A Riemannian Framework for Longitudinal Analysis of Resting-State Functional Connectivity

Qingyu Zhao[1], Dongjin Kwon[1,2], Kilian M. Pohl[2]

[1] Department of Psychiatry and Behavioral Sciences, Stanford University, Stanford, USA
[2] Center of Health Sciences, SRI International, Menlo Park, USA

Abstract. Even though the number of longitudinal resting-state-fMRI studies is increasing, accurately characterizing the changes in functional connectivity across visits is a largely unexplored topic. To improve characterization, we design a Riemannian framework that represents the functional connectivity pattern of a subject at a visit as a point on a Riemannian manifold. Geodesic regression across the 'sample' points of a subject on that manifold then defines the longitudinal trajectory of their connectivity pattern. To identify group differences specific to regions of interest (ROI), we map the resulting trajectories of all subjects to a common tangent space via the Lie group action. We account for the uncertainty in choosing the common tangent space by proposing a test procedure based on the theory of latent p-values. Unlike existing methods, our proposed approach identifies sex differences across 246 subjects, each of them being characterized by three rs-fMRI scans.

1 Introduction

Longitudinal resting-state(rs)-fMRI studies, in which participants are scanned at multiple visits, have been increasingly used for investigating functional connectivity changes and development in human brains [1,2]. However, current methods for rs-fMRI group analysis are mostly designed for cross-sectional studies [3,4]. In this paper, we propose a framework for performing group analysis on longitudinal rs-fMRI data.

Cross-sectional studies often encode the functional connectivity of a subject as an $n \times n$ covariance matrix $\boldsymbol{C}$ of BOLD (blood-oxygen-level dependent) time courses associated with n ROIs. To identify differences in functional connectivity between two groups, a univariate group test is typically applied to each connection (each element in the upper triangle of $\boldsymbol{C}$). One problem in this practice is that the univariate tests neglect the strong statistical dependence among matrix elements: a covariance matrix $\boldsymbol{C}$ is confined by the positive-definite constraint $\boldsymbol{C} = \boldsymbol{C}^{\mathbf{T}}$, $\mathbf{x}\boldsymbol{C}\mathbf{x}^{\mathbf{T}} > 0$ for all non-zero $\mathbf{x} \in \mathbb{R}^n$. One way to alleviate this problem is to leverage the fact that covariance matrices form a Riemannian manifold [5]. Previous rs-fMRI connectivity studies have reported improved detection sensitivity by performing group analysis directly on that manifold [4]. Based on

this observation, we design a Riemannian framework to analyze the change in connectivity patterns captured by longitudinal studies.

In a longitudinal study, each subject is characterized by a series of covariance matrices representing connectivity patterns at multiple visits. Motivated by previous works [6], our longitudinal framework is composed of two parts: (a) fitting a longitudinal trajectory on the covariance matrices of each subject via geodesic regression on the manifold [7] (Section 2); and (b) comparing subject-specific trajectories across groups (Section 3). The challenge in (b) is that trajectories of different subjects are not directly comparable as they are essentially tangent vectors defined in different tangent spaces. Several methods have been explored to handle this problem, either by directly performing group analysis on the 'tangent bundle space' [6] or by designing mixed-effect models for manifolds [8]. These methods, however, consider the object of interest (in our case a covariance matrix) as a single manifold-valued variable, such that group difference can only be identified with respect to the entire matrix (the whole brain connectivity) instead of each matrix element (connectivity between two ROIs). In order to enable ROI-specific analysis, we map all tangent vectors to a common tangent space of a template point on the manifold. This enables univariate testing to each matrix element across all mapped tangent vectors.

In addition, we define the mapping function based on the Lie group action and briefly discuss its favorable properties to the popular parallel transport mechanism [9,10]. We define the common tangent space via the the identity matrix [9] and the Fréchet mean [4], and we argue that the latter is preferred in the context of connectivity analysis. Finally, our group analysis accounts for the uncertainty in estimating the template via a robust test procedure based on the theory of latent p-values. We finally validate our proposed Riemannian framework using both synthetic and real rs-fMRI datasets.

