



HAL
open science

Local Optimal Transport for Functional Brain Template Estimation

Thomas Bazeille, Hugo Richard, Hicham Janati, Bertrand Thirion

► **To cite this version:**

Thomas Bazeille, Hugo Richard, Hicham Janati, Bertrand Thirion. Local Optimal Transport for Functional Brain Template Estimation. IPMI 2019 - 26th International Conference on Information Processing in Medical Imaging, Jun 2019, Hong Kong, China. 10.1007/978-3-030-20351-1_18 . hal-02278663

HAL Id: hal-02278663

<https://hal.science/hal-02278663>

Submitted on 17 Oct 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Local optimal transport for functional brain template estimation

T. Bazeille¹, H. Richard¹, H. Janati^{1,2}, B. Thirion¹

¹ INRIA, CEA Neurospin, France

² CREST ENSAE

Abstract. An important goal of cognitive brain imaging studies is to model the functional organization of the brain; yet there exists currently no functional brain atlas built from existing data. One of the main roadblocks to the creation of such an atlas is the functional variability that is observed in subjects performing the same task; this variability goes far beyond anatomical variability in brain shape and size. Function-based alignment procedures have recently been proposed in order to improve the correspondence of activation patterns across individuals. However, the corresponding computational solutions are costly and not well-principled. Here, we propose a new framework based on optimal transport theory to create such a template. We leverage entropic smoothing as an efficient means to create brain templates without losing fine-grain structural information; it is implemented in a computationally efficient way. We evaluate our approach on rich multi-subject, multi-contrasts datasets. These experiments demonstrate that the template-based inference procedure improves the transfer of information across individuals with respect to state of the art methods.

Keywords: Brain · Atlas Inference · fMRI · Functional Alignment.

1 Introduction

Brain anatomical and functional variability There is a very large biological variability between human brains, influenced by both genetic, developmental as well as environmental factors. It results in conspicuous anatomical differences, that have traditionally been characterized as diffeomorphic transformations and compensated by dedicated algorithms (see [1] for an application-oriented overview). Functional imaging, such as functional Magnetic Resonance Imaging (fMRI) also detects variations of activity across brain regions and thus provides a potential marker of variation in the functional organization of the brain i.e. the involvement of different neural modules when performing a given task. However, so far very little work has been dedicated to leverage fMRI contrasts to learn better correspondences between brains.

Capturing functional variability in brain responses The implicit tenet of most analyses is that only anatomical information can be leveraged to estimate accurate correspondences between brains, while the unmatched functional variability

is treated as a residual [2]. It might, at best, be reduced by smoothing; however, smoothing has been shown to create artificial overlaps of different functional territories, yielding biased models of functional organization (see e.g. [3]).

Meanwhile, the progressive improvement of fMRI contrast-to-noise ratio and resolution, as well as the development of *deep phenotyping* approaches, in which subjects were scanned under many conditions (e.g. [4, 5]) opens a new perspective, namely that of using the acquired functional contrasts to learn inter-individual correspondences. Diffeomorphic registrations were proposed to match functional variability across individuals [6] or features obtained in resting-state fMRI data [7], thus assuming a strong spatial regularity of the correspondence. This diffeomorphic constraint was abandoned in other works, in particular in the popular *hyperalignment* framework [8, 9], that attempts to identify and match activation patterns in the visual cortex without imposing strong regularity conditions. Following this line of research, direct functional correspondences obtained from resting-state fMRI data have been used to directly map functional areas across individuals [10, 11]. While Hyperalignment is currently used to pool data across individuals to boost brain activity classification (see for e.g. [12]), it is not applied in more traditional group studies. Even if free correspondences potentially yield good alignment for two individuals, this framework has a hard time producing a subject independent model — a template.

Specifically, two core questions are still open: *i)* What family of transformations should be allowed for cross-subject matching ? In particular, should the transformations have some spatial regularity ? *ii)* How to estimate templates under weakly constrained deformation models ?

Our contribution Following the intuitions that lead to the development of hyperalignment — map features across individuals under weak spatial constraints — we develop a novel framework based on optimal transport (OT), that is well suited to the estimation of explicit transformation models. In particular, we leverage the regularization procedures tied to this problem, namely entropic smoothing; we introduce spatial constraints in the matching by relying on local regions. Locality also clearly benefits to computational efficiency.

