas such as genomics and brain imaging generate vast amounts of structured data, which contain information about the relationships of genes or brain regions. Studying these relationships is essential for solving challenging scientific problems, but few computationally feasible statistical techniques incorporate both the structure and diversity of these data.

A common approach to understanding the behaviour of a complex biological or social system is to first discover blocks of highly interconnected units, also known as communities or clusters, that serve or contribute to a common function. These might be genes that are involved in a common pathway or areas in the brain with a common neurological function. Typically, we only observe the pairwise relationships between the units, which can be represented by a graph or network. Analysing networks has become an important part of the social and biological sciences. Examples of such networks include gene regulatory networks, friendship networks, and brain graphs. If we can discover the underlying block structure of such graphs, we can gain insight from the common characteristics or functions of the units within a block.

Existing research has extensively studied the algorithmic and theoretical aspects of finding node clusters within a graph, by Bayesian, maximum likelihood, and spectral approaches. Unlike model-based methods, spectral clustering is a relaxation of a cost minimization problem and has been shown to be effective in various settings (Ng et al., 2002; Von Luxburg, 2007). Modifications of spectral clustering, such as regularized spectral clustering, are accurate even for sparse networks (Chaudhuri et al., 2012; Amini et al., 2013; Qin & Rohe, 2013). On the other hand, certain Bayesian methods offer additional flexibility in how nodes are assigned to blocks, allowing for a single node to belong to multiple blocks or a mixture of blocks (Nowicki & Snijders, 2001; Airolidi et al., 2008). Maximum likelihood approaches can enhance interpretability by embedding nodes in a latent social space and providing methods for quantifying statistical uncertainty (Hoff et al., 2002; Handcock et al., 2007; Amini et al., 2013). For large graphs, spectral clustering is one of very few computationally feasible methods that has an algorithmic guarantee for finding the globally optimal partition.

The structured data generated by modern technologies often contain additional measurements that can be represented as graph node attributes or covariates. For example, these could be personal profile information in a friendship network or the spatial location of a brain region in a brain graph. There are two potential advantages of utilizing node covariates in graph clustering. First, if the covariates and the graph have a common latent structure, then the node covariates provide additional information to help estimate this structure. Even if the covariates and the graph do not share exactly the same structure, some similarity is sufficient for the covariates to assist in the discovery of the graph structure. Second, by using node covariates in the clustering procedure, we enhance the relative homogeneity of covariates within a cluster and filter out partitions that fail to align with the important covariates. This allows for easy contextualization of the clusters in terms of the member nodes' covariates, providing a natural way to interpret the clusters.

Methods that utilize both node covariates and the graph to cluster the nodes have previously been introduced, but many of them rely on ad hoc or heuristic approaches and none provide theoretical guarantees for statistical estimation. Most existing methods can be broadly classified into Bayesian approaches, spectral techniques, and heuristic algorithms. Many Bayesian models focus on categorical node covariates and are often computationally expensive (Chang & Blei, 2010; Balasubramanian & Cohen, 2011). A recent Bayesian model proposed by Yang et al. (2013) can discover multi-block membership of nodes with binary node covariates. This method has linear update time in the network size, but does not guarantee linear-time convergence. Heuristic algorithms use various approaches, including embedding the network in a vector space, at which point more traditional methods can be applied to the vector data (Gibert et al., 2012), or using the covariates to augment the graph and applying other graph clustering methods that tune the relative weights of node-to-node and node-to-covariate edges (Zhou et al., 2009). A commonly-used spectral approach to incorporate node covariates directly alters the edge weights based on the similarity of the corresponding nodes' covariates, and uses traditional spectral clustering on the weighted graph (Neville et al., 2003; Gunnemann et al., 2013).

This work introduces a spectral approach that performs well for assortative graphs and another that does not require this restriction. We give a standard definition of an assortative graph here and later define it in the context of a stochastic blockmodel.

DEFINITION 1 (Assortative graph). *A graph is assortative if nodes within the same cluster are more likely to share an edge than nodes in two different clusters.*

Assortative covariate-assisted spectral clustering adds the covariance matrix of the node covariates to the regularized graph Laplacian, boosting the signal in the top eigenvectors of the sum, which is then used for spectral clustering. This works well for assortative graphs, but performs poorly otherwise. Covariate-assisted spectral clustering, which uses the square of the regularized graph Laplacian, is presented as a more general method that performs well for assortative and non-assortative graphs. A tuning parameter is employed by both methods to adjust the relative weight of the covariates and the graph; in §2.3 we propose a way to choose this tuning parameter. Research on dynamic networks using latent space models has yielded an analogous form for updating latent coordinates based on a distance matrix and the latent coordinates from the previous time step (Sarkar & Moore, 2006). A similar framework can also be used to cluster multiple graphs (Eynard et al., 2015).

Variants of our methods previously introduced were derived by first considering the problem of minimizing the weighted sum of the k -means and graph cut objective functions and then solving a spectral relaxation of the original problem. Wang et al. (2009) decided against using an additive method similar to covariate-assisted spectral clustering because setting the method's tuning parameter is a nonconvex problem. They chose to investigate a method that uses the product of the generalized inverse of the graph Laplacian and the covariate matrix instead. Shiga et al. (2007) recognized the advantage of having a tuning parameter to balance the contribution of the graph and the covariates, but did not use the stochastic blockmodel to study their method. The full utility and flexibility of these types of approaches have not yet been presented, and neither paper derives any statistical results about the methods' performance. Furthermore, they do not consider the performance of these methods on non-assortative graphs. In contrast, we were initially motivated to develop covariate-assisted spectral clustering by its interpretation and propensity for theoretical analysis.

Very few of the clustering methods that employ both node covariates and the graph offer any theoretical results, and, to our knowledge, this paper gives the first statistical guarantee for these types of approaches. We define the node-contextualized stochastic blockmodel, which combines the stochastic blockmodel with a block mixture model for node covariates. Under this model, a bound on the misclustering rate of covariate-assisted spectral clustering is established in §3.2. The behaviour of the bound is studied for a fixed and an increasing number of covariates as a function of the number of nodes, and conditions for perfect clustering are derived. A general lower bound is also derived, demonstrating the conditions under which an algorithm using both the node covariates and the graph can give more accurate clusters than any algorithm using only the node covariates or the graph.

For comparison, an alternative method based on an adaptation of classical canonical correlation analysis is introduced (Hotelling, 1936), which uses the product of the regularized graph Laplacian and the covariate matrix as the input to the spectral clustering algorithm. Simulations indicate that canonical correlation performs worse than covariate-assisted spectral clustering under the node-contextualized stochastic blockmodel with Bernoulli covariates. However, canonical correlation analysis clustering is computationally faster than our clustering method and requires no tuning. In contrast, covariate-assisted spectral clustering depends on a single tuning parameter, which

interpolates between spectral clustering with only the graph and only the covariates. This parameter can be set without prior knowledge by using an objective function such as the within-cluster sum of squares. Some results for determining what range of tuning parameter values should be considered are provided in the description of the optimization procedure in §2.3. Alternatively, the tuning parameter can be set using prior knowledge or to ensure that the clusters achieve some desired quality, such as spatial cohesion. As an illustrative example, in §5 we study diffusion magnetic resonance imaging-derived brain graphs using two different sets of node covariates. The first analysis uses spatial location. This produces clusters that are more spatially coherent than those obtained using regularized spectral clustering alone, making them easier to interpret. The second analysis uses neurological region membership, which yields partitions that closely align with neurological regions, while allowing for patient-wise variability based on brain graph connectivity.

