

Compressed Online Dictionary Learning for Fast Resting-State fMRI Decomposition

Arthur Mensch, Gaël Varoquaux, Bertrand Thirion

► **To cite this version:**

Arthur Mensch, Gaël Varoquaux, Bertrand Thirion. Compressed Online Dictionary Learning for Fast Resting-State fMRI Decomposition. International Symposium on Biomedical Imaging, IEEE, Apr 2016, Prague, Czech Republic. pp.1282-1285, 10.1109/ISBI.2016.7493501 . hal-01271033v2

HAL Id: hal-01271033

<https://hal.archives-ouvertes.fr/hal-01271033v2>

Submitted on 12 May 2016

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.



COMPRESSED ONLINE DICTIONARY LEARNING FOR FAST RESTING-STATE FMRI DECOMPOSITION

Arthur Mensch⁽¹⁾ Gaël Varoquaux⁽¹⁾ Bertrand Thirion⁽¹⁾

⁽¹⁾Parietal team, Inria, CEA, Paris-Saclay University. Neurospin, 91191 Gif-sur-Yvette, France

ABSTRACT

We present a method for fast resting-state fMRI spatial decompositions of very large datasets, based on the reduction of the temporal dimension before applying dictionary learning on concatenated individual records from groups of subjects. Introducing a measure of correspondence between spatial decompositions of rest fMRI, we demonstrate that time-reduced dictionary learning produces result as reliable as non-reduced decompositions. We also show that this reduction significantly improves computational scalability.

Index Terms— resting-state fMRI, sparse decomposition, dictionary learning, online learning, range-finder

1. INTRODUCTION

Resting-state fMRI data analysis traditionally implies, as an initial step, to decompose a set of raw 4D records (time-series sampled in a volumic voxel grid) into a sum of spatially located *functional networks* that isolate a part of the brain signals. Functional networks, that can be seen as a set of brain *activation maps*, form a relevant basis for the experiment signals that captures its essence in a low-dimensional space. As such, they have been successfully used for feature extraction before statistical learning, *e.g.* in decoding tasks.

While principal component analysis (PCA) on image arrays has been the first method to be proposed for fMRI, independent component analysis (ICA) is presently the most popular decomposition technique in the field. It involves finding a spatial basis \mathbf{V} that is closest to a set of spatially *independent* sources. More recent work have shown that good results can be obtained imposing *sparsity* rather than *independence* to spatial decomposition [1], relying on *dictionary learning* formulation [2].

All these techniques suffer from their lack of scalability, as they were initially designed to be applied to small datasets. The recent increase in publicly available dataset size (*e.g.* HCP [3]) has revealed their limits in terms of memory usage and computational time. Efforts have been made to make decomposition methods available for large scale studies, possibly with several groups. They involve using a hierarchical model for dictionary learning [1] or incremental PCA techniques [4]. However, the former only proposes PCA+ICA based decomposition methods, which do not naturally yield sparse maps, and the latter suffers from its computational complexity. Running a satisfying decomposition algorithm on the full HCP dataset currently requires a very large workstation.

In this paper, we focus on dictionary learning methods for fMRI, and show how to make them more scalable in both time and memory. Uncovering the computational limitations of dictionary learning when analysing very large datasets, we propose to perform random-projection based hierarchical dimension reduction in the time direction before applying dictionary learning methods. As a result, time

and memory consumption are reduced, avoiding out-of-core computation. We introduce a measure of correspondence to relate results obtained from compressed data to those from non compressed data, and show that substantial gain in time and memory can be obtained with no significant loss in quality of the extract networks.

