

Predicting Activation Across Individuals with Resting-State Functional Connectivity Based Multi-Atlas Label Fusion

Georg Langs^{1,2}, Polina Golland², and Satrajit S. Ghosh^{3,4}

¹ Department of Biomedical Imaging and Image-guided Therapy, CIR Lab, Medical University of Vienna, Vienna, Austria

² CSAIL, MIT

³ McGovern Institute for Brain Research, MIT, Cambridge, MA, USA

⁴ Department of Otolaryngology and Laryngology, Harvard Medical School, Boston, MA, USA

Abstract. The alignment of brain imaging data for functional neuroimaging studies is challenging due to the discrepancy between correspondence of morphology, and equivalence of functional role. In this paper we map functional activation areas across individuals by a multi-atlas label fusion algorithm in a functional space. We learn the manifold of resting-state fMRI signals in each individual, and perform manifold alignment in an embedding space. We then transfer activation predictions from a source population to a target subject via multi-atlas label fusion. The cost function is derived from the aligned manifolds, so that the resulting correspondences are derived based on the similarity of intrinsic connectivity architecture. Experiments show that the resulting label fusion predicts activation evoked by various experiment conditions with higher accuracy than relying on morphological alignment. Interestingly, the distribution of this gain is distributed heterogeneously across the cortex, and across tasks. This offers insights into the relationship between intrinsic connectivity, morphology and task activation. Practically, the mechanism can serve as prior, and provides an avenue to infer task-related activation in individuals for whom only resting data is available.

1 Introduction

Establishing functional correspondence across the brains of individuals is a central prerequisite for neuroimaging group studies. Standard approaches rely on brain morphology to perform group-wise registration, and their improvement has brought a substantial boost to the specificity of neuroimaging results and their interpretation in light of neuroscientific questions. Recent results indicate that the variability of the functional architecture across individuals makes the concept of correspondence more challenging to grasp. Specifically, the link between anatomical location and functional role can be weak. This results in the decrease of specificity in group studies, and potential bias. In this paper we propose multi-atlas label fusion based on functional alignment. The method establishes correspondence of cortical positions based on resting-state functional magnetic resonance imaging (rs-fMRI) signals. Using this functional alignment with label-fusion of activations

observed during task fMRI (t-fMRI) in a population of source subjects, we predict task activations in a target, aligned subject. Transferring information using functional connectivity alignment results in higher accuracy of transferring task activation compared to morphological alignment. This method extends functional region based analyses [2] to functional networks.

Alignment of function across individuals. Neuroimaging group-studies typically rely on registering structural imaging data of all subjects to a common template using software such as FreeSurfer [5], FSL [8], or SPM [1]. This establishes spatial correspondence across the population, and allows for local comparison of activation, or connectivity. However, function exhibits a high degree of variability [12] and is not necessarily tightly linked to anatomy [2]. Approaches to match function across individuals beyond relying on anatomy have been proposed before. In [14] the cortical surfaces were aligned by maximizing the correlation among fMRI signals recorded in different subjects during watching synchronized movies. A common space representing visual stimulus responses was used to establish correspondence across individuals in [7]. In [10] a joint manifold representing the functional connectivity patterns recorded during language experiments in multiple individuals was used to align function across subjects independent of the anatomical anchors of functional units. Instead of using across-subject correlation, it relies on the within-subject correlation patterns to match shared network architecture across the population.

Contribution. In this paper we extend the functional connectivity alignment proposed in [10] to “resting state” data and multi-atlas label fusion. First, we establish correspondence between cortical surface points across individuals by functional connectivity alignment. Then, we predict task activations in a target subject, by evaluating the similarity of the matched embedding maps of all source subjects, and the target. Finally, we transfer predictions by selecting the most similar subject on a voxel-by-voxel basis. In the absence of any ground truth on functional connectivity, evaluating task-based correspondence by measuring the Dice coefficient between predicted and actual task-based activation provides an objective way to evaluate the alignment. The proposed approach can cope with variability, since it selects the most similar individual in the source population, based on functional connectivity, for activation label prediction. Furthermore, we can use it to study and compare the discrepancy between functional transfer of activation areas, and morphological transfer across the cortex. Finally, we gain insights into the relationship between the embedding structure, and the correspondence across the cortex.

