IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 3 January 2026, accepted 2 February 2026, date of publication 10 February 2026, date of current version 19 February 2026.

Digital Object Identifier 10.1109/ACCESS.2026.3663245

RESEARCH ARTICLE

Structure-Preserving EEG Augmentation via Riemannian Conditional Generative Adversarial Networks

MARTINA DOKU[1], IMAD EDDINE TIBERMACINE[1], (Member, IEEE), SAMUELE RUSSO[2], ABDELAZIZ RABEHI[3], MUSTAPHA HABIB[4], AND CHRISTIAN NAPOLI[1,5,6]

[1]Department of Computer, Automation and Management Engineering, Sapienza University of Rome, 00185 Rome, Italy
[2]Faculty of Psychology, Sapienza University of Rome, 00185 Rome, Italy
[3]Laboratory of Telecommunications and Smart Systems, Faculty of Sciences and Technologies, University of Djelfa, Djelfa 17000, Algeria
[4]Division of Building Technology and Design, Department of Civil and Architectural Engineering, KTH Royal Institute of Technology, 11428 Stockholm, Sweden
[5]Institute for Systems Analysis and Computer Science, Italian National Research Council, 00185 Rome, Italy
[6]Department of Computational Intelligence, Czestochowa University of Technology, 42201 Czestochowa, Poland

Corresponding author: Mustapha Habib (mushab@kth.se)

ABSTRACT Electroencephalography (EEG) data are pivotal in brain–computer interfaces (BCIs), yet their utility is hindered by data scarcity arising from high acquisition costs, noise susceptibility, and privacy constraints. Traditional augmentation methods, such as noise injection and signal transformations, often fail to preserve task-relevant structure in multichannel EEG, while deep generative models may suffer from mode collapse or produce physiologically inconsistent samples. To address these limitations, we propose a Riemannian Conditional Generative Adversarial Network (RC-GAN) that enforces geometric consistency during signal generation. RC-GAN leverages the manifold of symmetric positive definite (SPD) covariance matrices to regularize synthetic EEG trials according to covariance-based representations widely used in BCI decoding. Evaluated on the BNCI 2014-001 motor imagery dataset, the proposed method outperforms state-of-the-art augmentation techniques, achieving a 12.0% improvement in classification accuracy. Qualitative and quantitative analyses demonstrate that RC-GAN generates diverse and realistic EEG samples while enhancing robustness at different augmentation levels. These results highlight the benefit of incorporating Riemannian structure into generative models for EEG augmentation and provide a principled framework for improving the reliability of BCI systems.

INDEX TERMS EEG, GAN, Riemannian geometry, SPD matrices, motor imagery, BCI.

I. INTRODUCTION

EEG data are the core of BCIs, a technology that allows direct communication between the brain and an external device. The main issue in the field of BCIs is the lack of data, which makes it difficult to train deep learning models. This problem is due to a variety of factors [1]. A first limitation is the difficulty of collecting data as EEG acquisition requires specialized equipment that can be prohibitively expensive for many small research groups [2], [3]. Recording EEGs is a time-consuming process that involves human participants who must undergo extensive sessions under controlled conditions. This is especially challenging in clinical settings, where working with patients adds another layer of difficulty due to ethical and logistical constraints. Privacy concerns also play an important role, as EEG data contain sensitive neurological information, raising issues regarding data sharing and accessibility. Finally, the presence of artifacts in the data, such as noise [4], [5], muscle activity, and eye movements, can make it difficult to interpret the results, leading to a decrease in the quality of the data.

To solve this problem, data augmentation techniques are used to generate new samples from existing data, increase the size of the dataset, and ultimately improve the generalization

The associate editor coordinating the review of this manuscript and approving it for publication was Cesar Vargas-Rosales.

of the models trained on it. This solution to the lack of data problem is widely used in a variety of fields, such as computer vision, natural language processing, and speech recognition. In the field of computer vision, the most common data augmentation techniques involve geometric transformations such as rotation, scaling, translation, flipping, cropping, and color jittering [6]. However, these techniques are not directly applicable to EEG, as they can introduce artifacts into the data, making it difficult to interpret the results [7].

A first approach to EEG augmentation is the addition of noise to the signal [8]. This technique is widely used in the field of signal processing, as it is simple and effective, and different types of noise can be added to the signal to produce numerous variations [9]. However, it has some limitations, such as the introduction of artifacts in the data, which can make it difficult to interpret the results. Other approaches to EEG augmentation are the temporal scaling of the signal or the temporal flipping [10]. These techniques are simple to implement, but they distort the signal and can produce only a limited number of variations; therefore, they are not suitable for the generation of a sufficiently large dataset for training complex models. More advanced techniques are channel interpolation [11], Short-Time Fourier Transform (STFT) or Wavelet Transform [12], [13], [14]. These techniques effectively capture the time-frequency structure of the signal and can produce a large number of variations. However, these methods can introduce artifacts in the data, and they are computationally expensive, hence unsuitable for the generation of a large dataset [15].

The generation of realistic EEG signals has recently been addressed using Generative Models which report promising results, since they can produce a large number of variations of the data and capture the underlying structure, generalizing well to unseen data. In [10], the authors propose a method based on GANs to generate EEG signals. The problem with such models is that they are prone to mode collapse [16], a phenomenon in which the generator produces only a limited number of samples. To avoid this problem, Bouallegue et al. [17] exploits Wasserstein GANs (WGAN), a version of GANs that uses the Wasserstein distance to measure the distance between the real and generated samples, mitigating the mode collapse problem [18].

In [19] instead, the authors propose a method based on Diffusion Models. These models work by progressively transforming a structured dataset into noise using a diffusion process and then learn to reverse this process to generate realistic samples. Diffusion models capture intricate patterns using the learned sequence of probabilistic steps, and this makes them effective for high-dimensional data generation. The problem with this kind of models is that they suffer from the vanishing gradient phenomenon [20], that occurs when the gradients of the generator vanish, making the discriminator prevail over the generator in the learning process. An additional concern is the lack of control over the generated samples, as it is difficult to ensure that the generated samples are realistic and representative of the data distribution [21].

A promising approach in the field of data augmentation is the exploitation of Riemannian features, that have been used in many applications that range from classification, to encoding and to augmentation since they are able to capture relevant information about the structure of the data. In [22] the authors propose a training method that exploits the geometric structure of the images by forcing a bi-Lipschitz constraint between the input and the output of the generator. This method has shown promising results in the generation of realistic images. Another work that exploits the Riemannian structure of the data is the one from [23], where the authors propose a model that enhances the performance of a GAN model by enforcing a Riemannian constraint on the discriminator, again in the image domain. Finally, in [24] the Manifold Geometry Matching Generative Adversarial Network (MGM GAN) relies on the Riemannian structure of the data to generate realistic images, this time by enforcing a Riemannian constraint on the generator.

Most existing EEG augmentation pipelines either (i) rely on signal–space heuristics that yield only limited and often artifact-prone variations, or (ii) employ generic deep generative models that operate in Euclidean signal space without accounting for the geometric structure of covariance-based EEG representations. As a result, such models may suffer from mode collapse or generate samples whose second-order spatial statistics are inconsistent with those of real data, leading to reduced physiological plausibility and limited benefit for downstream decoding tasks.

Recent years have seen increasing interest in incorporating Riemannian geometry into deep generative modeling. Some works study the geometry induced by neural generators on latent or data manifolds to enable geodesic interpolation and curvature analysis of learned representations [25], [26], [27]. Others introduce geometry-aware regularization—primarily in image domains—by enforcing bi-Lipschitz constraints or Riemannian penalties on either the generator or the discriminator [22], [23], [24], [28]. In parallel, diffusion-based generative models have been extended to non-Euclidean manifolds, including Lie groups and SPD manifolds, to enable intrinsic generation of structured data [29], [30], [31]. While powerful, these approaches are typically developed for image or manifold-native data and do not directly address raw multichannel EEG generation under task-relevant SPD constraints [32], [33].

