

Discovering Voxel-Level Functional Connectivity Between Cortical Regions

Christopher Baldassano¹, Marius Cătălin Iordan¹, Diane M. Beck², and
Li Fei-Fei¹

¹ Department of Computer Science, Stanford University, Stanford, CA, USA

² Beckman Institute and Department of Psychology, University of Illinois at
Urbana-Champaign, Urbana, IL, USA

Abstract. Functional connectivity patterns are known to exist in the human brain at the millimeter scale, but the standard fMRI connectivity measure only computes functional correlations at a coarse level. We present a method for identifying fine-grained functional connectivity between any two brain regions by simultaneously learning voxel-level connectivity maps over both regions. We show how to formulate this problem as a constrained least-squares optimization, which can be solved using a trust region approach. Our method can automatically discover multiple correspondences between distinct voxel clusters in the two regions, even when these clusters have correlated timecourses. We validate our method by identifying a known division in the lateral occipital complex using only functional connectivity, thus demonstrating that we can successfully learn subregion connectivity structures from a small amount of training data.

1 Introduction

Untangling the connectivity structure of the incredibly complex networks underlying cognition is one of the key goals in modern neuroscience research. Functional magnetic resonance imaging (fMRI) has become an invaluable tool for understanding macro-scale connectivity networks, since it can sample brain activity at the millimeter scale across the entire cortex simultaneously. The standard approach for measuring functional connectivity using fMRI, however, ignores voxel-level information present in the fMRI signal, and can only investigate connectivity between large, predefined regions of interest (ROIs). This approach is unable to answer a more interesting question: “which specific voxels in each of these two regions are most strongly connected?”

We present a method for identifying voxel-level functional connectivity maps between any two regions, which can automatically detect multiple correspondences between subregions (see Fig. 1). Unlike previous approaches, our method has no limitations on the sizes of the regions and can find multiple solutions even when their timecourses are correlated. Our formulation also makes no assumptions about the anatomical locations of the connected clusters, making it much more widely applicable than previous methods.




Fig. 1: Functional connectivity methods. The standard measurement of functional connectivity between two regions averages together all voxels in each ROI, ignoring voxel-level connectivity differences. Recent CCRF/FF work produces a separate map over one region for each voxel in a seed region. Our method can learn connectivity structures over both ROIs simultaneously, and automatically identifies multiple connectivities between different sets of voxels.

Most fMRI studies measure functional connectivity between regions by simply computing the correlation between their mean timecourses [19], ignoring any connectivity differences at the subregion level. Methods that investigate *subregion* connectivity typically formulate the problem as learning “cortico-cortical receptive fields” (CCRFs) [12, 13] or “functional fingerprints” (FFs) [14]; one of the two regions is chosen as a seed region, and for each individual voxel (or cluster of voxels) in the seed region, these methods identify voxels in the second region that are most strongly functionally connected to the seed voxel. Seed voxels may then be grouped based on their connectivity signatures [5, 14, 16]. The connectivity weights are generally computed using linear regression [16, 22] or correlation [14], but have also been learned using support vector regression [13], mutual information mapping [4], and spatially-regularized regression [1]. All of these methods require treating the two regions asymmetrically, and cannot produce continuous-valued maps over both input regions, reducing their sensitivity to fine-grained connectivity differences.

There have been several proposed approaches for learning continuous-valued maps at the subregion level in both regions simultaneously. Some studies have performed multiple linear regression, measuring connectivity between multiple seed clusters in one region and all voxels in a second region [15, 20]. However, these methods require downsampling the seed region to a small number of clusters based on prior anatomical knowledge and manually comparing the connectivity maps for each seed cluster. Canonical correlation analysis can be used to learn voxel-level maps over both regions [8], but has a number of limitations. Multiple correspondences between subregions can only be identified if their timecourses are not (positively) correlated (as shown in section 3.2, this assumption is not typically valid) and the number of voxels in each region must be smaller than the number of timepoints (limiting the datasets to which this method can be applied). Our method can learn voxel-level maps for regions of any size, and can identify distinct subregions even if their timecourses are correlated, making it widely applicable for investigating connectivity between any arbitrary pair of ROIs.

