Assessing the relevance of the Individualized Neural Tuning model on IBC Data

Denis Fouchard
`denis.fouchard@telecom-paris.fr` https://team.inria.fr/mind/files/2022/04/mind$_{l}ogo.png$

INRIA de Saclay
MIND TEAM

october 2023 - February 2024

Supervised by Bertrand Thirion

Table des matières

1 Introduction

1.1 Purpose

Functional Brain Alignment is key to understanding individual idiosyncracies in brain functional disposition and is used in a variety of experimental setups such as brain decoding.

There are many ways to do brain alignment, that fall into two broad categories, pairwise alignment and template alignment. There are many ways to perform template alignment, such as the Shared Response Model [2], developed by Hugo Richard et. al. at INRIA Saclay. We are studying the Individual Neural Tuning model, developed by Feilong et. al. [1] at HaxbyLab. It aims at doing template alignment through hyperalignment, a method that uses searchlight decomposition of brains to do piecewise local alignments and combine the results.

Interesting claims are made in the paper regarding correlations, with impressive results such as Intra-Subjects correlations reaching 0.4 in some setups. Those numbers caught our attention and are the main reason of this analysis.

We aim to replicate some results from Feilong et. al (2023) [1], and further test the model on more data such as contrasts from the IBC Dataset. We will also compare the performances to other similar alignment techniques such as FastSRM [2].

Our objectives are the following :
- Verify the validity of the model and the claims made.
- Benchmark the model on higher resolution data, and with more timeframes
- Use other benchmarks

1.2 Main challenges

Trying to replicate and understand the model was not an easy task. We had to face many challenges among which :

- I could not access a complete implementation of the method, so we had to reimplement a great deal of the code by ourselves, trying to understand what the model should do.
- First versions of the preprint paper were unclear about the actual structure of the model. Although some parts were detailed (like the ensemble bagging method), there were many shadow zones, particularly regarding the way the matrix decomposition was done after denoising.
- Implementing the original proposal - using ensemble methods (bagging) on Ridge regression for each searchlight - was not realistic from a computational complexity point of view (even with multithreading), and the theoretical benefits were quite limited.
- There are lots of ambiguities, such as how the stimulus is obtained from the raw signal. We noticed that performance seems to increase when we scale temporal sources with singular values after performing the SVD on the concatenated raw data.
- HaxbyLab did not provide any code for Searchlight computation. Their Forrest Dataset consists of precomputed searchlights and distances, as well as pre-masked fMRI images (that is, raw matrices with no maser provided). We have no information about how the fMRI raw signal images were preprocessed.
- Plotting on their part relies on custom code, not standard and high-quality available libraries such as `nilearn`.
- No benchmarks were run on standard fMRI formats such as Nifti Images, or if so, no masker was provided.
- Benchmarking methods were not detailed properly, such as Pearson's correlation computing. We also suspect that some of the benchmarks being conducted on specific subject data were cherry-picked for the report.
- The code that we could get from that model was poorly commented and the variable names did not ease any understanding. We had to do a great deal of reverse-engineering which was frustrating.
- The way plots were done is not detailed. We suspect correlation was using some kind of optimization such as linear sum assignment, but this was never clearly specified.
- Related to the last point, there is no identification of sources in the model, so reconstructing the images is not possible. The benchmarked reconstruction.
- Lastly, while we had a good discussion with the team from which this method comes at the beginning, our conversation ended at some point and we had to continue alone.

1.3 Highlights

- Implementing our version of the INT and using it on IBC datasets revealed co-smoothing and prediction performance significantly better than our baseline, and similar results to FastSRM (with some caveats).
- While not studied in the original paper, the use of some dimension reduction in the latent source space yields better results, especially on large (many timeframes) datasets
- No mention is made in the original work regarding source identifiability. Identifiability can be achieved *a posteriori* via source synchronization.
- Searchlight parcelation is giving better marginally better results , but at the expense of much greater computational needs (see Annexe for details).

An implementation of the model is integrated in the `fmralign` Python package developped by T. Bazeille, E. DuPré, B. Thirion available at https://github.com/Parietal-INRIA/fmralign.
Benchmarks and plots are available **here**.

2 Datasets

2.1 Haxby Datasets

Name	Haxby-Raiders
#timepoints	1680
#subjects	23
#voxels	9675

Name	Haxby-Forrest
#timepoints	1818
#subjects	15
#voxels	9675

2.2 IBC Datasets

Name	IBC-Raiders
#timepoints	2119
#subjects	15
#voxels	46678

Name	IBC-Contrast
#timepoints	106
#subjects	7
#voxels	46678

!

Import point

For some experiments, we will use masking to work only on visual and movie-watching brain regions (such as the occipital cortex) for the IBC-Raiders dataset. This gives us a dataset with a voxel resolution of 6631.

3 Model analysis

For the following, we will denote :

- n (or n_s) the number of subjects for a given experimental setup
- t (or n_t) the number of acquisition time points for the experiment (we assume that all acquisitions are performed simultaneously for each subject)
- v (or n_v)the voxel resolution of each acquisition
- $B^{(i)}$ the raw signal acquisition of subject (i) for the whole experiment.
- k the number of brain regions (the same for each subject in one experiment)

That is, for one experiment we end up with raw signals $B^{(1)}, ..., B^{(n)}$, with $B^{(i)}$ being a $t \times v$ real signal matrix.

We aim here at benchmarking the different parts and modules of the broad "Individualized Neural Tuning" model, later referred to as "INT".
The model consists of :

- A denoising part : from raw fMRI data matrices $B^{(1)}, ..., B^{(n)}$ we compute their $B^{\hat{(1)}}, ..., B^{\hat{(n)}}$ denoised counterparts through either Searchlight-PCA / Searchlight-Ridge or Parcel-PCA / Parcel-Ridge

— A stimulus/tuning decomposition part. From $B^{(1)}, ..., B^{(n)}$ we extract a common temporal template called "stimulus" S as $t \times k$ matrix. Then from $B^{\hat{(1)}}, ..., B^{\hat{(n)}}$ and S we get tuning matrices $T^{(1)}, ..., T^{(n)}$.

