
Manifold-regression to predict from MEG/EEG brain signals without source modeling

David Sabbagh ^{*†}, Pierre Ablin, Gaël Varoquaux, Alexandre Gramfort, Denis A. Engemann [§]
Université Paris-Saclay, Inria, CEA, Palaiseau, 91120, France

Abstract

Magnetoencephalography and electroencephalography (M/EEG) can reveal neuronal dynamics non-invasively in real-time and are therefore appreciated methods in medicine and neuroscience. Recent advances in modeling brain-behavior relationships have highlighted the effectiveness of Riemannian geometry for summarizing the spatially correlated time-series from M/EEG in terms of their covariance. However, after artefact-suppression, M/EEG data is often rank deficient which limits the application of Riemannian concepts. In this article, we focus on the task of *regression* with rank-reduced covariance matrices. We study two Riemannian approaches that vectorize the M/EEG covariance between-sensors through projection into a tangent space. The Wasserstein distance readily applies to rank-reduced data but lacks affine-invariance. This can be overcome by finding a common subspace in which the covariance matrices are full rank, enabling the affine-invariant geometric distance. We investigated the implications of these two approaches in synthetic generative models, which allowed us to control estimation bias of a linear model for prediction. We show that Wasserstein and geometric distances allow perfect out-of-sample prediction on the generative models. We then evaluated the methods on real data with regard to their effectiveness in predicting age from M/EEG covariance matrices. The findings suggest that the data-driven Riemannian methods outperform different sensor-space estimators and that they get close to the performance of biophysics-driven source-localization model that requires MRI acquisitions and tedious data processing. Our study suggests that the proposed Riemannian methods can serve as fundamental building-blocks for automated large-scale analysis of M/EEG.

1 Introduction

Magnetoencephalography and electroencephalography (M/EEG) measure brain activity with millisecond precision from outside the head [23]. Both methods are non-invasive and expose rhythmic signals induced by coordinated neuronal firing with characteristic periodicity between minutes and milliseconds [10]. These so-called brain-rhythms can reveal cognitive processes as well as health status and are quantified in terms of the spatial distribution of the power spectrum over the sensor array that samples the electromagnetic fields around the head [3].

Statistical learning from M/EEG commonly relies on covariance matrices estimated from band-pass filtered signals to capture the characteristic scale of the neuronal events of interest [7, 22, 16]. However, covariance matrices do not live in an Euclidean space but a Riemannian manifold.

^{*} Additional affiliation: Inserm, UMRS-942, Paris Diderot University, Paris, France

[†] Additional affiliation: Department of Anaesthesiology and Critical Care, Lariboisière Hospital, Assistance Publique Hôpitaux de Paris, Paris, France

[‡] dav.sabbagh@gmail.com

[§] denis-alexander.engemann@inria.fr

Fortunately, Riemannian geometry offers a principled mathematical approach to use standard linear learning algorithms such as logistic or ridge regression that work with Euclidean geometry. This is achieved by projecting the covariance matrices into a vector space equipped with an Euclidean metric, the tangent space. The projection is defined by the Riemannian metric, for example the geometric affine-invariant metric [5] or the Wasserstein metric [6]. As a result, the prediction error can be substantially reduced when learning from covariance matrices using Riemannian methods [45, 14].

In practice, M/EEG data is often provided in a rank deficient form by platform operators but also curators of public datasets [32, 2]. Its contamination with high-amplitude environmental electromagnetic artefacts often render aggressive offline-processing mandatory to yield intelligible signals. Commonly used tools for artefact-suppression project the signal linearly into a lower dimensional subspace that is hoped to predominantly contain brain signals [40, 42, 34]. But this necessarily leads to inherently rank-deficient covariance matrices for which no affine-invariant distance is defined. One remedy may consist in using anatomically informed source localization techniques that can typically deal with rank deficiencies [17] and can be combined with source-level estimators of neuronal interactions [31]. However, such approaches require domain-specific expert knowledge, imply processing steps that are hard to automate (e.g. anatomical coregistration) and yields pipelines in which excessive amounts of preprocessing are not under control of the predictive model.

In this work, we focus on regression with rank-reduced covariance matrices. We propose two Riemannian methods for this problem. A first approach uses a Wasserstein metric that can handle rank-reduced matrices, yet is not affine-invariant. In a second approach, matrices are projected into a common subspace in which affine-invariance can be provided. We show that both metrics can achieve perfect out-of-sample predictions in a synthetic generative model. Based on the SPoC method [15], we then present a supervised and computationally efficient approach to learn subspace projections informed by the target variable. Finally, we apply these models to the problem of inferring age from brain data [33, 31] on 595 MEG recordings from the Cambridge Center of Aging (Cam-CAN, <http://cam-can.org>) covering an age range from 18 to 88 years [41]. We compare the data-driven Riemannian approaches to simpler methods that extract power estimates from the diagonal of the sensor-level covariance as well as the cortically constrained minimum norm estimates (MNE) which we use to project the covariance into a subspace defined by anatomical prior knowledge.

Notations We denote scalars $s \in \mathbb{R}$ with regular lowercase font, vectors $\mathbf{s} = [s_1, \dots, s_N] \in \mathbb{R}^N$ with bold lowercase font and matrices $\mathbf{S} \in \mathbb{R}^{N \times M}$ with bold uppercase fonts. \mathbf{I}_N is the identity matrix of size N . $[\cdot]^\top$ represents vector or matrix transposition. The Frobenius norm of a matrix will be denoted by $\|\mathbf{M}\|_F^2 = \text{Tr}(\mathbf{M}\mathbf{M}^\top) = \sum |M_{ij}|^2$ with $\text{Tr}(\cdot)$ the trace operator. $\text{rank}(\mathbf{M})$ is the rank of a matrix. The l_2 norm of a vector \mathbf{x} is denoted by $\|\mathbf{x}\|_2^2 = \sum x_i^2$. We denote by \mathcal{M}_P the space of $P \times P$ square real-valued matrices, $\mathcal{S}_P = \{\mathbf{M} \in \mathcal{M}_P, \mathbf{M}^\top = \mathbf{M}\}$ the subspace of symmetric matrices, $\mathcal{S}_P^{++} = \{\mathbf{S} \in \mathcal{S}_P, \mathbf{x}^\top \mathbf{S} \mathbf{x} > 0, \forall \mathbf{x} \in \mathbb{R}^P\}$ the subspace of $P \times P$ symmetric positive definite matrices, $\mathcal{S}_P^+ = \{\mathbf{S} \in \mathcal{S}_P, \mathbf{x}^\top \mathbf{S} \mathbf{x} \geq 0, \forall \mathbf{x} \in \mathbb{R}^P\}$ the subspace of $P \times P$ symmetric semi-definite positive (SPD) matrices, $\mathcal{S}_{P,R}^+ = \{\mathbf{S} \in \mathcal{S}_P^+, \text{rank}(\mathbf{S}) = R\}$ the subspace of SPD matrices of fixed rank R . All matrices $\mathbf{S} \in \mathcal{S}_P^{++}$ are full rank, invertible (with $\mathbf{S}^{-1} \in \mathcal{S}_P^{++}$) and diagonalizable with real strictly positive eigenvalues: $\mathbf{S} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$ with \mathbf{U} an orthogonal matrix of eigenvectors of \mathbf{S} ($\mathbf{U}\mathbf{U}^\top = \mathbf{I}_P$) and $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n)$ the diagonal matrix of its eigenvalues $\lambda_1 \geq \dots \geq \lambda_n > 0$. For a matrix \mathbf{M} , $\text{diag}(\mathbf{M}) \in \mathbb{R}^P$ is its diagonal. We also define the exponential and logarithm of a matrix: $\forall \mathbf{S} \in \mathcal{S}_P^{++}, \log(\mathbf{S}) = \mathbf{U} \text{diag}(\log(\lambda_1), \dots, \log(\lambda_n)) \mathbf{U}^\top \in \mathcal{S}_P$, and $\forall \mathbf{M} \in \mathcal{S}_P, \exp(\mathbf{M}) = \mathbf{U} \text{diag}(\exp(\lambda_1), \dots, \exp(\lambda_n)) \mathbf{U}^\top \in \mathcal{S}_P^{++}$. $\mathcal{N}(\mu, \sigma^2)$ denotes the normal (Gaussian) distribution of mean μ and variance σ^2 . Finally, $\mathbb{E}_s[\mathbf{x}]$ represents the expectation and $\text{Var}_s[\mathbf{x}]$ the variance of any random variable \mathbf{x} w.r.t. their subscript s when needed.