More broadly, geometry-aware generation has been explored through normalizing flows and score-based models on curved spaces, where probabilistic modeling is carried out intrinsically on manifolds and transport between tangent spaces is used to define valid dynamics [34], [35]. These developments highlight that geometric constraints can be built either by (i) analyzing the geometry induced by a generator, or by (ii) enforcing generation to match a prescribed geometric structure in a representation space.

Our work follows the second direction and targets the SPD geometry of covariance descriptors used in EEG decoding.

In the EEG domain, Riemannian geometry has been extensively adopted for feature extraction and classification through covariance representations on $\mathcal{S}_{++}^C$, yielding strong performance in motor imagery decoding [26], [36], [33]. However, most EEG generative models either ignore covariance geometry or operate directly in a feature space, weakening the link to signal-level realism [37], [38].

In this paper, we introduce RC-GAN, a conditional generator for raw EEG time-series that enforces geometric consistency on the manifold of symmetric positive definite covariance matrices. Our approach (i) couples a label-conditioned generator with a log–Euclidean consistency loss defined on $\mathcal{S}_{++}^C$, (ii) provides closed-form gradients for numerically stable training, and (iii) yields a bound that controls geodesic deviation under small regularization. Importantly, RC-GAN does not aim to reproduce individual EEG waveforms pointwise; instead, it generates synthetic trials whose covariance structure remains faithful to that of real data while allowing variability at the signal level.

In the context of this work, the term structure-preserving refers specifically to the preservation of second-order spatial structure in multichannel EEG signals. This structure is captured by regularized spatial covariance matrices, which encode channel-to-channel relationships arising from synchronized neural activity, shared cortical sources, and volume conduction. Such covariance representations form the basis of widely adopted EEG and BCI decoding pipelines, including common spatial patterns and Riemannian geometry–based classifiers, where discriminative information is predominantly conveyed by class-dependent variance and coupling patterns rather than exact waveform morphology.

II. PRELIMINARIES

This section fixes notation and recalls the differential–geometric and adversarial–learning concepts on which our model is built.

A. SYMMETRIC POSITIVE–DEFINITE MANIFOLD

Let

$$\mathcal{S}_{++}^C := \left\{ S \in \mathbb{R}^{C\times C} \,\middle|\, S = S^\top,\ S \succ 0 \right\}$$

denote the manifold of $C \times C$ SPD matrices. For $S \in \mathcal{S}_{++}^C$ the tangent space is $T_S\mathcal{S}_{++}^C = \{H \in \mathbb{R}^{C\times C} \mid H = H^\top\}$. Throughout, $\|\cdot\|_F$ is the Frobenius norm and $\mathrm{Tr}(\cdot)$ the trace.

1) LOG–EUCLIDEAN METRIC

Following [39], we equip $\mathcal{S}_{++}^C$ with the log–Euclidean Riemannian metric

$$g_S(A,B) = \left\langle S^{-\frac{1}{2}} A S^{-\frac{1}{2}},\ S^{-\frac{1}{2}} B S^{-\frac{1}{2}} \right\rangle_F, \tag{1}$$

whose induced geodesic distance is

$$d_{\mathrm{LE}}(S_1,S_2) = \left\| \log S_1 - \log S_2 \right\|_F. \tag{2}$$

Because the principal matrix logarithm $\log\colon \mathcal{S}_{++}^C \to Sym(C)$ is a diffeomorphism, $\mathcal{S}_{++}^C$ with (1) is isometric to the Euclidean vector space of symmetric matrices, endowing geodesics and gradients with closed forms.

Throughout this work, all distances, class anchors (Fréchet means), and gradients on $\mathcal{S}_{++}^C$ are defined under the *log–Euclidean* metric in (1)–(2). Equivalently, we work in the global chart $\log : \mathcal{S}_{++}^C \to \mathrm{Sym}(C)$, where computations reduce to Euclidean operations on $\log S$ and are mapped back to $\mathcal{S}_{++}^C$ via $\exp(\cdot)$.

a: EXPONENTIAL AND LOGARITHM MAPS (LOG–EUCLIDEAN)

For $S \in \mathcal{S}_{++}^C$ and $H \in T_S\mathcal{S}_{++}^C$,

$$\mathrm{Exp}_S^{\mathrm{LE}}(H) = \exp\big(\log S + H\big), \tag{3}$$

$$\mathrm{Log}_S^{\mathrm{LE}}(S') = \log S' - \log S, \tag{4}$$

where $\exp(\cdot)$ and $\log(\cdot)$ denote the matrix exponential and principal matrix logarithm, respectively.

2) FRÉCHET MEAN

Given SPD samples $\{S_i\}_{i=1}^N$, their intrinsic (Karcher) mean is the minimizer

$$\Pi_{\mathcal{M}} = \arg\min_{S\in\mathcal{S}_{++}^C} \sum_{i=1}^{N} d_{\mathrm{LE}}\big(S_i, S\big)^2, \tag{5}$$

solved efficiently by the Riemannian gradient flow [40]. We will use class-wise Fréchet means as Riemannian anchors in the loss of Sec. III-B.

B. GENERATIVE ADVERSARIAL NETWORKS

A Generative Adversarial Network (GAN) [41] is a two-player minimax game between

1) a **generator** $G_\theta : \mathcal{Z} \to \mathcal{X}$, mapping latent noise $\mathbf{z} \sim p_{\mathbf{z}}$ to synthetic data $\widehat{\mathbf{x}} = G_\theta(\mathbf{z})$;
2) a **discriminator** $D_\psi : \mathcal{X} \to (0,1)$, trained to distinguish real from generated samples.

The vanilla objective is

$$\min_\theta \max_\psi \mathbb{E}_{\mathbf{x}\sim p_{\mathrm{data}}}[\log D_\psi(\mathbf{x})] + \mathbb{E}_{\mathbf{z}}[\log(1 - D_\psi(G_\theta(\mathbf{z})))], \tag{6}$$

where p_{data} is the empirical data distribution.

C. CONDITIONAL GANs

In a conditional GAN (cGAN) Mirza2014 both networks receive side-information $c \in \mathcal{C}$ (e.g. a class label). The loss becomes

$$\min_\theta \max_\psi \mathbb{E}_{(\mathbf{x},c)\sim p_{\mathrm{data}}}[\log D_\psi(\mathbf{x}, c)] + \mathbb{E}_{\mathbf{z},c}[\log(1 - D_\psi(G_\theta(\mathbf{z}, c), c))]. \tag{7}$$

We will equip G_θ with a Riemannian regularizer so that its outputs match not only the conditional distribution $p_{\mathrm{data}}(\mathbf{x} \mid c)$ but also the geometry of $\mathcal{S}_{++}^C$.

D. NOTATION SUMMARY

C	Number of EEG channels.
T	Number of time samples per trial.
$\mathbf{x}$	Real EEG trial in $\mathbb{R}^{C\times T}$.
$\widehat{\mathbf{x}}$	Synthetic (generated) trial.
$\Sigma(\mathbf{x})$	Covariance of $\mathbf{x}$.
d_{LE}	Log–Euclidean distance.
$\Pi_{\mathcal{M}}$	Fréchet mean of a set/class in $\mathcal{S}_{++}^{C}$.
$p_{\mathbf{z}}$	Latent noise distribution (Gaussian).
λ	Weight of Riemannian loss term.
ε	Tikhonov covariance regularization parameter.
τ	Eigenvalue clipping threshold.
ρ	Minimum eigenvalue bound for SPD compactness.

These preliminaries allow us to define and analyze the Riemannian Conditional GAN introduced in Sec. III-A.

