Identification of functional clusters in the striatum using infinite relational modeling

Kasper Winther Andersen[1,2], Kristoffer H. Madsen[1,2], Hartwig Siebner[2], Lars Kai Hansen[1], and Morten Mørup[1]

[1] DTU Informatics, Technical University of Denmark
[2] Danish Research Centre for Magnetic Resonance, Copenhagen University Hospital Hvidovre
kwjo@imm.dtu.dk, stoffer@drcmr.dk, hartwig.siebner@drcmr.dk, lkh@imm.dtu.dk, mm@imm.dtu.dk

Abstract. In this paper we investigate how the Infinite Relational Model can be used to infer functional groupings of the human striatum using resting state fMRI data from 30 healthy subjects. The Infinite Relational Model is a non-parametric Bayesian method for infering community structure in complex networks. We visualize the solution found by performing evidence accumulation clustering on the maximum a posterior solutions found in 100 runs of the sampling scheme. The striatal groupings found are symmetric between hemispheres indicating that the model is able to group voxels across hemispheres, which are involved in the same neural computations. The reproducibility of the groupings found are assessed by calculating mutual information between half splits of the subject sample for various hyperparameter values. Finally, the model's ability to predict unobserved links is assessed by randomly treating links and non-links in the graphs as missing. We find that the model is performing well above chance for all subjects.

Keywords: complex network, graph theory, infinite relational model, basal ganglia, striatum

1 Introduction

Recently, graph theoretical network modeling has gained a lot of attention in neuroimaging, for reviews see, e.g, [3, 15]. Both functional networks (using modalities such as fMRI, EEG, and MEG) and anatomical brain networks (using DWI) have been analyzed using complex network methods. These studies cover both studies of the healthy brain as well as a wide range of neuropsychiatric and neurologic disorders [16]. In this work we use the Infinite Relational Model (IRM) [8, 17] to infer functional groupings of the human striatum. The IRM model is a non-parametric Bayesian network model, which assigns nodes into non-overlapping groups. The probability of a link between two nodes is determined by the groups the nodes are assigned to. During inference the number of groups and the group assignments are inferred, while the group link probabilities can be integrated out

of the model and are therefore not determined during inference. These probabilities are easily calculated afterwards given the group assignments. The IRM allows analysis of multi-graph networks and thus provides a natural framework for analyzing multiple subjects at once as demonstrated in [13].

The basal ganglia (BG) process information from the cerebral cortex in segregated parallel cortico-BG-thalamocortical loops [1]. The BG are involved in the adaptation of complex goal related behaviors [4, 6] and play a key role in the pathophysiology of many neurological (e.g., Parkinsons disease) and psychiatric (e.g., schizophrenia) disorders [11, 14]. The caudate nucleus and putamen (i.e., dorsal striatum) are the main input structures of the BG receiving topographically organized inputs from the cortex. Striatal sub-territory receives specific cortical inputs via corticostriatal feed-forward projections originating from largely segregated cortical input zones [6, 12]. The BG anatomy and function is largely symmetric between the two hemispheres.

2 Methods

2.1 Data

Resting state functional magnetic resonance imaging (rs-fMRI) data from $N = 30$ healthy controls was recorded for 20 min (482 volumes) per subject. The first two volumes were discarded to account for T1 equilibrium effects, the remaining 480 volumes were realigned to the time-series mean and spatially normalized to the MNI template using SPM. Nuisance effects related to residual movement or physiological effects were removed using a linear filter comprised of 24 motion related and a total of 64 physiological effects including cardiac, respiratory, respiration volume over time, and time series from left and right hemispheres CSF and white matter voxels. An anatomical mask consisting of the caudate nucleus and putamen, which was made in WFU PickAtlas [10] using the Talairach Daemon atlas [9], was used to extract the time series of the $J = 825$ voxels from all subjects.

The network graph representing functional connectivity in subject n is represented by the $[J \times J]$ adjacency matrix $\mathbf{A}^{(n)}$. Each graph is then composed of J nodes and $\mathbf{A}^{(n)}(i, j)$ is 1 if a link is present between voxels i and j and 0 elsewhere. $\mathbf{A}^{(n)}$ was obtained for each subject by first calculating the upper triangular part of the Pearson correlation matrix and then thresholding the matrix to include the highest 5000 positive correlations.

