



**HAL**  
open science

# Weakly supervised covariance matrices alignment through Stiefel matrices estimation for MEG applications

Antoine Collas, Rémi Flamary, Alexandre Gramfort

► **To cite this version:**

Antoine Collas, Rémi Flamary, Alexandre Gramfort. Weakly supervised covariance matrices alignment through Stiefel matrices estimation for MEG applications. 2024. hal-04422094

**HAL Id: hal-04422094**

**<https://hal.science/hal-04422094v1>**

Preprint submitted on 28 Jan 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Weakly supervised covariance matrices alignment through Stiefel matrices estimation for MEG applications

Antoine Collas

Université Paris-Saclay, Inria, CEA

Rémi Flamary

Ecole Polytechnique, Institut Polytechnique de Paris, CMAP, UMR 7641

Alexandre Gramfort

Université Paris-Saclay, Inria, CEA

## Abstract

This paper introduces a novel domain adaptation technique for time series data, called Mixing model Stiefel Adaptation (MSA), specifically addressing the challenge of limited labeled signals in the target dataset. Leveraging a domain-dependent mixing model and the optimal transport domain adaptation assumption, we exploit abundant unlabeled data in the target domain to ensure effective prediction by establishing pairwise correspondence with equivalent signal variances between domains. Theoretical foundations are laid for identifying crucial Stiefel matrices, essential for recovering underlying signal variances from a Riemannian representation of observed signal covariances. We propose an integrated cost function that simultaneously learns these matrices, pairwise domain relationships, and a predictor, classifier, or regressor, depending on the task. Applied to neuroscience problems, MSA outperforms recent methods in brain-age regression with task variations using magnetoencephalography (MEG) signals from the Cam-CAN dataset.

## 1 Introduction

Multivariate time series data is ubiquitous in a wide range of applications, such as remote sensing, finance, and neuroscience. In numerous scenarios, observed time series result from linear combinations of underlying signals of interest. In remote sensing, measured spectral bands are linear mixtures of endmember spectra (e.g., water, grass, wood) with coefficients corresponding to the proportions of these endmembers within the observed scene [27]. In neuroscience, a suitable generative model for magnetoencephalography (MEG) and electroencephalography (EEG) data is the linear instantaneous mixing model, leveraging the inherent linearity of Maxwell's equations [23].




Figure 1: **Illustration of MSA for a regression task.** Source and target covariance matrices,  $\Sigma_i^S$  and  $\Sigma_i^T$ , exhibit different patterns in their original spaces. Embedding them into  $x_i^S$  and  $x_i^T$  and then jointly learning optimal transport plan  $\pi$  and orthogonal bases  $U^S$  and  $U^T$  alleviate this problem by finding components that matter for prediction.

Covariance matrices have emerged as powerful descriptors for diverse data types, offering valuable insights into the underlying relationships within multivariate data. The application of Riemannian geometry to covariance matrices [42] has paved the way for innovative solutions in tasks such as classification [45, 21] and regression [40, 41, 7]. These models have been successfully employed in the analysis of multivariate time series data, including applications in remote sensing [10] and M/EEG [3, 47, 28] applications. Notably, recent studies have revealed the competitiveness of shallow covariance-based models when compared to deep learning techniques in neuroscience [14].

However, the challenge of domain adaptation poses a substantial hurdle in the analysis of multivariate time series data. For example, when data are measured using different sensors or at different periods, domain shifts occur, making it difficult to apply classifiers or regressors across domains directly. Domain adaptation addresses these shifts by aligning data [11, 30]. This challenge is particularly pronounced in M/EEG, where changes in recording devices, head morphology, experimental protocol, or data collection sites can induce shifts. It is especially crucial in applications like brain-computer interfaces (BCI) [25, 48, 31, 39, 6] and brain age prediction [33], where limited labeled data in the target domain necessitates effective adaptation techniques.

In this paper, we propose Mixing model Stiefel Adaptation (MSA), a novel domain adaptation approach for time series data under mixing models. We address the challenging problem of adapting models when the target domain has limited labeled data while a fully labeled source domain is available. Our approach involves the identification of two Stiefel matrices that are used to align the covariances of both source and target data. When applied to a Riemannian representation of the covariances of observed signals, these matrices facilitate the recovery of underlying signal variances. To achieve this, MSA simultaneously learns these matrices and a predictor. The proposed framework is general enough to be applied to both classification and regression tasks.

In the subsequent sections, we provide a detailed exposition of our novel approach. We show that MSA outperforms recent approaches in brain-age regression using magnetoencephalography (MEG) signals from the Cam-CAN dataset [44] while accommodating task variations.

## 2 Domain shift and covariance matrices under mixing models

This section presents time series under mixing models in the context of domain adaptation with mixing matrices that differ from one domain to another. These data are summarized with covariance matrices as discussed in the introduction. Then, the ‘‘optimal transport domain adaptation (OTDA) assumption’’ is introduced: we assume that, for large enough domains, there exists a permutation matrix that associates source variances with target variances. Furthermore, the main tools of Riemannian geometry associated with covariance matrices are presented.

### 2.1 Mixing models in a domain adaption context

Domain adaptation is a core machine learning technique used to enhance model performance when there is a distribution shift between the source and target domains denoted  $\mathcal{S}$  and  $\mathcal{T}$ , respectively. When the source domain provides labeled data for training, and only a fraction of the target domain is labeled, the problem is known as semi-supervised domain adaptation.

**Domain specific mixing model** We consider the covariance matrices of the signal as features for regression and classification. We assume that the signal is the sum of a mixed signal of interest plus an independent noise. Hence, for any domain  $\mathcal{D} \in \{\mathcal{S}, \mathcal{T}\}$ , the covariance matrix of the  $i^{\text{th}}$  signal is

$$\Sigma_i^{\mathcal{D}} \triangleq \mathbf{A}^{\mathcal{D}} \begin{bmatrix} \text{diag}(\mathbf{p}_i^{\mathcal{D}}) & \mathbf{0}_{q \times (p-q)} \\ \mathbf{0}_{(p-q) \times q} & \mathbf{N}_i^{\mathcal{D}} \end{bmatrix} (\mathbf{A}^{\mathcal{D}})^{\top} \in \mathbb{R}^{p \times p} \quad (1)$$

where  $\mathbf{A}^{\mathcal{D}}$  is the domain-specific mixing matrix,  $\mathbf{p}_i^{\mathcal{D}} = [p_{i,1}^{\mathcal{D}}, \dots, p_{i,q}^{\mathcal{D}}]^{\top} \in \mathbb{R}^q$  contains the variances of the unmixed signal of interest,  $\mathbf{0}$  is the zero matrix and  $\mathbf{N}_i^{\mathcal{D}} \in \mathbb{R}^{(p-q) \times (p-q)}$  is the covariance matrix of the noise.

### 2.2 Covariance matrices in the Riemannian geometry framework

We utilize the Riemannian geometry of symmetric positive definite matrices for both regressing outcomes and classifying covariance matrices. This approach is theoretically grounded, as supported by studies such as [42, 37]. In practical applications, the use of Riemannian geometry for covariance matrices demonstrates strong performance, notably in EEG [47, 7] and MEG [14, 33] applications.

**Riemannian geometry of  $\mathbb{S}_p^{++}$**  The covariance matrices belong to the set of  $p \times p$  symmetric positive definite matrices denoted  $\mathbb{S}_p^{++}$ . The latter is open in the set of  $p \times p$  symmetric matrices denoted  $\mathbb{S}_p$ , and thus is a smooth manifold [8]. A vector space is defined at each  $\Sigma \in \mathbb{S}_p^{++}$ , called the tangent space and denoted  $T_{\Sigma}\mathbb{S}_p^{++}$ . It is equal to  $\mathbb{S}_p$ , the ambient space. Equipped with a smooth inner product, a smooth manifold becomes a Riemannian manifold. To do so, we leverage the affine invariant Riemannian metric [42]. Indeed, it gives good theoretical properties to the  $\mathbb{S}_p^{++}$  manifold such as being geodesically complete and works well in practice [3, 47, 7]. Given

$\Gamma, \Gamma' \in \mathbb{S}_p$ , this metric writes  $\langle \Gamma, \Gamma' \rangle_{\Sigma} = \text{Tr}(\Sigma^{-1} \Gamma \Sigma^{-1} \Gamma')$  and the corresponding norm is  $\|\Gamma\|_{\Sigma} = \sqrt{\langle \Gamma, \Gamma \rangle_{\Sigma}}$ .

