

Spatial modeling of brain connectivity data via Bayesian latent distance models with nodes clustering

Emanuele Aliverti*

Department of Statistical Sciences, University of Padua, Padua, Italy

email: `aliverti@stat.unipd.it`

phone: +393475390146

Daniele Durante

Department of Decision Sciences, Bocconi University, Milan, Italy

email: `daniele.durante@unibocconi.it`

Abstract

Brain network data – measuring anatomical interconnections among a common set of brain regions – are increasingly collected for multiple individuals, and recent studies provide additional information on the brain regions of interest. These predictors typically include the 3-dimensional anatomical coordinates of the brain regions, and their membership to hemispheres and lobes. Although recent studies have explored the spatial effects underlying brain networks, there is still a lack of statistical analyses on the net connectivity topology, after controlling for spatial constraints. We answer this question via a Bayesian latent space model for network data, obtaining a meaningful representation for the net connectivity architecture via a set of latent positions which are assigned mixtures of Gaussian priors. This allows flexible

inference on brain network structures not explained by closeness in the anatomical space, and facilitates clustering among brain regions according to their latent positions.

Keywords Brain Network, Bayesian Latent Space Model, Spatial Effect, Mixture Prior

1 Introduction

Modern neuroimaging technologies, combined with state-of-the-art data pre-processing pipelines, have made it possible to reliably measure brain connectivity structures non-invasively in live humans [e.g., 6]. These data open new avenues to study underlying connectivity patterns within the brain, how they relate to brain anatomy, and if these structures change with subjects-specific phenotypes [e.g., 9, 24].

Our motivating application is drawn from a neuroimaging study on $n = 21$ subjects with no history of neurological disorders. The study provides, for each individual, information on the structural interconnections among a common set of $V = 68$ brain regions. In particular, the brain network for each individual i is available via a $V \times V$ symmetric adjacency matrix \mathbf{A}_i , with $A_{i[vu]} = A_{i[uv]} = 1$ if at least one white matter fiber has been observed between regions $v = 2, \dots, V$ and $u = 1, \dots, v - 1$ in subject $i = 1, \dots, n$, and $A_{i[vu]} = A_{i[uv]} = 0$ otherwise. In Figure 1 are reported data for subject 5 and subject 37, the individuals with minimum and maximum number of brain connections, respectively.

The $V = 68$ anatomical brain regions are those characterizing the Desikan atlas [7], and for each region v additional spatial information are available, covering the 3-dimensional anatomical centroid positions x_v, y_v, z_v , its membership to the left or right hemisphere, and in which anatomical lobe the region is located. See [17] for a classification of the Desikan atlas parcellation into anatomical lobes.

There has been an increasing focus in the neuroscience literature on exploiting the above information to learn spatial constraints underlying the wiring connectivity structures within the brain; see [3, 23], and the references cited therein for a general discussion on recent empirical findings. Consistent with results in these studies, Figure 2 highlights — in our dataset — a general preference



Figure 1: Adjacency matrices characterizing the brain network observed for two different subjects. Black refers to an edge; white to a non-edge.

for the brain regions to connect with others that are spatially closer, and belonging to the same hemisphere and lobe. However, according to Figure 2, these patterns, although informative, are not sufficient to explain the whole variability in regions connectivity, thereby motivating studies on the net wiring structures, not related to spatial constraints.

Statistical modeling of replicated brain networks is still in its infancy, and a statistical answer to the above questions is still lacking. We take a first step towards addressing this relevant goal with a Bayesian latent space models for social networks, to jointly estimate the effects of the spatial constraints in the brain network architecture, and the net connectivity properties after controlling for the anatomical brain organization. In particular, this net topology is modeled via a set of 3-dimensional latent positions accounting for the wiring structures not explained by the 3-dimensional anatomical positions. A flexible mixture of Gaussian priors over the latent coordinates is introduced to generalize the existing latent space approaches, facilitating a multivariate clustering scheme that leverages on local borrowing of information to improve the estimation of the latent positions.

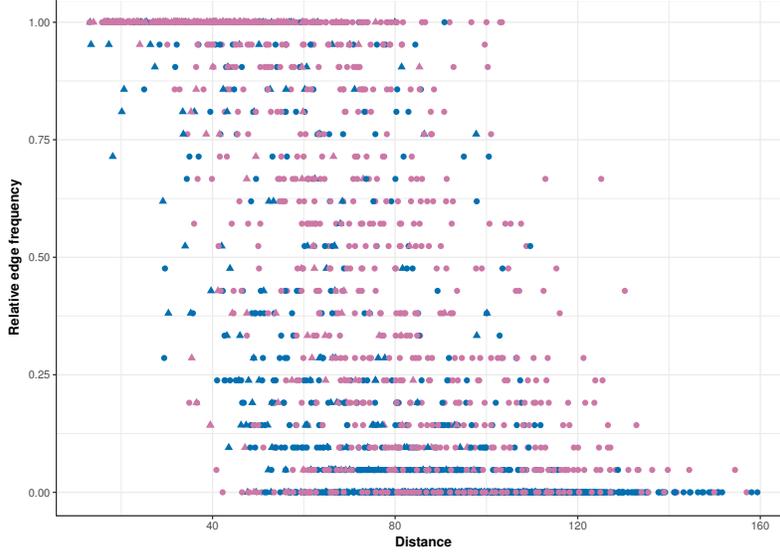


Figure 2: Plot of the relative edge frequencies $\sum_{i=1}^{n=21} A_{i[vu]}/n$ for each pair of brain regions $v = 2, \dots, V$ and $u = 1, \dots, v - 1$, versus their anatomical euclidean distance d_{vu} . Blue denotes pairs of brain regions in the same hemisphere, whereas purple characterizes pairs of regions in different hemispheres. Triangles and circles represent pairs of brain regions in the same and different lobes, respectively.

