A Geometry-Based Data Augmentation for EEG Motor Imagery Classification via Deep Riemannian Networks

João Guilherme Prado Barbon and Denis Gustavo Fantinato
The Department of Computer Engineering and Industrial Automation
State University of Campinas
Barão Geraldo - Campinas - SP - Brazil
j262760@dac.unicamp.br, denisf@unicamp.br

***Abstract*—This study introduces a novel data augmentation method for improving electroencephalography (EEG)-based motor imagery classification. Leveraging the inherent variability in EEG sensor placement, geometric augmentations are applied to EEG data to simulate realistic cap misalignments during data acquisition. The resulting augmented datasets are used to train Deep Riemannian Networks models, such as SPDNet and EE(G)-SPDNet. Results show the potential of the method for improving the robustness and generalizability of motor imagery brain-computer interfaces.**

***Index Terms*—Data Augmentation, Brain-computer Interfaces, Riemannian Networks.**

I. Introduction

Brain-computer interfaces (BCIs) represent a significant advancement in assistive technologies and neurorehabilitation, offering the potential to restore lost motor function and improve communication for individuals with severe motor impairments [3], [4]. However, the inherent complexity of electroencephalography (EEG) data, combined with the inter-subject and inter-session variability, poses a challenge for achieving reliable and robust BCI performance [12].

The field of machine learning, in special deep learning, has offered substantial advances in addressing these challenges [2]. The capacity of neural networks to learn complex representations from EEG data presents a considerable potential for improving the accuracy of BCI systems. This is especially true for models based on Riemannian geometry, which leverage the unique geometric properties of covariance matrices to effectively capture the complex structure of EEG data. These models have demonstrated promising results in classifying EEG signals in various applications [9].

However, electroencephalography signals are still inherently complex data due to their variability and redundancy. This variability arises from several sources: inter-subject and cognitive differences in the brain anatomy, intra-subject fluctuations in alertness and electrode placement, and volume conduction, including eye blinks and muscle movements. At the same time, signal redundancy, often marked by the spatial mixing of brain source signals and, in a temporal extent, the repeating wave segments, also adds complexity to EEG data. This redundancy, although initially hinders accurate classifications, can also serve as a discriminative feature. For instance, the redundancy rate may distinguish between healthy and epileptic EEG patterns [7]. Addressing these challenges is important for developing reliable interfaces, which demand methods that adapt to individual variability and handle possible redundancy effects.

In this context, this article presents a novel data augmentation strategy designed to mitigate the effects of EEG data variability in the acquisition process. Specifically, we leverage the variability introduced by the physical placement and possible misalignment of EEG sensors on the scalp, simulating them with geometric transformations applied directly to the data. This approach aims to generate synthetic EEG data that captures the fluctuations and imperfections generally found in real recordings, possibly leading to improvements in classification performance. We use Riemannian-based models, such as SPDNet [11] and EE(G)-SPDNet [10], taking advantage of the combination of its capabilities in handling the non-Euclidean structure of covariance matrices with the geometric transformations applied to the EEG data. The next sections detail our methodology, results and the implications of this method.

II. Deep Riemannian Networks

This study focuses on using Riemannian-based models for EEG classification, chosen for its particular suitability in handling the complexities of EEG data. Different from traditional Euclidean-based approaches, these models use the principles of Riemannian geometry to represent and process the data. This is a crucial advantage since this type of data often reside in a non-Euclidean manifold, exhibiting complex relations that are not entirely captured by Euclidean methods. The SPDNet approach maps high-dimensional EEG covariance matrices into a low dimensional Symmetric Positive Definite (SPD) manifold, resulting in a more compact and efficient representation. This reduction in dimensionality contributes

This study was financed, in part, by the São Paulo Research Foundation (FAPESP), Brazil. Process Number #2023/13300-4.

to improved computational efficiency and reduces the risk of overfitting.

