

Revisiting Riemannian Geometry-based EEG decoding through Approximate Joint Diagonalization

Fotis P Kalaganis¹, Nikos A Laskaris², Vangelis P Oikonomou¹, Spiros Nikopolopoulos¹ and Ioannis Kompatsiaris¹

¹ Centre for Research and Technology Hellas, Information Technologies Institute, Multimedia Knowledge and Social Media Analytics Laboratory, Thessaloniki - Thessaloniki, 57001, Greece

² Aristotle University of Thessaloniki, Department of Informatics, AIIA lab, Thessaloniki, 54124, Greece

E-mail: kalaganis@csd.auth.gr

Abstract. *Objective:* The wider adoption of Riemannian geometry in EEG processing is hindered by two factors: i) it involves the manipulation of complex mathematical formulations and, ii) it leads to computationally demanding tasks. The main scope of this work is to simplify particular notions of Riemannian geometry and provide an efficient and comprehensible scheme for neuroscientific explorations.

Approach: To overcome the aforementioned shortcomings, we exploit the concept of Approximate Joint Diagonalization in order to reconstruct the spatial covariance matrices assuming the existence of (and identifying) a common eigenspace in which the application of Riemannian Geometry is significantly simplified.

Main Results: The employed reconstruction process abides to physiologically plausible assumptions, reduces the computational complexity in Riemannian Geometry schemes and bridges the gap between rigorous mathematical procedures and computational neuroscience. Our approach is both formally established and experimentally validated by employing real and synthetic EEG data.

Significance: The implications of the introduced reconstruction process are highlighted by reformulating and re-introducing two signal processing methodologies, namely the “Symmetric Positive Definite (SPD) Matrix Quantization” and the “Coding over SPD Atoms”. The presented approach paves the way for robust and efficient neuroscientific explorations that exploit Riemannian Geometry schemes.

1. Introduction

Computational neuroscience is the field of study in which mathematical tools and theories are used in order to study and model brain functions. It may also incorporate diverse approaches, ranging from physics and chemistry to electrical engineering and computer science. The key purpose is to understand the information processing procedures that underpin the nervous system. This understanding typically passes through the employment of statistical methods and conventional mathematical

tools. More precisely, a typical study in macroscopic computational neuroscience (e.g. studying the brain at a higher level using neuroimaging technologies such as electroencephalography, fMRI, fNIRS, etc) involves, in one way or another, the extraction of feature descriptors that abide to a vector space. This is mainly attributed to the fact that these descriptors mostly rely on statistical modeling in addition to classical pattern recognition methods that are primarily based on features represented as vectors in an Euclidean space [1, 2].

However, classical approaches fail to capture very distinctive characteristics of neuroimaging data such as their inherent structure, the manifolds and the interrelation across the feature dimensions [3]. In the past decade, data analysis that exploits the concept of manifold analysis has been proved to be more robust and has been adopted in many applications. More specifically, Riemannian Geometry has recently entered the picture and is considered among the most promising fields in Neuroimaging data processing [4, 5]. The particular sub-fields that have taken great advantage by this methodological approach are electroencephalography (EEG) decoding and EEG-based Brain Computer Interfaces (BCIs). This is due to the fact that spatial covariance matrices, which are among the simplest functional brain connectivity estimators and a useful tool for brain state decoding, lie on a Riemannian manifold [6].

The key property of covariance matrices that is being exploited, concerns symmetry and positive definiteness[‡]. This property has led to a rapid growth of covariance matrix-based methods, in the fields of signal processing and machine learning, by making use of the well-established mathematical branch of differential geometry. Symmetric Positive Definite (SPD) matrices are known to live in smooth manifolds [7]. Each one of these manifolds, equipped with a Riemannian metric (i.e. an inner product for each point of the tangent space that varies smoothly from point to point), forms a Riemannian manifold [8, 9]. Riemannian geometry, therefore, provides a solid methodological framework for manipulating covariance matrices and performing computations. Its potential for mediating robust BCIs and EEG decoding has already been recognized and examples from the successful application to spatial covariance matrices derived from EEG measurements are reported in the corresponding literature [10–12].

In this context, the choice of the Riemannian metric for studying a particular manifold is of paramount importance. Although a wide variety of metrics have been proposed for studying the manifold of SPD matrices (e.g. Log-Euclidean, Cholesky, Procrustes, etc) [13], in EEG signal processing the Affine-Invariant Riemannian Metric (AIRM) is typically employed, which in turn induces a congruent invariant geodesic distance for SPD matrices. As its name states, the metric is invariant to affine transformations and therefore the sensor and the source spaces are equivalent under this setting. This equivalence is based on a widely accepted assumption, that the observed EEG is well approximated by a linear mixture of independent bipolar brain sources [14], where the mixing model is mathematically formulated as an invertible (i.e. full rank)

[‡] In rigorous terms, covariance matrices are positive semi-definite whereas strictly positive definite if and only if they are full rank.

mixing matrix.

Despite the advantages that Riemannian Geometry brings into the field, a critical issue that still needs further exploration when treating an instance of spatial covariance as an ensemble of features, is the high computational cost during the processing of the involved data [11]. This issue poses severe limitations, in particular for high density EEG recordings and real-time applications. Moreover, since our features are SPD matrices and not vectors, it is extremely difficult to adapt established machine learning and pattern recognition algorithms as it requires in-depth knowledge of this complicated mathematical branch.

It is exactly the two aforementioned shortcomings that we aim to address in this work. Our work relies on the previous, widely accepted, EEG model [15] according to which the electrical fields reflect the activity of a linear combination of presumably independent cortical bipolar sources lying in the brain [16]. Since the sources that produce the brain’s activity are considered anatomically fixed (at least for a certain/short period of time/lifespan) [17], a common unmixing matrix should exist that when applied to the EEG signal will decompose them into a set of mutually uncorrelated signals such that their covariance will be a diagonal matrix. It becomes evident that unmixing the recorded EEG signals and diagonalizing the respective covariance matrices are analogous problems [18, 19]. However, exact diagonalization is only the ideal scenario. In the actual case, the source signals are mixed with noise and then captured by the EEG device.

Therefore, we base our approach on the Approximate Joint Diagonalization (AJD) [20] of the spatial covariance matrices [21]. In practice, AJD finds a common unitary matrix that allows us to approximately diagonalize all the spatial covariance matrices (e.g. covariance matrices correspond to different segments of the same recording session) simultaneously. Then, the obtained dominantly diagonal matrices are used, by considering only their diagonal entries and discarding the almost zero off-diagonal elements. The advantage offered by this procedure is twofold. Firstly, the diagonal elements can be used to reconstruct the covariance matrices under a common mixing model (expressed by the common unitary matrix uncovered through AJD). Intuitively, this procedure serves as denoising on the spatial covariance matrices and leads to more robust estimates of the spatial covariance pattern. Secondly, since the mixing matrix is common for all the covariance matrices and the employed Riemannian metric is invariant to affine transformations, we are able to exploit the covariance matrices in their diagonal form and treat them as vectors. Therefore, the Riemannian Geometry tools typically employed in such settings are simplified, pave the way for computationally efficient matrix calculations while simultaneously maintain the structure of the original matrix manifold. The proposed approach is outlined in Figure 1.

Our approach is initially mathematically established and then experimentally validated by employing both real and synthetic EEG data. More specifically, we analytically prove that the reconstructed matrices are approximately similar to the original ones and maintain the SPD property. Hence, Riemannian Geometry tools

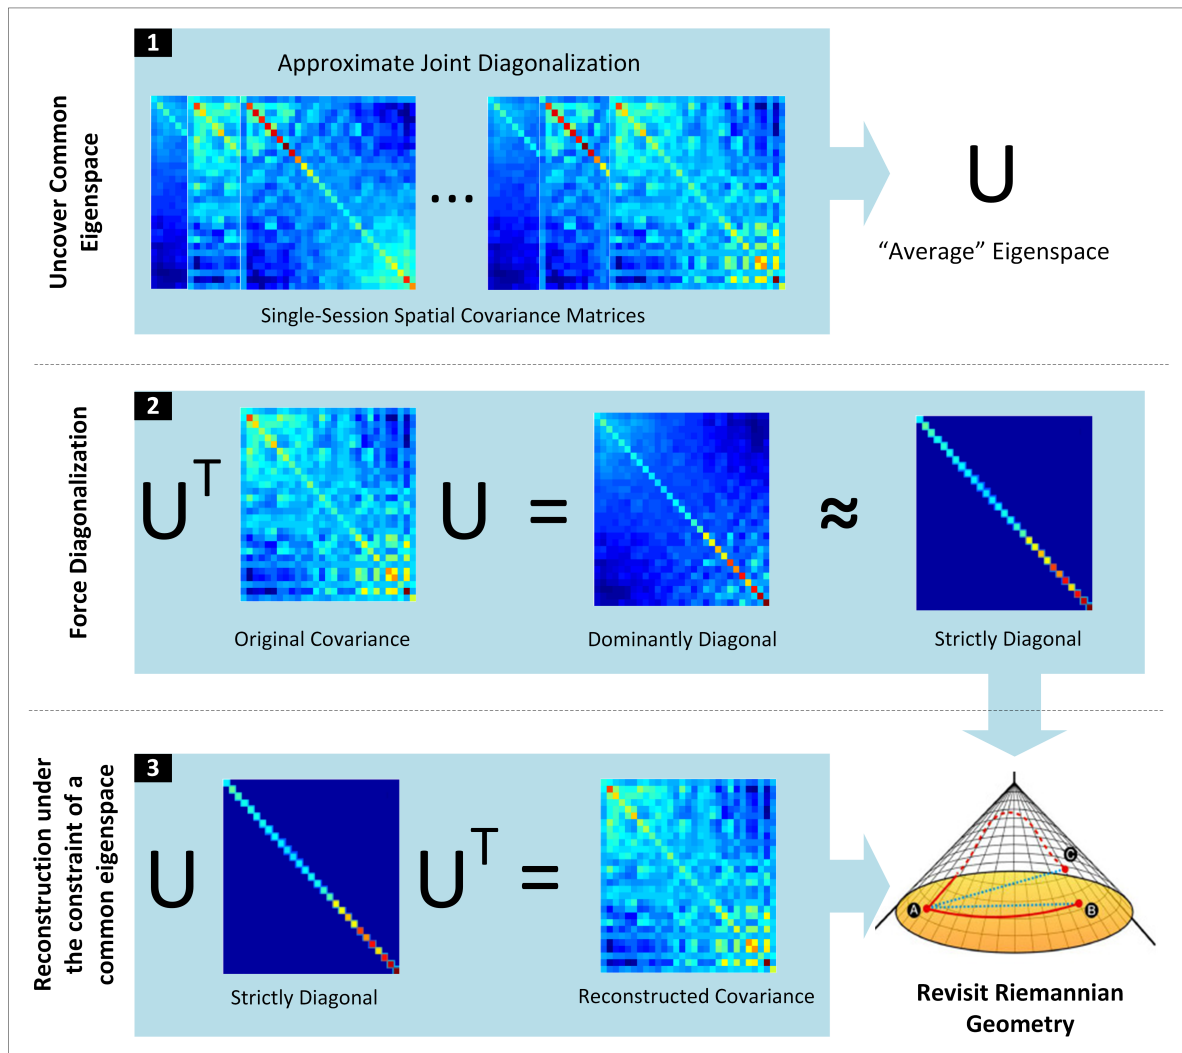


Figure 1. Outline of the proposed methodology. Initially the “average” eigenspace (i.e., the assumed common mixing matrix) is uncovered utilizing the AJD method. Then, this eigenspace is employed to diagonalize the covariance matrices. One may either exploit their strictly diagonal form (depicted in step 2) or reconstruct the covariance matrices (depicted in step 3).

