

HHS Public Access

Author manuscript Inf Process Med Imaging. Author manuscript; available in PMC 2020 July 22.

Published in final edited form as: Inf Process Med Imaging. 2019 June ; 11492: 867–879. doi:10.1007/978-3-030-20351-1 68.

Variational Autoencoder with Truncated Mixture of Gaussians for Functional Connectivity Analysis

Qingyu Zhao¹, Nicolas Honnorat², Ehsan Adeli¹, Kilian M. Pohl²

¹Stanford University

²SRI International

Abstract

Resting-state functional connectivity states are often identified as clusters of dynamic connectivity patterns. However, existing clustering approaches do not distinguish major states from rarely occurring minor states and hence are sensitive to noise. To address this issue, we propose to model major states using a non-linear generative process guided by a Gaussian-mixture distribution in a low-dimensional latent space, while separately modeling the connectivity patterns of minor states by a non-informative uniform distribution. We embed this truncated Gaussian-Mixture model in a Variational Autoencoder framework to obtain a general joint clustering and outlier detection approach, tGM-VAE. When applied to synthetic data with known ground-truth, tGM-VAE is more accurate in clustering connectivity patterns than existing approaches. On the rs-fMRI of 593 healthy adolescents, tGM-VAE identifies meaningful major connectivity states. The dwell time of these states significantly correlates with age.

1 Introduction

Functional connectivity refers to the functionally integrated relationship between spatially separated brain regions [1]. Recent work revealed that functional connectivity exhibits meaningful variations within the time series captured by resting-state fMRI [2,3]. As a consequence, a considerable amount of work has been directed to quantify dynamic functional connectivity. A popular way of quantification [2,4,5] is to first categorize the time-varying connectivity patterns of a subject or a population into several states. Then, the dwell time of each state and the transition probabilities across states are used to characterize functional dynamics and perform group analyses [4].

Most existing works identify *major connectivity states* (commonly observed states) by grouping the dynamic connectivity patterns into a few clusters [2,4,5]. However, some studies have indicated that there exist many *minor states* containing rare connectivity patterns that persist shortly (< 1% occupancy rate) [6]. These minor states often provide little merit to analysis because they may correspond to random individual brain variation, inaccurate connectivity pattern computed during state transitions or rs-fMRI noise. Instead of merging minor states into the major ones [4,5], recent work suggests to disentangle minor from major states by modeling an infinite number of clusters [6,7]. However, connectivity patterns in minor states may correspond to pure noise, so grouping them into clusters is not meaningful. In this paper, we address these concerns by developing a statistical framework

Motivated by the truncated stick-breaking representation of Dirichlet processes [8,9], our approach is guided by a Dirichlet prior when separating dynamic connectivity patterns into major and minor states. We define dynamic connectivity patterns by computing correlation matrices associated with sliding windows. We assume that the correlation matrices belonging to major states are generated by a non-linear process from a low-dimensional latent space, where the latent representations follow a Gaussian-mixture distribution. We assume that the rest of the correlation matrices, which correspond to minor states, are generated from a uniform distribution in the original space. To determine the optimal parameters of our model, we derive the variational lower-bound of its log marginal probability and find the maximum of that lower-bound by optimizing a variational autoencoder. As a result, our method, tGM-VAE, simultaneously achieves clustering and outlier-detection: tGM-VAE clusters dynamic connectivity patterns associated with major states and treats the minor states as outliers. In this work, we first demonstrate that tGM-VAE achieves higher accuracy in defining major clusters and outliers compared to traditional Gaussian-mixture-based approaches when clustering synthetic data with ground-truth. We then report that, for 15k correlation matrices derived from rs-fMRI scans of 593 adolescents in the the National Consortium on Alcohol and Neurodevelopment in Adolescence (NCANDA), tGM-VAE identifies meaningful connectivity states and a significant effect of age on their mean dwell time.