The SPDNet architecture is composed of three types of layers, as shown in Figure 1. First is the bilinear mapping (BiMap) layers, that transform the input SPD matrices into a new set of matrices. This transformation is mainly designed to maintain the properties while generating a more discriminative representation. The weight matrices within these layers are constrained to be semi-orthogonal, which ensures the output remains within the SPD manifold. Second, eigenvalue rectification (ReEig) layers introduce a non-linear activation function, similar to ReLU in convolutional neural networks. This layer rectifies the smaller eigenvalues in the matrices, addressing the issue of non-singular or non-positive definite matrices that frequently arise from the EEG signals. Finally, the eigenvalue logarithm (LogEig) layer maps the SPD matrices to a Euclidean space via a logarithm transformation, enabling the application of standard Euclidean operations for the final classification task.

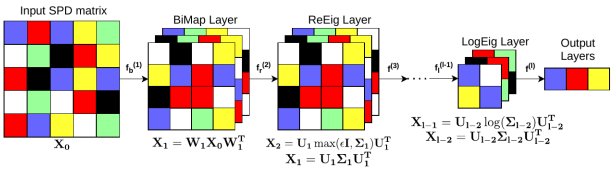

Fig. 1: Illustration of the SPDNet network.

The EE(G)-SPDNet model extends the standard SPDNet architecture by using a convolutional layer preceding the SPDNet layers, enabling an end-to-end processing of the raw EEG data. This convolutional layer functions as a learnable filter bank, potentially capturing ideal frequencies before computing the covariance matrices, thus improving generalizability. The model's deeper architecture, with multiple BiMap/ReEig layers, allows it to learn more complex representations from the data, potentially benefiting from larger amounts of data available, compared with a shallow architecture.

III. The Geometric Augmentation

In order to enhance the robustness of Riemannian models and take into account the effects of cap movement during data acquisition, we introduced a novel data augmentation strategy based on geometric transformations. This strategy simulates realistic variations in sensor placement that might occur during recording, specifically focusing on rotations and translations of the EEG cap.

To simulate rotations of the cap on the data, we applied matrix rotations to the EEG preprocessed signals using a general n-dimensional rotation algorithm [1]. This algorithm efficiently computes the rotation matrix by aligning the rotation axis $\mathbf{v} \in \mathcal{R}^{n-1\times n}$ with a specific subspace, performing the desired rotation, and then returning the axis to its initial position. For example, in a 3D rotation of angle θ around axis $\mathbf{x_1}$, we can define two points $\mathbf{a} = (a_1^0, a_2^0, a_3^0)$ and $\mathbf{b} = (b_1^0, b_2^0, b_3^0)$ as the vertices of a line segment $S = \overrightarrow{\mathbf{ab}}$. The matrix $\mathbf{v}^k = \begin{bmatrix}\mathbf{a}^k \\ \mathbf{b}^k\end{bmatrix}$ holds the vertices after k transformations, calculated iteratively as $\mathbf{v}^k = M_k \cdot \mathbf{v}^{k-1}$, where M_k is the transformation matrix. The rotation process begins by translating point $\mathbf{a}$ to the origin (Step 1 $\rightarrow$ 2 in Figure 2), resulting in $M_1 = T(-\mathbf{a})$:

$$\mathbf{v}^1 = M_1 \cdot \mathbf{v}^0, \tag{1}$$

$$= T(-\mathbf{a}) \cdot \begin{bmatrix}\mathbf{a}^0 \\ \mathbf{b}^0\end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 \\ b_1^0 - a_1^0 & b_2^0 - a_2^0 & b_3^0 - a_3^0 \end{bmatrix} \tag{2}$$

$$= \begin{bmatrix} 0 & 0 & 0 \\ b_1^1 & b_2^1 & b_3^1 \end{bmatrix} = \begin{bmatrix}\mathbf{0} \\ \mathbf{b}^1\end{bmatrix} \tag{3}$$

Next, $\mathbf{b}^1$ is rotated by angle θ_1 around $\mathbf{x_1}$ to align $\bar{ab}$ with the (x_1, x_2) plane. With $\theta_1 = \arctan 2(b_3^1, b_2^1)$, the transformation matrix is $M_2 = R_{3,2}(\theta_1)$, where $R_{a,b}(\theta)$ is defined as

$$R_{a,b}(\theta) = r_{ij} \begin{cases} r_{a,a} = \cos\theta, r_{b,b} = \cos\theta, \\ r_{a,b} = -\sin\theta, r_{b,a} = \sin\theta, \\ r_{i,j} = 1, j \neq a, j \neq b, \\ r_{i,j} = 0, \text{ else.} \end{cases} \tag{4}$$