III. PROPOSED MODEL

In this section we describe the proposed method, more specifically we focus on the architecture of the RC-GAN (Figure 1), its components, and the Riemannian loss that is used in the training of the model. The objective of this model is to produce sound signals that preserve the manifold structure of the real samples. We ensure this structure properties by adding a Riemannian constraint to the generator loss, ensuring that both the signal distance and the distance in the Riemannian domain are minimized.

A. RIEMANNIAN CONDITIONAL GENERATIVE ADVERSARIAL NETWORK (RC–GAN)

We cast EEG augmentation as a two–player game on the manifold of SPD covariance matrices (Figure 2). Let $\mathcal{X} = \mathbb{R}^{C\times T}$ be the space of single–trial EEG segments with C channels and T time steps, and $c \in \{1, \ldots, K\}$ the class label. For any trial $\mathbf{x} \in \mathcal{X}$ we define the (Tikhonov–regularized) spatial covariance

$$\Sigma(\mathbf{x}) \;=\; \frac{1}{T}\,\mathbf{x}\mathbf{x}^{\top} + \varepsilon I_C, \qquad \varepsilon > 0,$$

which lies in the SPD manifold $\mathcal{S}_{++}^{C} := \{S \in \mathbb{R}^{C\times C} \mid S = S^{\top},\ S \succ 0\}$.

1) RIEMANNIAN GEOMETRY OF COVARIANCES

Endow $\mathcal{S}_{++}^{C}$ with the log–Euclidean metric [39]. The geodesic distance between $S_1, S_2 \in \mathcal{S}_{++}^{C}$ is

$$d_{\mathrm{LE}}(S_1, S_2) \;=\; \big\|\log S_1 - \log S_2\big\|_F, \tag{8}$$

where $\|\cdot\|_F$ is the Frobenius norm and $\log(\cdot)$ the principal matrix logarithm. Because $\log\colon \mathcal{S}_{++}^{C} \to \mathbb{R}^{C\times C}$ is a diffeomorphism, (8) gives closed–form geodesics and gradients.

2) RC–GAN PLAYERS

The generator $G_\theta : \mathcal{Z}\times\mathcal{C} \to \mathcal{X}$ (with parameters θ) maps a latent vector $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, I)$ and the label c to a synthetic trial $\widehat{\mathbf{x}} = G_\theta(\mathbf{z}, c)$. The discriminator $D_\psi : \mathcal{X}\times\mathcal{C} \to (0, 1)$ (with parameters ψ) outputs the probability that its input is real. Architectural details (two-layer LSTM generator; three-block 1-D CNN discriminator) follow Figure 1 and Table 1.

3) LOSS WITH RIEMANNIAN REGULARISATION

Write $\mathbb{P}_{\mathrm{data}}$ for the empirical distribution of $(\mathbf{x}, c)$ and $\mathbb{P}_\theta$ for the model distribution generated by G_θ. The standard conditional-GAN (cGAN) objective is

$$\begin{aligned}\mathcal{L}_{\mathrm{cGAN}}(\theta, \psi) = \mathbb{E}_{(\mathbf{x},c)\sim\mathbb{P}_{\mathrm{data}}}\big[\log D_\psi(\mathbf{x}, c)\big] \\ + \mathbb{E}_{\mathbf{z},c}\big[\log\big(1 - D_\psi(G_\theta(\mathbf{z}, c), c)\big)\big].\end{aligned} \tag{9}$$

To force synthetic signals to respect covariance geometry we add the Riemannian consistency loss

$$\mathcal{L}_{\mathrm{R}}(\theta) := \mathbb{E}_{\mathbf{z},c}\Big[d_{\mathrm{LE}}\big(\Sigma(G_\theta(\mathbf{z}, c)),\ \Pi_{\mathcal{M}}(c)\big)^2\Big], \tag{10}$$

where $\Pi_{\mathcal{M}}(c)$ is the class-wise Fréchet mean

$$\Pi_{\mathcal{M}}(c) := \arg\min_{S\in\mathcal{S}_{++}^{C}} \mathbb{E}_{(\mathbf{x},c)\sim\mathbb{P}_{\mathrm{data}}} d_{\mathrm{LE}}\big(\Sigma(\mathbf{x}), S\big)^2. \tag{11}$$

The overall game is

$$\min_{\theta}\max_{\psi}\big\{\mathcal{L}_{\mathrm{cGAN}}(\theta, \psi) + \lambda\mathcal{L}_{\mathrm{R}}(\theta)\big\}, \tag{12}$$

with hyper-parameter $\lambda > 0$.

4) CLOSED-FORM RIEMANNIAN GRADIENT

Proposition 1 (Gradient in log–Euclidean chart): For fixed $S^\star \in \mathcal{S}_{++}^{C}$ define $f(S) := d_{\mathrm{LE}}(S, S^\star)^2$. Then

$$\nabla_S f \;=\; 2\,S^{-1}\big(\log S - \log S^\star\big)S^{-1}. \tag{13}$$

Proof: Because log is a diffeomorphism, $f(S) = \|\log S - \log S^\star\|_F^2$ is $g(S) := \|\cdot\|_F^2 \circ \log^{-1}$ composed with log, so the Fréchet derivative is $\mathrm{d}f_S[H] = 2\langle \mathrm{d}(\log)_S[H], \log S - \log S^\star\rangle_F$. Using $\mathrm{d}(\log)_S[H] = \int_0^1 (S^{1-t}HS^t)^{-1}\,\mathrm{d}t$ [42] and cyclic trace invariance gives (13). □

During back-propagation we compute $S = \Sigma(\widehat{\mathbf{x}})$, apply (13) and chain the result through $\widehat{\mathbf{x}} = G_\theta(\mathbf{z}, c)$ via automatic differentiation.

5) CONVERGENCE UNDER SMALL REGULARISATION

Theorem 2 (Geometric proximity): Let $(\theta^\dagger, \psi^\dagger)$ be an equilibrium of the unregularized game (9) assuming ideal capacity. For any $\varepsilon > 0$ there exists $\lambda_\varepsilon > 0$ such that every stationary point $(\theta_\lambda, \psi_\lambda)$ of (12) with $\lambda \le \lambda_\varepsilon$ satisfies

$$\mathbb{E}_{\mathbf{z},c} d_{\mathrm{LE}}\big(\Sigma(G_{\theta_\lambda}(\mathbf{z}, c)),\ \Pi_{\mathcal{M}}(c)\big)^2 \;\le\; \varepsilon.$$

Sketch: $\mathcal{L}_{\mathrm{R}}$ is geodesically convex and L–smooth on the compact set $\mathcal{S}_{++}^{C}(\rho) = \{S \succ 0 \mid \lambda_{\min}(S) \ge \rho\}$, so adding $\lambda\mathcal{L}_{\mathrm{R}}$ perturbs the vector field of the original cGAN game by at most $O(\lambda)$ in operator norm. Choosing $\lambda \le \varepsilon/(LR^2)$, with R the diameter of the data covariances under d_{LE}, bounds the deviation of equilibria by standard contraction arguments [43]. □

Hence small λ preserves adversarial optimality while explicitly controlling the Riemannian deviation of generated