indeed can be employed for appropriate handling of the reconstructed covariance matrices. Additionally, we show that typical Riemannian Geometry-based operations (i.e. Karcher mean calculation and geodesic mean re-centering) are simplified after the diagonalization procedure. Then, we show experimentally that the proposed transformation over a set of covariance matrices that belong to the same recording leads to a new representation that preserves the structure of the original manifold. Moreover, we introduce two approaches that exploit the fact that the reconstructed matrices are jointly diagonalizable in order to facilitate computationally efficient neuroscientific explorations. The first approach extends the idea of vector quantization to the Riemannian manifold of SPD matrices. This enables the realization of covariance matrix quantization in an effort to pave the way for a computationally efficient covariance-based

microstate analysis [22]. The second approach, termed as Coding over SPD Atoms, is based on the idea of expressing a diagonal SPD matrix as a positive sum of diagonal covariance-atoms so as to enable a computationally efficient classification approach for a demanding EEG-based BCI application.

2. Methodology

Here, we present our approach in a formal way. Emphasis is given on the reconstruction of spatial covariance matrices under a common eigenspace. This procedure can be interpreted as a robust estimation for covariance matrices under the constraint of a common mixing matrix and uncorrelated activity in the source space. Beyond that, it is our scope to re-introduce basic concepts of Riemannian Geometry, that are typically employed in EEG processing, tailored specifically for jointly diagonalizable spatial matrices. To further demonstrate the simplicity induced by our approach we present two methodologies, namely SPD Matrix Quantization and Coding over SPD Atoms, with the sole purpose being the demonstration of the capabilities that our approach brings. Most of the code that was deployed for this research is publicly available§.

2.1. Elements of Riemannian Geometry

Let $\mathbf{X}_i \in \mathbb{R}^{E \times T}, i = 1, \dots, n$ be a multichannel EEG segment, where E denotes the number of electrodes, T the number of time samples and n the number of available segments (or trials). Each segment (assuming zero mean signals) can also be described by the corresponding spatial covariance matrix $\mathbf{C}_i = \frac{1}{T-1} \mathbf{X}_i \mathbf{X}_i^\top \in \mathbb{R}^{E \times E}$, where $(\cdot)^\top$ denotes the transpose operator. By definition and under a sufficiently large T value to guarantee a full rank covariance matrix, spatial covariance matrices are SPD that lie on a Riemannian manifold instead of a vector space (e.g. scalar multiplication does not hold on the SPD manifold). In the field of differential geometry, a Riemannian manifold is a real, smooth manifold endowed with an inner product on the tangent space of each point that varies smoothly from point to point.

When treating EEG data, the manifold of SPD matrices denoted by $Sym_E^+ = \{\mathbf{C} \in \mathbb{R}^{E \times E} : \mathbf{x}^\top \mathbf{C} \mathbf{x} > 0, \text{ for all non-zero } \mathbf{x} \in \mathbb{R}^E\}$, is typically studied when it is equipped with the AIRM [7],

$$\langle \mathbf{A}, \mathbf{B} \rangle_{\mathbf{P}} \triangleq Trace(\mathbf{P}^{-1} \mathbf{A} \mathbf{P}^{-1} \mathbf{B}) \quad (1)$$

for $\mathbf{P} \in Sym_E^+$ and $\mathbf{A}, \mathbf{B} \in T_E^+(\mathbf{P})$, where $T_E^+(\mathbf{P})$ denotes the tangent space of Sym_E^+ at \mathbf{P} . Then, the following geodesic distance is induced

$$\delta(\mathbf{C}_i, \mathbf{C}_j) = \left\| \log m(\mathbf{C}_i^{-1/2} \mathbf{C}_j \mathbf{C}_i^{-1/2}) \right\|_F = \sqrt{\sum_{q=1}^E \log^2 \lambda_q} \quad (2)$$

where $\log m(\cdot)$ denotes the matrix logarithm operator and λ_q the eigenvalues of the matrix $\mathbf{C}_i^{-1/2} \mathbf{C}_j \mathbf{C}_i^{-1/2}$ or equivalently of the matrix $\mathbf{C}_i^{-1} \mathbf{C}_j$. We note that these two

§ https://github.com/fkalaganis/Riemannian_AJD

matrices are similar (i.e., hold the same eigenvalues) and that the indices i and j can be permuted. Among the other useful properties that are discussed thoroughly in [23], δ is congruent invariant for non singular matrices \mathbf{W} , i.e. $\delta(\mathbf{W}\mathbf{C}_i\mathbf{W}^\top, \mathbf{W}\mathbf{C}_j\mathbf{W}^\top) = \delta(\mathbf{C}_i, \mathbf{C}_j)$. This is an important property in EEG signal processing since it provides equivalence between the sensor and the source space [6]. According to the prevailing EEG model, the recorded activity is well approximated by a linear mixture of source signals. Hence, $\mathbf{X}_i = \mathbf{M}\mathbf{S}_i$ with \mathbf{M} denoting the mixing matrix and \mathbf{S}_i the source signals. Then, by substituting the observed signal with the equivalent mixing of sources, one may obtain the following covariance matrix, $\mathbf{C}_i = \frac{1}{T-1}\mathbf{M}\mathbf{S}_i\mathbf{S}_i^\top\mathbf{M}^\top$. Therefore, the mixing procedure in the time domain results in a congruent transformation in the corresponding covariance matrices. It becomes obvious that since δ is invariant to such transformations, the two spaces are considered equivalent. In a strict mathematical sense this is partially true (e.g., for certain forms of \mathbf{M}) and this topic is thoroughly discussed in [6]. Hereafter, the terms ‘‘AIRM-induced geodesic distance’’ or simply ‘‘geodesic distance’’ will be used interchangeably and will refer to Equation 2.

2.2. Revisiting Riemannian Geometry through AJD

The mixing matrix, \mathbf{M} , depends on the dipole position and orientation in the brain, the physical properties of the head and the positioning of the electrodes on the scalp. Therefore, it is sound to consider that \mathbf{M} is constant for a certain period of time (e.g. during a single recording session). Given that sources are independent and the corresponding activity (e.g., source signals) is mutually uncorrelated, the source spatial covariance matrices are diagonal.

The estimation of \mathbf{M} from the observed sensor-signals is an ill-posed problem referred to as Blind Source Separation (BSS). There are two approaches related to the BSS problem[24]: The most popular is known as Independent Component Analysis [25], where the goal is to transform the data such that the components become as independent as possible. An alternative approach is based on the notion of diagonality of certain characteristic matrices derived from the data [26]. In such a case, one may approximate \mathbf{M}^{-1} by exploiting the concept of AJD [27]. The joint diagonalization of a set of symmetric square matrices consists in finding the orthonormal change of basis, \mathbf{U} such that $\mathbf{U}^\top = \mathbf{U}^{-1}$, which makes the matrices as diagonal as possible. When the matrices in the set are ‘‘almost exactly jointly diagonalizable’’ (which is the prevailing assumption here), this approach intuitively defines the ‘‘average eigenspace’’ of the matrix set [20].

Following the notation of section 2.1, with \mathbf{C}_i we will denote the spatial covariance matrices that correspond to the EEG segments, \mathbf{X}_i . Let \mathbf{U} be the orthonormal matrix calculated by AJD over the set of \mathbf{C}_i with $i = 1, \dots, n$ that estimates the mixing matrix \mathbf{M} . Then, each \mathbf{C}_i can be converted to a dominantly diagonal matrix by following the transformation, $\mathbf{U}^\top\mathbf{C}_i\mathbf{U}$. Here we introduce the *diag*(\cdot) notation, which discards the non-diagonal elements of a matrix and results on a strictly diagonal matrix whose elements correspond to the elements of the main diagonal of the initial matrix.

Therefore, given that $\text{diag}(\mathbf{U}^\top \mathbf{C}_i \mathbf{U})$ is dominantly diagonal, we can reconstruct (e.g., re-estimate them under the constraint of a common eigenspace) all the spatial covariance matrices as $\tilde{\mathbf{C}}_i = \mathbf{U} \text{diag}(\mathbf{U}^\top \mathbf{C}_i \mathbf{U}) \mathbf{U}^\top$.

Lemma 1: The reconstructed covariance matrices are a “good” approximation (i.e., approximately similar matrices) of the original ones.

Proof: Let \mathbf{U} be an orthonormal matrix that approximately diagonalizes an SPD matrix \mathbf{C} . Since \mathbf{C} is an SPD matrix it has an eigen-decomposition $\mathbf{C} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^\top$ with $\mathbf{\Lambda}$ being a diagonal matrix with positive elements only. Since \mathbf{U} approximately diagonalizes \mathbf{C} then we have the approximately diagonal matrix as follows:

$$\begin{aligned}
\mathbf{L} &= \mathbf{U}^\top \mathbf{C} \mathbf{U} && \text{(substituting } \mathbf{C} \text{ with } \mathbf{V} \mathbf{\Lambda} \mathbf{V}^\top) \\
\text{diag}(\mathbf{L}) &\approx \mathbf{U}^\top \mathbf{V} \mathbf{\Lambda} \mathbf{V}^\top \mathbf{U} && \text{(multiplying left by } \mathbf{U} \text{ and right } \mathbf{U}^\top) \\
\mathbf{U} \text{diag}(\mathbf{L}) \mathbf{U}^\top &\approx \mathbf{U} \mathbf{U}^\top \mathbf{V} \mathbf{\Lambda} \mathbf{V}^\top \mathbf{U} \mathbf{U}^\top && \text{(recall that } \mathbf{U}^\top \mathbf{U} = \mathbf{U} \mathbf{U}^\top = \mathbf{I}) \\
\mathbf{U} \text{diag}(\mathbf{L}) \mathbf{U}^\top &\approx \mathbf{V} \mathbf{\Lambda} \mathbf{V}^\top && \text{(substituting back } \mathbf{V} \mathbf{\Lambda} \mathbf{V}^\top \text{ with } \mathbf{C}) \\
\mathbf{U} \text{diag}(\mathbf{L}) \mathbf{U}^\top &\approx \mathbf{C}
\end{aligned} \tag{3}$$

It becomes obvious from the penultimate equation that the reconstructed covariance matrices are almost similar to the initial ones (i.e., hold the same eigenvalues). This is an important aspect of the reconstruction process since the AIRM-induced geodesic distance is highly affected by the involved matrices’ eigenvalues. Moreover, since δ is invariant to congruence transformations, $\delta(\mathbf{C}_i, \mathbf{C}_j) \approx \delta(\text{diag}(\mathbf{L}_i), \text{diag}(\mathbf{L}_j))$, working with the reconstructed diagonal matrices is equivalent to working with the initial ones. We note that in supplementary material we show that $\delta(\mathbf{L}_i, \text{diag}(\mathbf{L}_i))$, which expresses the approximation error between the initial and the reconstructed matrix, is bounded by $\frac{\sqrt{E} \|\mathbf{L} - \text{diag}(\mathbf{L})\|}{\sqrt{\sum_i \min\{L_{ii}^2, \lambda_i^2\}}}$ where L_{ii} and λ_i are the diagonal elements and the eigenvalues of \mathbf{L} respectively. Additionally, the Results Section provide experimental evidence to support the validity of this reconstruction process.

