

Detection of Brain Functional-Connectivity Difference in Post-stroke Patients Using Group-Level Covariance Modeling

Gaël Varoquaux¹, Flore Baronnet², Andreas Kleinschmidt², Pierre Fillard¹,
and Bertrand Thirion¹

¹ Parietal project-team, INRIA Saclay-île de France

² INSERM, U562, CEA/Neurospin bât 145, 91191 Gif-Sur-Yvette, France

Abstract. Functional brain connectivity, as revealed through distant correlations in the signals measured by functional Magnetic Resonance Imaging (fMRI), is a promising source of biomarkers of brain pathologies. However, establishing and using diagnostic markers requires probabilistic inter-subject comparisons. Principled comparison of functional-connectivity structures is still a challenging issue. We give a new matrix-variate probabilistic model suitable for inter-subject comparison of functional connectivity matrices on the manifold of Symmetric Positive Definite (SPD) matrices. We show that this model leads to a new algorithm for principled comparison of connectivity coefficients between pairs of regions. We apply this model to comparing separately post-stroke patients to a group of healthy controls. We find neurologically-relevant connection differences and show that our model is more sensitive than the standard procedure. To the best of our knowledge, these results are the first report of functional connectivity differences between a single-patient and a group and thus establish an important step toward using functional connectivity as a diagnostic tool.

1 Introduction

The correlation structure of brain activity, measured via fMRI, reveals stable inter-subject networks of distant brain regions that can be the expression of cognitive function. In particular, some connectivity networks are present in the absence of stimuli. They can reveal intrinsic brain activity and are studied in the *resting-state* paradigm. These structures are of particular interest to study and diagnose brain diseases and disorders [1] as they can be used for deep probes of brain function on diminished subjects. Not only can they extract medically or cognitively relevant markers on subjects unconscious [2], or with limited cooperation [3], but they also give information on higher-level cognitive systems that are challenging to probe via medical imaging or behavioral clinical tests [4].

To use functional connectivity as a quantitative inference tool, principled probabilistic comparison of connectivity structures across subjects is needed. Unlike with stimuli-response studies used routinely in functional neuroimaging,

this comparison is challenging, as the underlying description of the signal is multivariate: each brain-activity time course is considered relative to others. Univariate group models, such as random effects or mixed effects, are in general not sound as they neglect the strong statistical dependence between parameters estimated from the data. Multivariate techniques have been successfully employed to single out outlying subjects [5], but have met little success: their results are difficult to interpret as they do not point to specific localized differences.

In this paper, we focus on the description of brain functional-connectivity using inter-regions correlation matrices. We first review the current practice in inter-subject functional covariance comparison and recall some results on the manifold of covariance matrices. Then, we introduce a probabilistic model at the group level for the different subjects' correlation matrices, and a corresponding algorithm to detect connectivity differences in a specific parametrization of the covariance matrices, as correlations are a form of covariance. We quantify on simulated data the performance of this detection. Finally, we apply the model to the individual comparison of the connectivity structure of stroke patients to a group of healthy controls, and show that it outperforms the current practice.

2 State of the Art

2.1 Problem Statement: Comparing Functional Brain Connectivity

We consider S subjects, represented by the correlation matrices between brain-activity time series extracted from n ROIs: $\{\Sigma^s \in \mathbb{R}^{n \times n}, s = 1 \dots S\}$. The challenge is to give a probabilistic description of the population of correlation matrices so as to find the significant differences between subjects or groups. The current practice in functional-connectivity studies is to compare the coefficients of the correlation matrices across subjects (see for instance [3,6]). This procedure can be expressed as a univariate additive linear model on the correlation matrix:

$$\Sigma^s = \Sigma^* + \mathbf{d}\Sigma^s \quad (1)$$

where Σ^* is a covariance matrix representative of the mean effect, or the group, and $\mathbf{d}\Sigma^s$ encode subject-specific contributions.

However, with this description it is difficult to isolate significant contributions to $\mathbf{d}\Sigma^s$. Indeed, for interpretation, some coefficients are zeroed out, eg by thresholding a test statistic, as in [3], which eventually leads to a non positive definite matrix, for which it is impossible to write a multivariate normal likelihood. As a result, the subject-variability description learned on a population cannot give probabilistic tests on new subjects.

