

Multi-Subject Manifold Alignment of Functional Network Structures via Joint Diagonalization

Karl-Heinz Nenning^{1*}, Kathrin Kollndorfer², Veronika Schöpf²,
Daniela Prayer², and Georg Langs¹

¹Computational Imaging Research Lab, ²Department of Biomedical Imaging and
Image-guided Therapy, Medical University of Vienna,
`karl-heinz.nenning@meduniwien.ac.at`

Abstract. Functional magnetic resonance imaging group studies rely on the ability to establish correspondence across individuals. This enables location specific comparison of functional brain characteristics. Registration is often based on morphology and does not take variability of functional localization into account. This can lead to a loss of specificity, or confounds when studying diseases. In this paper we propose multi-subject functional registration by manifold alignment via coupled joint diagonalization. The functional network structure of each subject is encoded in a diffusion map, where functional relationships are decoupled from spatial position. Two-step manifold alignment estimates initial correspondences between functionally equivalent regions. Then, coupled joint diagonalization establishes common eigenbases across all individuals, and refines the functional correspondences. We evaluate our approach on fMRI data acquired during a language paradigm. Experiments demonstrate the benefits in matching accuracy achieved by coupled joint diagonalization compared to previously proposed functional alignment approaches, or alignment based on structural correspondences.

1 Introduction

Functional magnetic resonance imaging (fMRI) is a valuable neuroscientific tool for deepening the understanding of the functional architecture of the brain, its relationship to cognition, or effects of disease and treatment. fMRI introduces numerous analytical challenges when seeking reliable answers to scientific or clinical questions. A particularly difficult issue is the link between function and anatomy. Relying on anatomical registration to establish correspondence before comparing function across individuals can introduce noise or confounds, if the location of functional regions varies or is correlated with disease. In this paper we propose a functional registration approach based on joint diagonalization of the functional connectivity relationship structure observed in multiple individuals.

A large body of research is focused on addressing fundamental challenges in analysing fMRI data such as low temporal resolution, noise, or motion artifacts. However, the potential disconnect between structure and function is

* This work was supported by the EU (FP7-ICT-2009-5/318068), and OeNB (15929).

only addressed in a small number of studies. It is relevant, because the brains anatomical structure can differ substantially across subjects, and the functional structure can even vary within the same subject [22]. Typical fMRI group studies assume functional correspondence after normalization to a common template space. However, since the localization and extent of functional areas varies across subjects within a certain spatial region, standard group analysis, such as the General Linear Model (GLM) [9] is degraded. Specifically, when studying the impact of damage to the brains functional architecture, the decoupling of changes in position and function is necessary. For instance, a brain tumor alters the anatomical structure, and might lead to functional reorganization, where other regions, possibly even on the opposite hemisphere, compensate for the function of the affected region [7].

In this work we propose a functional registration approach which overcomes the problem of structural and functional variability across individuals or pathologies, by uncoupling function from structure and aligning datasets based on their functional architecture. We propose to use coupled joint diagonalization [6] to learn correspondences, and mutually regularize the embedding representation of individual connectivity structures without using specific functional responses or localizers.

1.1 Related work

Frist, we briefly review relevant related work on functional alignment, and previous work on manifold alignment related to the method proposed in this paper.

Incorporating functional information into anatomical registration has been shown to account for functional variability between subjects, and improve the inter-subject registration. For example, in [4] spatial patterns of functional response are used to guide anatomical registration, [12] uses progressive matching of multi-range functional connectivity patterns for spatial normalization, and [21] performs fine-tuning of anatomical alignment via a nonrigid registration that maximizes inter-subject correlation. Decoupling function from structure can be achieved via manifold learning. It captures the underlying structure of the relationships within a dataset (e.g. of functional signals) by embedding it into a new space where the Euclidean distance represents these relationships. Manifold alignment then allows to align the embeddings of multiple datasets based on similar intrinsic characteristics. Two general approaches to manifold alignment have been proposed in the literature: one-step and two-step approaches. In one-step approaches, multiple datasets are aligned simultaneously, whereas in two-step approaches, alignment is performed on individually created embeddings. Extensions to the basic concept of one-step alignment [10, 25], have successfully been applied to align image datasets of different objects [23], textual datasets [27], or lung images of respiratory motion [1]. In [24], the authors proposed a basic two-step approach consisting of procrustes analysis of embedding maps with application for transfer learning. Two-step manifold alignment approaches were successfully used to map functionally coherent regions across a healthy population [14], to brain tumor patients [15], and for shape matching [18, 17, 16].