Applying M_2 to $\mathbf{v}^1$ yields $\mathbf{v}^2 = \begin{bmatrix}\mathbf{0} \\ \mathbf{b}^2\end{bmatrix}$ (Step 2 $\rightarrow$ 3 in Figure 2). This process continues one more time, rotating $\mathbf{b}^2$ to align with $\mathbf{x_1}$ until the desired orientation is achieved. Finally, the desired rotation of angle θ ($M_4 = R_{2,3}(\theta)$) is applied, followed by the inverse transformations (M_3^{-1}, M_2^{-1}, M_1^{-1}) to restore the original axis orientation.

For a general n-dimensional rotation of angle θ around axis $\mathbf{v}$, the overall transformation matrix M_{total} is a composition of a series of operations. First, an alignment matrix M is constructed as:

$$M = T(-\mathbf{a}) \cdot \prod_{k=2}^{n(n-1)/2} R_{c,c-1}\left(\arctan 2(v_{r,c}^{(k-1)}, v_{r,c-1}^{(k-1)})\right). \tag{5}$$

This matrix sequentially aligns the rotation axis with column c and row r, setting $\mathbf{v}_{r,c}^{(k)} = 0$. The final rotation matrix, M_{total}, then incorporates the desired rotation and reverses the alignment process:

$$M_{total} = M \cdot R_{(n-1),n}(\theta) \cdot M^{-1}. \tag{6}$$

The translation transformation, intended to simulate variations in the cap tightness (loose or tight fitting), was implemented by applying a translation matrix to the preprocessed EEG data, effectively adding or subtracting the signal's values. The transformations were then added into the respective preprocessed data, and the augmented dataset randomly shuffled, increasing the amount of samples for each subject analyzed.

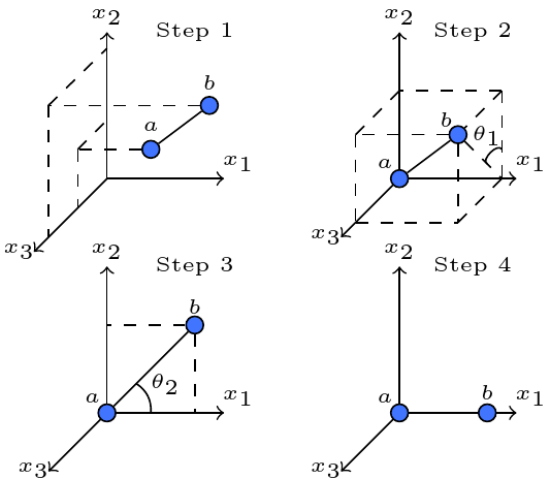

Fig. 2: General 3D transformation steps to align $\bar{ab}$ to axis x_1.

IV. Dataset and Methodology

A. Dataset and DRNs

This study uses the BCI Competition IV - dataset 2a [8], comprising EEG data from nine subjects performing four-class motor imagery tasks. The dataset features recordings of imagined movements of the left and right hands, both feet, and tongue. Each subject participated in two sessions on separate days, with each session consisting of multiple runs and a total of 288 trials. EEG data was acquired using a 22-channel Ag/AgCl montage, sampled at 250 Hz and bandpass filtered to reduce noise. The EEG data was preprocessed, with bandpass filtering between 4.0 Hz and 38.0 Hz. Trials were segmented into three seconds overlapping windows with a 0.5 second stride and the signals were converted to microvolts. A Sample Covariance Matrix layer (SCM) then transforms each epoched EEG data segment ($X \in \mathcal{R}^{N_{CH} \times N_T}$) into a covariance matrix ($C \in \mathcal{R}^{N_{CH} \times N_{CH}}$). This ensures the input of the model is symmetric and positive definite.