Background and M/EEG generative model MEG or EEG data measured on P channels are multivariate signals $\mathbf{x}(t) \in \mathbb{R}^P$. For each subject $i = 1 \dots N$, the data are a matrix $\mathbf{X}_i \in \mathbb{R}^{P \times T}$ where T is the number of time samples. For the sake of simplicity, we assume that T is the same for each subject, although it is not required by the following method. The *linear instantaneous mixing model* is a valid generative model for M/EEG data due to the linearity of Maxwell's equations [23]. Assuming the signal originates from $Q < P$ locations in the brain, at any time t , the measured signal

vector of subject $i = 1 \dots N$ is a linear combination of the Q source patterns $\mathbf{a}_j^s \in \mathbb{R}^P, j = 1 \dots Q$:

$$\mathbf{x}_i(t) = \mathbf{A}^s \mathbf{s}_i(t) + \mathbf{n}_i(t) , \quad (1)$$

where the patterns form the time and subject-independent source *mixing matrix* $\mathbf{A}^s = [\mathbf{a}_1^s, \dots, \mathbf{a}_Q^s] \in \mathbb{R}^{P \times Q}$, $\mathbf{s}_i(t) \in \mathbb{R}^Q$ is the *source vector* formed by the Q time-dependent sources amplitude, $\mathbf{n}_i(t) \in \mathbb{R}^P$ is a contamination due to noise. Note that the mixing matrix \mathbf{A}^s and sources \mathbf{s}_i are not known.

Following numerous learning models on M/EEG [7, 15, 22], we consider a regression setting where the target y_i is a function of the power of the sources, denoted $p_{i,j} = \mathbb{E}_t[s_{i,j}^2(t)]$. Here we consider the linear model:

$$y_i = \sum_{j=1}^Q \alpha_j f(p_{i,j}) , \quad (2)$$

where $\alpha \in \mathbb{R}^Q$ and $f : \mathbb{R}^+ \rightarrow \mathbb{R}$ is increasing. Possible choices for f that are relevant for neuroscience are $f(x) = x$, or $f(x) = \log(x)$ to account for log-linear relationships between brain signal power and cognition [7, 22, 11]. A first approach consists in estimating the sources before fitting such a linear model, for example using the Minimum Norm Estimator (MNE) approach [24]. This boils down to solving the so-called M/EEG inverse problem which requires costly MRI acquisitions and tedious processing [3]. A second approach is to work directly with the signals \mathbf{X}_i . To do so, models that enjoy some invariance property are desirable: these models are blind to the mixing \mathbf{A}^s and working with the signals \mathbf{x} is similar to working directly with the sources \mathbf{s} . Riemannian geometry is a natural setting where such invariance properties are found [18]. Besides, under Gaussian assumptions, model (1) is fully described by second order statistics [37]. This amounts to working with covariance matrices, $\mathbf{C}_i = \mathbf{X}_i \mathbf{X}_i^\top / T$, for which Riemannian geometry is well developed. One specificity of M/EEG data is, however, that signals used for learning have been rank-reduced. This leads to rank-deficient covariance matrices, $\mathbf{C}_i \in \mathcal{S}_{P,R}^+$, for which specific matrix manifolds need to be considered.

2 Theoretical background to model invariances on $\mathcal{S}_{P,R}^+$ manifold

2.1 Riemannian matrix manifolds

Endowing a continuous set \mathcal{M} of square matrices with a metric, that defines a local Euclidean structure, gives a Riemannian manifold with a solid theoretical framework. Let $\mathbf{M} \in \mathcal{M}$, a K -dimensional Riemannian manifold. For any matrix $\mathbf{M}' \in \mathcal{M}$, as $\mathbf{M}' \rightarrow \mathbf{M}$, $\xi_{\mathbf{M}} = \mathbf{M}' - \mathbf{M}$ belongs to a vector space $\mathcal{T}_{\mathbf{M}}$ of dimension K called the *tangent space* at \mathbf{M} .

The *Riemannian metric* defines an inner product $\langle \cdot, \cdot \rangle_{\mathbf{M}} : \mathcal{T}_{\mathbf{M}} \times \mathcal{T}_{\mathbf{M}} \rightarrow \mathbb{R}$ for each tangent space $\mathcal{T}_{\mathbf{M}}$, and as a consequence a norm in the tangent space $\|\xi\|_{\mathbf{M}} = \sqrt{\langle \xi, \xi \rangle_{\mathbf{M}}}$. Integrating this metric between two points gives a *geodesic distance* $d : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}^+$. It allows to define means on the manifold:

$$\text{Mean}_d(\mathbf{M}_1, \dots, \mathbf{M}_N) = \arg \min_{\mathbf{M} \in \mathcal{M}} \sum_{i=1}^N d(\mathbf{M}_i, \mathbf{M})^2 . \quad (3)$$