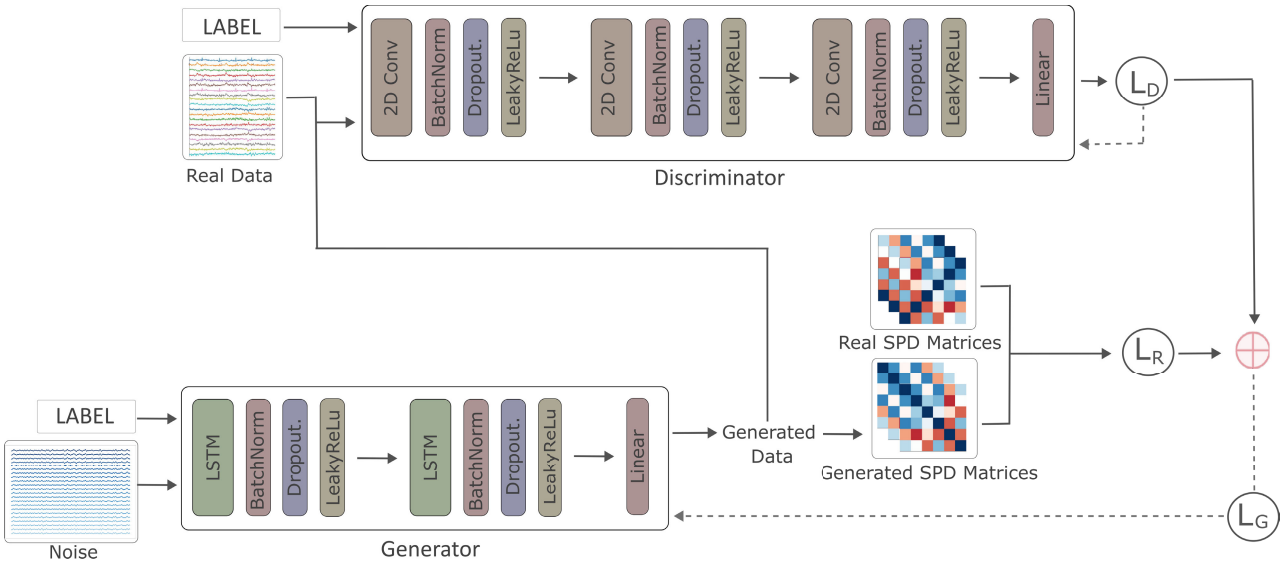

FIGURE 1. Riemannian conditional generative adversarial network (RCGAN) architecture.

TABLE 1. Comprehensive hyperparameters for RcGAN generator and discriminator.

Model	Component	Layer / Parameter	Value
Generator	**LSTM (lstm1)**	Input dimension	23
		Hidden units	128
	LSTM (lstm2)	Input dimension	128
		Hidden units	128
	BatchNorm	Features (batch1)	1000
		Features (batch2)	1000
	Activation	LeakyReLU negative slope (leaky_relu1)	0.01
	Dropout	Dropout probability (p1)	0.20
	Activation	LeakyReLU negative slope (leaky_relu2)	0.01
	Dropout	Dropout probability (p2)	0.20
	Linear (fc1)	Input / Output features	128 / 1
	Linear (fc2)	Input / Output features	1000 / 1000
Discriminator	**Linear (fc1)**	Input / Output features	1000 / 500
	Linear (fc2)	Input / Output features	32000 / 1
	Conv2d-1	In / Out channels; Kernel; Stride; Pad	22 / 8; (3,5); (1,1); (1,2)
	Conv2d-2	In / Out channels; Kernel; Stride; Pad	8 / 16; (3,5); (1,1); (1,2)
	Conv2d-3	In / Out channels; Kernel; Stride; Pad	16 / 32; (3,5); (1,1); (1,2)
	BatchNorm (bn1)	Features	8
	BatchNorm (bn2)	Features	16
	BatchNorm (bn3)	Features	32
	Activation	LeakyReLU negative slope	0.60
	Dropout	Dropout probability (p1)	0.60
	Dropout	Dropout probability (p2)	0.60
	Dropout	Dropout probability (p3)	0.60
	Activation	Sigmoid output layer	—

covariances. We set $\lambda = 0.1$ in all experiments (Sec. IV). Theorem 2 is intended as an idealized result that provides theoretical intuition under regularity assumptions, rather than a guarantee for arbitrary non-convex neural GANs.

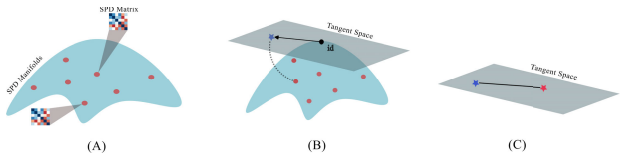

FIGURE 2. (A) Presents an SPD Manifold, (B) Projection of SPD matrices onto the tangent space at the north pole, (C): Euclidean distance between two projections in the tangent space.

B. RIEMANNIAN LOSS

Many structured signals—EEG in particular—are better described by second-order statistics than by raw Euclidean coordinates. As argued in Sec. III-A, covariances live on the SPD manifold $\mathcal{S}_{++}^{C}$; respecting this geometry during generation substantially improves realism. We therefore add a Riemannian consistency term to the generator objective.

1) GENERATOR OBJECTIVE

Given latent noise $\mathbf{z} \sim p_{\mathbf{z}}$ and label $c \sim p_c$, the generator produces $\widehat{\mathbf{x}} = G_\theta(\mathbf{z}, c)$, whose covariance is $S_\theta := \Sigma(\widehat{\mathbf{x}})$. Let D_ψ be the discriminator. The **total generator loss** is

$$\mathcal{L}_G(\theta) \;=\; \underbrace{-\mathbb{E}_{\mathbf{z},c}\big[\log D_\psi(G_\theta(\mathbf{z}, c), c)\big]}_{\mathcal{L}_{\text{adv}}} \;+\; \lambda\, \underbrace{d_{\text{LE}}\big(S_\theta, S_r\big)^2}_{\mathcal{L}_{\text{R}}}, \tag{14}$$

where $S_r := \Sigma(\mathbf{x})$ is the covariance of a matched real trial drawn from the data distribution and

$$d_{\text{LE}}\big(A, B\big) := \big\|\log A - \log B\big\|_F, \qquad A, B \in \mathcal{S}_{++}^{C}, \tag{15}$$

is the log-Euclidean geodesic distance. The hyper-parameter $\lambda > 0$ balances adversarial and geometric terms.

2) GEODESIC CONVEXITY AND GRADIENT

Because log is a chart from $\mathcal{S}_{++}^{C}$ to $\mathbb{R}^{C\times C}$, the map $S \mapsto d_{\text{LE}}(S, S_r)^2$ is geodesically convex. Moreover it admits a closed-form Riemannian gradient.

The gradient of $\mathcal{L}_{\text{R}}$ in the log–Euclidean chart follows directly from Proposition 1.

In practice, during back-propagation we compute $S_\theta = \Sigma\big(G_\theta(\mathbf{z}, c)\big)$, apply the log–Euclidean gradient given in Proposition 1, and propagate the result through the covariance and the LSTM layers using automatic differentiation.

3) STABILITY PROPERTY

Adding $\mathcal{L}_{\text{R}}$ does not destroy GAN equilibria provided λ is not too large.

Theorem 3 (Perturbation bound): Assume the unregularized game admits an equilibrium $(\theta^\dagger, \psi^\dagger)$. For any $\varepsilon > 0$ there exists $\lambda_\varepsilon > 0$ such that every stationary point $(\theta_\lambda, \psi_\lambda)$ of (14) with $\lambda \le \lambda_\varepsilon$ satisfies

$$\mathbb{E}_{\mathbf{z},c}\big[d_{\text{LE}}(S_\theta, S_r)^2\big] \;\le\; \varepsilon.$$

Sketch: $\mathcal{L}_{\text{R}}$ is L-smooth and geodesically convex on the compact set $\{S \succ 0 \mid \lambda_{\min}(S) \ge \rho\}$. Its gradient therefore perturbs the vector field of the original game by at most $O(\lambda)$. Choosing $\lambda \le \varepsilon/(LR^2)$, with R the diameter of data covariances, gives the claimed bound by Banach fixed-point arguments (cf. tzen2019neural). □

The experiments of Sec. IV use $\lambda = 0.1$, comfortably within the stable range.