One goal will be to assess the relevance of using a Searchlight-based approach over a Parcel-based approach. Indeed, the first is much more computationally heavy and memory intensive than the later, as demonstrated in Appendix B.

For this whole section, we will use the IBC-Raiders dataset.

3.1 Denoising

Recall : Denoising is performed in order to resolve the inverse problem. For each participant p, raw signal is defined as $B^{(p)} = \hat{B}^{(p)} + E = MW^{(p)} + E$ where E is random noise.
The specificity of this approach resides in the choice of doing a region-wise denoising instead of considering the whole brain.

The main idea is to use either a Searchlight or a Parcelation algorithm (such as Ward, K-Means) to decompose the brain into regions to work on separately. With a Searchlight, we get v greatly overlapping regions, while with a Parcelation we can get as many non-overlapping regions as we want. Empirically, for a resolution of $v = 46678$, values $k = 100$ and $k = 200$ work the best, as we will observe a decline in performance for higher values, as we will later see.

First, a local one-parcel denoising benchmark (Fig. 2). For a given parcel (437 voxels), we performed PCA (150 components), then Ridge with cross-validation). Correlation is computed between $B^{(1)}[pc]$ and $\hat{B}^{(1)}[pc] = MW_{(1)}[pc]$.

$$a[t_1, t_2] = \rho(\hat{B}^{(1)}[t_1], B^{(1)}[t_2])$$

$$b[s_1, s_2] = \frac{1}{n_t} \sum_{t=1}^{n_t} \rho(\hat{B}^{(s_1)}, B^{(s_2)})[i, i]$$

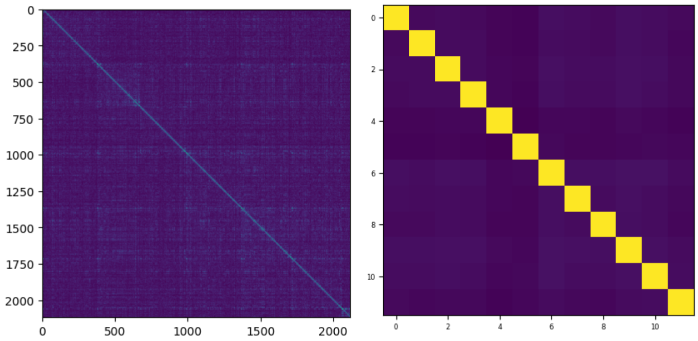

(a) Temporal Pearson Correlation between raw signal and hyperalignment denoised signal (2119 time frames, 437 voxels), for sub-01
. The diagonal average is 0.67.

(b) Mean voxelwise inter-subject Pearson correlation between raw signal and hyperalignment denoising (2119 time frames, 437 voxels). Diagonal average is 0.77

FIGURE 2 – Spacial and temporal correlations between raw signal and denoised signal (parcel hyperalignment).

3.2 Tuning-Stimulus decomposition

!

Import point

In this section, we will only work on visual regions.

3.2.1 Impact of scaling sources

We here try to assess the relevance of the decomposition on raw signal fMRI data. We therefore perform the decomposition as follows :

$$U, \Sigma, V = SVD([B_{(1)}, ..., B_{(n)}])$$

$$S = \frac{1}{\sqrt{n}} U$$

$$T_{(i)} = S^{-1} B_{(i)} \text{ for } i = 1, ..., n \tag{1}$$

We also try, as specified in the introduction, to scale temporal sources of S with their singular values. That changes the equation (**??**) into :

$$S_\Sigma = \frac{1}{\sqrt{n}} U\Sigma$$

$$T_{(i)} = S_\Sigma^{-1} B_{(i)} \text{ for } i = 1, ..., n \tag{2}$$

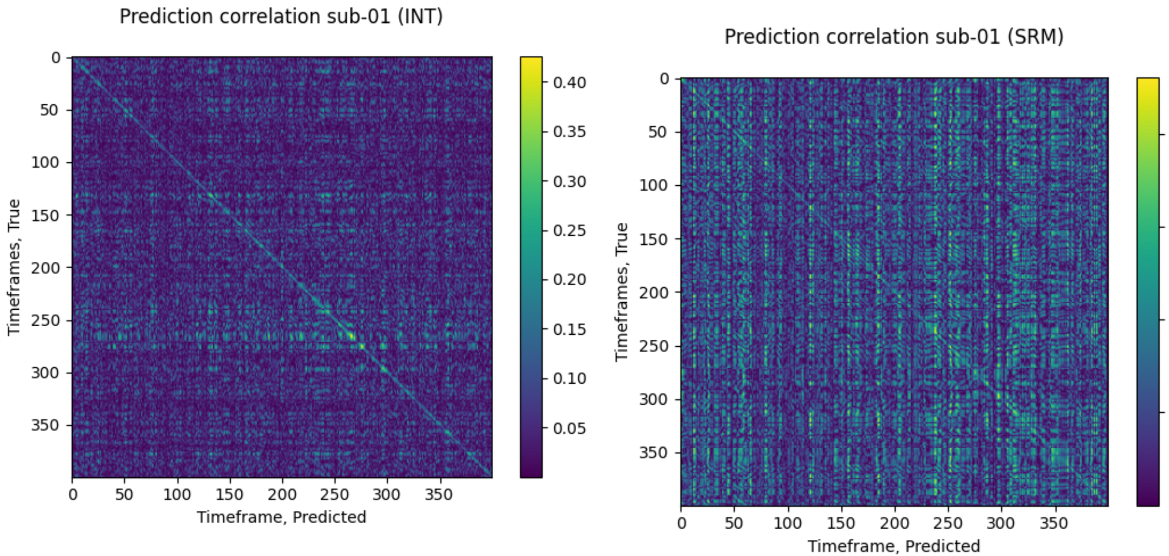

(a) Prediction for sub-01, INT(with *a posteriori* source identification).