In the following, we first review existing VAE-based clustering approaches in Section 2. We introduce in Section 3 the generative model of tGM-VAE, the variational lower bound of the resulting log marginal likelihood, and reformulate tGM-VAE into a joint clustering and outlier-detection approach. We present our synthetic experiments and clinical data analysis in Section 4.

2 Related Work

Traditional approaches to clustering connectivity patterns are mostly based on Gaussianmixture models [4,5]. These methods usually require fitting probability distributions in a high dimensional space, which is a challenging task. Moreover, it has been found that the underlying distributions of both fMRI measurements [6] and the derived correlation matrices [10] lie on a non-linear latent space. Therefore, modeling Gaussian-mixtures in the original space is suboptimal.

Generative models used in connection with neural-networks, such as VAEs, have recently attracted much attention for their capability of modeling latent representations of the data [11]. In VAE, the encoder approximates the intractable posterior distribution of the latent representation and the decoder aims to reconstruct the observation based on its latent representation. While traditional VAE assumes that latent variables follow a single Gaussian prior, recent works adopt mixture models in the latent space for semi-supervised learning [12] and clustering [13]. Dilokthanaku et al. [13] construct a two-level latent space that allows for a multi-modal prior of latent variables, but this model exhibits over-regularization

effects that require specific optimization procedures. Jiang et al. [14] explicitly define a generative process based on a mixture of Gaussians in the latent space, which achieves better clustering performance. Our tGM-VAE model is built upon a generative model similar to [14] to capture major states but also includes a non-informative distribution for modeling minor states.

Besides the above approaches for modeling fixed number of clusters, Bayesian nonparametric models have been adopted to model an infinite number of clusters. The semisupervised approach proposed in [15] uses multiple VAEs as a proxy of Gaussian-mixture models and automatically determines the number of VAEs by maximizing the reconstruction capability for the entire dataset. The stick-breaking construction [9] has also been adopted in VAE for semi-supervised classification, where the latent representation is a set of truncated categorical weights. While this approach is not intrinsically built for clustering, the truncation strategy motivates us to use the last category (remainder of the truncation) to capture all dynamic connectivity patterns that do not belong to major clusters. Contrary to the above two approaches, tGM-VAE only models the encoding/decoding process for major clusters and omits the latent representation for the remainder. This strategy is useful when the remainder corresponds to (a) minor clusters of no interest to analysis so modeling their latent presentations is redundant; (b) outliers whose latent representations are meaningless or do not form clusters.

3 Methods

3.1 The Generative Model

Let $\mathbf{X} = \{x_1, ..., x_N\}$ be a training dataset with *N* observations. Each x_i represents a dynamic connectivity pattern, e.g., the upper triangular part of an ROI-to-ROI correlation matrix derived from the rs-fMRI time series at a given sliding window [2]. We assume that each x_i belongs to a state, which, in our proposed generative process, is encoded by the categorical variable c_i . The first *K*-1 categories represent the major states and $c_i = K$ corresponds to the remainder (minor states). c_i is drawn from a categorical distribution $p_{\pi}(c_i) \sim \text{Cat}(\pi)$, where $\pi = [\pi^1, ..., \pi^K]$ belongs to the (K-1)-dimensional simplex and is generated from a Dirichlet prior with two parameters $p(\pi) \sim \text{Dir}(\alpha, ..., \alpha, \beta)$. By construction, a single parameter α controls the portion of the *K*-1 major clusters indifferently, and β separately controls for the portion of the remainder via a stick-breaking procedure Beta($(K-1) \cdot \alpha, \beta$) [8].