Two model configurations were used, SPDNet with a single BiMap/ReEig layer, and EE(G)-SPDNet with an initial convolutional layer followed by the SPDNet with three BiMap/ReEig blocks. The same parameters were used across all subject for each model: a learning rate of 0.01225 and a batch size of 64 were used for the SPDNet, trained for 600 epochs; a learning rate of 0.001 and a batch size of 216 were used for the EE(G)-SPDNet, trained for 50 epochs. The Adam optimizer was employed for both models, while only SPDNet had cosine annealing learning rate scheduling. The parameters for the augmentation transformations were optimized for each model, using a train/validation/test split, with validation size of 0.33. Following optimization, the best transformation parameters for each subject were used to compute the mean test accuracy across four independent runs, in a train/test split configuration. The Weights and Biases [5] and Hydra [13] frameworks were used for experiment logging and management.

B. Augmentation Hyperparameters Selection

The parameters used in these transformations – the rotation angle and axis, the translation vector, and the overall augmentation rate – were optimized for each subject and for each model using the Optuna framework [6]. For both SPDNet and EE(G)-SPDNet models, the optimization was leveraged with a total of 100 trials per subject to maximize the validation accuracy as the objective function. A Random Sampler was used to sort through the possible parameter values, and a Median Pruner criteria stopped unpromising runs. The search space for the rotation angle spanned $[-10°, 5°]$ with increments of $0.1°$, intended to capture possible asymmetries in the cap positioning. The rotation axis was represented as a matrix with dimension in terms of the number of channels N_{CH}, $(N_{CH} - 1, N_{CH})$, filled by the same value, searched within the interval $[-1.0, 1.0]$ with a 0.1 step size. The translation vector, with dimension $(1, N_{CH})$, similarly represented as a matrix with a single value per channel, was searched in the range $[-5.0, 5.0]$ with 0.1 increments. Finally, for the SPDNet model, the augmentation rate was optimized across the range [0.25, 1.0] with steps of 0.25, representing the proportion of augmented data added to the original dataset.

Table I shows the optimized rotation parameters for each subject on the SPDNet model trained on the BCI Competition IV - dataset 2a. These parameters reveal substantial subject-specific variations in the best configurations for rotation transformations. The selected augmentation rates for rotations ranged from 0.25 to 1.0, indicating that some subjects were able to exploit the augmented data more efficiently than others. The best rotation angles showed considerable variability, spanning from $-9.8°$ to $3.3°$, possibly reflecting different degrees of sensitivity to cap misalignment. The rotation axis values, also demonstrating variability within their defined range, highlight the complexity in capturing the nature of these misalignments with a single optimized value.

Subject	Rotation		
	Augmentation	Angle	Axis Value
1	1.0	-8.6	-0.2
2	1.0	-5.1	-0.3
3	0.5	3.3	0.0
4	1.0	-7.8	-0.7
5	0.75	-2.8	-0.1
6	0.75	0.4	0.4
7	0.25	0.2	0.8
8	0.75	-6.1	-0.3
9	1.0	-9.8	0.7

TABLE I: Optimized rotation parameters for the SPDNet model with BCI Competition IV - dataset 2a.

Table II presents the optimized translation parameters for the same dataset. The size of translation, with the best size spanning from -3.6 to 4.0, may reflect individual differences in the scalp's surface and skull shape, leading to different effects from a loose or tight EEG cap. These results emphasize the heterogeneous nature of EEG data and how individual subject-specific responses to geometric augmentation must be

taken into account, further reinforcing the need for adaptive feature extraction techniques involving cap movement.

Subject	Translation	
	Augmentation	Size
1	0.75	2.0
2	0.75	2.7
3	0.25	3.6
4	0.5	0.8
5	0.25	1.3
6	0.25	4.0
7	0.5	-3.6
8	0.5	2.0
9	0.75	1.7

TABLE II: Optimized translation parameters for the SPDNet model with BCI Competition IV - dataset 2a.

Finally, Table III combines the optimized rotation and translation parameters, providing a comprehensive view of the best configurations for each subject. The chosen parameters suggest that individual subjects may require different types of augmentation to maximize model performance, as some benefit from higher augmentation rates, whereas others need lower modification intensities.