Second, we introduce an efficient template estimation procedure based on this OT deformation model. We evaluate the models by their ability to predict unseen data in new subjects, and benchmark the OT-based model against alternatives, including diffeomorphic registration, on two datasets that contain rich functional information in groups of subjects.

2 Theory

Notations Let $p \in \mathcal{N}$, we denote $[p]$ the set of integers from 1 to p . Let $\mathbf{x} \in \mathbb{R}^d$, $\delta_{\mathbf{x}}$ will denote the Dirac mass at location \mathbf{x} . Given a brain region comprising p voxels, we consider the d -dimensional signals observed in these voxels $\mathbf{x} = \{\mathbf{x}_1, \dots, \mathbf{x}_p\}$. Here, these d -dimensional signals correspond to d activation maps observed in a given subject. We denote by \mathbf{X} the $p \times d$ matrix obtained by

concatenating the vectors in \mathbf{x} . $\|\cdot\|_F$ denotes Frobenius norm, and $tr(\cdot)$ the trace operator.

Correspondences from an optimal transport geometry perspective Let us consider the set \mathbf{x}^s of functional signals in a given subject s . Following the intuitions of [10], the p vectors $\{\mathbf{x}_1^s, \dots, \mathbf{x}_p^s\}$ together make up a measure μ^s lying on a latent manifold \mathcal{F}_s embedded in \mathbb{R}^d . The discrete measure μ^s with positive weights $\mathbf{w}^s > 0$ and support $\{\mathbf{x}_1^s, \dots, \mathbf{x}_p^s\} \in \mathcal{F}_s$ is defined as $\mu^s = \sum_{i=1}^p \mathbf{w}_i^s \delta_{\mathbf{x}_i^s}$. Note that in the present framework $\mathbf{w}^s = [\frac{1}{p} \dots \frac{1}{p}]$.

The difference between two measures $\mu^s = \sum_{i=1}^p \mathbf{w}_i^s \delta_{\mathbf{x}_i^s}$ and $\mu^t = \sum_{i=1}^q \mathbf{w}_i^t \delta_{\mathbf{x}_i^t}$, reflects the differences between individuals s and t . However, for $s \neq t$, the support of the \mathcal{F}_s and \mathcal{F}_t manifolds are distinct in general: two subjects do not exhibit the same set of responses, due to intrinsically different brain organization. Directly computing Kullback-Leibler divergence between μ^s and μ^t is useless, as non-coincident support leads to infinite values; fixing this mismatch by smoothing induces a loss of information. By contrast, the Wasserstein distance between μ^s and μ^t is well-defined [13]. In this framework, distance evaluation is tightly linked to functional alignment, as it is formulated as the task of finding an optimal coupling $\mathbf{R}^* \{\mathbf{x}_1^s, \dots, \mathbf{x}_p^s\} \rightarrow \{\mathbf{x}_1^t, \dots, \mathbf{x}_q^t\}$, $\mathbf{R} \in \mathbb{R}_+^{p \times q}$. Enforcing signal conservation and optimality of the alignment cost $\mathbf{C}(\mathbf{x}^s, \mathbf{x}^t)$ yields:

$$\mathbf{R}^* = \min_{\mathbf{R}} \sum_{i,j} \mathbf{R}_{i,j} \mathbf{C}(\mathbf{x}^s, \mathbf{x}^t)_{i,j} \quad (1)$$

$$\text{s.t.} \left(\sum_j \mathbf{R}_{i,j} \right)_i = w_i^s \text{ and } \left(\sum_i \mathbf{R}_{i,j} \right)_j = w_j^t \quad (2)$$

If both subjects functional data share a common number of voxels (i.e. $p = q$), and we search for a deterministic coupling (where each source has only one target voxel), this falls back to the optimal matching problem and makes \mathbf{R} a permutation matrix minimal with respect to $\mathbf{C}(\mathbf{x}^s, \mathbf{x}^t)$, that can be calculated through the Hungarian algorithm [14].

Kantorovich relaxation In general, admissible couplings in (1) are probabilistic, i.e. they can split the mass of a source location towards several target locations and their coefficients $\mathbf{R}_{i,j}$ encode the mass flow between points \mathbf{x}_i^s and \mathbf{x}_j^t [15].