2. SCALABILITY OF DICTIONARY LEARNING FOR FMRI

2.1. rfMRI decomposition existing formalism

We consider multi-subject rfMRI data: a set of matrices $(\mathbf{X}^s)_{s \in [1, t]}$ in $(\mathbb{R}^{n \times p})^t$, with p voxels per volume, n temporal samples per record, and t records. We seek to decompose it as :

$$\forall s \in [1, t], \mathbf{X}^s = \mathbf{U}^s \mathbf{V}^T \quad \text{with } \mathbf{U}^s \in \mathbb{R}^{n \times k}, \mathbf{V} \in \mathbb{R}^{p \times k} \quad (1)$$

Existing decomposition techniques vary in the criterion they optimize, and on the hierarchical model they propose. We focus on dictionary learning methods, that have been shown to obtain better results than ICA in [1]. To handle group studies, we choose the most simple hierarchical model, that consists in performing time concatenation of the records – first proposed by [5] for ICA. We write $\mathbf{U} \in \mathbb{R}^{nt \times k}$ and $\mathbf{X} \in \mathbb{R}^{nt \times p}$ the vertical concatenation of $(\mathbf{U}^s)_s$ and $(\mathbf{X}^s)_s$, and seek to decompose \mathbf{X} instead of \mathbf{X}^s .

A good decomposition should allow a good reconstruction of data while being spatially localized, *i.e.* *sparse* in voxel space. Such a decomposition setting can be formalized in a *dictionary learning* (DL) optimization framework, that combines a sparsity inducing penalty to a reconstruction loss. We seek to find k dense *temporal atoms*, *i.e.* time-series, that will constitute loadings for k sparse spatial maps with good signal recovery. In one of its original formulation [2], this leads to the following optimization problem:

$$\min_{\substack{\mathbf{U} \in \mathbb{R}^{nt \times k}, \\ \mathbf{V} \in \mathbb{R}^{p \times k}}} \|\mathbf{X} - \mathbf{U} \mathbf{V}^T\|_F^2 + \lambda \|\mathbf{V}\|_1 \quad \text{s.t. } \forall j, \|\mathbf{U}_j\|_2 \leq 1 \quad (2)$$

Each row $L_i(\mathbf{V})$ yields the sparse k loadings related to the k temporal atoms for a single voxel time-series, held in column X^i . [6] introduces an efficient online solver for this minimization problem, streaming on *voxel time-series*, *i.e.* loading \mathbf{X} columnwise: at iteration t , a voxel time-series batch $L_{b(t)}(\mathbf{V})$ is computed (using a Lasso solver) on the present dictionary \mathbf{U}_{t-1} , and \mathbf{U}_t is updated (using block coordinate descent) to best reconstruct previously seen time-series from previously computed sparse codes. The final spatial components are then obtained solving Lasso problems $\min_{\mathbf{V} \in \mathbb{R}^{p \times k}} \|\mathbf{X} - \mathbf{U}_{\text{end}} \mathbf{V}^T\|_F^2 + \lambda \|\mathbf{V}\|_1$.

This online algorithm provably converges towards a solution of Eq. 2 under conditions satisfied in neuro-imaging. A good initialization for temporal atoms is required to obtain an exploitable solution. It can typically be obtained by computing time-series associated to an initial guess on activation maps \mathbf{V}_{init} , *e.g.* obtained from

known brain networks. The temporal atoms are computed by solving $\min_{\mathbf{U}_i} \|\mathbf{X}_i - \mathbf{U}_i \mathbf{V}_{\text{init}}^\top\|_2$ for all $i \in \llbracket 1, n \rrbracket$.

2.2. Scalability challenge

Following [6], online dictionary learning has an overall complexity of $\mathcal{O}(n p k^2)$, as convergence is typically reached within one epoch on rfMRI. In theory, the dictionary learning problem is thus computationally scalable. However, on large rfMRI datasets, online dictionary learning faces two main challenges detailed below.

Out-of-core requirements for large datasets For datasets like HCP ($t = 2000$, $n = 1200$, $p = 20000$, 1.92TB), typical computers are unable to hold all data in memory. It is thus necessary to stream the data from disk, which is only reasonably efficient if the data are stored in the same direction as it is accessed. Yet online DL algorithm require to pass 3 times over data, during which it is streamed in different directions (row-wise for initialization, columnwise for DL and final Lasso solving), while fMRI images are naturally stored row-wise. For the sake of efficiency, storage copy and manipulation is required, which is a serious issue for neuroscientists dealing with over 1TB datasets. Going out-of-core sets a large *performance gap* between small datasets and large datasets.