2 Computing Subject-Specific Trajectory

Recall that the space of $n \times n$ covariance matrices forms a Riemannian manifold $\mathcal{M}$ [5]. Let $\boldsymbol{A}$ be a point on $\mathcal{M}$, $T_{\boldsymbol{A}}\mathcal{M}$ the tangent space at $\boldsymbol{A}$, and $\boldsymbol{X} \in T_{\boldsymbol{A}}\mathcal{M}$ a tangent vector. There is a unique geodesic curve γ (a locally length-minimizing curve on the manifold) with $\gamma(0) = \boldsymbol{A}$ and $\gamma(0)' = \boldsymbol{X}$. The analytical equation of a geodesic is defined by the *exponential map*,

$$\gamma(t) = Exp_{\boldsymbol{A}}(t\boldsymbol{X}) := \boldsymbol{A}^{\frac{1}{2}} expm(\boldsymbol{A}^{-\frac{1}{2}}(t\boldsymbol{X})\boldsymbol{A}^{-\frac{1}{2}})\boldsymbol{A}^{\frac{1}{2}}, \tag{1}$$

where *expm* is the matrix exponential operator. In other words, the exponential map $Exp_{\boldsymbol{A}}(\boldsymbol{X})$ at an initial point $\boldsymbol{A}$ projects a tangent vector $\boldsymbol{X}$ to a point on the manifold at $\gamma(1)$ along the geodesic γ defined by $(\boldsymbol{A}, \boldsymbol{X})$. The inverse mapping of $Exp_{\boldsymbol{A}}(\boldsymbol{X})$ is called the *log map*. It projects a point $\boldsymbol{B} \in \mathcal{M}$ back to a tangent vector at $\boldsymbol{A}$ via $Log_{\boldsymbol{A}}(\boldsymbol{B}) := \boldsymbol{A}^{\frac{1}{2}} logm(\boldsymbol{A}^{-\frac{1}{2}}\boldsymbol{B}\boldsymbol{A}^{-\frac{1}{2}})\boldsymbol{A}^{\frac{1}{2}}$.

Now let us consider M covariance matrices of a subject $\{\boldsymbol{C}^1, ..., \boldsymbol{C}^M\}$ measured at M visits. Let $\{t^1, ..., t^M\}$ be the time associated with those visits. Without loss of generality, we translate $\{t^i\}$ such that $t^1 = 0$. We then use *geodesic*

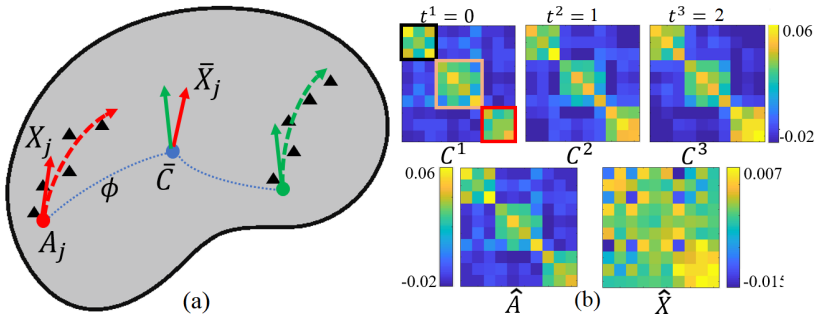

Fig. 1. (a) A Riemannian framework for longitudinal connectivity analysis: the j^{th} subject's trajectory (red dashed curve) is fitted via geodesic regression on subject-specific covariance matrices (triangles on the left). All resulting subject-specific tangent vectors are mapped via function ϕ to the tangent space of $\bar{\boldsymbol{C}}$ to perform group analysis. (b) Top: 3 covariance matrices $\{\boldsymbol{C}^1, \boldsymbol{C}^2, \boldsymbol{C}^3\}$ of a subject with $t = 0, 1, 2$. Bottom: The optimal geodesic $(\hat{\boldsymbol{A}}, \hat{\boldsymbol{X}})$ derived from geodesic regression.