2.2 Recent Developments on the Covariance-Matrix Manifold

The mathematical difficulty stems from the fact that the space of SPD matrices, Sym_n^+ , does not form a vector space: $\mathbf{A}, \mathbf{B} \in Sym_n^+ \not\Rightarrow \mathbf{A} - \mathbf{B} \in Sym_n^+$. The Fisher information matrix of the multivariate normal distribution can be

used to construct a metric on a parametrization of covariance matrices [7] and thus define Sym_n^+ as a Riemannian manifold that is well-suited for performing statistics on covariances [8]. Local differences on this manifold can be approximated by vectors of the tangent space: if \mathbf{B} is close enough to \mathbf{A} , the application: $\phi_{\mathbf{A}} : \mathbf{B} \rightarrow \log(\mathbf{A}^{-\frac{1}{2}}\mathbf{B}\mathbf{A}^{-\frac{1}{2}})$ maps locally the bipoint $\mathbf{A}, \mathbf{B} \in Sym_n^+ \times Sym_n^+$ to $\overrightarrow{\mathbf{A}\mathbf{B}} \in Sym_n$, the space of symmetric matrices¹. A convenient parametrization of $\overrightarrow{\mathbf{W}} \in Sym_n$ is $\text{Vec}(\mathbf{W}) = \{\sqrt{2}w_{i,j}, j < i, w_{i,i}, i = 1 \dots n\}$ that forms an orthonormal basis of the tangent space [8]. Finally, $\|\overrightarrow{\mathbf{A}\mathbf{B}}\|_{\mathbf{A}}^2 = \|\log(\mathbf{A}^{-\frac{1}{2}}\mathbf{B}\mathbf{A}^{-\frac{1}{2}})\|_2^2$ gives the intrinsic norm of $\overrightarrow{\mathbf{A}\mathbf{B}}$ on the Sym_n^+ manifold, according to the metric around point \mathbf{A} : the distance between \mathbf{A} and \mathbf{B} in the manifold.

3 Matrix-Variate Random Effects Model for Covariances

Multi-subject probability distribution for covariance matrices. Using the Riemannian parametrization of Sym_n^+ , we describe the individual correlation matrix population as a distribution of matrices scattered around a covariance matrix representative of the group, Σ^* . As this distribution must be estimated with a small number of observations compared to the feature space, we model it using the probability density function that minimizes the information with a given mean on the manifold, the generalized Gaussian distribution [8]:

$$p(\Sigma) = k(\sigma) \exp\left(-\frac{1}{2\sigma^2} \|\overrightarrow{\Sigma^* \Sigma}\|_{\Sigma^*}^2\right), \tag{2}$$

where σ encodes an isotropic variance on the manifold and k is a normalization factor. Given multiple observations of Σ corresponding to individual correlation matrices, Σ^s , the maximum likelihood estimate of Σ^* is independent of σ and given by the Fréchet mean of the observations [8], minimizing $\sum_s \|\overrightarrow{\Sigma^* \Sigma^s}\|_{\Sigma^*}^2$.

Parametrization in the tangent space. We express the individual covariance matrices as a perturbation of the group covariance matrix Σ^* :

$$\forall s = 1 \dots S, \quad \Sigma^s = \phi_{\Sigma^*}^{-1}(\mathbf{d}\Sigma^s) = \Sigma^{*\frac{1}{2}} \exp(\mathbf{d}\Sigma^s) \Sigma^{*\frac{1}{2}}, \tag{3}$$

thus, using (2),
$$p(\mathbf{d}\Sigma^s) = k'(\sigma) \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{d}\Sigma^s\|_2^2\right). \tag{4}$$

The parameters of $\text{Vec}(\mathbf{d}\Sigma^s)$ follow a normal distribution, with diagonal covariance σ , and the maximum-likelihood estimate of σ is given by $\hat{\sigma}_{MLE}^2 = \frac{1}{S} \sum_s \|\text{Vec}(\mathbf{d}\Sigma^s)\|_2^2$. The model can thus be interpreted as a random-effect model on the parametrization of $\text{Vec}(\mathbf{d}\Sigma^s)$, in the space tangent to the manifold Sym_n^+ .