One possibility, explored in [16] in the context of the MEG inverse problem, is to match the multi-dimensional signal distributions across subjects using a spatial-distance-based cost —matching preferentially voxels that are close in 3D space. The authors had to resort to unbalanced OT to deal with the variation of mass between subjects. With this formulation, μ^s and μ^t should be positive and of same norm, whereas functional signals usually have diverse amplitude across subjects and exhibit both positive and negative values, in proportions that vary across individuals. This approach requires to introduce several additional parameters to deal with unbalanced transport model. We chose another path here,

closer to the hyperalignment concept [8], which is to transport voxels according to the functional information they carry. To do so, we use the discrepancy of the functional features as a cost function between voxels :

$$\forall i, j \in [p] \times [q], \mathbf{C}(\mathbf{x}^s, \mathbf{x}^t)_{i,j} = \|\mathbf{x}_i^s - \mathbf{x}_j^t\|_2^2$$

and couple the input measures μ^s and μ^t , where all voxels have a constant weight, respectively $1/p$ and $1/q$.

Entropic regularization. We define the entropy of a coupling as

$$h(\mathbf{R}) = - \sum_{i,j} \mathbf{R}_{i,j} (\log(\mathbf{R}_{i,j}) - 1)$$

and use it as regularization function in equation 2, which becomes :

$$\mathbf{R}^{OT} = \min_{\mathbf{R}} \sum_{i,j} \mathbf{R}_{i,j} \mathbf{C}(\mathbf{X}^s, \mathbf{X}^t)_{i,j} - \epsilon h(\mathbf{R}) \text{ s.t. } \left(\sum_j \mathbf{R}_{i,j} \right)_i = \frac{1}{p}, \left(\sum_i \mathbf{R}_{i,j} \right)_j = \frac{1}{q} \quad (3)$$

The entropic term makes the objective function ϵ -strongly convex, hence leading to a unique optimal solution for a given ϵ . Besides making computation of transport faster using Sinkhorn algorithm [17], this entropic regularization also acts as a smoothing of the solution.

Other functional alignment methods We may consider other transformations ϕ that map μ^s to μ^t as linear couplings \mathbf{R} that predict the stacked feature values \mathbf{X}^t from \mathbf{X}^s , under some given constraints ($\mathbf{R} \in \mathcal{R}$):

$$\mathbf{R}^* = \underset{\mathbf{R} \in \mathcal{R}}{\operatorname{argmin}} \|\mathbf{R}(\mathbf{X}^s) - \mathbf{X}^t\|_F^2 \quad (4)$$

Hyperalignment[8], search a scaled orthogonal transformation $\mathbf{R} = \sigma \mathbf{Q}$ s.t. $\mathbf{Q}^T \mathbf{Q} = \mathbf{Id}, \sigma \in \mathbb{R}_+$. With this model, equation 4 is equivalent to the well-studied scaled Procrustes problem, solvable in closed form using the singular value decomposition of $\mathbf{X}^s \mathbf{X}^{tT}$: $(\mathbf{U}, \mathbf{\Sigma}, \mathbf{V}) = \operatorname{SVD}(\mathbf{X}^s \mathbf{X}^{tT})$:

$$\mathbf{R}^{so} = \sigma \mathbf{Q} \text{ where } (\sigma, \mathbf{Q}) = \left(\frac{\operatorname{tr}(\mathbf{\Sigma})}{\|\mathbf{X}^s\|_F^2}, \mathbf{U}\mathbf{V} \right)$$

In the case where $p = q$, permutations and every coupling acceptable in strict optimal transport sense (eq. 2) are part of this broader class of orthogonal transforms. By contrast, entropic regularization (eq. 3) yields a non-orthogonal solution in general. It is possible to further relax the orthogonality constraint, by simply looking for a linear coupling \mathbf{R} with a small norm. For computational efficiency, we consider here ℓ_2 norm penalization, yielding $\min_{\mathbf{R}} \|\mathbf{X}^t - \mathbf{R}\mathbf{X}^s\|_F^2 + \frac{\lambda}{2} \|\mathbf{R}\|_2^2$. Using such a model, the alignment problem boils down to a ridge regression also solvable in closed form:

$$\mathbf{R}^{ridge} = \mathbf{X}^t \mathbf{X}^{sT} (\mathbf{X}^{sT} \mathbf{X}^s + \lambda \mathbf{I}_d)^{-1}$$

Full brain alignment Although we aim at building a full-brain functional atlas, applying the above models to full brain suffers from two related issues: *i*) it may create some non-local correspondences (e.g. cross-hemisphere swap) that are not neuroanatomically plausible; *ii*) it is computationally heavy if not intractable.