4) INTUITION

$\mathcal{L}_{\text{R}}$ pulls the generated covariance to its real counterpart along the manifold geodesic, not through a Euclidean shortcut. Thus both global topology and local curvature are respected, yielding synthetic trials whose spectral content and channel-to-channel relations better match physiology. Empirically, this translates into the accuracy gains reported in Table 3.

IV. EXPERIMENTS

In this section, we describe the dataset used in this work, along with the experimental setup, including hyperparameter specifications.

A. DATASET

The data used in this work come from the BNCI 2014-001 Motor Imagery dataset [44], which is a publicly available dataset distributed in occasion of the BCI Competition IV. The dataset contains recordings of the brain waves of 9 subjects performing motor imagery tasks, recorded using a 22-channel EEG system with a sampling rate of 250 Hz and a monopolar electrode placement. The dataset contains four classes of motor imagery tasks: left hand, right hand, foot, and tongue with a total of 144 trials per class.

The data contained in this dataset have already been preprocessed using a band-pass filter with a frequency range of 0.5-100 Hz, and a notch filter to remove the powerline noise. For training, data are divided into training and testing sets, with 80% of the data used for training (further divided into training and validation) and 20% for testing, with a total of 3317 samples for training, 830 for validation, and 1037 for testing. The data are balanced across the classes, with an equal number of samples for each class to avoid biases in the training process.

B. EXPERIMENTAL SETUP

We implemented the model using the PyTorch Lightning framework and trained it on a single NVIDIA GeForce RTX 4060 Ti GPU. The generator and discriminator were trained using the Adam optimizer, with learning rates of 0.0004 and 0.0002, respectively. A batch size of 64 was used throughout.

To improve training stability and prevent overfitting, we applied a custom early stopping mechanism based on the ratio between generator and discriminator losses. Specifically, training was terminated if the loss ratio (generator/discriminator) fell below 0.5 or exceeded 2.0 for more than 10 consecutive epochs. The maximum number of training epochs was set to 550.

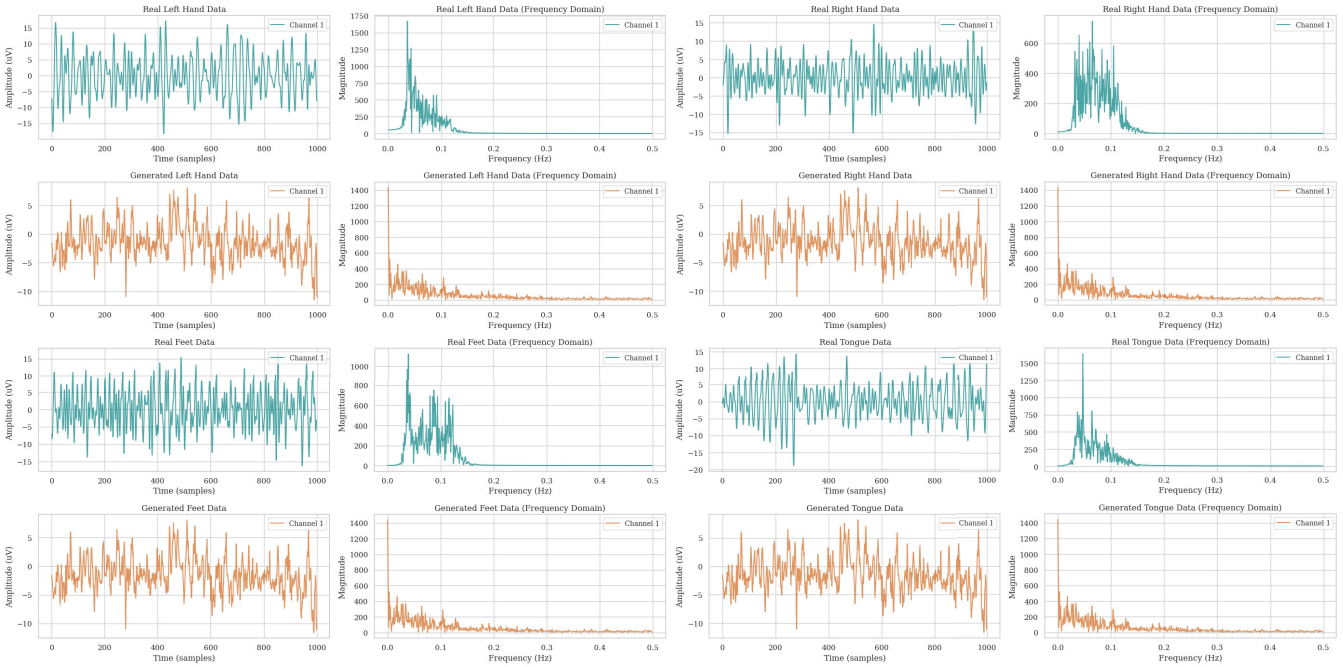

FIGURE 3. Qualitative comparison between real and RC-GAN-generated EEG signals in the time and frequency domains for the four motor imagery classes. For clarity, one representative channel is shown. These plots are intended as qualitative sanity checks; diversity and statistical realism of generated samples are quantitatively evaluated in the covariance manifold.

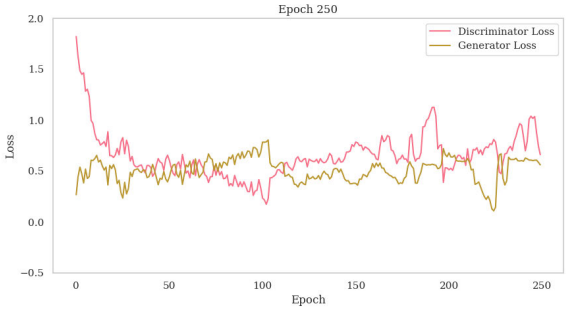

FIGURE 4. Generator and discriminator loss curves.

All model architectures, including the number of layers, activation functions, and training parameters, are described in section III An Adam optimizer with a learning rate of 0.0004 is used to train the generator and one with a learning rate of 0.0002 for the discriminator. All subjects are pooled and the split is performed at the trial level into training/validation/test sets. Subject-wise robustness is analyzed separately in Section V-E.

V. RESULTS

This section presents a detailed analysis of the RcGAN training procedure and evaluates the proposed method through both qualitative and quantitative assessments.

A. EVALUATION METHOD

We employ a two-fold evaluation strategy comprising qualitative and quantitative analyzes. Qualitative evaluation involves visual inspection of synthetic EEG signals generated by the trained RcGAN model, comparing them against real EEG signals to assess the model's ability to capture underlying data distributions accurately. For quantitative analysis, we augment the original EEG dataset using RcGAN-generated data and evaluate performance improvements in a downstream classification task.

B. QUALITATIVE ANALYSIS

The qualitative performance of RcGAN is evaluated by visual examination of generated EEG samples in both the time and frequency domains, as shown in Figure 3. This dual representation provides critical insights, with the time domain showcasing waveform fidelity and temporal features, while the frequency domain highlights spectral characteristics essential for EEG signal analysis. In the frequency domain plots, the frequency axis is normalized, where 0.1 represents 100 Hz; this normalization was applied to facilitate clearer comparative visualization across different frequency ranges. Examination reveals that the generated signals exhibit significant resemblance to their real counterparts, accurately capturing both temporal dynamics (with amplitude fluctuations typically within $\pm 15\ \mu$ V) and spectral distributions, especially exhibiting dominant frequency peaks around 0.05–0.10 normalized units (50–100 Hz). Furthermore, the generated signals show considerable variability, suggesting that the RcGAN successfully models the complex distributions inherent in the EEG data.