2 Latent space modeling of brain networks

Let $\Pr(A_{i[vu]} = 1) = \pi_{vu} \in (0, 1)$, denote the population probability of an edge between the brain regions v and u , for every $v = 2, \dots, V$, $u = 1, \dots, v - 1$ and $i = 1, \dots, n$. We seek a flexible representation for the edge-specific probability which allows learning of anatomical and latent constraints in the the brain regions. Taking the lead from latent variable models for networks, we will assume that edges are distributed according to conditional independent Bernoulli random variables,

$$A_{i[vu]} \mid \pi_{vu} \sim \text{Bern}(\pi_{vu}), \text{ for each } i = 1, \dots, n, \text{ and } v = 2, \dots, V, u = 1, \dots, v - 1. \quad (1)$$

To learn π_{vu} , will consider a latent space model representation capable to infer anatomical and latent constraints underlying brain connectivity structures. Such representation characterise heterogeneity among nodes by placing them in an d -dimensional Euclidean latent space, and modelling

the edge specific probability through a function of a suitable distance in the latent space. In order to favour the interpretation of the model, we will fix the size of the latent space to 3 and measure distance between points in the latent space through the Euclidean distance. This specification favours a natural analogy between physical and latent space, and improves the interpretation of the latent constraints of the brain connectivity patterns.

More precisely, we model

$$\text{logit}(\pi_{vu}) = \beta_0 + \beta_1 \mathbf{hem}_{vu} + \beta_2 \mathbf{lobe}_{vu} + \beta_3 d_{vu} - \bar{d}_{vu}, \quad (2)$$

for each $v = 2, \dots, V$ and $u = 1, \dots, v - 1$, where \mathbf{hem}_{vu} and \mathbf{lobe}_{vu} are binary predictors indicating shared membership to the same hemisphere and lobe, respectively, whereas d_{vu} and \bar{d}_{vu} denote the anatomical and latent Euclidean distances between brain regions v and u , respectively. Hence, according to equation (2) — and consistent with Figure 2 — the edge probability between regions v and u is allowed to change with their shared hemisphere and lobe membership, and their anatomical distance $d_{vu} = \sqrt{(x_v - x_u)^2 + (y_v - y_u)^2 + (z_v - z_u)^2}$. The latent distance $\bar{d}_{vu} = \sqrt{(\bar{x}_v - \bar{x}_u)^2 + (\bar{y}_v - \bar{y}_u)^2 + (\bar{z}_v - \bar{z}_u)^2}$ adds instead a layer of flexibility, which allows modeling of net connectivity architectures not related to brain anatomy. In particular, according to equation (2), the closer two brain regions are in this latent space, the more likely is to observe a connection among them, after controlling for anatomical constraints. Hence, by providing inference on the latent positions \bar{x}_v , \bar{y}_v and \bar{z}_v for each brain region $v = 1, \dots, V$, we allow a deeper understanding of brain connectivity.

The statistical model in equations (1)–(2) is in the same spirit of the flexible latent space model proposed by [16], and therefore is characterized by similar properties and theoretical support. However, differently from [16], we focus on joint modeling of multiple adjacency matrices $\mathbf{A}_1, \dots, \mathbf{A}_n$, instead of just a single one. Although this difference may apparently require novel computational methods and inference procedures, note that, under representation (1), a sufficient statistic $A_{[vu]}$ in the joint likelihood for the data $A_{i[vu]}$, $i = 1, \dots, n$, is $A_{[vu]} = \sum_{i=1}^n A_{i[vu]} \sim \text{Binom}(n, \pi_{vu})$, for every pair of regions $v = 2, \dots, V$ and $u = 1, \dots, v - 1$. Hence, joint modeling of multiple

adjacency matrices $\mathbf{A}_1, \dots, \mathbf{A}_n$ under (1)–(2), coincides with providing inference on the weighted network $\mathbf{A} = \sum_{i=1}^n \mathbf{A}_i$ with binomial edges having probabilities factorized as in (2).

From a scientific perspective, it is desirable to obtain a richer interpretation of the estimated model, and this can be obtained, among many other approaches, by estimating groups of brain regions devoted to locally specialized processing [2]. In our application, this aim can be achieved estimating clusters over the latent space, since this is equivalent to detect groups of brain regions characterised by similar behaviours in terms of connectivity patterns, after controlling for the anatomical properties of the brains. Although it is possible to cluster the latent coordinates after estimation, according to their closeness in the latent space, the inclusion of clusters estimation in the inferential process is generally preferred from a statistical perspective, since it allows to borrow information among different groups, exploiting cluster membership in order to improve estimation of latent positions. An important justification for model based clustering, its benefits and theoretical properties can be found in [10] and references cited therein.

A Bayesian latent space model that generalize [16] is introduced in [14] assuming a mixture of multivariate isotropic Gaussian prior for the latent positions, where each component encodes a different group of nodes. In our motivating example, this specification would assume that the 3-dimensional latent space can be partitioned into spherical groups of brain regions that are similar with respect to *all* the elements of their latent coordinates. Although this strategy improves the flexibility of the standard latent space model of [16], the assumption of a single cluster membership could be too restrictive to characterize and describe complex connectivity patterns. Two generic brain regions could be assigned to location in the latent space such that their position is very similar with respect to some coordinates, but significantly far apart over the remaining, and the clustering strategy mentioned above would be unable to detect any local similarity, since it groups only brain regions that are globally close. In order to improve the analysis of the unobservable connectivity patterns, it is important to capture the local proximity information, avoiding either an inefficient allocation of brain regions in too many clusters and an oversimplified representation with few groups. The idea to cluster observations according to a subset of parameters is related to the methods for analysing dependency structure in latent class analyses [5], and has proven to

improve the flexibility and performance of complex models, for example in functional data analysis [8, 20] or tensor data [1].