2. METHODOLOGY

2.1. Notation

Let $G(E, V)$ be a graph, where V is the set of vertices or nodes and E is the set of edges, which represent relationships between the nodes. Let N be the number of nodes. Index the nodes in V by $\{1, \dots, N\}$; then E contains a pair (i, j) if there is an edge between nodes i and j . A graph's edge set can be represented as the adjacency matrix $A \in \{0, 1\}^{N \times N}$, where $A_{ij} = A_{ji} = 1$ if $(i, j) \in E$ and $A_{ij} = A_{ji} = 0$ otherwise. We restrict ourselves to studying undirected and unweighted graphs, although with small modifications most of our results also apply to directed and weighted graphs.

Define the regularized graph Laplacian as

$$L_\tau = D_\tau^{-1/2} A D_\tau^{-1/2},$$

where $D_\tau = D + \tau I$ and D is a diagonal matrix with $D_{ii} = \sum_{j=1}^N A_{ij}$. The regularization parameter τ is treated as a constant, and is included to improve spectral clustering performance on sparse graphs (Chaudhuri et al., 2012). Throughout, we shall set $\tau = N^{-1} \sum_{i=1}^N D_{ii}$, i.e., the average node degree (Qin & Rohe, 2013).

For the graph $G(E, V)$, let each node in the set V have an associated bounded covariate vector $X_i \in [-J, J]^R$, and let $X \in [-J, J]^{N \times R}$ be the covariate matrix where each row corresponds to a node covariate vector. Let $\|\cdot\|$ denote the spectral norm and $\|\cdot\|_F$ the Frobenius norm. Let $I(\cdot)$ denote the indicator function. For sequences $\{a_N\}$ and $\{b_N\}$, $a_N = O(b_N)$ if and only if $b_N = O(a_N)$ and $a_N = O(b_N)$.

2.2. Spectral clustering for a graph with node covariates

The spectral clustering algorithm has been employed to cluster graph nodes using various functions of the adjacency matrix. For instance, applying the algorithm to L_τ corresponds to regularized spectral clustering, where the value of the regularization parameter is set prior to running the algorithm. All of the methods we consider will employ this algorithm, but will use a different input matrix such as L_τ , \tilde{L} or L^{CCA} as defined later.

Algorithm 1. Spectral clustering.

Given input matrix W and number of clusters K :

Step 1. Find eigenvectors $U_1, \dots, U_K \in \mathbb{R}^N$ corresponding to the K largest eigenvalues of W .

Step 2. Use the eigenvectors as columns to form the matrix $U = (U_1, \dots, U_K) \in \mathbb{R}^{N \times K}$.

Step 3. Form the matrix U^* by normalizing each of U 's rows to have unit length.

Step 4. Run k -means clustering with K clusters, treating each row of U^* as a point in \mathbb{R}^K .

Step 5. If the i th row of U^* falls in the k th cluster, assign node i to cluster k .

Step 4 of the spectral clustering algorithm uses k -means clustering, which is sensitive to initialization. In order to reduce this sensitivity, we use multiple random initializations. To take advantage of available graph and node covariate data in graph clustering, it is necessary to employ methods that incorporate both of these data types. As discussed in § 1, spectral clustering has many advantages over other graph clustering methods. Hence, we propose three approaches that use the spectral clustering framework and utilize both the graph structure and the node covariates.

Assortative covariate-assisted spectral clustering uses the leading eigenvectors of

$$\tilde{L}(\alpha) = L_\tau + \alpha XX^\top,$$

where $\alpha \in [0, \infty)$ is a tuning parameter. When using $\{0, 1\}$ -Bernoulli covariates, the covariate term can be interpreted as adding to each element (i, j) a value proportional to the number of covariates equal to 1 for both i and j . In practice, the covariate matrix X should be parameterized as in linear regression; specifically, categorical covariates should be re-expressed with dummy variables. For continuous covariates, it can be beneficial to centre and scale the columns of X before performing the analysis. The scaling helps satisfy condition (i) in Lemma 1, which ensures that the top K eigenvectors of $\tilde{L}(\alpha)$ contain block information. As demonstrated in the simulations in § 4, this method is robust and has good performance for assortative graphs, but does not perform well for non-assortative graphs.

Covariate-assisted spectral clustering uses the leading eigenvectors of

$$\tilde{L}(\alpha) = L_\tau L_\tau + \alpha XX^\top.$$

This approach performs well for non-assortative graphs and nearly as well as our assortative clustering method for assortative graphs. When there is little chance of confusion, \tilde{L} will be used for notational convenience.

To run covariate-assisted spectral clustering on the large graphs, such as the brain graphs in § 5, the top K eigenvectors of \tilde{L} are computed using the implicitly restarted Lanczos bidiagonalization algorithm (Baglama & Reichel, 2006). At each iteration, the algorithm only needs to compute the product $\tilde{L}v$, where v is an arbitrary vector. For computational efficiency, the product is calculated as $L_\tau(L_\tau v) + \alpha X(X^\top v)$. This takes advantage of the sparsity of L_τ and the low-rank structure of XX^\top . Ignoring log terms and any special structure in X , it takes $O\{(|E| + NR)K\}$ operations to compute the required top K eigenvectors of \tilde{L} , where R is the number of columns in X . The graph clusters are obtained by iteratively employing the spectral clustering algorithm on $\tilde{L}(\alpha)$ while varying the tuning parameter α until an optimal value is obtained. The details of this procedure are described in the next section.

As an alternative, we propose a modification of classical canonical correlation analysis (Hotelling, 1936) whose similarity matrix is the product of the regularized graph Laplacian and the covariate matrix,

$$L^{\text{CCA}} = L_\tau X.$$

The spectral clustering algorithm is employed on L^{CCA} to obtain node clusters when the number of covariates, R , is greater than or equal to the number of clusters, K . This approach inherently provides a dimensionality reduction in the common case where the number of covariates is much

less than the number of nodes. If $R \ll N^{-1} \sum_i D_{ii}$, then spectral clustering with L^{CCA} has a faster running time than covariate-assisted spectral clustering.

2.3. Setting the tuning parameter

In order to perform spectral clustering with $\tilde{L}(\alpha)$, it is necessary to determine a specific value for the tuning parameter, α . The tuning procedure presented here presumes that both the graph and the covariates contain some block information, as demonstrated by the simulations in § 4. In practice, an initial test can be used to determine if the graph and the covariates contain common block information, and such a test will be presented in future work. The tuning parameter should be chosen to achieve a balance between L_τ and X such that the information in both is captured in the leading eigenspace of \tilde{L} . For large values of α , the leading eigenspace of \tilde{L} is approximately the leading eigenspace of XX^T . For small values of α , the leading eigenspace of \tilde{L} is approximately the leading eigenspace of L_τ . A good initial choice of α is the value which makes the leading eigenvalues of $L_\tau L_\tau$ and αXX^T equal, namely $\alpha_0 = \lambda_1(L_\tau L_\tau) / \lambda_1(XX^T)$.