Having shown the aforementioned, we need to show that the reconstructed matrices (as obtained through force diagonalization) still hold the SPD property. The proof lies on the concept that a matrix is SPD if and only if its eigenvalues are strictly positive.

Lemma 2: The reconstructed covariance matrices (i.e., under the constraint of a common eigenspace) maintain the SPD property.

Proof: Let \mathbf{U} be an orthonormal matrix and \mathbf{C} a full-rank spatial covariance matrix. Since \mathbf{C} is a real SPD matrix, it admits an eigen-decomposition $\mathbf{C} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^\top$ with \mathbf{V} being orthonormal and $\mathbf{\Lambda}$ being a diagonal matrix with positive elements on its diagonal. We transform \mathbf{C} such that $\mathbf{L} = \mathbf{U}^\top \mathbf{C} \mathbf{U}$. By setting $\mathbf{Q} = \mathbf{U}^\top \mathbf{V}$, we may rewrite \mathbf{L} as $\mathbf{L} = \mathbf{Q} \mathbf{\Lambda} \mathbf{Q}^\top$, with \mathbf{Q} being also orthonormal (the product of two orthonormal matrices is also orthonormal). Then, the diagonal entries of \mathbf{L} are in the form of:

$$L_{ii} = Q_{i1} \lambda_1 Q_{i1} + Q_{i2} \lambda_2 Q_{i2} + \dots + Q_{in} \lambda_n Q_{in} = Q_{i1}^2 \lambda_1 + Q_{i2}^2 \lambda_2 + \dots + Q_{in}^2 \lambda_n \tag{4}$$

Since all λ_i are positive (eigenvalues of an SPD matrix) and there exists at least one non-zero matrix element Q_{ij} for each i (\mathbf{Q} constitutes a base in \mathbb{R}^E), then the above

summation is strictly positive. Since these elements serve as the eigenvalues of the reconstructed covariance matrix, this matrix still holds the SPD property. This fact allows us to employ Riemannian Geometry in order to manipulate the reconstructed covariance matrices.

2.3. Exploiting Jointly Diagonalizable Matrices in Riemannian Geometry

In this section we present two approaches that exploit the fact that the reconstructed spatial covariance matrices (as obtained through the process described in Section 2.2) admit a common eigenspace defined by \mathbf{U} . Before describing the advanced neuroscientific explorations, namely i) SPD Matrix Quantization, and ii) Coding over SPD Atoms, we will prove that the geodesic mean (i.e., Karcher mean) of diagonal SPD matrices is also a diagonal SPD matrix. Moreover, we will show that the typical Riemannian Geometry-based transfer learning procedure followed to align different sets of spatial covariance matrices (each stemming from different subjects and/or recording) so that each set of matrices has the identity matrix as its geodesic mean [28], serves as a scaling factor in order to normalize the energy of the source signals (each signal normalized by a different scaling factor).

Let us denote by $\{\mathbf{C}_i : i = 1, \dots, n\}$, a set of SPD matrices that admit a common eigenspace. Each \mathbf{C}_i admits an eigendecomposition in the form of $\mathbf{C}_i = \mathbf{U}\mathbf{D}_i\mathbf{U}^\top$ with \mathbf{D}_i being a diagonal matrix with positive elements only and \mathbf{U} being an orthonormal matrix that expresses their common eigenspace. We know that $\mathcal{K}(\mathbf{U}^\top\mathbf{C}_i\mathbf{U}) = \mathbf{U}^\top\mathcal{K}(\mathbf{C}_i)\mathbf{U}$, where $\mathcal{K}(\cdot)$ denotes the Karcher mean over a set of matrices. Hence, $\mathcal{K}(\mathbf{D}_i) = \mathbf{U}^\top\mathcal{K}(\mathbf{C}_i)\mathbf{U}$, or equivalently, $\mathcal{K}(\mathbf{C}_i) = \mathbf{U}\mathcal{K}(\mathbf{D}_i)\mathbf{U}^\top$. By showing that the matrix $\mathcal{K}(\mathbf{D}_i)$ is also diagonal we are able to infer that the geodesic mean of matrices that admit a common eigenspace also admits the same eigenspace.

Lemma 3: The Karcher mean, under the AIRM, of diagonal matrices is also a diagonal matrix.

Proof: The computation of the Karcher mean relies on the convergence of the following iterative formula

$$\mathcal{K}^{(\nu+1)}(\mathbf{D}_i) = \mathcal{K}^{(\nu)}(\mathbf{D}_i) \exp \left(-\theta \sum_{i=1}^n \log m(\mathbf{D}_i^{-1} \mathcal{K}^{(\nu)}(\mathbf{D}_i)) \right) \quad (5)$$

with ν expressing the iteration step, θ a positive tuning parameter and $\mathcal{K}^{(0)}(\mathbf{D}_i) = \mathbf{I}$. It is known that the power, the logarithm and exponentiation of a diagonal matrix lead to a diagonal matrix. Moreover, the product of two diagonal matrices is a diagonal matrix too. Therefore, we may deduce from Equation 5 that the Karcher mean of diagonal matrices is also a diagonal matrix.

Having the Karcher mean over $\{\mathbf{C}_i\}$ calculated, which for sake of notation simplicity will be denoted as \mathcal{K}_C , we may align the set of covariance matrices by applying the transformation $\mathcal{K}_C^{-1/2}\mathbf{C}_i\mathcal{K}_C^{-1/2}$. Then, the geodesic mean of the new, transformed covariance matrices, will be the identity matrix. This, transfer learning, procedure is

typically employed in order to align multiple recordings so that each set of covariance matrices (i.e., each set stems from a different recording or subject) is centered at (i.e., has as geodesic mean) the identity matrix. Such a procedure offers the ability to use data from previous sessions or from a dataset of past users to calibrate and initialize the decoding scheme, allowing a calibration-less mode of operation.

It is important to notice that the previous algorithmic procedure of alignment, which operates on the covariance matrices, is equivalent to applying the $\mathcal{K}_C^{-1/2}$ as a spatial filter in the time domain. Therefore, the sensor signals \mathbf{X}_i are spatially filtered as $\mathcal{K}_C^{-1/2}\mathbf{X}_i$. Exploiting Lemma 3, this is equal to $\mathbf{U}\mathcal{K}_D^{-1/2}\mathbf{U}^\top\mathbf{X}_i$. Then, replacing \mathbf{X}_i by its linear mixture of sources (where \mathbf{U} serves as the mixing matrix) leads to $\mathbf{U}\mathcal{K}_D^{-1/2}\mathbf{U}^\top\mathbf{U}\mathbf{S}_i$. Since \mathbf{U} is orthonormal the spatial filtering operation is equal to $\mathbf{U}\mathcal{K}_D^{-1/2}\mathbf{S}_i$ where the term $\mathcal{K}_D^{-1/2}\mathbf{S}_i$ expresses the scaling of sources. More precisely, each row of \mathbf{S}_i , which is a source signal, is multiplied by the corresponding diagonal element of $\mathcal{K}_D^{-1/2}$. After the source signals have been scaled, they are mixed by \mathbf{U} .

2.3.1. SPD Matrix Quantization In this section we exploit jointly diagonalizable SPD matrices in order to reformulate the concept of matrix quantization in a computationally efficient way. As in the typical vector quantization, it allows us to assign each SPD matrix to its closest prototype and consequently estimate the discretized description of a probability density function over a set of predefined prototypical matrices. Therefore, the following approach gives rise to a computationally efficient framework for functional connectivity microstate analysis that abides to Riemannian Geometry. The advantage of this approach is that brain's functional connectivity is investigated as whole, instead of studying connectivity features individually.

The most straightforward way to obtain the prototypical matrices, collectively referred to as the codebook, is through the k-Means algorithm. The classical k-Means algorithm has been extended particularly for the SPD manifold under several Riemannian metrics, including the AIRM [29]. Given that the initial set of matrices admits a common eigenspace and the prototypes are calculated via the Karcher mean, these prototypes abide to the same eigenspace too. In order to assign each covariance to the closest prototype, it is essential to calculate its distance from each of the available prototypes.

Formally, the codebook $\{\mathbf{A}_i\}$ with $i = 1, \dots, k$ divides the manifold into k voronoi regions as follows

$$\mathbf{R}_i = \{\mathbf{C} \in \text{Sym}_E^+ : \delta(\mathbf{C}, \mathbf{A}_i) \leq \delta(\mathbf{C}, \mathbf{A}_j) \forall j = 1, \dots, i-1, i+1, \dots, k\} \quad (6)$$

Then, to assign each spatial covariance matrix \mathbf{C}_i to its closest \mathbf{A}_i , one has to calculate $\delta(\mathbf{C}_i, \mathbf{A}_j) \forall j = 1, \dots, k$. However, both of the involved matrices admit a common eigenspace \mathbf{U} . By exploiting the congruence invariance property of δ we may simply calculate $\delta(\mathbf{U}^\top\mathbf{C}_i\mathbf{U}, \mathbf{U}^\top\mathbf{A}_j\mathbf{U})$. Since under this transformation the matrices are diagonal, the calculation of δ reduces from $\mathcal{O}(E^3)$ to $\mathcal{O}(E)$ with E being the number

of sensors. Therefore, the computational cost, in the context of Riemannian geometry-based microstate analysis, is significantly reduced. This is supported by the fact that the distance function, δ , has to be computed k times (k denoting the number of prototypes that form the codebook) for each of the given covariance matrices in order to uncover the closest prototype.

