

Hyperalignment of Multi-subject fMRI Data by Synchronized Projections

Raif M. Rustamov^(✉) and Leonidas Guibas

Computer Science Department, Stanford University, Stanford, USA
rustamov@research.att.com

Abstract. Group analysis of fMRI data via multivariate pattern methods requires accurate alignments between neuronal activities of different subjects in order to attain competitive inter-subject classification rates. Hyperalignment, a recent technique pioneered by Haxby and collaborators, aligns the activations of different subjects by mapping them into a common abstract high-dimensional space. While hyperalignment is very successful in terms of classification performance, its “anatomy free” nature excludes the use of potentially helpful information inherent in anatomy. In this paper, we present a novel approach to hyperalignment that allows incorporating anatomical information in a non-trivial way. Experiments demonstrate the effectiveness of our approach over the original hyperalignment and several other natural alternatives.

1 Introduction

Apart from being a fundamental issue in cognitive neuroscience, “the problem of conceptual similarity across neural diversity” [3] has a direct practical manifestation when analyzing fMRI data. Namely, group analysis of fMRI data via multivariate pattern methods requires aligning activations of different subjects. While, pragmatically, the goal of alignment is to attain inter-subject classification (ISC) rates comparable to within subject classification (WSC) rates, ideally, such alignments should take into account both anatomical and functional features of the brain.

Existing spatial alignment approaches are based either purely on anatomical features [6, 12], or on a combination of anatomical features with features extracted from fMRI data, such as activations directly [9] or connectivity derived from activations [4, 5]. However, these approaches do not consistently yield ISC rates comparable to WSC rates [5]. On the opposite end of the spectrum is the recently introduced class of methods summed under the name “hyperalignment” [7, 8, 14]. Hyperalignment essentially finds linear combinations of voxel activations that agree across the subjects, yielding subject specific linear maps (matrices) that transform their activations into a common abstract high-dimensional space. While hyperalignment works well in practice achieving ISC rates on par with or even better than WSC rates, in the current form, it lacks a mechanism for incorporating anatomical information that potentially may lead to even better classification performance.

The goal of this work is to introduce an approach to hyperalignment that allows the use of anatomical information. We start by computing pair-wise (hyper-) alignments between subjects by setting up an optimization problem containing terms involving both anatomical and functional features. Next, we need to aggregate these pair-wise alignments into an overall alignment of all subjects. To achieve this, inspired by the recent work on synchronization [10, 11, 13], we introduce the method of synchronized projections, which yields the final maps of activations into a common space shared between all subjects.

Our approach has a number of advantages. First, while our approach shares the same core idea with hyperalignment – mapping activations into a common space – yet our maps are heavily guided by anatomical information. Second, we do not make any restrictions on the choice of pair-wise alignments; the synchronized projections can be applied more generally to any set of pair-wise alignments that can be expressed as linear maps. Third, experimental results confirm the superiority of synchronized projections in terms of ISC rates over more straightforward approaches that align all subjects to a reference subject or to a floating subject that is iteratively refined.

The paper is organized as follows. We introduce our computation procedure for pair-wise alignments in Sect. 2.1. The main technical contribution, the method of synchronized projections, is described in Sect. 2.2. We present experimental evaluation of our approach on a multi-subject category perception data in Sect. 3.

2 Approach

The input to our algorithm is fMRI data elicited from n_{subj} subjects exposed to a common synchronous stimulus, such as viewing a number of images in the same order. The data for i -th subject is recorded in $n_{\text{TR}} \times n_{\text{vox}}$ matrix X^i , where each row corresponds to a time point, and each column to a voxel in the subject’s brain. Note that each row-vector represents a spatially-varying fMRI activation at some time, and the rows in X^i are ordered consistently across all subjects. On the other hand, the columns – each containing the time course of a particular voxel – are *not* assumed to be in correspondence across the subjects. Since the activations of different subjects are not directly comparable, we cannot train a single multi-voxel pattern classifier that would work for all the subjects at once.

Our goal is to provide a way of computing features/projections of fMRI activations that are consistent across subjects. To this end, our algorithm computes projection matrices – one for each subject – which can be used to map activations of that subject into a common space shared between all subjects.