There is a finite range of α for which the leading eigenspace of $\tilde{L}(\alpha)$ is not a continuous function of α ; outside this range, the leading eigenspace is always continuous in α . In simulations, the clustering results are exceedingly stable in the continuous range of α . Hence, only the values of α inside a finite interval need to be considered. This section gives an interval $\alpha \in [\alpha_{\min}, \alpha_{\max}]$ that is computed with only the eigenvalues of $L_\tau L_\tau$ and XX^T . Within this interval, α is chosen to minimize an objective function. Empirical results demonstrating these properties are given in the Supplementary Material.

Let $\lambda_i(M)$ be the i th eigenvalue of matrix M . To find the initial range $[\alpha_{\min}, \alpha_{\max}]$, define a static vector $v \in \mathbb{R}^N$ as a vector that satisfies conditions (1) and (2) below. For $\epsilon \geq 0$,

$$v^T L_\tau L_\tau v \geq \lambda_K(L_\tau L_\tau), \quad v^T XX^T v \leq \epsilon, \quad (1)$$

$$v^T XX^T v \geq \lambda_K(XX^T), \quad v^T L_\tau L_\tau v \leq \epsilon. \quad (2)$$

The eigenspaces of XX^T and $L_\tau L_\tau$ have little overlap along a static vector, v ; perhaps there is a cluster in the graph that does not appear in the covariates, or vice versa. These static vectors produce discontinuities in the leading eigenspace of $\tilde{L}(\alpha)$.

For example, if v_* is an eigenvector of $L_\tau L_\tau$ and a static vector of type (1), then as α changes, it will remain a slightly perturbed eigenvector of $\tilde{L}(\alpha)$. When $v_*^T \tilde{L}(\alpha_*) v_*$ is close to $\lambda_K\{\tilde{L}(\alpha_*)\}$, then, in some neighbourhood of α_* , the slightly perturbed version of v_* will transition into the leading eigenspace of \tilde{L} . This transition corresponds to a discontinuity in the leading eigenspace.

As shown in the Supplementary Material, the concept of static vectors with $\epsilon = 0$ can be used to find a limited range of α for possible discontinuities. The range of α values for which discontinuities can occur is $[\alpha_{\min}, \alpha_{\max}]$, where

$$\alpha_{\min} = \frac{\lambda_K(L_\tau L_\tau) - \lambda_{K+1}(L_\tau L_\tau)}{\lambda_1(XX^T)},$$

$$\alpha_{\max} = \frac{\lambda_1(L_\tau L_\tau)}{\lambda_R(XX^T)I(R \leq K) + \{\lambda_K(XX^T) - \lambda_{K+1}(XX^T)\}I(R > K)}.$$

The tuning parameter $\alpha \in [\alpha_{\min}, \alpha_{\max}]$ is chosen to be the value which minimizes the k -means objective function, the within-cluster sum of squares,

$$\Phi(\alpha) = \sum_{i=1}^K \sum_{u_j \in F_i} \|u_j(\alpha) - C_i(\alpha)\|^2,$$

where u_j is the j th row of U , C_i is the centroid of the i th cluster from k -means clustering, and F_i is the set of points in the i th cluster. Hence, the tuning parameter is $\alpha = \arg \min_{\alpha \in [\alpha_{\min}, \alpha_{\max}]} \Phi(\alpha)$.

3. THEORY

3.1. Node-contextualized stochastic blockmodel

To illustrate what covariate-assisted spectral clustering estimates, this section proposes a statistical model for a network with node covariates and shows that covariate-assisted spectral clustering is a weakly consistent estimator of certain parameters in the proposed model. To derive statistical guarantees for covariate-assisted spectral clustering, we assume a joint mixture model for the graph and the covariates. Under this model, each node belongs to one of K blocks and each edge in the graph corresponds to an independent Bernoulli random variable. The probability of an edge between any two nodes depends only on the block membership of those nodes (Holland et al., 1983). In addition, each node is associated with R independent covariates with bounded support, where expectation depends only on the block membership and R can grow with the number of nodes.

DEFINITION 2 (Node-contextualized stochastic blockmodel). *Consider a set of nodes $\{1, \dots, N\}$. Let $Z \in \{0, 1\}^{N \times K}$ assign each of the N nodes to one of the K blocks, where $Z_{ij} = 1$ if node i belongs to block j . Let $B \in [0, 1]^{K \times K}$ be of full rank and symmetric, where B_{ij} is the probability of an edge between nodes in blocks i and j . Conditional on Z , the elements of the adjacency matrix are independent Bernoulli random variables. The population adjacency matrix $A = E(A | Z)$ fully identifies the distribution of A and $\mathcal{A} = ZBZ^T$.*

Let $X \in [-J, J]^{N \times R}$ be the covariate matrix and $M \in [-J, J]^{K \times R}$ the covariate expectation matrix, where $M_{i,j}$ is the expectation of the j th covariate of a node in the i th block. Conditional on Z , the elements of X are independent and the population covariate matrix is

$$\mathcal{X} = E(X | Z) = ZM. \quad (3)$$

Under the node-contextualized stochastic blockmodel, covariate-assisted spectral clustering seeks to estimate the block membership matrix Z . In the next section, we show that this estimate is consistent. If B is assumed to be positive definite, the same results hold for assortative covariate-assisted spectral clustering up to a constant factor. These results motivate the definition of an assortative graph in the context of the node-contextualized stochastic blockmodel.

DEFINITION 3 (Assortative graph). *A graph generated under the node-contextualized stochastic blockmodel is said to be assortative if the block probability matrix B corresponding to the graph is positive definite. Otherwise, it is said to be non-assortative.*

Many common networks are assortative, such as friendship networks or brain graphs. Dating networks are one example of a non-assortative network. Most relationships in a dating network are heterosexual, comprised of one male and one female. In a stochastic blockmodel, where the blocks are constructed by gender, B will have small diagonal elements and large off-diagonal elements, producing more relationships between genders than within genders. Such a matrix is not positive definite. More generally, non-assortative stochastic blockmodels will tend to generate more edges between blocks and fewer edges within blocks. These non-assortative blocks appear in the spectrum of L_τ as large negative eigenvalues. By squaring the matrix L_τ , the eigenvalues become large and positive, matching the positive eigenvalues in XX^T .

3.2. *Statistical consistency under the node-contextualized stochastic blockmodel*

The proof of consistency for covariate-assisted spectral clustering under the node-contextualized stochastic blockmodel requires three results. Lemma 1 expresses the eigendecomposition of the population version of the covariate-assisted Laplacian,

$$\tilde{\mathcal{L}}(\alpha) = (\mathcal{D} + \tau I)^{-1/2} \mathcal{A}(\mathcal{D} + \tau I)^{-1} \mathcal{A}(\mathcal{D} + \tau I)^{-1/2} + \alpha E(\mathcal{X}\mathcal{X}^\top),$$

in terms of Z . Theorem 1 bounds the spectral norm of the difference between $\tilde{\mathcal{L}}$ and $\tilde{\mathcal{L}}$. Then, the Davis–Kahan theorem (Davis & Kahan, 1970) bounds the difference between the sample and population eigenvectors in Frobenius norm. Finally, Theorem 3 combines these results to establish a bound on the misclustering rate of covariate-assisted spectral clustering. The argument largely follows Qin & Rohe (2013). The results provided here do not include the effects of Step 3 in Algorithm 1. The proofs are in the Supplementary Material.

