

Single-Subject Structural Networks with Closed-Form Rotation Invariant Matching Improve Power in Developmental Studies of the Cortex

Benjamin M. Kandel¹, Danny JJ Wang², James C. Gee¹, and Brian B. Avants¹

¹ Penn Image Computing and Science Laboratory, University of Pennsylvania,
Philadelphia, Pennsylvania

² Department of Neurology, University of California,
Los Angeles, California

Abstract. Although much attention has recently been focused on single-subject functional networks, using methods such as resting-state functional MRI, methods for constructing single-subject structural networks are in their infancy. Single-subject cortical networks aim to describe the self-similarity across the cortical structure, possibly signifying convergent developmental pathways. Previous methods for constructing single-subject cortical networks have used patch-based correlations and distance metrics based on curvature and thickness. We present here a method for constructing similarity-based cortical structural networks that utilizes a rotation-invariant representation of structure. The resulting graph metrics are closely linked to age and indicate an increasing degree of closeness throughout development in nearly all brain regions, perhaps corresponding to a more regular structure as the brain matures. The derived graph metrics demonstrate a four-fold increase in power for detecting age as compared to cortical thickness. This proof of concept study indicates that the proposed metric may be useful in identifying biologically relevant cortical patterns.

1 Introduction

Brain connectivity has emerged as a dominant trend in recent neuroimaging research. Connectivity can be measured by correlation of function [6], diffusion-based structural connections [7], and covariation of cortical structure across populations [1]. Covariance patterns of cortical structure often recapitulate functional connectivity patterns [11], although this analysis is complicated in part because while functional networks can be derived on a per-subject basis, structural networks are commonly derived on a group basis, making statistical analysis of structural networks challenging [3]. Therefore, although [11] showed that ICA components of cortical covariance are similar to ICA components of rs-fMRI networks, statistical analysis on a per-subject basis is not straightforward.

As a result of these difficulties with group-wise structural networks, several groups have begun to pursue single-subject cortical networks. These methods

have coalesced around two camps: Those that generate a network based on the difference between some derived scalar metrics from cortical morphology, such as cortical thickness or curvature, and those that use a patch-based correlation between two different voxels. In the first camp, [10] used a combination of cortical thickness and curvature-based metrics to construct networks. Similarly, Dai [5] used differences between regional cortical thickness measurements to create cortical networks. In the second camp, Tijms et al. [12] construct similarity networks based on the correlation between patches centered around different voxels.

Previous methods for constructing single-subject cortical networks suffer from several drawbacks. Fundamentally, cortical thickness, although an important measure of cortical structure, does not capture all the information of the surrounding neighborhood of a voxel; similarities in cortical thickness do not necessarily imply similarities in structure as a whole. Furthermore, when combining separate scalar values, such as thickness and one or several curvature measurements, how to combine the features into a meaningful distance metric is not at all straightforward, and previous methods have constructed complex and highly specialized models for individual diseases [10]. Examining correlations of patches centered around given voxels is an intuitive and straightforward approach; the only parameter to choose is the patch size, which can be chosen based on principled methods.

However, the method by Tijms for computing correlations between voxels has some technical drawbacks. To obtain rotation invariance, the method rotated cubes in increments of 45 degrees to obtain maximal correlation with the test patch. This method suffers from several flaws: First, the choice of 45 degrees is arbitrary; there is no fundamental reason to only rotate patches in increments of 45 degrees. Second, it appears that [12] “rotates” cubes by permuting the entries in the patch. This will lead to distortion in the shape, as the distance from voxel in the center of a square to the voxel immediately above it is a factor of $\sqrt{2}$ smaller than the distance between the center of the square to the corner. Third, the operation is only defined for the somewhat arbitrary shape of cubes of size $3 \times 3 \times 3$ voxels. Finally, even if the rotation were defined for arbitrary angles, an exhaustive search in three dimensions as [12] does would be computationally infeasible.

As opposed to the previously proposed approaches, we propose a closed-form, truly rotation-invariant approach to computing structural similarity across brain regions. Reorienting two images so that their orientations match is a well-posed problem that has known solutions. Leveraging these methods, we construct a rotation-invariant representation of the patches surrounding individual voxels. The correlation between different patches for different voxels give the adjacency weights in the graph. An overview is shown in Figure 1.