The *manifold exponential* at $\mathbf{M} \in \mathcal{M}$, denoted $\text{Exp}_{\mathbf{M}}$, is a smooth mapping from $\mathcal{T}_{\mathbf{M}}$ to \mathcal{M} that preserves local properties. In particular, $d(\text{Exp}_{\mathbf{M}}(\xi_{\mathbf{M}}), \mathbf{M}) = \|\xi_{\mathbf{M}}\|_{\mathbf{M}}$ for $\xi_{\mathbf{M}}$ small enough. Its inverse is the *manifold logarithm* $\text{Log}_{\mathbf{M}}$ from \mathcal{M} to $\mathcal{T}_{\mathbf{M}}$, with $\|\text{Log}_{\mathbf{M}}(\mathbf{M}')\|_{\mathbf{M}} = d(\mathbf{M}, \mathbf{M}')$ for $\mathbf{M}, \mathbf{M}' \in \mathcal{M}$. Finally, since $\mathcal{T}_{\mathbf{M}}$ is Euclidean, there is a linear invertible mapping $\phi_{\mathbf{M}} : \mathcal{T}_{\mathbf{M}} \rightarrow \mathbb{R}^K$ such that for all $\xi_{\mathbf{M}} \in \mathcal{T}_{\mathbf{M}}$, $\|\xi_{\mathbf{M}}\|_{\mathbf{M}} = \|\phi_{\mathbf{M}}(\xi_{\mathbf{M}})\|_2$. This allows to define the *vectorization operator* at $\mathbf{M} \in \mathcal{M}$, $\mathcal{P}_{\mathbf{M}} : \mathcal{M} \rightarrow \mathbb{R}^K$, defined by $\mathcal{P}_{\mathbf{M}}(\mathbf{M}') = \phi_{\mathbf{M}}(\text{Log}_{\mathbf{M}}(\mathbf{M}'))$. Fig. 1 illustrates these concepts.

The vectorization explicitly captures the local Euclidean properties of the Riemannian manifold:

$$d(\mathbf{M}, \mathbf{M}') = \|\mathcal{P}_{\mathbf{M}}(\mathbf{M}')\|_2 \quad (4)$$




Figure 1: Tangent Space, exponential and logarithm on Riemannian manifold illustration.

Hence, if a set of matrices M_1, \dots, M_N is located in a small portion of the manifold, denoting $\bar{M} = \text{Mean}_d(M_1, \dots, M_N)$, it holds:

$$d(M_i, M_j) \simeq \|\mathcal{P}_{\bar{M}}(M_i) - \mathcal{P}_{\bar{M}}(M_j)\|_2 \quad (5)$$

For additional details on matrix manifolds, see [1], chap. 3.

Regression on matrix manifolds The vectorization operator is key for machine learning applications: it projects points in \mathcal{M} on \mathbb{R}^K , and the distance d on \mathcal{M} is approximated by the distance ℓ_2 on \mathbb{R}^K . Therefore, those vectors can be used as input for any standard regression technique, which often assumes a Euclidean structure of the data. More specifically, throughout the article, we consider the following regression pipeline. Given a training set of samples $M_1, \dots, M_N \in \mathcal{M}$ and target continuous variables $y_1, \dots, y_N \in \mathbb{R}$, we first compute the mean of the samples $\bar{M} = \text{Mean}_d(M_1, \dots, M_N)$. This mean is taken as the reference to compute the vectorization. After computing $v_1, \dots, v_N \in \mathbb{R}^K$ as $v_i = \mathcal{P}_{\bar{M}}(M_i)$, a linear regression technique (e.g. ridge regression) with parameters $\beta \in \mathbb{R}^K$ can be employed assuming that $y_i \simeq v_i^\top \beta$.

2.2 Distances and invariances on positive matrices manifolds

We will now introduce two important distances: the geometric distance on the manifold \mathcal{S}_P^{++} (also known as affine-invariant distance), and the Wasserstein distance on the manifold $\mathcal{S}_{P,R}^+$.

The geometric distance Seeking properties of covariance matrices that are invariant by linear transformation of the signal, leads to endow the positive definite manifold \mathcal{S}_P^{++} with the *geometric* distance [18]:

$$d_G(\mathcal{S}, \mathcal{S}') = \|\log(\mathcal{S}^{-1}\mathcal{S}')\|_F = \left[\sum_{i=1}^P \log^2 \lambda_k \right]^{\frac{1}{2}} \quad (6)$$

where $\lambda_k, k = 1 \dots P$ are the real eigenvalues of $\mathcal{S}^{-1}\mathcal{S}'$. The affine invariance property writes:

$$\text{For } \mathbf{W} \text{ invertible, } d_G(\mathbf{W}^\top \mathcal{S} \mathbf{W}, \mathbf{W}^\top \mathcal{S}' \mathbf{W}) = d_G(\mathcal{S}, \mathcal{S}') . \quad (7)$$

This distance gives a Riemannian-manifold structure to \mathcal{S}_P^{++} with the inner product $\langle \mathbf{P}, \mathbf{Q} \rangle_{\mathcal{S}} = \text{Tr}(\mathbf{P} \mathcal{S}^{-1} \mathbf{Q} \mathcal{S}^{-1})$ [18]. The corresponding manifold logarithm at \mathcal{S} is $\text{Log}_{\mathcal{S}}(\mathcal{S}') = \mathcal{S}^{\frac{1}{2}} \log(\mathcal{S}^{-\frac{1}{2}} \mathcal{S}' \mathcal{S}^{-\frac{1}{2}}) \mathcal{S}^{\frac{1}{2}}$ and the vectorization operator $\mathcal{P}_{\mathcal{S}}(\mathcal{S}')$ of \mathcal{S}' w.r.t. \mathcal{S} : $\mathcal{P}_{\mathcal{S}}(\mathcal{S}') = \text{Upper}(\mathcal{S}^{-\frac{1}{2}} \text{Log}_{\mathcal{S}}(\mathcal{S}') \mathcal{S}^{-\frac{1}{2}}) = \text{Upper}(\log(\mathcal{S}^{-\frac{1}{2}} \mathcal{S}' \mathcal{S}^{-\frac{1}{2}}))$, where $\text{Upper}(\mathbf{M}) \in \mathbb{R}^K$ is the vectorized upper-triangular part of \mathbf{M} , with unit weights on the diagonal and $\sqrt{2}$ weights on the off-diagonal, and $K = P(P+1)/2$.

The Wasserstein distance Unlike \mathcal{S}_P^{++} , it is hard to endow the $\mathcal{S}_{P,R}^+$ manifold with a distance that yields tractable or cheap-to-compute logarithms [43]. This manifold is classically viewed as $\mathcal{S}_{P,R}^+ = \{\mathbf{Y}\mathbf{Y}^\top | \mathbf{Y} \in \mathbb{R}_*^{P \times R}\}$, where $\mathbb{R}_*^{P \times R}$ is the set $P \times R$ matrices of rank R [30]. This view allows to write $\mathcal{S}_{P,R}^+$ as a quotient manifold $\mathbb{R}_*^{P \times R} / \mathcal{O}_R$, where \mathcal{O}_R is the orthogonal group of size R . This means that each matrix $\mathbf{Y}\mathbf{Y}^\top \in \mathcal{S}_{P,R}^+$ is identified with the set $\{\mathbf{Y}\mathbf{Q} | \mathbf{Q} \in \mathcal{O}_R\}$.