LEMMA 1. *Under the node-contextualized stochastic blockmodel, let $\mathcal{D}_B = \text{diag}(BZ^\top \mathbf{1}_n + \tau)$, $\tilde{P} = Z^\top Z$, and $\tilde{B} = \mathcal{D}_B^{-1/2} BZ^\top (\mathcal{D} + \tau I)^{-1} Z B \mathcal{D}_B^{-1/2} + \alpha M M^\top$. Let $\varkappa = \max_l |c_l - \bar{c}|$, where $c_l = \sum_i \text{var}(X_{il} | Z_i = l)$ and $\bar{c} = \sum_l c_l / K$. Define $\mathcal{U} \in \mathbb{R}^{N \times K}$ with columns containing the top K eigenvectors of $\tilde{\mathcal{L}}$. Assume (i) $\lambda_K(\tilde{B}\tilde{P}) > 2\alpha\varkappa$; then there exists an orthogonal matrix $V \in \mathbb{R}^{K \times K}$ such that $\mathcal{U} = Z(ZZ^\top)^{-1/2} V$. Furthermore, $Z_i(ZZ^\top)^{-1/2} V = Z_j(ZZ^\top)^{-1/2} V$ if and only if $Z_i = Z_j$, where Z_i is the i th row of the block membership matrix.*

Under assumption (i) the rows of the population eigenvectors are equal if and only if the corresponding nodes belong to the same block. This assumption requires the population eigengap to be greater than the maximum of the absolute difference between the sum of covariate variances within a block and the mean of the sums across all blocks. If all the covariates have equal variance in all blocks, the assumption is trivially true. Since $\mathcal{X}\mathcal{X}^\top$ is effectively being used as a measure of similarity between nodes, if the covariate variances across blocks are unequal, the difference in scale makes the blocks more difficult to distinguish. This is evidenced by a reduction in the eigengap proportional to this difference. In practice, this condition is not restrictive since X can be centred and normalized. To derive a bound on the misclustering rate, we will need a bound on the difference between the population eigenvectors and the sample eigenvectors. In order to establish this bound, the following theorem bounds the spectral norm of the difference between $\tilde{\mathcal{L}}$ and $\tilde{\mathcal{L}}$.

THEOREM 1. *Let $d = \min \mathcal{D}_{ii}$, $\mathcal{X}_{ik}^{(p)} = E(\mathcal{X}_{ik}^p)$, $\varpi = 8\alpha^2 \sum_k \{ \sum_i \mathcal{X}_{ik}^{(2)} \sum_l (\mathcal{X}_{lk}^{(2)} - \mathcal{X}_{lk}^2) + \mathcal{X}_{ik}^{(4)} \}$, $\delta = 12(d + \tau)^{-1/2} + \varpi^{1/2}$, and $S = 3\alpha N J^2$. For any $\epsilon > 0$, if (ii) $d + \tau > 3 \log(8N/\epsilon)$ and (iii) $\varpi/S^2 > 3 \log(8N/\epsilon)$, then with probability at least $1 - \epsilon$,*

$$\|\tilde{\mathcal{L}} - \tilde{\mathcal{L}}\| \leq \delta \{3 \log(8N/\epsilon)\}^{1/2}.$$

Consider a node-contextualized stochastic blockmodel with two blocks, within-block probabilities p , and between-block probabilities q . Condition (ii) holds when $p + q > \Theta(\log(N)/N)$ and condition (iii) holds when $R \geq \Theta(\log N)$. Hence, condition (ii) restricts the sparsity of the graph, while condition (iii) requires that the number of covariates grow with the number of nodes. Now we use Theorem 1 and the Davis–Kahan theorem to bound the difference between the sample and population eigenvectors.

THEOREM 2. *Let λ_K be the K th largest eigenvalue of $\tilde{\mathcal{L}}$ and let \mathcal{O} be a rotation matrix. Let the columns of U and \mathcal{U} contain the top K eigenvectors of $\tilde{\mathcal{L}}$ and $\tilde{\mathcal{L}}$, respectively. Under assumptions*

(i) in Lemma 1, (ii) and (iii) in Theorem 1, and (iv) $\delta\{3 \log(8N/\epsilon)\}^{1/2} \leq \lambda_K/2$, with probability at least $1 - \epsilon$,

$$\|U - \mathcal{U}\mathcal{O}\|_F \leq \frac{8\delta\{3K \log(8N/\epsilon)\}^{1/2}}{\lambda_K}.$$

The next theorem bounds the proportion of misclustered nodes. In order to define misclustering, recall that the spectral clustering algorithm uses k -means clustering to cluster the rows of U . Let C_i and \mathcal{C}_i be the cluster centroids of the i th node generated using k -means clustering on U and \mathcal{U} , respectively. A node i is correctly clustered if C_i is closer to \mathcal{C}_i than \mathcal{C}_j for all j such that $Z_j \neq Z_i$. In order to avoid identifiability problems and since clustering only requires the estimation of the correct subspace, the formal definition is augmented with a rotation matrix \mathcal{O} . The following definition formalizes this intuition.

DEFINITION 4. Let \mathcal{O} be a rotation matrix that minimizes $\|U\mathcal{O}^T - \mathcal{U}\|_F$. Define the set of misclustered nodes as

$$\mathcal{M} = \{i : \text{there exists } j \neq i \text{ such that } \|C_i\mathcal{O}^T - C_i\|_2 > \|C_i\mathcal{O}^T - C_j\|_2\}.$$

Using the definition of misclustering and the result from Theorem 2, the next theorem bounds the misclustering rate, $|\mathcal{M}|/N$.

THEOREM 3. Let $P = \max_i (Z^T Z)_{ii}$ denote the size of the largest block. Under assumptions (i) in Lemma 1, (ii) and (iii) in Theorem 1, and (iv) in Theorem 2, with probability at least $1 - \epsilon$, the misclustering rate satisfies

$$\frac{|\mathcal{M}|}{N} \leq \frac{c_0 K P \delta^2 \log(8N/\epsilon)}{N \lambda_K^2}.$$

The asymptotics of the misclustering rate depend on the number of covariates and the sparsity of the graph. This is demonstrated by Corollary 1, which provides insight into how the number of covariates and graph sparsity affect the misclustering rate and the choice of tuning parameter.

COROLLARY 1. Assume $B_{i,i} = p$ and $B_{i,j} = q$ for all $i \neq j$; in addition, let $M_{i,i} = m_1$, $M_{i,j} = m_2$ for all $i \neq j$, and $R > 1$. For computational convenience, assume that each block has the same number of nodes N/K and R is a multiple of K , where K is fixed. Let $R = \Theta\{(\log N)^{a+1}\}$, $d + \tau = \Theta\{(\log N)^{b+1}\}$, and $\alpha = \Theta\{N^{-1}(\log N)^{-1-c}\}$, where $a, b, c \geq 0$. Then, the misclustering bound from Theorem 3 becomes

$$\frac{|\mathcal{M}|}{N} \leq c_2 \frac{(\log N)^{a-2c} + (\log N)^{(a-b)/2-c} + (\log N)^{-b}}{(\log N)^{2(a-c)} + (\log N)^{a-c} + \Theta(1)}.$$

If $a \leq b$, then the minimum misclustering rate is $O\{(\log N)^{-b}\}$ when $c = (a + b)/2$. If $a > b$, then the minimum misclustering rate is $O\{(\log N)^{-a}\}$ when $c = 0$.