2 Functional Connectivity as an Optimization Problem

2.1 Traditional Method

Functional connectivity between two ROIs is often measured by computing the Pearson product-moment correlation coefficient (r value) between the mean timecourses of the two ROIs [19]. Pairs of ROIs with a high r^2 value are then said to be strongly functionally connected. In order to generalize this method, we first recast it as an equivalent linear regression problem, in which we measure the similarity of the two mean timecourses, up to a scaling factor w :

$$\underset{w}{\text{minimize}} \quad \|w \cdot \text{mean}_v(\mathbf{A}^1) - \text{mean}_v(\mathbf{A}^2)\|_2^2$$

where \mathbf{A}^1 and \mathbf{A}^2 are the (# voxels x # timepoints) data matrices from two ROIs, and $\text{mean}_v(\cdot)$ denotes an average across voxels. The r^2 value is then equivalent to the fraction of variance in $\text{mean}_v(\mathbf{A}^2)$ explained by our predictor $w \cdot \text{mean}_v(\mathbf{A}^1)$ [21]. We can therefore think of the traditional correlation method as minimizing an objective of the form $\|\mathbf{a}^1 \mathbf{A}^1 - \mathbf{a}^2 \mathbf{A}^2\|_2^2$, in which \mathbf{a}^1 and \mathbf{a}^2 are constant-weight *connectivity maps*.

2.2 Voxel-Level Method

Although the traditional problem can describe functional connectivity at the coarse scale of ROIs, it makes the simplistic assumption that all voxels within each region have the same functional connectivity properties. This prevents us from using the traditional method to explore connectivity differences at the voxel level, which are often of scientific interest [13–16, 20, 22]. To learn voxel-level connectivity weights, we would like to relax the constraints on both \mathbf{a}^1 and \mathbf{a}^2 and allow the connectivity maps to be nonconstant. Note that a CCRF/FF method would relax only one of these constraints, learning a connectivity map over only one of the regions.

As will be shown in section 3, simply allowing each voxel to be chosen independently can lead to severe overfitting on the small datasets typical of fMRI experiments. It is possible to avoid overfitting by imposing a spatial regularization term that penalizes the average squared difference between every voxel i and its neighbors $n(i)$. This type of regularization encourages the maps to be spatially smooth, reflecting a common view of cortical organization, and has been applied in a variety of MRI and fMRI experiments [1, 6–8, 11, 17]. The neighborhoods $n(i)$ can be defined in a number of ways, with neighbors chosen based on physical distance between voxels or distance along the cortical surface. For the experiments in this paper, we choose the 10 voxels that are closest to i along the cortical surface (varying the number of neighbors between 5 and 15 has little effect). Adding these regularization terms to our objective function, we obtain:

$$\frac{1}{T} \|\mathbf{a}^1 \mathbf{A}^1 - \mathbf{a}^2 \mathbf{A}^2\|_2^2 + \lambda \left[\sum_{i \in v_1} \sum_{j \in n(i)} \frac{1}{|n(i)|} (\mathbf{a}_i^1 - \mathbf{a}_j^1)^2 + \sum_{i \in v_2} \sum_{j \in n(i)} \frac{1}{|n(i)|} (\mathbf{a}_i^2 - \mathbf{a}_j^2)^2 \right]$$

where T is the number of timepoints in our dataset, v_k is the set of all voxels in ROI k , and λ is a hyperparameter that controls the regularization strength. We can write this objective compactly as

$$\left\| \begin{bmatrix} \frac{1}{\sqrt{T}} \mathbf{A}^1 T & -\frac{1}{\sqrt{T}} \mathbf{A}^2 T \\ \sqrt{\lambda} \mathbf{D}_1 & 0 \\ 0 & \sqrt{\lambda} \mathbf{D}_2 \end{bmatrix} \begin{bmatrix} \mathbf{a}^1 \\ \mathbf{a}^2 \end{bmatrix} \right\|_2^2$$

where \mathbf{D}_k is a sparse connectivity matrix; each row represents an edge from a voxel i to a voxel j , with nonzero entries in column i ($1/\sqrt{|n(i)|}$) and column j ($-1/\sqrt{|n(i)|}$). Our objective therefore has the form $\|\mathbf{X}_\lambda \cdot \boldsymbol{\beta}\|_2^2$, where $\boldsymbol{\beta} = [\mathbf{a}^1 \ \mathbf{a}^2]^T$ is the concatenation of the connectivity maps in both regions.