2.2 The Infinite Relational Model

Following the notation in [13] the IRM generative model can be written as:

$$\begin{aligned}
\mathbf{Z}|\alpha &\sim \mathrm{CRP}(\alpha) \\
\boldsymbol{\rho}^{(n)}|\beta &\sim \mathrm{Beta}(\beta, \beta) \\
\mathbf{A}^{(n)}(i, j)|\mathbf{Z}, \boldsymbol{\rho}^{(n)} &\sim \mathrm{Bernoulli}(\mathbf{z}_{i_r} \boldsymbol{\rho}^{(n)} \mathbf{z}_{j_r}^{\top}).
\end{aligned}$$

As such, the probability of a link between two voxels is determined by the groups in which the voxels are members of. $\boldsymbol{\rho}^{(n)}$ is the subject specific group link probability matrix and defines the probability of links between groups. $\mathbf{Z}$ is a $[J \times D]$ binary matrix indicating group membership for each voxel and is shared across all subjects. We use symmetric Beta functions with hyperparameter β as priors for the group link probabilities and the Chinese Restaurant Process (CRP) is used as prior for the voxel group assignments. By integrating $\boldsymbol{\rho}$ out the posterior can be written as:

$$\begin{aligned} P(\mathbf{A}^{(n)}|\mathbf{Z},\beta) &= \int P(\mathbf{A}^{(n)}|\mathbf{Z},\boldsymbol{\rho}^{(n)})P(\boldsymbol{\rho}^{(n)}|\beta)\mathrm{d}\boldsymbol{\rho}^{(n)} \\ &= \prod_{a\geq b} \frac{\mathrm{Beta}(\mathbf{M}_{+}^{(n)}(a,b)+\beta, \mathbf{M}_{-}^{(n)}(a,b)+\beta)}{\mathrm{Beta}(\beta,\beta)}, \end{aligned}$$

where $\mathbf{M}_{+}^{(n)}(a,b) = (1-\frac{1}{2}\delta_{a,b})\mathbf{z}_a^{\top}(\mathbf{A}^{(n)}+\mathbf{A}^{(n)^{\top}})\mathbf{z}_b$ is the number of links and $\mathbf{M}_{-}^{(n)}(a,b) = (1-\frac{1}{2}\delta_{a,b})\mathbf{z}_a^{\top}(\boldsymbol{e}\boldsymbol{e}^{\top}-\boldsymbol{I})\mathbf{z}_b - \mathbf{M}_{+}^{(n)}(a,b)$ is the number of non-links between group a and b. $\boldsymbol{e}$ is a vector of length J with ones in all entries. The subjects' adjacency matrices are assumed independent thus their joint distribution is:

$$P(\mathbf{A}^{(1)},...,\mathbf{A}^{(N)}|\mathbf{Z},\beta) = \prod_{n}\prod_{a\geq b} \frac{\mathrm{Beta}(\mathbf{M}_{+}^{(n)}(a,b)+\beta, \mathbf{M}_{-}^{(n)}(a,b)+\beta)}{\mathrm{Beta}(\beta,\beta)}.$$

Using Bayes' theorem the posterior likelihood can be found as:

$$\begin{aligned} &P(\mathbf{Z}|\mathbf{A}^{(1)},...,\mathbf{A}^{(N)},\beta,\alpha) \propto P(\mathbf{A}^{(n)}|\mathbf{Z},\beta,)P(\mathbf{Z}|\alpha) = \\ &\left[\prod_{n}\prod_{a\geq b} \frac{\mathrm{Beta}(\mathbf{M}_{+}^{(n)}(a,b)+\beta, \mathbf{M}_{-}^{(n)}(a,b)+\beta)}{\mathrm{Beta}(\beta,\beta)}\right]\left[\alpha^{D}\frac{\Gamma(\alpha)}{\Gamma(J+\alpha)}\prod_{a}\Gamma(n_a)\right]. \end{aligned}$$