(b) Prediction for sub-01, Indentifiable SRM

Figure 3 – IBC-Raiders, (400 timepoints) prediction vs. ground truth temporal correlation for sub-01. Training is performed on the first 400 time points, and testing on the next 400 time points. Reading : the diagonal represents the same timeframe correlation for every voxel of sub-01. For both methods, the M correlation matrix is computed as follows : $M[t_1, t_2] = \rho(B^{(1)}[t_1], B^{(1)}_{pred}[t_2])$

Note : standardization/normalization did not change the results.

Model \Metric	Same TR correlation (avg/std)	Inter TR correlation (avg/std)
Identifiable SRM 100TR	0.275/0.145	0.067/0.052
INT 100TR	0.185/0.055	0.091/0.074

3.2.2 Impact of reducing source components

Model \Metric	Same TR correlation (avg/std)	Inter TR correlation (avg/std)
SRM 400TR (50 components)	**0.232**/0.066	0.077/0.059
SRM 400TR (5 components)	**0.242**/0.085	0.11/0.080
INT 400TR (full sources)	0.178/0.053	0.055/0.044
INT 400TR (5 components)	0.231/0.086	0.099/0.075
INT 400TR (20 components)	**0.238**/0.068	0.080/0.068
INT 400TR (50 components)	0.209/0.062	0.068/0.053

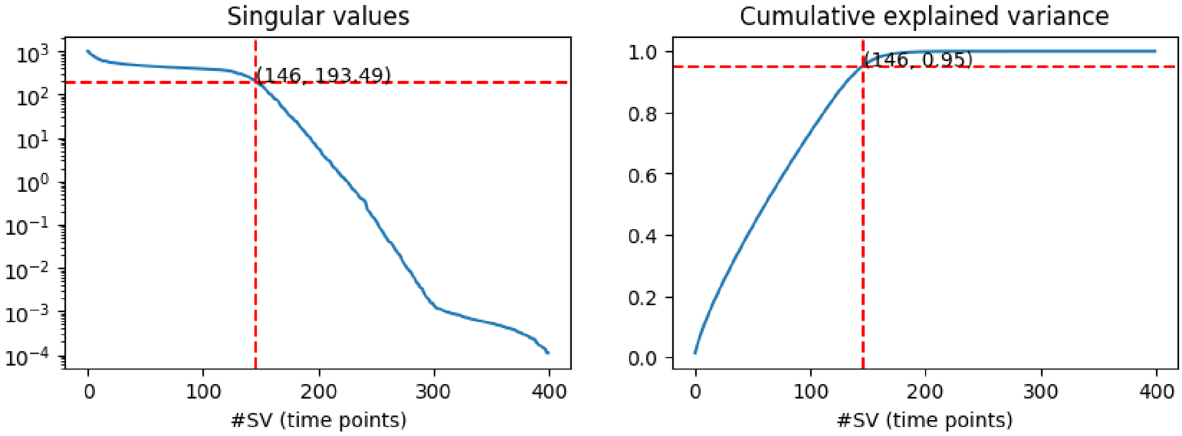

FIGURE 4 – IBC-Raiders, all subjects, 400 timepoints - Plots of Singular values and explained variance in stimulus sources (the U matrix from concatenated raw signal SVD) (0.95-threshold)

3.2.3 Source identification

On all correlation benchmarks, we had to compensate for the lack of source identification in the model, by performing *a posteriori* identification through linear sum assignment. This gives a representation of *how the model would work if there was identification.* For instance, the Identifiable Shared Model Response does not need such an assignment. See 4.2 for more details.

3.2.4 Diving on each model shared stimulus

We ran the INT model and FastSRM on one parcel (357 voxels) and extracted for both models the stimulus data (called *shared response* in SRM, and *stimulus* in the INT).

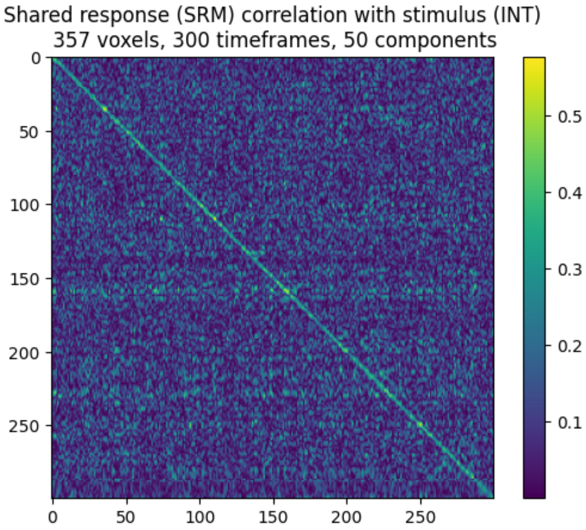

FIGURE 5 – Temporal Pearson correlation (300-time frames) between the first 50 components of stimulus signal from INT (S_{INT}) and shared response from FastSRM (S_{SRM}). Both were computed from the combination of 12 subjects' data. The hereby M matrix is computed as follows :

$$M[t_1, t_2] = \rho(S_{SRM}[t_1], S_{INT}[t_2]), (t_1, t_2) \in [1, n_t]^2$$

We obtain a mean inter-frame correlation of 0.11 ($+$- 0.083) and a mean same-frame correlation of 0.36 ($+$-0.065).

For the whole watching session (2119 frames), we got to 0.44 $+=$ 0.06 intra-TR correlation and 0.11 $+=$ 0.084 inter-TR correlation.

4 Results

4.1 Prediction Correlation

Task - We perform a correlation between ground truth fMRI images and their reconstruction using the INT model.

We train our model on the first half of the timestamps (400 time-stamps) with all subjects, thus building tuning matrices for each subject. We then extract the stimulus signal from the second half (with all the subjects again), and use the training tuning data to rebuild the brain response signals for each subject.