In [9] functional alignment is used with a *searchlight* approach. In this popular procedure in neuroimaging an algorithm is applied on a set of 3D balls of radius r covering the brain. Regions described by multiple balls are predicted by averaging the prediction in each ball. Although this procedure does produce local correspondences, it is computationally costly. One needs approximately $\frac{p}{r^3}$ balls to ensure that all voxels are covered. Furthermore, the aggregation of several balls destroys the local structure enforced by the algorithms within each ball (sparsity, smoothness, orthogonality...).

We propose instead parcellation-based alignment. In a first step we do a functional clustering of data to find local non-overlapping clusters c_1, \dots, c_C of voxels with common activity patterns. In each of these clusters we find the optimal alignment transform and concatenate these local transforms to recover a full-brain transform with the desired regularities. Formally the optimal alignment transform to align two subjects \mathbf{x}_{train}^s and subject \mathbf{x}_{train}^t on the training session is obtained by solving the problem in each cluster $c \in \{c_1, \dots, c_C\}$:

$$\mathbf{R}^*[c] = \mathbf{R}^*(\mathbf{X}_{train}^s[c], \mathbf{X}_{train}^t[c]), \text{ with } * \in \{OT, so, ridge\} \quad (5)$$

On the test session \mathbf{X}_{test}^t is predicted using \mathbf{X}_{test}^s by:

$$\forall c \in \{c_1, \dots, c_C\} \mathbf{X}_{test}^t[c] = \mathbf{R}^*[c] \mathbf{X}_{test}^s[c] \quad (6)$$

Several algorithms perform well for functional clustering, in this study we used a computationally efficient recursive K-means method, where, for a target number of k regions, a first clustering into \sqrt{k} pieces is obtained, and each of them is clustered in turn into \sqrt{k} parts. Solving problem 3 or 4 within a local region constrains the solution to remain local and acts as a regularizer.

Template inference Pairwise correspondences do not scale well with many individuals. A template measure \mathbf{T} is needed, which can be obtained by solving

$$\min_{\mathbf{T}, \mathbf{R}_1, \dots, \mathbf{R}_n} \sum_{i=1}^n \|\mathbf{R}_s(\mathbf{T}) - \mathbf{X}^s\|_*^2 \quad (7)$$

for the chosen loss $\|\cdot\|_*$ (Wasserstein or Frobenius). We solve it through alternate minimization iterating over :

- a \mathbf{R} -step of independent alignment of the current template to every sample, thus estimating $\mathbf{R}_i, i = 1..n$
- a T -step where \mathbf{T} is regressed to minimize jointly its distance to the samples.

$$\min_{\mathbf{R}_s} \|\mathbf{R}_s(\mathbf{T}) - \mathbf{X}^s\|_*^2, \forall s \in [n] \quad \mathbf{R}\text{-step}$$

$$\min_{\mathbf{T}} \sum_{s=1}^n \|\mathbf{R}_s(\mathbf{T}) - \mathbf{X}^s\|_*^2 \quad T\text{-step}$$

Note that, for all norms considered here, the T -step results in a quadratic problem solved by conjugate gradient. As initialization, we define first alignment operators list as identity, thus the first T -step is $\min_{\mathbf{T}} \sum_{s=1}^n \|\mathbf{T} - \mathbf{X}^s\|_2^2$ and the first template is the sample mean. In practice, we run 4 iterations of the alternate minimization, as we found that was sufficient for convergence.

Complexity analysis. We recapitulate the algorithmic complexity of the OT algorithm for $p = q$. While the Hungarian algorithm finds an optimal (deterministic) permutation in a complexity of $O(p^3)$, Sinkhorn find a τ -approximate solution of the unregularized problem in $O(p^2 \log(p) \tau^{-3})$ operations [13]. If the brain contains P voxels divided evenly into k clusters, the overall complexity of the problem is thus $O(\frac{P^2}{k} \log(\frac{P}{k}) \tau^{-3})$. Increasing k thus reduces the computational burden; on the other hand, choosing k too large yields small clusters, hence reduces the benefit of OT. In the following experiments, we set $k = 200$.

3 Experiments

Datasets To benchmark these methods and assess their prediction accuracy, we ran experiments on two datasets, where individual data were previously registered in MNI-space following standard procedure (SPM12 software called though Nipype for IBC, HCP pipeline for HCP). The "Individual Brain Charting" (IBC) [5] contains scans of the same 13 participants for a wide variety of cognitive tasks. The data were acquired using a 3T scanner (acquisition resolution of $1.5mm$ resampled at $3mm$ after spatial normalization). We worked directly on activation maps: for a given functional contrast, they associate an activation statistic with each voxel.