We generalize the contribution of [14] introducing a Bayesian latent space model in which the latent positions are characterized by coordinate-specific mixture priors, in order to induce a flexible clustering scheme that allows brain regions to belong to one group along each latent coordinate, improving inference on the latent position through parsimonious local borrowing of information. The introduction of a shrinkage prior over the clusters mixing probabilities facilitates the regulation of model complexity, allowing a different number of mixture components over each latent coordinate and integrating model regularization into the inferential process. Our approach unifies the explicative power of the standard latent space model, where the latent coordinates are characterised by a continuous domain, with the great flexibility of the cluster methods, where the unobservable patterns are characterised by a discrete nature.

3 Bayesian inference

We will rely on a full Bayesian specification of the proposed model, since it facilitates to borrow information among latent coordinates and cluster components and can integrate the estimation of the number of mixture components in the inferential process via an appropriate prior distribution. In order to represent groups in the latent space, we will assume that the elements of the vectors $\bar{\mathbf{x}} = [\bar{x}_1, \dots, \bar{x}_V]$, $\bar{\mathbf{y}} = [\bar{y}_1, \dots, \bar{y}_V]$ and $\bar{\mathbf{z}} = [\bar{z}_1, \dots, \bar{z}_V]$ are drawn from mixtures of univariate Gaussian distributions, independently across each coordinate. More formally,

$$\bar{x}_v \sim \sum_{h=1}^{H_{\mathbf{x}}} \nu_{\mathbf{x}_h} \mathcal{N}(\mu_{\mathbf{x}_h}, \sigma_{\mathbf{x}_h}^2), \quad \bar{y}_v \sim \sum_{h=1}^{H_{\mathbf{y}}} \nu_{\mathbf{y}_h} \mathcal{N}(\mu_{\mathbf{y}_h}, \sigma_{\mathbf{y}_h}^2), \quad \bar{z}_v \sim \sum_{h=1}^{H_{\mathbf{z}}} \nu_{\mathbf{z}_h} \mathcal{N}(\mu_{\mathbf{z}_h}, \sigma_{\mathbf{z}_h}^2), \quad (3)$$

for all $v = 1, \dots, V$, where $\nu_{\mathbf{x}_h}$, $\nu_{\mathbf{y}_h}$ and $\nu_{\mathbf{z}_h}$ are the probability that a brain region belongs to the generic group h over the latent coordinate $\bar{\mathbf{x}}$, $\bar{\mathbf{y}}$ and $\bar{\mathbf{z}}$ respectively, and with

$$\nu_{\mathbf{x}_h} \geq 0, \sum_{h=1}^{H_{\mathbf{x}}} \nu_{\mathbf{x}_h} = 1, \quad \nu_{\mathbf{y}_h} \geq 0, \sum_{h=1}^{H_{\mathbf{y}}} \nu_{\mathbf{y}_h} = 1, \quad \nu_{\mathbf{z}_h} \geq 0, \sum_{h=1}^{H_{\mathbf{z}}} \nu_{\mathbf{z}_h} = 1.$$

Note that the specification of the mixture prior as in equation (3) and the likelihood in equations (1) and (2) effectively generalizes the existing Bayesian latent space model for networks, in particular [16, 14], since the standard latent space model of [16] is obtained as a degenerate case if $H_{\mathbf{x}} = H_{\mathbf{y}} = H_{\mathbf{z}} = 1$, while the mixture model specification of [14] can be regarded as a particular case of equation (3) in which cluster membership is identical across each latent coordinates. In agreement with the mixture representation specified in equation (3), every brain region is allowed to belong to a different group over each latent dimension, creating local clusters with respect to the each coordinate. Moreover, since every dimension of the latent space can be characterized by a different number of mixture components, it is important to control the model complexity with a proper prior distribution over the mixtures weights, to avoid coordinates characterised by too many components or too few. With this objectives in mind, it is natural to take the lead from [22] and specify conjugate Dirichlet distributions for the mixture weights, setting the number of mixture components to an upper conservative bound H . This specification facilitates the adaptive property of the Dirichlet distribution in deleting redundant components, estimating the appropriate number of clusters over every latent coordinate. Hence, we fix $H_{\mathbf{x}} = H_{\mathbf{y}} = H_{\mathbf{z}}$ to an upper bound H , let $\boldsymbol{\nu}_{\mathbf{x}} = (\nu_{\mathbf{x}_1}, \dots, \nu_{\mathbf{x}_H})$, $\boldsymbol{\nu}_{\mathbf{y}} = (\nu_{\mathbf{y}_1}, \dots, \nu_{\mathbf{y}_H})$, $\boldsymbol{\nu}_{\mathbf{z}} = (\nu_{\mathbf{z}_1}, \dots, \nu_{\mathbf{z}_H})$ and specify

$$\boldsymbol{\nu}_{\mathbf{x}} \sim \text{Dirichlet}\left(\frac{1}{H}, \dots, \frac{1}{H}\right), \quad \boldsymbol{\nu}_{\mathbf{y}} \sim \text{Dirichlet}\left(\frac{1}{H}, \dots, \frac{1}{H}\right), \quad \boldsymbol{\nu}_{\mathbf{z}} \sim \text{Dirichlet}\left(\frac{1}{H}, \dots, \frac{1}{H}\right). \quad (4)$$