Metric - Pearson's correlation coefficients between brain signals (ground truth v. prediction), for each subject and for each timestamp. We plotted the distributions of correlations we obtained.
We get 4 different distributions :

- Correlation within the same subject for the same time frame ($C(S, T)$)
- Correlation within the same subject but for different time frames ($C(S, \bar{T})$)
- Correlation between two different subjects for the same time frame ($C(\bar{S}, T)$)
- Correlation between two different subjects for different time frames ($C(\bar{S}, \bar{T})$)

$$C(S, T) = \{\rho(B^{(s)}_{pred}[t], B^{(s)}[t]) \ s.t. \ t \in [1, n_t]; s \in [1, n_s]\}$$

$$C(S, \bar{T}) = \{\rho(B^{(s)}_{pred}[t], B^{(s)}[\bar{t}]) \ s.t. \ t \neq \bar{t}\}$$

$$C(S, T) = \{\rho(B^{(s)}_{pred}[t], B^{(\bar{s})}[t]) \ s.t. \ s \neq \bar{s}\}$$

$$C(S, T) = \{\rho(B^{(s)}_{pred}[t], B^{(s)}[t]) \ s.t. \ s \neq \bar{s}, t \neq \bar{t}\}$$

Note : On plots, a timeframe is denoted 'TR' and subject are denoted 'subs'

4.1.1 Haxby-Forrest

Note : Searchlights were precomputed and bundled into the `neuroboros` package. We don't have its implementation so we had to reimplement our own using code from the Python `nilearn` library. Moreover, a choice was made in the original paper to plot correlation distribution for only one subject, which could be cherry-picked. We chose instead to plot distribution for all subjects. As such, the results were a bit lower than his results.

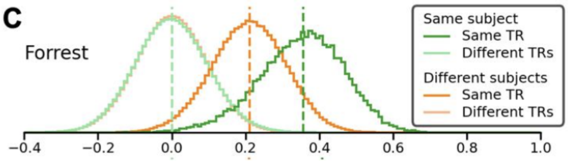

FIGURE 6 – Prediction vs. ground-truth voxels correlation distributions (same/different subjects, same/different time points) (all subjects, Haxby-Forrest, Feilong Ma)

4.1.2 IBC-Raiders

In this section, we will only work on visual regions.

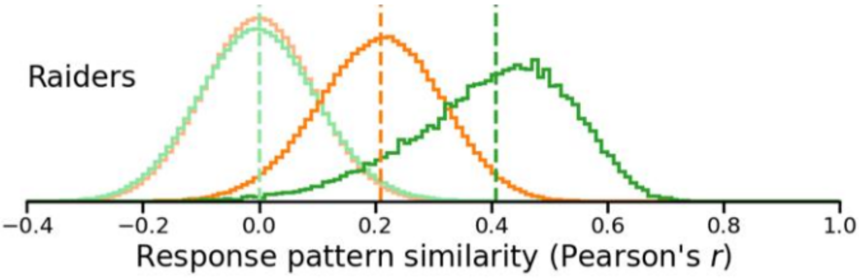

FIGURE 7 – Prediction vs. ground-truth voxels correlation distributions (same/different subjects, same/different time points) (all subjects, Haxby-Raiders, Feilong Ma)

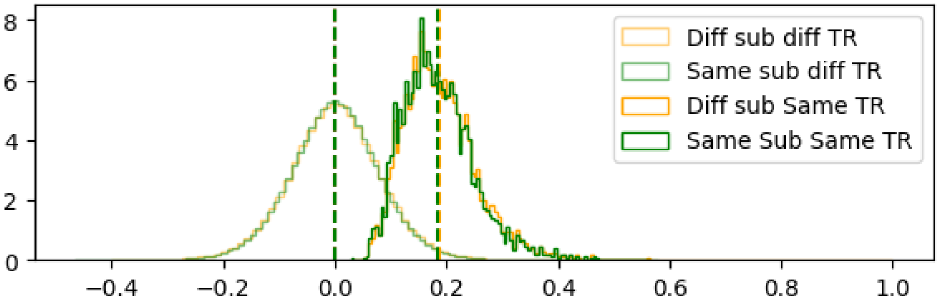

FIGURE 8 – INT - Prediction vs. ground-truth voxels correlation distributions (same/different subjects, same/different time points) (all subjects, IBC-Raiders visual areas, 400 time points, full sources)

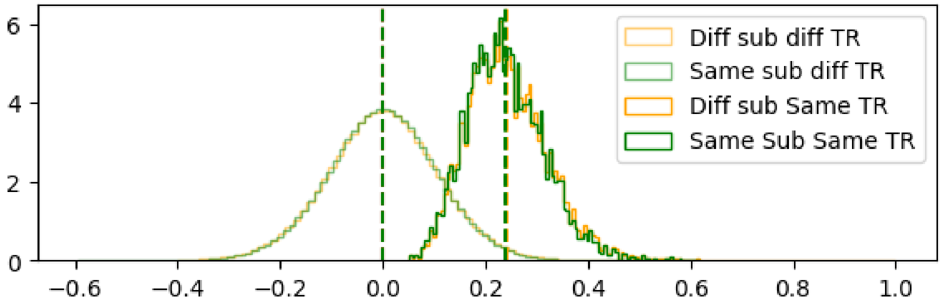

Figure 9 – INT - Prediction vs. ground-truth voxels correlation distributions (same/different subjects, same/different time points) (all subjects, IBC-Raiders visual areas, 400-time points, 20 source components)

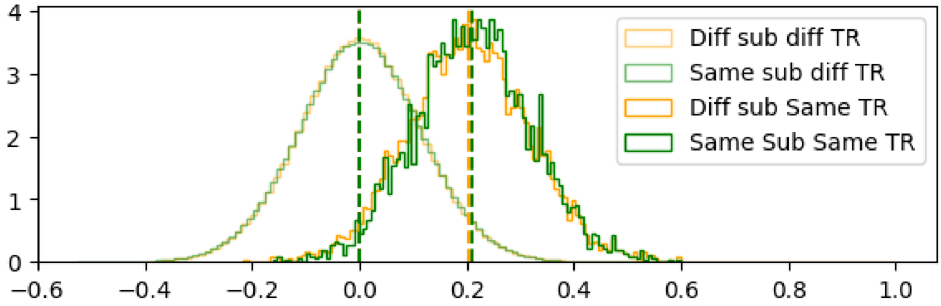

Figure 10 – IdentifiableSRM - reduction vs. ground-truth voxels correlation distributions (same/different subjects, same/different time points) (all subjects, IBC-Raiders visual areas, 400-time points, 50 source components)

4.2 Identifiability issues

One of the big drawbacks of the *IndividualNeuralTuning* method is the identifiability of sources. When the S stimulus matrix is computed (1) through SVD, there is a random permutation of the temporal sources compared to the raw signal, that is, for t a time-point in the original time series, the stimulus row associated with $B[t]$ is $S[\sigma(t)]$, with σ a random permutation.