For simplicity, let c_i^k denote $c_i = k$. We assume that when c_i^k with k < K, x_i is generated from a latent representation z_i through a non-linear process modeled by a neural-network f with parameter θ . $p_{\theta}(\mathbf{x}_i | \mathbf{z}_i) \sim \mathcal{N}(f_{\theta}(\mathbf{z}_i), \sigma_x^2)$, where σ_x^2 is the fixed standard deviation of noise. We further assume z_i is drawn from a Gaussian distribution with mean μ_k and an identity covariance: $p_{\mu}(\mathbf{z}_i | c_i^k) \sim \mathcal{N}(\mu_k, \mathbf{I})$ with $\mu = \{\mu_k | k < K\}$. In other words, the marginal distribution of z_i follows a Gaussian mixture in the latent space. On the other hand, when $c_i = K$, we assume x_i is simply drawn from a uniform distribution in a unit domain ξ embedded in the original space containing all observations after normalization: $p(\mathbf{x}_i | c_i^K) \sim \mathcal{U}(\xi)$. Based on the above generative model parameters $\theta = \{\pi, \mu, \theta\}$, we have

3.2 Variational Lower Bound

Given the training dataset *X* and the two parameters $\{a, \beta\}$ of the Dirichlet prior, the generative model parameters Θ are determined by maximizing the marginal probability $p(\mathbf{X}, \pi | \mu, \Theta, a, \beta)$. Assuming i.i.d for each x_i , the log marginal probability can be written as:

graphical diagram of this model is given in Fig. 1a.

$$\log p(\mathbf{X}, \boldsymbol{\pi} | \boldsymbol{\mu}, \boldsymbol{\theta}, \boldsymbol{\alpha}, \boldsymbol{\beta}) = \sum_{i=1}^{N} \log p_{\boldsymbol{\Theta}}(\boldsymbol{x}_i) + \log p(\boldsymbol{\pi} | \boldsymbol{\alpha}, \boldsymbol{\beta})$$
(1)

In the above equation, the log likelihood $\log p_{\Theta}(x_i)$ can not be directly optimized, so variational inference is used to maximize its lower-bound. Typically, lowerbounds for graphical models are derived by approximating an intractable posterior $p(z_i, c_i|x_i)$ on the latent variables with a tractable function $q(z_i, c_i|x_i)$. Here we make the common mean-field assumption: $q(z_i, c_i|x_i) = q(z_i|x_i)q(c_i|x_i)$. When omitting the subscripts *i* to simplify notations, it reads:

$$\log p_{\Theta}(\mathbf{x}) = \log \left(p_{\Theta}(\mathbf{x}, c^{K}) + \sum_{k=1}^{K-1} \int_{\mathbf{z}} p_{\Theta}(\mathbf{x}, \mathbf{z}, c^{k}) \right)$$
(2)

$$= \log\left(p_{\Theta}(\mathbf{x}, c^{K})\frac{q(c^{K}|\mathbf{x})}{q(c^{K}|\mathbf{x})} + \sum_{k=1}^{K-1}\int_{\mathbf{z}}p_{\Theta}(\mathbf{x}, \mathbf{z}, c^{k})\frac{q(\mathbf{z}, c^{k}|\mathbf{x})}{q(\mathbf{z}, c^{k}|\mathbf{x})}\right)$$
(3)

$$= \log\left(q(c^{K}|\mathbf{x})\frac{p_{\mathbf{\pi}}(c^{K})}{q(c^{K}|\mathbf{x})} + \sum_{k=1}^{K-1} q(c^{k}|\mathbf{x})\mathbb{E}_{q(\mathbf{z}|\mathbf{x})}\left[\frac{p_{\Theta}(\mathbf{x}, \mathbf{z}, c^{k})}{q(\mathbf{z}, c^{k}|\mathbf{x})}\right]\right)$$
(4)

$$\geq q(c^{K}|\mathbf{x})\log\frac{\pi^{K}}{q(c^{K}|\mathbf{x})} + \sum_{k=1}^{K-1} q(c^{k}|\mathbf{x})\mathbb{E}_{q(\mathbf{z}|\mathbf{x})}\left[\log\frac{p_{\Theta}(\mathbf{x}, \mathbf{z}, c^{k})}{q(\mathbf{z}, c^{k}|\mathbf{x})}\right]$$
(5)

$$= \sum_{k=1}^{K} q(c^{k} | \mathbf{x}) \log \frac{\pi^{k}}{q(c^{k} | \mathbf{x})} + \sum_{k=1}^{K-1} q(c^{k} | \mathbf{x}) \mathscr{L}^{k}(\mathbf{x}), \text{ where }$$
(6)