Related work. The work is closely linked to surface matching algorithms based on curvature [11]. However, instead of relying on morphological features, we inject functional information to map similarity across the cortex. Multi-atlas label fusion transfers information such as labels from a set of atlases to target data. Instead of building a single model from the atlas population, it first fits all atlas templates to the target data. Then it transfers labels from atlases, or groups of atlases to the target based on a similarity function that reflects the suitability of an atlas

for predicting the labeling of the target [15,9,18]. Reducing the dimensionality of data can capture underlying structure that is not apparent in its native space. It can be achieved through linear models (e.g., PCA [13], ICA [3]) or using non-linear embedding approaches. The latter assume that the data of interest *lives* on a low dimensional nonlinear manifold and estimate its intrinsic coordinate structure by *embedding* the data.

2 Method

The proposed method first performs embedding of individual resting state functional connectivity graphs. The embedding maps are aligned and labels are transferred from a source population to a target individual based on their fit in the embedding space after alignment in a multi-atlas label fusion approach.

Embedding the Intrinsic Connectivity Structure of rs-fMRI Data. We view an fMRI sequence $\mathbf{I} \in \mathbb{R}^{T \times N}$ as a graph of N voxels (or cortical surface vertices). Each voxel (vertex) \mathbf{v}_i carries an fMRI signal over T time points. We calculate a pairwise similarity matrix $\mathbf{W} \in \mathbb{R}^{+N \times N}$ that assigns the correlation $\mathbf{W}(i, j)$ of the time-courses to each pair of voxels (i, j) (edge) [4]. Following [4] this graph defines a Markov chain with transition matrix $\mathbf{P} = \mathbf{D}^{-\frac{1}{2}} \mathbf{W} \mathbf{D}^{-\frac{1}{2}}$, where \mathbf{D} is a diagonal normalization matrix such that $d_i = \mathbf{D}(i, i) = \sum_j w(i, j)$ is the strength of node i . The eigenvectors of the transition matrix scaled by their eigenvalues (λ^t) define an embedding that results in a representation of each voxel as a point in the embedding space: $i \mapsto \Phi_i$ [10] (where t is the diffusion time parameter). To capture positive and negative correlations and to create an affinity matrix, the correlation matrix was scaled between 0 and 1 and then converted to a sparse graph representation (A) using a nearest neighbor approach. The 100 closest neighbors of each vertex were retained. The graph was checked to ensure that it was a connected graph. The diffusion map embedding was then computed on the normalized Laplacian of this graph. The eigenvalues λ are divided by $1 - \lambda$. The division of the λ parameter provides noise robustness and allows variation from the standard form potentially eliminating the use of the diffusion time (t) parameter. In many empirically tested cases, where the embedding is known, setting $t = 0$ returns a result that is close to using the optimal diffusion time parameter. Given that the optimal solution is generally unknown, setting ($t = 0$) is often a practical choice. The resulting embedding is a lower dimensional representation of the intrinsic functional connectivity of the brain.

Aligning Embeddings. For a target \mathbf{I}^T and each source \mathbf{I}_s^S , we find an orthonormal alignment via Procrustes analysis [16]: $\mathbf{Q}_{S,T}$. Given target- $\Phi^T = [\phi_1^T, \dots, \phi_N^T]^\top$ and source embedding coordinates $\Phi^S = [\phi_1^S, \dots, \phi_1^S]^\top$, $\mathbf{Q}_{S,T} = \mathbf{V}\mathbf{U}^\top$, where \mathbf{U} and \mathbf{V} are constructed via the singular value decomposition $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top = \Phi^T \text{diag}(\mathbf{w}_m) \Phi_{r_m}^S$. See [10] for detailed description of the approach. Figure 1 shows each of the first 5 individual components of the average embedding from 40 participants and the same components from 3 individual participants (we enlarge one for better visibility in the bottom row). We hypothesise