**Riemannian mean on  $\mathbb{S}_p^{++}$**  The Riemannian distance (or geodesic distance) associated with the affine invariant metric is  $d_{\mathbb{S}_p^{++}}(\Sigma, \Sigma') = \|\log(\Sigma^{-1/2} \Sigma' \Sigma^{-1/2})\|_F$  with  $\log : \mathbb{S}_p^{++} \rightarrow \mathbb{S}_p$  being the matrix logarithm. It should be noted that, since the Riemannian metric is affine invariant,  $d_{\mathbb{S}_p^{++}}$  is also affine invariant, i.e. for every  $\mathbf{A} \in \mathbb{R}^{p \times p}$  invertible, we have  $d_{\mathbb{S}_p^{++}}(\mathbf{A} \Sigma \mathbf{A}^{\top}, \mathbf{A} \Sigma' \mathbf{A}^{\top}) = d_{\mathbb{S}_p^{++}}(\Sigma, \Sigma')$ . In the following, we vectorize covariance matrices  $\Sigma_i^{\mathcal{D}}$  in a tangent space  $T_{\bar{\Sigma}} \mathbb{S}_p^{++}$ . A classical choice for the base point is the Riemannian mean  $\bar{\Sigma}$ . This mean, denoted as  $\bar{\Sigma}$ , is defined for a set  $\Sigma_1, \dots, \Sigma_n \in \mathbb{S}_p^{++}$  as

$$\bar{\Sigma} \triangleq \arg \min_{\Sigma \in \mathbb{S}_p^{++}} \sum_{i=1}^n d_{\mathbb{S}_p^{++}}(\Sigma, \Sigma_i)^2 \quad (2)$$

and is computed with a Riemannian gradient descent [37, 49].

**Linearization of the  $\mathbb{S}_p^{++}$  manifold** The affine invariant metric induces the Riemannian exponential  $\exp_{\Sigma} : \mathbb{S}_p \rightarrow \mathbb{S}_p^{++}$  which is such that, for  $t \in [0, 1]$ ,  $\exp_{\Sigma}(t\Gamma)$  is the geodesic with initial position  $\Sigma$  and speed  $\Gamma$ . The inverse operator is the Riemannian logarithm  $\log_{\Sigma} : \mathbb{S}_p^{++} \rightarrow \mathbb{S}_p$ . This operator is preponderant in the rest of the paper since it transforms covariance matrices that belong to a Riemannian manifold into elements of  $\mathbb{S}_p$ , a vector space. Furthermore, the Riemannian logarithm can be vectorized using the operator  $\phi_{\Sigma}$  defined such that  $\|\log_{\Sigma}(\Sigma')\|_{\Sigma} = \|\phi_{\Sigma}(\Sigma')\|_2$  and that is

$$\phi_{\Sigma}(\Sigma') \triangleq \text{vec} \left( \log \left( \Sigma^{-1/2} \Sigma' \Sigma^{-1/2} \right) \right). \quad (3)$$

Depending on the context, the operator  $\text{vec}$  is either the full vectorization (hence  $\phi_{\Sigma}(\Sigma') \in \mathbb{R}^{p^2}$ ) or the vectorization of the upper triangular part with off-diagonal elements multiplied by  $\sqrt{2}$  to preserve the norm (hence  $\phi_{\Sigma}(\Sigma') \in \mathbb{R}^{p(p+1)/2}$ ).

**Riemannian geometry for domain adaptation** For any  $\mathcal{D} \in \{\mathcal{S}, \mathcal{T}\}$ , we denote the source and target embeddings by

$$\mathbf{x}_i^{\mathcal{D}} \triangleq \phi_{\bar{\Sigma}^{\mathcal{D}}}(\Sigma_i^{\mathcal{D}}) \quad (4)$$

where  $\bar{\Sigma}^{\mathcal{D}}$  is the Riemannian mean of the  $\Sigma_i^{\mathcal{D}}$ . Finally, it should be noted that, due to the equation (8) and the affine invariance of the Riemannian distance, we have that

$$\bar{\Sigma}^{\mathcal{D}} = \mathbf{A}^{\mathcal{D}} \bar{\mathbf{E}}^{\mathcal{D}} (\mathbf{A}^{\mathcal{D}})^{\top} \quad (5)$$

where  $\bar{\mathbf{E}}^{\mathcal{D}}$  is a block diagonal matrix with the upper left block equal to  $\text{diag}(\bar{\mathbf{p}}) \in \mathbb{R}^{q \times q}$  with elements  $\bar{p}_l = (\prod_{i=1}^n p_{i,l}^{\mathcal{S}})^{1/n} = (\prod_{i=1}^n p_{i,l}^{\mathcal{T}})^{1/n}$ .

### 2.3 Regression and classification with mixed and domain shifted signals

We consider two classical models in M/EEG on the decision function for regression and classification problems. These models relate outcomes  $y_i^{\mathcal{D}}$  with the variances

$\{p_{il}^{\mathcal{D}}\}_{l=1}^q$  [19, 5, 12]. We extend these models to domain adaptation, and notably, we assume parameters  $\beta_0, \dots, \beta_q \in \mathbb{R}$  of these decision functions to be equal between source and target domains.

**Prediction models: regression and classification** For regression problems, a classical assumption is the existence of a log-linear relationship between the value to regress  $y_i^{\mathcal{D}}$  of the  $i^{\text{th}}$  subject in domain  $\mathcal{D}$  and the associated variances  $p_{il}^{\mathcal{D}}$  (1), i.e., there exists  $\beta_0, \dots, \beta_q$  such that

$$y_i^{\mathcal{D}} = \sum_{l=1}^q \beta_l \log(p_{il}^{\mathcal{D}}) + \beta_0 + \varepsilon_i^{\mathcal{D}} \quad (6)$$

where  $\varepsilon_i^{\mathcal{D}}$  is a random noise. Considering a binary classification problem<sup>1</sup> with outcomes in  $\{-1, 1\}$ , the second model assumes that  $y_i^{\mathcal{D}}$  is the sign of a log-linear regression on  $\{p_{il}^{\mathcal{D}}\}_{l=1}^q$ , i.e.,

$$y_i^{\mathcal{D}} = \text{sign} \left( \sum_{l=1}^q \beta_l \log(p_{il}^{\mathcal{D}}) + \beta_0 + \varepsilon_i^{\mathcal{D}} \right) \quad (7)$$

where  $\text{sign}(a) = 1$  if  $a \geq 0$  and  $-1$  otherwise.

**OTDA assumption** An additional assumption is the optimal transport domain adaptation (OTDA) assumption. In our setup, we assume there exists a matrix  $\boldsymbol{\pi} \in \boldsymbol{\Pi}(n, m) \triangleq \{\boldsymbol{\pi} \in \mathbb{R}^{n \times m} : \pi_{ij} \geq 0, \boldsymbol{\pi} \mathbf{1}_m = \frac{1}{n} \mathbf{1}_n, \boldsymbol{\pi}^\top \mathbf{1}_n = \frac{1}{m} \mathbf{1}_m\}$ , the set of discrete joint distributions of uniform marginals, such that

$$\mathbf{p}_i^{\mathcal{S}} = \mathbf{p}_j^{\mathcal{T}} \quad \text{if } \pi_{ij} > 0. \quad (8)$$

This assumption is tightly related to the assumption of preserved conditional distribution stated in [11]. Indeed, the authors assume the label information is preserved by the transformation  $\boldsymbol{\pi}$ , i.e., for every  $i$ ,  $\mathbb{P}_{\mathcal{S}}(y|\mathbf{p}_i^{\mathcal{S}}) = \mathbb{P}_{\mathcal{T}}(y|m(\mathbf{p}_i^{\mathcal{S}}))$  with  $\mathbb{P}_{\mathcal{D}}$  the conditional probability distribution on the domain  $\mathcal{D}$  and  $m(\mathbf{p}_i^{\mathcal{S}})^\top = (\boldsymbol{\pi}[\mathbf{p}_1^{\mathcal{T}}, \dots, \mathbf{p}_m^{\mathcal{T}}]^\top)_{[i,:]}$ . This assumption holds in our setup if the noise  $\varepsilon_i^{\mathcal{D}}$  has the same distribution for the two domains.