$$\mathscr{L}^{k}(\boldsymbol{x}):=\mathbb{E}_{q(\boldsymbol{z}\mid\boldsymbol{x})}[p_{\theta}(\boldsymbol{x}\mid\boldsymbol{z})]-D_{\mathrm{KL}}(q(\boldsymbol{z}\mid\boldsymbol{x})\parallel p_{\mu}(\boldsymbol{z}\mid\boldsymbol{c}^{k}))$$
(7)

where D_{KL} denotes the KL divergence between two probability distributions.

Interpretation of the Lower Bound— \mathscr{D}^k corresponds to the formulation of the traditional single-Gaussian VAE [11] with respect to the k^{th} cluster. Specifically, $\mathbb{E}_{q(z|x)}[p_{\theta}(x|z)]$ encourages the decoded reconstruction of the latent variable to resemble the

observation. The D_{KL} term is commonly interpreted as a regularizer encouraging the approximate posterior q(z|x) to resemble the cluster-specific Gaussian prior $p_{\mu}(z|c^k)$.

The right term of the lower-bound (Eq. 6) sums the losses of single-Gaussian VAEs over the K-1 major clusters and weighs them by cluster-assignment probability q(c|x). Maximizing this term improves the encoding/decoding capability for patterns in major states while keeping their latent variables to form clusters. The left term of the lower-bound corresponds to the KL-divergence between q(c|x) and Cat(π) and encourages the posterior categorical distribution to approximate the categorical prior. It is important to note that latent representations are only modeled for the K-1 clusters but not for the remainder. The portion of the remainder is controlled by the left term of the lower-bound.

3.3 Reformulation

In this section, we reformulate our model to demonstrate, by re-organizing the lower-bound of Eq. 6, that tGM-VAE can be interpreted as a joint outlier-detection and clustering framework. Given the generative process described in Section 3.1, the categorical variable *c* can be constructed by first differentiating the major clusters from the remainder. Let *b* denote a Bernoulli variable generated by $p(b) \sim \text{Ber}(\gamma)$, where $\gamma \in [0,1]$ defines the portion of the remainder. When b^0 (b = 0 for major clusters), a cluster assignment variable *m* is drawn from a categorical distribution $p_{\psi}(m) \sim \text{Cat}(\psi)$, where $\psi = [\psi^1, ..., \psi^{K-1}]$ follows Dir(*a*,..., *a*). This construction also involves two parameters, {*a*, γ }. The graphical diagram of this model is given in Fig. 1b.

For posterior inference, different q functions are constructed for the reformulated generative process. Let q(b|x) denote the approximate posterior of assigning x to either major clusters or the remainder and let $q(m|x, b^0)$ denote the major cluster assignment given b^0 . Then q(c|x) and π in Section 3.2 become

$$q(c^{k} | \mathbf{x}) = q(m^{k} | \mathbf{x}, b^{0})q(b^{0} | \mathbf{x}) \text{ for } k < K, \text{ and } q(c^{K} | \mathbf{x}) = q(b^{1} | \mathbf{x})$$
(8)

$$\pi^{k} = \psi^{k}(1 - \gamma) \text{ for } k < K, \text{ and } \pi^{K} = \gamma$$
⁽⁹⁾

Replacing the terms in Eq. 6 with Eq. 8,9 leads to the following lower bound

$$\log p_{\Theta}(\mathbf{x}) \ge q(b^0 | \mathbf{x}) \mathscr{G}(\mathbf{x}, m) - D_{\mathrm{KL}}(q(b | \mathbf{x}) \| \mathrm{Ber}(\gamma)), \tag{10}$$

where $\mathscr{G}(\mathbf{x}, m) = \sum_{k=1}^{K-1} q(m^k | \mathbf{x}, b^0) \mathscr{L}^k(\mathbf{x}) - D_{\mathrm{KL}}(q(m | \mathbf{x}, b^0) \| \operatorname{Cat}(\psi))$ is exactly the formulation of Gaussian-mixture VAE with K-1 clusters [14]. From Eq. (10) we can see that $q(b^0 | \mathbf{x})$ essentially gives the probability of \mathbf{x} being an inlier. Data with high inlier-probability are then clustered by $\mathscr{G}(\mathbf{x}, m)$, while the right term in Eq. (10) regularizes the portion of outliers with parameter γ . In practice, we use an additional weight λ to balance the two types of losses in Eq. (10), a common practice in VAE frameworks [13,16].