Since this is a homogeneous least-squares problem, it is clear that we must impose some constraint on the voxel weights $\boldsymbol{\beta}$ to avoid the degenerate solution $\boldsymbol{\beta} = 0$ (intuitively, the best timecourse prediction will always occur when the weight maps on both ROIs are identically zero, since this allows for perfect matching between the two regions). A standard method for choosing nonzero-weight solutions to homogeneous least-squares problems is to constrain the norm of $\boldsymbol{\beta}$ to be a constant, so we choose the constraint $\|\boldsymbol{\beta}\|_2 = 1$. In addition, we impose the elementwise constraint $\boldsymbol{\beta} \succeq 0$; allowing negative connectivity weights makes the maps hard to interpret, since multiple solutions can be superimposed (with different signs) and inter-region connectivity can be confounded with intra-region connectivity (since weights in the same region can have opposite signs).

2.3 Solving the Voxel-Level Optimization Problem

Due to the constraint $\|\boldsymbol{\beta}\|_2 = 1$, this optimization problem is not convex and may have multiple local minima. Although this makes the problem more complicated to analyze, the existence of multiple optima actually matches our intuition about functional connectivity structure; we know that for some pairs of regions (such as in early visual cortex [13]) there are multiple distinct connectivities between different subregions.

To find a locally optimal $\boldsymbol{\beta}$, we use a trust region approach [2, 3, 18]. This optimization method searches for local extrema by iteratively taking small steps in the parameter space. On each iteration, we create a convex approximation to the optimization problem by linearizing the norm constraint around the current set of parameters, and then find the optimal solution within a local trust region. Given our current solution $\boldsymbol{\beta}_i$, we minimize $\|\mathbf{X}_\lambda \cdot (\boldsymbol{\beta}_i + \mathbf{s})\|_2^2$ over \mathbf{s} , subject to $(\|\boldsymbol{\beta}_i\|_2^2 - 1) + 2\boldsymbol{\beta}_i^T \mathbf{s} \leq \theta$, $\|\mathbf{s}\|_2 \leq \Delta$, and $\boldsymbol{\beta}_i + \mathbf{s} \succeq 0$; the trust region is defined by the radius Δ and the constraint tolerance θ , and we iterate by setting $\boldsymbol{\beta}_{i+1} \leftarrow \boldsymbol{\beta}_i + \mathbf{s}$ until $\|\mathbf{s}\|_2 < \epsilon$. In our experiments, we use $\theta = 0.05$, $\Delta = \sqrt{\theta}$, $\epsilon = \Delta/100$, and $\lambda = 10^2$ (our results are not sensitive to this choice) and use CVX (a package for specifying and solving convex programs) to solve the convex problem at each iteration [9]. We obtain multiple solutions by trying 20 initializations of $\boldsymbol{\beta}_0$, each of which assigns all the connectivity weight to a single random voxel (in either region).

3 Results

3.1 Experimental Design

To validate our method, we use an objects-in-context block-design dataset consisting of 10 subjects performing 4 runs (for full details, see [1]). We trained our model on only a single run (306 timepoints) to demonstrate the efficiency of our approach. In each subject, LOC was defined in an independent set of localizer scans as the top 500 voxels responding more to objects than scrambled images.

Imaging data were acquired with a 3 Tesla G.E. Healthcare scanner. A gradient echo, echo-planar sequence was used to obtain functional images (TR=2s, 1.56x1.56x4 mm³). The functional data were motion-corrected and each voxel’s timecourse was z-scored to have zero mean and unit variance. A high-resolution (1x1x1mm³) SPGR structural scan was collected in each scanning session.

3.2 ILOC - rLOC Connectivity

We used the objects-in-context dataset to learn the connectivity between left LOC (lLOC) and right LOC (rLOC). It has been previously shown that LOC consists of two functionally distinct subregions: a posterior-dorsal subdivision (LO), and an anterior-ventral subdivision (pFs) [10]. Since these two subregions have different functional response patterns, we expect distinct functional connections between the anterior side of lLOC and rLOC and/or between the posterior side of lLOC and rLOC. For comparison purposes, we also apply a CCRF/FF correlation clustering method, which can only learn maps over one region at a time, from [14].