We apply the method to tracking the network dynamics of cortical structure in a pediatric dataset. We find that network measurements increase the power of detecting age changes by a factor of four as compared to using cortical thickness. In sum, our contributions are: 1) Method for constructing rotation-invariant structural similarity metrics; 2) Method for combining these sensibly

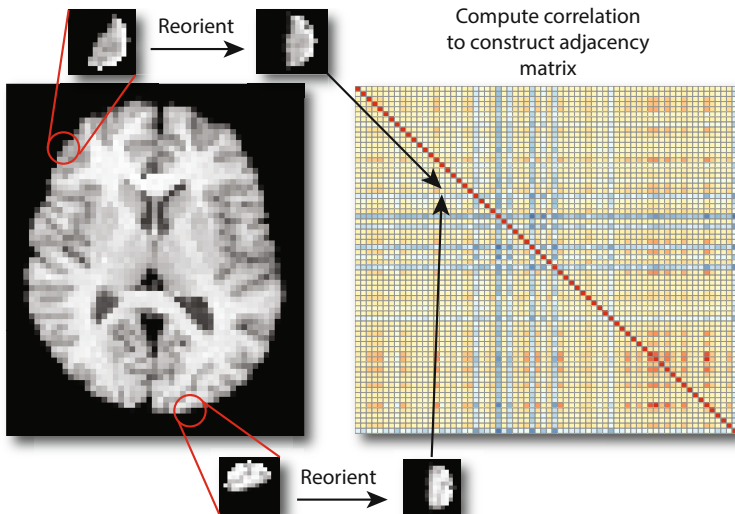


Fig. 1. Overview of adjacency matrix construction. The patches surrounding each voxel are extracted and aligned to a common reference frame. The correlation between the patches is entered into the adjacency matrix.

into a smaller-dimensional graph; 3) Demonstration that pediatric development and gender is closely correlated with node closeness, which is more predictive of age than scalar ROI values; and 4) Demonstration that the proposed method is superior to thickness distance-based cortical networks for predicting age and gender.

2 Methods

We consider a undirected graph \mathcal{G} with edges \mathcal{K} and nodes \mathcal{N} . The edges in the graph correspond to the strength of connection between different parts of the brain. Given an image \mathcal{I} with J scalar-valued voxels at locations $x_j \in \mathcal{I}, j = 1, \dots, J$, we seek a function $d : \mathcal{I}(x_i) \times \mathcal{I}(x_j) \mapsto \mathbb{R}^+$ to map from the input image space to graph edge weights. In fMRI, this function can simply be the correlation of the time-series at the different voxels, but because we consider scalar voxels, this option is not available to us. Instead, we consider the similarity between the neighborhood surrounding the voxels of interest. We denote the neighborhood of a voxel x_i as $\mathcal{N}_i = \{x_j \mid \|x_j - x_i\|_2^2 \leq r\}$, where r is the radius of the neighborhood. The edge weight between voxels x_i, x_j is then described by

$$\mathcal{K}_{i,j} = d(\mathcal{N}_i, \mathcal{N}_j). \quad (1)$$

A naïve approach to generating the function $d(x_i, x_j)$ would be to simply compute the correlation of the vector representation of \mathcal{N}_i with \mathcal{N}_j , but this would

not account for the curved structure of the brain. Ensuring that the metric between two neighborhoods is rotation-invariant is not trivial.

Rotation-Invariant Correlation: We use a closed-form solution to align voxel neighborhoods to a canonical reference frame. Although the choice of orientation is arbitrary, we must choose one orientation as a base for reorienting all the patches. Instead of choosing one patch, which could bias our results, we first generate an $n \times m$ matrix, where each row i is the vector representation of \mathcal{N}_i , each of which consists of m voxels. For computational feasibility, we take a random sampling of n voxels from around the cortex. We found that no benefit was achieved by sampling more than 5000 sample voxels. The first singular vector of the sample patch matrix serves as our canonical reference frame.