The benefit of two-step alignment is computational feasibility, whereas one-step manifold alignment techniques can only be applied on a small number of datasets with few sampling points, or with a sparsity constraint. Moreover, one-step alignment is particularly sensitive to subtle differences in the intrinsic structure of each dataset, affecting their applicability on many real-world problems. However, they have the favorable property to mutually regularize the individual embedding representations. The drawbacks of one-step alignment can be overcome with manifold alignment based on joint diagonalization of Laplacians [6, 13]. The joint diagonalization approach finds the eigenbases of multiple Laplacian matrices of different datasets simultaneously. These eigenbases can be approximated by a small number of eigenvectors of the dataset specific Laplacians, reducing computational complexity of aligning multiple datasets with a large number of sampling points.

Most recent manifold alignment techniques require at least a small subset of correspondences between the datasets to be known. However, in the neurological context at the center of this paper, this information is generally not available, or not reliable. In this case, unsupervised manifold alignment, i.e. alignment without known correspondences, establishes a similarity measure between the datasets. As proposed in [26], correspondence weights are established based on a local pattern matching technique, and successfully applied to aligned protein data. An unsupervised two-step alignment approach was introduced in [20], where distances between sample points in the embedding space are used to establish parameterized distances curves and subsequently similarities are calculated between those. However, such correspondence estimation techniques are only reliable when there are distinct comparable intrinsic structures present, e.g. image datasets of different objects with comparable rotations.

1.2 Contribution

We propose a manifold alignment approach to build an atlas of functional network characteristics establishing functional correspondences between individuals. Diffusion maps [3] encode the functional architecture of each individual and initial correspondences are established via a two-step manifold alignment approach [15]. A subset of correspondence estimates between dense areas in this initial alignment is used to initiate coupled joint diagonalization that simultaneously finds common eigenbases of the functional connectivity structures and refined correspondences across all subjects [6].

2 Method

Diffusion maps [3] are used to encode the functional connectivity structure of each individual. Based on these subject-specific embeddings initial correspondences between fMRI datasets are established by a two-step manifold alignment approach [15]. It rearranges the spectral components based on spatial constraints, resulting in a population atlas of functional structure. Initial correspondence estimate pairs between datasets are drawn randomly from dense areas in

this joint atlas. They initiate coupled joint diagonalization [6]. This results in refined correspondences, and a common eigenbasis, which represent the aligned underlying connectivity structure of a study group.

2.1 Spectral representation of functional connectivity

We embed the whole-brain functional connectivity structure of each individual by diffusion maps [3], resulting in a functional geometry reflecting the functional organization of the brain. The fMRI data $\mathbf{I} \in \mathbb{R}^{T \times N}$ consists of N voxels \mathbf{v}_i observed at T time points. The functional architecture of a brain state can be modeled as a connectivity graph G with vertices V representing the voxels, and edges E between these vertices. Let W denote the affinity matrix assigning weights to the edges, where W_{ij} is the Pearson correlation coefficient between the time-series of vertex i and j . The spectral representation of the functional connectivity structure can be established via eigendecomposition of the normalized Laplacian matrix L , defined as $L = D^{-1/2}WD^{-1/2}$, where D is the diagonal matrix of node degrees, i.e. $d_i = \sum_j W(i, j)$. The normalized Laplacian is a symmetric version of the random walk matrix $L = D^{-1}W$, which is scaling invariant, and can be viewed as a diffusion process, where the transition probabilities for a diffusion time t between nodes can be expressed by L^t . Eigendecomposition of L results in eigenvalues Λ and corresponding eigenvectors Φ , where $\lambda_1 = 1 > \lambda_2 > \dots > \lambda_n$. The dataset can be represented by a diffusion map with new embedding coordinates Ψ , where each vertex i in the original graph is represented by