Fig. 2a shows the solutions for a representative subject. Little spatial structure is evident when using the CCRF/FF correlation clustering method with either left or right LOC as the seed region (top, results shown using left LOC as seed region) or using our method without regularization (middle). However, after adding the regularization term (bottom), we obtain a clear posterior-anterior segregation of the two solutions in this subject. These two clusters partition left and right LOC into two functional units, corresponding to LO and pFs [10]. Note that our algorithm did not use any prior knowledge about the number of subregions or their spatial configuration. Across subjects, the number of solutions in the regularized case ranged from one to three.

To quantify the correspondence between the left and right connectivity maps, we then measure the correlation between these posterior-anterior connectivity profiles for the left and right hemispheres, shown in Fig. 2b. Using correlation clustering (averaged for the choice of lLOC or rLOC as the seed region) or applying our method without regularization (NR), we obtain only a slight correlation between the posterior-anterior profiles. When adding regularization (R), however, our method produces maps which are highly correlated between hemispheres ($t_9 = 3.95, p < 0.01$, two-tailed t-test), giving a significant improvement over both the CCRF/FF method ($t_9 = 3.98, p < 0.01$, two-tailed paired t-test) and the unregularized method ($t_9 = 3.82, p < 0.01$, two-tailed paired t-test).




Fig. 2: ILOC-rLOC results. (a) In this representative subject, a CCRF/FF correlation clustering approach (top) fails to find anterior-posterior connectivity maps in left and right LOC, as does our method without spatial regularization (middle). Adding regularization (bottom) produces a separate posterior and anterior correspondence between hemispheres. (b) Only our proposed method (with spatial regularization) yields a strong correspondence between the hemispheres’ weight maps along the posterior-anterior axis. ** $p < 0.01$. (Best viewed in color)

Note that the representative timecourses for the anterior and posterior clusters are strongly positively correlated ($r = 0.83$ left, 0.81 right, averaged among subjects with exactly two clusters); the subtle distinction between these clusters therefore could not be identified by a CCA method [8]. Unlike [10], which used a specialized adaptation design, we are able to identify this anterior-posterior difference using only a single run from a dataset that was not tailored for this purpose.

4 Conclusion

We have shown that learning smooth voxel-level functional connectivity maps can be formulated as a constrained least-squares problem, and have demonstrated that we can recover the ground truth subregional connectivity structure of LOC without using specialized datasets or a large number of training timepoints. By simultaneously learning weight maps over two regions and by including a spatial smoothness term, our method is much more sensitive to fine-grained connectivity differences than previous methods.

This work is funded by National Institutes of Health Grant 1 R01 EY019429 (to L.F.-F. and D.M.B.), a National Science Foundation Graduate Research Fellowship under Grant No. DGE-0645962 (to C.B.) and a William R. Hewlett Stanford Graduate Fellowship (to M.C.I.).