For model inference we use a Gibbs sampling scheme in combination with split-merge sampling [7, 8, 13], requiring the posterior likelihood for a node's assignment given the assignment of the remaining nodes:

$$\begin{aligned} &P(\mathbf{Z}(i,a)=1|\mathbf{Z}\backslash\mathbf{z}_{i_r},\mathbf{A}^{(1)},...,\mathbf{A}^{(N)}) \\ &\propto \begin{cases} m_a \prod_n \prod_b \frac{\mathrm{Beta}(\mathbf{M}_{+}^{(n)}(a,b)+\beta, \mathbf{M}_{-}^{(n)}(a,b)+\beta)}{\mathrm{Beta}(\beta,\beta)} & \text{if } m_a > 0 \\ \alpha \prod_n \prod_b \frac{\mathrm{Beta}(\mathbf{M}_{+}^{(n)}(a,b)+\beta, \mathbf{M}_{-}^{(n)}(a,b)+\beta)}{\mathrm{Beta}(\beta,\beta)} & \text{otherwise.} \end{cases} \end{aligned}$$

$m_a = \sum_{j\neq i}\mathbf{Z}(j,a)$ is the size of the a^{th} functional group disregarding the assignment of the i^{th} node. This posterior likelihood can be evaluated efficiently since we only need to compute $\mathbf{M}_{+}^{(n)}$ and $\mathbf{M}_{-}^{(n)}$ and evaluate the Beta function for entries affected by the considered assignment change.

3 Results and discussion

3.1 Group membership visualization

The Gibbs samling result in a posterior distribution of group assignments which makes visualization hard. Thus, here we use the Evidence Accumulation Clustering (EAC) framework [5] to summarize and visualize the MAP solutions from $r = 100$ runs, each ran for 1000 iterations and with α and β fixed to 1. From the MAP solutions we generated the voxel by voxel co-occurrence matrix $\mathbf{C} = \frac{1}{r}\sum_{i=1}^{r}(\mathbf{Z}^{(i)}\mathbf{Z}^{(i)\top} - \mathbf{I})$ where $\mathbf{C}(i,j)$ is the empirical probability that voxels i and j were observed in the same group. Using $\mathbf{C}$ agglomerative hierarchical clustering based on average linkage was performed. We compare the clustering found by IRM with a simpler approach where the mean adjacency matrix $\mathbf{S} = \frac{1}{N}\sum_{n=1}^{N}\mathbf{A}^{(n)}$ was used for agglomerative hierarchical clustering in place of $\mathbf{C}$.

The median number of groups found in the 100 runs was 16 (range 14-18). The average normalized mutual information (NMI) between each pair of the MAP solutions was 0.78 (std=0.04) and all pairs were highly significant ($p < 0.001$, as tested using permutation testing) indicating that the groupings found by IRM are stable across runs. The cophenetic correlation can be used to assess the dispersion of the co-occurrence matrix $\mathbf{C}$ and is therefore also a measure of the stability of IRM solutions across random initializations[2]. The cophenetic correlation was 0.98.

The left part of Fig. 1 shows the grouping found by EAC of the IRM MAP solutions. The voxel groups are rendered on the anatomical mask used (gray: putamen; purple: caudate nucleus) and shown next to the dendrogram. The groups are symmetric between hemispheres, i.e. same sub-territory in left and right striatum are grouped together, suggesting that the IRM is able to group voxels in bilateral hemispheres which are involved in the same neural computations. Using the hierarchical clustering one can assess relations between groups. Again, the model is able to extract meaningful anatomical information, since here the green part of the tree defines the putamen, the red part defines the ventral part of caudate head/tail, while the blue part defines the dorsal part of caudate head/tail.

The right part of Fig. 1 show the voxel groups and dendrogram found by average linkage clustering of $\mathbf{S}$. Besides cluster 10, which is a large cluster of both right and left caudate nucleus, the clustering of $\mathbf{S}$ does not show the symmetry of the IRM grouping but here the groups are in general lateralized to either left or right striatum and are in general groups of nearby voxels. The green part of the dendrogram defines right posterior putamen, blue is bilateral caudate nucleus and anterior putamen, while red reflects posterior putamen. The cophenetic correlation of the hierarchical clustering based on $\mathbf{S}$ was 0.67 indicating the dendrogram is representing less of the information in the data compared to the dendrogram of the IRM model.