The algorithm proceeds in two steps. First, for each pair of subjects, we compute a linear map that transports the activation vectors of one subject to the reference frame of the other. In the second step, we compute the projection matrices by setting up an optimization problem which essentially requires the following: the projection of an activation should be roughly the same if one were to transport the activation to another subject and then project. This leads to

a matrix eigenvalue problem for some symmetric positive semi-definite matrix. The details of our construction are provided in the remainder of this section.

2.1 Pair-Wise Alignment Maps

As in original hyperalignment, we will use linear maps to align fMRI responses of different subjects. Thus, the first step of our algorithm is to compute, for every pair of subjects $i, j = 1, \dots, n_{\text{subj}}$, a linear map (matrix) C^{ij} that transforms fMRI activations of subject i to the reference frame of subject j , namely by achieving $X^i C^{ij} \approx X^j$. What this means is that while the voxel activations in two different subjects may not be directly comparable, yet one can make linear combinations from activations of some voxels in the i -th subject’s brain that will be compatible with the activation of a given voxel in the j -th subject’s brain. Thus, the matrix entry C_{pq}^{ij} captures the coefficient with which voxel q of subject i appears in the linear combination for voxel p of subject j .

These alignment matrices are learned from the training data by posing an optimization problem of the following form: $C^{ij} = \arg \min_C \|X^i C - X^j\|_F$, where $\|\cdot\|_F$ is the Frobenius norm. Since the amount of training data is limited, this optimization problem is overly under-determined and needs some kind of regularization. For example, the original hyperalignment [7] requires the matrices C^{ij} to be orthogonal.

Here we propose a different regularization that incorporates the anatomical information. Remember that the brains can be anatomically aligned using a number of approaches; here we will use the Talairach alignment [12]. As a result of such alignment, all of the brain images are placed into a common 3D space, and one can compute the Euclidean distance D_{pq}^{ij} between voxel q of subject i and voxel p of subject j . We now seek the pairwise alignment matrix via the following optimization problem:

$$C^{ij} = \arg \min_C \|X^i C - X^j\|_F^2 + \mu \sum_{p,q} (D_{pq}^{ij} C_{pq})^2. \quad (1)$$

The proposed regularization term has an important advantage over the orthogonality requirement of original hyperalignment. Orthogonality requirement makes it possible for spatially distant voxels to take part in the linear combination for a given voxel, rendering hyperalignment “anatomy free”. Our regularizer, on the other hand, penalizes spatially distant voxels, effectively imposing the prior that the anatomical alignment is not too far from truth.

2.2 Synchronized Projections

The second step of our algorithm uses the pairwise alignment maps in order to construct projection matrices into a d -dimensional common space shared between all subjects. To this end, for each subject i we construct an $n_{\text{vox}} \times d$ matrix P^i , such that the projected activations $X^i P^i$ are consistent between subjects and can be used to train a single multi-voxel pattern classifier that would work for all subjects at once.

The main idea is as follows: if we have an activation row-vector \mathbf{v} of subject i , then it can be transported to subject j by computing $\mathbf{v}C^{ij}$; since the activation before and after transport represents the same stimulus, then their projections (using the respective subject's projection matrix) should be roughly the same: $\mathbf{v}C^{ij}P^j \approx \mathbf{v}P^i$. Since this should hold for all activation vectors and all pairs of subjects, we can setup an optimization problem that minimizes the discrepancies between projections. One way to formalize this is to seek the projection matrices as minimizers of $\sum_{i,j} \|C^{ij}P^i - P^j\|_F^2$, subject to normalization constraints to avoid trivial solutions.

To put the problem into a more familiar form, let us denote by \mathbb{P} the $n_{\text{subj}}n_{\text{vox}} \times d$ matrix obtained by stacking together all of the matrices P^i , $i = 1, \dots, n_{\text{subj}}$. We can rewrite the optimization objective as follows:

$$\sum_{i,j} \|C^{ij}P^i - P^j\|_F^2 = \mathbb{P}^\top \mathbb{L} \mathbb{P}, \quad (2)$$

where \mathbb{L} is $n_{\text{subj}}n_{\text{vox}} \times n_{\text{subj}}n_{\text{vox}}$ matrix. This matrix consists of $n_{\text{subj}} \times n_{\text{subj}}$ blocks L^{ij} of size $n_{\text{vox}} \times n_{\text{vox}}$. Namely, letting I be the $n_{\text{vox}} \times n_{\text{vox}}$ identity matrix, we have

$$L^{ij} = \begin{cases} -(C^{ij} + C^{ji^\top}) & , i \neq j \\ (n_{\text{subj}} - 1)I + \sum_{k, k \neq i} C^{ki^\top} C^{ki} & , i = j \end{cases}$$

As can be seen directly from Eq. (2), \mathbb{L} is a symmetric positive semi-definite matrix. This is a generalized notion of graph Laplacian to the setting where edges are decorated by pair-wise mappings [10, 11, 13].