It has recently been proposed [35] to use the standard Frobenius metric on the total space $\mathbb{R}_*^{P \times R}$. This metric in the total space is equivalent to the *Wasserstein* distance [6] on $\mathcal{S}_{P,R}^+$:

$$d_W(\mathcal{S}, \mathcal{S}') = \left[\text{Tr}(\mathcal{S}) + \text{Tr}(\mathcal{S}') - 2\text{Tr}((\mathcal{S}^{\frac{1}{2}} \mathcal{S}' \mathcal{S}^{\frac{1}{2}})^{\frac{1}{2}}) \right]^{\frac{1}{2}} \quad (8)$$

This provides cheap-to-compute logarithms:

$$\text{Log}_{\mathbf{Y}\mathbf{Y}^\top}(\mathbf{Y}'\mathbf{Y}'^\top) = \mathbf{Y}'\mathbf{Q}^* - \mathbf{Y} \in \mathbb{R}_*^{P \times R} , \quad (9)$$

where $\mathbf{U}\Sigma\mathbf{V}^\top = \mathbf{Y}^\top \mathbf{Y}'$ is a singular value decomposition and $\mathbf{Q}^* = \mathbf{V}\mathbf{U}^\top$. The vectorization operator is then given by $\mathcal{P}_{\mathbf{Y}\mathbf{Y}^\top}(\mathbf{Y}'\mathbf{Y}'^\top) = \text{vect}(\mathbf{Y}'\mathbf{Q}^* - \mathbf{Y}) \in \mathbb{R}^{PR}$, where the *vect* of a matrix is the vector containing all its coefficients.

This framework offers closed form projections in the tangent space for the Wasserstein distance, which can be used to perform regression. Importantly, since $S_P^{++} = S_{P,P}^+$, we can also use this distance on the positive definite matrices. This distance possesses the *orthogonal invariance* property:

$$\text{For } \mathbf{W} \text{ orthogonal, } d_W(\mathbf{W}^\top \mathbf{S} \mathbf{W}, \mathbf{W}^\top \mathbf{S}' \mathbf{W}) = d_W(\mathbf{S}, \mathbf{S}') . \quad (10)$$

This property is weaker than the affine invariance of the geometric distance (7). A natural question is whether such an affine invariant distance also exists on this manifold. Unfortunately, it is shown in [8] that the answer is negative for $R < P$ (proof in appendix 6.3).

3 Manifold-regression models for M/EEG

3.1 Generative model and consistency of linear regression in the tangent space of S_P^{++}

Here, we consider a more specific generative model than (1) by assuming a specific structure on the noise. We assume that the additive noise $\mathbf{n}_i(t) = \mathbf{A}^n \boldsymbol{\nu}_i(t)$ with $\mathbf{A}^n = [\mathbf{a}_1^n, \dots, \mathbf{a}_{P-Q}^n] \in \mathbb{R}^{P \times (P-Q)}$ and $\boldsymbol{\nu}_i(t) \in \mathbb{R}^{P-Q}$. This amounts to assuming that the noise is of rank $P - Q$ and that the noise spans the same subspace for all subjects. Denoting $\mathbf{A} = [\mathbf{a}_1^s, \dots, \mathbf{a}_Q^s, \mathbf{a}_1^n, \dots, \mathbf{a}_{P-Q}^n] \in \mathbb{R}^{P \times P}$ and $\boldsymbol{\eta}_i(t) = [s_{i,1}(t), \dots, s_{i,Q}(t), \nu_{i,1}(t), \dots, \nu_{i,P-Q}(t)] \in \mathbb{R}^P$, this generative model can be compactly rewritten as $\mathbf{x}_i(t) = \mathbf{A} \boldsymbol{\eta}_i(t)$.

We assume that the sources s_i are decorrelated and independent from $\boldsymbol{\nu}_i$: with $p_{i,j} = \mathbb{E}_t[s_{i,j}^2(t)]$ the powers, *i.e.* the variance over time, of the j -th source of subject i , we suppose $\mathbb{E}_t[\mathbf{s}_i(t) \mathbf{s}_i^\top(t)] = \text{diag}((p_{i,j})_{j=1 \dots Q})$ and $\mathbb{E}_t[\mathbf{s}_i(t) \boldsymbol{\nu}_i(t)^\top] = 0$. The covariances are then given by:

$$\mathbf{C}_i = \mathbf{A} \mathbf{E}_i \mathbf{A}^\top , \quad (11)$$

where $\mathbf{E}_i = \mathbb{E}_t[\boldsymbol{\eta}_i(t) \boldsymbol{\eta}_i(t)^\top]$ is a block diagonal matrix, whose upper $Q \times Q$ block is $\text{diag}(p_{i,1}, \dots, p_{i,Q})$.

In the following, we show that different functions f from (2) yield a linear relationship between the y_i 's and the vectorization of the \mathbf{C}_i 's for different Riemannian metrics.

Proposition 1 (Euclidean vectorization). *Assume $f(p_{i,j}) = p_{i,j}$. Then, the relationship between y_i and $\text{Upper}(\mathbf{C}_i)$ is linear.*

Proof. Indeed, if $f(p) = p$, the relationship between y_i and the $p_{i,j}$ is linear. Rewriting Eq. (11) as $\mathbf{E}_i = \mathbf{A}^{-1} \mathbf{C}_i \mathbf{A}^{-\top}$, and since the $p_{i,j}$ are on the diagonal of the upper block of \mathbf{E}_i , the relationship between the $p_{i,j}$ and the coefficients of \mathbf{C}_i is also linear. This means that there is a linear relationship between the coefficients of \mathbf{C}_i and the variable of interest y_i . In other words, y_i is a linear combination of the vectorization of \mathbf{C}_i w.r.t. the standard Euclidean distance. \square

Proposition 2 (Geometric vectorization). *Assume $f(p_{i,j}) = \log(p_{i,j})$. Denote $\bar{\mathbf{C}} = \text{Mean}_G(\mathbf{C}_1, \dots, \mathbf{C}_N)$ the geometric mean of the dataset, and $\mathbf{v}_i = \mathcal{P}_{\bar{\mathbf{C}}}(\mathbf{C}_i)$ the vectorization of \mathbf{C}_i w.r.t. the geometric distance. Then, the relationship between y_i and \mathbf{v}_i is linear.*

The proof is given in appendix 6.1. It relies crucially on the affine invariance property that means that using Riemannian embeddings of the \mathbf{C}_i 's, is equivalent to working directly with the \mathbf{E}_i 's.