The prior distributions for the mean and variance of the Gaussian kernels are specified as conjugate Normal-Inverse Gamma, with common hyperparameter settings. Taking as a reference a generic mixture component over the latent coordinate $\bar{\mathbf{x}}$, we specify

$$\mu_{\mathbf{x}_h} \mid \sigma_{\mathbf{x}_h}^2 \sim \text{N}(\alpha_0, \sigma_{\mathbf{x}_h}^2 / \kappa_0) \quad \frac{1}{\sigma_{\mathbf{x}_h}^2} \sim \text{Gamma}\left(\frac{\eta_0}{2}, \frac{\eta_0 \xi_0}{2}\right), \quad (5)$$

for every value of h . Since the component-specific mean parameters are centred on a common value, it is important to induce substantial variability across mixture components via suitable choice of the hyperparameters to avoid degenerate situations. Lastly, we define a multivariate Gaussian prior distribution for the joint vector of coefficients,

$$\boldsymbol{\beta} = [\beta_0, \dots, \beta_3]^\top \sim \text{N}_4(\mathbf{0}, \Psi_0), \quad \Psi_0 = \text{diag}(\psi_0, \dots, \psi_3).$$

[21] introduced an ergodic [4] Gibbs sampling strategy in the context of Bayesian logistic regression, showing that the binomial distribution, with probability of success parametrized in terms of log-odds, can be expressed as a scale mixture of a Gaussian with Pólya-Gamma mixing measure. This result imply that an intermediate Pólya-Gamma data-augmentation step induce conjugancy between the Gaussian prior and binomial likelihood, and this scheme can be adapted for our application with opportune variations. However, the Euclidean distance of equation (2) is not a linear functional, and an additional Metropolis-Hastings step is required to update the latent positions. The full conditional distributions are available in closed form for the remaining parameters.

We outline here the main features of the algorithm for posterior computation, while the complete algorithm is reported in Appendix A. The Gibbs sampling iterates over the following steps:

1. Augment the observed data, sampling the latent Pólya-Gamma random variable
2. Sample the joint covariates coefficient from its full conditional Normal distribution
3. Cycle over the dimensions of the latent space:
 - 3.1. Sample the latent class indicator from its full conditional Categorical distribution
 - 3.2. Update each latent positions via a Metropolis-Hastings step

3.3. Sample the Gaussian kernels parameters of the mixture components from their Normal-Inverse Gamma full conditional distribution

3.4. Sample the cluster weights from the full conditional Dirichlet distribution, updated with cluster sizes

Since the latent coordinates are non identifiable up to translation, rotations and reflections, [14] impose unity-norm constraints in their approach. However, we prefer to avoid any constraint during the estimation, and post-process the unconstrained sample through a Procrustean transformation, since our empirical findings suggest that unconstrained estimation speeds up the mixing time of the Gibbs sampler, as noted in several Bayesian hierarchical model [e.g., 13, 18]. Since the Procrustean rotation is unique, the transformed sample coincides with a sample from the constrained posterior distribution of the latent positions [16]. An additional issue regarding the identifiability of the parameters is due to the known label-switching of mixture components. When this problem occurs, the relabelling algorithm of [26] can be employed, and the moderate dimensionality of our motivating example allows a feasible implementation, in terms of computational time and required additional storage.

4 Simulation

In order to assess the empirical performance of the proposed methodology, we consider a simple simulation study under three different settings. Artificial data are created to resemble realistic brain structures with 20 brain regions for the ease of computation and graphical representation. Data are generated according to the model outlined in equation (2), with one dummy covariate that mimic hemisphere membership. The latent space structures is assumed of growing complexity across the settings, in order to create challenging situations. In the first setting, the first coordinate is sampled from a mixture of two normal distributions, while the second and third are sampled from a standard normal distribution. In the second setting, the first and second coordinates are sampled from two independent mixtures of two normal distributions, while the third is sampled from a standard normal. In the third and last setting, each coordinate is sampled from a mixture



Figure 3: Results for the simulation study. True coordinates are reported as light gray cross. Estimated latent coordinates (posterior mean) are reported as full points, with color and shape according to the cluster membership (MAP or after relabeling) estimated over the indicated latent coordinate.

of two normal distributions. The main challenge is determined by different local complexity and varying cluster membership across each dimension, even when the number of clusters is equal across the latent dimensions.

The hyperparameters were specified, in every scenario, as follow: $\psi_0 = \psi_1 = 2$, $\kappa_0 = 2$, $\alpha_0 = 0$, $\eta_0 = 30$ and $\xi_0 = 1$. This values were selected to induce a non-informative prior on the coefficient β , and to centre the prior of the mixture kernels roughly around a non informative normal distribution,

in order to create substantial variability across mixture components. The upper bound for the mixture components was set as $H = 5$. We simulated 5000 iterations, discarding the first 2500 as burn-in, and the Metropolis-Hastings steps was conducted via a normal random walk, tuning the step size to obtain an acceptance ratio close to 0.2 [e.g., 12, 11]. Label-switching problem, when encountered, was faced with the method of [26] as implemented in the R package `label.switching` [19]. Graphical results are reported in Figure 3, and compared to the original values. As the complexity of the latent space increases, it becomes more difficult to detect with great precision the exact location of the latent positions, however the empirical results are satisfactory, in particular as far as the clustering identification is concerned. Our algorithm was able to detect the correct number of components over each component in every scenario, and to allocate nodes coherently with the data generating mechanism.

5 Results

We applied our procedure to the dataset described in section 1. Posterior inference is performed with the same settings as in the simulation study, obtaining similar results in terms of convergence and mixing of the chains.

We additionally conducted sensitivity analysis focusing on the parameters η_0, κ_0, ξ_0 of the Gaussian kernels of the mixture component, and moderate changes in these parameters do not change inferential conclusions. The posterior summaries referred to covariates coefficients and latent coordinates were particularly robust even against severe changes of the hyperparameters, while cluster memberships was more sensible. This issue is well known in the Bayesian literature of mixture models for clustering [e.g. 12, Chapter 22], and a common approach consists in using informative prior based on data driven justifications, as outlined for example in [18, sec 2.4].