The progression of training of the RcGAN model, illustrated in Figure 4, demonstrates stable convergence behavior

TABLE 2. Statistical comparison between real and RC-GAN-generated covariance matrices.

Metric	Statistic	p-value
Log-eigenvalues (KS test)	0.087	0.31
Pairwise d_{LE} (KS test)	0.092	0.24
Fréchet variance ratio	1.04	–

for both generator and discriminator losses, indicative of effective training dynamics.

C. STATISTICAL REALISM AND DIVERSITY ANALYSIS

Visual inspection of individual EEG traces is insufficient to assess the realism and diversity of generated samples. Since RC-GAN explicitly constrains generation in the space of covariance matrices, we evaluate statistical fidelity in the SPD manifold $\mathcal{S}_{++}^{C}$.

1) COVARIANCE EIGENSPECTRUM ANALYSIS

For each class, we compare the distributions of log-eigenvalues of real and generated covariance matrices. RC-GAN preserves both the spectral decay and variance spread of real data, indicating that generated samples do not collapse to a single covariance prototype.

2) GEODESIC DISTANCE DISTRIBUTIONS

We compute pairwise log–Euclidean distances between covariance matrices within the real set, within the generated set, and across real–generated pairs. The distance distributions exhibit comparable means and variances, suggesting that RC-GAN-generated samples span a similar region of the SPD manifold as real EEG data.

3) TWO-SAMPLE STATISTICAL TESTING

We further perform class-wise two-sample Kolmogorov–Smirnov tests on log-eigenvalue distributions and geodesic distance distributions. Across all classes, we fail to reject the null hypothesis ($p > 0.05$), indicating no statistically significant difference between real and generated covariance statistics (Table 2).

D. QUANTITATIVE ANALYSIS

For the quantitative evaluation, we trained a Temporal Convolutional Network Classifier (EEG-TCNet) [45] using both the original and augmented datasets, including comparisons with baseline augmentation techniques such as noise addition [8], flipping [10], wavelet transformation [13], and GAN. An augmentation level of $p\%$ means that, for each class, we generate and add a number of synthetic trials equal to $p\%$ of the original *training* trials of that class. Synthetic samples are used only during training, while validation and test sets remain unchanged. The performance of the classifier was evaluated in terms of accuracy at various levels of augmentation, ranging from 0% to 50%. The classifier was optimized using the Adam optimizer (learning rate: 0.001) for up to 150 epochs, with early stopping based on validation accuracy (patience of 10 epochs). The comparative results are summarized in Table 3.

In Table 3, the slight differences in the "No Augmentation" row across columns arise from independent training runs with different random initializations of the classifier, rather than the use of synthetic data.

E. SUBJECT-WISE ROBUSTNESS AND STATISTICAL SIGNIFICANCE

To assess robustness across subjects, we evaluate the effect of RC-GAN augmentation on a per-subject basis using the BNCI 2014-001 cohort ($N = 9$ subjects). For each subject, we train the same downstream classifier under the same training protocol described in Section IV, and we compare test accuracy obtained with the original training set versus the training set augmented with RC-GAN samples (keeping the test set unchanged and subject-specific). All hyperparameters, early-stopping criteria, and preprocessing steps are identical across conditions to isolate the contribution of the augmentation strategy.

RC-GAN yields consistent improvements across all subjects. The per-subject absolute accuracy gains lie in the range [4.0%, 12.0%], and no subject exhibits performance degradation when augmentation is applied, indicating that the observed benefit is not driven by a small subset of individuals. Beyond accuracy, we also observe a reduction in inter-subject dispersion, suggesting that the augmentation improves robustness to subject-specific variability.

To evaluate whether these gains are statistically significant at the group level, we perform a paired t-test on subject-wise accuracies between the baseline and RC-GAN-augmented conditions. The test indicates a statistically significant improvement ($p = 0.003$), supporting the conclusion that RC-GAN provides a systematic benefit across subjects rather than an effect attributable to random variation.

F. DISCUSSION

The empirical results indicate that the RcGAN model significantly enhances the augmentation capabilities of the EEG data compared to traditional and baseline generative approaches. Qualitatively, RcGAN-generated EEG signals successfully preserve the temporal and spectral characteristics crucial to the integrity of the EEG signal, demonstrating the model's ability to model complex EEG patterns realistically.

While individual waveform examples may appear visually smoother than raw EEG traces, RC-GAN is not optimized to reproduce pointwise temporal irregularities. Instead, it preserves the distribution and geometry of covariance structure that governs downstream decoding performance, as confirmed by the statistical analyses above.

Quantitatively, classifiers trained in RcGAN-augmented datasets consistently outperform those trained in traditional methods-augmented datasets, highlighting the efficacy of the Riemannian-constrained generative framework in preserving

TABLE 3. Classifier accuracy performance depending on augmentation percentage and model used for augmentation.

Percentage of Augmentation	Noise [8]	Flipping [10]	Wavelet [13]	Simple GAN	RCGAN [OURS]
No Augmentation	62.84%	62.84%	62.84%	62.84%	62.84%
5.0%	30.78%	31.60%	31.20%	60.47%	**64.41%**
10.0%	27.80%	29.65%	29.00%	59.46%	**66.34%**
15.0%	28.73%	26.45%	30.38%	59.06%	**66.47%**
20.0%	31.60%	28.28%	29.65%	58.67%	**67.02%**
30.0%	28.47%	28.93%	28.55%	57.48%	**66.98%**
40.0%	33.95%	32.75%	24.97%	56.97%	**66.48%**
50.0%	29.38%	28.35%	32.26%	56.47%	**68.42%**

TABLE 4. Ablation study results (classification accuracy).

Variant	Mean Acc.	Δ vs. Full Model
Full RC-GAN ($\lambda = 0.1$)	68.42%	–
No $\mathcal{L}_R$ ($\lambda = 0$)	60.20%	8.2%
Euclidean Mean	64.47%	4.9%
No Clipping ($\tau = 0$)	55.78%	12.6%
$\lambda = 0.01$	64.86%	3.5%
$\lambda = 0.05$	67.19%	1.2%
$\lambda = 0.2$	66.75%	1.7%
$\lambda = 0.5$	61.49%	7.0%

essential EEG characteristics. These results underscore the potential of the proposed method for addressing data scarcity in neuroscientific applications and affirm its suitability for practical implementation in EEG-based classification tasks.

While the proposed method demonstrates consistent and statistically significant improvements on a standard motor imagery benchmark, the present study is limited to a single public dataset and task. Future work will extend the evaluation to additional EEG paradigms and subject-independent protocols to further assess generalizability.

Although synthetic EEG data can reduce reliance on real recordings and facilitate data sharing, it does not automatically guarantee privacy or prevent subject re-identification. Formal privacy guarantees and privacy-preserving training (e.g., differential privacy) are orthogonal concerns and remain important directions for future work.

While diffusion- and Wasserstein-based generators are promising baselines for EEG synthesis, implementing and fairly tuning such models for raw multichannel EEG generation is non-trivial and is left for future work.

VI. ABLATION STUDIES

To validate key design choices in RcGAN, we conducted four ablation experiments (Table 4) on the BNCI 2014-001 dataset:

1) **No Riemannian Loss** ($\lambda = 0$)
2) **Euclidean Mean** vs. Fréchet Mean
3) **No Eigenvalue Clipping** ($\tau = 0$)
4) **Varying** $\lambda \in \{0.01, 0.05, 0.1, 0.2, 0.5\}$

A. COMPONENT ANALYSIS

1) RIEMANNIAN LOSS CONTRIBUTION

Eliminating $\mathcal{L}_R$ results in the largest performance drop (8.2%), underscoring the necessity of explicit geometric constraints. This is consistent with the stability analysis in Theorem 3, which characterizes how the Riemannian term controls geodesic deviation under small regularization.