Therefore, while the tuning data extracted from this method might indeed retain individual functional idiosyncrasies as claimed in the original paper, from one decomposition to another, we end up with mixed sources. This fact makes it virtually impossible to get two different tuning matrices from the same subject from different runs with sources in the same order.

What we have to do is some sort of *a posteriori* source identification by solving a linear assignment problem on the temporal correlation matrix between two tuning matrices from the same individual.

This allowed us to estimate as best as we could the σ permutation between two tuning matrices, which is essential while doing co-smoothing or predictions.

4.2.1 Prediction correlation on IBC-Raiders with *a posteriori* source identification

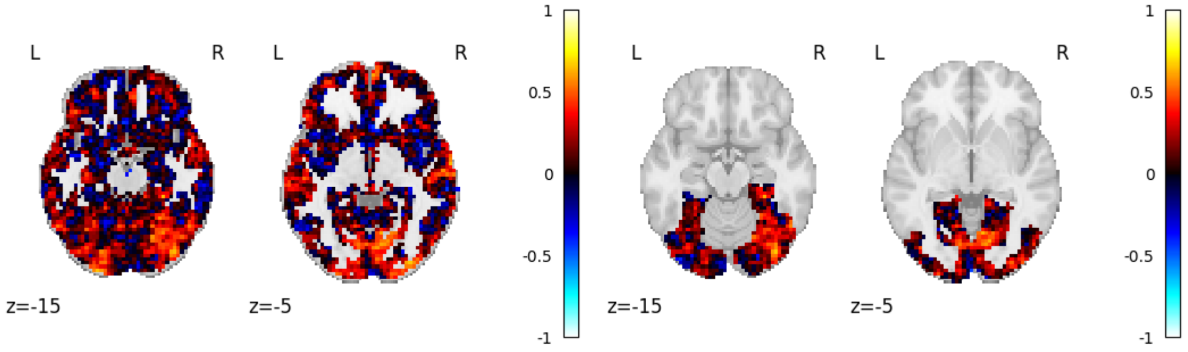

(a) Baseline (average) (whole brain) (b) Baseline (average) (visual areas)

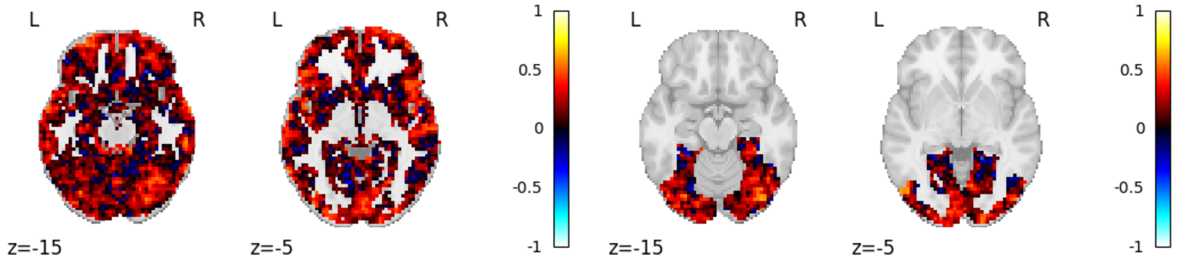

(c) INT - 200 parcels (whole brain) (d) INT - 200 parcels (visual areas)

FIGURE 11 – IBC-Raiders - Benchmark of INT with parcellation (200 parcels) both on whole brain and visual area - co-smoothing with stimulus from the first 200-time framers data on all subjects except `sub-07`, tuning on `sub-07` first 100 time-frames. Prediction/ground-truth voxel correlation map for `sub-07`. The resulting map C is given for a voxel v by $C[v] = \rho(B^t_{pred,sub-07}[v], B^t_{sub-07}[v])$

Method	**Mean correlation**
Baseline (average) (whole brain)	0.087
Baseline (average) (visual area)	0.180
INT - 200 parcels (whole brain) - A posteriori identification	0.158
INT - 200 parcels (visual area) - A posteriori identification	0.185
INT - 200 parcels (whole brain) - No identification	- 0.02
INT - 200 parcels (visual area) - No identification	0.01

TABLE 1 – Compilation of average correlations voxelwise (53 timepoints)

4.3 Co-smoothing - IBC-Contrasts

Task - We did a co-smoothing on 7 subjects with 6 subjects used for the template and one left-out subject as the target. For each of the 7 subjects, we have two sets of acquisitions 'AP' and 'PA', 53 time points each. We take a subject (`sub-07`) as the target subject. We train the model on the 6 other subjects to get a global stimulus on both 'AP' and 'PA', while we extract a specific tuning matrix for the left-out subject on the 'AP' part using the first half of the shared stimulus matrix. We obtain a reconstruction image of `sub-07` for the 'PA' acquisition by multiplying this tuning matrix with the second part of the stimulus. We perform all of that on masked images with only the visual regions of the brain.