We reported in Table 1 the posterior summaries for the elements of the vector β , labelled with the respective covariate name. The covariate `distance` has been rescaled to avoid numerical issue during the estimating procedure, a standard practice in supervised learning. Current empirical findings indicate a general preference for the brain regions to connect with others that are spatially

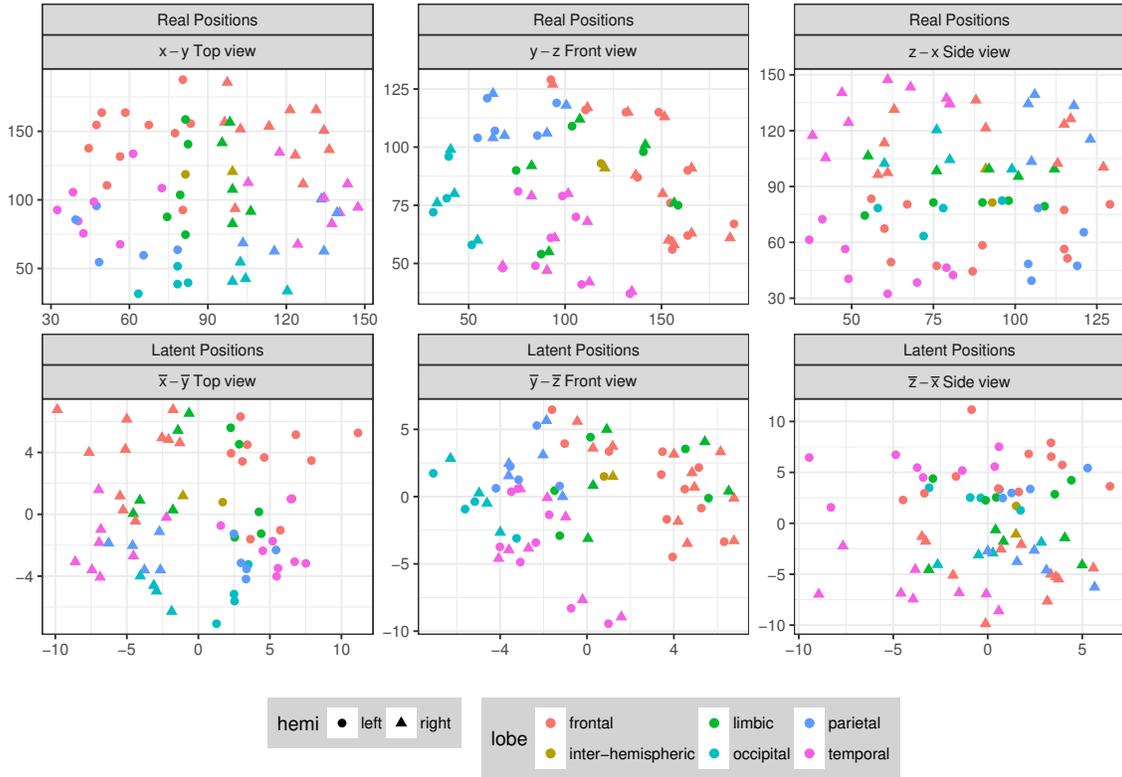


Figure 4: Graphical representation of real coordinates (upper panel) and posterior means of the corresponding latent positions (lower panel), with lobes and hemisphere information.

closer, and belonging to the same hemisphere and lobe. This results are coherent with the current neuroscientific knowledge [e.g., 3, 23] and with our discussion on section 1.

According to Figure 4, the 3-dimensional anatomical and latent positions are interestingly related. However, the coordinates of the brain regions in the latent space do not coincide with the anatomical positions, and characterize those net connectivity patterns after controlling for anatomical constraints. For example, regions in the temporal lobe have a more peripheral position in the latent space, whereas regions in the parietal and occipital lobes become more central. Despite our model explicitly takes into account lobes and hemisphere membership, the latent coordinates still capture several signal of dependency from this fundamental anatomical properties.

In order to interpret the estimated clustering information, Figure 5 compare the results from

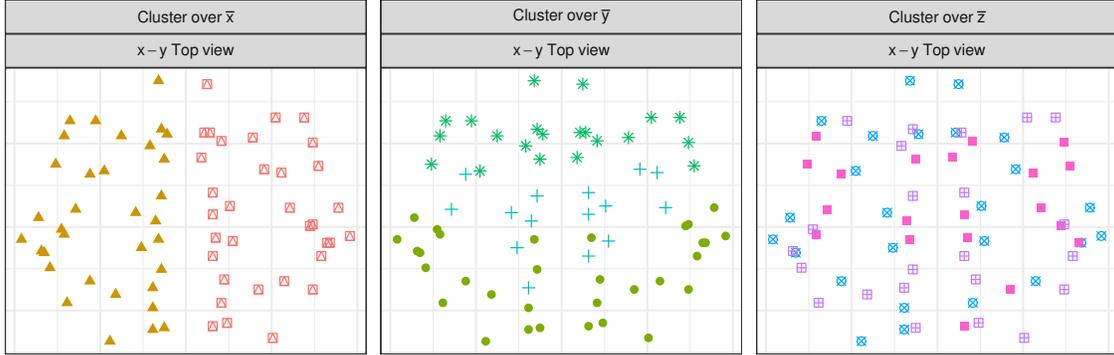


Figure 5: Graphical representation of the cluster partition. Colours and shapes are defined by cluster membership, in order estimated over the latent coordinate \bar{x} (left plot), over \bar{y} (central plot) and over \bar{z} (right plot). Location of the brain regions are given by real coordinate.