3.2 Varying hyperparameters

To test the reproducibility of groupings found for various choices of hyperparameters we split the subject sample in half and ran the IRM on each subsample and calculated the mutual information (MI) and normalized mutual information (NMI) between the MAP solutions found from each subsample. This was repeated for 10 different splits for each hyperparameters value. Fig. 2(a) show the mean (std) log likelihood, MI, NMI and number of groups found when varying $\log_{10}(\alpha)$ from -15 to 15 keeping $\beta = 1$. Likewise, Fig. 2(b) shows the same when varying $\log_{10}(\beta)$ from -6 to 2 keeping $\alpha = 1$. The IRM is seen to be very robust for the choice of α, which controls the prior belief on the group distributions, where both the MI and NMI remain constant over the wide α-range investigated. The choice of β, which controls the prior belief on group link probabilities, have a stronger influence of the MI and NMI. Here, NMI and log likelihood peaks at $\beta = 0.1$. At $\beta = 100$ perfect repetition is found between the two splits (NMI=1) although only 2 groups are found, so the information maintained in this grouping is low as reflected by the low MI.

3.3 Link prediction

To evaluate the model's ability to predict unobserved links we treated at random 2.5% of the links and an equivalent number of non-links in each adjacency matrix as missing. The area under the curve (AUC) of the receiver-operator characteristics was used as performance measure to evaluate how well the model was able to predict these unobserved links. Fig. 3 shows the mean (std) AUC for each of the 30 subjects of 100 model runs. Across subjects the mean (std) AUC was found to be 0.83 (0.06) which is well above chance for all subjects.

4 Conclusion

In this work we used the Infinite Relational Model to infer functional groupings in the human striatum. We show that the groups found are symmetric between hemispheres indicating that the IRM is able to find groups of voxels which are involved in the same neural computations. We evaluate the model's reproducibility by splitting the subject sample in half and compute mutual information and normalized mutual information between splits when varying the hyperparameters of the model. We show that the solutions are very robust to the choice of α, which controls the grouping distribution, while the solutions are more sensitive to the choice of β, which controls the prior belief in group link probabilities. Further, we show that the model is able to predict missing links in the graph well above chance.

References

1. Alexander, G.E., Crutcher, M.D., DeLong, M.R.: Basal ganglia-thalamocortical circuits: Parallel substrates for motor, oculomotor, "prefrontal" and "limbic" functions. Progress in Brain Research 85, 119–146 (1991)

2. Brunet, J.P., Tamayo, P., Golub, T.R., Mesirov, J.P.: Metagenes and molecular pattern discovery using matrix factorization. Proceedings of the National Academy of Sciences of the United States of America 101(12), 4164–9 (2004)
3. Bullmore, E.T., Bassett, D.S.: Brain graphs: graphical models of the human brain connectome. Annual review of clinical psychology 7, 113–40 (Apr 2011)
4. Doyon, J., Bellec, P., Amsel, R., Penhune, V., Monchi, O., Carrier, J., Lehéricy, S., Benali, H.: Contributions of the basal ganglia and functionally related brain structures to motor learning. Behavioural brain research 199(1), 61–75 (2009)
5. Fred, A.L.N., Jain, A.K.: Combining multiple clusterings using evidence accumulation. IEEE transactions on pattern analysis and machine intelligence 27(6), 835–50 (2005)
6. Haber, S.: The primate basal ganglia: parallel and integrative networks. Journal of Chemical Neuroanatomy 26(4), 317–330 (2003)
7. Jain, S., Neal, R.M.: A Split-Merge Markov chain Monte Carlo Procedure for the Dirichlet Process Mixture Model. Journal of Computational and Graphical Statistics 13(1), 158–182 (2004)
8. Kemp, C., Tenenbaum, J., Griffiths, T., Yamada, T., Ueda, N.: Learning systems of concepts with an infinite relational model. In: Proceedings of the National Conference on Artificial Intelligence. vol. 21, pp. 381–388. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999 (2006)
9. Lancaster, J.L., Woldorff, M.G., Parsons, L.M., Liotti, M., Freitas, C.S., Rainey, L., Kochunov, P.V., Nickerson, D., Mikiten, S.A., Fox, P.T.: Automated Talairach atlas labels for functional brain mapping. Human brain mapping 10(3), 120–31 (2000)
10. Maldjian, J.A., Laurienti, P.J., Kraft, R.A., Burdette, J.H.: An automated method for neuroanatomic and cytoarchitectonic atlas-based interrogation of fMRI data sets. NeuroImage 19(3), 1233–1239 (2003)
11. Mehler-Wex, C., Riederer, P., Gerlach, M.: Dopaminergic dysbalance in distinct basal ganglia neurocircuits: implications for the pathophysiology of parkinsons disease, schizophrenia and attention deficit hyperactivity disorder. Neurotoxicity research 10, 167–179 (2006)
12. Middleton, F.A., Strick, P.L.: Basal-ganglia 'projections' to the prefrontal cortex of the primate. Cerebral cortex (New York, N.Y. : 1991) 12(9), 926–35 (Sep 2002)
13. Mørup, M., Madsen, K., Dogonowski, A.M., Siebner, H., Hansen, L.: Infinite relational modeling of functional connectivity in resting state fmri. In: Advances in Neural Information Processing Systems 23, pp. 1750–1758 (2010)
14. Obeso, J.a., Rodríguez-Oroz, M.C., Benitez-Temino, B., Blesa, F.J., Guridi, J., Marin, C., Rodriguez, M.: Functional organization of the basal ganglia: therapeutic implications for Parkinson's disease. Movement disorders : official journal of the Movement Disorder Society 23 Suppl 3, S548–59 (2008)
15. Sporns, O.: The human connectome: a complex network. Annals of the New York Academy of Sciences 1224(1), 109–25 (Apr 2011)
16. Xia, M., He, Y.: Magnetic Resonance Imaging and Graph Theoretical Analysis of Complex Brain Networks in Neuropsychiatric Disorders. Brain Connectivity 1(5), 349–365 (2011)
17. Xu, Z., Tresp, V., Yu, K., Kriegel, H.: Infinite hidden relational models. In Proceedings of the 22nd International Conference on Uncertainity in Artificial Intelligence (2006)