2.3.2. Coding over SPD Atoms In the context of Riemannian geometry, brain’s functional connectivity is typically captured by a single covariance matrix and studied as a whole. However, this functional connectivity estimator might squeeze, in a fuzzified summary, distinct underlying neural processes that take place in the brain simultaneously or sequentially. In other words, the estimator may result in an artificial superposition of connectivity patterns (“sub-connectivities”), with each one having the potential to uniquely be ascribed to a different neural process. Therefore, when investigating a particular neurophysiological response, it is expected that among the involved sub-connectivities only few will have a dominant presence, presumably those directly related to the phenomenon of interest. The following approach paves the way for neuroscientific explorations of such kind.

The key concept here is to re-write a spatial covariance as a positive sum of prototypical covariance matrices (will be referred to as covariance atoms or simply atoms). Although the process of obtaining the covariance atoms is a critical aspect, it is beyond the scope of this work to propose such a methodology. Here we exploit the fact that covariance matrices can be reconstructed by admitting common eigenvectors and we focus on the representation over the atoms (e.g., as obtained through the Riemannian k-Means). We show that this problem is reduced to the problem of obtaining a representation of the vectors holding the eigenvalues under a positivity constraint.

Let $\{\mathbf{A}_i\}$ with $i = 1, \dots, k$ be the set of atom matrices and \mathbf{C} a spatial covariance matrix. We note that each of \mathbf{A}_i and \mathbf{C} are SPD matrices that admit an eigendecomposition $\mathbf{U}\mathbf{D}_i\mathbf{U}^\top$ and $\mathbf{U}\mathbf{A}\mathbf{U}^\top$ respectively. We wish to find the appropriate weights $w_i \geq 0$, $i = 1, \dots, k$ such that \mathbf{C} and $w_1\mathbf{A}_1 + w_2\mathbf{A}_2 + \dots + w_k\mathbf{A}_k$ are as close as possible under the AIRM-induced geodesic distance. We note that the positivity constraint imposed for w_i ensures that the resulting matrix maintains its SPD property. By substituting the matrices with their eigendecomposition forms, the above summation can be written as $\mathbf{U}(w_1\mathbf{D}_1 + w_2\mathbf{D}_2 + \dots + w_k\mathbf{D}_k)\mathbf{U}^\top$. Since δ is invariant to congruence, we want to minimize $\delta(\mathbf{A}, w_1\mathbf{D}_1 + w_2\mathbf{D}_2 + \dots + w_k\mathbf{D}_k)$. By exploiting δ ’s analytic form, we obtain the following optimization problem:

$$\underset{w_1, \dots, w_k}{\operatorname{argmin}} \sum_{j=1}^E \left[\log \left(\frac{\Lambda_j}{w_1 D_{1j} + \dots + w_k D_{kj}} \right) \right]^2 \quad s.t. \quad w_1, \dots, w_k \geq 0 \quad (7)$$

where Λ_j denotes the j -th diagonal element of \mathbf{A} and D_{kj} to the j -th diagonal element of \mathbf{D}_k (i.e., the element that lies at position (j, j)). Since the involved matrices in the optimization problem are strictly diagonal the formula can be expressed using vectors

that contain their diagonal elements. Formally, Equation 7 is a non-linear optimization problem imposed by linear constraints, therefore, a wide variety of solvers can be employed to reach an optimal solution [30, 31]. The simplicity of our solution comes in contrast to the existing frameworks for SPD coding over atoms. Existing approaches found in recent literature [32], which do not exploit the concept of a common eigenspace, are based on complex mathematical manipulations and require specialized software for reaching a solution. A procedure that comes with an extremely high computational cost.

3. Datasets

In order to demonstrate and validate the proposed methodology we employ three distinct EEG datasets with each one enabling us to present its implications in different aspects of brain signal analytics. The first dataset concerns multi-channel EEG recordings of event related responses and it is included here in an attempt to differentiate between attentive and passive recording condition. The second, is a motor imagery dataset and our methodology will be employed so as to differentiate between distinct mental states (imaging of closing left or right fist). Finally, the proposed approach is also validated on synthetic EEG data. Synthetic data generation allows us to construct EEG signals that exhibit particular phenomena of interest.

3.1. Dataset A: Attention

Six subjects (four female, two male) participated in a pc-game, where they were engaged into the scenario of controlling a racing-car by means of eye-saccades while their brain activity was being recorded with 64-channel EEG device [33]. While the car was moving ahead, a wall appeared suddenly on the left (right) of the road and the participant had to avoid it by performing an anti-saccade (i.e. moving his gaze towards the opposite direction). The wall had a checkerboard pattern embedded on it, like the ones traditionally employed to elicit the transient visual evoked potential (tVEP) responses in neurophysiology research. During a second round, the subjects performed the same task but this time without having to perform the anti-saccade. The subjects should simply watch the racing car passively while it was avoiding automatically the wall. The first condition is referred to as attentive whereas the second as passive and the brain-decoding task was defined as “predict the attention level” based merely on the early visual evoked response. This is a challenging task since the spatio-temporal characteristics of the EEG do not offer any significant separability [34]. With the scope of isolating the neural patterns that are mostly associated with the conditions we aim to separate, the multichannel signal was first band-limited within the alpha-rhythm range (i.e. 8-13 Hz) [35, 36] and covariance patterns were derived for the latency range [-50 400] msec, with zero denoting the latency of wall appearance. This time interval includes a well-defined brain activation pattern that appears in both cases (i.e., attentive and

passive), known as P100 response, and excludes the anti-saccade event which takes place 3.6 seconds later.

3.2. Dataset B: Physionet Motor Imagery

The original dataset (DOI: 10.13026/C28G6P) contains the brain activity of 109 healthy participants, recorded via a 64-sensors EEG array at a sampling frequency of 160 Hz [37]. The participants' brain activity was recorded while they were performing a series of pseudo-randomized cue-triggered Motor Imagery (MI) tasks. Initially, we excluded data from six participants (subjects 88, 89, 92, 100, 104 and 106) due to inconsistencies in either the sampling frequency or the duration of the performed tasks. Then, we selected runs corresponding to the MI of left/right hand. Finally, we extracted trials corresponding to MI hand movements in the form of segments starting with the visual cue and lasting for 4.1 seconds.

3.3. Dataset C: Simulated Data

Simulated data were employed to validate the introduced framework with specific benchmark connectivity patterns that follow certain spatial or/and temporal modulations. Briefly, an EEG forward model was generated using the OpenMEEG package [38]. The model consists of 2004 dipoles placed over the cortex that generate electrical fields registered, on the scalp, via 64 electrodes following the 10–20 system. Projecting dipole time series to the scalp through a particular mixing matrix is a physiologically plausible model of EEG signals [39]. Dipole time series were generated as pink noise apart from few specific dipoles that were assigned activations that resulted from the band-pass filtering of white gaussian noise (so as to simulate narrowband nonstationary timeseries) [40]. The generated timeseries however may vary from case to case (e.g., based on the aspect of the introduced methodology we wish to demonstrate), and will be described with more detail within the Results section accordingly.

4. Results

Our initial and primary concern is to demonstrate that the reconstruction of the spatial covariance matrices under a common eigenspace preserves the original data structure. In addition, we dedicate two subsections, for SPD Matrix Quantization and Coding over SPD Atoms respectively, in order to showcase the implications when spatial covariance matrices that admit a common eigenspace are being employed. We note that when a classification accuracy is reported, the eigenspace (i.e., an estimation of the mixing matrix) is uncovered based on the training set only. Additionally, it should be noted that in order to depict the spatial covariance matrices in two dimensions, by means of geodesic semantic maps, we employ the classical Multi-Dimensional Scaling (cMDS) algorithm [41] combined with the AIRM induced geodesic distance unless it is otherwise stated.

4.1. Manifold Structure Preservation

Undeniably, forcing a set of matrices to be jointly diagonalizable leads to some sort of distortion and to loss of information. To quantify this loss we employ a “diagonal dominance” metric [42] which, intuitively, expresses the percentage of matrix energy that lies on its main diagonal. Therefore, after the AJD of matrices (i.e., prior to strict diagonalization and reconstruction), we calculated the ratio of the L2-norm of the diagonal elements over the matrix Frobenius norm. For the datasets A and B, the obtained average diagonal dominance values were exceeding 80% in both cases. Whereas, in the case of dataset C (as described in this Subsection), the average diagonal dominance value was approximately 88%. Beyond that, it is important to demonstrate experimentally that this process neither provides an altered view of the underlying physiological processes nor interferes with the structure of the data manifold. To this end, we exploit all of the datasets described in Section 3.

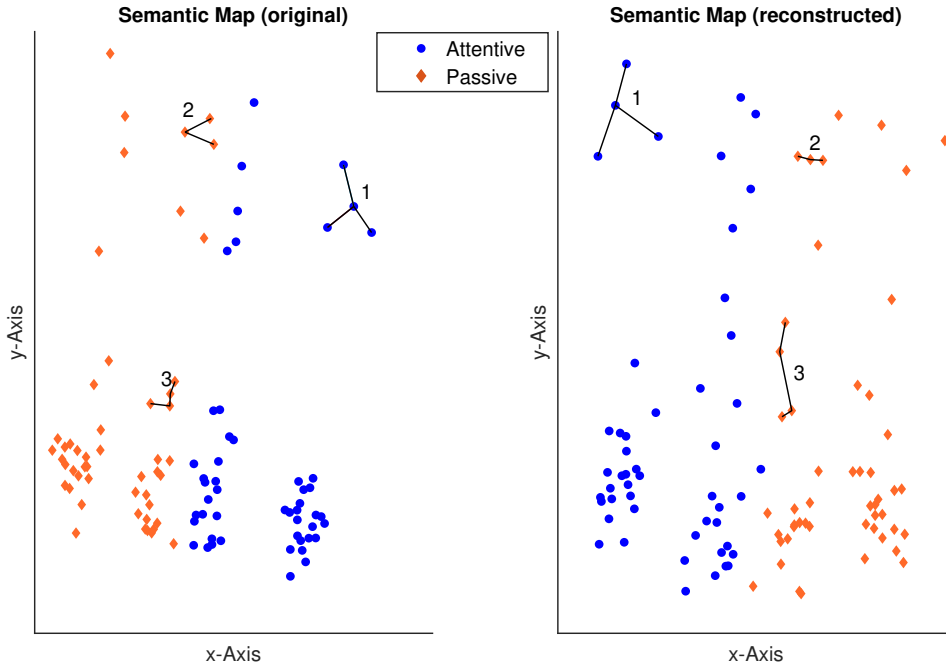


Figure 2. 2D semantic geodesic map, obtained via the cMDS algorithm, of spatial covariance matrices under the AIRM-induced geodesic distance. Each dot corresponds to a single-trial spatial covariance matrix with its color indicating the corresponding condition. The left and right map correspond to the original and the reconstructed covariance space respectively. Some particular neighborhoods are marked and indexed in order to demonstrate the differences between the two spaces.