regression [7] to characterize the change of $\{\boldsymbol{C}^i\}$ over time (the relationship between t^i and $\boldsymbol{C}^i$). Specifically, to find a geodesic curve $(\hat{\boldsymbol{A}}, \hat{\boldsymbol{X}})$ that optimally fits the data (Fig. 1), we minimize the following objective function:

$$(\hat{\boldsymbol{A}}, \hat{\boldsymbol{X}}) = \arg\min_{\boldsymbol{A}, \boldsymbol{X}} \sum_{i=1}^{M} d(Exp_{\boldsymbol{A}}(t^i \boldsymbol{X}), \boldsymbol{C}^i)^2, \tag{2}$$

where $d(\boldsymbol{A}, \boldsymbol{B})$ measures the *geodesic distance* between $\boldsymbol{A}, \boldsymbol{B} \in \mathcal{M}$ via the Riemannian metric $d(\boldsymbol{A}, \boldsymbol{B}) := ||Log_{\boldsymbol{A}}(\boldsymbol{B})||_{\boldsymbol{A}} := \sqrt{\text{tr}(logm(\boldsymbol{A}^{-\frac{1}{2}} \boldsymbol{B} \boldsymbol{A}^{-\frac{1}{2}}))}$.

In particular, the initial point $\hat{\boldsymbol{A}}$ characterizes the connectivity pattern at baseline. The tangent vector $\hat{\boldsymbol{X}}$ characterizes the longitudinal trajectory of that subject's connectivity. Note that the tangent space $T_{\hat{\boldsymbol{A}}}\mathcal{M}$ is the vector space of $n \times n$ symmetric matrices. Therefore, unlike $\hat{\boldsymbol{A}}$, the matrix elements in the upper triangle of $\hat{\boldsymbol{X}}$ are mutually independent for univariate tests.

3 Group Analysis for Trajectories

Now let $\{(\boldsymbol{A}_j, \boldsymbol{X}_j) | j = 1, ..., N\}$ represent the geodesics of N different subjects. While cross-sectional analysis is interested in analyzing baseline matrices $\{\boldsymbol{A}_j\}$ [4], here we aim to identify group differences in $\{\boldsymbol{X}_j\}$, *i.e.*, the difference across subjects in their longitudinal changes. As mentioned, these tangent vectors are defined in different tangent spaces (e.g. the red and green vector in Fig. 1a), therefore not directly comparable. To handle this problem, we map all tangent vectors $\{\boldsymbol{X}_j | j = 1, ..., N\}$ to a common tangent space $T_{\bar{\boldsymbol{C}}}\mathcal{M}$. This step essentially requires the definition of (a) a mapping function $\phi_{\boldsymbol{A}_j \to \bar{\boldsymbol{C}}}(\boldsymbol{X}_j) = \bar{\boldsymbol{X}}_j$ and (b) a template point $\bar{\boldsymbol{C}}$. Finally, we estimate the p-values of univariate testing for all $n(n+1)/2$ matrix elements across $\{\bar{\boldsymbol{X}}_j | j = 1, ..., N\}$.