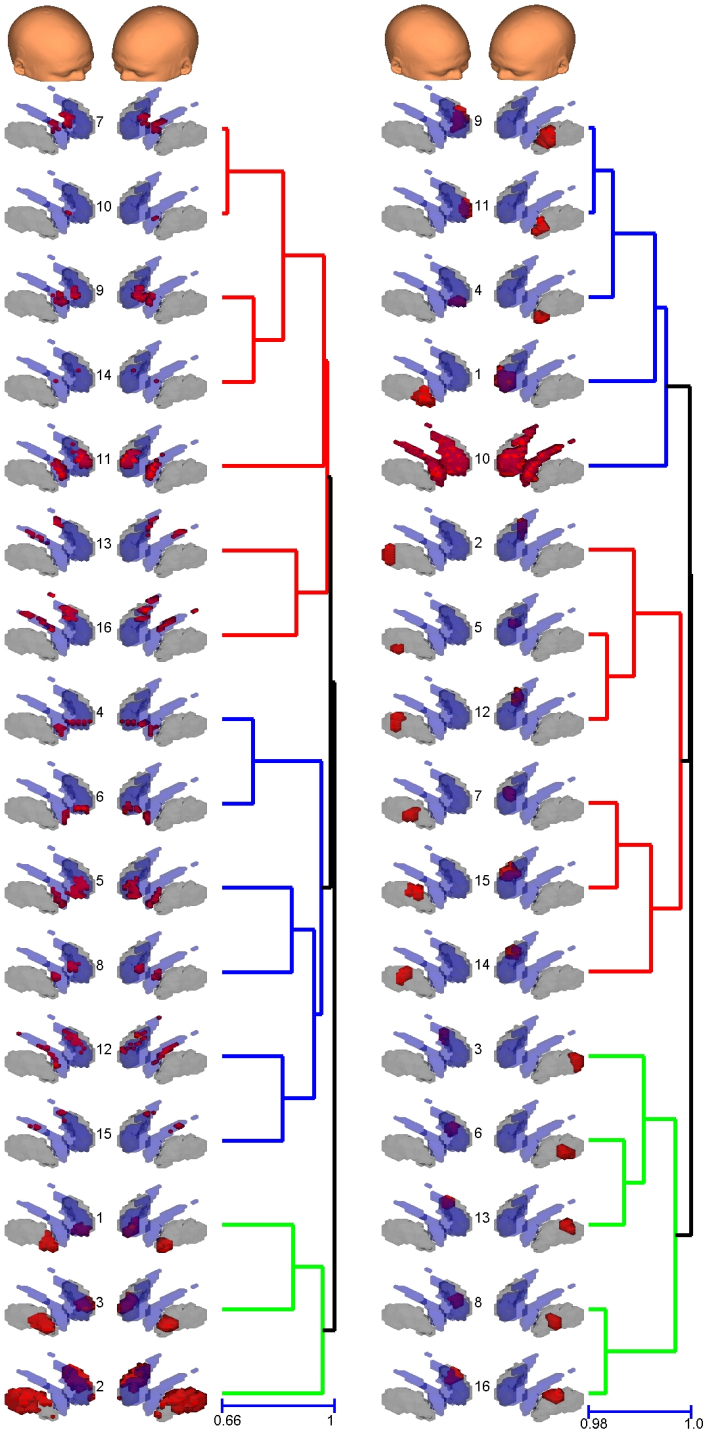

Fig. 1. Left; agglomerative hierarchical clustering of the co-occurence matrix **C** of the MAP solutions found from 100 starts of the IRM inference. Right; for comparison a simpler approach where the average adjacency matrix **S** was used for agglomerative hierarchical clustering in place of **C**.

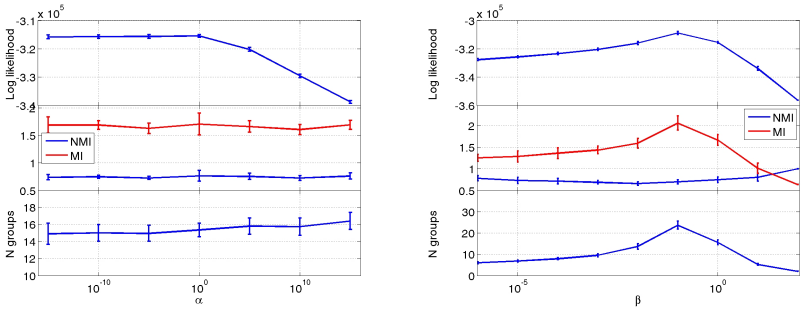

(a) Varying α (b) Varying β

Fig. 2. Log likelihood (top panel), mutual information (MI) and normalized mutual information (NMI) (middle panel) and number of groups (lower panel) between the MAP solutions found by randomly splitting the subjects in two half. Mean and standard deviations are shown for 10 splits