$$\Psi_i^t = (\lambda_1^t \psi_1(x), \lambda_2^t \psi_2(x), \dots, \lambda_d^t \psi_d(x)), \quad (1)$$

for every data point x in d dimensions of the embedding space. This translates the functional organization of a brain to a functional geometry Ψ , where the functional relations of the fMRI signals are captured by the diffusion distance D_t [15]. The diffusion distance $D_t(i, j)$ is defined as the probability of traveling from node i to node j in t steps, by considering all possible paths between these two nodes. The Euclidean distance in this new embedding space approximates the diffusion distance

$$D_t(x_i, x_j) = \|\psi_t(x_i) - \psi_t(x_j)\|^2. \quad (2)$$

Thus, the functional relationship between voxels of a brain is translated to Euclidean distances in a new embedding space, capturing the functional geometry of the dataset. This uncouples function from anatomy, and is the basis for functional registration of corresponding functional structures across individuals [15].

2.2 Initial correspondence estimates across individuals

We aim at aligning functionally equivalent areas across individuals. Assuming that for a majority of functional areas the spatial difference is small we use

spatial correspondence for initialization. We perform an initial two-step manifold alignment following the initialization procedure in [15]. First, we align brain morphology to a common template [8]. Then, we orthonormally align individual embedding maps representing each subject based on correspondence defined by spatial position in the common template. For every subject specific embedding Ψ_s and template embedding Ψ_t , an orthonormal transformation matrix Q_{st} is calculated, accounting for rotations, sign flips and reordering of the spectral dimensions. Q_{st} is defined as

$$Q_{st} = \underset{Q}{\operatorname{argmin}} \left(\sum_{c=1}^C w_c \| Q\Psi_{sc} - \Psi_{tc} \|^2 \right), \quad (3)$$

where Ψ are the embeddings of s and t and w is the Euclidean distance between a pair c of corresponding voxels. After this initial alignment of embeddings, we perform nonrigid point cloud registration to refine the alignment and close potential gaps between the embedded sampling point distributions by Coherent Point Drift [19].

While preserving the functional relationship between brain regions, diffusion maps typically encode a distinct functionally coherent cluster of brain regions along a specific dimension, far from the origin in the leading dimensions of the embedding space. We use this property to obtain initial pairwise point-to-point correspondence estimates between two orthonormally aligned diffusion maps Ψ'_s and Ψ_t , resulting in n correspondence pairs $C = \{\langle i, j(i) \rangle\}_{i=1, \dots, n}$ with $j(i) = \operatorname{argmin}_j \| \Psi_t(i) - \Psi'_s(j) \|^2$. For each pair of embeddings, a subset of q correspondence pairs C is selected based on their distance from the origin and used as initial coupling for joint diagonalization, described in the following section. Note that the purpose of this initial alignment is only to determine a subset of initial correspondence estimates. Sub-sequent joint diagonalization is based on the individual Laplacians and the correspondence pairs.

2.3 Joint diagonalization

Joint diagonalization simultaneously finds the common eigenbases of the Laplacians of multiple datasets [2, 13, 6]. In theory, if multiple datasets share similar intrinsic structures, their Laplacians have similar eigenbases with variations in rotations, coefficient sign or ordering of spectral components. Joint diagonalization results in coupled eigenbases, which make the eigenbases of each individual dataset consistent. The basic joint diagonalization problem with full coupling can be formulated as an optimization problem of minimizing off-diagonal elements:

$$\min_V \sum_{i=1}^M \mathbf{off}(V^T L_i V) \text{ with } V^T V = I, \quad (4)$$

where M is the number of different datasets, V is their common eigenbases, L_i are the dataset specific Laplacians. The off-diagonal penalty $\mathbf{off}(X)$ is typically defined as the sum of squared off diagonal elements $\mathbf{off}(X) = \| X - \operatorname{diag}(X) \|_F^2$.

This problem can be solved with the generalized Jacobi method (JADE method) [2]. It assumes full coupling with an equal amount of samples for every dataset, i.e. every sampling point in a dataset has a corresponding point in all other datasets. In our case we cannot assume identical numbers of sampling points or complete given correspondences between datasets. Instead we only use a sub-set of paired vertices to guide the joint diagonalization.