Choosing the Mapping Function $\boldsymbol{\phi}$. A popular choice of ϕ is the *parallel transport* [9,10], which transports tangent vectors on a manifold such that they stay parallel with respect to the affine connection. Despite its appealing geometrical meaning, the major drawback of parallel transport is its path-dependency: transporting a tangent vector along two different curves with the same start and end point generally results in two different 'copies' of the vector. This phenomenon can lead to ambiguity in choosing the template for group analysis. To show this effect, let $\phi^p_{\boldsymbol{A}\to\boldsymbol{B}}(\boldsymbol{X})$ denote the parallel transport of $\boldsymbol{X} \in T_{\boldsymbol{A}}\mathcal{M}$ along the geodesic from $\boldsymbol{A}$ to $\boldsymbol{B}$. Let $\boldsymbol{X}_1 \in T_{\boldsymbol{A}_1}\mathcal{M}$ and $\boldsymbol{X}_2 \in T_{\boldsymbol{A}_2}\mathcal{M}$ be two subject-specific tangent vectors. We further assume that they are equivalent when transported to $\bar{\boldsymbol{C}}$, *i.e.*, $\bar{\boldsymbol{X}}_1 = \phi^p_{\boldsymbol{A}_1\to\bar{\boldsymbol{C}}}(\boldsymbol{X}_1) = \phi^p_{\boldsymbol{A}_2\to\bar{\boldsymbol{C}}}(\boldsymbol{X}_2) = \bar{\boldsymbol{X}}_2$. Now if we perturb $\bar{\boldsymbol{C}}$ by ε via $\bar{\boldsymbol{C}}^* = Exp_{\bar{\boldsymbol{C}}}(\varepsilon)$ and transport $\boldsymbol{X}_1, \boldsymbol{X}_2$ to $\bar{\boldsymbol{C}}^*$, in general,

$$\bar{\boldsymbol{X}}^*_1 = \phi^p_{\boldsymbol{A}_1\to\bar{\boldsymbol{C}}^*}(\phi^p_{\bar{\boldsymbol{C}}\to\boldsymbol{A}_1}(\bar{\boldsymbol{X}}_1)) \neq \phi^p_{\boldsymbol{A}_2\to\bar{\boldsymbol{C}}^*}(\phi^p_{\bar{\boldsymbol{C}}\to\boldsymbol{A}_2}(\bar{\boldsymbol{X}}_2)) = \bar{\boldsymbol{X}}^*_2 \tag{3}$$

as ϕ^p is path-dependent. Contradictory to our previous assumption that the two subject-specific tangents are equivalent, Eq. 3 reveals that they are different at $\bar{\boldsymbol{C}}^*$. In other words, (in)equality relationships among $\{\boldsymbol{X}_j\}$ are variant to the template selection, which can lead to serious ambiguity in the comparison of subject-specific trajectories.

To resolve this problem, we exploit the fact that $\mathcal{M}$ is equipped with an affine-invariant Riemannian metric [5]. To be specific, let GL_n denote the Lie group of all $n \times n$ invertible matrices. This group acts on $\mathcal{M}$ via a smooth mapping function $\psi : GL_n \times \mathcal{M} \Rightarrow \mathcal{M}$, $\psi_{\boldsymbol{G}}(\boldsymbol{A}) := \boldsymbol{G}\boldsymbol{A}\boldsymbol{G}^{\mathbf{T}} = \boldsymbol{B}$, where $\boldsymbol{G} \in GL_n$ and $\boldsymbol{A}, \boldsymbol{B} \in \mathcal{M}$. This group action can be naturally extended to tangent vectors via its derivative map $d\psi_{\boldsymbol{G}}(\boldsymbol{X}) := \boldsymbol{G}\boldsymbol{X}\boldsymbol{G}^{\mathbf{T}} = \boldsymbol{Y}$, where $\boldsymbol{X} \in T_{\boldsymbol{A}}\mathcal{M}$ and $\boldsymbol{Y} \in T_{\boldsymbol{B}}\mathcal{M}$. In other words, $d\psi_{\boldsymbol{G}}$ achieves the mapping of tangent vectors across different tangent spaces based on the aforementioned smooth group action. With this construction, we propose the following mapping function

$$\phi^g_{\boldsymbol{A}\to\boldsymbol{B}}(\boldsymbol{X}) = d\psi_{\boldsymbol{G}}(\boldsymbol{X}),\ \boldsymbol{G} = \boldsymbol{B}^{\frac{1}{2}}\boldsymbol{A}^{-\frac{1}{2}}. \tag{4}$$

Since the group action ψ is transitive [5], so is ϕ^g, *i.e.*, $\phi^g_{\boldsymbol{B}\to\boldsymbol{C}}(\phi^g_{\boldsymbol{A}\to\boldsymbol{B}}(\boldsymbol{X})) = \phi^g_{\boldsymbol{A}\to\boldsymbol{C}}(\boldsymbol{X})$. This property avoids the path-dependent assumption as required in parallel transport, so that (in)equality relationships among $\{\bar{\boldsymbol{X}}_j\}$ are invariant to the choice of the template.