$$U\Sigma V^t = [B^{(1)}_{AP+PA}, .., B^{(6)}_{AP+PA}]$$

$$S_{tot} := \sqrt{7} * U$$

$$S_{tot} := \begin{pmatrix} S_{tot,AP} \\ S_{tot,PA} \end{pmatrix}$$

$$T_{(target)} = S^{-1}_{tot,AP}\hat{B}^{(target)}_{AP}$$

$$B_{pred,PA} = S_{tot,PA}T_{(target)}$$

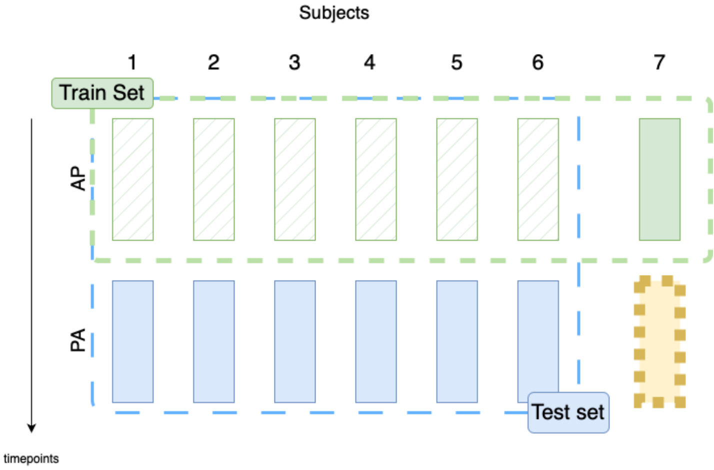

FIGURE 12 – Diagram of Co-smoothing experiment conducted on IBC-Contrast. 'AP' and 'PA' acquisitions both contain 53 time frames. The goal is to reconstruct sub-07's brain response for the 'PA' run from its tuning data acquired by template denoising on the train set and stimulus extracted from the test set.

Metric - We compute voxelwise Pearson's Correlation coefficient across all testing time frames (namely the 'PA' acquisition), between the ground truth `sub-07` and the prediction we got from the task. We then plot the resulting correlation map which has the same resolution as the original data.

Baseline - We used as a baseline for comparison an average image of all subjects except the target subject. This took the role of baseline prediction.

What should be noted is that even though both FastSRM and INT did worse than the Baseline average method, they managed to do quite well on some parts of the brain, especially the visual cortex.

We tried to use only Parcel/Searchlight Hyperalignment without the Neural Tuning phase (that is, decomposing the denoised signal between a stimulus matrix S and tuning matrices $T_1, ...T_7$. As a recall, before doing such a decomposition, $B_{(1)}, .B_{(7)}$ subjects fMRI data is decomposed between a Piecewise PCA template M_{PC} and individual transposition matrices $W_1, ..W_7$ through Piecewise Ridge regression with $\alpha = 10^{-3}, ..., \alpha = 10^3$ (Grid Search). If searchlights are used, we perform weighting on each searchlight PCA as searchlights overlap. This is done to compute a denoised version of subject data $\hat{B}_{(i)} = M_{PC}W_i$. What we did was just keep M_{PC} the template from denoising as the global template, and individual transposition matrices $W_1, ..W_7$ as "tuning information". We got some pretty interesting results there using this method (denoted as Parcel/Searchlight Hyperalignment). The drawback is that, by construction, rotations matrices are not invertible matrices (although square matrices), nor are they orthogonal. They do not represent any interesting transformation apart from "making the template stick to individual data". They are obtained with a basic Ridge regression in the original model, although, choosing a Procrustes rotation instead of Ridge regression for computing transposition matrices $W_1, ..W_7$ might be something to explore, as those matrices would be orthogonal.

4.3.1 Whole brain correlation

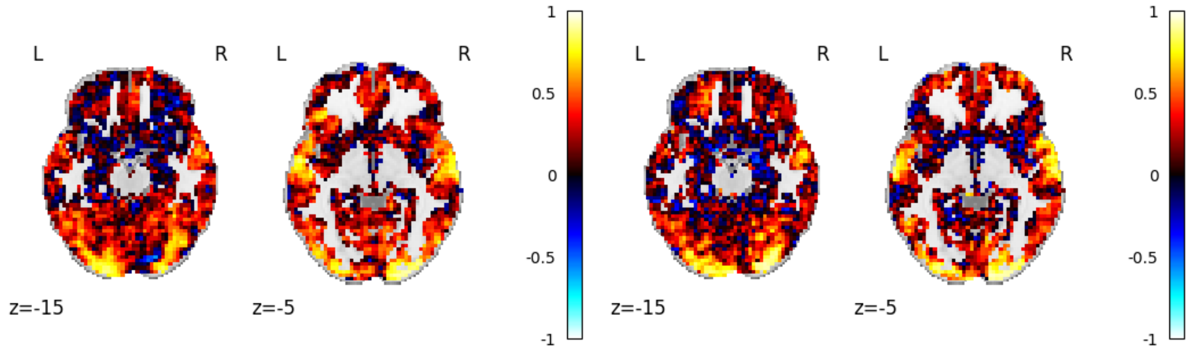

(a) Baseline (average)