Grid search in parameter setting The sparsity of the maps obtained depends critically on parameter λ , that scales non trivially with p . It is therefore impossible to set it independently from the *experiment size*, and several runs must be performed to obtain best maps, relative to their neurological relevance or a validation criterion. Grid search should be run in parallel for efficiency, which is a serious issue when doing out-of-core computation, as simultaneous access to the disk from different processes makes the pipeline IO-bound. Reducing dataset size therefore reduces disk and memory usage, which permits the efficient use of more CPUs.

Both issues suggest to reduce memory usage by reducing datasets size while keeping the essential part of its signal: being able to keep data in memory avoids drastic loss in performance.

3. TIME-COMPRESSED DICTIONARY LEARNING

Reducing time dimension Good quality maps are already obtained using small datasets with standard number of samples (ADHD dataset, $n = 150$). For this reason, we investigated how large datasets can be reduced to fit in memory while keeping reasonable map accuracy compared to the non-reduced version.

Indeed, the n time samples per subject are not uniformly scattered in voxel space, and should exhibit some low dimension structure: we expect them to be scattered close to some low rank subspace of \mathbb{R}^p , spanned by a set of m vector $\mathbf{X}_r^s \in \mathbb{R}^{m \times p}$. We thus perform a *hierarchical rank reduction*: \mathbf{X}^s is first approximated by a rank m surrogate $\mathbf{P}^\top \mathbf{X}_r^s$, and a final rank k decomposition is computed over concatenated data. We show that such reduction is conservative enough to allow good map extraction. Geometrically, we project \mathbf{X} on a low rank subset of $\mathbb{R}^{n \times p}$:

$$\mathbf{P} = \underset{\mathbf{Q} \in \mathbb{R}^{n \times m}}{\operatorname{argmin}} \left\| \mathbf{X}^s - \mathbf{Q} \mathbf{Q}^\top \mathbf{X}^s \right\|_F \quad \mathbf{X}_r^s = \mathbf{P}^\top \mathbf{X}^s \quad (3)$$

Then $\mathbf{X}^s = \mathbf{P} \mathbf{X}_r^s + \mathbf{E}^s$ where \mathbf{E}^s is a residual full rank noise matrix.

We approximate \mathbf{X}^s with \mathbf{X}_r^s at *subject level* to retain subject variability. Hence, replacing \mathbf{X} with \mathbf{X}_r , the concatenation of (\mathbf{X}_r^s) , in Eq. 2, we obtain a *reduced* dictionary learning objective.

Importantly, we must have $m t > k$ so that \mathbf{X}_r is at least of rank k to recover k sparse activation maps. On the other hand, we show that reducing \mathbf{X}_r^s matrix beyond $m < k$ can still provide good results.

In our reduced dictionary learning algorithm, time and memory complexity are reduced by a factor $\alpha = \frac{m}{n}$, where m should typically be of the same order than k . This linear speed-up becomes much more dramatic when reduction allows to go from out-of-core to in-core computation. It comes to the cost of the time required for matrix reduction that we study in the following paragraph.

While Eq. 3 can be seen as another way of decomposing $(\mathbf{X}^s)_s$, let us stress that this decomposition is performed in voxel space, in contrast with dictionary learning itself, that identify a good basis in *time* space. The objective is to quickly find a good *summary* of each $(\mathbf{X}^s)_s$ prior to applying dictionary learning, so as to reduce the dimensionality of the dictionary learning problem.