for each α (a) and β (b). The IRM is seen to be very robust to the choice of α where both the MI and NMI remain constant over the wide α-range investigated. The choice of β have a stronger influence of the MI and NMI. Here, NMI and log likelihood peaks at $\beta = 0.1$. At $\beta = 100$ perfect repetition is found between the two splits (NMI=1) although here the number of groups found is 2, so the information maintained in this grouping is low as reflected by the low MI.

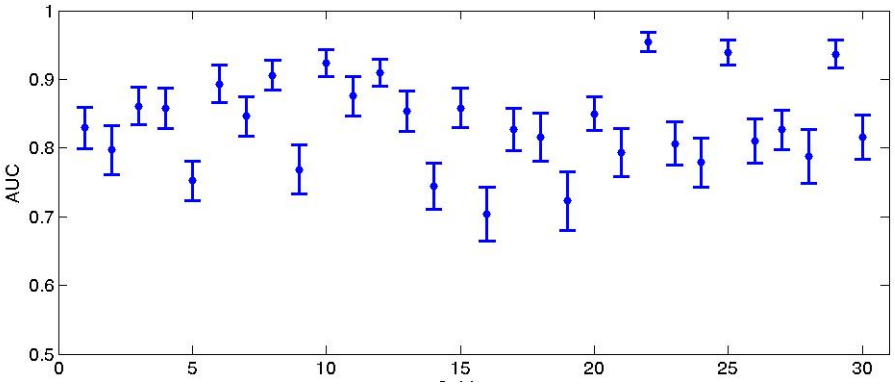

Fig. 3. The area under the curve (AUC) of the receiver-operator characteristics for the link prediction for each of the 30 subjects. The mean (std) across subjects were 0.83 (0.06).