Published in final edited form as:

Inf Process Med Imaging. 2017 June ; 10265: 275–286. doi:10.1007/978-3-319-59050-9_22.

Riccati-Regularized Precision Matrices for Neuroimaging

Nicolas Honnorat and Christos Davatzikos

Department of Radiology, University of Pennsylvania, 3700 Hamilton Walk, Richards Building, 7th Floor, Philadelphia, PA 19104, USA

Abstract

The introduction of graph theory in neuroimaging has provided invaluable tools for the study of brain connectivity. These methods require the definition of a graph, which is typically derived by estimating the effective connectivity between brain regions through the optimization of an ill-posed inverse problem. Considerable efforts have been devoted to the development of methods extracting sparse connectivity graphs.

The present paper aims at highlighting the benefits of an alternative approach. We investigate low-rank L2 regularized matrices recently introduced under the denomination of Riccati regularized precision matrices. We demonstrate their benefits for the analysis of cortical thickness map and the extraction of functional biomarkers from resting state fMRI scans. In addition, we explain how speed and result quality can be further improved with random projections. The promising results obtained using the Human Connectome Project dataset, as well as, the numerous possible extensions and applications suggest that Riccati precision matrices might usefully complement current sparse approaches.

Keywords

rs-fMRI; Precision; Sparse inverse covariance

1 Introduction

Resting-state functional MRI (rs-fMRI) studies of brain connectivity have received a considerable amount of interest. Thanks to the continuous improvement of imaging techniques and the development of big data infrastructures, large datasets are now available for conducting these investigations [17]. An increasing effort has been devoted to the development of mathematical frameworks able to summarize these humongous datasets into robust and concise causal models and biomarkers, with the hope of better describing cognition, teasing out the mechanisms underlying brain diseases and defining novel clinical dimensions [12]. Different measures of connectivity were proposed [18] and graph theoretical approaches have been widely spread. [2]. The most straightforward approaches assume that the time series observed during a rs-fMRI scan are generated by a multivariate Gaussian process and attempt to analyze its structure. These studies typically start with the definition of locations in the gray matter once rs-fMRI have been registered, motion

corrected, denoised and normalized. The connectivity between these nodes is measured by computing the covariance of their time series and a sparse graph built from the inverse of the covariance matrix, also known as precision matrix. These last two steps can be performed jointly, by directly estimating sparse precision matrices [6,18,19]. However, and despite very impressive recent development, the estimation of sparse precision matrices is still time-consuming for large matrices [6,11].

In this paper, we propose an alternative approach based on low-rank Riccati regularized precision matrices, introduced first by Witten and Tibshirani [20] and formalized by Honorio and Jaakkola [10]. As for sparse matrices, we measure network characteristics directly from these matrices [16,19]. Our approach offers several benefits such as a very competitive computational efficiency, which deteriorates only linearly with data dimension, straightforward practical and theoretical extensions. We demonstrate that reducing the dimension of the input neuroimaging signals via random projections [9] can simultaneously improve test-retest performances and reduce computational burden, and we present two extensions: the estimation of precision matrices at a population level, and the adaptation of Riccati penalties to regions of interests. These results were established using the data available for the hundred unrelated subjects of the HCP dataset [17]. An in-depth test-retest validation was carried out by reducing the spatial dimension of the resting state fMRI (rs-fMRI) scans with Glasser et al. parcellation [8]. In addition, we demonstrate that our approach can handle full resolution data and other modalities by analyzing cortical thickness maps.

The remainder of the paper is organized as follows. The methods combined in this work are presented in Sect. 2. Section 3 presents several variants of our approach addressing related neuroscience applications. The experimental results are presented in Sect. 4 and discussion concludes the paper.

2 Methods

The random projection method, described in Sect. 2.1, was used as a preprocessing step for reducing the dimensionality of our imaging data and filtering noise. We present, in Sect. 2.2, a generalization of Honorio and Jaakkola Riccati regularized precision matrices [10]. The Gaussian entropy introduced by Tononi, Sporns and Edelman [16] for measuring functional network integration is presented in Sect. 2.3. We used this measure for extracting biomarkers from Riccati regularized precision matrices.