We commence with results concerning Dataset A. Figure 2 presents an indicative 2D semantic geodesic map [43], achieved through the cMDS algorithm. Each dot corresponds to a spatial covariance matrix of a single-trial response (all belonging to the same participant) and as it can be seen, the semantic geodesic maps are similar in both

cases (x-Axis is inverted). This fact serves as an indication that despite the error that is being introduced by the reconstruction process, the structure that underpins these spatial covariance matrices remains unaltered. Moreover, some particular neighborhoods have need marked and indexed in order to demonstrate how the reconstruction process affects the data structure in more detailed manner. Specifically, neighborhood #1 is identical in both of the semantic maps. On the other hand, neighborhoods #2 and #3 have been affected (mostly exhibiting orientation and scaling deformations).

To quantify the aforementioned tendency, we adopted a structure preservation metric that is based on the k-Nearest Neighbor (kNN) [44]. In the original covariance space we calculate for each spatial covariance matrix its five nearest neighbours. This calculation is also performed on the reconstructed covariance space. The metric indicates the number of neighbours being common (their order is not being taken into account) in the two spaces. On average, over the six subjects, we calculated an 82% neighbourhood preservation. In other terms, this means that approximately, four out of the five neighbours in the original covariance space continue to belong to the 5-neighbourhood in the reconstructed covariance space. Finally, it is important to study how classification accuracy is affected by the reconstruction process. To this end, a typical Riemannian Geometry classifier is employed to measure the classification accuracy in the two spaces (original and reconstructed). Considering the passive vs attentive task we managed to obtain 87% and 89% on the original space when the chessboard appeared on the right and left side respectively. The corresponding accuracy values for the reconstructed space appeared as 91% and 92% under the same cross-validation folds. We note that no statistical significance was found in the reported classification accuracies ($p > 0.1$). With respect to the computational cost of classifying a covariance matrix, the reconstructed matrices require on average 75% less time than the original ones.

Then, we examine our approach on Dataset B, a dataset that deals with a popular BCI paradigm and contains a large amount (i.e., 103) of subjects. Previously, we presented evidence that the reconstruction process does not affect the covariance matrices that originate from a single recording session. However, the BCI domain typically deals with signals that stem from a variety of recording sessions and/or subjects. Hence, transfer learning procedures are employed in order to facilitate subject-independent schemes. Here, a different eigenspace is uncovered for each subject independently and then the reconstructed covariance matrices are aligned accordingly (refer to Subsection 2.3). It is therefore crucial to demonstrate that the original geometrical relations between the covariance matrices of different subjects are not tampered by the reconstruction procedure.

In the left panels of Figure 3 we present the 2D semantic geodesic maps, as obtained with the cMDS method, for the original and the reconstructed covariance matrices. For each subject a pair of prototypical covariance matrices has been calculated, corresponding to the Karcher mean of MI-left and MI-right trials. These prototypical covariance matrices are presented in blue and orange color respectively. Since the data have been aligned, each blue-red pair is centered at the identity matrix which is located

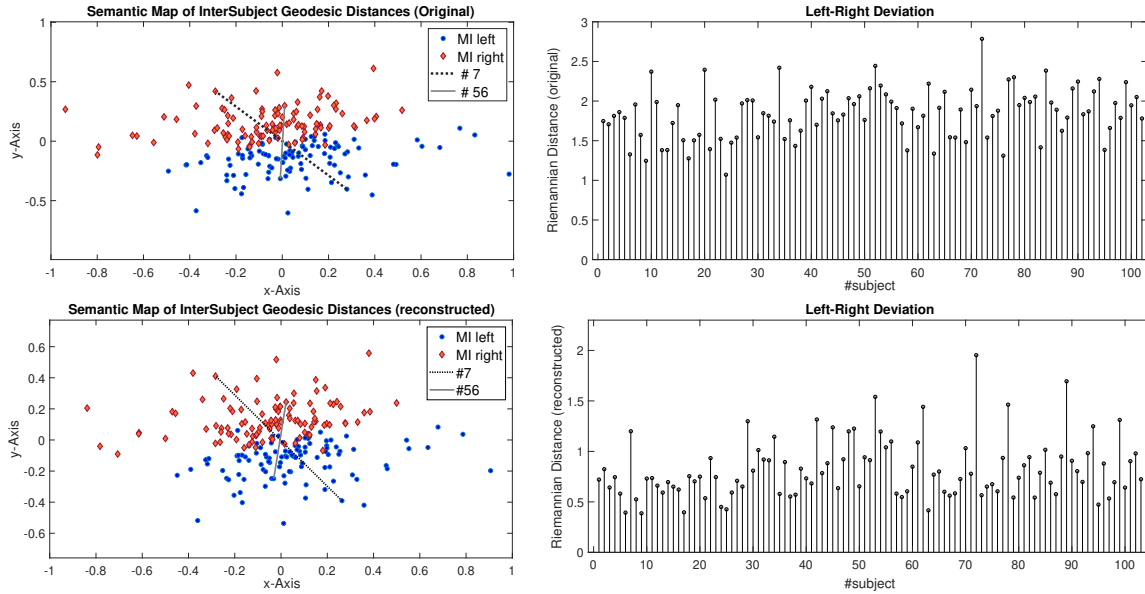


Figure 3. In the left panels shown are the semantic geodesic maps of the original (top) and reconstructed (bottom) covariance patterns corresponding to Dataset B. For every subject, two prototypical covariance patterns (left/right hand accordingly) had been mined first, and then all these patterns were compared against each-other based on AIRM induced geodesic distance. Blue/orange dots correspond to motor imagery of left/right hand, while two such pairs of prototypes coming from two randomly selected subjects are joined via lines. In the right panels, shown are the subject specific geodesic distances for the whole Dataset B. These distances quantify the differences in the two recording conditions (MI-left vs MI-right).

at $(0,0)$ in the 2D plane. As it can be seen, the distribution of prototypical covariance matrices over the two spaces (original and reconstructed) is almost identical. Moreover, two straight lines connect the MI-left/right pairs of two randomly chosen subjects. This illustration serves as a strong indication that the relation between covariance patterns belonging to different subjects is only slightly affected. Another important aspect that should be investigated in such a dataset concerns the separability of the two conditions (left vs right hand imagery movement). To this end, we employ the multivariate Wald–Wolfowitz (WW) test and calculate the corresponding score [45]. The obtained values (-11.47 and -10.76 for the original and reconstructed space respectively) are similar in both cases and indicate that data separability is well maintained. Additionally, we estimated the distances between the prototypical covariance matrices for every subject independently (i.e., in total 103 task-related deviations, shown in the rightmost panels of Figure 3). As it can be seen, the calculated distances in the two spaces exhibit similar characteristics. Ignoring the DC-offset that is being absent in the reconstructed space, the distances maintain their original trend. These measures provide experimental evidence that the reconstructed space maintains the separability of the original one. To quantify the aforementioned, the pairwise linear correlation coefficient was calculated with a value exceeding 70% value (statistically significant; $p < 0.001$).

Moreover, in terms of classification performance (subject specific; 10-fold cross validation scheme), under the Minimum Distance to Mean (MDM) classifier [6], the obtained average accuracy for the original and reconstructed covariance matrices were 62.27% and 61.97% respectively (same folds; no statistically significant difference). Concerning the computational cost, the classification time of a single covariance matrix is, again, approximately 75% less in the case of the reconstructed than the original. It is noted that during the testing phase of the MDM classification scheme, the geodesic distance of a new test sample from the prototypical centers (each corresponding to the geodesic mean of each class) is calculated and the label of the closest prototype is assigned to the testing sample. Hence, the reported reduction in the computational cost incorporates the gain from the estimation of two geodesic distances (of test sample to either of the prototypes) in an efficient way.

Finally, we examine the validity of our method on simulated EEG data. To this end, we generated simulated data that belong to two distinct classes, namely class A and class B. Initially, we selected a kernel dipole for each class and assigned an alpha-band activation signal, by means of filtering a white noise signal. Then, to its 30 closest neighbours we assigned a similar activation pattern with the distance modulating the strength of the activation (i.e., the strength was fading while moving away from the kernel dipole). Pink-noise activations were assigned to remaining dipoles. This simulated data generation strategy allows us to examine our approach when a particular brain region is governed by a non-diagonal covariation pattern (i.e., source signals of this region are not independent). We note that in class A the kernel dipole was located in the left occipital brain region whereas in class B it was located in the right occipital brain region. The neurophysiological characteristics of the simulated single-trial signals are depicted in Figure 4 (panels A, B and D). As it can be seen in panel C, where the 2D semantic geodesic map of the reconstructed single-trial covariance matrices (as derived from 1.5 seconds long signals) is depicted, the reconstruction process clearly preserves the separability of data.

4.2. SPD Matrix Quantization

To demonstrate the new possibilities opened by the reconstruction of covariance matrices under a common eigenspace, we apply the method presented on Subsection 2.3.1. Using the Riemannian k-Means method over a set of spatial covariance matrices, we initially uncover a codebook. Since we aim to employ this method in order to perform connectivity-based microstate analysis, for each single-trial a series of covariance patterns is calculated. Then, each of these patterns is assigned to its closest prototype (which is part of the codebook) and in this way the single-trial covariation dynamics are represented by means of a symbolic timeseries.