Coupled diagonalization allows to overcome these limitations [6, 13]. Only a subset of sparse point-wise correspondences is used to establish a coupling relation F . F_i is a $n_i \times q$ coupling matrix, where n_i is the number of data points and q is the number of correspondences. Based on this sparse subset, eigenbases V_i are established, such that $V_i^T L_i V_i$ is approximately diagonal and the eigenbases correspond for corresponding samples $F_i V_i \approx F_j V_j$. This results in multiple dataset specific eigenbases, contrary to the original joint diagonalization approach [2].

To incorporate the coupling between datasets, the formulation in Equation 4 can be rewritten as

$$\min_{V_i} \sum_{i=1}^M \text{off}(V_i^T L_i V_i) + \mu \sum_{i,j=1}^M \| F_i^T V_i - F_j^T V_j \|_F^2 \quad \text{with } V_i^T V_i = I, \quad (5)$$

where V_i are the dataset specific eigenbases, L_i the data specific Laplacians and μ is a coupling weight [6]. The sparse coupling between two datasets i and j is defined by corresponding points k_i and k_j and encoded as $F_i(k_i, q) = 1$ and $F_j(k_j, q) = 1$, with zeros elsewhere in the q^{th} column. However, the coupling is not restricted to such point-wise coupling, and can contain any correspondence weighting. We use the leading k dimensions of the embedding, since they hold the majority of information. This reduces the computational complexity to an optimization problem with $k^2 M$ variables, instead of $\sum^M n_i^2$ variables for full eigenbases. Then, the formulation in Equation 5 can be rewritten as

$$\min_{A_i} \sum_{i=1}^M \| A_i^T \bar{L}_i A_i \|_F^2 + \mu \sum_{i,j=1}^M \| F_i^T \bar{U}_i A_i - F_j^T \bar{U}_j A_j \|_F^2 \quad \text{with } A_i^T A_i = I, \quad (6)$$

where \bar{U}_i are the first k eigenvectors of L_i and A_i is a $k \times k$ matrix with linear combination coefficients to be estimated [6, 13]. The specific eigenbases can be reconstructed with $\bar{V}_i = \bar{U}_i A_i$.

To solve Equation 6, we employ an optimization on a Stiefel manifold with orthogonality constraints, introduced in [28]. Hereby, an unconstrained problem is solved by building the orthogonality constraint into the optimization method in the form of a projected descend. Following [6, 13], the cost function is rewritten to

$$\min_{A_i} \| A_i^T \bar{L}_i A_i \|_F^2 + \mu \sum_{j=1}^M \| F_i^T \bar{U}_i A_i - F_j^T \bar{U}_j A_j \|_F^2, \quad (7)$$

and solved for all A_i alternatingly, with the gradients for the off diagonal penalty given by

$$\nabla_{A_i} \| A_i^T \bar{L}_i A_i \|_F^2 = 4(\bar{L}_i A_i A_i^T \bar{L}_i A_i - \bar{L}_i A_i \bar{L}_i), \quad (8)$$




Fig. 1: The encoded information is reflected by the eigenvalues of the Laplacians. They indicate a dominance of the task structure, and the relation between task-negative (blue) and task-positive (red) networks can be observed as a smooth transition in the diffusion maps.

and the gradient for the coupling term given by

$$\nabla_{A_i} \| F_i^T \bar{U}_i A_i - F_j^T \bar{U}_j A_j \|_F^2 = 2 \bar{U}_i^T F_i (F_i^T \bar{U}_i A_i - F_j^T \bar{U}_j A_j). \quad (9)$$

Coupled joint diagonalization can be viewed as a mixture between a one-step and a two-step manifold alignment approach. First, like in two-step alignment techniques, Laplacians are computed for each dataset. The eigenvectors and eigenvalues of each individual Laplacian are then used to find common eigenbases, following the principle of one-step manifold alignment approaches.

After coupled joint alignment every voxel in every individual is represented by a point in the joint embedding space. The final correspondences are determined by nearest neighbor search. For every voxel i in a dataset X , a corresponding voxel \hat{j} in a dataset Y can be found by the minimal Euclidean distance between the aligned embedding maps Ψ_X and Ψ_Y : $\hat{j} = \operatorname{argmin}_j \| \Psi_X(i) - \Psi_Y(j) \|^2$.