¹ Note that we do not use the same definition of the mapping as in [8,7], as we are interested in mapping to Sym_n , the tangent space around \mathbf{I}_n , and not the tangent space in \mathbf{A} . It extracts a statistically independent parametrization (Eq. (3) and (4)).

Algorithm 1. Estimation of the group model

- 1: **Input:** individual time series $\mathbf{X}^1 \dots \mathbf{X}^s$.
 - 2: **Output:** estimated group covariance matrix $\widehat{\Sigma}^*$, group variance $\widehat{\sigma}$.
 - 3: **for** $s = 1$ **to** S **do**
 - 4: Compute $\widehat{\Sigma}^s \leftarrow \text{LedoitWolf}(\mathbf{X}^s)$.
 - 5: **end for**
 - 6: Compute $\widehat{\Sigma}^* \leftarrow \text{intrinsic mean}(\widehat{\Sigma}^1 \dots \widehat{\Sigma}^s)$.
 - 7: **for** $s = 1$ **to** S **do**
 - 8: Compute $\mathbf{d}\widehat{\Sigma}^s \leftarrow \widehat{\Sigma}^{*-1/2} \Sigma^s \widehat{\Sigma}^{*-1/2} - \mathbf{I}_n$.
 - 9: **end for**
 - 10: Compute $\widehat{\sigma} \leftarrow \sqrt{\frac{1}{S} \sum_s \|\text{Vec}(\mathbf{d}\widehat{\Sigma}^s)\|_2^2}$.
-

Assuming that the distribution is narrow on the manifold, $\|\mathbf{d}\Sigma^s\|_2 \ll 1$, eq. (3) can be seen as the application of the placement function to move a noise $\mathbf{d}\Sigma^s$ isotropic around \mathbf{I}_n to Σ^* (see [8], section 3.5):

$$\Sigma^s \simeq \Sigma^{*\frac{1}{2}} (\mathbf{I}_n + \mathbf{d}\Sigma^s) \Sigma^{*\frac{1}{2}}. \tag{5}$$

Model estimation from the data. We start from individual time-series of brain activity in selected regions of interest, $\mathbf{X} \in \mathbb{R}^{n \times t}$. We use the Ledoit-Wolf shrinkage covariance estimator [9] for a good bias-variance compromise when estimating correlation matrices from t time points with $n < t < n^2$. From this estimate of individual correlation matrices, we compute the intrinsic mean on Sym_n^+ using algorithm 3 of [10]. Finally, we estimate σ from the residuals of individual correlation matrices in the space tangent in Σ^* (see algorithm 1).

4 Testing Pair-Wise Correlations Statistics

The multivariate probabilistic model for correlations between regions exposed above enables us to define an average correlation matrix of a group, as well as the dispersion of the group in the covariance matrix space. Thus it can be used to test if a subject is different to a control group. However, to aid diagnosis, it is paramount to pin-point why such a subject may be different. In the tangent space, the parameters $d\Sigma_{i,j}^s$ of $\text{Vec}(\mathbf{d}\Sigma^s)$ are mutually independent. We can thus conduct univariate analysis on these parameters to test which significantly differs from the control group. However, the independence of the parameters is true only in the space tangent at the population average Σ^* , of which we only have an estimate $\widehat{\Sigma}^*$. Thus, to account for projection error, we resort to non-parametric sampling of the control population to define a null distribution for the parameters $d\Sigma_{i,j}^s$.

Specifically, we are interested in testing if a difference observed for a subject in one of the $d\Sigma_{i,j}^s$ can be explained by variation of the control population. As the control population is typically small, we generate the null distribution by leave