In order to examine the efficiency of this application, we exploit the recordings of Dataset A. Our aim here, is to demonstrate that the concept of SPD Matrix Quantization is capable to uncover the differences between the passive and the attentive

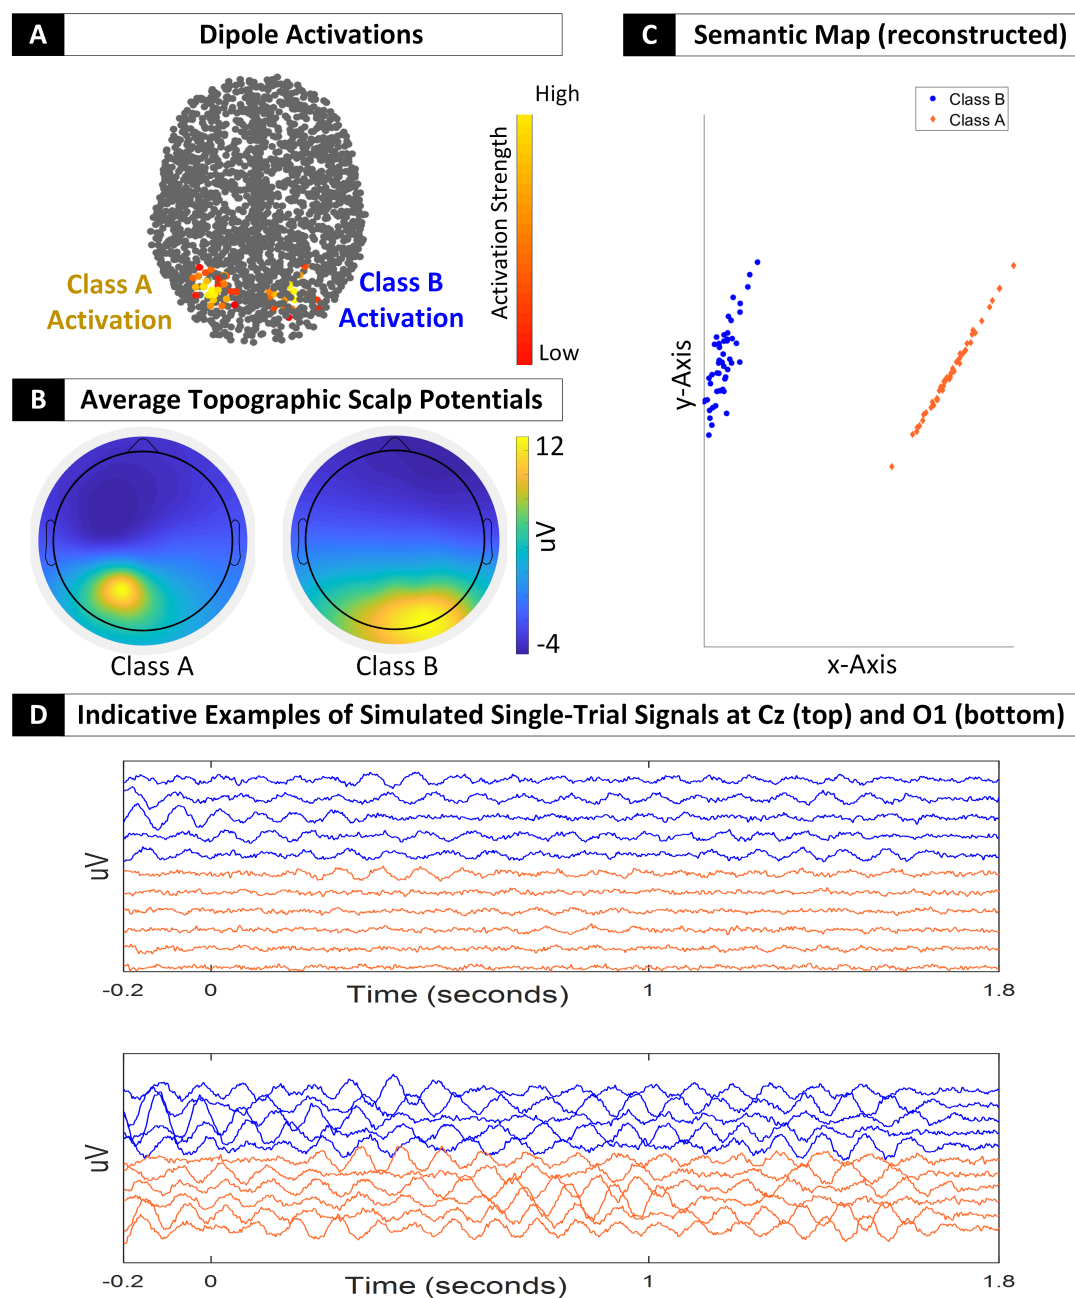


Figure 4. A) Dipole activations for the two classes of simulated single-trial signals. In the first class an activation pattern appeared on the left occipital dipoles whereas in the second on the right. Colorcoding (from red to yellow) indicates the strength of the activation pattern. Dipoles colored in gray correspond to pink noise. B) Topographic scalp potentials for the two classes. They were obtained by calculating the leading eigenvector for the average covariance matrix of each class. C) 2D semantic geodesic map of the Single-trial covariance patterns. Blue and orange color indicate covariance matrices of Class A and B respectively. D) Indicative examples of the simulated signals at two sensors (Cz and O1). Blue and orange color indicate single-trial signals of Class A and B respectively.

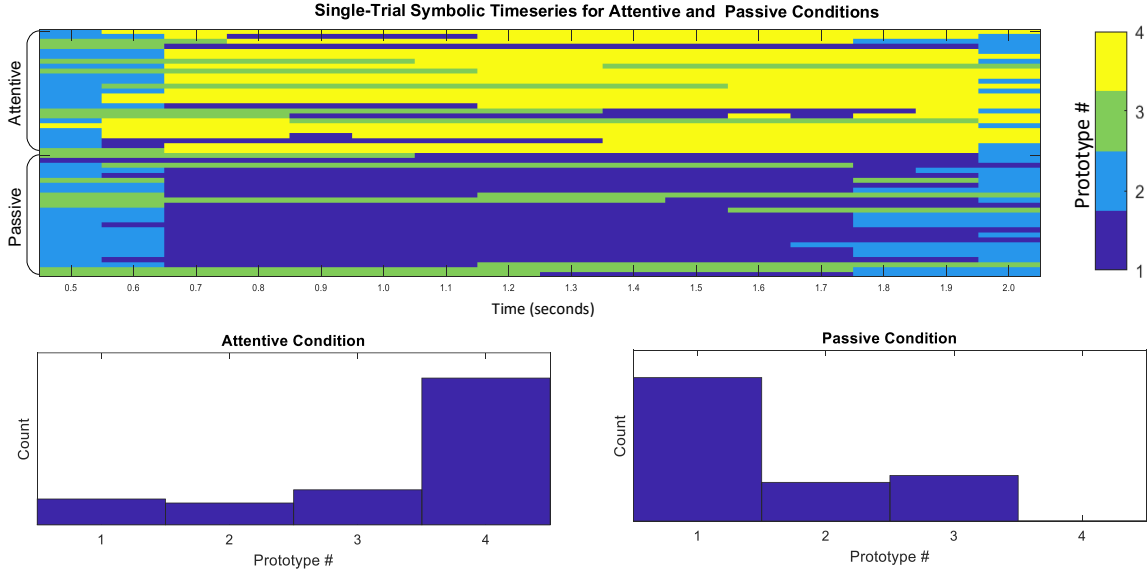


Figure 5. In the top panel, we present the symbolic timeseries for single-subject trials belonging to Dataset A. Each color represents a different prototype. In the bottom panel, the distribution of each condition over the four prototypes is illustrated.

conditions. The time-varying covariance patterns were calculated with a step size equal to 0.1 seconds and 1 second long window size. Then, the Riemannian k-Means algorithm [29] was employed in order to uncover the four (as decided in preliminary analysis) prototypical matrices that formed the codebook. Figure 5 depicts the symbolic times-series for the passive and attentive conditions that belong to the same, indicative, subject. Moreover, we present the aggregated distribution of the symbolic-timeseries over the codebook. As it can be seen, the prototypes #1 and #4 are capable of discriminating the attentive and passive single trials. To quantify the aforementioned, we were able to achieve an average classification accuracy of 93% and 94% when the wall appeared on the right and left side respectively. The reported accuracy values are based on a 10-fold cross validation scheme, using linear SVMs, in a subject-specific manner. Although similar classification results were achieved when the original covariance matrices were employed (inferior by 1%, but without a statistically significant difference), it should be noted that the classification time of a reconstructed covariance matrix requires approximately 95% less time compared to an original one.

To further demonstrate this approach we generated 100 artificial signals, 3 seconds long each, according to the model described in Subsection 3.3. All dipole signals were generating pink noise, apart from one that was generating nonstationary alpha-band activity. Half of the simulated trials were characterized by their alpha-band energy which was modulated by a sigmoid function (i.e., increasing energy) whereas the other half had their alpha-band modulated by an inverse sigmoid function (i.e., decreasing energy in alpha band over time). Again, the time-varying covariance matrices were calculated with a step size equal to 0.1 seconds and 1 second long window size. The key concept

here is to illustrate that SPD matrix quantization on the time-varying reconstructed matrices is capable to detect such subtle, time-varying, changes. Figure 6 depicts the colorcoded symbolic times-series for the two aforementioned conditions alongside with their corresponding distributions over the codebook. As expected, the distribution diagrams do not exhibit any significant separability for the two conditions, since, the two conditions differ in the order of appearance of the prototypes. However, the presented quantization, illustrated in the form of colorcoded timeseries, follows the trend of the source signals (i.e., energy modulated by a sigmoid function). This serves as a clear evidence that the introduced reconstruction process does not tamper the underlying, simple yet subtle, physiological processes.

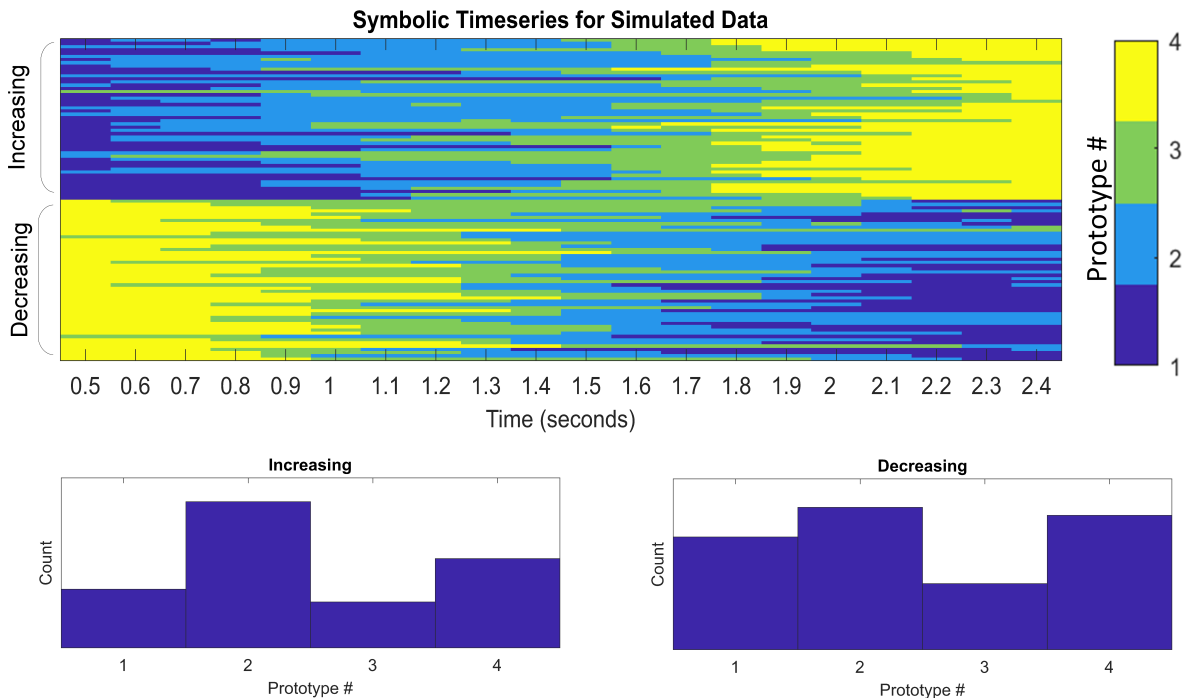


Figure 6. On the top side of the figure we present the symbolic timeseries for the simulated EEG data. Each color represents a different prototype. On the bottom the distribution of each condition over the four prototypes is illustrated.