(b) FastSRM - 5 components

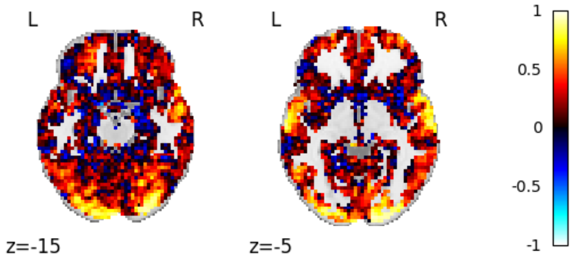

(c) FastSRM - 20 components

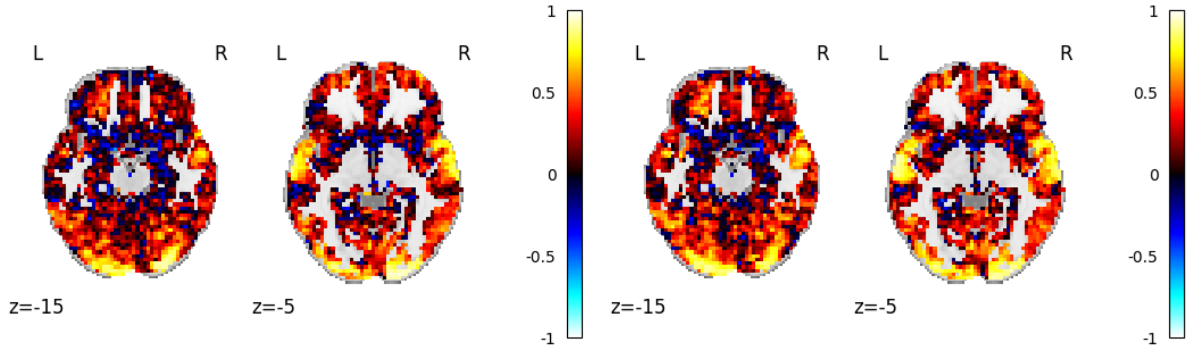

(d) INT - 200 parcels

(e) INT - 5mm searchlight

FIGURE 13 – IBC-Contrasts - Benchmark of INT vs. other methods on the whole brain (46678 voxels) - co-smoothing with stimulus from full "AP" + "PA" data on all subjects except `sub-07`, tuning on `sub-07` "AP" (53 time points). Prediction/ground-truth voxel correlation map for `sub-07`. The resulting map C is given for a voxel v by $C[v] = \rho(B^t_{pred,sub-07}[v], B^t_{sub-07}[v])$

Method	**Mean correlation**
Baseline (average)	0.203
FastSRM - 5 components	0.202
FastSRM - 20 components	0.194
INT - 100 parcels	**0.204**
INT - 5mm searchlight	**0.228**

TABLE 2 – Compilation of average correlations voxelwise (53 timepoints)

4.3.2 Movie watching regions correlation

Since the task is movie watching, we averaged the correlation only on the brain zones involved in movie watching. We used a ROI masker provided by the `fmralign` Python package.

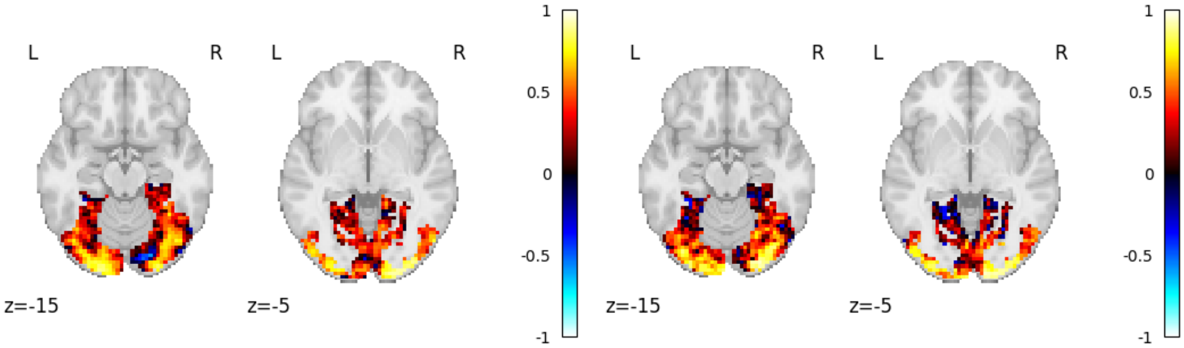

(a) Baseline (average)

(b) FastSRM - 5 components

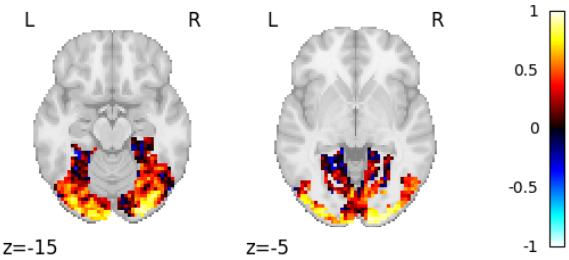

(c) FastSRM - 20 components

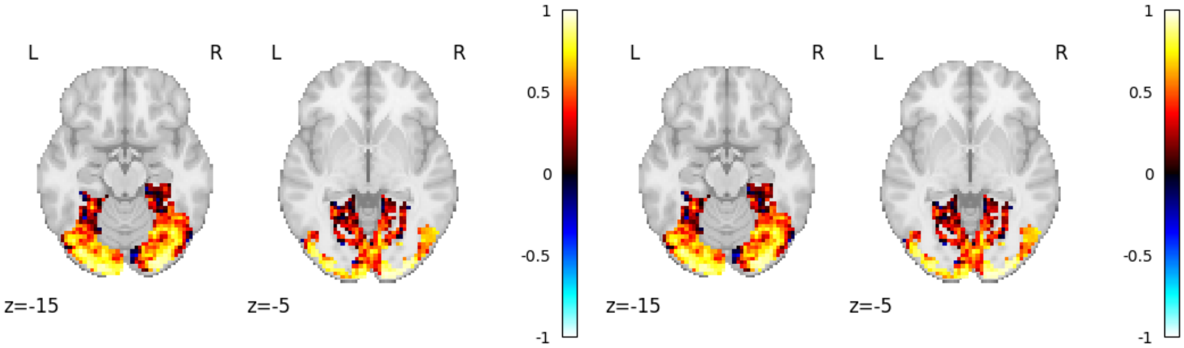

(d) INT - 200 parcels

(e) INT - 5mm searchlight

FIGURE 14 – IBC-Contrasts - Benchmark of INT vs. other methods on the visual area (6631 voxels) - co-smoothing with stimulus from full "AP" + "PA" data on all subjects except `sub-07`, tuning on `sub-07` "AP" (53 time points). Prediction/ground-truth voxel correlation map for `sub-07`. The resulting map C is given for a voxel v in the subset of voxel defining the visual area by $C[v] = \rho(B^t_{pred,sub-07}[v], B^t_{sub-07}[v])$

Method	**Mean correlation**
Baseline (average)	0.344
FastSRM - 5 components	0.384
FastSRM - 20 components	0.395
INT - 100 parcels	**0.400**
INT - 5mm searchlight	**0.428**

Table 3 – Compilation of average correlations voxelwise (53 timepoints)

5 Conclusion

We have asserted that although the model seems to work and give some results, we were not able to replicate exactly the results obtained by Feilong et. al., moreover, with higher resolution fMRI images, the model seems to lose in performance.