Choosing the Template $\bar{\boldsymbol{C}}$. In the context of connectivity analysis, we argue that the Fréchet mean [4] is a more appropriate template compared to the identity matrix $\mathbf{I}$ [9]. Recall that $\boldsymbol{X}(u,v)$, the $(u,v)^{th}$ element in a subject-specific tangent vector $\boldsymbol{X}$, encodes the longitudinal information about connectivity between the u^{th} and v^{th} ROI. Since $\phi^g_{\boldsymbol{A}\to\bar{\boldsymbol{C}}}$ (or $\phi^p_{\boldsymbol{A}\to\bar{\boldsymbol{C}}}$) is a general linear transformation, $\bar{\boldsymbol{X}}(u,v)$ is a linear combination of $\boldsymbol{X}(u,v)$ and other matrix elements. Consequently, $\bar{\boldsymbol{X}}(u,v)$ no longer precisely relates the two ROIs. This reveals one critical trade-off: only $\boldsymbol{X}$ models true subject-specific ROI information, whereas only $\bar{\boldsymbol{X}}$ is geometrically comparable with other tangent vectors. To alleviate this issue, we realize that when $\bar{\boldsymbol{C}}$ is close to $\boldsymbol{A}$, $\bar{\boldsymbol{C}}^{\frac{1}{2}}\boldsymbol{A}^{-\frac{1}{2}}$ is close to

Algorithm 1 Latent p-value map

1: Resample $\{\boldsymbol{A}_j | j = 1, ..., N\}$ 500 times to compute the empirical bootstrap distribution of the Fréchet mean $\{\bar{\boldsymbol{C}}^1, ..., \bar{\boldsymbol{C}}^{500}\}$.
2: For each bootstrapped $\bar{\boldsymbol{C}}^i$, map $\{\boldsymbol{X}_j | j = 1, ..., N\}$ to it to achieve $\{\bar{\boldsymbol{X}}_j^i | j = 1, ..., N\}$. Apply univariate tests to each matrix element to get a p-value map $\boldsymbol{P}^i$.
3: The final latent p-value map $\boldsymbol{P} = \frac{1}{500} \sum_{i=1}^{500} \boldsymbol{P}^i$.

an identity transformation ($\bar{\boldsymbol{X}}(u, v) \approx \boldsymbol{X}(u, v)$). This motivates us to choose $\bar{\boldsymbol{C}}$ as the Fréchet mean, because it is the 'closest' matrix to all $\{\boldsymbol{A}_j\}$ so that true ROI information can be optimally preserved for all subjects.