4.3. Coding over SPD Atoms

In this section we demonstrate the employment of Coding over SPD Atoms on Dataset A. Our effort is mostly dedicated to demonstrating how the reconstruction of covariance under a common eigenspace paves the way for efficient and easily conceivable neuroscientific explorations that lie on the principles of Riemannian Geometry. The key concept here is to approximate each covariance matrix as a supersposition of covariance atoms. Hence, each covariance matrix can be now represented solely by its membership (i.e., the non-negative weights) to the covariance atoms. Therefore, the number of atoms characterizes the dimension of the new, weight space. Although the

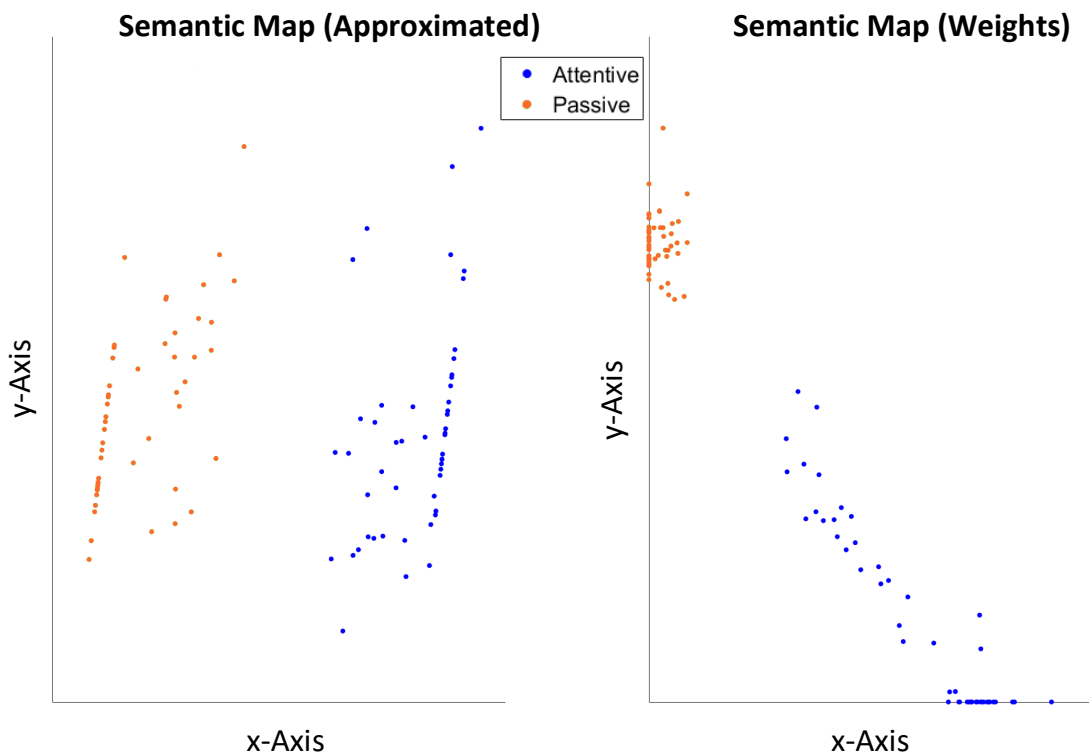


Figure 7. 2D semantic geodesic map of the approximated covariance matrices on the left and right respectively. On the right, the semantic map of the weight space is depicted under the Euclidean distance. Since only two atoms were employed in this indicative example, the weight space illustration depicts the obtained approximation weights unaltered. Blue dots correspond to the attentive condition whereas orange to the passive condition of a randomly selected participant.

process of uncovering suitable atoms is very important for neuroscientific explorations, in the presented case the atoms are uncovered by exploiting the Riemannian k-Means algorithm.

Figure 7 depicts the 2D semantic geodesic map of the approximated covariance matrices alongside with the 2D semantic map of the weight space under the Euclidean distance. As it can be seen, the 2D semantic geodesic map of the approximated covariance matrices is a highly distorted version of the reconstructed one. This is due to the fact that only two covariance atoms were used for the approximation. However, the approximated covariance matrices exhibit clearer separability characteristics than the reconstructed ones. Moreover, the inspection of the weight space allows us to infer that trials belonging to the passive condition, mostly located close to the y-Axis, are clearly characterized by their membership to the second atom. On the other hand, attentive trials cover a more wide spectrum of membership over both of the atoms. We should note, that in both conditions there are several trials (those abiding on the axes) that can be characterized solely by their membership to one of the uncovered

atoms. Despite that such interpretations are bound to the method that was employed to uncover the atoms, they clearly demonstrate the implications of this approach on neuroscientific explorations. In practical terms, using linear SVMs on the weight space which is Euclidean and under a 10-fold cross validation scheme, we achieved an average classification accuracy (over the six subjects) of 97% and 96% when the chessboard appeared on the right and left side respectively. These results are on par with state-of-the-art methodologies [34].

5. Discussion

In this work we presented a methodology that allows the reconstruction of spatial covariance matrices under the assumption of a common eigenspace. The vast majority of our effort is dedicated to showing that the introduced reconstruction process does not tamper the structure of the original covariance manifold. Given its validity, our research aims to bridge the gap between Riemannian Geometry and neuroscience. Indeed, Riemannian Geometry has been proved to be an invaluable asset in EEG signal analysis (e.g., [46]), however it requires rigorous mathematical procedures that may hinder its wider adoption. Moreover, this problem becomes more acute given that classical pattern recognition and machine learning schemes cannot be readily employed in Riemannian schemes.

In many cases, EEG decoding algorithms are agnostic with respect to the physiological processes that accompany the brain signals (e.g., their origin or the information they may carry) [47, 48]. Here, we aim to offer a valuable insight on the connection of the neural mechanisms that underpin the EEG signals in the context of Riemannian geometry. Although AJD has been previously exploited in this context (particularly for blind source separation [49, 50]), the pervasive and elusive EEG variability among different individuals have remained, among others, insufficiently explored. Towards this end, we established, that under the assumption of common mixing matrix, one of the most typical transfer learning approaches has a strong neurological basis, that of normalizing the energy of the source signals. Ultimately we believe that our work may further foster neuroscientifically-aware brain decoding algorithms.

In the presented work we show that when the AIRM is employed to deal with spatial covariance matrices, one may reconstruct (exploiting the concept of AJD) these matrices under a common eigenspace and equivalently work with the corresponding eigenvalues. Such a reconstruction offers the advantage of vectorization, where the spatial covariance matrices can be completely characterized by their diagonalized form and consequently be treated as vectors. However, in order to maintain the original structure of the matrices a reformulation of the AIRM and its induced geodesic distance should be employed for vectors, as described in Section 2.1.

In addition to the simplifications that the presented methodology brings, it also comes with a second advantage; that of computational efficiency. As a matter of

fact inner products and distance calculations under the AIRM have a computational complexity of $\mathcal{O}(E^3)$. However, working with the equivalent diagonalized forms the corresponding computational complexity reduces to $\mathcal{O}(E)$. This is a large improvement, especially when dealing with dense EEG recordings or incorporating Riemannian Geometry schemes in real-time applications. However, the reduction in computational complexity (e.g., in terms of time cost) can be observed only during the testing phase or when complicated manipulations are required (e.g., in Riemannian geometry-based microstate analysis) subsequent to the diagonalization process. In cases where the classification task is applied only on a few samples (compared to the training set), the reduced testing time cannot compensate for the increased cost induced by AJD. Hence, it becomes obvious that AJD, which is a computationally demanding operation, may cancel out any computational gain achieved by the proposed reconstruction process. A straightforward solution to this shortcoming, concerns the calculation of the common eigenspace only on a small sample of data. This will allow the proposed approach to generalize, at the cost of a less accurate mixing matrix estimation, effectively for the entire dataset. Moreover, the AJD concept can be incorporated as a calibration procedure, since it has to be performed only once, leading to computationally efficient online applications.

The introduced methodology is based on physiologically plausible assumptions that underpin the EEG recordings. Our major assumption concerns the existence of a constant mixing matrix. Despite that this idea is supported by a wide variety of literature, we experimentally show that a “physiologically ideal” unmixing matrix indeed exists within each recording session. Since the scalp position is stable with respect to the electrodes, it is sound to assume that there should exist a matrix which can approximately and jointly diagonalize the independent physiological sources for all the trials within each session.

One important aspect that has not been discussed particularly in the manuscript concerns the dimensionality of the mixing matrix and its mathematical connections with the employed Riemannian Geometry schemes. In a strict mathematical sense, the AIRM-induced geodesic distance is invariant to congruence when the matrix is full-rank. Therefore, equivalence between the sensor space and the source space can be achieved when the number of sources and the number of electrodes are equal (i.e., the mixing matrix is a non-singular square matrix). In practice however, Riemannian Geometry has offered significant achievements in EEG signal processing and is still an evolving research field. As discussed in [6], Riemannian Geometry in the EEG field is more robust when the number of available sensors is sufficiently enough. This, still unexplored, aspect may limit the implications of both Riemannian Geometry generally and the introduced methodology particularly.

As mentioned earlier, bridging the gap between rigorous mathematical procedures and computational neuroscience is among the most profound contributions of this work. Hence, it was within our scope to show the implications of the introduced methodology. To this end, we reformulated and re-introduced two approaches, namely SPD Matrix

Quantization and Coding over SPD Atoms, in order to showcase the benefits that the reconstruction of spatial covariance matrices under a common eigenspace brings. One important aspect in these two approaches concerns the codebook and atom learning respectively. In both cases we employed a naive algorithm, the Riemannian k-Means. This of course may confine the reported results but it was beyond the scope of our work to propose new algorithms for unsupervised learning. Our objective here was to demonstrate the capabilities that open via the introduced reconstruction procedure. We believe, that the presented work may pave the way for robust and efficient neuroscientific explorations that abide to Riemannian Geometry.

6. Acknowledgments

This work was supported by the NeuroMkt project, co-financed by the European Regional Development Fund of the EU and Greek National Funds through the Operational Program “Competitiveness, Entrepreneurship and Innovation”, under RESEARCH CREATE INNOVATE (T2EDK-03661).