Aligning the orientation of two vectors has a well-known analytical solution [9]. Aligning two images corresponds to aligning the orientations of the first eigenvector (or two eigenvectors for a 3D image) of the covariance matrix of the gradient of the image. We denote the gradient operator $g : \mathcal{N} \mapsto \mathbb{R}^D$, where D is the number of dimensions in the image. We compute the gradient by convolving our image with the derivative of a Gaussian ($\sigma = 1$ voxel). The covariance matrix $C(\mathcal{N}_i)$ of the gradient of the neighborhood \mathcal{N}_i is then given by

$$C(\mathcal{N}_i) = \sum_{x_i \in \mathcal{N}_i} g(\mathcal{I}(x_i)) g(\mathcal{I}(x_i))^T \in \mathbb{R}^{D \times D}. \quad (2)$$

To align the patches of two voxels x_i and x_j , we denote the k 'th eigenvector of $C(\mathcal{N}_i)$ as w_k and the k 'th eigenvector of $C(\mathcal{N}_j)$ as v_k and calculate the rotation matrix Q that best aligns them:

$$\arg \min_Q \sum_{k \in \{1,2\}} \|w_k - Qv_k\|^2 \quad (3)$$

Denoting $B = w_k v_k^T$, we compute the singular value decomposition (SVD) of B : $B = USV^T$. Then the analytical solution to Equation 3 is given by $Q = UMV^T$, where $M = \text{diag}[1 \ 1 \ \det(U) \ \det(V)]$. We then rotate the voxel coordinates x_i by Q and use a linear interpolator to regenerate the neighborhood image after the rotation. Because the eigenvalues are unsigned, they can sometimes result in an alignment that is flipped by 180 degrees from the correct alignment. To eliminate this possibility, we check for a negative correlation between the sample patch and the reference patch and flip the rotation matrix if necessary. A more computationally expensive alternative is to use the Radon transform to estimate orientation [8].

Correlation Matrix Construction: In most connectome construction schemes, data is first averaged over some brain parcellation and those averaged values are then used for calculating correlations [13]. In this case, however, the average of a series of patches is ill-defined, and we found that constructing correlation matrices in this manner did not yield meaningful results. Instead, we first calculated the correlation of the vector representation of the reoriented

neighborhood of each voxel in the cortex with every other voxel in the cortex. As is standard, we constructed an $I \times I$ correlation matrix, where there are a total of I regions (nodes) for each subject. The correlation between region i and j was then calculated as the mean of the correlation of the vector representation of the reoriented neighborhood of each voxel in region i with each voxel in region j . Once the correlation matrix was constructed, normalized closeness was calculated using the `igraph` package in R [4]. For node i , closeness is defined as $\sum_{j \neq i} K_{i,j}$, with the normalization running over all nodes.

We compared our results to the method of [5], which uses the difference in cortical thickness between two regions to construct the network. Given I total regions, each with cortical thickness $t(i)$, we construct an $I \times I$ distance matrix D , where $D(i, j) = \exp - \left(\frac{(t(i) - t(j))^2}{\sigma} \right)$. We set σ to 0.015, as recommended in [5].

Clinical Pediatric Data: Our pediatric data consists of 119 subjects, with mean age 12.42, range 7.07-17.99 years, 61 females and 58 males. Magnetization-Prepared Rapid Acquisition Gradient Echo (MPRAGE) images were acquired on a Siemens Trio Tim scanner (3T) using a 3D inversion recovery sequence with TR/TE/TI = 2170/4.33/1100 ms. The resolution was $1 \times 1 \times 1 \text{mm}^2$ with a matrix size of $256 \times 256 \times 192$. Flip angle = 7° and total scan time was 8:08 minutes. Image preprocessing, including bias correction, skull-stripping, segmentation, and warping of the AAL label set to the subject space was performed with ANTs [2], and the AAL label set was used for generating ROI's to construct the graphs.

Computation Considerations and Parameters: One of the advantages of a correlation-based approach to similarity evaluation is the simplicity and lack of parameters in the method. The only free parameter in this method is the patch size, which can be set based on the scale of features to be matched. Matching small patches will find similarities between small features, such as position on sulcus or gyrus, whereas matching large patches will find regional similarities. We found that downsampling images to 3mm and using a patch radius of 3 voxels was appropriate for looking at correlations between ROI's on the scale of AAL labels. Reorientation of patches takes under 20 minutes on an Intel Xeon CPU at 2.40GHz with 2 GB of memory. The code for constructing the adjacency matrix is open-source and is available at <https://github.com/bkandel/PatchAnalysis>.