2.1 Random Projection

Random projections were proposed for compressing high-dimensional measurement while preserving their Euclidean distance. The random projections proposed by Halko, Martinsson and Tropp in [9] achieve performances close to a truncated singular value decomposition (TSVD): when a data matrix X of size $N \times T$, $T < N$ is projected for creating a thinner matrix Y of size $N \times t$, $t < T$, the t non-zero singular values of Y are close to the t largest singular values of X . These random projections were proposed to accelerate the computation of singular value decompositions (SVD) [9].

[9] is straightforward to implement. Figure 1 provides the pseudo code of the random projection algorithm used for this work. This algorithm generates an orthogonal projection matrix by randomly combining matrix rows and orthonormalizing the basis obtained through the Gram-Schmidt process. As explained in [9] and reported in Sect. 4.1 for the HCP data, the quality of the projection basis can be significantly improved by running a few power iterations, but the cost of computation grows rapidly with the number of matrix rows.

2.2 Riccati Regularized Precision Matrices

Let X denote a matrix of size $N \times T$ containing N time series of T time points, normalized to zero mean and unit variance. The associated covariance matrix will be referred as

$C = \frac{1}{T} X X^T$. Sparse, Tikhonov and Riccati regularized precision matrices are obtained by solving the following optimization problem:

$$\operatorname{argmax}_{Q \succ 0} [\log \det Q - \langle C, Q \rangle - \rho R(Q)] \quad (1)$$

where an L1 norm is chosen for R for generating sparse precision matrices [19], the trace of Q in the case of Tikhonov regularization [10], and R is the square of the Frobenius norm for Riccati regularized precision matrices [10]. As explained in [10,20], Riccati regularization is a ridge penalty on the components of the precision matrix whereas, when precision matrices are computed for solving a linear regression, Tikhonov regularization corresponds to a ridge penalty on the coefficients of a linear regression. In this work, we generalize the Riccati regularization described in [20] and [10] by introducing an invertible matrix V and working with the penalty:

$$R(Q) = \frac{1}{2} \|VQV^T\|_2^2 \quad (2)$$

When V is a diagonal matrix, $V = \operatorname{diag}(v)$, the penalty R becomes a squared weighted Frobenius norm, which can be expressed as follows:

$$R(Q) = \frac{1}{2} \|B \odot Q\|_2^2 \quad (3)$$

$$B = vv^T \quad (4)$$

where $\odot$ denotes the Hadamard product. This specific case, which is easy to interpret and interesting for applications, will be referred as Hadamard-Riccati regularization.

An analytical solution of (1) is obtained by following the Honorio and Jaakkola derivation [10], which bears similarities with the derivation of the *Scout*(2..) method of Witten and Tibshirani [20]. More precisely, the extrema of the objective (1) are found by solving:

$$Q^{-1} - C - \rho V^T V Q V^T V = 0 \quad (5)$$

Following [10], Q is obtained as a matrix geometric mean:

$$V Q V^T = P = \left(\frac{1}{\rho} D \right) \# \left(D^{-1} + \frac{1}{4\rho} D \right) - \frac{1}{2\rho} D \quad (6)$$

$$\text{where } D = V^{-T} C V^{-1} = \left(\frac{1}{\sqrt{T}} V^{-T} X \right) \left(\frac{1}{\sqrt{T}} V^{-T} X \right)^T \quad (7)$$

According to the properties of geometric means of matrices $\#$, recalled in [13], the eigenvectors of D are also eigenvectors of P , and an eigenvalue p of P depends only on the eigenvalue d of D associated with the same eigenvector:

$$p(d) = \sqrt{\frac{d}{\rho} \left(\frac{1}{d} + \frac{d}{4\rho} \right)} - \frac{d}{2\rho}. \quad (8)$$

This property leads to the efficient computation of Q presented in Fig. 2. The computation of Hadamard-Riccati regularized precision matrices, for which matrices V are diagonal, is almost as fast as the original [10]. We found that random projections [9] are of prominent interest for the computation of Riccati regularized inverses. First, because they accelerate all the computations by reducing matrices dimensions. Second, because they provide a direct control of the rank of the rank-deficient part of the precision matrix Q . Lastly, because they reduce precision matrices noise by truncating the small singular values of the covariance matrices C . This protective effect is illustrated in Sect. 4.1.