3 Evaluation

3.1 Data

We evaluate the alignment algorithm on fMRI data collected in study with a language related block design paradigm, comprising 12 healthy individuals. During the activation blocks, subjects are instructed to read nouns and mentally generate associated verbs. During the baseline block hash signs are presented. The fMRI was acquired on a 3T unit (Philips Medical Systems, Best, The Netherlands) equipped with a 12 channel head coil, using a single shot echo planar imaging (EPI) sequence with TR/TE 3000/35 ms, $128 \times 128 \times 32mm$ Matrix, voxel size $1.8 \times 1.8 \times 4mm^3$, 100 volumes and a duration of 5 minutes comprising 5 task - rest alterations. Standard fMRI preprocessing is performed with FSL [11] and Freesurfer [8], including co-registration, motion-correction and removal of time-points with extensive motion. The fMRI data is normalized to a common space via registration to the Freesurfer average cortical surface template in MNI space [5]. The time-series are projected onto this cortical surface and are normalized to have zero mean and unit variance. The preprocessing results in a total of 4718 cortical brain regions used for further analysis.

3.2 Evaluation

We evaluate if the alignment accurately matches areas that exhibit the same language task response activation. A standard GLM [9] activation analysis is performed on the fMRI data of each individual, and the obtained β parameters are transformed to z-scores. Since manifold alignment is performed without any direct activation information, we use the activation analysis as reference for validation of matching accuracy. In a leave-one-out cross validation, we label the brain regions of a target subject with the z-scores of the closest region in the manifold atlas of the remaining subjects after alignment ('prediction'). We compare the z-scores of the subject specific GLM analysis on the target subject with the predicted z-scores by applying a cut-off and calculating the Dice-coefficient between predicted and measured z-scores. To illustrate the relationship between accuracy and z-score, we report results over a range of z-score values. We compare the alignment of two variants of the proposed approach and 2 alternative approaches. (1) Coupled joint diagonalization using initial correspondences in high density regions of the embedding space (*DG*), (2) coupled joint diagonalization using initial correspondences randomly distributed across the space (*DGrand*), (3) the orthonormal Procrustes alignment of functional connectivity manifolds based on spatial correspondences as used for initialization in [15] (*MA ortho*), and we establish correspondence based on the MNI coordinates of cortical vertices, i.e, relying entirely on the anatomical registration (*MNI*). We perform embedding of the fMRI data of each individual with a correlation threshold of zero, retaining only positive correlation values between brain regions.

3.3 Results: alignment of language networks

Examples of embedded functional network structures are shown in Figure 1 with the first two dimensions of the diffusion maps. The activation z-score is visualized by the color. Note that z-score information does not enter embedding or alignment. While the overall structure varies, the functional architecture is comparable: a common smooth transition between task-positive regions (red) and task-negative regions (blue). This can be explained by the fact that the intrinsic functional structure of each individual, modeled via correlation between time-series, is dominated by the task paradigm. This is supported by the eigenvalues of the Laplacians (Figure 1 left), which reveal a rapid decrease in information encoded along each additional dimension, consistent across the study group. Therefore, we limit our further analysis only to the first 5 dimensions of the diffusion map embeddings.

Figure 2a shows the alignment accuracy of *MA ortho* using 1, 3 and 5 eigenvectors, analogously Figure 2b shows the results for *DG*. A more liberal threshold results in more voxels labeled as activated, which implicates a trend towards a higher dice score. A strict univariate GLM analysis might miss some voxels involved in a task, which can be found via network matching. Overall, diffusion maps with 3 dimensions achieve the best alignment results, which can be explained by the basic task-driven intrinsic functional connectivity structure of