References

- [1] K. Glomb, J. Cabral, A. Cattani, A. Mazzoni, A. Raj, B. Franceschiello, Computational models in electroencephalography, *Brain topography* (2021) 1–20.
- [2] J. D. Cohen, N. Daw, B. Engelhardt, U. Hasson, K. Li, Y. Niv, K. A. Norman, J. Pillow, P. J. Ramadge, N. B. Turk-Browne, et al., Computational approaches to fmri analysis, *Nature neuroscience* 20 (3) (2017) 304–313.
- [3] C. H. Nguyen, P. Artemiadis, Eeg feature descriptors and discriminant analysis under riemannian manifold perspective, *Neurocomputing* 275 (2018) 1871–1883.
- [4] M. Shahbazi, A. Shirali, H. Aghajan, H. Nili, Using distance on the riemannian manifold to compare representations in brain and in models, *NeuroImage* 239 (2021) 118271.
- [5] D. Sabbagh, P. Ablin, G. Varoquaux, A. Gramfort, D. A. Engemann, Predictive regression modeling with meg/eeg: from source power to signals and cognitive states, *NeuroImage* 222 (2020) 116893.
- [6] M. Congedo, A. Barachant, R. Bhatia, Riemannian geometry for eeg-based brain-computer interfaces; a primer and a review, *Brain-Computer Interfaces* 4 (3) (2017) 155–174.
- [7] X. Pennec, P. Fillard, N. Ayache, A riemannian framework for tensor computing, *International Journal of computer vision* 66 (1) (2006) 41–66.
- [8] R. Vemulapalli, D. Jacobs, Riemannian metric learning for symmetric positive definite matrices, *arXiv preprint arXiv:1501.02393* (2015).
- [9] V. Arsigny, P. Fillard, X. Pennec, N. Ayache, Log-euclidean metrics for fast and simple calculus on diffusion tensors, *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 56 (2) (2006) 411–421.
- [10] X. Navarro-Sune, A. Hudson, F. D. V. Fallani, J. Martinerie, A. Witon, P. Pouget, M. Raux, T. Similowski, M. Chavez, Riemannian geometry applied to detection of respiratory states from eeg signals: The basis for a brain–ventilator interface, *IEEE Transactions on Biomedical Engineering* 64 (5) (2017) 1138–1148.
- [11] F. P. Kalaganis, N. A. Laskaris, E. Chatzilari, S. Nikolopoulos, I. Kompatsiaris, A riemannian geometry approach to reduced and discriminative covariance estimation in brain computer interfaces, *IEEE Transactions on Biomedical Engineering* 67 (1) (2019) 245–255.
- [12] A. Barachant, S. Bonnet, M. Congedo, C. Jutten, Multiclass brain–computer interface classification by riemannian geometry, *IEEE Transactions on Biomedical Engineering* 59 (4) (2012) 920–928.
- [13] I. L. Dryden, A. Koloydenko, D. Zhou, Non-euclidean statistics for covariance matrices, with applications to diffusion tensor imaging, *The Annals of Applied Statistics* 3 (3) (2009) 1102–1123.
- [14] R. Vigário, J. Sarela, V. Jousmiki, M. Hamalainen, E. Oja, Independent component approach to the analysis of eeg and meg recordings, *IEEE transactions on biomedical engineering* 47 (5) (2000) 589–593.
- [15] J. Kybic, M. Clerc, O. Faugeras, R. Keriven, T. Papadopoulo, Generalized head models for meg/eeg: boundary element method beyond nested volumes, *Physics in Medicine & Biology* 51 (5) (2006) 1333.
- [16] A. Delorme, J. Palmer, J. Onton, R. Oostenveld, S. Makeig, Independent eeg sources are dipolar, *PloS one* 7 (2) (2012) e30135.
- [17] J. J. Ermer, J. C. Mosher, S. Baillet, R. M. Leahy, Rapidly recomputable eeg forward models for realistic head shapes, *Physics in Medicine & Biology* 46 (4) (2001) 1265.
- [18] A. Belouchrani, K. Abed-Meraim, J.-F. Cardoso, E. Moulines, A blind source separation technique using second-order statistics, *IEEE Transactions on signal processing* 45 (2) (1997) 434–444.
- [19] A. Cichocki, P. Georgiev, Blind source separation algorithms with matrix constraints, *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences* 86 (3) (2003) 522–531.
- [20] J.-F. Cardoso, A. Souloumiac, Jacobi angles for simultaneous diagonalization, *SIAM journal on*

- matrix analysis and applications 17 (1) (1996) 161–164.
- [21] M. Congedo, B. Afsari, A. Barachant, M. Moakher, Approximate joint diagonalization and geometric mean of symmetric positive definite matrices, *PloS one* 10 (4) (2015) e0121423.
 - [22] C. M. Michel, T. Koenig, Eeg microstates as a tool for studying the temporal dynamics of whole-brain neuronal networks: a review, *Neuroimage* 180 (2018) 577–593.
 - [23] W. Förstner, B. Moonen, A metric for covariance matrices, in: *Geodesy-The Challenge of the 3rd Millennium*, Springer, 2003, pp. 299–309.
 - [24] K.-R. Müller, R. Vigario, F. Meinecke, A. Ziehe, Blind source separation techniques for decomposing event-related brain signals, *International Journal of Bifurcation and Chaos* 14 (02) (2004) 773–791.
 - [25] A. Hyvärinen, Survey on independent component analysis (1999).
 - [26] A. Ziehe, Blind source separation based on joint diagonalization of matrices with applications in biomedical signal processing, Ph.D. thesis, Universität Potsdam (2005).
 - [27] A. Yeredor, A. Ziehe, K.-R. Müller, Approximate joint diagonalization using a natural gradient approach, in: *International Conference on Independent Component Analysis and Signal Separation*, Springer, 2004, pp. 89–96.
 - [28] P. Zanini, M. Congedo, C. Jutten, S. Said, Y. Berthoumieu, Transfer learning: a riemannian geometry framework with applications to brain–computer interfaces, *IEEE Transactions on Biomedical Engineering* 65 (5) (2018) 1107–1116.
 - [29] K. You, H.-J. Park, Re-visiting riemannian geometry of symmetric positive definite matrices for the analysis of functional connectivity, *NeuroImage* 225 (2021) 117464.
 - [30] R. H. Byrd, M. E. Hribar, J. Nocedal, An interior point algorithm for large-scale nonlinear programming, *SIAM Journal on Optimization* 9 (4) (1999) 877–900.
 - [31] R. H. Byrd, J. C. Gilbert, J. Nocedal, A trust region method based on interior point techniques for nonlinear programming, *Mathematical programming* 89 (1) (2000) 149–185.
 - [32] A. Cherian, S. Sra, Riemannian dictionary learning and sparse coding for positive definite matrices, *IEEE transactions on neural networks and learning systems* 28 (12) (2016) 2859–2871.
 - [33] D. Liparas, S. I. Dimitriadis, N. A. Laskaris, A. Tzelepi, K. Charalambous, L. Angelis, Exploiting the temporal patterning of transient vep signals: A statistical single-trial methodology with implications to brain–computer interfaces (bcis), *Journal of neuroscience methods* 232 (2014) 189–198.
 - [34] F. P. Kalaganis, N. A. Laskaris, E. Chatzilari, D. A. Adamos, S. Nikolopoulos, I. Kompatsiaris, A complex-valued functional brain connectivity descriptor amenable to riemannian geometry, *Journal of Neural Engineering* 17 (2) (2020) 024001.
 - [35] J. J. Foxe, A. C. Snyder, The role of alpha-band brain oscillations as a sensory suppression mechanism during selective attention, *Frontiers in psychology* 2 (2011) 154.
 - [36] W. Klimesch, Alpha-band oscillations, attention, and controlled access to stored information, *Trends in cognitive sciences* 16 (12) (2012) 606–617.
 - [37] G. Schalk, D. J. McFarland, T. Hinterberger, N. Birbaumer, J. R. Wolpaw, Bci2000: a general-purpose brain-computer interface (bci) system, *IEEE Transactions on biomedical engineering* 51 (6) (2004) 1034–1043.
 - [38] A. Gramfort, T. Papadopoulo, E. Olivi, M. Clerc, Openmeeg: opensource software for quasistatic bioelectromagnetics, *Biomedical engineering online* 9 (1) (2010) 1–20.
 - [39] A. Gramfort, T. Papadopoulo, E. Olivi, M. Clerc, Forward field computation with openmeeg, *Computational intelligence and neuroscience* 2011 (2011).
 - [40] M. X. Cohen, A data-driven method to identify frequency boundaries in multichannel electrophysiology data, *Journal of Neuroscience Methods* 347 (2021) 108949.
 - [41] A. Mead, Review of the development of multidimensional scaling methods, *Journal of the Royal Statistical Society: Series D (The Statistician)* 41 (1) (1992) 27–39.
 - [42] J. Xu, M. Grosse-Wentrup, V. Jayaram, Tangent space spatial filters for interpretable and efficient riemannian classification, *Journal of Neural Engineering* 17 (2) (2020) 026043.

- [43] N. Laskaris, A. Ioannides, Semantic geodesic maps: a unifying geometrical approach for studying the structure and dynamics of single trial evoked responses, *Clinical Neurophysiology* 113 (8) (2002) 1209–1226.
- [44] H.-U. Bauer, K. Pawelzik, T. Geisel, A topographic product for the optimization of self-organizing feature maps, *Advances in Neural Information Processing Systems* 4 (1991).
- [45] N. A. Laskaris, A. A. Ioannides, Exploratory data analysis of evoked response single trials based on minimal spanning tree, *Clinical Neurophysiology* 112 (4) (2001) 698–712.
- [46] K. Georgiadis, F. P. Kalaganis, V. P. Oikonomou, S. Nikolopoulos, N. A. Laskaris, I. Kompatsiaris, Rneumark: A riemannian eeg analysis framework for neuromarketing (2022).
- [47] M. Xu, F. He, T.-P. Jung, X. Gu, D. Ming, Current challenges for the practical application of electroencephalography-based brain–computer interfaces, *Engineering* 7 (12) (2021) 1710–1712.
- [48] S. Khademi, M. Neghabi, M. Farahi, M. Shirzadi, H. R. Marateb, A comprehensive review of the movement imaginary brain-computer interface methods: Challenges and future directions, *Artificial Intelligence-Based Brain-Computer Interface* (2022) 23–74.
- [49] M. Congedo, C. Gouy-Pailler, C. Jutten, On the blind source separation of human electroencephalogram by approximate joint diagonalization of second order statistics, *Clinical Neurophysiology* 119 (12) (2008) 2677–2686.
- [50] F. Bouchard, J. Malick, M. Congedo, Riemannian optimization and approximate joint diagonalization for blind source separation, *IEEE Transactions on Signal Processing* 66 (8) (2018) 2041–2054.