Algorithm 2. Coefficient-level tests

- 1: **Input:** individual time series for controls $\mathbf{X}^1 \dots \mathbf{X}^s$ and a patient \mathbf{X}^k , p-value p , number of bootstraps, m .
 - 2: **Output:** Pair-wise p-values $p_{i,j}$ controlling for the difference in $d\Sigma_{i,j}$ between the patient and the control group.
 - 3: Initialize $P_{i,j}^o \leftarrow$ empty lists, for $i, j \in \{1 \dots n\}, j < i$.
 - 4: **for** 1 **to** m **do**
 - 5: Choose a surrogate patient $\tilde{s} \in 1 \dots S$.
 - 6: Choose a subset \tilde{S} of $\{1 \dots S\} \setminus \tilde{s}$ of S surrogate controls.
 - 7: Compute $\widetilde{\Sigma}^*$ and $\widetilde{\mathbf{d}\Sigma}^s$ for $s \in \tilde{S}$ using algorithm 1 on the surrogate controls.
 - 8: Compute $\widetilde{\mathbf{d}\Sigma}^s$ for $s = \tilde{s}$, using $\widetilde{\Sigma}^*$ and eqn 5.
 - 9: For all i, j , append to $P_{i,j}^o$ the T test comparing $\widetilde{\mathbf{d}\Sigma}_{i,j}^s$ for $s \in \tilde{S}$ and for $s = \tilde{s}$.
 - 10: **end for**
 - 11: Compute $\widehat{\Sigma}^*$ and $\widehat{\mathbf{d}\Sigma}^s$ for $s \in \tilde{S}$ using algorithm 1 on the complete control group.
 - 12: Compute $\widehat{\mathbf{d}\Sigma}^k$, using $\widehat{\Sigma}^*$ and eqn 5.
 - 13: For all i, j , compute $t_{i,j}$ the T test comparing $\widehat{\mathbf{d}\Sigma}_{i,j}^s$ for $s \in \tilde{S}$ and for $s = \tilde{s}$.
 - 14: $p_{i,j} = 1 - \text{quantile}(t_{i,j} \text{ in } P_{i,j}^o)$
-

one out: for each control subject, we generate surrogate control populations \tilde{S} by bootstrap from the other control subjects and estimate the corresponding average covariance $\widetilde{\Sigma}^*$. We use $\widetilde{\Sigma}^*$ to project all the individual correlation matrices, including the left out subject, to compute $\widetilde{\mathbf{d}\Sigma}_{i,j}^s$, and we do a one sample T test of the difference between $\widetilde{\mathbf{d}\Sigma}_{i,j}^s$ for the left out subject with regards to the resampled control group \tilde{S} . We record the values of this T test as an estimate of the null distribution $P_{i,j}^o$ of the T test on the corresponding coefficient between the controls and a patient. Finally, we estimate the average covariance $\widehat{\Sigma}^*$ for the complete group of controls and, for each k patient to investigate, we perform a T test of the difference between $\widehat{\mathbf{d}\Sigma}_{i,j}^s$ for the patient and the control group. We use $P_{i,j}^o$ to associate a p-value to each coefficient per subject. We correct for multiple comparisons using Bonferroni correction: for each patient, $\frac{1}{2}n(n-1)$ tests are performed. The procedure is detailed in algorithm 2.

5 Algorithm Evaluation on Simulated Data

Algorithm 2 relies on approximations of the exact problem for coefficient-level detection of differences. In order to quantify the performance of this detection, we study Receiver Operator Characteristic (ROC) on simulated data: we draw a population of control covariances using eq. (5) with the parameters of $\text{Vec}(\mathbf{d}\Sigma)$ normally distributed with deviation σ . For simulated patients, we add differences of amplitude $d\Sigma$ to a few coefficients (~ 20) of this variability noise. For Σ^* and σ , we use the parameters estimated on real data (section 6). We investigate the performance of algorithm 2 to recover these differences for a variety of parameters. We observe good recovery for $d\Sigma > \sigma$ (Fig 1), and find that the comparison in the tangent space (eq. 5) outperforms a comparison in $\mathbb{R}^{n \times n}$ (eq. 1).

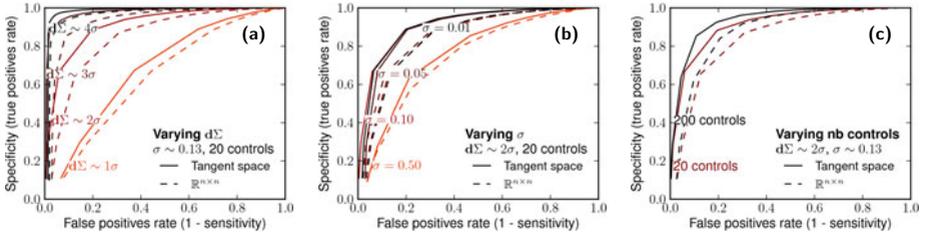


Fig. 1. ROC curves on data simulated according to the variability model given by eqn 5. **(a)** for different values of patient differences $d\Sigma$. **(b)** for different values of control variability σ . **(c)** for different number of controls.