Since our goal is to minimize $\mathbb{P}^\top \mathbb{L} \mathbb{P}$, we need to impose some constraints on \mathbb{P} in order to avoid trivial solutions. We require the columns of \mathbb{P} to be orthonormal because this leads to an eigenvalue problem. Namely, it is easy to see that then the columns of optimal \mathbb{P} are simply the eigenvectors corresponding to the smallest d eigenvalues of \mathbb{L} , and that the optimal objective value is the sum of the smallest d eigenvalues of \mathbb{L} .

It follows that when the dimension d of the common space is increased from one value to another, all the previous projected coordinates are kept intact and new coordinates are added. In a sense, the projected coordinates are naturally ordered by their corresponding eigenvalues – the smaller the eigenvalue, the stronger is the inter-subject commonality (i.e. the smaller is its contribution to the discrepancy as measured by our objective) captured by the corresponding projected coordinate. Therefore, to obtain a low-dimensional common representation space, we do not need to start with a high-dimensional space and then select the principal component directions as in original hyperalignment [7]. Our eigenvalue based ordering of coordinates provides a more principled criterion than the maximum variance directions criterion of the PCA, because large variance could in fact be due to the absence of commonality along a direction.

3 Results

Our goal is to show the benefit of synchronization in comparison to two other natural approaches, and also to compare it with anatomical alignment and the




Fig. 1. ISC performance comparison of a number of approaches to multi-subject fMRI data alignment.




Fig. 2. Dependence of ISC performance on the dimension d of the common space.

original hyperalignment of [7]. Our experiments are based on category perception (faces and objects) data from [7] that is distributed together with hyperalignment module [1] of PyMVPA package. This dataset is challenging as evidenced, for example, by the inability of a previously introduced generalization of hyperalignment [14] to improve over the original hyperalignment [1]. This can be attributed in part to the small size of the dataset, which limits the number of training samples.

Our evaluation protocol follows directly the one described in [7]: first, for all subjects, all runs except one are used for voxel selection, pair-wise map computation, and determining common space; second, a linear multi-class SVM classifier [2] is trained on these runs of all subjects except one subject; the classifier is tested on the held-out run of the held-out subject. The obtained classifier accuracy is a measure of inter-subject classification success; this accuracy, averaged over held-out subjects and held-out runs, constitutes our performance metric.

Figure 1 shows the performance comparison of five different approaches to alignment of multi-subject fMRI data. The horizontal axis on this graph is the value of the regularization weight μ appearing in the optimization problem for computing pair-wise maps, Eq. (1). Of course, the performances of Talairach alignment (taken directly from the hyperalignment module website [1]) and original hyperalignment [7] (re-implemented in MATLAB for consistency; results are in agreement with hyperalignment module [1]) are independent of the parameter μ . In accordance with [7], voxel selection is done by retaining a fixed number ($n_{\text{vox}} = 200$) of voxels with highest F -scores. For all methods in this figure except synchronization, the dimension of common space is tied (equal) to the number of voxels; for fair comparison, we set $d = n_{\text{vox}}$ for synchronization as well.

The performance of our synchronization approach is also compared to two other natural approaches, labeled “direct” and “iterated direct” in the graph. The direct approach picks one of the subjects, say r , as a reference, and then uses the pair-wise maps C^{ir} to map the activations of all of the subjects to the

frame of this reference subject; these mapped activations are used as features in machine learning step. The iterated direct approach starts out in exactly the same manner, except that mapping process is repeated. Namely, after the first mapping is complete, for each TR, the average of mapped activations are computed, and the new pair-wise maps (from all subjects to the reference subject) are computed to match these averaged activations on reference subject. This iterative process is similar to the original hyperalignment technique of Haxby et al. [7], except that the pair-wise maps are computed using Eq. (1). The performances of direct and iterated direct approaches are averaged over all the reference subject choices.

As indicated by Fig. 1, the synchronization approach consistently improves over the natural alternatives – the direct and iterated direct approaches – using the same type of pair-wise maps. In addition, for some parameter settings, synchronization provides a non-negligible improvement over the original hyperalignment approach of Haxby et al. [7].