Subject	Rot+Tr			
	Augmentation	Angle	Axis Value	Tr. Size
1	1.0	-8.6	0.0	2.0
2	1.0	-5.1	-0.3	2.7
3	0.5	3.3	0.0	3.6
4	1.0	-7.8	-0.7	0.8
5	0.75	-2.8	-0.1	1.3
6	0.75	0.4	0.4	4.0
7	0.5	0.2	0.8	-3.6
8	0.75	-6.1	-0.3	2.0
9	1.0	-9.8	0.7	1.7

TABLE III: Optimized rotation and translation parameters combined for the SPDNet model with BCI Competition IV - dataset 2a.

The transformation parameters with the EE(G)-SPDNet model were optimized following the same configuration as described before, except that the augmentation rate was searched across a fixed set of values [0.5, 1.0, 1.5, 2.0, 3.0, 4.0, 5.0, 6.0, 7.0, 8.0, 9.0, 10.0], allowing for augmentations up to ten times the original dataset size. This strategy aimed at leveraging the deeper architecture of the EE(G)-SPDNet model, potentially improving its ability to learn more complex representations and achieve better generalization. To ensure diversity in the augmented data and prevent excessive repetition of transformations, in augmentations above 1.0, a small perturbation (0.001) was added to each parameter value. Figure 3 shows the effects of -10° and 5° rotations around the axis $\mathbf{v} = [[0, 1, 0], [0, 0, 0]]$ applied to a topographic map of subject one of BCI Competition IV - dataset 2a.

V. Results and Discussion

Table IV presents classification accuracies for the SPDNet and EE(G)-SPDNet models, respectively, trained on the BCI Competition IV - dataset 2a with and without geometric data augmentation. In the SPDNet model, the results show an inter-subject variability in response to augmentation: some subjects exhibited minor accuracy improvements (e.g., subjects 2 and 3), while others showed minor decreases or no change. It's important to note that the average accuracy remained consistent and within the statistical uncertainty (around 1-2% difference) across baseline and augmented conditions for both models, suggesting the augmentation strategy did not negatively impact overall performance. Also, the augmentation strategy combining rotation and translation showed classification accuracies similar to those obtained using rotation alone, indicating that the translation component of the augmentation may not significantly contribute to improving model performance. Hence, rotation was selected for the next step.

Subject	SPDNet				EE(G)-SPDNet	
	Baseline	Rotation	Translation	Rot. + Transl.	Baseline	Rotation
1	73.3	71.6	72.3	72.2	79.5	**85.6**
2	44.6	**45.8**	**45.4**	**45.5**	51.7	**53.7**
3	74.5	**74.9**	74.5	**75.0**	77.2	**84.7**
4	41.8	**42.0**	41.5	**41.9**	61.5	**69.9**
5	38.6	37.7	35.5	37.8	52.5	**65.3**
6	45.2	44.9	**45.5**	44.4	48.0	**52.5**
7	72.5	70.6	70.2	70.2	76.1	**83.5**
8	75.2	73.2	74.3	72.9	77.1	76.3
9	73.1	69.6	69.4	70.0	75.5	74.2
Average	59.9	58.9	58.7	58.9	66.6	**71.7**

TABLE IV: Classification accuracies using SPDNet and EE(G)-SPDNet models on the BCI Competition IV - dataset 2a with geometric augmentations.

The EE(G)-SPDNet outperformed the SPDNet in both baseline and rotation conditions, demonstrating that the convolutional layer and deeper BiMap/ReEig architecture improve data generalization. The exceptions were subjects 8 and 9, that showed close baseline accuracy and a small variation with rotation augmentation compared to the SPDNet. In general, the rotation augmentation under the EE(G)-SPDNet model showed considerable improvements across many subjects, with accuracy gains up to 12% with subject 5. The higher augmentation rates selected for the EE(G)-SPDNet model supports the hypothesis that larger datasets could enhance this model's ability to fully utilize its deeper architecture and achieve better generalization. The subject's optimal augmentation parameters highlight the inherent variability in individual responses and the potential influence of EEG cap movement during data acquisition. This variability, while initially a challenge, could be leveraged through a more refined optimization of the augmentation parameters.