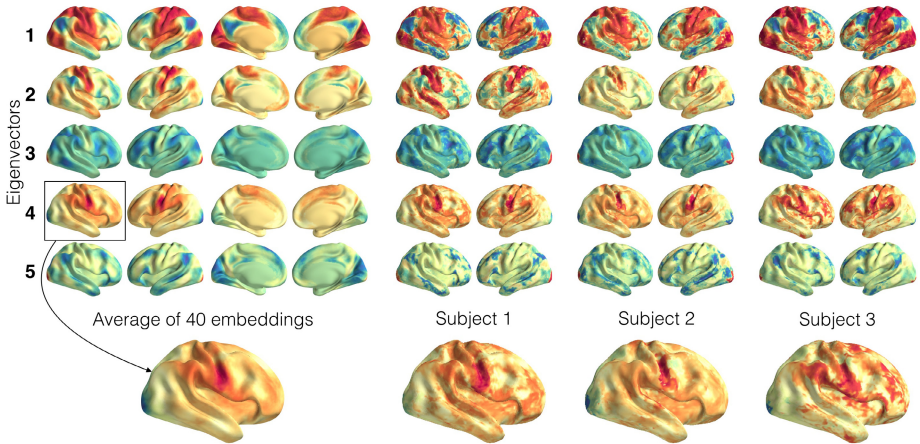


Fig. 1. The first 5 components of an average embedding computed from the 40 participants and three individual embeddings. The differences across individuals reflect different spatial distributions of resting state functional networks. We use the similarity of these surface markers to (locally) select the best source subjects during prediction by label fusion.

that these components form an intrinsic functional basis of brain activity and show aligned embedding coordinates on the cortical surface. The coefficients of the first 5 eigenvectors projected to the surface after functional alignment in the embedding space mark comparable systems on the cortex. Their fine-grained cortical distribution varies across individuals. The proposed label transfer is based on the assumption that activation can be transferred among individuals with similar cortical functional eigenvector profiles.

Predicting Activation. The proposed multi-atlas label fusion method chooses source subjects for the prediction of activation on each cortical vertex based on how similar those eigenvector coefficients (Figure 1) are at the vertex location. That is, we select source subjects based on how similar their resting-state connectivity aligns with the target subject connectivity. The alignment in the embedding space enables a straight-forward calculation of this similarity. Specifically, we predict activation maps in a target subject based on the known activation in a set of source subjects. We know the activation $f(\mathbf{v}_{i,s}^S)$ of each voxel $\mathbf{v}_{i,s}^S$ in each source volume s . Given the target embedding Φ^T and all aligned source embeddings Φ_s^{AS} for each voxel (or surface vertex) we calculate a score by the Euclidean distance between the target point ϕ_i^T and the anatomically corresponding source point $\phi_{i,s}^{AS}$ in the embedding space: $k_{i,s} = \|\phi_i^T - \phi_{i,s}^{AS}\|$. Then the predictor for the activation $f(\mathbf{v}_i^T)$ at voxel \mathbf{v}_i^T is

$$\forall i : f(\mathbf{v}_i^T) = f(\mathbf{v}_{i,s^*}^S), \text{ where } s^* = \underset{s=1,\dots,S}{\operatorname{argmin}} k_{i,s}. \quad (1)$$

We compare this prediction guided by the functional alignment to two alternatives that are based solely on the anatomical position. As a first comparison,

we predict $f(\mathbf{v}_i^T) = \text{mean}_{s=1, \dots, S} f(\mathbf{v}_{i,s}^S)$, i.e., the average activation f at the anatomically corresponding positions in the source subjects. Secondly, we predict $f(\mathbf{v}_i^T) = f(\mathbf{v}_{i,s}^S)$, where s is chosen randomly. For the evaluation, we average this prediction-accuracy over all individuals.

3 Experiments

Data We used data from 40 randomly selected participants from the Human Connectome Project (HCP; [17]) 500-Subject data release. For each participant, we used the data from one of the preprocessed functional runs [6] together with FIX cleanup. For each rs-fMRI run, data projected onto the average cortical surface (59412 vertices \times 1200 timepoints) were used to construct a correlation matrix (59412 \times 59412). This matrix is the basis for the embedding and functional alignment. We used task fMRI data of 7 paradigms from the same participants (motor, language, working memory, social, gambling relational, emotion). For each paradigm we used z-score maps calculated for each contrast within each paradigm. We evaluated the accuracy of the activation prediction by leave-one-out cross validation across 40 subjects. During each prediction by label fusion, we chose one contrast, and predicted the Z-score on a hold-out target subject. The prediction was performed by label-fusion from the Z-scores of the remaining source subjects. We calculated the Dice coefficient between predicted-, and actual activation region in the target subject using a Z-score cutoff of $Z > 3.09$, corresponding to $p < 0.001$ as a marker for *activation region*.

Impact of Embedding Parameters. To evaluate the effect of embedding parameters of the diffusion map embedding, on the prediction accuracy, two of the embedding parameters were varied: diffusion time was set to 0 and 5, while the number of nearest neighbors was varied across 50, 100, and 500. In our experiments, these parameters had a minimal effect on the Dice coefficients, with variations ($\sigma > 0.01$) being observed only at values of $Z > 5$ for the emotion and gambling tasks.