**Prediction from covariance matrices** These models motivate the recovery of  $\mathbf{p}_i^{\mathcal{D}}$  from  $\mathbf{x}_i^{\mathcal{D}}$ . The following proposition indicates that, under the mixing model (1), this recovery is possible through orthogonal projections of  $\mathbf{x}_i^{\mathcal{S}}$  and  $\mathbf{x}_j^{\mathcal{T}}$  with well-chosen matrices. The latter belongs to the Stiefel manifold, the set of  $q$ -dimensional orthogonal basis in  $\mathbb{R}^{p^2}$  denoted  $\text{St}(p^2, q) \triangleq \{\mathbf{U} \in \mathbb{R}^{p^2 \times q} : \mathbf{U}^\top \mathbf{U} = \mathbf{I}_q\}$ .

**Proposition 2.1** (Stiefel projections). *Given source and target embeddings,  $\mathbf{x}_i^{\mathcal{S}}$  and  $\mathbf{x}_j^{\mathcal{T}}$  (4), following mixing models (1), with  $\pi_{ij} > 0$  (8), there exist  $\mathbf{U}^{\mathcal{S}}, \mathbf{U}^{\mathcal{T}} \in \text{St}(p^2, q)$  such that*

$$(\mathbf{U}^{\mathcal{S}})^\top \mathbf{x}_i^{\mathcal{S}} = (\mathbf{U}^{\mathcal{T}})^\top \mathbf{x}_j^{\mathcal{T}} = [\log(p_{i,1}/\bar{p}_1), \dots, \log(p_{i,q}/\bar{p}_q)]^\top$$

where  $p_{i,l} \triangleq p_{i,l}^{\mathcal{S}} = p_{j,l}^{\mathcal{T}}, \forall l \in \llbracket 1, q \rrbracket$ ,

<sup>1</sup>Throughout the paper, classification problems are presented with two classes only to simplify the exposition. The extension to multiple classes is straightforward.

The proof is available in the Appendix A.1. Intuitively,  $U^S$  and  $U^T$  extract the signal of interest and unmix it, i.e., invert  $A^S$  and  $A^T$ . Thus, these Stiefel matrices remove the shift between source and target domains. Combining Proposition 2.1 with models (6) and (7) gives two new models that relate  $\mathbf{x}_i^D$  with  $y_i^D$ .

**Corollary 2.2** (Realined predictive models). *For all  $D \in \{S, T\}$  and under mixing models (1), there exists  $U^D \in \text{St}(p^2, q)$  and  $\beta_0, \dots, \beta_q \in \mathbb{R}$  such that the data  $(\mathbf{x}_i^D, y_i^D)$  follow the models:*

$$y_i^D = \sum_{l=1}^q \beta_l ((U^D)^\top \mathbf{x}_i^D)_l + \beta_0 + \varepsilon_i^D,$$

in regression, and

$$y_i^D = \text{sign} \left( \sum_{l=1}^q \beta_l ((U^D)^\top \mathbf{x}_i^D)_l + \beta_0 + \varepsilon_i^D \right)$$

in classification.

We finish this section by pointing out that until now,  $\mathbf{x}_i^D \in \mathbb{R}^{p^2}$ . In practice, we use the half vectorization as explained in section 2.2 to reduce the dimension of the embeddings to  $d \triangleq p(p+1)/2$ .

### 3 Joint learning of $U^S$ , $U^T$ and $\beta$

In the last section, we defined our data, their underlying mixing model, and the domain adaptation setup. We now present an algorithm to learn  $U^S$ ,  $U^T$  and  $\beta = [\beta_0, \dots, \beta_q] \in \mathbb{R}^{q+1}$  from data.

**Setup** We assume having access from a fully labeled source domain and a partially labeled target domain. Indeed,  $n$  source data,  $\{(\mathbf{x}_i^S, y_i^S)\}_{i=1}^n$ , are concatenated into  $(\mathbf{X}^S, \mathbf{y}^S) \in \mathbb{R}^{n \times d} \times \mathbb{R}^n$ . Among the  $m$  targets data,  $m_k$  are labeled,  $\{(\mathbf{x}_i^T, y_i^T)\}_{i=1}^{m_k}$ , and concatenated into  $(\mathbf{X}^{T_\ell}, \mathbf{y}^{T_\ell}) \in \mathbb{R}^{m_k \times d} \times \mathbb{R}^{m_k}$ . Overall,  $(\mathbf{X}^S, \mathbf{y}^S, \mathbf{X}^{T_\ell}, \mathbf{y}^{T_\ell})$  constitutes the training set. The remaining target data  $\{\mathbf{x}_i^T\}_{i=m_k+1}^m$  are unlabeled, and the goal is to predict their associated outcomes  $\{y_i^T\}_{i=m_k+1}^m$ .

#### 3.1 Learning problem

**Overall loss** Based on models from section 2, we propose a loss function to jointly estimate  $U^S$ ,  $U^T$ , and  $\beta$  for regression and classification problems. Given  $\gamma, \rho \geq 0$ , the overall loss is

$$\mathcal{L}(U^S, U^T, \beta, \pi) \triangleq \mathcal{L}_{\text{sup.}}(U^S, U^T, \beta) + \gamma \mathcal{L}_{\text{met.}}(U^S, U^T, \pi) + \rho \mathcal{L}_{\text{Gr}}(U^S, U^T). \quad (9)$$

This loss is a sum of three terms: an empirical risk  $\mathcal{L}_{\text{sup.}}$  on the score functions, a metric learning loss  $\mathcal{L}_{\text{met.}}$  between the embeddings  $(U^S)^\top \mathbf{x}_i^S$  and  $(U^T)^\top \mathbf{x}_i^T$  and a regularization loss  $\mathcal{L}_{\text{Gr}}$  to control how far  $\text{span}(U^T)$  is from  $\text{span}(U^S)$ .

**Empirical risk  $\mathcal{L}_{\text{sup}}$ .** The loss function on the score function is a regularized empirical risk on the source data and the few labeled target data

$$\mathcal{L}_{\text{sup.}}(\mathbf{U}^S, \mathbf{U}^T, \boldsymbol{\beta}) \triangleq \sum_{\mathcal{D} \in \{\mathcal{S}, \mathcal{T}_\ell\}} \sum_{i=1}^{|\mathcal{D}|} \mathcal{L}(y_i^{\mathcal{D}}, \boldsymbol{\beta}^\top (\mathbf{U}^{\mathcal{D}})^\top \mathbf{x}_i^{\mathcal{D}} + \beta_0) + \varepsilon \|\boldsymbol{\beta}\|_2^2 \quad (10)$$

where  $\mathcal{L} : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ , and  $\varepsilon \geq 0$ . For a regression problem, we leverage the classic square error  $\mathcal{L}(y, \hat{y}) = (y - \hat{y})^2$ . For a binary classification problem in  $\{-1, 1\}$ , the chosen loss function is the logistic loss  $\mathcal{L}(y, \hat{y}) = \log(1 + \exp(-y\hat{y}))$ .