2.3 Tononi-Sporns-Edelman Entropy

Tononi, Sporns and Edelman introduced a measure of functional integration derived from precision matrices [16]. This measure will be referred as Tononi-Sporns-Edelman entropy (TSEe) in the sequel. Under the standard assumption that functional time series are Gaussian, TSEe measures the Gaussian entropy of a functional networks $\mathcal{N}$ as follows:

$$TSEe(Q, \mathcal{N}) = \frac{1}{2} \log \det([Q]_{\mathcal{N}}) \quad (9)$$

where $[Q]_{\mathcal{N}}$ denotes the restriction of the precision matrix Q to the nodes in the network $\mathcal{N}$. TSEe is a standard measure of functional integration and has already been used for measuring the integration of the networks derived from sparse precision matrices [19]. In this work, TSEe was measured for Riccati regularized precision matrices as well. When the

penalty R is constant over the network $\mathcal{N}$, the structure of the Riccati precision matrix can be exploited for accelerating TSEe computation. A constant penalty R over $\mathcal{N}$ corresponds indeed to a matrix V proportional to the identity for the nodes in $\mathcal{N}$:

$$[V]_{\mathcal{N}} = \alpha I \quad (10)$$

Under this assumption and following the notations of Fig. 2:

$$[Q]_{\mathcal{N}} = [W]_{\mathcal{N},:} \Omega [W]_{\mathcal{N},:}^T + \frac{1}{\alpha^2 \sqrt{\rho}} I \quad (11)$$

where $[W]_{\mathcal{N},:}$ denotes the restriction of the rows of W to the nodes in the network $\mathcal{N}$. Because Ω is a diagonal matrix with strictly positive diagonal components the following matrix can be computed in a single pass over $[W]_{\mathcal{N},:}$:

$$\bar{W} = [W]_{\mathcal{N},:} \Omega^{1/2} \quad (12)$$

Let s_i denote one of the m singular values of $\bar{W}$ and n the number of nodes of the network $\mathcal{N}$. The left singular vector of $\bar{W}$ associated to s_i is an eigenvector of $[Q]_{\mathcal{N}}$ and the associated eigenvalue λ_i is equal to $s_i^2 + \frac{1}{\alpha^2 \sqrt{\rho}}$. The remaining eigenvectors of $[Q]_{\mathcal{N}}$ are associated with the same eigenvalue: $\frac{1}{\alpha^2 \sqrt{\rho}}$. As a result, TSEe can be computed at the cost of a single SVD, as shown in Fig. 3.

3 Applications

3.1 Robust Structural Distances

Cortical thickness (CT) is a scalar measured from structural MRI describing local cortical gray matter geometry [1,4]. The structural covariance matrix is obtained by computing CT covariance across a population, for all pairs of brain locations. Large structural covariances indicate that brain regions develop, age or suffer from a disease in similar ways across a population [1]. The inverse of a structural covariance $C_{\mathcal{S}}$, obtained for a healthy population, can be used for defining a Mahalanobis distance $d_{\mathcal{S}}$ teasing out abnormal CT maps:

$$d_{\mathcal{S}}(a, b) = \sqrt{(a-b)^T C_{\mathcal{S}}^{-1} (a-b)} \quad (13)$$

This distance is small when the difference between CT maps a and b is likely to be observed in the healthy population, whereas large distances correspond to unusual CT variations. In this work, we introduce Riccati regularized structural precision matrices. We show experimentally that regularization and random projections improve structural distance robustness, in Sect. 5.

3.2 Shared Functional Networks

An increasing effort has been dedicated to the extraction of biomarkers capturing the specificities of individual rs-fMRI scans, with the aim of developing rs-fMRI based diagnostic tools. Because these scans are strongly affected by noise and subject motion, several regularization strategies were proposed such as the introduction of population averages [19].

In this work, rather than introducing a group average precision matrix and penalizing the differences between individual scans and the group average [19], we propose to perform a joint SVD (JSVD) when computing Riccati regularized precision matrices. This JSVD forces Riccati regularized precision matrices to share their eigenvectors. As a result, scan specificities are encoded in a reduced set of values, the eigenvalues of the Riccati regularized precision matrices, which can be interpreted as scan-specific loadings. This modeling offers many advantages for investigating neurodevelopment, aging and brain diseases [5,19]. The shared eigenvectors will be referred as shared functional networks in the sequel.