3.4 Clustering Correlation Matrices

The design of our VAE network is based on the above inference procedure. More specifically, all the approximate posteriors are modeled by neural networks. Similar to the traditional VAE [11], q(z|x) is an encoder network (Fig. 1c red blocks) with parameters ϕ_{z^2} , which encodes the posterior as a multivariate Gaussian with an identity covariance $q(z|x) = \mathcal{N}(z; \tilde{\mu}, \mathbf{I})$. While allowing for a diagonal or full covariance are both reasonable practices, we simply rely on the non-linear neural network to capture the covariance structure, and we only use the mean to capture the clustering effects in the latent space. The encoder has 3 densely connected hidden layers with *tanh* activation. The dimensions of the 3 layers are (*D*, 16, 3), where *D* is the leading "power of two" that is smaller than the input dimension (e.g., D=64 for a 15×15 correlation matrix with 105 upper triangular elements). The decoder network $f_{\theta}(z)$ has an inverse structure as the encoder and uses MSE reconstruction loss. For the optimization of these two networks, the SGVB estimator and reparameterization trick are adopted [11].

Contrary to previous work [14,17], we also use neural networks to model the categorical posteriors q(b|x) and q(m|x) (Fig. 1c orange blocks). Their first two layers were shared from the encoder of q(z|x) and the last layer is densely connected with *soft-max* activation. This construction rigorously reflects the structure of the generative model described in Section 3.3 (Fig. 1b) and allows for two separate mechanisms for detecting outliers with q(b|x) and assigning clusters with q(m|x). By comparison, a single neural network for q(c|x) would be obtained from the model described in Section 3.1 (Fig. 1a), but this network would treat the clusters and outliers indifferently.

4 Results

tGM-VAE was first validated and compared to traditional clustering approaches based on synthetic experiments, where rs-fMRI series and time-varying correlation matrices were simulated according to a ground-truth state sequence. We measured, in particular, the accuracy of tGM-VAE in connectivity states estimation. Then, tGM-VAE was used to cluster 15k correlation matrices obtained from the rs-fMRI scans of 593 adolescents in the NCANDA study [18]. The relation between the age of a subject and the mean dwell time of the connectivity states was finally examined.

4.1 Synthetic Experiments

Data Simulation—We followed the simulation procedure presented in [2,6] by first generating a state sequence of 50000 time points associated with 10 connectivity states, among which 5 states were major states. The transition probability from the *i*th state to the *j*th state was set to $0.9\delta_{ij} + 0.1b_i^j$, where δ is the Kronecker Delta function, and $b_i = [b_i^1, ..., b_i^{10}]$ was randomly generated from Dir(10,..., 10, 1,..., 1). This process led to self-transition probabilities varying between 0.9 and 0.95, and cross-state transition probabilities between 1e-4 and 0.05. The mean dwell time of a state (average time that a state continuously persists before switching to another state) varied between 8 and 15 time points. The occupancy rate of a major state (percentage of a state occupying the sequence) varied between 8% to 30%,

and the total occupancy of the 5 minor states varied between 5% to 10%. These metrics are similar for real rs-fMRI data reported in [6].

Next, a connectivity pattern was simulated for each state. In the first experiment, we assumed that there were 15 regions of interest (ROI) in the brain, so each state was associated with a 15×15 matrix, known as the community matrix [6]. For the *i*th state, a 1D loading vector $u_i \in \mathbb{R}^{15}$ consisted of {1,-1,0} (representing positive/negative or no activation of each ROI) was randomly generated. Then, the *i*th community matrix was computed by $u_i u_i^T$ [2].