**Metric learning loss  $\mathcal{L}_{\text{met}}$ .** Then, we introduce the metric learning loss that enhances the estimation of  $\mathbf{U}^S$  and  $\mathbf{U}^T$  leveraging a geometrical point of view. First, we leverage optimal transport (OT) [38, 18] to estimate a mapping  $\boldsymbol{\pi} \in \boldsymbol{\Pi}(n, m)$  between the source and target domains and hence bring closer points from the domains with a high value of  $\pi_{ij}$ . Second, we compute intra-domain affinity matrices  $\mathbf{K}(\mathbf{y})$  so that two points with similar  $y_i$  should be brought closer. Overall, this metric learning loss function is

$$\begin{aligned} \mathcal{L}_{\text{met.}}(\mathbf{U}^S, \mathbf{U}^T, \boldsymbol{\pi}) &\triangleq \overbrace{\langle \mathbf{C}(\mathbf{X}^S \mathbf{U}^S, \mathbf{X}^T \mathbf{U}^T), \boldsymbol{\pi} \rangle}^{\text{OT loss}} \\ &+ \underbrace{\sum_{\mathcal{D} \in \{\mathcal{S}, \mathcal{T}_\ell\}} \frac{|\mathcal{D}|}{n + m_k} \langle \mathbf{C}(\mathbf{X}^{\mathcal{D}} \mathbf{U}^{\mathcal{D}}, \mathbf{X}^{\mathcal{D}} \mathbf{U}^{\mathcal{D}}), \mathbf{K}(\mathbf{y}^{\mathcal{D}}) \rangle}_{\text{similarity loss}} \end{aligned} \quad (11)$$

where  $\langle \boldsymbol{\xi}, \boldsymbol{\eta} \rangle \triangleq \text{Tr}(\boldsymbol{\xi}^\top \boldsymbol{\eta})$ ,  $(\mathbf{C}(\mathbf{X}, \mathbf{X}'))_{ij} \triangleq \|\mathbf{x}_i - \mathbf{x}'_j\|_2^2$  and  $\mathbf{K}(\mathbf{y})$  is an affinity matrix: for a regression problem, we leverage the normalized Gaussian kernel whose elements are  $(\mathbf{K}(\mathbf{y}))_{ij} = \exp(-\frac{1}{2}(y_i - y_j)^2) / \sum_{l,m} \exp(-\frac{1}{2}(y_l - y_m)^2)$ ; for classification problems, we use the matrix with elements  $(\mathbf{K}(\mathbf{y}))_{ij} = 1_{y_i=y_j} / \sum_{l,m} 1_{y_l=y_m}$ . It should be noted that the joint distribution  $\boldsymbol{\pi}$  estimates the permutation matrix from the assumption of equation (8).

**Grassmann loss  $\mathcal{L}_{\text{Gr}}$**  We now move on to the last term of the total loss (9). Unfortunately, learning  $\mathbf{U}^T$  from data seems infeasible since we have few labels in the target domain. To alleviate this problem, we regularize  $\text{span}(\mathbf{U}^S)$  and  $\text{span}(\mathbf{U}^T)$  to be not too far apart. To do so, we leverage the Riemannian distance of the Grassmann manifold [13], i.e.

$$\begin{aligned} \mathcal{L}_{\text{Gr}}(\mathbf{U}^S, \mathbf{U}^T) &\triangleq \frac{1}{\sqrt{2}} \|\mathbf{U}^S (\mathbf{U}^S)^\top - \mathbf{U}^T (\mathbf{U}^T)^\top\|_F \\ &= \left[ \sum_{i=1}^q \sin(\theta_i)^2 \right]^{1/2} \end{aligned} \quad (12)$$

where  $\theta_1, \dots, \theta_k$  are the principle angles between  $\mathbf{U}^S$  and  $\mathbf{U}^T$  [13]. It's important to note that as  $\rho \rightarrow +\infty$  in the overall loss (9),  $\text{span}(\mathbf{U}^T) \rightarrow \text{span}(\mathbf{U}^S)$ , and hence  $\mathbf{U}^T \rightarrow \mathbf{U}^S \mathbf{O}$  for some  $q \times q$  orthogonal matrix  $\mathbf{O}$ . Thus, for  $q$  much lower than  $d$ ,  $\mathcal{L}_{\text{Gr}}$  controls the complexity to estimate  $\mathbf{U}^T$  as desired.



---

**Algorithm 1** Alternate optimization

---

- 1: Initialize:  $(\mathbf{U}^S, \mathbf{U}^T) \in \text{St}(d, q)^2$
  - 2: **repeat**
  - 3:   Compute  $\boldsymbol{\pi}^*$  and  $\boldsymbol{\beta}^*$  as in (14)
  - 4:   Perform a Riemannian gradient descent step on  $\text{St}(d, q)^2$  of  $\mathcal{L}^*$  (15)
  - 5: **until** convergence
- 

### 3.2 Reduction to a problem on the Stiefel manifold $\text{St}(d, q)^2$ and optimization

We now discuss the optimization of the loss function (9). The optimization problem is written

$$\underset{(\mathbf{U}^S, \mathbf{U}^T, \boldsymbol{\beta}, \boldsymbol{\pi}) \in \text{St}(d, q)^2 \times \mathbb{R}^q \times \boldsymbol{\Pi}(n, m)}{\text{minimize}} \quad \mathcal{L}(\mathbf{U}^S, \mathbf{U}^T, \boldsymbol{\beta}, \boldsymbol{\pi}). \quad (13)$$

Given  $(\mathbf{U}^S, \mathbf{U}^T) \in \text{St}(d, q)^2$  fixed,  $\mathcal{L}$  is easily minimized with respect to  $(\boldsymbol{\beta}, \boldsymbol{\pi}) \in \mathbb{R}^q \times \boldsymbol{\Pi}(n, m)$ . Indeed, denoting

$$(\boldsymbol{\beta}^*, \boldsymbol{\pi}^*) = \underset{(\boldsymbol{\beta}, \boldsymbol{\pi}) \in \mathbb{R}^q \times \boldsymbol{\Pi}(n, m)}{\text{arg min}} \quad \mathcal{L}(\mathbf{U}^S, \mathbf{U}^T, \boldsymbol{\beta}, \boldsymbol{\pi}), \quad (14)$$

$\boldsymbol{\beta}^*$  is computed with a ridge or logistic solver and  $\boldsymbol{\pi}^*$  with an OT solver. Hence, (13) reduces to

$$\underset{(\mathbf{U}^S, \mathbf{U}^T) \in \text{St}(d, q)^2}{\text{minimize}} \quad \left\{ \mathcal{L}^*(\mathbf{U}^S, \mathbf{U}^T) \triangleq \mathcal{L}(\mathbf{U}^S, \mathbf{U}^T, \boldsymbol{\beta}^*, \boldsymbol{\pi}^*) \right\}. \quad (15)$$

The latter is solved with a Riemannian gradient descent algorithm [1, 8] leveraging a Riemannian Adam optimizer [4] on the Stiefel manifold [32]. This alternate optimization procedure is summed up in algorithm 1.

### 3.3 Related work

**Unsupervised subspace learning** Numerous methods to address domain shifts exist [50]. Focusing on subspace-based methods, Transfer Component Analysis (TCA) [34] learns some transfer components across domains in a reproducing kernel Hilbert space using maximum mean discrepancy to align source and target distributions. Other methods such as Subspace Alignment [15], joint cross-domain classification and subspace learning for unsupervised adaptation [16] and principle components-based methods [31] learn source components leveraging target principal components to align data from both domains. These methods are unsupervised in the sense they do not use the target labels. In contrast, MSA benefits from target labels to simultaneously learn embeddings, with the two Stiefel matrices, and the predictor, which helps with low signal-to-noise ratio M/EEG applications. Indeed, [33] showed the inefficiency of unsupervised domain adaptation methods such as [31] for brain age applications.