3 Results

Validation of Rotation Invariance: We first checked that our output graph metrics are indeed rotation invariant. We rotated the images of ten subjects chosen at random in increments of ten degrees and constructed graphs from the rotated images. We plotted the deviation of the closeness value from the subject-wise mean closeness value vs. rotation (Figure 2 for a sample ROI).




Fig. 2. Verification of rotation invariance: Difference between closeness of left hippocampus and subject-wise mean vs. rotation. The means over all rotations did not show a correlation with rotation, although there were slightly more low outliers at 90 and 270 degrees and slightly more high outliers at 0 and 180 degrees.

T-tests between the closeness values at each rotation and the mean closeness did not reveal a significant difference from the mean for any rotation (minimum FDR-corrected p-value 0.13). It may still be possible to achieve even less dependence on angle by scaling the patches to minimize the effect of outlying voxels.

Sample Subject: A thresholded correlation matrix overlaid on the MNI template brain shown in Figure 3. Cortical thickness-derived graphs tend to have many cliques, corresponding to regions with similar cortical thickness, that are not connected to each other; graphs using our method tended to have more central nodes.

Pediatric Data: In pediatric data, graph closeness was found to be highly correlated with age in most regions, whereas cortical thickness was not found to be as correlated in as many regions. To evaluate correlation of closeness with age in an ROI-wise basis, we performed an ANOVA comparing the models (in R notation) $\text{ROI.Closeness} \sim \text{Sex} + \text{BrainVolume}$ and $\text{ROI.Closeness} \sim \text{Sex} + \text{BrainVolume} + \text{Age} + \text{Age}:\text{Sex} + \text{Age}^2 + \text{Age}^2:\text{Sex}$, where $:$ signifies an interaction term. Analogous ANOVA's, using the same covariates, were performed for the thickness-derived graphs and cortical thickness. Patch closeness was found to be significantly correlated with age (after FDR correction) in 68 out of 68 cortical regions, whereas for thickness, only 17 were found to be correlated, and for thickness-derived structural graphs, 0 regions were correlated with age. Results for global mean measurements, with p-values computed with the same models, is shown in Figure 4. To recover the regression coefficient for age, with power 0.99 and alpha level 0.015, we would need 58 subjects using our method; 223 using cortical thickness; and 506 using cortical thickness-derived graph metrics.

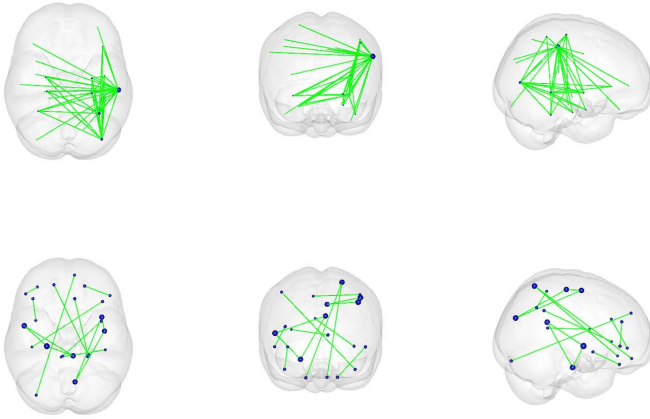


Fig. 3. Correlation between nodes (thresholded to reveal only top 2% of edges), overlaid on MNI brain. Top: Representative subject graph using our method. Bottom: Representative subject graph using cortical thickness-derived graph.