6 Application to Post-stroke Connectivity Reorganization

Standard clinical scores, such as the NIHSS, as well as fMRI studies can be used to assess the consequences of cerebral strokes, but they test specific cognitive functions and have little sensitivity to higher-order cognitive malfunctions. Resting-state functional-connectivity is thus a valuable tool to study post-stroke reorganization. We apply our model to stroke patients.

Resting-state fMRI dataset. After giving informed consent, ischemic-stroke patients, as well as age-matched healthy controls, underwent MRI scanning. Subjects with existing neurology, psychiatry, or vascular pathologies were excluded from the study. 10 patients and 20 controls were scanned during a resting-state task: subjects were given no other task than to stay awake but keep their eyes closed. 2 sessions of 10 minutes of fMRI data were acquired on a Siemens 3T Trio scanner (245 EPI volumes, TR=2.46s, 41 slices interlaced, isotropic 3 mm voxels). After slice-timing, motion correction, and inter-subject normalization using SPM5, 33 ROIs were defined in the main resting-state networks by intersecting a segmentation of the gray matter with correlation maps from seeds selected from the literature. For each subject, the BOLD time series corresponding to the regions were extracted and orthogonalized with respect to confound time series: time courses of the white matter and the cerebro-spinal fluid, and movement regressors estimated by the motion-correction algorithm. Covariance modeling was performed on the resulting 33 time series.

Separating patients from controls with the matrix-variate covariance model. To measure the discriminative power of the matrix-variate model introduced in section 3, we test the likelihood of patient data in a model learned on controls. Specifically, we evaluate by *leave one out* the likelihood of each control in the model learned on the other controls. We compare this value to the average likelihood of patients in the 20 models obtained by *leave one out*. We perform this comparison both using the group model isotropic on the tangent space (eq. 5), and the group model isotropic in $\mathbb{R}^{n \times n}$ (eq. 1). We find that the model in the tangent space separates better patients from controls (Fig 2).

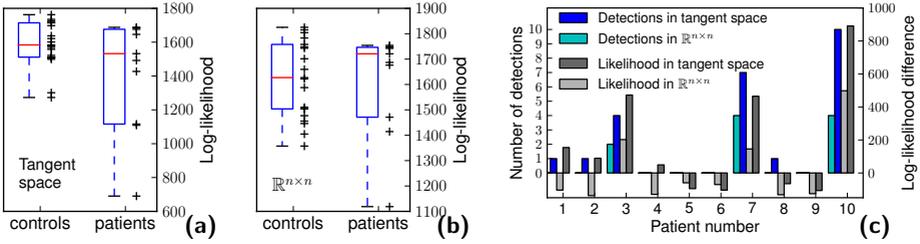


Fig. 2. (a) Likelihood of the controls and the patients in the model parametrized in the tangent space. (b) Likelihood of the controls and the patients in the model parametrized in $\mathbb{R}^{n \times n}$. (c) Number of coefficients detected as significantly different from the control group per patient, for the model parametrized in the tangent space, as well as in $\mathbb{R}^{n \times n}$.

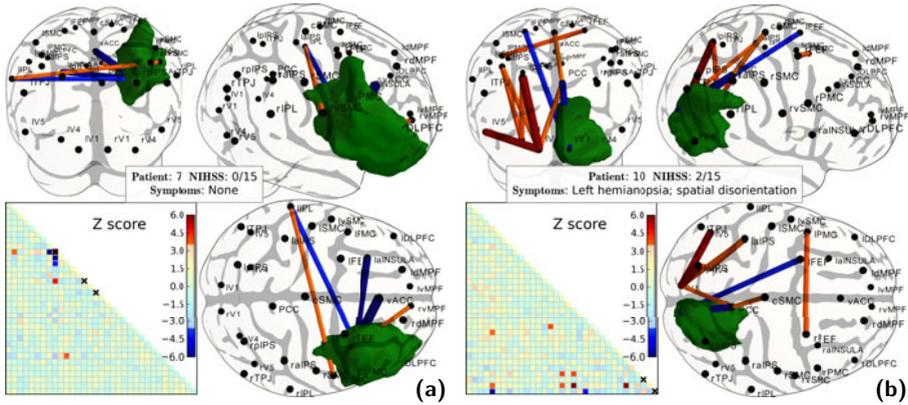


Fig. 3. Significant differences on two patients ($p < 0.05$ uncorrected), represented as connections between regions: increased connectivity appears in red, and decreased in blue. The lesion, manually segmented from anatomical images, is represented in green. ROIs fully covered by a lesion are marked with a black cross on the correlation matrix.