3.3 Functional Network Biomarkers

TSEe is an interesting functional biomarker. However, when small brain networks are investigated TSEe might be corrupted by a noise induced by the random variation of the other components in the precision matrix. To address this issue, we suggest to penalize more the components of the precision matrix corresponding to nodes outside the network of interest. We design a simple penalty to achieve this goal by choosing for V a diagonal matrix equal to the identity when restricted to the nodes of the network and to α times the identity, $\alpha > 1$, when restricted to the other nodes. As explained in Sect. 2.2 such a penalty is an Hadamard-Riccati penalty. When α is increased, this penalty gradually isolates the network of interest from the rest of the brain. As illustrated in Sect. 4.4, this effect can improve test-retest performances for some functional networks.

4 Experimental Validation

4.1 HCP Dataset

All the experiments presented in this work were carried out with the hundred unrelated subjects of the HCP dataset [17]. For each subject, four 15 min long rs-fMRI scans of 1200 time points are available, and several maps describing the local geometry of the cortex such as cortical thickness [7]. We used the rs-fMRI scans processed with the ICA+FIX pipeline with MSMAll registration and the cortical thickness map registered in the 32k Conte69 atlas, also registered with MSMAll [7]. Rs-fMRI scans were bandpass-filtered between 0.05 and 0.1Hz by an equiripple finite impulse response filter and the first two hundred timepoints impacted by the temporal filtering were discarded. Cortical thickness maps outliers were discarded by thresholding each map independently at ± 4.4478 median absolute deviation from the median. This thresholding can be interpreted as a counterpart of the standard thresholding of Gaussian variable to three standard deviations from the mean robust to the presence of outliers [15]. All the time series and concatenated cortical thickness maps were normalized to zero mean and unit variance. The spatial dimension of the data was reduced by averaging the neuroimaging signals over Glasser et al. multi-modal parcellation [8]. The

hundred eighty time series obtained for each hemisphere were normalized again to zero mean and unit variance.

The quality of the random projections (RP) was estimated by (1) concatenating the four rs-fMRI scans of each subject, (2) measuring for each subject the proportion of the squared Frobenius norm of the signal kept by the random projections, and (3) comparing the singular values of the time series before and after random projections. The results presented in Fig. 4 demonstrate that RP behaves almost like a perfect truncated SVD (TSVD) after only three power iterations. The results also suggest that the 4000 timepoints time series can be randomly projected into a dimension 200 with negligible information loss, even without power iterations.

4.2 Robust Structural Distances

Riccati regularized structural precision matrices reliability was measured by the split sample negative log likelihood, a measure decreasing with reproducibility. More precisely, the dataset was randomly split a hundred times into two groups of fifty subjects. For each split, the CT maps of the two groups were concatenated and normalized to zero mean and unit variance separately. A precision matrix Q was computed for the first group and its negative log likelihood was measured by:

$$NLL(Q) = \langle C, Q \rangle - \log \det Q \quad (14)$$

where C is the structural covariance obtained from the second group. This test-retest procedure estimates the ability of the precision matrices learned for the first group to fit/generalize to the remaining HCP subjects. The results reported in Fig. 5(a) demonstrate that the reliability of structural precision matrices is improved by TSVD and RP and reaches an optimum at small dimension and for a moderate penalty $\rho = 0.5$. RP and TSVD results are very close, for large dimensions and large penalties. For the sake of simplicity, V was set to the identity for these experiments.

We measured the ability of our method to handle large data by computing structural precision matrices at full Conte69 32 k atlas resolution and both hemisphere simultaneously (59412 nodes total). On a standard office computer running an Intel Core i5-200 CPU 3.3 GHz and 8Gb RAM, without random projections, the Riccati precisions were obtained in 12.47 s on average (over 100 runs). A random projection to dimension seven followed by the computation of the Riccati precision required 0.28 s on average (over 100 runs) and captured CT variation modes similar to the one presented in Fig. 5(b). By comparison, sparse precision matrices are typically obtained in two hours for 20000 nodes without GPU acceleration [11].