Afterwards, synthetic rs-fMRI signals at each time point were randomly sampled from a Gaussian distribution with the covariance being the state-specific community matrix at that time point. Gaussian noise of standard deviation 0.1 was further added to the synthetic rs-fMRI series. Finally, dynamic correlation matrices were generated using a sliding window of length 11. These different steps are summarized in Fig. 2.

Clustering Accuracy—tGM-VAE clustered the dynamic correlation matrices into 5 major states with the following parameter settings: $\gamma = 0.075$, $\beta = 1.1$, and $\lambda = 200$. These settings corresponded to an accurate estimate of the portion of the remainder (γ), a rather non-informative Dirichlet prior (β) and a strong regularization on the portion of the remainder (λ). Fig. 3 presents the 3D latent space associated with tGM-VAE. Only the 5 major states are displayed as the latent representations of the remainder were not modeled. We can observe that the latent representations were reasonably clustered by states, thanks to the Gaussian-mixture modeling in the latent space [9,13,14].

To associate the 5 estimated clusters with the 5 ground-truth major states, the correlation matrices in an estimated cluster were first averaged and linked to the closest community matrix with respect to the Frobenius norm. As there was no interest in differentiating minor connectivity states, the clustering accuracy was measured with respect to the 6 classes (5 clusters + remainder). tGM-VAE was compared with three other clustering approaches as indicated by Fig. 4. Both Gaussian-Mixture Model (GMM) and Gaussian-Mixture VAE (GM-VAE) clustered the entire dataset into 5 clusters (merging minor states into major ones); The non-parametric Dirichlet Process (DP) Gaussian-mixture approach modeled an infinite number of clusters, so the 5 largest clusters estimated by DP were considered major states and the rest was considered the remainder. The clustering accuracy of these approaches was 68.4% (GMM), 69.0% (DP), 74.8% (GM-VAE) and 78.5% (tGM-VAE). Fig. 3b shows the estimated state sequence produced by tGM-VAE (most accurate) and GMM (least accurate). We observe that the two VAE-based methods produced significantly improved clustering accuracy than the two traditional Gaussian-Mixture methods (GMM and DP). This improvement indicates that the modeling of latent representations and the associated non-linear generative processes as provided by the VAE framework were helpful in analyzing correlation matrices. Moreover, the truncation of tGM-VAE could accurately capture the minor states and provided 3.7% improvement over GM-VAE, whereas explicitly clustering minor states was a less effective strategy (DP only 0.6% improvement over GMM).

Next, the above comparison was repeated for different simulation settings (Fig. 4). To demonstrate that tGM-VAE can generalize to brain parcellations of different scales, the number of ROIs was varied between 10 and 50, which covered the typical range used in existing analyses of functional dynamics [2,4,6,7]. In all settings the two VAE-based approaches produced more accurate clustering, and tGM-VAE was the most accurate approach. This was also the case when the standard deviation of noise in synthetic rs-fMRI time series was varied between 0.05 to 1. Another important parameter (not relevant to clustering approaches) in the analysis of functional dynamics is the length of the sliding window for computing correlation matrices. Previous works often use a window size longer than the mean dwell time of connectivity states in order to reliably compute correlation values, but this strategy could potentially fail to differentiate dynamic connectivity patterns across neighboring states because the long window often covers multiple state transitions. While the analysis of window length is not the focus of the presented work, our experimental results (Fig. 4c) indicate that choosing a window size longer than the mean dwell time does not guarantee accurate clustering.

Note that the shallow neural networks tested here are a simplification choice and not a limitation of the method. Further exploration in the network structure would lead to better results for tGM-VAE. For instance, setting the dimension of latent space larger than 3 would produce higher accuracy for large correlation matrices (Fig. 4d).