The range-finding approach \mathbf{X}_r^s can be computed exactly with truncated SVD, following Eckart–Young–Mirsky theorem. However, exact SVD computation is typically $\mathcal{O}(p n^2)$, which is above dictionary learning complexity and makes prior data reduction useless when trying to reduce *both* computation time and memory usage. Fortunately, we show that we do not need exact m rank best approximation of \mathbf{X} to obtain a satisfying \mathbf{V} . Following [7] formalism, we seek $(\hat{\mathbf{P}}^s)_s \in (\mathbb{R}^{n \times m})^t$ such that

$$\|\mathbf{X}^s - \hat{\mathbf{P}}^s \hat{\mathbf{P}}^{s\top} \mathbf{X}^s\|_F \approx \min_{\operatorname{rank}(\mathbf{Y}^s) \leq m} \|\mathbf{E}^s\| = \|\mathbf{X}^s - \mathbf{Y}^s\|_F \quad (4)$$

In [7], Alg. 4.4, Halko proposes a fast, randomized algorithm to compute such $\hat{\mathbf{P}}^s$, with measurable precision $\|\hat{\mathbf{E}}^s - \mathbf{E}^s\|$. Setting $\hat{\mathbf{P}} = \operatorname{Diag}((\hat{\mathbf{P}}^s)_s)$, $\hat{\mathbf{X}}_r = \hat{\mathbf{P}} \mathbf{X}$, we use this random range-finding (rf) algorithm to solve Eq. 2, where we replace \mathbf{X} with $\hat{\mathbf{X}}_r$:

$$\min_{\substack{\mathbf{U}_r \in \mathbb{R}^{m t \times k} \\ \mathbf{V} \in \mathbb{R}^{p \times k}}} \left\| \hat{\mathbf{P}} \mathbf{X} - \mathbf{U}_r \mathbf{V}^\top \right\|_F^2 + \lambda \|\mathbf{V}\|_1 \quad \text{s.t.} \quad \|(\mathbf{U}_r)_j\|_2 \leq 1 \quad (5)$$

The randomized range finding algorithm has a complexity of $\mathcal{O}(n p m)$, which is of same order as dictionary learning algorithm. In practice, we show in Sec. 5 that its cost becomes negligible with respect to the reduction of dictionary learning cost, when the reduction ratio is high enough.

In a more straightforward way, we can set $\mathbf{X}_r^s = \mathbf{X}_I^s$, with I subset (ss) of $\llbracket 1, n \rrbracket$ of size m . This category of reduction includes time subsampling of records. In this case, $\|\hat{\mathbf{E}}_{\text{ss}}^s - \mathbf{E}^s\|$ cannot be controlled, and is expected to be larger than $\|\hat{\mathbf{E}}_{\text{rf}}^s - \mathbf{E}^s\|$. Subsampling, for example, is expected to alias high frequency signal in records, preventing the recovery of activation maps with high frequency loadings in final dictionary learning application.

4. VALIDATION

Reference result-set Validation of dictionary learning methods for rfMRI is challenging, as there is no ground truth to assess the quality of resulting map sets. However, we can assess how much a result-set \mathbf{V} obtained on a reduced dataset \mathbf{X}_r from Eq. 5 is comparable to a result-set \mathbf{V}^0 obtained on \mathbf{X} from Eq. 2.

Result-set comparison Two sets of maps \mathbf{V}^0 and \mathbf{V} can only be compared with an indicator invariant to map ordering. Two sets are *comparable* if each map from the first set is comparable to a map in the second set. We find the best one-to-one coupling between

these two sets of maps and compute correlation between each *best assigned* couple of maps: $\text{corr}(\mathbf{v}_i^0, \mathbf{v}_j) = \frac{|(\mathbf{v}_i^0)^\top \mathbf{v}_j|}{\|\mathbf{v}_i^0\|_2 \|\mathbf{v}_j\|_2}$ to measure similarity between two maps \mathbf{v}_j (held in column $C_j(\mathbf{V})$) and \mathbf{v}_i^0 . We set d to be the mean correlation between best assigned maps:

$$d(\mathbf{V}, \mathbf{V}^0) = \max_{\Omega \in S_k} \text{Tr}(\mathbf{V}^\top \Omega \mathbf{V}^0) \quad (6)$$

where S_k is the set of permutation matrices. Ω can be computed efficiently using the Hungarian algorithm.