These results demonstrate that the tuning parameter α is determined by the balance between the number of covariates and the sparsity of the graph. A nonzero optimal α value signifies that including covariates improves the misclustering bound, although it might not improve the asymptotics of the bound. Furthermore, the asymptotic misclustering rate is determined by the

asymptotic behaviour of the number of covariates or the mean number of edges, whichever is greater as determined by a and b , respectively. For example, if we allow the number of covariates to grow with the number of nodes such that $a = 0$ or $R = \Theta(\log N)$ and let the mean number of edges increase such that $b = 0$ or $d + \tau = \Theta(\log N)$, then the covariates and the graph contribute equally to the asymptotic misclustering rate.

Remark 1. It is instructive to compare the value of α suggested by the results in Corollary 1 with the possible values of α based on the optimization procedure in § 2.3. Computing α_{\min} and α_{\max} with $\tilde{\mathcal{L}}$ instead of \tilde{L} , for convenience, gives $\alpha_{\min} = \Theta\{(NR)^{-1}\}$ and $\alpha_{\max} = \Theta\{(NR)^{-1}\}$. Therefore, the optimization procedure will yield $\alpha = \Theta\{(NR)^{-1}\} = \Theta\{(N)^{-1}(\log N)^{-a-1}\}$. This agrees with the results of Corollary 1 when the mean node degree and the number of covariates grow at the same rate with respect to the number of nodes or $a = b$.

COROLLARY 2. *Based on Theorem 3, perfect clustering requires $\delta\{c_0KP \log(8N/\epsilon)\}^{1/2} < \lambda_K$. Under the simplifying assumptions given in Corollary 1, perfect clustering is achieved when the number of covariates is $R \geq \Theta(N \log N)$.*

3.3. General lower bound

The next theorem gives a lower bound for clustering a graph with node covariates. This bound uses Fano’s inequality and is similar to that shown in Chaudhuri et al. (2012) for a graph without node attributes. We restrict ourselves to a node-contextualized stochastic blockmodel with $K = 2$ blocks, but allow an arbitrary number of covariates R .

THEOREM 4. *Consider the node-contextualized stochastic blockmodel with $K = 2$ blocks and B such that $B_{1,1} \geq B_{2,2} \geq B_{1,2}$. Let the Kullback–Leibler divergence of the covariates be $\Gamma = \sum_{i=1}^R KL(\gamma_i, \gamma'_i)$, where γ_i and γ'_i are the distributions of the i th covariate under opposite block assignments, and $\Delta = B_{1,1} - B_{1,2}$. For a fixed $B_{1,1}$ and $N \geq 8$, in order to correctly recover the block assignments with probability at least $1 - \epsilon$, Δ must satisfy*

$$\Delta \geq \frac{B_{1,1}(1 - B_{1,1})}{\left[\frac{2}{N} \left\{ \frac{\log 2}{2}(1 - \epsilon) - \Gamma - \frac{\log 2}{N} \right\} \right]^{-1/2} + (1 - B_{1,1})}.$$

Remark 2. If

$$B_{1,1} - B_{1,2} < \frac{B_{1,1}(1 - B_{1,1})N^{-1/2}}{\left\{ (1 - \epsilon - \frac{2}{N}) \log 2 \right\}^{-1/2} + (1 - B_{1,1})}, \tag{4}$$

$$\Gamma < \left(\frac{1}{2} - \frac{\epsilon}{2} - \frac{1}{N} \right) \log 2, \tag{5}$$

then only an algorithm that uses both the graph and node covariates can yield correct blocks with high probability. Condition (4) specifies when the graph is insufficient and condition (5) specifies when the covariates are insufficient to individually recover the block membership with high probability.

Remark 3. The upper bound for covariate-assisted spectral clustering in Theorem 3 can be compared with the general lower bound. Simplifying the general lower bound gives the condition $\Delta \geq \Theta(N^{-1/2})$ for perfect clustering with probability $1 - \epsilon$. This is the same condition as for

regularized spectral clustering. According to Theorem 3, for this method to achieve perfect clustering with probability $1 - \epsilon$ requires $\delta\{c_0 K p \log(8N/\epsilon)\}^{1/2} < \lambda_K$. As highlighted in Corollary 2, this condition cannot be satisfied for a fixed R , so it cannot be shown that covariate-assisted spectral clustering achieves perfect clustering for a fixed number of covariates. This is consistent with similar results for regularized spectral clustering.

4. SIMULATIONS

4.1. Varying graph or covariate signal

In these simulations, consider a node-contextualized stochastic blockmodel with $K = 3$ blocks and $R = 3$ node Bernoulli covariates. Define the block probabilities for the assortative graph, the non-assortative graph, and the covariates as

$$B = \begin{pmatrix} p & q & q \\ q & p & q \\ q & q & p \end{pmatrix}, \quad B' = \begin{pmatrix} q & p & p \\ p & q & p \\ p & p & q \end{pmatrix}, \quad M = \begin{pmatrix} m_1 & m_2 & m_2 \\ m_2 & m_1 & m_2 \\ m_2 & m_2 & m_1 \end{pmatrix}, \quad (6)$$

where $p > q$ and $m_1 > m_2$. This implies that for the assortative graph the probability of an edge within a block is p , which is greater than q , the probability of an edge between two blocks. The opposite is true for the non-assortative graph. In the k th block, the probability of the k th covariate being 1 is m_1 and the probability of the other covariates being 1 is m_2 .

These simulations compare five methods. The first three are canonical correlation analysis clustering, covariate-assisted spectral clustering, and assortative covariate-assisted spectral clustering, which utilize node edges and node covariates to cluster the graph. The other two methods utilize either the node edges or the node covariates. For the node edges, regularized spectral clustering is used; for the node covariates, spectral clustering on the covariate matrix is used.

The first set of simulations investigates the effect of varying the block signal in the graph on the misclustering rate. This is done by varying the difference in the within- and between-block probabilities, $p - q$. The simulations are conducted for the assortative and non-assortative graphs, using B and B' in (6), shown in Fig. 1(a) and (b), respectively. In the assortative case, our assortative clustering method performs better than any of the other methods. Covariate-assisted spectral clustering performs slightly worse than the assortative variant, but still outperforms the other methods. In the non-assortative case, our clustering method has the best performance, while the assortative version always does worse than using only the covariates or the graph.

The second set of simulations investigates the effect of varying the block signal of the covariates on the misclustering rate by changing the difference between the block-specific covariate probabilities, $m_1 - m_2$. As shown in Fig. 1(c), assortative covariate-assisted spectral clustering tends to have a better misclustering rate than the other methods. Only when the difference in the covariate block probabilities is very small and X effectively becomes a noise term does regularized spectral clustering outperform our assortative clustering method. For the non-assortative case shown in Fig. 1(d), assortative covariate-assisted spectral clustering performs poorly, while covariate-assisted spectral clustering is able to outperform all other methods for a sufficiently large difference in the covariate block probabilities. This is expected since the covariates in the assortative variant effectively increase the edge weights within a block, which will smooth out the block structure specified by B' .