the clustering scheme with the real spatial coordinates of the brain regions. Each plot represent a different segmentation the brain architecture, and correspond to clusters estimated over a different latent coordinate. The leftmost plot, representing clusters estimated over the latent \bar{x} , clearly encodes the brain segmentation into left and right hemisphere. This results is coherent with the result reported in Figure 4, where the distinction between hemispheres was still persistent in the latent space, and in particular over the first latent coordinate. Over the second latent coordinate \bar{y} , cluster membership is remarkably different from the one previously described, being more related to lobes segmentation. A further partition scheme is observed for clusters estimated over the third coordinate \bar{z} , in which groups are coherent with lobes segmentation although characterized by a completely different fashion. This results highlight the great flexibility of the latent space specification, being able both to identify the proper number of clusters in every latent coordinate — 2 over \bar{x} and 3 over \bar{y} and \bar{z} — and to induce cluster membership, over each latent dimension,

Table 1: Posterior summaries for the parameters β . Mean, median, standard deviation and High Posterior probability interval of level 95%. Parameters are labelled with the covariate they refer to.

	Mean	Median	Std. dev	HPD _{95%}
(Intercept)	7.27	7.27	0.18	(6.94, 7.60)
lobes	0.24	0.24	0.06	(0.13, 0.35)
hemisphere	0.60	0.61	0.18	(0.29, 0.92)
distance	-0.35	-0.35	0.06	(-0.47, -0.23)

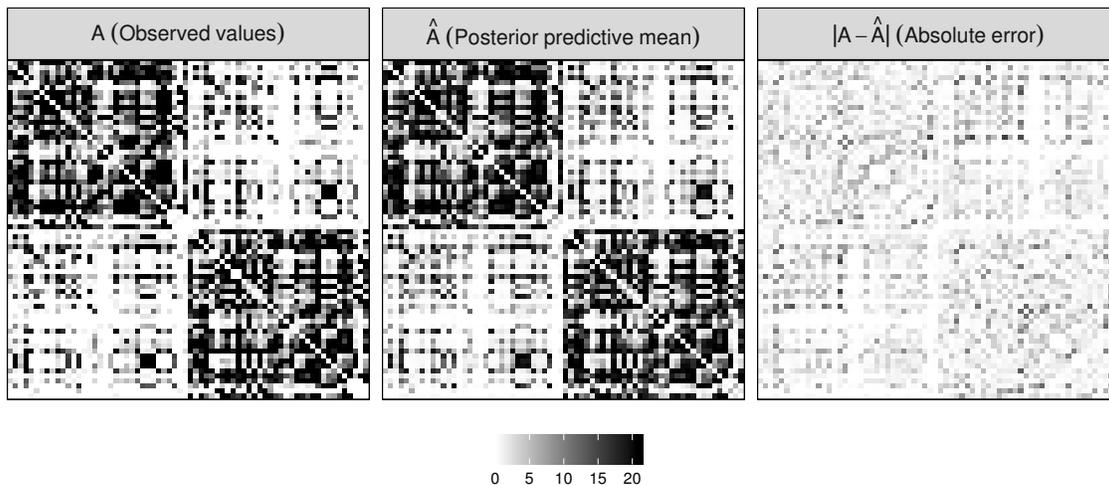


Figure 6: Graphical representation of the observed adjacency matrix \mathbf{A} , the mean of its posterior predictive density $\hat{\mathbf{A}}$, and the absolute error. Colors range from white to black as the values go from 0 to $n = 21$.

not necessarily equal.

The empirical performance of the proposed model is evaluated in Figure 6. The most natural estimator for a Bayesian model is the posterior predictive density, since it provides estimates that take into account the uncertainty deriving from the estimation procedure [e.g., 12]. The posterior predictive density for the matrix \mathbf{A} is summarised via its posterior mean, namely $\hat{\mathbf{A}}$, computed via Monte Carlo simulation through the Markov chains previously stored. Results seem satisfactory in terms of predictive performance, and the proposed estimator captures the most relevant structural patterns of the modelled network, in particular hemisphere and lobes specific wiring mechanisms. The residual absolute error matrix does not indicate pathological disagreement between the observed data and their predictions.

6 Discussion

We have proposed a Bayesian latent space model for network data to investigate the topology of the brain connectivity patterns. Our proposal estimates jointly the effect of the spatial constraints and the structure of the connectivity patterns after controlling for the anatomical constraints, and

can be regarded as a generalization of the latent space models available in literature, in particular [16, 14]. It improves the flexibility and the explanatory power of the previous approaches introducing a multivariate clustering scheme over the latent coordinates to improve the detection of local similarities among brain regions. Empirical findings on the anatomical constraints are coherent with the current state of the art of neuroscientific literature, confirming that brain regions are more likely to be connected with regions that belong to the same lobe, hemisphere and that are spatially closer. The latent coordinates, moreover, are interestingly related to the anatomical one, and the multivariate clustering is characterized by a clear anatomical interpretation, visually highlighted reporting the estimated groups in the original space.

Other distances, instead of the Euclidean, can be employed in our specification. The choice of this metric for the latent space was mainly justified to create a direct analogy with the real space, where the Euclidean distance is the most natural one. For example, the weighted distance of the latent eigenmodel of [15] could be employed with minor modification, and its form is particularly appealing since it would avoid the Metropolis-Hastings step involved in our algorithm. A similar justification holds as far as the size of the latent space is concerned, since it was set in advance equal to 3 to motivate the analogy with the real coordinates, but in principle can be regarded as a parameter, hence assigned a prior distribution and estimated from the data.