**Semi-supervised rotation correction** Methods tailored for domain adaptation in M/EEG applications have emerged alongside domain-specific whitening (4). Riemannian Procrustes Analysis (RPA) [39] aligns covariance matrix class centers by

learning with a gradient descent a  $p \times p$  orthogonal matrix, involving  $\mathcal{O}(p^2)$  parameters. In the same spirit, Tangent Space Alignment (TSA) [6] aligns tangent space vector class centers with a  $d \times d$  orthogonal matrix, requiring  $\mathcal{O}(p^4)$  parameters, learned through singular value decomposition. These methods apply only to classification datasets. In comparison, MSA learns two Stiefel matrices, that is,  $\mathcal{O}(p^2q)$  parameters, where  $q$  is user-defined. This hyperparameter has the benefit of controlling the cost computation and the complexity of the model depending on the task. Furthermore, it applies to both regression and classification tasks.

## 4 Numerical experiments

In this section, we assess the performance of MSA on a regression problem on MEG data. This problem is subject-specific, i.e., each subject has one outcome. We have two sets of subjects without any overlap from datasets where a shift has occurred, such as variations in tasks performed by the subjects. In order to promote research reproducibility, code is available on [github](https://github.com/antoinecollas/MSA)<sup>2</sup>.

### 4.1 Regression on MEG data

**Cam-CAN dataset** We utilize the well-established Cambridge Center of Aging and Neuroscience (Cam-CAN) dataset [44]. Our objective is to predict the age of individuals based on their MEG recordings, a regression problem of significant interest in neuroscience research. Indeed, results from [26] suggest that the prediction error of models trained to learn age from brain data of healthy populations provides clinically relevant information related to neurodegenerative anomalies physical and cognitive decline. In this dataset, 646 subjects (319 female, 327 male) were collected with the same 306-channel VectorView MEG system (Elekta Neuromag, Helsinki) at 1kHz frequency. Their age distribution is from 18.5 to 88.9 years with an average of  $54.9 \pm 18.4$  years with an almost uniform spread over the age range. The dataset is divided into two sets of subjects with no overlap but with the same age distribution. Notably, the subjects perform different tasks, depending on the domain, during data acquisition. These tasks are rest (eyes closed), passive (audio-visual stimuli), and somatosensory tasks (audio-visual plus a manual response). Hence, we get one time series and one outcome, the age, per subject. Preprocessing is applied to these time series; see Appendix B.1 for the details. To get the semi-supervised setting, 10% of the target labels are kept in addition to the source labels during training. Then, we predict ages with several methods, all involving ridge regression.

**Methods** First, we compare MSA with ridge regressions that are trained on embeddings either of the source domain or of the 10% labeled target domain, as in [40]. These embeddings (4) are computed at the Riemannian mean (2) of the full dataset. We denote these two methods  $\text{Ridge} - \mathcal{S}$  and  $\text{Ridge} - \mathcal{T}_\ell$  respectively. Then, we leverage recently proposed methods for brain age prediction under domain shifts: ridge regressions are trained either on the source domain or on the concatenation of the source and

<sup>2</sup><https://github.com/antoinecollas/MSA>

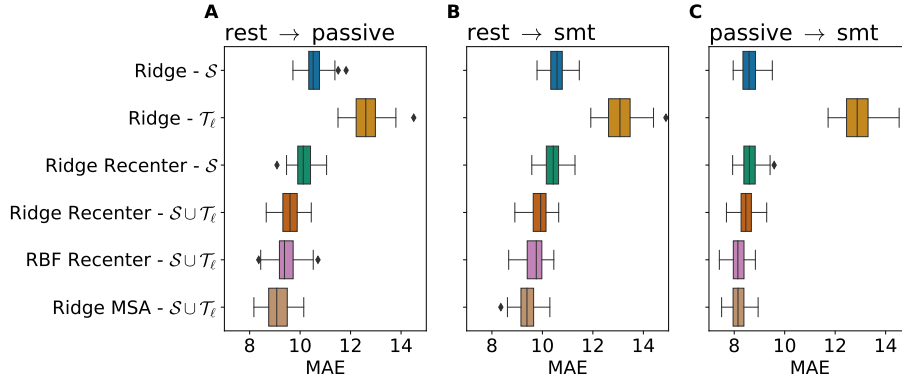


Figure 2: **Mean absolute errors (MAE) of different regressors on the brain age prediction problem of the Cam-CAN dataset (the lower, the better).** The 646 subjects ( $p = 65$ ) are split into source and target domains (323 subjects each) and are associated with two different tasks. The latter are reported over each subfigure, e.g., the subfigure (A) presents results with the rest for the source task and the passive task for the target. 10% of the target labels are kept during training, and 100 different data splits are performed.

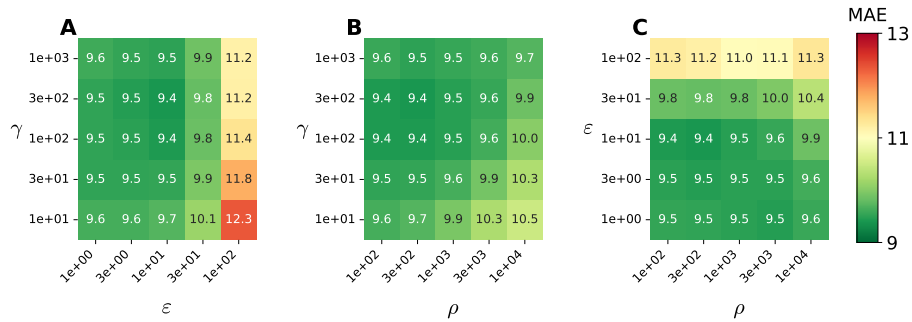


Figure 3: **Mean absolute errors (MAE) of MSA for different values of hyperparameters on the brain age prediction problem of the Cam-CAN dataset (the lower, the better).** The source task is rest, and the target task is somatosensory. 10% of the target labels are kept during training, and 100 different data splits are performed.

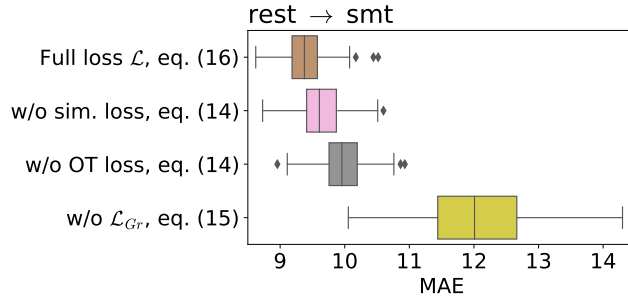


Figure 4: **Ablation study of the proposed loss (9) on the brain age prediction problem of the Cam-CAN dataset.** Mean absolute errors (MAE) of the full loss versus when one of the terms is removed are reported (the lower, the better). The source task is rest, and the target is somatosensory. 10% of the target labels are kept during training, and 100 different data splits are performed.

labeled target domains with embeddings with domain-specific Riemannian means as proposed in [33]. They are denoted  $\text{Ridge Recenter} - \mathcal{S}$  and  $\text{Ridge Recenter} - \mathcal{S} \cup \mathcal{T}$  respectively. Recently, [7] proposed to use kernel ridge regression with an RBF kernel for brain age prediction and showed state-of-the-art results. One kernel per frequency band is computed with embeddings, i.e.,  $K_{i,j} = \exp(-\frac{1}{2\sigma^2} \|\mathbf{x}_i - \mathbf{x}_j\|_2^2)$  for  $\mathbf{x}_i, \mathbf{x}_j$  in  $\mathcal{X}^{\mathcal{S}}$  and/or  $\mathcal{X}^{\mathcal{T}_\ell}$  and where  $\sigma^2$  is the hyperparameter of length scale. Then, kernels are summed to get a single kernel. All labeled data (source and labeled target) are used to compute the kernels. We denote this method  $\text{RBF Recenter} - \mathcal{S} \cup \mathcal{T}_\ell$ . Finally, MSA is denoted  $\text{Ridge MSA} - \mathcal{S} \cup \mathcal{T}_\ell$  and is used with rank  $q = 30$ . Performance with the best-performing hyperparameters is reported for all methods.