4.2 The NCANDA Dataset—We applied tGM-VAE to the rs-fMRI data of 593 normal adolescents (age 12–21; 284 boys and 309 girls) from the NCANDA study [18] to investigate dynamic connectivity states in young brains. The rs-fMRI time series was preprocessed using the publicly available pipeline as described in the NCANDA study [18]. For each subject, functional time series were extracted from 45 cerebral regions (averaged bilaterally) as defined by the sri24 atlas [19]. Dynamic correlation matrices of size 45×45 were then derived for each subject based on a sliding-window approach [2] and improved by a linear shrinkage operation [20]. As mentioned, there is no consensus on the optimal length of the sliding-window. In the present work, we selected the length that produced the largest number of strong correlations (absolute value 0.5) to maximize the information contained in the training data. Our experiments suggest that the optimum was achieved at 10 time points (22s) regardless of the parcellation used to produce correlation matrices (Fig. 6). Afterwards, a total of 153587 matrices were derived for the entire cohort and clustered by tGM-VAE into 5 major states [4]. The dimension of the latent space was set to 6. Other parameters were set as in the synthetic experiments. Fig. 5 shows the mean correlation matrices associated with the 5 major states detected by tGM-VAE and visualizes their graph structures. These 5 states correspond to well-known functional networks: auditory network (State 1), limbic and thalamostriatal network (State 2), visual network (State 3), salience network (State 4) and the default mode network (State 5).

Based on the clustering results, the state sequence was recovered for each subject and the mean dwell time over all states was computed. A group analysis was then performed to investigate the aging effect on the mean dwell time. First, sex and scanner-type were removed as confounding factors from mean dwell time using regression analysis [18,21]. The residuals were then correlated with age, resulting in a significant positive correlation

(one-tailed *p*=.0006, Fig 6). This age-related increase of mean dwell time could also be observed when the analysis was repeated with the dimension of latent space varying between 3 to 7. These results essentially indicate each connectivity state tends to persist longer in older adolescents, which converges with current concept of neurodevelopment that variation of dynamic functional connectivity declines with age [22].

5 Conclusion

In this paper, we have presented a novel joint clustering and outlier-detection approach to identify functional connectivity states. Our model, tGM-VAE, introduces for the first time a truncated Gaussian-mixture model in the variational autoencoder framework. This approach allows us to cluster data corrupted by noise, outliers and minor clusters of no interest to analysis. We used tGM-VAE to extract major functional connectivity states from resting-state fMRI scans and characterize their dynamics. We showed that modeling latent representations of correlation matrices improves clustering accuracy compared to traditional Gaussian-mixture approaches and that our truncation strategy is useful in disentangling minor and major connectivity states. In the future, we will expand our framework to improve the modeling of state transitions.

References

- Buckner R, Krienen F, Yeo B: Opportunities and limitations of intrinsic functional connectivity MRI. Nat Neurosci. 16(7) (2013) 832–837 [PubMed: 23799476]
- 2. Allen E, Damaraju E, Plis S, Erhardt E, Eichele T, Calhoun V: Tracking whole-brain connectivity dynamics in the resting state. Cerebral Cortex 24(3) (2014) 663–676 [PubMed: 23146964]
- Zalesky A, Fornito A, Cocchi L, Gollo LL, Breakspear M: Time-resolved resting-state brain networks. PNAS 111(28) (2014) 10341–10346 [PubMed: 24982140]
- Damarajua E, Allen E, Belgerc A, et al.: Dynamic functional connectivity analysis reveals transient states of dysconnectivity in schizophrenia. NeuroImage: Clinical 5 (2014) 298–308 [PubMed: 25161896]
- Yu Q, Erhardt EB, Sui J, et al.: Assessing dynamic brain graphs of time-varying connectivity in fmri data: Application to healthy controls and patients with schizophrenia. NeuroImage 107 (2015) 345– 355 [PubMed: 25514514]
- Taghia J, Ryali S, et al.: Bayesian switching factor analysis for estimating time-varying functional connectivity in fmri. NeuroImage 155 (2017) 271–290 [PubMed: 28267626]
- Nielsen S, Madsen K, Røge R, Schmidt MN, Mørup M: Nonparametric modeling of dynamic functional connectivity in fmri data. In: Machine Learning and Interpretation in Neuroimaging Workshop (2015)
- Blei DM, Jordan MI: Variational inference for dirichlet process mixtures. Bayesian Analysis 1(1) (2017) 121–144
- 9. Nalisnick E, Smyth P: Stick-breaking variational autoencoders. In: ICLR. (2017)
- 10. Zhao Q, Kwon D, Pohl KM: A riemannian framework for longitudinal analysis of resting-state functional connectivity. In: MICCAI. (2018)
- 11. Kingma D, Welling M: Auto-encoding variational bayes. In: ICLR. (2013)
- 12. Kingma DP, Rezende DJ, Mohamed S, Welling M: Semi-supervised learning with deep generative models. In: NIPS. (2014)
- 13. Dilokthanakul N, Mediano PA, Garnelo M: Deep unsupervised clustering with gaussian mixture variational autoencoder. (2017)
- 14. Jiang Z, Zheng Y, Tan H, Tang B, Zhou H: Variational deep embedding: An unsupervised and generative approach to clustering. In: IJCAI. (2017)