Proposition 3 (Wasserstein vectorization). *Assume $f(p_{i,j}) = \sqrt{p_{i,j}}$. Assume that \mathbf{A} is orthogonal. Denote $\bar{\mathbf{C}} = \text{Mean}_W(\mathbf{C}_1, \dots, \mathbf{C}_N)$ the Wasserstein mean of the dataset, and $\mathbf{v}_i = \mathcal{P}_{\bar{\mathbf{C}}}(\mathbf{C}_i)$ the vectorization of \mathbf{C}_i w.r.t. the Wasserstein distance. Then, the relationship between y_i and \mathbf{v}_i is linear.*

The proof is given in appendix 6.2. The restriction to the case where \mathbf{A} is orthogonal stems from the orthogonal invariance of the Wasserstein distance. In the neuroscience literature square root rectifications are however not commonly used for M/EEG modeling. Nevertheless, it is interesting to see that the Wasserstein metric that can naturally cope with rank reduced data is consistent with this particular generative model.

These propositions show that the relationship between the samples and the variable y is linear in the tangent space, motivating the use of linear regression methods (see simulation study in Sec. 4). The argumentation of this section relies on the assumption that the covariance matrices are full rank. However, this is rarely the case in practice.




Figure 2: Proposed regression pipeline. The considered choices for each sequential step are detailed below each box. Identity means no spatial filtering $\mathbf{W} = \mathbf{I}$. Only the most relevant combinations are reported. For example Wasserstein vectorization does not need projections as it directly applies to rank-deficient matrices. Geometric vectorization is not influenced by the choice of projections due to its affine-invariance. Choices for vectorization are depicted by the colors used for visualizing subsequent analyses.

3.2 Learning projections on \mathcal{S}_R^{++}

In order to use the geometric distance on the $\mathbf{C}_i \in \mathcal{S}_{P,R}^+$, we have to project them on \mathcal{S}_R^{++} to make them full rank. In the following, we consider a linear operator $\mathbf{W} \in \mathbb{R}^{P \times R}$ of rank R which is common to all samples (*i.e.* subjects). For consistency with the M/EEG literature we will refer to rows of \mathbf{W} as *spatial filters*. The covariance matrices of ‘spatially filtered’ signals $\mathbf{W}^\top \mathbf{x}_i$ are obtained as: $\Sigma_i = \mathbf{W}^\top \mathbf{C}_i \mathbf{W} \in \mathbb{R}^{R \times R}$. With probability one, $\text{rank}(\Sigma_i) = \min(\text{rank}(\mathbf{W}), \text{rank}(\mathbf{C}_i)) = R$, hence $\Sigma_i \in \mathcal{S}_R^{++}$. Since the \mathbf{C}_i ’s do not span the same image, applying \mathbf{W} destroys some information. Recently, geometry-aware dimensionality reduction techniques, both supervised and unsupervised, have been developed on covariance manifolds [28, 25]. Here we considered two distinct approaches to estimate \mathbf{W} .

Unsupervised spatial filtering A first strategy is to project the data into a subspace that captures most of its variance. This is achieved by Principal Component Analysis (PCA) applied to the averaged covariance matrix computed across subjects: $\mathbf{W}_{\text{UNSUP}} = \mathbf{U}$, where \mathbf{U} contains the eigenvectors corresponding to the top R eigenvalues of the average covariance matrix $\overline{\mathbf{C}} = \frac{1}{N} \sum_{i=1}^N \mathbf{C}_i$. This step is blind to the values of y and is therefore unsupervised. Note that under the assumption that the time series across subjects are independent, the average covariance $\overline{\mathbf{C}}$ is the covariance of the data over the full population.

Supervised spatial filtering We use a supervised spatial filtering algorithm [15] originally developed for intra-subject Brain Computer Interfaces applications, and adapt it to our cross-person prediction problem. The filters \mathbf{W} are chosen to maximize the covariance between the power of the filtered signals and y . Denoting by $\mathbf{C}_y = \frac{1}{N} \sum_{i=1}^N y_i \mathbf{C}_i$ the weighted average covariance matrix, the first filter \mathbf{w}_{SUP} is given by:

$$\mathbf{w}_{\text{SUP}} = \arg \max_{\mathbf{w}} \frac{\mathbf{w}^\top \mathbf{C}_y \mathbf{w}}{\mathbf{w}^\top \overline{\mathbf{C}} \mathbf{w}}.$$

In practice, all the other filters in \mathbf{W}_{SUP} are obtained by solving a generalized eigenvalue decomposition problem (see the proof in Appendix 6.4).

The proposed pipeline is summarized in Fig. 2.

4 Experiments

4.1 Simulations

We start by illustrating Prop. 2. Independent identically distributed covariance matrices $\mathbf{C}_1, \dots, \mathbf{C}_N \in \mathcal{S}_P^{++}$ and variables y_1, \dots, y_N are generated following the above generative model. The matrix \mathbf{A} is taken as $\exp(\mu \mathbf{B})$ with $\mathbf{B} \in \mathbb{R}^{P \times P}$ a random matrix, and $\mu \in \mathbb{R}$ a scalar controlling the distance from \mathbf{A} to identity ($\mu = 0$ yields $\mathbf{A} = \mathbf{I}_P$). We use the log function for f to link the source powers (*i.e.* the variance) to the y_i ’s. Model reads $y_i = \sum_j \alpha_j \log(p_{ij}) + \varepsilon_i$, with $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$ a small additive random perturbation.

We compare three methods of vectorization: the geometric distance, the Wasserstein distance and the non-Riemannian method “log-diag” extracting the \log of the diagonals of C_i as features. Note that the diagonal of C_i contains the powers of each sensor for subject i . A linear regression model is used following the procedure presented in Sec. 2. We take $P = 5$, $N = 100$ and $Q = 2$. We measure the score of each method as the average mean absolute error (MAE) obtained with 10-fold cross-validation. Fig. 3 displays the scores of each method when the parameters σ controlling the noise level and μ controlling the distance from A to I_p are changed. We also investigated the realistic scenario where each subject has a mixing matrix deviating from a reference: $A_i = A + E_i$ with entries of E_i sampled i.i.d. from $\mathcal{N}(0, \sigma^2)$.

The same experiment with $f(p) = \sqrt{p}$ yields comparable results, yet with Wasserstein distance performing best and achieving perfect out-of-sample prediction when $\sigma \rightarrow 0$ and A is orthogonal.




Figure 3: Illustration of Prop.2. Data is generated following the generative model with $f = \log$. The regression pipeline consists in projecting the data in the tangent space, and then use a linear model. The left plot shows the evolution of the score when random noise of variance σ^2 is added to the variables y_i . The MAE of the geometric distance pipeline goes to 0 in the limit of no noise, indicating perfect out-of-sample prediction. This illustrates the linearity in the tangent space for the geometric distance (Prop. 2). The middle plot explores the effect of the parameter μ controlling the distance between A and I_p . Riemannian geometric method is not affected by μ due to its affine invariance property. Although the Wasserstein distance is not affine invariant, its performance does not change much with μ . On the contrary, the log-diag method is sensitive to changes in A . The right plot shows how the score changes when mixing matrices become sample dependent. We can see then only when $\sigma = 0$ supervised + log-diag and Riemann reach perfect performance. Geometric Riemann is uniformly better and indifferent to projection choice. Wasserstein, despite model mismatch, outperforms supervised + log-diag with high σ .