A Robust Test Procedure. In the end, a univariate test is applied to each connection (upper triangular matrix element) across $\{\bar{\boldsymbol{X}}_j\}$. However, with different choices of $\bar{\boldsymbol{C}}$, both the null distribution and observed values of $\{\bar{\boldsymbol{X}}_j(u, v)\}$ vary accordingly to the associated $\phi^g_{\boldsymbol{A} \to \bar{\boldsymbol{C}}}$. Therefore, the uncertainty in estimating the Fréchet mean can lead to unstable p-values. To solve this problem, we resort to the theory of latent p-values [11]. Recall that the p-value is defined as $p = \text{pr}(\bar{\mathcal{X}}_0 > \bar{\boldsymbol{X}} | \mathcal{H}_0)$, *i.e.*, the probability of obtaining a result larger (right-tailed) than the observed $\bar{\boldsymbol{X}}$ under the null hypothesis $\mathcal{H}_0$. Recently, statisticians also interpret p-values as random variables [11]. Specifically, the latent p-value considers both the strength of the evidence against the null hypothesis, and the uncertainty in the evidence and null distribution. Here we regard both $\bar{\mathcal{X}}_0$ and $\bar{\boldsymbol{X}}$ as unobservable latent variables. As described in [11], the latent p-value can be approximated by generating Monte-Carlo realizations of $\bar{\mathcal{X}}_0$. We then let $\bar{\mathcal{X}}_0$ and $\bar{\boldsymbol{X}}$ be dependent on a latent template $\bar{\boldsymbol{C}}$, and we marginalize $\bar{\boldsymbol{C}}$ by $p = \int_{\bar{\boldsymbol{C}}} \text{pr}(\bar{\mathcal{X}}_0(\bar{\boldsymbol{C}}) > \bar{\boldsymbol{X}}(\bar{\boldsymbol{C}}) | \bar{\boldsymbol{C}}, \mathcal{H}_0) \text{pr}(\bar{\boldsymbol{C}})$. Finally, we can approximate this integration by sampling $\bar{\boldsymbol{C}}$ via bootstrapping. Intuitively, instead of regarding $\bar{\boldsymbol{X}}$ and $\bar{\mathcal{X}}_0$ as deterministic variables, we perform multiple test procedures based on templates sampled from $\text{pr}(\bar{\boldsymbol{C}})$ (Algorithm 1).

4 Experiments

Synthetic Data. In this experiment, we simulated covariance matrices for 2 groups of subjects. Each group has $N = 20$ subjects, and each subject has $M = 3$ covariance matrices. We added longitudinal changes to the 3 matrices of each subject in *Group B*. We validated our framework based on the accuracy in identifying those effects between the two groups.

We started from simulating BOLD time series using SimTB [12]. We randomized simulation parameters for each subject and derived covariance matrices from simulated BOLD signals. To be specific, we defined $n = 10$ ROIs and randomly grouped them into 3 independent networks (Fig. 1b). Only ROIs of the same network activated simultaneously, thus having non-zero covariances. The activation amplitude of each ROI was sampled from the standard Gaussian $\mathcal{N}(0, 1)$. The unique activation (functional noise) probability of each ROI was $u = 0.35$. We added Riccan noise (imaging noise) with a Contrast-to-Noise Ratio (CNR)

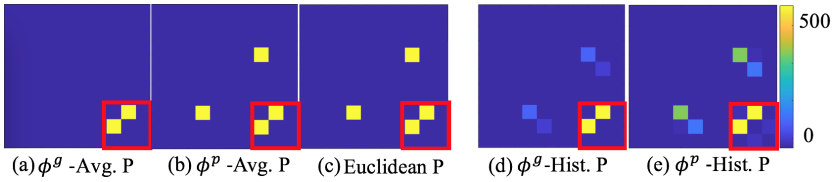

Fig. 2. (a)(b)(c) Group differences (yellow) identified by the three methods. (d)(e) Histogram of identified elements counted using the 500 p-value maps.

of 1.5. The above simulation parameters were then shared across subjects. To synthesize subject variability, we perturbed the activation amplitudes of each subject with $\mathcal{N}(0, 0.05)$. Then, for each subject in *Group A*, BOLD time courses at 3 time points ($t = 0, 1, 2$) were generated using those simulation parameters. The only difference in simulating *Group B* subjects was that the activation amplitudes of the 3^{rd} network successively changed at $\{t^1, t^2, t^3\}$, with a changing rate sampled from $\mathcal{N}(0, 0.25)$. Consequently, only the covariances within that network could change longitudinally. Fig. 1b shows one example, where the 9 elements within the red square (the 3^{rd} network) are true positives to be identified by group analysis. Unlike the simulation in [4], our strategy relies only on the basic fMRI signal generation mechanism, so that the simulated matrices do not bias any particular covariance modeling technique.