Comparison of Prediction Accuracy. We compare the proposed label-fusion with functional alignment based prediction with an approach that does not take the functional rs-fMRI information into account. Figure 2 provides an overview of prediction accuracy with functional alignment label-fusion versus the average accuracy when predicting the activation based on the anatomical position. Each row corresponds to one of 7 paradigms, and each column to one of the contrasts in these paradigms. We only show the first 6 contrasts for each paradigm. Label-fusion that takes the distance in the functional embedding map into account when choosing a source subject for each vertex, consistently yields higher accuracy than uninformed transfer. We observe similar differences when comparing the label-fusion with prediction based on the average Z-score in the population. Note that the proposed approach primarily improves accuracy, if the anatomical mapping accuracy is poor. Figure 3 provides a detailed comparison for 4 illustrative paradigms, and a range of Z-score cut-off values. The plots show the

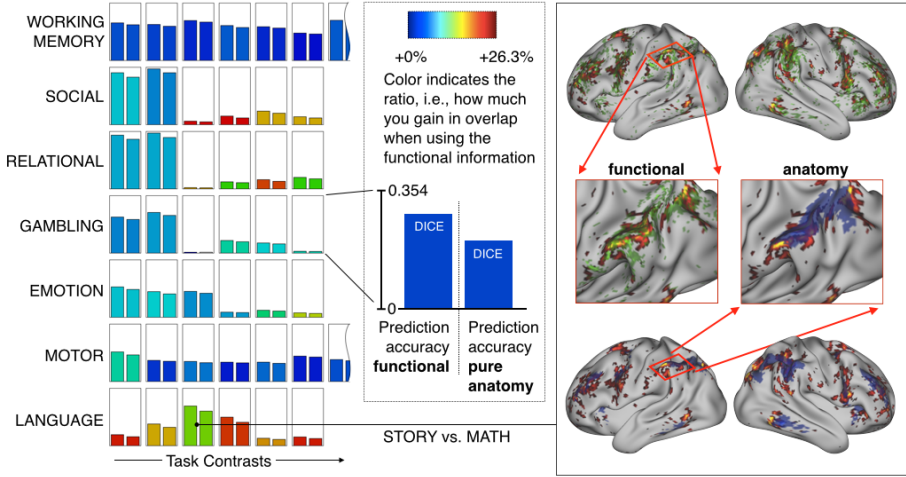


Fig. 2. Improvement of accuracy by functional alignment label-fusion over anatomical alignment varies across different experiment conditions. For 7 tasks, we show the prediction accuracy for the first 6 contrasts. On the right the actual individual activation and predictions based on functional multi-atlas fusion (green) is shown in comparison to prediction based on the average activation in the source subjects (blue).




Fig. 3. Ratio between functional alignment label-fusion and predictions based on anatomy for example contrasts for 4 tasks: (1) averaged accuracy when transferring from random individuals (red, right scale), (2) transferring the average Z-score of the source population (blue, left scale).

ratio between the proposed alignment accuracy and the two morphology based comparison methods. Values higher than 1 mean that label-fusion accuracy is higher. Results show that the improvement varies across paradigms and Z-score cut-offs, but that the majority of contrasts exhibit improvement for the proposed label-fusion.

The Heterogeneous Distribution of the Impact across the Cortex. Is the advantage of the functional alignment label-fusion clustered in certain areas? Figure 4 shows the average ratio between prediction accuracy (Dice of $p < 0.001$ areas) of the proposed approach and anatomical alignment mapped to the cortical surface. We only provide values in those areas where the 7 available paradigms exhibit any activation (we chose a more liberal $p < 0.01$ to allow for estimates of this ratio in larger areas). The ratio is different across the cortex, and exhibits strong




Fig. 4. Improvement by functional alignment label-fusion varies across the cortex. Only areas that are covered by at least one of the 7 paradigms with $p < 0.01$ are plotted.

symmetry across the two hemispheres. Areas such as motor cortex, and visual cortex show little improvement by functional alignment, indicating that the location of function varies little across the population. In contrast specifically those areas close to language network, and temporal regions exhibit the most gain.

4 Conclusion

In this paper we propose the prediction of the areas active during a task response in a target subject by multi-atlas label fusion from a source population based on functional alignment. The label-fusion algorithm selects predictors for the Z-score in a target subject among the Z-score at the corresponding position in a population of source subjects. The selection is based on the fit between the aligned target and the source in the embedding space, and extends the rationale behind alignment across individuals using functional regions to functional networks. The embedding represents the functional connectivity observed during rs-fMRI. Results suggest that the resting-state functional connectivity structure is a reliable basis to guide the mapping of task response activations across individuals. It highlights the strong link between the functional connectivity architecture of the brain, and the location of specific functional units that serve individual tasks. The distribution of this gain across the cortex enables insights into the locally varying link between morphology and function.