4.2 MEG data

Predicting biological age from MEG on the Cambridge center of ageing dataset In the following, we apply our methods to infer age from brain signals. Age is a dominant driver of cross-person variance in neuroscience data and a serious confounder [39]. As a consequence of the globally increased average lifespan, ageing has become a central topic in public health that has stimulated neuropsychiatric research at large scales. The link between age and brain function is therefore of utmost practical interest in neuroscientific research.

To predict age from brain signals, here we use the currently largest publicly available MEG dataset provided by the Cam-CAN [38]. We only considered the signals from magnetometer sensors ($P = 102$) as it turns out that once SSS is applied (detailed in Appendix 6.6), magnetometers and gradiometers are linear combination of approximately 70 signals ($65 \leq R_i \leq 73$), which become redundant in practice [19]. We considered task-free recordings during which participants were asked to sit still with eyes closed in the absence of systematic stimulation. We then drew $T \simeq 520,000$ time samples from $N = 595$ subjects. To capture age-related changes in cortical brain rhythms [4, 44, 12], we filtered the data into 9 frequency bands: low frequencies $[0.1 - 1.5]$, $\delta[1.5 - 4]$, $\theta[4 - 8]$, $\alpha[8 - 15]$, $\beta_{low}[15 - 26]$, $\beta_{high}[26 - 35]$, $\gamma_{low}[35 - 50]$, $\gamma_{mid}[50 - 74]$ and $\gamma_{high}[76 - 120]$ (Hz unit). These frequencies are compatible with conventional definitions used in the Human Connectome Project [32]. We verify that the covariance matrices all lie on a small portion of the manifold, justifying projection in a common tangent space. Then we applied the covariance pipeline independently in each frequency band and concatenated the ensuing features.

Data-driven covariance projection for age prediction Three types of approaches are here compared: Riemannian methods (Wasserstein or geometric), methods extracting log-diagonal of matrices (with or without supervised spatial filtering, see Sec. 3.2) and a biophysics-informed method based on the MNE source imaging technique [24]. The MNE method essentially consists in a standard Tikhonov regularized inverse solution and is therefore linear (See Appendix 6.5 for details). Here it serves as gold-standard informed by the individual anatomy of each subject. It requires a T1-weighted MRI and the precise measure of the head in the MEG device coordinate system [3] and the coordinate alignment is hard to automate. We configured MNE with $Q = 8196$ candidate dipoles. To obtain spatial smoothing and reduce dimensionality, we averaged the MNE solution using a cortical parcellation encompassing 448 regions of interest from [31, 21]. We then used ridge regression and tuned its regularization parameter by generalized cross-validation [20] on a logarithmic grid of 100 values in $[10^{-5}, 10^3]$ on each training fold of a 10-fold cross-validation loop. All numerical experiments were run using the Scikit-Learn software [36], the MNE software for processing M/EEG data [21] and the PyRiemann package [13]. We also ported to Python some part of the Matlab code of Manopt toolbox [9] for computations involving Wasserstein distance. The proposed method, including all data preprocessing, applied on the 500GB of raw MEG data from the Cam-CAN dataset, runs in approximately 12 hours on a regular desktop computer with at least 16GB of RAM. The preprocessing for the computation of the covariances is embarrassingly parallel and can therefore be significantly accelerated by using multiple CPUs. The actual predictive modeling can be performed in less than a minute on standard laptop. Code used for data analysis can be found on GitHub⁵.




Figure 4: Age prediction on Cam-CAN MEG dataset for different methods, ordered by out-of-sample MAE. The y-axis depicts the projection method, with identity denoting the absence of projection. Colors indicate the subsequent embedding. The biophysics-driven MNE method (blue) performs best. The Riemannian methods (orange) follow closely and their performance depends little on the projection method. The non-Riemannian methods log-diag (green) perform worse, although the supervised projection clearly helps.

Riemannian projections are the leading data-driven methods Fig. 4 displays the scores for each method. The biophysically motivated MNE projection yielded the best performance (7.4y MAE), closely followed by the purely data-driven Riemannian methods (8.1y MAE). The chance level was 16y MAE. Interestingly, the Riemannian methods give similar results, and outperformed the non-Riemannian methods. When Riemannian geometry was not applied, the projection strategy turned out to be decisive. Here, the supervised method performed best: it reduced the dimension of the problem while preserving the age-related variance.

Rejecting a null-hypothesis that differences between models are due to chance would require several independent datasets. Instead, for statistical inference, we considered uncertainty estimates of paired differences using 100 Monte Carlo splits (10% test set size). For each method, we counted how often it was performing better than the baseline model obtained with identity and log-diag. We observed for supervised log-diag 73%, identity Wasserstein 85%, unsupervised geometric 96% and biophysics 95% improvement over baseline. This suggests that inferences will carry over to new data.

Importantly, the supervised spatial filters and MNE both support model inspection, which is not the case for the two Riemannian methods. Fig. 5 depicts the marginal patterns [27] from the supervised filters and the source-level ridge model, respectively. The sensor-level results suggest predictive dipolar patterns in the theta to beta range roughly compatible with generators in visual, auditory and motor cortices. Note that differences in head-position can make the sources appear deeper than

⁵ https://www.github.com/DavidSabbagh/NeurIPS19_manifold-regression-meeeg

they are (distance between the red positive and the blue negative poles). Similarly, the MNE-based model suggests localized predictive differences between frequency bands highlighting auditory, visual and premotor cortices. While the MNE model supports more exhaustive inspection, the supervised patterns are still physiologically informative. For example, one can notice that the pattern is more anterior in the β -band than the α -band, potentially revealing sources in the motor cortex.

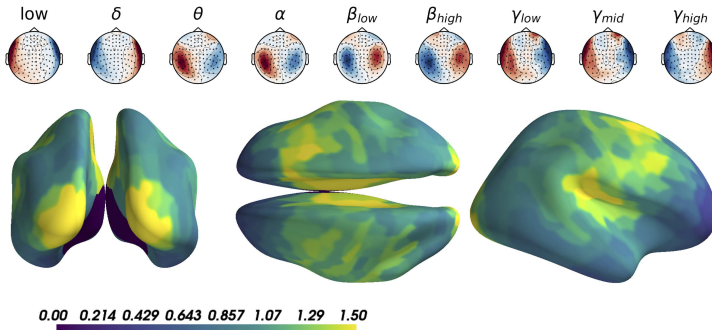


Figure 5: Model inspection. Upper panel: sensor-level patterns from supervised projection. One can notice dipolar configurations varying across frequencies. Lower panel: standard deviation of patterns over frequencies from MNE projection highlighting bilateral visual, auditory and premotor cortices.