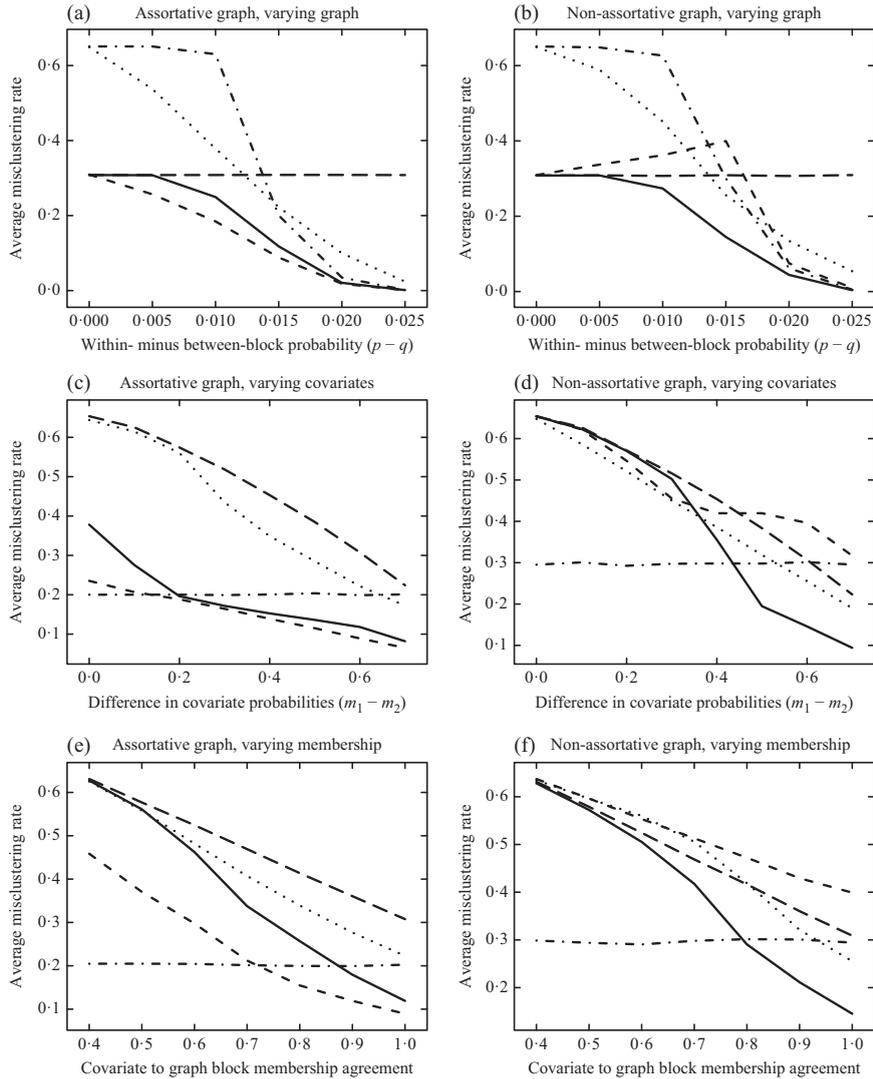


Fig. 1. Average misclustering rate of five clustering methods: covariate-assisted spectral clustering (solid), assortative covariate-assisted spectral clustering (dash), canonical correlation analysis clustering (dot), regularized spectral clustering (dot-dash), and spectral clustering on the covariate matrix (long dash). The fixed parameters are $N = 1500$, $p = 0.03$, $q = 0.015$, $m_1 = 0.8$, and $m_2 = 0.2$.

4.2. Model misspecification

The final simulation considers the case where the block membership in the covariates is not necessarily the same as the block membership in the graph. The node Bernoulli covariates no longer satisfy (3) in Definition 2, but $\mathcal{X} = YM$, where $Y \in \{0, 1\}^{N \times K}$ is a block membership matrix that differs from Z . As such, the underlying clusters in the graph do not align with the clusters in the covariates. This simulation varies the proportion of block assignments in Y which agree with the block assignments in Z to investigate the robustness of the methods with respect to this form of model misspecification. The results in Fig. 1(e) show that assortative covariate-assisted spectral clustering is robust with respect to covariate block membership model misspecification for the assortative graph. The misclustering rate shown is computed relative to the block membership of the graph. For this case, our assortative clustering method is able to

achieve a lower misclustering rate than regularized spectral clustering when the proportion of agreement between the block membership of the graph and the covariates is greater than 0.7. Since a three-block model is used, the lowest proportion of agreement possible is one third due to identifiability. For the non-assortative graph, Fig. 1(f), covariate-assisted spectral clustering requires a higher level of agreement at 0.8.

5. CLUSTERING DIFFUSION MRI CONNECTOME GRAPHS

Assortative covariate-assisted spectral clustering was applied to brain graphs recovered from diffusion magnetic resonance imaging (Craddock et al., 2013). Each node in a brain graph corresponds to a voxel in the brain. The edges between nodes are weighted by the number of estimated fibres that pass through both voxels. The centre of a voxel is treated as the spatial location of the corresponding node. These spatial locations were centred and used as the first set of covariates in the analysis. The dataset used in this analysis contains 42 brain graphs obtained from 21 different individuals. Only the largest connected components of the brain graphs were used, ranging in size from 707 000 to 935 000 nodes, with a mean density of 744 edges per node. In addition, the brain graphs contain brain atlas labels corresponding to 70 different neurological brain regions, which were treated as a second set of covariates.

Whereas the simulations attempted to demonstrate the effectiveness of our clustering method in utilizing node covariates to help discover the underlying block structure of the graph, this analysis focuses on the ability of our clustering method to discover highly connected clusters with relatively homogeneous covariates. The node covariates contextualize the brain clusters and improve their interpretability. Like other clustering methods, covariate-assisted spectral clustering is mainly an exploratory tool which may or may not provide answers directly but can often provide insight into relationships within the data. In this example, it is used to examine the relationships between brain graph connectivity, spatial location, and brain atlas labels. The utility of covariate-assisted spectral clustering was explored by partitioning the brain graphs into 100 clusters. The brain graphs in this dataset are assortative, so our assortative clustering method was used in this analysis. Since the brain graphs have heterogeneous node degrees, the rows of the eigenvector matrix were normalized when applying the spectral clustering algorithm to improve the clustering results (Qin & Rohe, 2013). Figure 2 shows a section of a sample brain graph with nodes plotted at their corresponding spatial locations and coloured by cluster membership. For reference, the neurological brain atlas clusters with 70 different regions and an additional category for unlabelled nodes are also plotted. The brain graphs were clustered using three different approaches: regularized spectral clustering, and assortative covariate-assisted spectral clustering with spatial location and with brain atlas membership. The tuning parameter α was set using the procedure in § 2.3, and the values were $\alpha = 0.0004$ with spatial location covariates and $\alpha = 0.0708$ with brain atlas membership covariates.