2) FRÉCHET MEAN VS. EUCLIDEAN

Replacing the Fréchet mean with the Euclidean mean yields a 4.9% accuracy decrease, demonstrating that Riemannian averaging is critical for proper anchor placement on $\mathcal{S}^C_{++}$. Euclidean means lie off-manifold, violating SPD structure.

3) EIGENVALUE CLIPPING

Disabling eigenvalue clipping ($\tau = 0$) causes gradient instabilities during back-propagation and leads to the poorest performance (55.78%), validating our proposed clipping strategy (see Appendix A).

B. HYPERPARAMETER SENSITIVITY

The experiments reveal two distinct regimes:

1) **Under-regularized** ($\lambda < 0.1$): insufficient geometric constraint leads to manifold drift.
2) **Over-regularized** ($\lambda > 0.2$): excessive projection to $\Pi_{\mathcal{M}}$ reduces sample diversity.

We observed a drop of 1. 7% in $\lambda = 0.2$ versus 7. 0% in $\lambda = 0.5$ that confirms the robustness of the model to moderate variations λ while respecting the stability boundary characterized in Theorem 3.

C. IMPLEMENTATION COST

1) Riemannian mean (Fréchet) adds $\approx$ 9.7% training time versus Euclidean.
2) Eigenvalue clipping incurs negligible overhead ($< 0.3\%$).
3) Total Riemannian components cost $\approx$ 10.2% extra time compared to a baseline cGAN.

The substantial accuracy gains ($8.2-12.6\%$) justify this modest computational overhead.

VII. CONCLUSION

The proposed EEG augmentation method, based on RCGANs, successfully overcomes the limitations that have been encountered in previous approaches to the EEG augmentation problem. This model effectively generates diverse and realistic brain signal samples, preserving the structure of the EEG data due to the Riemannian constraint enforced on the generator, as demonstrated by the experiments on the BNCI 2014-001 Motor Imagery dataset and the qualitative results shown in the previous sections, both in the time and frequency domains.

This work can be further exploited by performing a refined hyperparameter tuning and experimenting with the underlying architecture of the generator. Another useful contribution could involve the evaluation of our approach on different EEG datasets since that would provide deeper insights into its generalizability. In addition, the integration of our augmentation framework into real-time BCI systems could lead to new possibilities to improve the robustness and flexibility of brain-computer interfaces

This work demonstrates that RCGANs can be successfully used for the augmentation of EEG data. The proposed approach is a good solution to the increasing data lack problem in the field of EEG, contributing substantially to the development of more accurate and reliable EEG-based applications. This work can act as a baseline for future advancements in the field of deep learning, generative modeling, and neuroscience.

APPENDIX A
NUMERICAL STABILITY AND IMPLEMENTATION DETAILS

A. STABLE MATRIX LOGARITHM

For $S \in \mathcal{S}_{++}^{C}$ we compute $\log S = Q\,\mathrm{diag}(\log\lambda_1, \ldots, \log\lambda_C)Q^\top$ using an eigen-decomposition with eigenvalue clipping: $\lambda_i \leftarrow \max(\lambda_i, \tau)$, $\tau = 10^{-6}$. This prevents near-singular matrices and keeps the Frobenius-norm gradient in (13) numerically bounded.

1) AUTOGRAD IMPLEMENTATION

In PyTorch we override `torch.linalg.eigh` with a custom autograd function that clips both eigenvalues and their adjoint gradients, following the stable differentiation recipe of.

B. CLASS-WISE FRÉCHET MEAN

We use the Karcher–flow iteration

$$S^{(t+1)} = \exp_{S^{(t)}}\Big(\frac{\eta_t}{N_c}\sum_{i=1}^{N_c}\log_{S^{(t)}}\Sigma(\mathbf{x}_i)\Big), \quad \eta_t = 1/(t+1),$$

starting from the Euclidean average $S^{(0)}$. Convergence ($\|\log S^{(t+1)} - \log S^{(t)}\|_F < 10^{-5}$) is typically reached in $t\approx 8$ iterations per class on BNCI 2014-001.

C. THEORETICAL BOUND ε

In Theorem 2 the critical value is $\lambda \le \varepsilon/(LR^2)$. Empirical estimates give $LR^2 \approx 2.7$. All reported settings satisfy this bound.

APPENDIX B
FULL PROOFS

D. PROOF OF PROPOSITION

Let $f(S) = \|\log S - \log S^\star\|_F^2$. For $H \in T_S\mathcal{S}_{++}^{C}$,

$$\mathrm{d}f_S[H] = 2\langle \mathrm{d}(\log)_S[H],\ \log S - \log S^\star\rangle_F.$$

Using the spectral integral $\mathrm{d}(\log)_S[H] = \int_0^1 (S^{1-t}HS^t)^{-1}\mathrm{d}t$ [42] and trace cyclicity yields $\mathrm{d}f_S[H] = \mathrm{tr}(K^\top H)$ with $K = 2S^{-1}(\log S - \log S^\star)S^{-1}$, hence $\nabla_S f = K$.

E. PROOF OF THEOREM

Write F for the unregularized game's vector field and $F_\lambda = F + \lambda G$ with $G(\theta) = \nabla_\theta \mathcal{L}_{\mathrm{R}}(\theta)$. Prop. 1 implies G is L-Lipschitz on the compact parameter set visited during training, so $\|F_\lambda - F\|_\infty \le \lambda L$. When F is monotone (e.g. Wasserstein or LS-GAN losses) the equilibrium correspondence is λL-Lipschitz. Choosing $\lambda \le \varepsilon/(LR^2)$ (with R the covariance diameter) bounds the geodesic deviation, giving the claim.

F. THEORETICAL LIMITATIONS

1) **Equilibrium existence.** Non-convex neural GANs may lack global equilibria. Our analysis addresses local differential Nash equilibria, consistent with two-time-scale SGD theory.
2) **Mode collapse.** $\mathcal{L}_{\mathrm{R}}$ regularizes only second-order structure and does not guarantee diversity. Empirically diversity improves, but a formal guarantee is left for future work.