We learn alignment between subjects between $d = 53$ contrasts derived from data acquired with antero-posterior(AP) EPI phase encoding; to assess the quality of our predictions, we also use 53 contrasts acquired in separate sessions using the same experimental paradigms with posterior-anterior(PA) EPI phase encoding. Note that the resulting AP/PA distortions were estimated and corrected with FSL's topup software prior to image pre-processing. Human Connectome Project (HCP) [18]. is the collection of neuroimaging and behavioral data on 1,200 normal young adults, aged 22-35. For our experiment we focused on 20 randomly chosen subjects. For each, we used the $d = 25$ statistical maps available in both left-to-right(LR) and right-to-left(RL) phase encoding, resampled at $3mm$ after spatial normalization. We learn alignment between subjects using LR images and assess prediction on RL acquisitions.

Pairwise alignment benchmarks In pairwise prediction, we first learn the optimal alignment operator between a source and a target subject on training data. We then use this alignment and supplementary images of the source subject to




Fig. 1: Pairwise prediction of target subject from aligned source subject

predict additional data of the target subject and score this prediction against the true target subject images using a prediction metric (see Fig. 1).

We use this straightforward set-up on IBC and HCP datasets to compare prediction performance of the alignment methods presented above, applied on a parcellation of the brain: (i) Scaled-Orthogonal Transform, (ii) Ridge Regression, (iii) Optimal Permutation, (iv) OT with entropic smoothing. We compare these to two baselines, the identity transform (that predicts the target subject as the source subject data), and a multi-purpose state-of-the-art diffeomorphic medical image registration algorithm: symmetric image normalization (SyN)[19]. SyN yields a diffeomorphic mapping maximizing Mattes mutual information between local regions. Since it works only for scalar images, it was applied only on the principal components of the training set of images. In a second experiment, we study the influence of the amount entropic regularization used in OT loss on both datasets. Obviously some pairs of subjects have more similar functional data than others. To make our evaluation process robust to this variability, we tested every method on the same set of 20 pairs of subjects chosen randomly in each dataset. Our implementation relies on Nilearn for data handling, open-source solvers (Scikit-learn, POT for Sinkhorn, antspsy for SyN).




Fig. 2: Template-based prediction of left-out subject

Template-based alignment benchmarks To evaluate template-based prediction accuracy, we split the IBC dataset randomly into two folds of 7 and 6 subjects.

We inferred the template from the train subjects across every AP and PA contrast. We then learned an alignment operator between each of the test subject and the template using AP contrasts and try to predict their PA contrasts that we scored using predictions metrics (see Figure 2). We compare the results of the same methods quantitatively in terms of prediction loss. In a third experiment, we infer a template on which we learn alignment for all subjects using AP data. We then apply these alignments to left out PA data to bring all our subject in a common space. In this common space, we a run a one sample-test and compare the group effects detected by each method, for specific conditions.

Prediction metrics To measure the quality of our prediction $\mathbf{R}_i(\mathbf{X})$, at the voxel level, against the ground truth \mathbf{Y} , we defined η^2 , the normalized reconstruction error, as :

$$\eta_{\star}^2(\mathbf{Y}, \mathbf{Y}_i, \mathbf{X}) = 1 - \frac{\sum_{i=1}^n (\mathbf{Y}_i - \mathbf{R}_i \mathbf{X})^2}{\sum_{i=1}^n \mathbf{Y}_i^2},$$

where \star stands for identity, ridge, scaled orthogonal, OT. An η^2 value of 0 means that the quality of the prediction is equivalent to predicting 0 along all the dimensions. A perfect prediction yields a value of 1. To focus on the prediction improvement that can be made through alignment - independently of the pre-existing distance between data to align - we assess performance quantitatively using a reconstruction ratio R_{η^2} . This ratio is also defined at voxel level and is superior to 0 if the voxel is predicted better by aligned data than by raw data.

$$R_{\eta_{\star}^2}(\mathbf{Y}, R, \mathbf{X}) = 1 - \frac{\sum_{i=1}^n (\mathbf{Y}_i - \mathbf{R}_i \mathbf{X})^2}{\sum_{i=1}^n (\mathbf{Y}_i - \mathbf{X}_i)^2} = 1 - \frac{1 - \eta_{\star}^2(\mathbf{Y}, \mathbf{R}_i, \mathbf{X})}{1 - \eta_{id}^2(\mathbf{Y}, Id, \mathbf{X})}$$