One important future direction is to adapt our specification with larger networks, characterized by a huge number of nodes V , since modern technologies are capable to divide the brain structure in thousands of nodes [25]. Since our approach relies on an Gibbs sampler, the required computational time when the number of nodes increase becomes prohibitive rapidly, and there are great margins of improvement. A promising future direction could be to integrate the recent developments in MCMC theory in order to construct an approximate Markov Chain that converge more rapidly to the stationary distribution.

References

- [1] A. Bhattacharya and D.B. Dunson. “Simplex factor models for multivariate unordered categorical data”. In: *Journal of the American Statistical Association* 107.497 (2012), pp. 362–377.
- [2] E. Bullmore and O. Sporns. “Complex brain networks: graph theoretical analysis of structural and functional systems”. In: *Nature Reviews Neuroscience* 10.3 (2009), pp. 186–198.
- [3] E. Bullmore and O. Sporns. “The economy of brain network organization”. In: *Nature Reviews Neuroscience* 13 (2012), pp. 336–349.
- [4] H.M. Choi and J.P. Hobert. “The Polya-Gamma Gibbs sampler for Bayesian logistic regression is uniformly ergodic”. In: *Electronic Journal of Statistics* 7 (2013), pp. 2054–2064.
- [5] H. Chung, S.T. Lanza, and E. Loken. “Latent transition analysis: Inference and estimation”. In: *Statistics in medicine* 27.11 (2008), pp. 1834–1854.
- [6] R. Craddock et al. “Imaging human connectomes at the macroscale”. In: *Nature Methods* 10 (2013), pp. 524–539.
- [7] R.S. Desikan et al. “An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest”. In: *Neuroimage* 31 (2006), pp. 968–980.
- [8] D.B. Dunson. “Multivariate kernel partition process mixtures”. In: *Statistica Sinica* 20.4 (2010), p. 1395.
- [9] D. Durante and David B.D. Dunson. “Bayesian inference and testing of group differences in brain networks”. In: *Bayesian Analysis* (2017), in press.
- [10] C. Fraley and A.E. Raftery. “Model-based clustering, discriminant analysis, and density estimation”. In: *Journal of the American statistical Association* 97.458 (2002), pp. 611–631.
- [11] A. Gelman, G.O. Roberts, and W.R. Gilks. “Efficient Metropolis jumping rules”. In: *Bayesian statistics* 5.599-608 (1996), p. 42.
- [12] A. Gelman et al. *Bayesian data analysis*. Vol. 2. CRC press Boca Raton, FL, 2014.

- [13] A. Gelman et al. “Using redundant parameterizations to fit hierarchical models”. In: *Journal of Computational and Graphical Statistics* 17.1 (2008), pp. 95–122.
- [14] M.S. Handcock, A.E. Raftery, and J.M. Tantrum. “Model-based clustering for social networks”. In: *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 170.2 (2007), pp. 301–354.
- [15] P.D. Hoff. “Modeling homophily and stochastic equivalence in symmetric relational data”. In: *Advances in neural information processing systems*. 2008, pp. 657–664.
- [16] P.D. Hoff, A.E. Raftery, and M.S. Handcock. “Latent space approaches to social network analysis”. In: *Journal of the American Statistical Association* 97.460 (2002), pp. 1090–1098.
- [17] X. Kang et al. “Hemispherically-unified surface maps of human cerebral cortex: reliability and hemispheric asymmetries”. In: *PloS one* 7 (2012), e45582.
- [18] P.N. Krivitsky and M.S. Handcock. “Fitting latent cluster models for networks with latent-net”. In: *Journal of Statistical Software* 24.05 (2008).
- [19] P. Papastamoulis. “label.switching: An R Package for Dealing with the Label Switching Problem in MCMC Outputs”. In: *Journal of Statistical Software, Code Snippets* 69.1 (2016), pp. 1–24.
- [20] M. Petrone S.and Guindani and A.E. Gelfand. “Hybrid Dirichlet mixture models for functional data”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 71.4 (2009), pp. 755–782.
- [21] N.G. Polson, J.G. Scott, and J. Windle. “Bayesian inference for logistic models using Pólya–Gamma latent variables”. In: *Journal of the American statistical Association* 108.504 (2013), pp. 1339–1349.
- [22] J. Rousseau and K. Mengersen. “Asymptotic behaviour of the posterior distribution in overfitted mixture models”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 73.5 (2011), pp. 689–710.

- [23] O. Sporns. “Structure and function of complex brain networks”. In: *Dialogues in Clinical Neuroscience* 15 (2013), pp. 247–262.
- [24] C.J. Stam. “Modern network science of neurological disorders”. In: *Nature Reviews Neuroscience* 15 (2014), pp. 683–695.
- [25] M.L. Stanley et al. “Defining nodes in complex brain networks”. In: *Frontiers in computational neuroscience* 7 (2013).
- [26] M. Stephens. “Dealing with label switching in mixture models”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 62.4 (2000), pp. 795–809.

A Appendix

Gibbs sampler

The quantities involved in the Gibbs sampling algorithm can be expressed in a more convenient vectorised form. Denote with $\mathcal{L}(\cdot)$ the operator $\mathcal{L} : \mathbb{R}^{V \times V} \rightarrow \mathbb{R}^{V(V-1)/2}$ that vectorises the lower-triangular elements of a matrix columnwise. Let $\tilde{\mathbf{a}} = \mathcal{L}(\mathbf{A}) - \frac{n}{2} \mathbf{1}_{V(V-1)/2}$, where $\mathbf{1}$ is a constant vector equal to 1. Denote by \mathbf{X} the $V(V-1)/2 \times 4$ matrix with the first column constant and equal to 1, and the remaining columns equal to the vectorised version of the edge covariates `hem`, `lobe`, and `d`. Let $\bar{\mathbf{d}}$ the $V(V-1)/2$ vector corresponding to the vectorised version of \bar{d}_{vu} .