**Performance comparison** Box plots of mean absolute errors (MAE) on the 90% unlabeled data are reported in Figure 2 with training and testing on 100 different data splits. First of all, we observe that training solely on the target set ( $\text{Ridge} - \mathcal{T}_\ell$ ) leads to the worst MAE. Hence, applying the method of [40] is not recommended when too few labels are available. Then,  $\text{Ridge Recenter} - \mathcal{S}$  gives lower MAE than  $\text{Ridge} - \mathcal{S}$ , which shows the importance of considering domain adaptation techniques when considering changes of tasks. The three other methods,  $\text{Ridge Recenter} - \mathcal{S} \cup \mathcal{T}$ ,  $\text{RBF Recenter} - \mathcal{S} \cup \mathcal{T}_\ell$  and  $\text{Ridge MSA} - \mathcal{S} \cup \mathcal{T}_\ell$  all leverage these domain adaptation techniques and the few labels from the test set. Since they use these additional labels, they outperform the other methods. The  $\text{RBF Recenter} - \mathcal{S} \cup \mathcal{T}_\ell$  is the only non-linear model and performs slightly better than  $\text{Ridge Recenter} - \mathcal{S} \cup \mathcal{T}$ . Overall,  $\text{Ridge MSA} - \mathcal{S} \cup \mathcal{T}$  performs better than all other methods, regardless of the considered pair of tasks.

**Parameters sensitivity and ablation study** Then, we study the hyperparameter sensitivity of MSA. The three considered hyperparameters are  $\gamma$ ,  $\rho$  and  $\varepsilon$  from the loss (9). 100 data splits are performed for each triplet value, and MAE is reported in Figure 3.

This shows a certain robustness to the choice of hyperparameters. Finally, to assess the need for the different losses in the overall loss (9), we ablate each one, search for the best-remaining hyperparameters  $\gamma$ ,  $\rho$ , and  $\varepsilon$ , and report MAE in Figure 4. We observe that each ablation leads to poorer MAE, which empirically justifies using all terms together in the full loss (9). In particular, removing the Grassmann loss drops performance a lot because it resumes to estimate two different Stiefel matrices. Notably, removing the optimal transport also leads to a high-performance drop, which motivates the “OTDA assumption” introduced in section 2.1.

## 5 Conclusions

In conclusion, this paper has introduced a new domain adaptation approach for time series data under mixing models. The key innovation lies in utilizing a domain-dependent mixing matrix (1) and the OTDA assumption (8). Indeed, we identified two Stiefel matrices that, applied to a Riemannian representation of observed signal covariances, recover variances from the underlying signal. An integrated cost function enabled simultaneous learning of these matrices, establishing pairwise relationships between source and target domains and constructing task-specific predictors. Applied to MEG problems, this approach has outperformed recent methods in brain-age regression.

## 6 Acknowledgments

The scientific Python ecosystem enabled numerical computation: Geoopt [29], Matplotlib [22], Numpy [20], POT [17], Pyriemann [2], Pytorch [35], Seaborn [46], and Sklearn [36].

## References

- [1] Pierre-Antoine Absil, Robert Mahony, and Rodolphe Sepulchre. Optimization algorithms on matrix manifolds. Princeton University Press, 2008.
- [2] Alexandre Barachant, Quentin Barthélemy, Jean-Rémi King, Alexandre Gramfort, Sylvain Chevallier, Pedro L. C. Rodrigues, Emanuele Olivetti, Vladislav Goncharenko, Gabriel Wagner vom Berg, Ghiles Reguig, Arthur Lebeurier, Erik Bjäreholt, Maria Sayu Yamamoto, Pierre Clisson, and Marie-Constance Corsi. `pyRiemann/pyRiemann: v0.5`. In Zenodo. Zenodo, June 2023. Version Number: v0.5.
- [3] Alexandre Barachant, Stéphane Bonnet, Marco Congedo, and Christian Jutten. Multiclass Brain–Computer Interface Classification by Riemannian Geometry. IEEE Transactions on Biomedical Engineering, 59(4):920–928, 2012.
- [4] Gary Becigneul and Octavian-Eugen Ganea. Riemannian Adaptive Optimization Methods. In International Conference on Learning Representations, 2019.
- [5] Benjamin Blankertz, Ryota Tomioka, Steven Lemm, Motoaki Kawanabe, and Klaus-robert Muller. Optimizing Spatial filters for Robust EEG Single-Trial Analysis. IEEE Signal Processing Magazine, 25(1):41–56, 2008.
- [6] Alexandre Bleuzé, Jérémie Mattout, and Marco Congedo. Transfer Learning for the Riemannian Tangent Space: Applications to Brain-Computer Interfaces. In 2021 International Conference on Engineering and Emerging Technologies (ICEET), pages 1–6, 2021.
- [7] Clément Bonet, Benoit Malézieux, Alain Rakotomamonjy, Lucas Drumetz, Thomas Moreau, Matthieu Kowalski, and Nicolas Courty. Sliced-Wasserstein on Symmetric Positive Definite Matrices for M/EEG Signals. In Proceedings of the 40th International Conference on Machine Learning, pages 2777–2805. PMLR, July 2023.
- [8] Nicolas Boumal. An introduction to optimization on smooth manifolds. Cambridge University Press, 2023.
- [9] Yilun Chen, Ami Wiesel, Yonina C. Eldar, and Alfred O. Hero. Shrinkage Algorithms for MMSE Covariance Estimation. IEEE Transactions on Signal Processing, 58(10):5016–5029, 2010.
- [10] Antoine Collas, Arnaud Breloy, Chengfang Ren, Guillaume Ginolhac, and Jean-Philippe Ovarlez. Riemannian Optimization for Non-Centered Mixture of Scaled Gaussian Distributions. IEEE Transactions on Signal Processing, 71:2475–2490, 2023.
- [11] Nicolas Courty, Rémi Flamary, Devis Tuia, and Alain Rakotomamonjy. Optimal Transport for Domain Adaptation. IEEE Transactions on Pattern Analysis and Machine Intelligence, 39(9):1853–1865, September 2017. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.

- [12] 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, February 2014.
- [13] Alan Edelman, Tomás A Arias, and Steven T Smith. The geometry of algorithms with orthogonality constraints. SIAM journal on Matrix Analysis and Applications, 20(2):303–353, 1998. Publisher: SIAM.
- [14] Denis A. Engemann, Apolline Mellot, Richard Höchenberger, Hubert Banville, David Sabbagh, Lukas Gemein, Tonio Ball, and Alexandre Gramfort. A reusable benchmark of brain-age prediction from M/EEG resting-state signals. NeuroImage, 262:119521, November 2022.
- [15] Basura Fernando, Amaury Habrard, Marc Sebban, and Tinne Tuytelaars. Unsupervised Visual Domain Adaptation Using Subspace Alignment. In 2013 IEEE International Conference on Computer Vision, pages 2960–2967, December 2013. ISSN: 2380-7504.
- [16] Basura Fernando, Tatiana Tommasi, and Tinne Tuytelaars. Joint cross-domain classification and subspace learning for unsupervised adaptation. Pattern Recognition Letters, 65:60–66, November 2015.
- [17] Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z. Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, Léo Gautheron, Nathalie T. H. Gayraud, Hicham Janati, Alain Rakotomamonjy, Ievgen Redko, Antoine Rolet, Antony Schutz, Vivien Seguy, Danica J. Sutherland, Romain Tavenard, Alexander Tong, and Titouan Vayer. POT: Python Optimal Transport. Journal of Machine Learning Research, 22(78):1–8, 2021.
- [18] Edouard Grave, Armand Joulin, and Quentin Berthet. Unsupervised Alignment of Embeddings with Wasserstein Procrustes. In Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics, pages 1880–1890. PMLR, April 2019.
- [19] Moritz Grosse-Wentrup and Martin Buss. Multiclass Common Spatial Patterns and Information Theoretic Feature Extraction. IEEE Transactions on Biomedical Engineering, 55(8):1991–2000, August 2008.
- [20] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. Nature, 585(7825):357–362, September 2020. Publisher: Springer Science and Business Media LLC.