Additionally, from the experiments we did on both Searchlight hyperalignment denoising and simple parcellation denoising, we can not deduce a significative improvement from this technique, although the computations needed are much higher resulting in a considerably longer real-life runtime. Finally, the INT model does not seem to perform better than FastSRM, while being considerably more computationally intensive.

A Diagrams

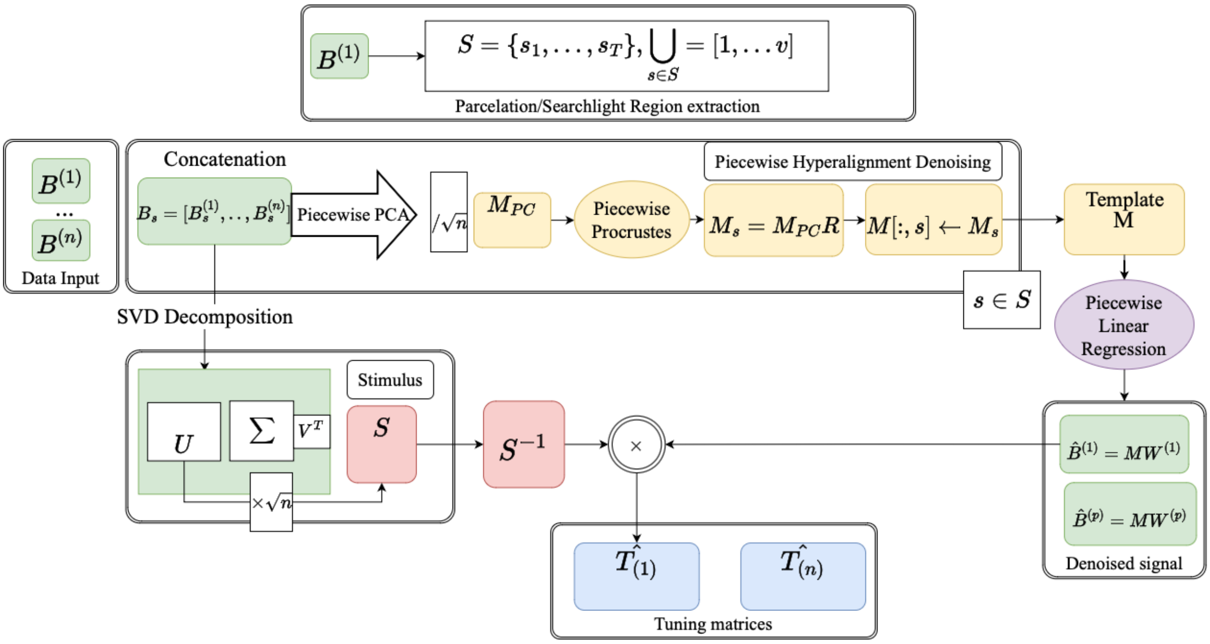

Figure 15 – Diagram of the Individualized Tuning Model from HaxbyLab. Note that if searchlight region extraction is considered, a weighting step is added before Piecewise Linear Regression

B Memory Complexity

We denote for this purpose :

- n_s : number of subjects
- n_v : number of voxels
- n_t : number of time points

B.1 Global Template M with searchlights

M is of size $n_t \times n_v$. There are always n_v searchlights, of average size s.

We iterate over all searchlights to compute a PCA over all participants. We therefore perform n_s PCAs.

In each searchlight : - We extract the desired data across all participants for all time points -> we get a matrix of size $n_s \times n_t k$. After PCA reduction, we end up with a matrix of size $n_t \times k$. - We therefore have a memory complexity of $O(n_s * n_t * k)$

So a total memory complexity of

$$C_M = O(n_v * n_s * n_t * k)$$

B.2 Individual transformations $W^{(i)}$

We do have to compute n_s such matrices of size $n_v \times n_v$.

In each searchlight : - We perform a Ridge regression with input matrix X of size $n_t * k$ and objective matrix Y of size $k \times k$. Each Ridge has a memory complexity of $O(n_t * k + k^2)$

- We do that on each searchlight for each subject so a total complexity of :

$$n_s * \sum_{sl} 2|sl|^2 = O(n_s * n_v * k^2 + n_s * n_v * n_t * k)$$

$$C_W = O(n_s * n_v * k(k + n_t))$$

B.2.1 Numerical examples for phase 2

For the IBC-Sound dataset, we have :

- $n_s = 10$

— $n_t = 53$
— $n_v = 46407$
— $k = 704$

We store matrices as `ndarray` with `dtype=float64`. The raw data has size `197 MB`.

If we were to store all the data while computing, we would end up with at least (conservative estimations not taking all data into account) 2TB.

If we trust Python and its garbage collector it is supposed to be far less, but if we multitask, and do only one searchlight at a time on each subject we have an instantaneous complexity divided by the number of searchlights, so approx. 43 MB. If we want to cap the mem usage at let's say 10 GB, we can compute 232 searchlights at the same time.

B.3 Comparaison with simple parcelation

As we have the nice property that

$$\sum_{p \in parcels} |p| = n_v$$

We directly have the average size of a parcel $k = \frac{n_v}{n_p}$ with n_p the number of parcels.

Thus complexity for phase 1 (computation of M) :

$$C_M = O(n_s * n_t * n_v)$$

And for phase 2 :

$$C_W = O(n_s * n_v * (n_t + k))$$

So a complexity improvement of a factor $O(k)$ in both.

In application, generally $k \sim 1000$ so a massive improvement.

Références

[1] Ma Feilong, Samuel A. Nastase, Guo Jiahui, Yaroslav O. Halchenko, M. Ida Gobbini, James V. Haxby ;
The individualized neural tuning model : Precise and generalizable cartography of functional architecture in individual brains. Imaging Neuroscience 2023

[2] Hugo Richard, Lucas Martin, Ana Luısa Pinho, Jonathan Pillow, Bertrand Thirion ;
Fast shared response model for fMRI data