Acknowledgments. Data provided by the Human Connectome Project, WU-Minn Consortium (1U54MH091657) funded by NIH Blueprint for Neuroscience Research, and McDonnell Center for Systems Neuroscience at WU. Work was partly supported by CDMRP grant PT100120 (SG), by NIH NICHD R01HD067312 and NIH NIBIB NAC P41EB015902, OeNB 14812 and 15929, and EU FP7 2012-PIEF-GA-33003.

References

1. Ashburner, J.: Computational anatomy with the spm software. *Magnetic Resonance Imaging* 27(8), 1163–1174 (2009)
2. Brett, M., Johnsrude, I.S., Owen, A.M.: The problem of functional localization in the human brain. *Nat. Rev. Neurosci.* 3(3), 243–249 (2002)

3. Cardoso, J.F.: Source separation using higher order moments. In: 1989 International Conference on Acoustics, Speech, and Signal Processing, ICASSP 1989, pp. 2109–2112. IEEE (1989)
4. Coifman, R.R., Lafon, S.: Diffusion maps. *App. Comp. Harm. An.* 21, 5–30 (2006)
5. Fischl, B., Salat, D., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., van der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., et al.: Whole Brain Segmentation: Automated Labeling of Neuroanatomical Structures in the Human Brain. *Neuron* 33(3), 341–355 (2002)
6. Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Van Essen, D.C., Jenkinson, M.: WU-Minn HCP Consortium: The minimal preprocessing pipelines for the human connectome project. *NeuroImage* 80, 105–124 (2013). <http://europepmc.org/articles/PMC3720813>
7. Haxby, J.V., Guntupalli, J.S., Connolly, A.C., Halchenko, Y.O., Conroy, B.R., Gobbini, M.I., Hanke, M., Ramadge, P.J.: A common, high-dimensional model of the representational space in human ventral temporal cortex. *Neuron* 72(2), 404–416 (2011)
8. Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W., Smith, S.M.: FSL. *NeuroImage* 62(2), 782–790 (2012)
9. Langerak, T.R., van der Heide, U.A., Kotte, A.N., Viergever, M.A., van Vulpen, M., Pluim, J.P.: Label fusion in atlas-based segmentation using a selective and iterative method for performance level estimation (simple). *IEEE Transactions on Medical Imaging* 29(12), 2000–2008 (2010)
10. Langs, G., Sweet, A., Lashkari, D., Tie, Y., Rigolo, L., Golby, A.J., Golland, P.: Decoupling function and anatomy in atlases of functional connectivity patterns: Language mapping in tumor patients. *NeuroImage* 103, 462–475 (2014)
11. Lombaert, H., Sporring, J., Siddiqi, K.: Diffeomorphic spectral matching of cortical surfaces. In: Gee, J.C., Joshi, S., Pohl, K.M., Wells, W.M., Zöllei, L. (eds.) *IPMI 2013*. LNCS, vol. 7917, pp. 376–389. Springer, Heidelberg (2013)
12. Mueller, S., Wang, D., Fox, M.D., Yeo, B.T.T., Sepulcre, J., Sabuncu, M.R., Shafee, R., Lu, J., Liu, H.: Individual variability in functional connectivity architecture of the human brain. *Neuron* 77(3), 586–595 (2013)
13. Pearson, K.: Principal components analysis. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 6(2), 559 (1901)
14. Sabuncu, M., Singer, B., Conroy, B., Bryan, R., Ramadge, P., Haxby, J.: Function-based intersubject alignment of human cortical anatomy. *Cerebral Cortex* 20(1), 130–140 (2010)
15. Sabuncu, M., Yeo, B., Van Leemput, K., Fischl, B., Golland, P.: A generative model for image segmentation based on label fusion. *IEEE Transactions on Medical Imaging* 29(10), 1714–1729 (2010)
16. Scott, G., Longuet-Higgins, H.: An algorithm for associating the features of two images. *Proceedings: Biological Sciences* 244(1309), 21–26 (1991)
17. Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Uğurbil, K.: The wu-minn human connectome project: an overview. *Neuroimage* 80, 62–79 (2013)
18. Wachinger, C., Golland, P.: Spectral label fusion. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) *MICCAI 2012, Part III*. LNCS, vol. 7512, pp. 410–417. Springer, Heidelberg (2012)