4.3 Shared Functional Networks

The joint SVD (JSVD) method [3] was used in this work for defining shared functional networks. We compared the ability of JSVD, TSVD, and RP to robustly capture individual function by computing first Riccati regularized precision matrices for all the rs-fMRI scans of the hundred unrelated HCP subjects, for different dimensions and penalties ρ . Because

functional networks are usually described for correlations or partial correlations, we derived partial correlations from all these precision matrices. We compared the methods by measuring the average intraclass correlation coefficient (ICC) of the partial correlations. We measured if the repeated scans of a single subject were producing partial correlations more similar to each other than scans of different subjects. We measured an $ICC(C, 1)$ [14]. For the sake of simplicity, V was set to the identity for these experiments. The results of Fig. 6(a) clearly demonstrate that JSVD better captures the specificities of subjects brain function.

We checked the reliability/reproducibility of JSVD results by concatenating the first two and last two scans of each subject, computing JSVD, TSVD and RP Riccati regularized matrices for the first scans and measuring the negative log likelihood obtained with the last scans. As indicated in Fig. 6(b), we observed that JSVD matrices generalize slightly less than their TSVD and RP counterparts. These results suggest that a given population is much better described using JSVD, at the cost of a small decrease of generalizability to other populations.

4.4 Functional Network Biomarkers

TSEe measures the integration of a functional subnetwork, and can, therefore, be considered as a biomarker. We observed that when TSEe is computed for Riccati regularized precision matrices, the test-retest reproducibility of this biomarker can sometimes be improved by penalizing the precisions involving nodes not part of the subnetwork of interest. During our experiments, the visual cortex was considered as the network of interest and we compared the ICC measured for different Riccati Hadamard penalizations. As explained in Sect. 3.3, the vector v defining the Hadamard Riccati penalty was set to 1 for the nodes inside the visual cortex and α for the other nodes. The original Riccati penalty [10] corresponds to $\alpha = 1$. Figure 7(a) and (b) illustrates the effects of parameter α . For large α values, the precisions outside the visual cortex are almost discarded and the Hadamard-Riccati penalization has the same effect as a restriction of the entire analysis to the visual cortex. This effect was beneficial in terms of biomarker ICC for small penalties, and detrimental for large penalties.

5 Discussion

In this work, we present several neuroimaging applications of Riccati regularized precisions matrices. Because these precision matrices are low rank, stored efficiently, and the SVD required for their computation is fast, they can be computed at full brain resolution very efficiently, contrary to sparse precision matrices [10,19]. However, we do not believe that confronting these two approaches would be fully relevant. Sparse precision matrices elegantly capture the connectivity between brain regions, which is sparse by nature. By contrast, Riccati regularized matrices are designed for extracting the connectivity of large graphs where some homogeneity/redundancy is present, and hence suitable for a low-rank description. We could claim that the first approach captures the integration of brain regions, whereas the second exploits the segregation of brain function. For this reason, we think that a combined framework, generating precisions matrices sparse for long range connections

and low-rank for small range connections, should ideally leverage the benefits of both approaches.

Because Riccati regularized and Tikhonov regularized precision matrices are computed in a similar fashion [10], their main differences reside in the larger flexibility offered by the Riccati regularized matrices. Contrary to the Tikhonov penalization which acts only on the diagonal of the precision, the Riccati regularization penalizes all the components of the matrix, which offers more freedom for designing penalties. A comparison of the eigenvalue transformation induced by the two penalties also suggest that the information corresponding to the large covariance eigenvalues is slightly better preserved into Riccati regularized precision matrices. The possibility of merging both penalties into a larger analytic framework is an interesting open question.

The experiments presented in this paper have the potential to stimulate novel applications. For instance, similarly to Sect. 3.1, robust structural distances could be derived from the other cortical measures provided by Freesurfer [4] such as areal distortion and cortical curvature, and for HCP myelin maps obtained by combining T1 and T2 weighted MRI scans [7,17]. In addition, we emphasize that, by considering symmetric Riccati penalties only, we have restricted our investigations to optimization problems that can be solved efficiently but we have missed large families of applications. Asymmetric penalties would involve more elaborate algebraic Riccati equations and hopefully stimulate novel neuroimaging applications of control theory.

6 Conclusion

In this paper, we propose an integrated approach for the extraction of neuroimaging biomarkers. We measure the entropy of brain networks defined by computing Riccati penalized precision matrices. We demonstrate how these biomarkers can be improved by reducing data dimension via random projection. We highlight several neuroscience applications for which Riccati regularized precision matrices offer novel perspectives. These applications were all validated by processing the hundred unrelated subjects of the HCP dataset. We hope that the promising results obtained, both in terms of speed and test-retest performances and the broad range of possible theoretical refinements, will encourage further developments and additional neuroimaging applications.