VI. Conclusions

This study presents a novel geometric data augmentation strategy designed to improve the robustness and accuracy of motor imagery classification BCIs based on EEG, using Deep Riemannian Networks (DRNs) models. Our study demonstrates that, for a shallow model such as SPDNet, the proposed augmentation method consistently maintained performance comparable to the baseline results, while for a deeper model

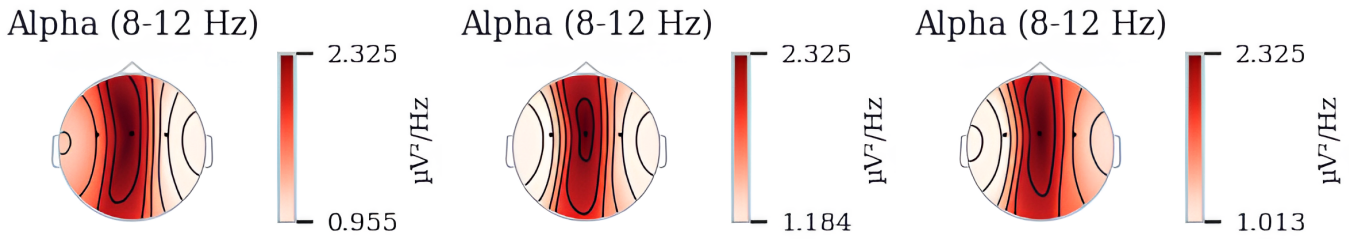

(a) Rotation of $-10°$ (b) No rotation (c) Rotation of $5°$

Fig. 3: Rotation transformation effect on topographic map of EEG data from the subject one of BCI Competition IV - dataset 2a. The data was selected for the right hand movement and C3, Cz and C4 channels.

such as EE(G)-SPDNet, there was a considerable improvement across many subjects. This suggests the augmentation strategy is robust and highlights the potential of DRNs to effectively process and interpret EEG data with inherent variability. Further investigation into refining the methods, including different datasets and a more detailed exploration of different geometric transformations is also intended.

REFERENCES

[1] Antonio Aguilera and Ricardo Pérez-Aguila. General n-dimensional rotations. 02 2004.

[2] Ali Al-Saegh, Shefa A. Dawwd, and Jassim M. Abdul-Jabbar. Deep learning for motor imagery EEG-based classification: A review. *Biomedical Signal Processing and Control*, 63:102172, January 2021.

[3] Kai Keng Ang and Cuntai Guan. Brain-computer Interface for Neurorehabilitation of Upper Limb After Stroke. *Proceedings of the IEEE*, 103(6):944–953, 2015.

[4] Mahdi Bamdad, Homayoon Zarshenas, and Mohammad A. Auais. Application of BCI systems in neurorehabilitation: a scoping review. *Disability and Rehabilitation: Assistive Technology*, 10(5):355–364, January 2015.

[5] Lukas Biewald. Experiment Tracking with Weights and Biases, 2020. Software available from wandb.com.

[6] Akiba et al. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019.

[7] Amin et al. A Novel Approach Based on Data Redundancy for Feature Extraction of EEG Signals. *Brain Topography*, 29(2):207–217, March 2016.

[8] Tangermann et al. Review of the BCI Competition IV. *Frontiers in Neuroscience*, 6, 2012.

[9] Tibermacine et al. Riemannian Geometry-Based EEG Approaches: A Literature Review, 2024.

[10] Wilson et al. Deep Riemannian Networks for EEG Decoding. (arXiv:2212.10426), August 2023. arXiv:2212.10426 [cs, eess, stat].

[11] Zhiwu Huang and Luc Van Gool. A Riemannian Network for SPD Matrix Learning. (arXiv:1608.04233), December 2016. arXiv:1608.04233 [cs].

[12] Simanto Saha and Mathias Baumert. Intra- and Inter-subject Variability in EEG-Based Sensorimotor Brain Computer Interface: A Review. *Frontiers in Computational Neuroscience*, 13, January 2020.

[13] Omry Yadan. Hydra - A framework for elegantly configuring complex applications. Github, 2019.