Comparing random results Eq. 2 and 5 admits many local minima that depend on algorithm initialization, and on the order used for streaming dataset columns. For any dataset $\mathbf{Y} \in \{\mathbf{X}, (\mathbf{X}_r)_{\text{method}}^{\text{reduction}}\}$, we expect obtained maps $\mathbf{V}_i = \text{DL}_i(\mathbf{Y})$ to capture a neurological/physical phenomenon for any run i corresponding to a streaming order. As in [8], we perform l runs numbered on $S_l \subset \mathbb{N}$ of the algorithm to obtain different maps, and compare the *concatenation* $\mathcal{V}_l(\mathbf{Y}) = [(\mathbf{V}_i)_{i \in S_l}]$ of these maps to the concatenation of reference maps $\mathcal{V}_l^0(\mathbf{X}) = [(\mathbf{V}_i^0)_{i \in S_l^0}]$ with runs numbered on S_l^0 :

$$d_l(\mathbf{X}, S_l^0, \mathbf{Y}, S_l) = d(\mathcal{V}_l^0(\mathbf{X}), \mathcal{V}_l(\mathbf{Y})) \quad (7)$$

We thus take into account non unicity of DL solutions: different maps are obtained when performing the dictionary learning algorithm over the *same* data with the *same* parameters. We model result maps $(\mathbf{v}_i)_i$ to be part of a larger *full* result-set \mathcal{V} :

$$\mathcal{V}(\mathbf{Y}) = \left\{ \mathbf{v}_i = C_i(\mathbf{V}) \in \mathbb{R}^p \text{ s.t. } \mathbf{V} \in \mathbb{R}^{p \times k}, \exists \mathbf{U} \in \mathbb{R}^{n \times k}, \right. \\ \left. (\mathbf{U}, \mathbf{V}) \in \underset{\mathbf{U}, \mathbf{V}}{\text{argmin}} \left\| \mathbf{Y} - \mathbf{U}\mathbf{V}^\top \right\|_F^2 + \lambda \|\mathbf{V}\|_1 \right\} \quad (8)$$

When result-sets are concatenated over all possible streaming orders, we expect d_p to converge toward a $S_p^{(0)}$ independent measure:

$$d_\infty(\mathcal{V}(\mathbf{X}), \mathcal{V}(\mathbf{Y})) = \lim_{l \rightarrow \infty} d_l(\mathbf{X}, S_l^0, \mathbf{Y}, S_l) \quad (9)$$

It is expected that $d_\infty(\mathcal{V}(\mathbf{X}), \mathcal{V}(\mathbf{X})) = 1$, but p is finite in practice. Ensuring $S_l^0 \cap S_l = \emptyset$, we measure mean result-set correspondence $d_l(\mathbf{X}, S_l^0, \mathbf{X}, S_l)$ over different runs on the same dataset \mathbf{X} , and compare it to $d_l(\mathbf{X}, S_l^0, \mathbf{Y}, S_l)$ to assess the *reduction* effect.

5. RESULTS

Tools and datasets We validate our reduction framework over two different datasets with different size: ADHD data, with 40 records, $n = 150$ time steps per record; a subset of HCP dataset, using 40 subjects, 2 records per subject, subsampling records from $n = 1200$ to $n = 400$ to obtain reference \mathbf{X} .

Dictionary learning output depends on its initialization, and the problem of choosing the *best* number of components k is very ill-posed. We bypass these problems by choosing $k = 70$ for HCP dataset, $k = 20$ for ADHD dataset, and use reference ICA-based maps RSN20 and RSN70 from [9] as initialization – we prune unused dictionary atoms on HCP dataset.

For benchmarking, we measure CPU time only, *i.e.* ignore IO time as it is very platform dependent. To limit disk access in out-of-core computation, small memory usage is crucial for IO time.

We use *scikit-learn* for computation, along with the *Nilearn* neuro-imaging library. Code for the methods and experiments is available at <http://github.com/arthurmensch/nilearn/tree/isisbi>.