Acknowledgments

Data were provided by the Human Connectome Project, WU-Minn Consortium (PI: D. Van Essen and K. Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University. This work was supported by NIH grant R01 EB022573.

References

1. Alexander-Bloch A, Giedd J, Bullmore E. Imaging structural co-variance between human brain regions. *Nat Rev Neurosci.* 2013; 14(5):322–336. [PubMed: 23531697]
2. Bullmore S. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci.* 2009; 10:186–198. [PubMed: 19190637]

3. Cardoso J, Souloumiac A. Jacobi angles for simultaneous diagonalization. *SIAM J Matrix Anal Appl.* 1996; 17:161–164.
4. Dale A, Fischl B, Sereno M. Cortical surface-based analysis. I segmentation and surface reconstruction. *Neuroimage.* 1999; 9:179–194. [PubMed: 9931268]
5. Eavani, H., Satterthwaite, TD., Gur, RE., Gur, RC., Davatzikos, C. Unsupervised learning of functional network dynamics in resting state fMRI. In: Gee, JC, Joshi, S, Pohl, KM, Wells, WM., Zöllei, L., editors. *IPMI 2013. LNCS. Vol. 7917.* Springer; Heidelberg: 2013. p. 426–437.
6. Fan J, Liao Y, Liu H. An overview of the estimation of large covariance and precision matrices. *Econom J.* 2016; 19(1):C1–C32.
7. Glasser MF, Sotiropoulos SN, Wilson JA, Coalson TS, Fischl B, Andersson JL, Xu J, Jbabdi S, Webster M, Polimeni JR, Van Essen DC, Jenkinson M. The minimal preprocessing pipelines for the human connectome project. *Neuroimage.* 2013; 80:105–124. [PubMed: 23668970]
8. Glasser M, Coalson T, Robinson E, Hacker C, Harwell J, Yacoub E, Ugurbil K, Andersson J, Beckmann C, Jenkinson M, Smith S, Van Essen D. A multi-modal parcellation of human cerebral cortex. *Nature.* 2016; 536:171–178. [PubMed: 27437579]
9. Halko N, Martinsson P, Tropp JA. Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.* 2011; 53(2):217–288.
10. Honorio, J., Jaakkola, T. Inverse covariance estimation for high-dimensional data in linear time and space: spectral methods for riccati and sparse models. *Proceedings of the 2013 Conference on Uncertainty in Artificial Intelligence*; 2013;
11. Hsieh CJ, Sustik MA, Dhillon IS, Ravikumar PK, Poldrack R. Big and quic: sparse inverse covariance estimation for a million variables. *Advances in Neural Information Processing Systems.* 2013:3165–3173.
12. Insel T, Cuthbert B, Garvey M, Heinssen R, Pine DS, Quinn K, Sanislow C, Wang P. Research domain criteria (rdoc): toward a new classification framework for research on mental disorders. *Am J Psychiatry.* 2010; 167(7):748–751. [PubMed: 20595427]
13. Lim Y. The matrix golden mean and its applications to riccati matrix equations. *SIAM J Matrix Anal Appl.* 2006; 29(1):54–66.
14. McGraw K, Wong S. Forming inferences about some intraclass correlation coefficients. *Psychol Methods.* 1996; 1(1):30–46.
15. Rousseeuw PJ, Croux C. Alternatives to the median absolute deviation. *J Am Stat Assoc.* 1993; 88(424):1273–1283.
16. Tononi G, Sporns O, Edelman G. A measure for brain complexity: relating functional segregation and integration in the nervous system. *Proc Nat Acad Sci USA (PNAS).* 1994; 91:5033–5037. [PubMed: 8197179]
17. Van Essen D, Smith S, Barch D, Behrens T, Yacoub E, Ugurbil K. WU-Minn HCP Consortium. The WU-Minn human connectome project: an overview. *NeuroImage.* 2013; 80:62–79. [PubMed: 23684880]
18. Varoquaux G, Craddock RC. Learning and comparing functional connectomes. *NeuroImage.* 2013; 80:405–415. [PubMed: 23583357]
19. Varoquaux G, Gramfort A, Poline J, Thirion B. Brain covariance selection: better individual functional connectivity models using population prior. *Advances in Neural Information Processing Systems.* 2010
20. Witten D, Tibshirani R. Covariance-regularized regression and classification for highdimensional problems. *J R Stat Soc.* 2009; 71(3):615–636.