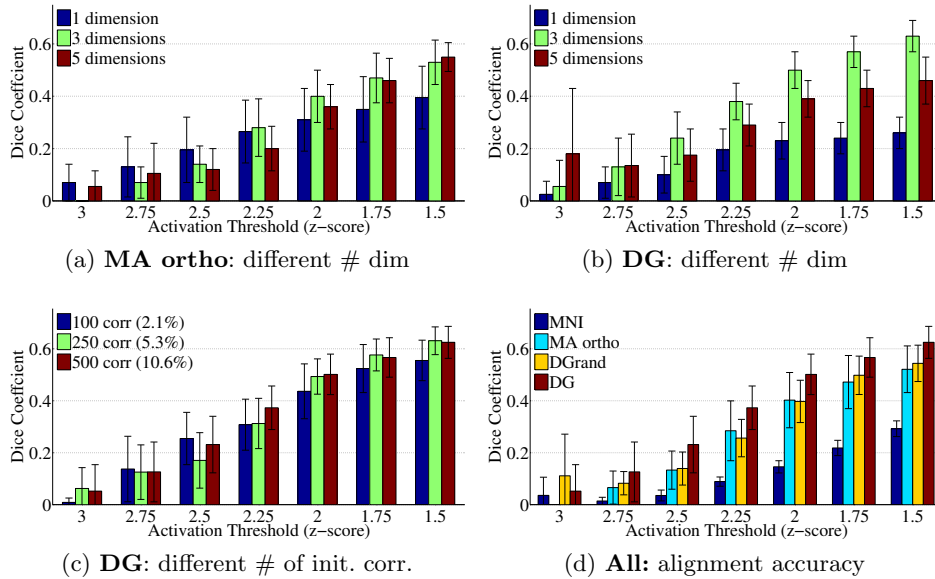


Fig. 2: Alignment accuracy: **(a)** Two step orthonormal manifold alignment for different z-score cut-offs and numbers of eigenvectors. **(b)** Alignment results for *DG* and 500 correspondences, shows highest alignment accuracy with the first 3 eigenvectors, although the first 5 eigenvectors are more sensitive at higher z-score thresholds. **(c)** As expected, *DG* with 500 correspondences (10.5%) yields overall best results, however comparable performance can be achieved with only 100 couplings (2.1%). **(d)** The overall comparison shows that our application of joint diagonalization yields better results compared to the two-step manifold alignment (*MA ortho*). The utilization of weighted correspondences (*DG*), yields better results compared to joint diagonalization with random pairs of correspondences (*DGrand*). As expected, performing transfer learning with the closest region on the cortex (*MNI*) performs underwhelming.

the language dataset. This is especially noticeable for *DG*, indicating that 3 common eigenbases are shared across the dataset. The impact on the number of initial correspondences during coupled joint diagonalization is shown in Figure 2c. While 500 correspondences, 10.6% of points in the datasets, yield better results, even with 100 correspondences, which correspond to only 2.1% of sampling points, matching accuracy is comparable. Finally Figure 2d compares all 4 alignment approaches. The coupled diagonalization approach (*DG*) is more accurate than the orthonormal two-step manifold alignment (*MA ortho*). Moreover, we observe the benefit of drawing initial correspondences during coupled joint diagonalization from dense regions in the atlas. When drawing random correspondences between closest points (*DGrand*), the joint diagonalization approach performs similar to the two-step manifold alignment method with spatial con-

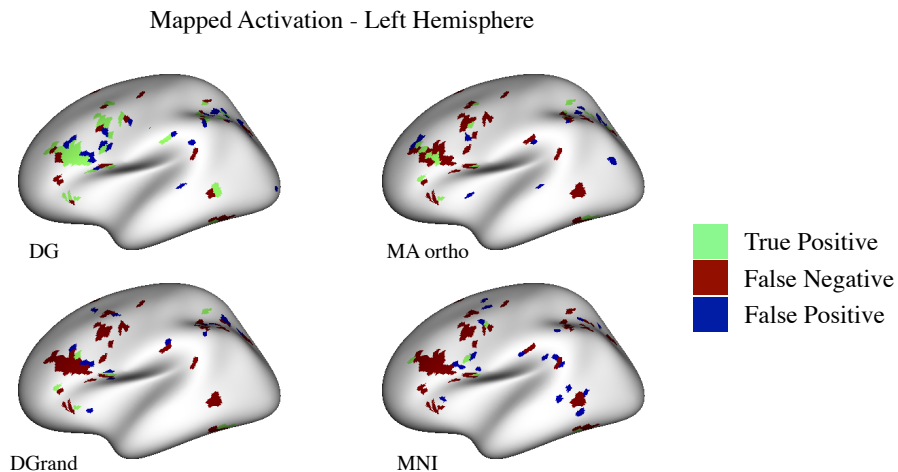


Fig. 3: A single subject example of mapped activation with a z-score threshold of 2.5 on the left hemisphere. Our application of coupled diagonalization (DG) maps core regions related to language function successfully, with false positive and false negative mappings on the border areas.