Fig. 1. Result-set correspondence with non-reduced DL result-set, using different methods with different reduction ratios, increasing number of runs to show d_l stabilization; variance over runs computed using 4 different result-sets $S_l^{(0)}$; ADHD dataset.

Indicator and reduction validity Fig. 1 shows d_l behavior as l increases. The results demonstrate the relevance of random range-finding as it out-performs simple subsampling. We first obtain a *reference* set of maps \mathcal{V}_p^0 from non-reduced \mathbf{X} , choosing λ to obtain little overlapping maps ($\lambda = 1, 6$ for ADHD, HCP). Secondly, we compute $d_l(\mathbf{X}, S_l^0, \mathbf{Y}, S_l)$ setting $\mathbf{Y} = \{\mathbf{X}, (\mathbf{X}_r)_{\text{rf}, m}, (\mathbf{X}_r)_{\text{ss}, m}\}$, for various $m \in [n/40, n]$. As the relationship between λ and a given level of sparsity depends on m , we run DL on \mathbf{Y} on a range of λ so as to find the value that matches best the reference run.

We observe that running DL several times does produce sets of maps that overlap more and more, as they cover a larger part of the result-sets \mathcal{V} defined in Eq. 8, and stabilizes for $l > 10$. This suggest that d_l does cater for randomness in DL algorithms and constitutes a good indicator for comparing two DL methods.

For $\alpha > .025$, and $l \geq 2$, Fig. 1 shows that compressed DL produces maps that are *as comparable* with non-reduced DL maps as non-reduced DL maps obtained streaming on different orders:

$$d_l(\mathbf{X}, S_l^0, \mathbf{X}, S_l) \approx \hat{d}_l(\mathbf{X}, S_l^0, \mathbf{X}_r, S_l) \quad (10)$$

Overlap between $\mathcal{V}_l(\mathbf{X})$ and $\mathcal{V}_l(\mathbf{X}_r)$ is thus comparable to overlap between $\mathcal{V}_l(\mathbf{X})$ and $\mathcal{V}_l(\mathbf{X})$ for different runs from S_l, S_l^0 . They are therefore of the same inner quality for neuroscientists as it not possible to tell one apart from the other.

For large compression factors – typically with $m < k$, for $\alpha < .1$ on ADHD, $\alpha < .05$ on HCP – range finding reduction performs significantly better than subsampling. Both methods perform similarly for small compression factors, which shows that subsampling already provides good *large* low-rank approximation of \mathbf{X} . Using a range-finding algorithm is therefore useful when drastically reducing data size, typically when loading very large datasets in memory.

Qualitative accuracy We validate qualitatively our results, as this is crucial in DL decomposition: maps obtained from reduced data should capture the same underlying neurological networks as reference maps. In Fig. 2, we display matched maps when comparing two result-sets. For this, we find matchings between sets $(\mathcal{V}_l, \mathcal{V}_l^0)$, and we display the maps corresponding to the median-value of this matching. Maps are strongly alike from a neurological perspective. In particular, maps do not differ more between our reduced dictionary learning approach and the reference algorithm than across two runs of the reference algorithm.




Fig. 2. Median aligned maps with various methods; HCP dataset; reduction ratio $\alpha = .025$.




Fig. 3. Time/accuracy using range-finder projectors and subsampling before DL; blue stripe recalls correspondence of results when performing different runs on non-reduced \mathbf{X} . $l = 10, 3$ for ADHD, HCP. Variance over runs computed using 4 distinct subject sets $S_l^{(0)}$.

Time and accuracy tradeoff For efficient neuroimaging data analysis, the important quantity is the tradeoff between quality of the results and computation time. On Fig. 3, we plot $d_l(\mathbf{X}, \mathbf{Y})$ – omitting S_l, S_l^0 in notation – against computational CPU time, for various \mathbf{Y} . Using range-finding algorithm and to a lesser extent time subsampling on data before map decomposition does not significantly deteriorate results up to large reduction factor, while allowing large gains in time *and* memory. Compression can be higher for larger datasets: we can reduce our HCP subset up to 40 times, ADHD up to 20 times, keeping $d_l(\mathbf{X}, \mathbf{X}_r)$ within the standard deviation of $d_l(\mathbf{X}, \mathbf{X})$.