4 Results

Pairwise alignment Figure 3 shows that functional alignment methods generally improve prediction quality from one subject to another with respect to the identity, though not uniformly over the cortex. Sensory and motor regions typically obtain high scores, showing the stability of the signals across subjects in these areas; by contrast, other regions obtain low scores overall. Syn offers no improvement of prediction scores, nor does the optimal permutation of voxels. This means that a strict one-to-one mapping of voxels is not suitable for functional alignment. For the three other methods, we clearly see different behavior between regions with high signal-to-noise ratio (SNR) and regions with lower SNR. Figure 4 (a-b), report the compared distributions of predictions ratios R_{η^2} on IBC and HCP datasets and are consistent with previous observations. Ridge and OT outperform all other methods on IBC dataset whereas Scaled Orthogonal and Ridge perform slightly better on HCP dataset. Figure 4 (c), shows that entropic regularization strongly improves prediction scores up to an inflexion point.


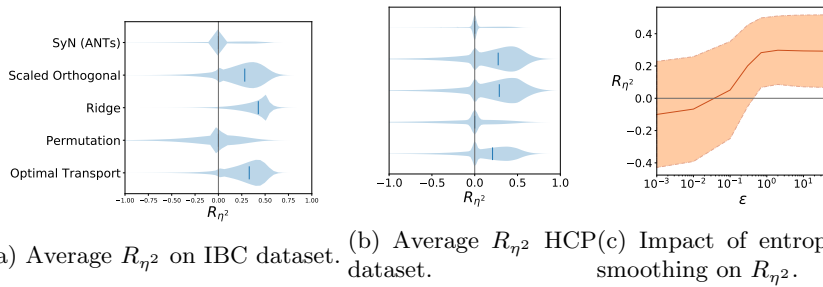


Fig. 3: Prediction score $\eta^2(\mathbf{Y}, \mathbf{R}, \mathbf{X})$ of target subject 9 using alignment with subject 15 (IBC dataset, $z=-17\text{mm}$).



(a) Average R_{η^2} on IBC dataset. (b) Average R_{η^2} HCP dataset. (c) Impact of entropic smoothing on R_{η^2} .

Fig. 4: Pairwise prediction : (a), (b) Alignment score of methods on 20 pairs of IBC/HCP datasets. (c) Effect of entropic regularization (OT) IBC dataset.

Template alignment In Figure 5, we can observe that prediction accuracy is strongly improved by the use of a learned template in most brain regions, which establishes that functional alignment performs well to estimate cross-subject correspondences and identify a latent brain activity template. It also validates that our template estimation procedure manages to capture some of the inter-subject variability in its mapping process.

Fig. 5 shows that Scaled orthogonal, Ridge and Optimal transport are equally accurate at predicting new subject data overall. However they do so in dissimilar ways. Ridge regression tend to predict 0 in regions with low SNR. This strong smoothing effect comes at the expense of providing predictions that are not very precise for high SNR regions. On the contrary, Optimal Transport makes large mistakes in low SNR regions but predicts high-SNR regions more accurately. This behavior is desirable in a template-building procedure since it better preserves specificity and structure in functional signal.

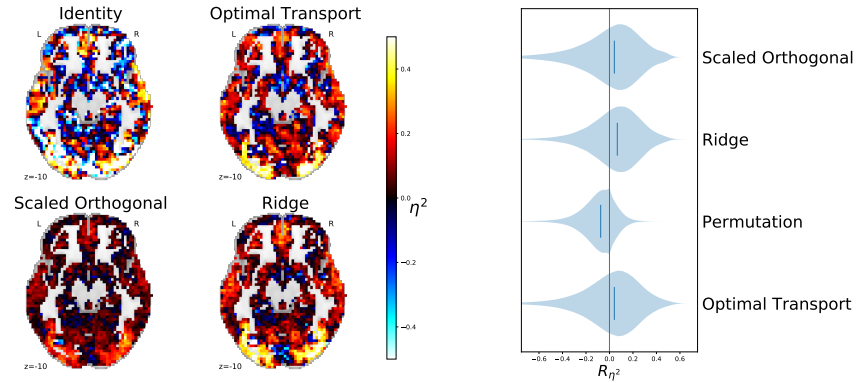


Fig. 5: Template-based prediction: map yields the score of IBC-subject 11 missing data prediction. Plot shows prediction ratio distributions across test subjects, i.e. gain from alignment over prediction from group average.