- [21] Zhiwu Huang and Luc Van Gool. A Riemannian Network for SPD Matrix Learning. In Association for the Advancement of Artificial Intelligence (AAAI), 2017.
- [22] John D. Hunter. Matplotlib: A 2D graphics environment. Computing in Science & Engineering, 9(3):90–95, 2007. Publisher: IEEE COMPUTER SOC.
- [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):413–497, April 1993.
- [24] 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, October 2017.
- [25] Vinay Jayaram, Morteza Alamgir, Yasemin Altun, Bernhard Scholkopf, and Moritz Grosse-Wentrup. Transfer Learning in Brain-Computer Interfaces. IEEE Computational Intelligence Magazine, 11(1):20–31, February 2016.
- [26] Tobias Kaufmann, Dennis van der Meer, Nhat Trung Doan, Emanuel Schwarz, Martina J. Lund, Ingrid Agartz, Dag Alnæs, Deanna M. Barch, Ramona Baur-Streubel, Alessandro Bertolino, Francesco Bettella, Mona K. Beyer, Erlend Bøen, Stefan Borgwardt, Christine L. Brandt, Jan Buitelaar, Elisabeth G. Celius, Simon Cervenka, Annette Conzelmann, Aldo Córdova-Palomera, Anders M. Dale, Dominique J. F. de Quervain, Pasquale Di Carlo, Srdjan Djurovic, Erlend S. Dørum, Sarah Eisenacher, Torbjørn Elvsåshagen, Thomas Espeseth, Helena Fatouros-Bergman, Lena Flyckt, Barbara Franke, Oleksandr Frei, Beathe Haatveit, Asta K. Håberg, Hanne F. Harbo, Catharina A. Hartman, Dirk Heslenfeld, Pieter J. Hoekstra, Einar A. Høgestøl, Terry L. Jernigan, Rune Jonassen, Erik G. Jönsson, Peter Kirsch, Iwona Kłoszewska, Knut K. Kolskår, Nils Inge Landrø, Stephanie Le Hellard, Klaus-Peter Lesch, Simon Lovestone, Arvid Lundervold, Astri J. Lundervold, Luigi A. Maglanoc, Ulrik F. Malt, Patrizia Mecocci, Ingrid Melle, Andreas Meyer-Lindenberg, Torgeir Moberget, Linn B. Norbom, Jan Egil Nordvik, Lars Nyberg, Jaap Oosterlaan, Marco Papalino, Andreas Papassotiropoulos, Paul Pauli, Giulio Pergola, Karin Persson, Geneviève Richard, Jaroslav Rokicki, Anne-Marthe Sanders, Geir Selbæk, Alexey A. Shadrin, Olav B. Smeland, Hilkka Soininen, Piotr Sowa, Vidar M. Steen, Magda Tsolaki, Kristine M. Ulrichsen, Bruno Vellas, Lei Wang, Eric Westman, Georg C. Ziegler, Mathias Zink, Ole A. Andreassen, and Lars T. Westlye. Common brain disorders are associated with heritable patterns of apparent aging of the brain. Nature Neuroscience, 22(10):1617–1623, October 2019. Number: 10 Publisher: Nature Publishing Group.
- [27] Nirmal Keshava and John F Mustard. Spectral unmixing. IEEE signal processing magazine, 19(1):44–57, 2002. Publisher: IEEE.
- [28] Reinmar J. Kobler, Jun-ichiro Hirayama, Qibin Zhao, and Motoaki Kawanabe. SPD domain-specific batch normalization to crack interpretable unsupervised domain adaptation in EEG, October 2022.



- [29] Max Kochurov, Rasul Karimov, and Serge Kozlukov. Geoopt: Riemannian Optimization in PyTorch, 2020. [\\_eprint: 2005.02819](#).
- [30] Wouter M. Kouw and Marco Loog. A Review of Domain Adaptation without Target Labels. IEEE Transactions on Pattern Analysis and Machine Intelligence, 43(3):766–785, 2021.
- [31] Gal Maman, Or Yair, Danny Eytan, and Ronen Talmon. Domain Adaptation Using Riemannian Geometry of Spd Matrices. In ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 4464–4468, May 2019. ISSN: 2379-190X.
- [32] Jonathan H Manton. Optimization algorithms exploiting unitary constraints. IEEE transactions on signal processing, 50(3):635–650, 2002. Publisher: IEEE.
- [33] Apolline Mellot, Antoine Collas, Pedro LC Rodrigues, Denis A Engemann, and Alexandre Gramfort. Harmonizing and aligning M/EEG datasets with covariance-based techniques to enhance predictive regression modeling. [bioRxiv](#), pages 2023–04, 2023. Publisher: Cold Spring Harbor Laboratory.
- [34] Sinno Jialin Pan, Ivor W. Tsang, James T. Kwok, and Qiang Yang. Domain Adaptation via Transfer Component Analysis. IEEE Transactions on Neural Networks, 22(2):199–210, February 2011. Conference Name: IEEE Transactions on Neural Networks.
- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc., 2019.
- [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] Xavier Pennec, Pierre Fillard, and Nicholas Ayache. A Riemannian Framework for Tensor Computing. International Journal of Computer Vision, 66(1):41–66, January 2006.
- [38] Gabriel Peyré and Marco Cuturi. Computational Optimal Transport: With Applications to Data Science. Foundations and Trends® in Machine Learning, 11(5-6):355–607, February 2019. Publisher: Now Publishers, Inc.
- [39] Pedro Luiz Coelho Rodrigues, Christian Jutten, and Marco Congedo. Riemannian Procrustes Analysis: Transfer Learning for Brain–Computer Interfaces. IEEE Transactions on Biomedical Engineering, 66(8):2390–2401, August 2019.

- [40] David Sabbagh, Pierre Ablin, Gaël Varoquaux, Alexandre Gramfort, and Denis A. Engemann. Manifold-regression to predict from MEG/EEG brain signals without source modeling. Advances in Neural Information Processing Systems, 32, 2019.
- [41] David Sabbagh, Pierre Ablin, Gaël Varoquaux, Alexandre Gramfort, and Denis A. Engemann. Predictive regression modeling with MEG/EEG: from source power to signals and cognitive states. NeuroImage, 222:116893, November 2020.
- [42] Lene Theil Skovgaard. A Riemannian Geometry of the Multivariate Normal Model. Scandinavian Journal of Statistics, 11(4):211–223, 1984. Publisher: [Board of the Foundation of the Scandinavian Journal of Statistics, Wiley].
- [43] S. Taulu, J. Simola, and M. Kajola. Applications of the signal space separation method. IEEE Transactions on Signal Processing, 53(9):3359–3372, September 2005. Conference Name: IEEE Transactions on Signal Processing.
- [44] Jason R. Taylor, Nitin Williams, Rhodri Cusack, Tibor Auer, Meredith A. Shafto, Marie Dixon, Lorraine K. Tyler, Cam-CAN, and Richard N. Henson. 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.
- [45] Oncel Tuzel, Fatih Porikli, and Peter Meer. Pedestrian detection via classification on riemannian manifolds. IEEE transactions on pattern analysis and machine intelligence, 30(10):1713–1727, 2008. Publisher: IEEE.
- [46] Michael L. Waskom. seaborn: statistical data visualization. Journal of Open Source Software, 6(60):3021, 2021. Publisher: The Open Journal.
- [47] Florian Yger, Maxime Berar, and Fabien Lotte. Riemannian Approaches in Brain-Computer Interfaces: A Review. IEEE Transactions on Neural Systems and Rehabilitation Engineering, 25(10):1753–1762, October 2017.
- [48] Paolo Zanini, Marco Congedo, Christian Jutten, Salem Said, and Yannick Berthoumieu. Transfer Learning: A Riemannian Geometry Framework With Applications to Brain-Computer Interfaces. IEEE Transactions on Biomedical Engineering, 65(5):1107–1116, May 2018.
- [49] Hongyi Zhang and Suvrit Sra. First-order Methods for Geodesically Convex Optimization. In Conference on Learning Theory, pages 1617–1638. PMLR, June 2016. ISSN: 1938-7228.
- [50] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A Comprehensive Survey on Transfer Learning. Proceedings of the IEEE, 109(1):43–76, January 2021. Conference Name: Proceedings of the IEEE.