input: data X of size $N \times T$, parameter t and number of power iteration q

1. form the $N \times t$ matrix Ω by sampling from Gaussian distribution $\mathcal{N}(0, 1)$
2. $U \leftarrow (X^T X)^q X^T \Omega$
3. get W by orthonormalizing the columns of U with Gram-Schmidt process
4. $Y \leftarrow XW$

output: projected data Y , of size $N \times t$

Fig. 1.
Random projection for dimensionality reduction [9].

input: time series X of size $N \times T$, parameter ρ and $N \times N$ invertible matrix V

1. $[U, S, Z] \leftarrow SVD\left(\frac{1}{\sqrt{T}}V^{-T}X\right)$; U left singular vectors, S contains singular values
2. form diagonal matrix $\Omega = \text{diag}(\omega_1, \dots, \omega_n)$ from singular values $s_1, \dots, s_m$:

$$\omega_i = \sqrt{\frac{1}{\rho} + \frac{s_i^4}{4\rho^2}} - \frac{s_i^2}{2\rho} - \frac{1}{\sqrt{\rho}}$$
3. $W \leftarrow V^{-1}U$
4. $Q \leftarrow W\Omega W^T + \frac{1}{\sqrt{\rho}}V^{-1}V^{-T}$

output: Riccati-penalized precision matrix Q

Fig. 2.

Computation of Riccati-penalized precision matrix Q , for input time series X , the penalization ρ and a $N \times N$ invertible matrix V .

input: network $\mathcal{N}$ containing n nodes, Riccati precision matrix $Q = W\Omega W^T + \frac{1}{\sqrt{\rho}}V^{-1}V^{-T}$ with $[V]_{\mathcal{N}} = \alpha I$, Riccati penalty ρ

1. $\bar{W} \leftarrow [W]_{\mathcal{N},:} \Omega^{1/2}$
2. $[U, S, V] \leftarrow SVD(\bar{W})$ where $S = \text{diag}(s_1, \dots, s_m)$ and $m \leq n$ necessarily
3. $TSEe(Q, \mathcal{N}) \leftarrow \frac{(m-n)\log(\alpha^4 \rho)}{4} + \frac{1}{2} \sum_{i=1}^m \log \left(s_i^2 + \frac{1}{\alpha^2 \sqrt{\rho}} \right)$

output: TSEe for the network $\mathcal{N}$

Fig. 3.
Efficient computation of the TSEe for constant Riccati penalties.

Fig. 4. RP for the functional data

Remaining spectrum for increasing number of random projections (a) without power iterations (b) three power iterations ($q = 3$) (c) For the first subject: singular values before and after random projections ($q = 3$).

Fig. 5. Structural precision

(a) Average split sample negative log-likelihood (100 repetitions) of Riccati regularized precision matrices built for the cortical thickness (CT) averaged over the 180 parcellation, with respect to the “dimension”: the number of singular values kept by the truncated SVD or by the random projection (RP). Dimension 50 corresponds to the original data, without RP or TSVD. Four different Riccati penalties ρ were tested. (b) one of the seven modes of CT variation obtained at full brain resolution for the left hemisphere, $\rho = 100.0$ and RP into seven dimensions. This map corresponds to a column of the matrix W defined in Fig. 2.

Fig. 6. Shared functional networks better capture subject specificities but generalize slightly less (a) average ICC observed for the partial correlations derived from the Riccati precision matrices of the entire dataset, for different dimensions and penalties ρ . TSVD and RP results differ only at small dimension. (b) for each subject: negative log likelihood of the precision matrices obtained for the first two scans of the subject, evaluated with the last two scans. RP and TSVD results are close. JointSVD precisions, obtained for all the subjects simultaneously, generalize slightly less. The dimension was set to 25 and ρ to 0.25.

Fig. 7. Biomarkers extracted from Hadamard-Riccati precision matrices
 (a) Riccati regularized precision matrix (b) Hadamard Riccati regularized precision matrix
 (c) Visual cortex TSE ICC w.r.t Riccati penalties p and non ROI suppression α