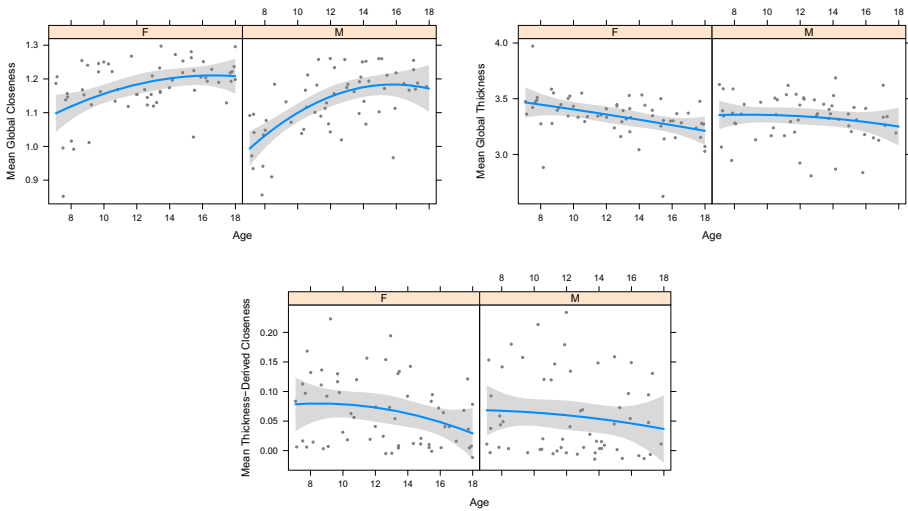


Fig. 4. Upper left: Mean closeness vs. age using our method; p-value (details in text) 5.46×10^{-8} . Upper right: Mean thickness vs. age, p-value 0.03. Bottom: Mean thickness-derived graph closeness vs. age, p-value 0.29.

4 Conclusion

We have presented a principled and closed-form method to generate single-subject cortical graphs and shown that the graphs are more sensitive to age changes than cortical thickness or cortical thickness-derived graphs are. This method has only one free parameter and shows a biologically meaningful trend with age. The method may also be used for tracking cortical changes in Alzheimer's disease and other neurodegenerative conditions.

References

1. Alexander-Bloch, A., Giedd, J.N., Bullmore, E.: Imaging structural co-variance between human brain regions. *Nature Reviews Neuroscience* 14(5), 322–336 (2013)
2. Avants, B.B., Tustison, N.J., Song, G., Cook, P.A., Klein, A., Gee, J.C.: A reproducible evaluation of ANTs similarity metric performance in brain image registration. *NeuroImage* 54(3), 2033–2044 (2011)
3. Bassett, D.S., Bullmore, E., Verchinski, B.A., Mattay, V.S., Weinberger, D.R., Meyer-Lindenberg, A.: Hierarchical organization of human cortical networks in health and schizophrenia. *The Journal of Neuroscience* 28(37), 9239–9248 (2008)
4. Csardi, G., Nepusz, T.: The igraph software package for complex network research. *InterJournal* 2006. *Complex Systems* 1695, 1–9 (2006)
5. Dai, D., He, H., Vogelstein, J., Hou, Z.: Network-based classification using cortical thickness of AD patients. In: Suzuki, K., Wang, F., Shen, D., Yan, P. (eds.) *MLMI 2011*. LNCS, vol. 7009, pp. 193–200. Springer, Heidelberg (2011)
6. Greicius, M.D., Krasnow, B., Reiss, A.L., Menon, V.: Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences* 100(1), 253–258 (2003)
7. Greicius, M.D., Supekar, K., Menon, V., Dougherty, R.F.: Resting-state functional connectivity reflects structural connectivity in the default mode network. *Cerebral Cortex* 19(1), 72–78 (2009)
8. Jafari-Khouzani, K., Soltanian-Zadeh, H.: Radon transform orientation estimation for rotation invariant texture analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 27(6), 1004–1008 (2005)
9. Kabsch, W.: A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A* 32(5), 922–923 (1976)
10. Raj, A., Mueller, S.G., Young, K., Laxer, K.D., Weiner, M.: Network-level analysis of cortical thickness of the epileptic brain. *NeuroImage* 52(4), 1302–1313 (2010)
11. Segall, J.M., Allen, E.A., Jung, R.E., Erhardt, E.B., Arja, S.K., Kiehl, K., Calhoun, V.D.: Correspondence between structure and function in the human brain at rest. *Frontiers in Neuroinformatics* 6, 10 (2012)
12. Tijms, B.M., Seris, P., Willshaw, D.J., Lawrie, S.M.: Similarity-based extraction of individual networks from gray matter MRI scans. *Cerebral Cortex* 22(7), 1530–1541 (2012)
13. Varoquaux, G., Craddock, R.C.: Learning and comparing functional connectomes across subjects. *NeuroImage* 80, 405–415 (2013)