## A Technical details

### A.1 Proof of Proposition 2.1

Without loss of generality, we assume that  $\boldsymbol{\pi}$  is the identity, and thus, from equation (8)  $\text{diag}(\boldsymbol{p}_i^S) = \text{diag}(\boldsymbol{p}_i^T)$ . Given  $\boldsymbol{W}^S = (\bar{\boldsymbol{\Sigma}}^S)^{-1/2} \boldsymbol{A}^S (\bar{\boldsymbol{E}}^S)^{1/2} \in \mathcal{O}(p)$ , we have

$$\log \left( (\bar{\boldsymbol{\Sigma}}^S)^{-1/2} \boldsymbol{\Sigma}_i^S (\bar{\boldsymbol{\Sigma}}^S)^{-1/2} \right) = \boldsymbol{W}^S \log \left( (\bar{\boldsymbol{E}}^S)^{-1/2} \boldsymbol{E}_i^S (\bar{\boldsymbol{E}}^S)^{-1/2} \right) (\boldsymbol{W}^S)^\top.$$

This implies that

$$\begin{aligned} \boldsymbol{x}_i^S &= \text{vec} \left( \log \left( (\bar{\boldsymbol{\Sigma}}^S)^{-1/2} \boldsymbol{\Sigma}_i^S (\bar{\boldsymbol{\Sigma}}^S)^{-1/2} \right) \right) \\ &= (\boldsymbol{W}^S \otimes \boldsymbol{W}^S) \text{vec} \left( \log \left( (\bar{\boldsymbol{E}}^S)^{-1/2} \boldsymbol{E}_i^S (\bar{\boldsymbol{E}}^S)^{-1/2} \right) \right). \end{aligned}$$

Since  $\log \left( (\bar{\boldsymbol{E}}^S)^{-1/2} \boldsymbol{E}_i^S (\bar{\boldsymbol{E}}^S)^{-1/2} \right)$  is a block diagonal matrix whose upper  $q \times q$  block is  $\text{diag}(\log(p_{i,1}/\bar{p}_1), \dots, \log(p_{i,q}/\bar{p}_q))$ , there exists a basis  $\boldsymbol{Q} = [\boldsymbol{U}, \boldsymbol{U}_\perp] \in \mathcal{O}(d)$  such that

$$\text{vec} \left( \log \left( (\bar{\boldsymbol{E}}^S)^{-1/2} \boldsymbol{E}_i^S (\bar{\boldsymbol{E}}^S)^{-1/2} \right) \right) = \boldsymbol{U} [\log(p_{i,1}/\bar{p}_1), \dots, \log(p_{i,q}/\bar{p}_q)]^\top + \boldsymbol{U}_\perp \boldsymbol{n}_i^S$$

where  $\boldsymbol{n}_i^S = \boldsymbol{U}_\perp^\top \text{vec} \left( \log \left( (\bar{\boldsymbol{E}}^S)^{-1/2} \boldsymbol{E}_i^S (\bar{\boldsymbol{E}}^S)^{-1/2} \right) \right)$ . It follows that for  $\boldsymbol{U}^S \triangleq (\boldsymbol{W}^S \otimes \boldsymbol{W}^S) \boldsymbol{U}$ , we have

$$(\boldsymbol{U}^S)^\top \boldsymbol{x}_i^S = [\log(p_{i,1}/\bar{p}_1), \dots, \log(p_{i,q}/\bar{p}_q)]^\top.$$

Applying the same reasoning to the target embeddings, we get that for  $\boldsymbol{U}^T \triangleq (\boldsymbol{W}^T \otimes \boldsymbol{W}^T) \boldsymbol{U}$ , we have  $(\boldsymbol{U}^T)^\top \boldsymbol{x}_i^T = [\log(p_{i,1}/\bar{p}_1), \dots, \log(p_{i,q}/\bar{p}_q)]^\top$ .

## B Numerical experiments details

### B.1 Datasets preprocessing

We follow the preprocessing steps proposed in [14, 33] for the Cam-CAN dataset. We apply a FIR band-pass filter between 0.1 and 49Hz. We decimate the signals to get a sampling frequency of 200Hz. To compensate for environmental noise, we perform a temporal signal-space-separation (tSSS) method [43] with a chunk duration of 10 seconds and a correlation threshold of 98%. Then, we retain channels corresponding to magnetometers. Each filtered recording is segmented in 10s epochs without overlap. Then, epochs are filtered into 7 frequency bands: 0.1-1, 1-4, 4-8, 8-15, 15-26, 26-35, and 35-49Hz. We performed artifact rejection by thresholding extreme peak-to-peak amplitudes on single epochs using the local autoreject method [24]. We average the epochs and compute one covariance matrix per subject and per frequency band using the Oracle Approximating Shrinkage (OAS) estimator [9]. Then, principal component analysis transforms them to matrices of size  $p = 65$ . We compute embeddings  $\boldsymbol{x}_i^D$ , still for each frequency band, using the function  $\phi$  from equation (3). Finally, these embeddings of all frequency bands are concatenated into one embedding per subject.

## B.2 MEG additional figures

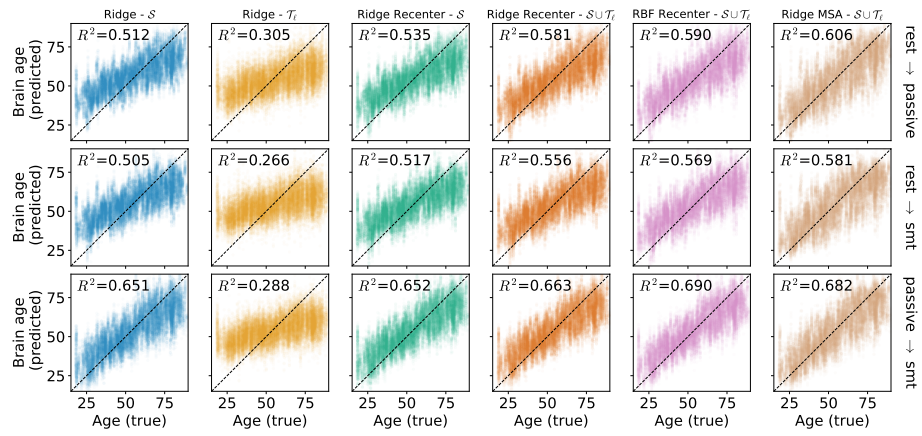


Figure 5: **Scatter plots of different regressors on the brain age prediction problem of the Cam-CAN dataset.**  $R^2$  scores are reported for each method on each pair of tasks (the higher, the better).




Figure 6: **Evolution of different quantities involved in the loss (9) on the brain age prediction problem of the Cam-CAN dataset.** The source task is rest, and the target task is somatosensory. Each quantity is computed 100 times with different data splits and then averaged.

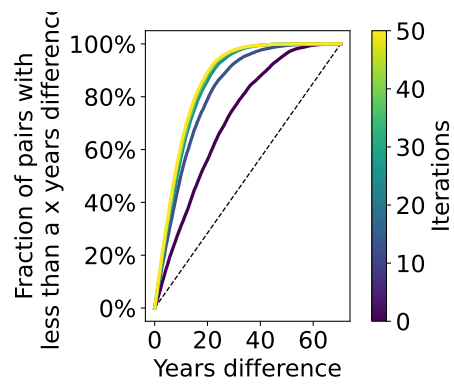


Figure 7: **Fraction of pairs  $\pi_{ij}$  with less than a given number of years difference on the brain age prediction problem of the Cam-CAN dataset for different numbers of iterations of the Algorithm 1.** The source task is rest, and the target task is somatosensory. Given a number of iterations, the optimal transport plan  $\pi$  is computed 100 times with different data splits and then averaged.