5 Discussion

In this contribution, we proposed a mathematically principled approach for regression on rank-reduced covariance matrices from M/EEG data. We applied this framework to the problem of inferring age from neuroimaging data, for which we made use of the currently largest publicly available MEG dataset. To the best of our knowledge, this is the first study to apply a covariance-based approach coupled with Riemannian geometry to regression problem in which the target is defined across persons and not within persons (as in brain-computer interfaces). Moreover, this study reports the first benchmark of age prediction from MEG resting state data on the Cam-CAN. Our results demonstrate that Riemannian data-driven methods do not fall far behind the gold-standard methods with biophysical priors, that depend on manual data processing. One limitation of Riemannian methods is, however, their interpretability compared to other models that allow to report brain-region and frequency-specific effects. These results suggest a trade-off between performance and explainability. Our study suggests that the Riemannian methods have the potential to support automated large-scale analysis of M/EEG data in the absence of MRI scans. Taken together, this potentially opens new avenues for biomarker development.

Acknowledgement

This work was supported by a 2018 “*médecine numérique*” (for digital Medicine) thesis grant issued by Inserm (French national institute of health and medical research) and Inria (French national research institute for the digital sciences). It was also partly supported by the European Research Council Starting Grant SLAB ERC-YStG-676943.

References

- [1] P-A Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization algorithms on matrix manifolds*. Princeton University Press, 2009.
- [2] Anahit Babayan, Miray Erbey, Deniz Kumral, Janis D. Reinelt, Andrea M. F. Reiter, Josefin Röbbig, H. Lina Schaare, Marie Uhlig, Alfred Anwander, Pierre-Louis Bazin, Annette Horstmann, Leonie Lampe, Vadim V. Nikulin, Hadas Okon-Singer, Sven Preusser, André Pampel, Christiane S. Rohr, Julia Sacher, Angelika Thöne-Otto, Sabrina Trapp, Till Nierhaus, Denise Altmann, Katrin Arelin, Maria Blöchl, Edith Bongartz, Patric Breig, Elena Cesnaite, Sufang Chen, Roberto Cozatl, Saskia Czerwonatis, Gabriele Dambrauskaitė, Maria Dreyer, Jessica Enders, Melina Engelhardt, Marie Michele Fischer, Norman Forschack, Johannes Golchert, Laura Goltz, C. Alexandrina Guran, Susanna Hedrich, Nicole Hentschel, Daria I. Hoffmann, Julia M. Huntenburg, Rebecca Jost, Anna Kosatschek, Stella Kunzendorf, Hannah Lammers, Mark E. Lauckner, Keyvan Mahjoory, Ahmad S. Kanaan, Natacha Mendes, Ramona Menger, Enzo Morino, Karina

- Näthe, Jennifer Neubauer, Handan Noyan, Sabine Oligschläger, Patricia Panczyszyn-Trzewik, Dorothee Poehlchen, Nadine Putzke, Sabrina Roski, Marie-Catherine Schaller, Anja Schieferbein, Benito Schlaak, Robert Schmidt, Krzysztof J. Gorgolewski, Hanna Maria Schmidt, Anne Schrimpf, Sylvia Stasch, Maria Voss, Annett Wiedemann, Daniel S. Margulies, Michael Gehler, and Arno Villringer. A mind-brain-body dataset of MRI, EEG, cognition, emotion, and peripheral physiology in young and old adults. *Scientific Data*, 6:180308 EP –, 02 2019.
- [3] Sylvain Baillet. Magnetoencephalography for brain electrophysiology and imaging. *Nature Neuroscience*, 20:327 EP –, 02 2017.
- [4] Luc Berthouze, Leon M James, and Simon F Farmer. Human eeg shows long-range temporal correlations of oscillation amplitude in theta, alpha and beta bands across a wide age range. *Clinical Neurophysiology*, 121(8):1187–1197, 2010.
- [5] Rajendra Bhatia. *Positive Definite Matrices*. Princeton University Press, 2007.
- [6] Rajendra Bhatia, Tanvi Jain, and Yongdo Lim. On the bures–wasserstein distance between positive definite matrices. *Expositiones Mathematicae*, 2018.
- [7] B. Blankertz, R. Tomioka, S. Lemm, M. Kawanabe, and K. Muller. Optimizing spatial filters for robust eeg single-trial analysis. *IEEE Signal Processing Magazine*, 25(1):41–56, 2008.
- [8] Silvere Bonnabel and Rodolphe Sepulchre. Riemannian metric and geometric mean for positive semidefinite matrices of fixed rank. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1055–1070, 2009.
- [9] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research*, 15:1455–1459, 2014.
- [10] György Buzsáki and Rodolfo Llinás. Space and time in the brain. *Science*, 358(6362):482–485, 2017.
- [11] György Buzsáki and Kenji Mizuseki. The log-dynamic brain: how skewed distributions affect network operations. *Nature Reviews Neuroscience*, 15(4):264, 2014.
- [12] C Richard Clark, Melinda D Veltmeyer, Rebecca J Hamilton, Elena Simms, Robert Paul, Daniel Hermens, and Evian Gordon. Spontaneous alpha peak frequency predicts working memory performance across the age span. *International Journal of Psychophysiology*, 53(1):1–9, 2004.
- [13] M. Congedo, A. Barachant, and A. Andreev. A new generation of brain-computer interface based on Riemannian geometry. *arXiv e-prints*, October 2013.
- [14] Marco Congedo, Alexandre Barachant, and Rajendra Bhatia. Riemannian geometry for EEG-based brain-computer interfaces; a primer and a review. *Brain-Computer Interfaces*, 4(3):155–174, 2017.
- [15] Sven Dähne, Frank C Meinecke, Stefan Haufe, Johannes Höhne, Michael Tangermann, Klaus-Robert Müller, and Vadim V Nikulin. Spoc: a novel framework for relating the amplitude of neuronal oscillations to behaviorally relevant parameters. *NeuroImage*, 86:111–122, 2014.
- [16] Jacek Dmochowski, Paul Sajda, Joao Dias, and Lucas Parra. Correlated components of ongoing eeg point to emotionally laden attention – a possible marker of engagement? *Frontiers in Human Neuroscience*, 6:112, 2012.
- [17] Denis A Engemann and Alexandre Gramfort. Automated model selection in covariance estimation and spatial whitening of meg and eeg signals. *NeuroImage*, 108:328–342, 2015.
- [18] Wolfgang Förstner and Boudewijn Moonen. A metric for covariance matrices. In *Geodesy-The Challenge of the 3rd Millennium*, pages 299–309. Springer, 2003.
- [19] Pilar Garcés, David López-Sanz, Fernando Maestú, and Ernesto Pereda. Choice of magnetometers and gradiometers after signal space separation. *Sensors*, 17(12):2926, 2017.
- [20] Gene H. Golub, Michael Heath, and Grace Wahba. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, 21(2):215–223, 1979.