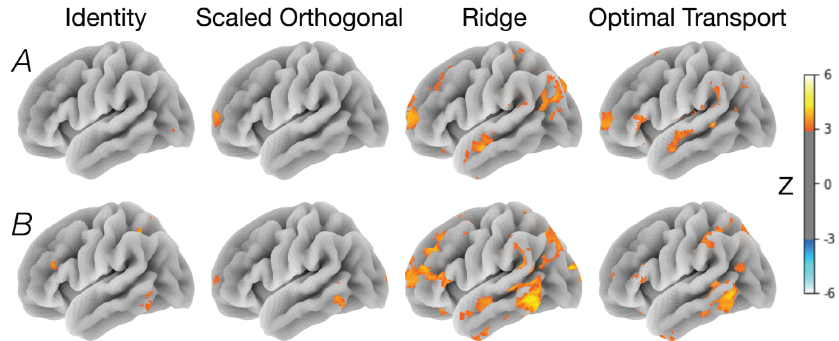


Fig. 6: 1 sample test Z scores after aligning IBC subjects to template. 2 conditions : (A) *Match* (*hcp-relational task*)(B) *0 back place*(*hcp-wm task*)

We finally consider a realistic use case, where test data would be functionally aligned in view of a group study (one-sample t-test). We present the ensuing group-level brain maps, projected on the cortical surface, in Fig.6, for 2 contrasts. The *Match* contrast (borrowed from a visual comparison task on artificial textures), shows that both ridge and optimal transport recover regions in the anterior and posterior segments of the superior temporal sulcus that would not be detected using standard approaches, or scale-orthogonal functional alignment. For the *0-back place* contrast, borrowed from a visual matching task on place images, OT- and Ridge-based alignment recover regions in the temporo occipital junction and inferior parietal sulcus that would not be detected using standard approaches, nor by scale-orthogonal functional alignment. In particular, place sensitive regions of the ventral and dorsal visual cortex seem to have been successfully recovered by the functional alignment approach.

5 Discussion

We demonstrated in this work that functional alignment can be used to infer a template that captures a good-enough level of detail to generalize to data from new subjects. This confirms prior results from [8, 9], but we moreover introduced a principled way to build brain templates in this framework. Among the possible approaches to identify correspondences and infer a template, weakly-regularized ones proved to be particularly efficient: local ridge regression is clearly the less structured method, yet it yields high accuracy, typically higher than the more constrained local Procrustes alignment; smooth OT performed better than non-smooth OT. Eventually, the most constrained approach, Syn, did not perform well on this task. Importantly, these results do not reflect mere overfit, as the predictions are made on images not used for alignment and template inference. Additional extensions can be brought to the current framework:

- First, remove the reliance on a fixed parcellation of the brain, that is probably suboptimal to identify cross-subjects correspondences. This can be done by ensembling results obtained from multiple parcellations [20]. This should improve accuracy, albeit at a higher computation cost.
- Second, learn correspondences between subjects and toward a common template from resting-state data, following the ideas in [11, 7]. This will be especially useful since resting-state data are more and more frequently used in group studies and represent the standard for brain organization studies.

Acknowledgments This project has received funding from the European Unions Horizon 2020 Research and Innovation Programme under Grant Agreement No. 785907 (HBP SGA2).

References

1. Arno Klein, Jesper Andersson, Babak A Ardekani, John Ashburner, Brian Avants, MingChang Chiang, Gary E Christensen, D Louis Collins, James Gee, Pierre Hellier, et al. Evaluation of 14 non linear deformation algorithms applied to human brain mri registration. *Neuroimage*, 46(3):786–802, 2009.
2. Thirion B. Functional neuroimaging group studies. In Thompson W Ombao H, Lindquist M and Aston J, editors, *Handbook of Neuroimaging Data Analysis*, chapter 12, pages 335–354. Chapman and Hall and CRC, 2016.
3. Evelina Fedorenko, Michael K Behr, and Nancy Kanwisher. Functional specificity for highlevel linguistic processing in the human brain. *Proceedings of the National Academy of Sciences*, 108(39):16428–16433, 2011.
4. D. M. Barch, G. C. Burgess, M. P. Harms, S. E. Petersen, B. L. Schlaggar, M. Corbetta, M. F. Glasser, S. Curtiss, S. Dixit, and C. Feldt. Function in the human connectome: task-fMRI and individual differences in behavior. *Neuroimage*, 80:169–189, 2013.
5. Ana Luísa Pinho, Alexis Amadon, Torsten Ruest, Murielle Fabre, Elvis Dohmatob, Isabelle Denghien, Chantal Ginisty, Séverine Becuwe-Desmidt, Séverine Roger, Laurence Laurier, Véronique Joly-Testault, Gaëlle Médiouni-Cloarec, Christine Doublé, Bernadette Martins, Philippe Pinel, Evelyn Eger, gael varoquaux,