Figs. 2a,b,c show the identified group differences (yellow) derived from three longitudinal analysis methods: (a) our proposed Riemannian framework with the mapping function ϕ^g (Eq. 4); (b) Riemannian framework with parallel transport ϕ^p (defined in [13]); and (c) linear regression on each element across subject-specific covariance matrices $\{\boldsymbol{C}^i\}$ (Euclidean). The final univariate test was the two-sample t-test. The latent p-value map (Section 3.3) was used in *Method a* and *b*, whereas *Method c* only required a deterministic p-value map as it directly compared the trajectories (linear slopes) of different subjects. Significant matrix elements were identified at $p \leq 0.05$ after correcting for multiple comparison via the Bonferroni procedure. Due to the influence of ROI unique activation and the Riccan noise, none of the three methods identified all true positives. However, only *Method a* yielded no false positive. To show the non-deterministic nature of the univariate tests in *Method a* and *b*, we identified group differences using each of the 500 p-value maps and counted the frequency of each element being significant. The spread of the two resulting histograms (Figs. 2d,e) indicates the latent p-value map is preferred over any particular deterministic p-value map. Moreover, our proposed mapping function ϕ^g (Fig. 2d) achieved more consistent group test results compared to the parallel transport (Fig. 2e). Next, to show the importance of using the Fréchet mean as the template, we sampled 1000 $\bar{\boldsymbol{C}}$ via Principal Geodesic Analysis, and generated a deterministic p-value map for each sampled $\bar{\boldsymbol{C}}$. We sorted and binned the 1000 test results according to the distance from the sampled $\bar{\boldsymbol{C}}$ to the Fréchet mean. Fig. 3a indicates when $\bar{\boldsymbol{C}}$ was near the Fréchet mean, the test accuracy significantly increased. Finally, we tested the

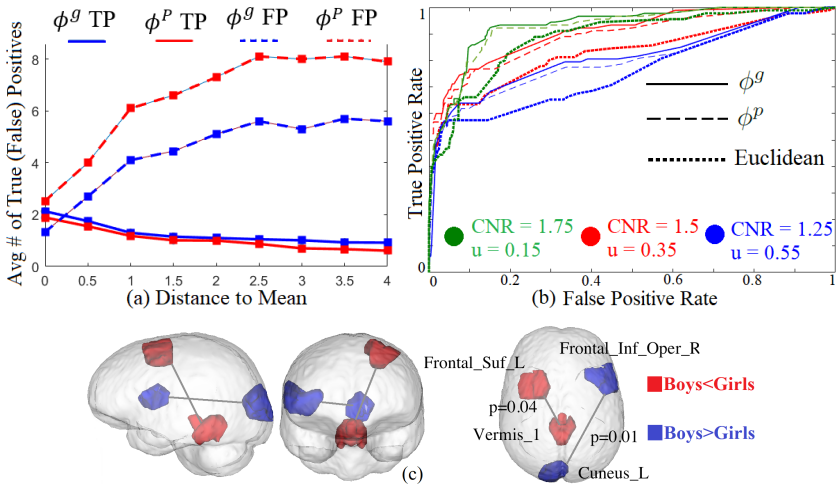

Fig. 3. (a) Number of true (false) positives averaged over templates sampled within a certain distance range. (b) ROC of the three methods under three noise levels. (c) The application of our proposed framework to the NCANDA dataset revealed sex effects in two connections. The p-values were corrected for multiple comparison.

three methods' robustness against noise (u and CNR). Under each noise level, we repeated the experiments (including re-grouping ROIs and re-randomizing activation amplitudes) 10 times. Fig. 3b shows that the Riemannian framework always outperformed the traditional Euclidean method, and our proposed mapping function ϕ^g always achieved the best result.

The NCANDA Dataset. We applied our framework to rs-fMRI data of 246 normal adolescents (age 12-21; 117 boys and 129 girls) from the National Consortium on Alcohol and Neurodevelopment in Adolescence (NCANDA) [14]. Each participant in this dataset was scanned three times (baseline, 1-year follow-up and 2-year follow-up). We adopted the same preprocessing procedure as described in [14]. The brain was parcelated into 100 ROIs using [15]. The final univariate test was a general linear model (GLM) accounting for sex, age, site and race. Using *Method a*, we identified sex effects in two connections (Fig. 3c). Specifically, the connectivity between *Frontal_Suf_L* and *Vermis_1* increases faster in girls. The connectivity between *Frontal_Inf_Oper_R* and *Cuneus_L* increases faster in boys. *Method b* and *c* did not yield any significant finding.