- [21] Alexandre Gramfort, Martin Luessi, Eric Larson, Denis A. Engemann, Daniel Strohmeier, Christian Brodbeck, Lauri Parkkonen, and Matti S. Hämäläinen. MNE software for processing MEG and EEG data. *NeuroImage*, 86:446–460, Feb. 2014.
- [22] M. Grosse-Wentrup* and M. Buss. Multiclass common spatial patterns and information theoretic feature extraction. *IEEE Transactions on Biomedical Engineering*, 55(8):1991–2000, Aug 2008.
- [23] Matti Hämäläinen, Riitta Hari, Risto J Ilmoniemi, Jukka Knuutila, and Olli V Lounasmaa. Magnetoencephalography—theory, instrumentation, and applications to noninvasive studies of the working human brain. *Reviews of modern Physics*, 65(2):413, 1993.
- [24] MS Hämäläinen and RJ Ilmoniemi. Interpreting magnetic fields of the brain: minimum norm estimates. Technical Report TKK-F-A559, Helsinki University of Technology, 1984.
- [25] Mehrtash Harandi, Mathieu Salzmann, and Richard Hartley. Dimensionality reduction on spd manifolds: The emergence of geometry-aware methods. *IEEE transactions on pattern analysis and machine intelligence*, 40(1):48–62, 2017.
- [26] Riitta Hari and Aina Puce. *MEG-EEG Primer*. Oxford University Press, 2017.
- [27] Stefan Haufe, Frank Meinecke, Kai Gørgen, Sven Dähne, John-Dylan Haynes, Benjamin Blankertz, and Felix Bießmann. On the interpretation of weight vectors of linear models in multivariate neuroimaging. *NeuroImage*, 87:96 – 110, 2014.
- [28] Inbal Horev, Florian Yger, and Masashi Sugiyama. Geometry-aware principal component analysis for symmetric positive definite matrices. *Machine Learning*, 106, 11 2016.
- [29] Mainak Jas, Denis A Engemann, Yousra Bekhti, Federico Raimondo, and Alexandre Gramfort. Autoreject: Automated artifact rejection for MEG and EEG data. *NeuroImage*, 159:417–429, 2017.
- [30] Michel Journée, Francis Bach, P-A Absil, and Rodolphe Sepulchre. Low-rank optimization on the cone of positive semidefinite matrices. *SIAM Journal on Optimization*, 20(5):2327–2351, 2010.
- [31] Sheraz Khan, Javeria A Hashmi, Fahimeh Mamashli, Konstantinos Michmizos, Manfred G Kitzbichler, Hari Bharadwaj, Yousra Bekhti, Santosh Ganesan, Keri-Lee A Garel, Susan Whitfield-Gabrieli, et al. Maturation trajectories of cortical resting-state networks depend on the mediating frequency band. *NeuroImage*, 174:57–68, 2018.
- [32] Linda J Larson-Prior, Robert Oostenveld, Stefania Della Penna, G Michalareas, F Prior, Abbas Babajani-Feremi, J-M Schoffelen, Laura Marzetti, Francesco de Pasquale, F Di Pompeo, et al. Adding dynamics to the Human Connectome Project with MEG. *Neuroimage*, 80:190–201, 2013.
- [33] Franziskus Liem, Gaël Varoquaux, Jana Kynast, Frauke Beyer, Shahrzad Kharabian Masouleh, Julia M. Huntenburg, Leonie Lampe, Mehdi Rahim, Alexandre Abraham, R. Cameron Craddock, Steffi Riedel-Heller, Tobias Luck, Markus Loeffler, Matthias L. Schroeter, Anja Veronica Witte, Arno Villringer, and Daniel S. Margulies. Predicting brain-age from multimodal imaging data captures cognitive impairment. *NeuroImage*, 148:179 – 188, 2017.
- [34] Scott Makeig, Anthony J. Bell, Tzyy-Ping Jung, and Terrence J. Sejnowski. Independent component analysis of electroencephalographic data. In *Proceedings of the 8th International Conference on Neural Information Processing Systems*, NIPS’95, pages 145–151, Cambridge, MA, USA, 1995. MIT Press.
- [35] Estelle Massart and Pierre-Antoine Absil. Quotient geometry with simple geodesics for the manifold of fixed-rank positive-semidefinite matrices. Technical report, UCLouvain, 2018. preprint on webpage at <http://sites.uclouvain.be/absil/2018.06>.
- [36] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

- [37] Pedro Luiz Coelho Rodrigues, Marco Congedo, and Christian Jutten. Multivariate time-series analysis via manifold learning. In *2018 IEEE Statistical Signal Processing Workshop (SSP)*, pages 573–577. IEEE, 2018.
- [38] Meredith A Shafto, Lorraine K Tyler, Marie Dixon, Jason R Taylor, James B Rowe, Rhodri Cusack, Andrew J Calder, William D Marslen-Wilson, John Duncan, Tim Dalgleish, et al. The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) study protocol: a cross-sectional, lifespan, multidisciplinary examination of healthy cognitive ageing. *BMC neurology*, 14(1):204, 2014.
- [39] Stephen M Smith and Thomas E Nichols. Statistical challenges in “big data” human neuroimaging. *Neuron*, 97(2):263–268, 2018.
- [40] Samu Taulu and Matti Kajola. Presentation of electromagnetic multichannel data: the signal space separation method. *Journal of Applied Physics*, 97(12):124905, 2005.
- [41] Jason R Taylor, Nitin Williams, Rhodri Cusack, Tibor Auer, Meredith A Shafto, Marie Dixon, Lorraine K Tyler, Richard N Henson, et al. The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) data repository: structural and functional MRI, MEG, and cognitive data from a cross-sectional adult lifespan sample. *Neuroimage*, 144:262–269, 2017.
- [42] Mikko A Uusitalo and Risto J Ilmoniemi. Signal-space projection method for separating MEG or EEG into components. *Medical and Biological Engineering and Computing*, 35(2):135–140, 1997.
- [43] Bart Vandereycken, P-A Absil, and Stefan Vandewalle. Embedded geometry of the set of symmetric positive semidefinite matrices of fixed rank. In *2009 IEEE/SP 15th Workshop on Statistical Signal Processing*, pages 389–392. IEEE, 2009.
- [44] Bradley Voytek, Mark A Kramer, John Case, Kyle Q Lepage, Zechari R Tempesta, Robert T Knight, and Adam Gazzaley. Age-related changes in 1/f neural electrophysiological noise. *Journal of Neuroscience*, 35(38):13257–13265, 2015.
- [45] F. Yger, M. Berar, and F. Lotte. Riemannian approaches in brain-computer interfaces: A review. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 25(10):1753–1762, Oct 2017.