straints. An example of matched activation maps is illustrated in Figure 3. For one subject and a z-score threshold of 2.5 we compare the actual activation with the predicted activation. The core regions of the language related networks are mapped successfully, with false positive and false negative mappings located on the edges of these regions, which is analogous to the typical gaussian structure of GLM activation analysis in fMRI studies. Analogous to the quantitative results the coupled joint diagonalization (DG) most accurately predicts the language regions.

4 Conclusion

We propose coupled joint diagonalization to perform alignment of embedding representations of functional network structures. Sparse coupling reduces the computational complexity of the problem and facilitates manifold alignment for multiple datasets with a large amount of sampling points. In contrast to two-step approaches, joint alignment allows to draw upon multiple subject data simultaneously to learn a robust joint embedding. Estimating a subset of initial correspondences between functional datasets is challenging, and sub-sequent refinement and reestimation of correspondences after joint diagonalization exploits shared functional connectivity architecture, in order to establish correspondence despite spatial variability of functionally active regions. Manifold alignment allows for spectral embedding and clustering, or corresponding transfer learning on a group level, by taking advantage of common intrinsic structure present in multiple datasets. The proposed manifold alignment is suited to support the

analysis of functional characteristics of the brain by accounting for functional variability, and has the potential to quantify the spatial properties of functional reorganization.

References

1. C.F. Baumgartner, C. Kolbitsch, J.R. McClelland, D. Rueckert, and A.P. King. Groupwise simultaneous manifold alignment for high-resolution dynamic mr imaging of respiratory motion. In J.C. Gee, S. Joshi, K.M. Pohl, W.M. Wells, and L. Zillei, editors, *Information Processing in Medical Imaging*, volume 7917 of *Lecture Notes in Computer Science*, pages 232–243. Springer Berlin Heidelberg, 2013.
2. Jean-Francois Cardoso and Antoine Souloumiac. Jacobi angles for simultaneous diagonalization. *SIAM J. Matrix Anal. Appl.*, 17(1):161–164, January 1996.
3. Ronald R. Coifman and Stphane Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006. Special Issue: Diffusion Maps and Wavelets.
4. Bryan R. Conroy, Benjamin D. Singer, J. Swaroop Guntupalli, Peter J. Ramadge, and James V. Haxby. Inter-subject alignment of human cortical anatomy using functional connectivity. *NeuroImage*, 81(0):400 – 411, 2013.
5. A.C. Evans, D.L. Collins, S.R. Mills, E.D. Brown, R.L. Kelly, and T.M. Peters. 3d statistical neuroanatomical models from 305 mri volumes. In *Nuclear Science Symposium and Medical Imaging Conference, 1993., 1993 IEEE Conference Record.*, pages 1813–1817 vol.3, Oct 1993.
6. D. Eynard, A. Kovnatsky, M.M. Bronstein, K. Glashoff, and A.M. Bronstein. Multimodal manifold analysis by simultaneous diagonalization of laplacians. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, PP(99):1–1, 2015.
7. Javier Fandino, Spyros S. Kollias, Heinz G. Wieser, Anton Valavanis, and Yasuhiro Yonekawa. Intraoperative validation of functional magnetic resonance imaging and cortical reorganization patterns in patients with brain tumors involving the primary motor cortex. *Journal of Neurosurgery*, 91(2):238–250, 1999.
8. Bruce Fischl. Freesurfer. *NeuroImage*, 62(2):774 – 781, 2012. 20 YEARS OF fMRI 20 YEARS OF fMRI.
9. K.J. Friston, A.P. Holmes, K. J. Worsley, J.-P. Poline, C. D. Frith, and R. S. J. Frackowiak. Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, 2(4):189–210, 1994.
10. Ham, Lee, and Saul. Semisupervised alignment of manifolds. In *Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics*, 2005.
11. Mark Jenkinson, Christian F. Beckmann, Timothy E.J. Behrens, Mark W. Woolrich, and Stephen M. Smith. Fsl. *NeuroImage*, 62(2):782 – 790, 2012. 20 YEARS OF fMRI 20 YEARS OF fMRI.
12. Di Jiang, Yuhui Du, Hewei Cheng, Tianzi Jiang, and Yong Fan. Groupwise spatial normalization of fmri data based on multi-range functional connectivity patterns. *NeuroImage*, 82(0):355 – 372, 2013.
13. A. Kovnatsky, M. M. Bronstein, A. M. Bronstein, K. Glashoff, and R. Kimmel. Coupled quasi-harmonic bases. *Computer Graphics Forum*, 32(2pt4):439–448, 2013.
14. Georg Langs, Danial Lashkari, Andrew Sweet, Yanmei Tie, Laura Rigolo, AlexandraJ. Golby, and Polina Golland. Learning an atlas of a cognitive process in its