Next, we fix the parameter $\mu = 1$, and study the dependence of ISC performance on the dimension d of the common space. Figure 2 shows that even for the dimensionality as low as 10, our approach yields performance competitive with Haxby et al. hyperalignment. This is in agreement with the finding in [7] that keeping a limited number of principal components of the common space is sufficient for obtaining improved ISC rates. However, here we do not need to apply principal component analysis, because the coordinates of common space obtained via our algorithm are already ordered by the degree of inter-subject commonality; see discussion at the end of Sect. 2.2.

Finally, we investigate what happens if one were to change the type of pair-wise alignments used in synchronization. Following the idea of original hyperalignment [7], we require that the pair-wise alignment matrices are orthogonal. More precisely, in optimization problem of Eq. (1) we drop the anatomy based regularizer, and instead require that C^{ij} is orthogonal, which reduces the problem to Procrustes analysis as in [7]. The curve in Fig. 2 labeled “Synch. Haxby et al.” shows the performance of synchronization applied to these new pair-wise maps. It can be seen that the performance is equivalent to the original hyperalignment starting at around dimension $d = 35$, which is in good agreement with the dimension of reduced common space identified in [7] via PCA.

4 Conclusion

We have introduced an approach allowing to inject anatomical information into hyperalignment. Experiments demonstrated the effectiveness of our approach over the original hyperalignment and several other natural alternatives.

Acknowledgments. The authors acknowledge the support of NSF grants FODAVA 808515 and DMS 1228304, AFOSR grant FA9550-12-1-0372, ONR grant N00014-13-1-0341, a Google research award, and the Max Plack Center for Visual Computing and Communications.

References

1. PyMVPA User Manual: Hyperalignment for Between-Subject Analysis. http://dev.py_mvpa.org/examples/hyperalignment.html
2. Chang, C.C., Lin, C.J.: LIBSVM: a library for support vector machines. *ACM Trans. Intell. Syst. Technol.* **2**, 27:1–27:27 (2011). <http://www.csie.ntu.edu.tw/~cjlin/libsvm>
3. Churchland, P.M.: Conceptual similarity across sensory and neural diversity: the Fodor/Lepore challenge answered. *J. Philos.* **95**(1), 5–32 (1998)
4. Conroy, B.R., Singer, B., Haxby, J.V., Ramadge, P.J.: fMRI-based inter-subject cortical alignment using functional connectivity. In: NIPS, pp. 378–386 (2009)
5. Conroy, B.R., Singer, B.D., Guntupalli, J.S., Ramadge, P.J., Haxby, J.V.: Inter-subject alignment of human cortical anatomy using functional connectivity. *NeuroImage* **81**, 400–411 (2013)
6. Fischl, B., Sereno, M.I., Tootell, R.B., Dale, A.M.: High-resolution intersubject averaging and a coordinate system for the cortical surface. *Hum. Brain Mapp.* **8**(4), 272–284 (1999)
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. Lorbert, A., Ramadge, P.J.: Kernel hyperalignment. In: NIPS, pp. 1799–1807 (2012)
9. Sabuncu, M.R., Singer, B.D., Conroy, B., Bryan, R.E., Ramadge, P.J., Haxby, J.V.: Function-based intersubject alignment of human cortical anatomy. *Cereb. Cortex* **20**(1), 130–140 (2010)
10. Singer, A.: Angular synchronization by eigenvectors and semidefinite programming. *Appl. Comput. Harmon. Anal.* **30**(1), 20–36 (2011)
11. Singer, A., Wu, H.T.: Vector diffusion maps and the connection Laplacian. *Comm. Pure Appl. Math.* **65**(8), 1067–1144 (2012)
12. Talairach, J., Tournoux, P.: *Co-planar Stereotaxic Atlas of the Human Brain: 3-D Proportional System: An Approach to Cerebral Imaging*. Thieme, Stuttgart (1988)
13. Wang, F., Huang, Q., Guibas, L.: Image co-segmentation via consistent functional maps. In: *Proceedings of the International Conference on Computer Vision (ICCV)* (2013)
14. Xu, H., Lorbert, A., Ramadge, P., Guntupalli, J., Haxby, J.: Regularized hyperalignment of multi-set fmri data. In: *IEEE Statistical Signal Processing Workshop*, pp. 229–232 (2012)