The Gibbs sampler iterates over the following steps:

1. Pólya-Gamma update

The introduction of the Pólya-Gamma latent variable induce conjugancy between the Normal prior for $\boldsymbol{\beta}$ and the binomial likelihood.

Update from the *full conditional* latent Pólya-Gamma distribution:

$$\omega_l | - \sim \text{PG}(n, \mathbf{X}_l \boldsymbol{\beta} - \bar{\mathbf{d}}_l),$$

with $l = 1, \dots, V(V-1)/2$

2. Coefficients update

Denote by $\mathbf{\Omega}$ the $V(V-1)/2 \times V(V-1)/2$ diagonal matrix with the Pólya-Gamma sampled value, Sample β from the *full-conditional* normal distribution:

$$\beta | - \sim N_4(\mu_\beta, \Sigma_\beta), \quad \Sigma_\beta = \left(\mathbf{X}^\top \mathbf{\Omega} \mathbf{X} + \Psi_0^{-1} \right)^{-1}, \quad \mu_\beta = \Sigma_\beta \left(\mathbf{X}^T \tilde{\mathbf{a}} + \mathbf{X}^T \mathbf{\Omega} \bar{\mathbf{d}} \right)$$

3. Cluster assignment update

The algorithm now proceeds separately across every latent coordinate. We report the detailed steps for the first latent coordinate $\bar{\mathbf{x}}$.

Denote as $k_{\mathbf{x}v}$ the latent cluster indicator relative to the brain region v with respect to the latent coordinate \mathbf{x} . Sample the latent cluster indicator from the conjugate Categorical distribution, with probability

$$P[k_{\mathbf{x}v} = h] = \frac{\nu_{\mathbf{x}_h} \phi(\bar{x}_v; \mu_{\mathbf{x}_h}, \sigma_{\mathbf{x}_h}^2)}{\sum_{h=1}^{H_{\mathbf{x}}} \nu_{\mathbf{x}_h} \phi(\bar{x}_v; \mu_{\mathbf{x}_h}, \sigma_{\mathbf{x}_h}^2)},$$

for $v = 1, \dots, V$ and with $\phi(\bar{x}_v; \mu_{\mathbf{x}_h}, \sigma_{\mathbf{x}_h}^2)$ indicating the normal density with parameters $(\mu_{\mathbf{x}_h}, \sigma_{\mathbf{x}_h}^2)$ evaluated in \bar{x}_v .

4. Latent coordinates update

Denote with $P(\tilde{\mathbf{a}} | \beta, \bar{\mathbf{d}}, \mathbf{X})$ the full data binomial likelihood, with log-odds as defined in equation (2).

Update the generic latent coordinate \bar{x}_v with a Metropolis-Hastings step, with normal random walk and considering the posterior distribution proportional to

$$\bar{x}_v | k_{\mathbf{x}v} = h, - \propto \phi(\bar{x}_v; \mu_{\mathbf{x}_h}, \sigma_{\mathbf{x}_h}^2) P(\tilde{\mathbf{a}} | \beta, \bar{\mathbf{d}}, \mathbf{X})$$

for $v = 1, \dots, V$.

5. Cluster kernel update

Denote as

$$n_{\mathbf{x}h} = \sum_{v=1}^V \mathbb{I}[k_{\mathbf{x}v} = h], \quad \bar{m}_{\mathbf{x}h} = n_{\mathbf{x}h}^{-1} \sum_{v=1}^V \bar{x}_v \mathbb{I}[k_{\mathbf{x}v} = h], \quad s_{\mathbf{x}h}^2 = \sum_{v=1}^V (\bar{x}_v - \bar{m}_{\mathbf{x}h})^2 \mathbb{I}[k_{\mathbf{x}v} = h]$$

the cluster size, intra-cluster mean and intra-cluster sum of squares deviation from the mean, respectively.

Sample the cluster mean and variance from the conjugate Normal-inverse Gamma full conditional

$$\begin{aligned} \frac{1}{\sigma_{\mathbf{x}h}^2} \mid - &\sim \text{Gamma}\left(\frac{\eta_{\mathbf{x}h}}{2}, \frac{\eta_{\mathbf{x}h} \xi_{\mathbf{x}h}}{2}\right), & \mu_{\mathbf{x}h} \mid \sigma_{\mathbf{x}h}^2, - &\sim \text{N}\left(\alpha_{\mathbf{x}h}, \frac{\sigma_{\mathbf{x}h}^2}{\kappa_{\mathbf{x}h}}\right) \\ \text{with } \eta_{\mathbf{x}h} &= \eta_0 + n_{\mathbf{x}h}, \quad \kappa_{\mathbf{x}h} = \kappa_0 + n_{\mathbf{x}h}, \quad \alpha_{\mathbf{x}h} = \frac{\kappa_0 \alpha_0 + n_{\mathbf{x}h} \bar{m}_{\mathbf{x}h}}{\kappa_{\mathbf{x}h}} \\ \text{and } \xi_{\mathbf{x}h} &= \frac{1}{\eta_{\mathbf{x}h}} \left[\eta_0 \xi_0 + s_{\mathbf{x}h}^2 + \frac{\kappa_0 n_{\mathbf{x}h}}{\kappa_{\mathbf{x}h}} (\bar{m}_{\mathbf{x}h} - \alpha_0)^2 \right], \end{aligned}$$

for $h = 1, \dots, H$.

6. Update the cluster weights

Sample from the full conditional Dirichlet distribution, updated with cluster sizes

$$\boldsymbol{\nu}_{\mathbf{x}} \mid - \sim \text{Dirichlet}\left(\frac{1}{H} + n_{\mathbf{x}1}, \dots, \frac{1}{H} + n_{\mathbf{x}H}\right)$$