- functional geometry. In G. Székely and H.K. Hahn, editors, *Information Processing in Medical Imaging*, volume 6801 of *Lecture Notes in Computer Science*, pages 135–146. Springer Berlin Heidelberg, 2011.
15. Georg Langs, Andrew Sweet, Danial Lashkari, Yanmei Tie, Laura Rigolo, Alexandra J. Golby, and Polina Golland. Decoupling function and anatomy in atlases of functional connectivity patterns: Language mapping in tumor patients. *NeuroImage*, 103(0):462 – 475, 2014.
 16. H. Lombaert, L. Grady, J.R. Polimeni, and F. Chériet. Focusr: Feature oriented correspondence using spectral regularization—a method for precise surface matching. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(9):2143–2160, Sept 2013.
 17. Herve Lombaert, Leo Grady, Jonathan R. Polimeni, and Farida Chériet. Fast brain matching with spectral correspondence. In Gbor Székely and Horst K. Hahn, editors, *Information Processing in Medical Imaging*, volume 6801 of *Lecture Notes in Computer Science*, pages 660–673. Springer Berlin Heidelberg, 2011.
 18. D. Mateus, R. Horaud, D. Knossow, F. Cuzzolin, and E. Boyer. Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8, June 2008.
 19. A. Myronenko and Xubo Song. Point set registration: Coherent point drift. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(12):2262–2275, 2010.
 20. Y. Pei, Fengchun Huang, F. Shi, and Hongbin Zha. Unsupervised image matching based on manifold alignment. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34(8):1658–1664, Aug 2012.
 21. Mert R. Sabuncu, Benjamin D. Singer, Bryan Conroy, Ronald E. Bryan, Peter J. Ramadge, and James V. Haxby. Function-based intersubject alignment of human cortical anatomy. *Cerebral Cortex*, 20(1):130–140, 2010.
 22. Stephen M. Smith, Christian F. Beckmann, Narender Ramnani, Mark W. Woolrich, Peter R. Bannister, Mark Jenkinson, Paul M. Matthews, and David J. McGonigle. Variability in fmri: A re-examination of inter-session differences. *Human Brain Mapping*, 24(3):248–257, 2005.
 23. M. Torki, A. Elgammal, and Chan Su Lee. Learning a joint manifold representation from multiple data sets. In *Pattern Recognition (ICPR), 2010 20th International Conference on*, pages 1068–1071, Aug 2010.
 24. Chang Wang and Sridhar Mahadevan. Manifold alignment using procrustes analysis. In *Proceedings of the 25th International Conference on Machine Learning, ICML '08*, pages 1120–1127, 2008.
 25. Chang Wang and Sridhar Mahadevan. A general framework for manifold alignment. *AAAI Fall Symposium Series*, 2009.
 26. Chang Wang and Sridhar Mahadevan. Manifold alignment without correspondence. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI'09*, pages 1273–1278, 2009.
 27. Chang Wang and Sridhar Mahadevan. Multiscale manifold learning. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, 2013.
 28. Zaiwen Wen and Wotao Yin. A feasible method for optimization with orthogonality constraints. *Mathematical Programming*, 142(1-2):397–434, 2013.