The range-finder algorithm adds a time overhead that shift performance curve towards higher time for large compression. However, it allows 4 times lower memory usage and thus higher overall efficiency when considering IO. Moreover, benchmarks were performed on a single core, while reduction can be parallelized over subjects to reduce its overhead.

We outline best time/accuracy trade-off reduction ratios in Fig. 3 and Table 1. They depend on chosen k and on dataset, but any reasonably low reduction (with $m \gtrsim k$) ratio is likely to produce good results with little accuracy loss. Following this strategy, we set $\alpha = .025$ and performed the entire processing of 100 subjects of the HCP dataset (384GB) on a single workstation (64GB RAM) in less than 7 hours.

Data	RF α	CPU Time		Corresp. $d_l(\mathbf{X}, \mathbf{Y})$	
		Red.	N-red.	Reduced	Non-red.
HCP	.025	849 s	7425 s	$.703 \pm .141$	$.628 \pm .105$
ADHD	.05	71 s	186 s	$.796 \pm .020$	$.801 \pm .016$

Table 1. Time/accuracy with most interesting method for each dataset, comparing to reference DL run. RF α : range-finder ratio

6. ACKNOWLEDGEMENT

The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 604102 (Human Brain Project).

7. CONCLUSION

We introduce the use of a randomized range finding algorithm to reduce large scale datasets before performing dictionary learning and extract spatial maps. To prove efficiency of time reduction before dictionary learning, we have designed a meaningful indicator to measure result maps *correspondence*, and have demonstrated that fMRI time samples have a low rank structure that allows range finding projection to be more efficient than simple subsampling.

This approach enables a 40-fold data reduction upon loading of each subjects. It thus makes processing large datasets such as the HCP (1.92TB) tractable on a single workstation, time-wise and memory-wise.

8. REFERENCES

- [1] G. Varoquaux, A. Gramfort, F. Pedregosa, V. Michel, and B. Thirion, “Multi-subject dictionary learning to segment an atlas of brain spontaneous activity,” *IPMI*, vol. 22, pp. 562, 2011.
- [2] K. Kreutz-Delgado, J. F. Murray, Bhaskar D. Rao, et al., “Dictionary learning algorithms for sparse representation,” *Neural computation*, vol. 15, pp. 349, 2003.
- [3] D.C. Van Essen, K. Ugurbil, E. Auerbach, et al., “The Human Connectome Project: A data acquisition perspective,” *NeuroImage*, vol. 62, pp. 2222, 2012.
- [4] S. M. Smith, A. Hyvärinen, G. Varoquaux, K. L. Miller, and C. F. Beckmann, “Group-PCA for very large fMRI datasets,” *NeuroImage*, vol. 101, pp. 738, 2014.
- [5] V.D. Calhoun, T. Adali, G.D. Pearlson, and J.J. Pekar, “A method for making group inferences from functional mri data using independent component analysis,” *Hum. Brain Mapp.*, 2001.
- [6] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, “Online learning for matrix factorization and sparse coding,” *The Journal of Machine Learning Research*, vol. 11, pp. 19–60, 2010.
- [7] N. Halko, P. G. Martinsson, and J. A. Tropp, “Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions,” *SIAM Rev.*, vol. 53, 2011.
- [8] J. Himberg, A. Hyvärinen, and F. Esposito, “Validating the independent components of neuroimaging time series via clustering and visualization,” *Neuroimage*, vol. 22, pp. 1214, 2004.
- [9] S.M. Smith, P.T. Fox, K.L. Miller, et al., “Correspondence of the brain’s functional architecture during activation and rest,” *Proc. Nat. Acad. Sci.*, vol. 106, pp. 13040, 2009.