As shown in Fig. 2, regularized spectral clustering yielded spatially diffuse clusters of densely connected nodes. By adding spatial location using covariate-assisted spectral clustering, we obtained densely connected and spatially coherent clusters. Regularized spectral clustering had two clusters of about 80 000 nodes and four clusters with fewer than 1000 nodes, while the largest cluster from our clustering method had fewer than 50 000 nodes and no clusters had fewer than 1000 nodes. Both greater spatial coherence and increased uniformity in cluster size demonstrated by covariate-assisted spectral clustering are important qualities for interpreting the partition. In addition, the clusters have a greater similarity with the brain atlas labels, though this similarity is still not very substantial. This suggests that brain graph connectivity is governed by more than just the neurological regions in the brain atlas.

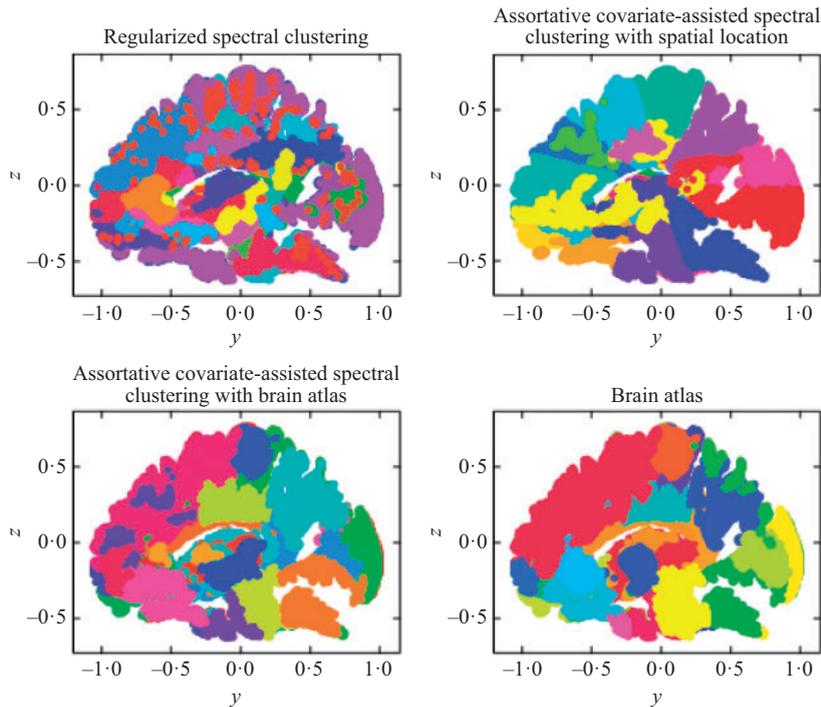


Fig. 2. Brain graph cross-sections with nodes plotted spatially and coloured by cluster membership for three different clustering methods and a brain atlas.

The relation between the brain atlas and the brain graph was studied further by treating brain atlas membership as the node covariates. This allowed the discovery of highly connected regions with relatively homogeneous graph atlas labels. As shown in Fig. 2, relative to the brain atlas, some of the clusters are broken up, a few are joined together, and others overlap with multiple brain atlas regions, but the high similarity is clearly visible. Importantly, this approach gives us clusters that are highly aligned with known neurological regions while allowing for individual variability of the partitions based on brain graph connectivity. The adjusted Rand index was used to quantify the similarity of the partitions of a brain graph specified by the different clustering methods and the brain atlas in Table 1. The alignment with the partitions based only on spatial location and either covariate-assisted spectral clustering with spatial location or the brain atlas is greater than between the two methods. This indicates that both the clusters from our method and the brain atlas are spatially coherent yet not highly overlapping.

Brain graph connectivity appears to be giving the clusters that use spatial location a different configuration from the brain atlas, as seen in Fig. 2. As expected, covariate-assisted spectral clustering with brain atlas membership has the highest adjusted Rand index partition similarity with the brain atlas but low similarity with the regularized spectral clustering partitions. If a more balanced partition alignment is desired, the tuning parameter can be adjusted accordingly.

The relationship between all 42 brain graphs was analysed by using the adjusted Rand index to compare partitions between them, as shown in Fig. 3. To conduct the comparison, the nodes of each brain graph were matched by spatial location, and any nonmatching nodes were ignored. Both regularized spectral clustering and covariate-assisted spectral clustering with spatial location distinguish clearly between individuals based on their brain graph partitions, but the latter gave partitions which are more homogeneous both within and between individuals. This increased

Table 1. *The adjusted Rand index between different partitions*

	ACASC-X	Brain atlas	ACASC-BA	SC-X
RSC	0.095	0.082	0.085	0.092
ACASC-X	—	0.169	0.189	0.278
Brain atlas	—	—	0.838	0.226
ACASC-BA	—	—	—	0.227

RSC, regularized spectral clustering; ACASC-X, assortative covariate-assisted spectral clustering with spatial location; ACASC-BA, assortative covariate-assisted spectral clustering with brain atlas membership; SC-X, spectral clustering using spatial location.

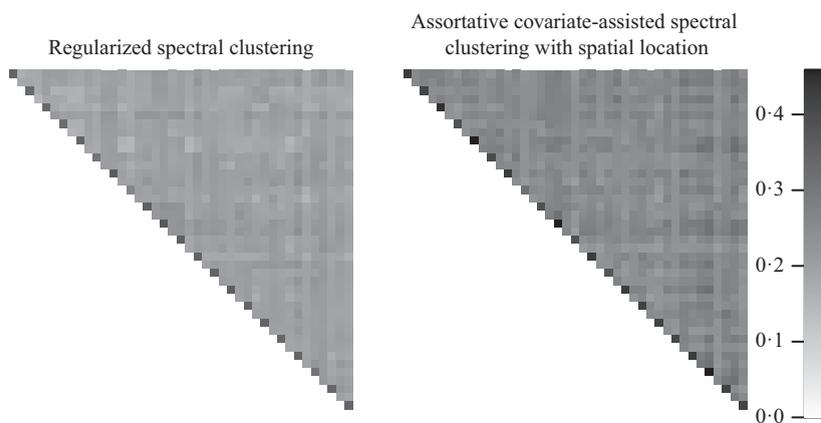


Fig. 3. Heat maps of the adjusted Rand index comparing the partitions of 42 brain graphs. The separate heat maps are based on partitions generated with different methods. Each row and column corresponds to a brain scan and adjacent rows or columns correspond to two scans of the same individual.

partition consistency is favourable since a high degree of variation in the clusters between individuals would make them more difficult to interpret.

6. DISCUSSION

Although the node-contextualized stochastic blockmodel is useful for studying graph clustering methods, data can deviate from the model's assumptions. More generally, covariate-assisted spectral clustering can be used to find highly connected communities with relatively homogeneous covariates, where the balance between these two objectives is controlled by the tuning parameter and can be set empirically or decided by the analyst. Relatively homogeneous covariates contextualize the clusters, making them easier to interpret and allowing the analyst to focus on partitions that align with important covariates. Beyond its scientific interest, the brain graph analysis demonstrates the computational efficiency of our clustering method, since the analysis could not have been feasibly conducted with existing methods. Nevertheless, determining an optimal tuning parameter still presents a computational burden. Using a low-rank update algorithm for eigenvector decomposition can further reduce this cost.