5 Conclusion

Based on the Riemannian modeling of covariance matrices, we introduced a framework for performing group analysis on longitudinal rs-fMRI. Importantly, our proposed tangent vector mapping function and latent p-value strategy aim to best compromise between the geometry of data and clinical interpretation. Both synthetic and real-data experiments indicated the potential of our proposed

framework in longitudinal connectivity analysis. Nevertheless, a theoretical combination of Riemannian geometry and consistent statistics still remains an open topic. In addition, we need to further explore the challenging issue of characterizing multi-level longitudinality, *i.e.*, dynamics in the time courses of each single visit vs. development across visits.

Acknowledgement. This research was supported in part by NIH grants U24AA021697-06, AA005965, AA013521, AA017168.

References

1. Odish, O., et al.: Longitudinal resting state fMRI analysis in healthy controls and premanifest huntington's disease gene carriers: a three-year follow-up study. Hum Brain Mapp. **36**(1) (2015) 110–119
2. van der Horn, H., et al.: The default mode network as a biomarker of persistent complaints after mild traumatic brain injury: A longitudinal functional magnetic resonance imaging study. J Neurotrauma. **34**(23) (2017) 3262–3269
3. Beckmann, C., Mackay, C., Filippini, N., Smith, S.: Group comparison of resting-state fmri data using multi-subject ICA and dual regression. OHBM (2009)
4. Varoquaux, G., Baronnet, F., Kleinschmidt, A., Fillard, P., Thirion, B.: Detection of brain functional-connectivity difference in post-stroke patients using group-level covariance modeling. In: MICCAI. LNCS (2010) 200–208
5. Pennec, X., Fillard, P., Ayache, N.: A Riemannian framework for tensor computing. IJCV **66**(1) (2006) 41–66
6. Hong, Y., Singh, N., Kwitt, R., Niethammer, M.: Group testing for longitudinal data. In: IPMI. LNCS (2015) 139–151
7. Fletcher, P.T.: Geodesic regression and the theory of least squares on Riemannian manifolds. IJCV **105**(2) (2013) 171–185
8. Kim, H.J., Adluru, N., Suri, H., Vemuri, B.C., Johnson, S.C., Singh, V.: Riemannian nonlinear mixed effects models: Analyzing longitudinal deformations in neuroimaging. In: CVPR. (2017) 172–181
9. Ng, B., Dressler, M., Varoquaux, G., Poline, J.B., Greicius, M., Thirion, B.: Transport on Riemannian manifold for functional connectivity-based classification. In: MICCAI. LNCS (2014) 405–412
10. Campbell, K.M., Fletcher, P.T.: Efficient parallel transport in the group of diffeomorphisms via reduction to the Lie algebra. In: MICGen. (2017)
11. Thompson, E.A., Geyer, C.J.: Fuzzy p-values in latent variable problems. Biometrika **94**(1) (2007) 49–60
12. Erhardt, E., Allen, E., Wei, Y., Eichele, T., Calhoun, V.: SimTB, a simulation toolbox for fMRI data under a model of spatiotemporal separability. NeuroImage **59**(4) (2012) 4160–4167
13. Ferreira, R., Xavier, J., Costeira, J.P., Barroso, V.: Newton method for riemannian centroid computation in naturally reductive homogeneous spaces. In: ICASSP. (2006)
14. Müller-Oehring, E., et al.: Influences of age, sex, and moderate alcohol drinking on the intrinsic functional architecture of adolescent brains. Cereb. Cortex **28**(3) (2018) 1049–1063
15. Craddock, R., James, G., Holtzheimer, P., Hu, X., Mayberg, H.: A whole brain fMRI atlas generated via spatially constrained spectral clustering. Hum Brain Mapp. **33**(8) (2012) 1914–1928