- 15. Abbasnejad E, Dick AR, van den Hengel A: Infinite variational autoencoder for semi-supervised learning. In: CVPR. (2017)
- 16. Higgins I, Matthey L, Pal A, et al.: beta-vae: Learning basic visual concepts with a constrained variational framework. In: ICLR. (2017)
- 17. Ebbers J, et al.: Hidden markov model variational autoencoder for acoustic unit discovery. In: InterSpeech. (2017)
- Müller-Oehring E, Kwon D, Nagel B, Sullivan E, et al.: Influences of age, sex, and moderate alcohol drinking on the intrinsic functional architecture of adolescent brains. Cerebral Cortex 28(3) (2017) 1049–1063
- Rohlfing T, Zahr N, Sullivan E, Pfefferbaum A: The sri24 multichannel atlas of normal adult human brain structure. Hum Brain Mapp. 31(5) (2014) 798–819
- 20. Chen Y, Wiesel A, Eldar YC, Hero AO: Shrinkage algorithms for mmse covariance estimation. IEEE Trans. on Sign. Proc 58(10) (2010)
- Pfefferbaum A, Kwon D, et al.: Altered brain developmental trajectories in adolescents after initiating drinking. American Journal of Psychiatry. 175(4) (2017) 370–380 [PubMed: 29084454]
- 22. Chen Y, Wang W, et al.: Age-related decline in the variation of dynamic functional connectivity: A resting state analysis. Front Aging Neurosci. 9(23) (2017)







(a) Bayesian model associated with tGM-VAE. Solid lines denote the generative model and dashed lines the inference model. Gray nodes denote observed variables and given parameters. (b) Reformulated model. (c) Neural network implemented from (b).







Fig.3.

(a) Latent representations of correlation matrices computed by tGM-VAE color-coded by ground-truth states (left) and estimated states (right). (b) State sequences estimated by tGM-VAE and GMM overlaid with the ground-truth sequence.



Fig.4.

Clustering accuracy scores measured on synthetic data by varying (a) noise level; (b) size of correlation matrix; (c) sliding window length. (d) tGM-VAE accuracy as a function of latent space dimension.



Fig.5.

Functional connectivity patterns of 5 major states derived from the NCANDA rs-fMRI data. Top: mean correlation matrices; Bottom: Graph visualization of the mean correlation matrices. Node color corresponds to lobe names. Node size corresponds to sum of positive correlations associated with that node. White edges correspond to correlations 0.25 and black edges -0.25. Edge thickness corresponds to absolute value of correlation.



Fig.6.

Left: The number of strong correlations (absolute value 0.5) depends on sliding window length but not on the number of ROIs in a parcellation. Right: For the NCANDA cohort, aging effect in the mean dwell time corrected for sex and scanner.