This work is meant as a step towards statistical understanding of graphs with node covariates. Further work is needed to better understand the use of covariate-assisted spectral clustering for

network contextualization. Methods for determining the relative contribution of the graph and the covariates to a graph partition and tests to signify which covariates are informative would be useful. Ultimately, a thorough examination of the relationship between graph structure and node covariates is essential for a deep understanding of the underlying system.

ACKNOWLEDGEMENT

This research was supported by the U.S. National Institutes of Health, National Science Foundation, and Defense Advanced Research Projects Agency. The authors would like to thank Yilin Zhang, Tai Qin, Jun Tao, Soumendu Sundar Mukherjee, and Zoe Russek for helpful comments.

SUPPLEMENTARY MATERIAL

Supplementary material available at *Biometrika* online includes proofs of all the theorems and details about selecting the tuning parameter α .

REFERENCES

- AIROLDI, E. M., BLEI, D. M., FIENBERG, S. E. & XING, E. P. (2008). Mixed membership stochastic blockmodels. *J. Mach. Learn. Res.* **9**, 1981–2014.
- AMINI, A. A., CHEN, A., BICKEL, P. J. & LEVINA, E. (2013). Pseudo-likelihood methods for community detection in large sparse networks. *Ann. Statist.* **41**, 2097–122.
- BAGLAMA, J. & REICHEL, L. (2006). Restarted block Lanczos bidiagonalization methods. *Numer. Algor.* **43**, 251–72.
- BALASUBRAMANYAN, R. & COHEN, W. W. (2011). Block-LDA: Jointly modeling entity-annotated text and entity-entity links. In *Proc. 2011 SIAM Int. Conf. Data Mining*, B. Liu, H. Liu, C. Clifton, T. Washio & C. Kamath, eds., vol. 11. Philadelphia, PA: SIAM, pp. 450–61.
- CHANG, J. & BLEI, D. M. (2010). Hierarchical relational models for document networks. *Ann. Appl. Statist.* **4**, 124–50.
- CHAUDHURI, K., CHUNG, F. & TSIATAS, A. (2012). Spectral clustering of graphs with general degrees in the extended planted partition model. *J. Mach. Learn. Res.* **2012**, 1–23.
- CRADDOCK, R. C., JBABDI, S., YAN, C.-G., VOGELSTEIN, J. T., CASTELLANOS, F. X., DI MARTINO, A., KELLY, C., HEBERLEIN, K., COLCOMBE, S. & MILHAM, M. P. (2013). Imaging human connectomes at the macroscale. *Nature Meth.* **10**, 524–39.
- DAVIS, C. & KAHAN, W. M. (1970). The rotation of eigenvectors by a perturbation. *SIAM J. Numer. Anal.* **7**, 1–46.
- EYNARD, D., KOVNATSKY, A., BRONSTEIN, M., GLASHOFF, K. & BRONSTEIN, A. (2015). Multimodal manifold analysis by simultaneous diagonalization of laplacians. *IEEE Trans. Pat. Anal. Mach. Intel.* **37**, 2505–17.
- GIBERT, J., VALVENY, E. & BUNKE, H. (2012). Graph embedding in vector spaces by node attribute statistics. *Pat. Recogn.* **45**, 3072–83.
- GUNNEMANN, S., FARBER, I., RAUBACH, S. & SEIDL, T. (2013). Spectral subspace clustering for graphs with feature vectors. In *Proc. IEEE 13th Int. Conf. Data Mining*, W. Ding, T. Washio, H. Xiong, G. Karypis, B. Thuraisingham, D. Cook & X. Wu, eds. Los Alamitos, CA: IEEE, pp. 231–40.
- HANDCOCK, M. S., RAFTERY, A. E. & TANTRUM, J. M. (2007). Model-based clustering for social networks. *J. R. Statist. Soc. A* **170**, 301–54.
- HOFF, P. D., RAFTERY, A. E. & HANDCOCK, M. S. (2002). Latent space approaches to social network analysis. *J. Am. Statist. Assoc.* **97**, 1090–8.
- HOLLAND, P. W., LASKEY, K. B. & LEINHARDT, S. (1983). Stochastic blockmodels: First steps. *Social Networks* **5**, 109–37.
- HOTELLING, H. (1936). Relations between two sets of variates. *Biometrika* **28**, 321–77.
- NEVILLE, J., ADLER, M. & JENSEN, D. (2003). Clustering relational data using attribute and link information. In *Proc. Text Mining and Link Analysis Workshop, 18th Int. Joint Conf. on Artificial Intelligence*, A. Cohn, ed. San Francisco, CA: Morgan Kaufmann Publishers, pp. 9–15.
- NG, A. Y., JORDAN, M. I. & WEISS, Y. (2002). On spectral clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems 14*, T. Dietterich, S. Becker & Z. Ghahramani, eds. Cambridge, MA: MIT Press, pp. 849–56.
- NOWICKI, K. & SNIJDERS, T. A. B. (2001). Estimation and prediction for stochastic blockstructures. *J. Am. Statist. Assoc.* **96**, 1077–87.

- QIN, T. & ROHE, K. (2013). Regularized spectral clustering under the degree-corrected stochastic blockmodel. In *Advances in Neural Information Processing Systems 26*, C. Burges, L. Bottou, M. Welling, Z. Ghahramani & K. Weinberger, eds. Red Hook, NY: Curran Associates, Inc., pp. 3120–8.
- SARKAR, P. & MOORE, A. W. (2006). Dynamic social network analysis using latent space models. In *Advances in Neural Information Processing Systems 18*, Y. Weiss, B. Schölkopf & J. Platt, eds. Cambridge, MA: MIT Press, pp. 1145–52.
- SHIGA, M., TAKIGAWA, I. & MAMITSUKA, H. (2007). A spectral clustering approach to optimally combining numerical vectors with a modular network. In *Proc. 13th ACM SIGKDD Int. Conf. Know. Disc. Data Mining*, P. Berkhin, R. Caruana, X. Wu & S. Gaffney, eds. New York, NY: ACM, pp. 647–56.
- VON LUXBURG, U. (2007). A tutorial on spectral clustering. *Stat. Comp.* **17**, 395–416.
- WANG, F., DING, C. H. & LI, T. (2009). Integrated KL (K-means-Laplacian) clustering: A new clustering approach by combining attribute data and pairwise relations. In *Proc. 2009 SIAM Int. Conf. Data Mining*, C. Apte, H. Park, K. Wang & M. J. Zaki, eds., vol. 9. Philadelphia, PA: SIAM, pp. 38–48.
- YANG, J., MCAULEY, J. & LESKOVEC, J. (2013). Community detection in networks with node attributes. In *Data Mining (ICDM), 2013 IEEE 13th Int. Conf.*, W. Ding, T. Washio, H. Xiong, G. Karypis, B. Thuraisingham, D. Cook & X. Wu, eds. Los Alamitos, CA: IEEE, pp. 1151–6.
- ZHOU, Y., CHENG, H. & YU, J. X. (2009). Graph clustering based on structural/attribute similarities. In *Proceedings of the Very Large Databases Endowment*, H. V. Jagadish, S. Abiteboul, T. Milo, J. Patel & P. Rigaux, eds., vol. 2. San Jose, CA: VLDB Endowment, pp. 718–29.

[Received on 15 June 2015. Editorial decision on 13 December 2016]