Detected connection differences. We apply algorithm 2 to detect the significant coefficient-level differences for each subject. We compare to a similar univariate procedure applied to the parametrization in $\mathbb{R}^{n \times n}$ given by eq. (1), rather than the tangent space. We find that coefficient-level analysis detects more differences between ROI pairs when applied on the tangent-space parametrization (Fig 2c).

7 Discussion

Interpretation of the tangent space. Projecting on the space tangent to the group mean corresponds to applying a whitening matrix $\Sigma^*^{-\frac{1}{2}}$ learned on the group (eq. 5) that converts the Gaussian process described by the group covariance

to an independent and identically distributed (iid) process. In other words, the *coloring* of the time series common to the group is canceled out to compare subjects on iid coefficients on the correlation matrix.

Probing neurological processes. For certain subjects, both procedures fail to detect a single connection that makes a significant difference. Indeed, the variability of resting-state activity in the control group induces some variability in the projection to the tangent space. For patients with small lesions, this variability is larger than the univariate differences. On the other hand, for patients with important lesions, the functional connectivity analysis reveals profound differences in the correlation structure that reflect functional reorganization. Some express a direct consequence of the lesion, for example when the gray matter in one of two ROIs has been damaged by the lesion, as can be seen on Fig 3a. Others reflect functional reorganization. For instance, patient 10 has a right visual cortex damaged by a focal lesion, but the analysis shows increased connectivity in his left visual cortex (Fig 3b). Functional connectivity analysis thus reveals modifications that go beyond the direct anatomical consequences of the lesion.

8 Conclusion

We have presented a matrix-variate probabilistic model for covariances, and have shown that it can be expressed as a random effect model on a particular parametrization of the covariance matrix. The ability to draw conclusions on the connectivity between pairs of regions is important because it is a natural representation of the problem. We applied this model to the comparison of functional brain connectivity between subjects. We were able to detect significant differences in functional connectivity between a single stroke patient and 20 controls. A controlled detection of network-wide functional-connectivity differences between subjects opens the door to new markers of brain diseases as well as new insights in neuroscience, as functional connectivity can probe phenomena that are challenging to access via stimuli-driven studies.

References

1. Greicius, M.: Resting-state functional connectivity in neuropsychiatric disorders. *Current opinion in neurology* 21, 424 (2008)
2. Vanhaudenhuyse, A., et al.: Default network connectivity reflects level of consciousness in non-communicative brain-damaged patients. *Brain* 133, 161 (2010)
3. Fair, D., et al.: Development of distinct control networks through segregation and integration. *PNAS* 104, 13507 (2007)
4. Uddin, L., Menon, V.: The anterior insula in autism: Under-connected and under-examined. *Neuroscience and Biobehavioral Reviews* 33, 1198–1203 (2009)
5. Sato, J., et al.: An fMRI normative database for connectivity networks using one-class support vector machines. *Human brain mapping* 30, 1068 (2008)
6. Rocca, M., et al.: Altered functional and structural connectivities in patients with MS: a 3-T study. *Neurology* 69, 2136 (2007)

7. Lenglet, C., et al.: Statistics on the manifold of multivariate normal distributions: Theory and application to diffusion tensor MRI processing. *Journal of Mathematical Imaging and Vision* 25, 423–444 (2006)
8. Pennec, X., Fillard, P., Ayache, N.: A Riemannian framework for tensor computing. *International Journal of Computer Vision* 66, 41–66 (2006)
9. Ledoit, O., Wolf, M.: A well-conditioned estimator for large-dimensional covariance matrices. *J. Multivar. Anal.* 88, 365–411 (2004)
10. Fletcher, P., Joshi, S.: Riemannian geometry for the statistical analysis of diffusion tensor data. *Signal Processing* 87, 250–262 (2007)