- Christophe Pallier, Stanislas Dehaene, Lucie Hertz-Pannier, and Bertrand Thirion. Individual Brain Charting, a high-resolution fMRI dataset for cognitive mapping. *Scientific Data*, 5:180105, June 2018.
6. M.R. Sabuncu, B.D. Singer, B. Conroy, R.E. Bryan, P.J. Ramadge, and J.V. Haxby. Function-based intersubject alignment of human cortical anatomy. *Cereb Cortex*, 20:130, 2010.
 7. KarlHeinz Nenning, Hesheng Liu, Satrajit S Ghosh, Mert R Sabuncu, Ernst Schwartz, and Georg Langs. Diffeomorphic functional brain surface alignment: Functional demons. *NeuroImage*, 2017.
 8. James V Haxby, J Swaroop Guntupalli, Andrew C Connolly, Yaroslav O Halchenko, Bryan R Conroy, M Ida Gobbini, Michael Hanke, and Peter J Ramadge. A common, highdimensional model of the representational space in human ventral temporal cortex. *Neuron*, 72(2):404–416, 2011.
 9. J Swaroop Guntupalli, Michael Hanke, Yaroslav O Halchenko, Andrew C Connolly, Peter J Ramadge, and James V Haxby. A model of representational spaces in human cortex. *Cerebral cortex*, 26(6):2919–2934, 2016.
 10. Georg Langs, Danial Lashkari, Andrew Sweet, Yanmei Tie, Laura Rigolo, Alexandra J. Golby, and Polina Golland. Learning an atlas of a cognitive process in its functional geometry. In *Information Processing in Medical Imaging - 22nd International Conference, IPMI 2011, Kloster Irsee, Germany, July 3-8, 2011. Proceedings*, pages 135–146, 2011.
 11. Georg Langs, Danhong Wang, Polina Golland, Sophia Mueller, Ruiqi Pan, Mert R. Sabuncu, Wei Sun, Kuncheng Li, and Hesheng Liu. Identifying shared brain networks in individuals by decoupling functional and anatomical variability. *Cerebral Cortex*, 26(10):4004–4014, 2016.
 12. Umut Güçlü and Marcel A. J. van Gerven. Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream. *Journal of Neuroscience*, 35(27):10005–10014, 2015.
 13. Gabriel Peyré and Marco Cuturi. Computational Optimal Transport. *arXiv e-prints*, page arXiv:1803.00567, March 2018.
 14. J. Munkres. Algorithms for the assignment and transportation problems. *Journal of the Society of Industrial and Applied Mathematics*, 5(1):32–38, March 1957.
 15. Leonid Kantorovitch. On the translocation of masses. *Management Science*, 5(1):1–4, 1958.
 16. Alexandre Gramfort, Gabriel Peyré, and Marco Cuturi. Fast optimal transport averaging of neuroimaging data. *CoRR*, abs/1503.08596, 2015.
 17. Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in neural information processing systems*, pages 2292–2300, 2013.
 18. David C. Van Essen, Stephen M. Smith, Deanna M. Barch, Timothy E.J. Behrens, Essa Yacoub, Kamil Ugurbil, and for the WU-Minn HCP Consortium. The wu-minn human connectome project: An overview. 80:62–79, 2013.
 19. Brian B. Avants, Charles L. Epstein, Murray Grossman, and James C. Gee. Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis*, 12(1):26–41, 2008.
 20. Andrés Hoyos-Idrobo, Gaël Varoquaux, and Bertrand Thirion. Towards a Faster Randomized Parcellation Based Inference. In *PRNI 2017 - 7th International Workshop on Pattern Recognition in NeuroImaging*, Toronto, Canada, June 2017.