References

1. Baldassano, C., Iordan, M.C., Beck, D.M., Fei-Fei, L.: Voxel-level functional connectivity using spatial regularization. *NeuroImage* (2012), <http://www.sciencedirect.com/science/article/pii/S1053811912007756>
2. Byrd, R.H., Schnabel, R.B., Shultz, G.A.: A trust region algorithm for nonlinearly constrained optimization. *SIAM Journal on Numerical Analysis* 24(5), 1152–1170 (October 1987)
3. Celis, M.R., Jr., J.E.D., Tapia, R.A.: A trust region strategy for equality constrained optimization. Tech. Rep. 84-1, Mathematical Sciences Department, Rice University (September 1984)
4. Chai, B., Walther, D.B., Beck, D., Fei-Fei, L.: Exploring functional connectivity of the human brain using multivariate information analysis. In: *Advances in Neural Information Processing Systems* 22 (2009)
5. Cohen, A.L., Fair, D.A., Dosenbach, N.U., Miezin, F.M., Dierker, D., Van Essen, D.C., Schlaggar, B.L., Petersen, S.E.: Defining functional areas in individual human brains using resting functional connectivity MRI. *Neuroimage* 41, 45–57 (May 2008)
6. Conroy, B., Singer, B., Haxby, J., Ramadge, P.: fMRI-based inter-subject cortical alignment using functional connectivity. In: *Advances in Neural Information Processing Systems* 22 (2009)
7. Cuignet, R., Chupin, M., Benali, H., Colliot, O.: Spatial and anatomical regularization of SVM for brain image analysis. In: *Advances in Neural Information Processing Systems* 23 (2010)
8. Deleus, F., Van Hulle, M.M.: Functional connectivity analysis of fMRI data based on regularized multiset canonical correlation analysis. *J. Neurosci. Methods* 197(1), 143–157 (Apr 2011)
9. Grant, M., Boyd, S.: CVX: Matlab software for disciplined convex programming, version 1.21. <http://cvxr.com/cvx> (Apr 2011)
10. Grill-Spector, K., Kushnir, T., Edelman, S., Avidan, G., Itzchak, Y., Malach, R.: Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron* 24(1), 187 – 203 (1999), <http://www.sciencedirect.com/science/article/pii/S0896627300808326>
11. Grosenick, L., Klingenberg, B., Knutson, B., Taylor, J.E.: A family of interpretable multivariate models for regression and classification of whole-brain fMRI data. *ArXiv e-prints* (Oct 2011)
12. Haak, K.V., Winawer, J., Harvey, B.M., Dumoulin, S.O., Wandell, B.A., Cornelissen, F.W.: Cortico-cortical population receptive field modeling. In: *Perception* 40 ECVF Abstract Supplement. p. 49 (2011)
13. Heinzle, J., Kahnt, T., Haynes, J.: Topographically specific functional connectivity between visual field maps in the human brain. *NeuroImage* 56(3), 1426 – 1436 (2011), <http://www.sciencedirect.com/science/article/pii/S1053811911002540>
14. Kim, J.H., Lee, J.M., Jo, H.J., Kim, S.H., Lee, J.H., Kim, S.T., Seo, S.W., Cox, R.W., Na, D.L., Kim, S.I., Saad, Z.S.: Defining functional SMA and pre-SMA

- subregions in human MFC using resting state fMRI: functional connectivity-based parcellation method. *Neuroimage* 49, 2375–2386 (Feb 2010)
15. Margulies, D.S., Kelly, A.M., Uddin, L.Q., Biswal, B.B., Castellanos, F.X., Milham, M.P.: Mapping the functional connectivity of anterior cingulate cortex. *Neuroimage* 37, 579–588 (Aug 2007)
 16. Margulies, D.S., Vincent, J.L., Kelly, C., Lohmann, G., Uddin, L.Q., Biswal, B.B., Villringer, A., Castellanos, F.X., Milham, M.P., Petrides, M.: Precuneus shares intrinsic functional architecture in humans and monkeys. *Proc. Natl. Acad. Sci. U.S.A.* 106(47), 20069–20074 (Nov 2009)
 17. Ng, B., Abugharbieh, R.: Generalized sparse regularization with application to fmri brain decoding. In: *IPMI*. pp. 612–623 (2011)
 18. Powell, M.J.D., Yuan, Y.: A trust region algorithm for equality constrained optimization. *Mathematical Programming* 49, 189–211 (1991)
 19. Rogers, B.P., Morgan, V.L., Newton, A.T., Gore, J.C.: Assessing functional connectivity in the human brain by fMRI. *Magn Reson Imaging* 25, 1347–1357 (Dec 2007)
 20. Roy, A.K., Shehzad, Z., Margulies, D.S., Kelly, A.M., Uddin, L.Q., Gotimer, K., Biswal, B.B., Castellanos, F.X., Milham, M.P.: Functional connectivity of the human amygdala using resting state fMRI. *Neuroimage* 45, 614–626 (Apr 2009)
 21. Stockburger, D.W.: *Introductory statistics: Concepts, models, and applications* (1996), <http://www.psychstat.missouristate.edu/IntroBook/sbk00.htm>
 22. Zhang, D., Snyder, A.Z., Fox, M.D., Sansbury, M.W., Shimony, J.S., Raichle, M.E.: Intrinsic functional relations between human cerebral cortex and thalamus. *J. Neurophysiol.* 100, 1740–1748 (Oct 2008)