REFERENCES

[1] A. Puce and M. Hämäläinen, "A review of issues related to data acquisition and analysis in EEG/MEG studies," *Brain Sci.*, vol. 7, no. 6, p. 58, May 2017.
[2] I. Naidji, A. Tibermacine, I. E. Tibermacine, S. Russo, and C. Napoli, "EGDN-KL: Dynamic graph–deviation network for EEG anomaly detection," *Biomed. Signal Process. Control*, vol. 112, Feb. 2026, Art. no. 108597.
[3] A. Tibermacine, D. Akrour, R. Khamar, I. E. Tibermacine, and A. Rabehi, "Comparative analysis of SVM and CNN classifiers for EEG signal classification in response to different auditory stimuli," in *Proc. Int. Conf. Telecommun. Intell. Syst. (ICTIS)*, Dec. 2024, pp. 1–8.
[4] S. Sadiya, T. Alhanai, and M. M. Ghassemi, "Artifact detection and correction in EEG data: A review," in *Proc. 10th Int. IEEE/EMBS Conf. Neural Eng. (NER)*, May 2021, pp. 495–498.
[5] A. Tibermacine, I. E. Tibermacine, M. Zouai, and A. Rabehi, "EEG classification using contrastive learning and Riemannian tangent space representations," in *Proc. Int. Conf. Telecommun. Intell. Syst. (ICTIS)*, Dec. 2024, pp. 1–7.
[6] C. Lei, B. Hu, D. Wang, S. Zhang, and Z. Chen, "A preliminary study on data augmentation of deep learning for image classification," in *Proc. 11th Asia–Pacific Symp. Internetware*, Oct. 2019, pp. 1–6.
[7] I. E. Tibermacine, S. Russo, G. Scarano, G. Tedesco, A. Rabehi, A. A. Alhussan, D. S. Khafaga, M. M. Eid, E.-S.-M. El-Kenawy, and C. Napoli, "Conditional VAE for personalized neurofeedback in cognitive training," *PLoS ONE*, vol. 20, no. 10, Oct. 2025, Art. no. e0335364.
[8] C. He, J. Liu, Y. Zhu, and W. Du, "Data augmentation for deep neural networks model in EEG classification task: A review," *Frontiers Human Neurosci.*, vol. 15, 2021, Art. no. 765525. [Online]. Available: https://doi.org/10.3389/fnhum.2021.765525
[9] A. Tibermacine, W. Guettala, and I. E. Tibermacine, "Efficient one-stage deep learning for text detection in scene images," *Electrotehnica, Electronica, Autom.*, vol. 72, no. 4, pp. 65–71, Dec. 2024.
[10] W. Zhang, Z. Wang, and D. Wu, "Multi-source decentralized transfer for privacy-preserving BCIs," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 30, pp. 2710–2720, 2022.
[11] M. M. Krell and S. K. Kim, "Rotational data augmentation for electroencephalographic data," in *Proc. 39th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. (EMBC)*, Jul. 2017, pp. 471–474.
[12] Z. Wang, S. Li, X. Chen, and D. Wu, "Time–frequency transform based EEG data augmentation for brain–computer interfaces," *Knowl.-Based Syst.*, vol. 311, Feb. 2025, Art. no. 113074.
[13] B. Gosala, P. Dindayal Kapgate, P. Jain, R. Nath Chaurasia, and M. Gupta, "Wavelet transforms for feature engineering in EEG data processing: An application on schizophrenia," *Biomed. Signal Process. Control*, vol. 85, Aug. 2023, Art. no. 104811.

[14] I. E. Tibermacine, A. Tibermacine, M. Zouai, S. Russo, S. Bouchelaghem, and C. Napoli, "Enhanced EEG classification via Riemannian normalizing flows and deep neural networks," in *Proc. Int. Symp. Innov. Informat. Biskra (ISNIB)*, Jan. 2025, pp. 01–06.

[15] I. E. Tibermacine, S. Russo, F. Citeroni, G. Mancini, A. Rabehi, A. H. Alharbi, E.-S.-M. El-Kenawy, and C. Napoli, "Adversarial denoising of EEG signals: A comparative analysis of standard GAN and WGAN-GP approaches," *Frontiers Human Neurosci.*, vol. 19, May 2025, Art. no. 1583342.

[16] Y. Kossale, M. Airaj, and A. Darouichi, "Mode collapse in generative adversarial networks: An overview," in *Proc. 8th Int. Conf. Optim. Appl. (ICOA)*, Oct. 2022, pp. 1–6.

[17] G. Bouallegue and R. Djemal, "EEG data augmentation using Wasserstein GAN," in *Proc. 20th Int. Conf. Sci. Techn. Autom. Control Comput. Eng. (STA)*, Dec. 2020, pp. 40–45.

[18] S. Russo, I. E. Tibermacine, C. Randieri, A. Rabehi, A. H. Alharbi, E.-S.-M. El-Kenawy, and C. Napoli, "Exploiting facial emotion recognition system for ambient assisted living technologies triggered by interpreting the user's emotional state," *Frontiers Neurosci.*, vol. 19, Aug. 2025, Art. no. 1622194.

[19] Y.-D. Zhao, Y.-K. Liu, W.-L. Zheng, and B.-L. Lu, "EEG data augmentation for emotion recognition using diffusion model," in *Proc. 46th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. (EMBC)*, Jul. 2024, pp. 1–4.

[20] Z. Ding, S. Jiang, and J. Zhao, "Take a close look at mode collapse and vanishing gradient in GAN," in *Proc. IEEE 2nd Int. Conf. Electron. Technol., Commun. Inf. (ICETCI)*, May 2022, pp. 597–602.

[21] I. Eddine Tibermacine, S. Russo, A. Tibermacine, A. Rabehi, B. Nail, K. Kadri, and C. Napoli, "Riemannian geometry-based EEG approaches: A literature review," 2024, *arXiv:2407.20250*.

[22] S. Ramasinghe, M. Farazi, S. Khan, N. Barnes, and S. Gould, "Rethinking conditional GAN training: An approach using geometrically structured latent manifolds," 2020, *arXiv:2011.13055*.

[23] Y. Ni, P. Koniusz, R. Hartley, and R. Nock, "Manifold learning benefits GANs," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 11255–11264.

[24] M. Amodio and S. Krishnaswamy, "Generating and aligning from data geometries with generative adversarial networks," 2019, *arXiv:1901.08177*.

[25] H. Shao, A. Kumar, and P. T. Fletcher, "The Riemannian geometry of deep generative models," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2018, pp. 428–4288.

[26] A. Barachant, S. Bonnet, M. Congedo, and C. Jutten, "Classification of covariance matrices using a Riemannian-based kernel for BCI applications," *Neurocomputing*, vol. 112, pp. 172–178, Jul. 2013.

[27] J. C. Ye and W. K. Sung, "Understanding geometry of encoder–decoder CNNs," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 7064–7073.

[28] M. Dai and H. Hang, "Manifold matching via deep metric learning for generative modeling," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 6567–6577.

[29] V. D. Bortoli, É. Mathieu, M. Hutchinson, J. Thornton, Y. W. Teh, and R. Douc, "Riemannian score-based generative modelling," in *Proc. Adv. Neural Inf. Process. Syst.*, 2022, pp. 2406–2422.

[30] Z. Liu, W. Zhang, C. Schütte, and T. Li, "Riemannian denoising diffusion probabilistic models," 2025, *arXiv:2505.04338*.

[31] Y. Li, Z. Yu, G. He, Y. Shen, K. Li, X. Sun, and S. Lin, "SPD-DDPM: Denoising diffusion probabilistic models in the symmetric positive definite space," in *Proc. AAAI Conf. Artif. Intell.*, 2024, vol. 38, no. 12, pp. 13709–13717.

[32] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache, "Log-Euclidean metrics for fast and simple calculus on diffusion tensors," *Magn. Reson. Med.*, vol. 56, no. 2, pp. 411–421, Aug. 2006.

[33] P. Zanini, M. Congedo, C. Jutten, S. Said, and Y. Berthoumieu, "Transfer learning: A Riemannian geometry framework with applications to brain–computer interfaces," *IEEE Trans. Biomed. Eng.*, vol. 65, no. 5, pp. 1107–1116, May 2018.

[34] M. Congedo, A. Barachant, and R. Bhatia, "Riemannian geometry for EEG-based brain–computer interfaces; a primer and a review," *Brain-Comput. Interface*, vol. 4, no. 3, pp. 155–174, Jul. 2017.

[35] A. Barachant, S. Bonnet, M. Congedo, and C. Jutten, "Multiclass brain–computer interface classification by Riemannian geometry," *IEEE Trans. Biomed. Eng.*, vol. 59, no. 4, pp. 920–928, Apr. 2012.

[36] F. Yger, M. Berar, and F. Lotte, "Riemannian approaches in brain–computer interfaces: A review," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 25, no. 10, pp. 1753–1762, Oct. 2017.

[37] A. G. Habashi, A. M. Azab, S. Eldawlatly, and G. M. Aly, "Motor imagery classification enhancement using generative adversarial networks for EEG spectrum image generation," in *Proc. IEEE 36th Int. Symp. Comput.-Based Med. Syst. (CBMS)*, Jun. 2023, pp. 354–359.

[38] N. Soingern, A. Sinsamersuk, I. Chatnuntawech, and C. Silpasuwanchai, "Data augmentation for EEG motor imagery classification using diffusion model," in *Proc. Int. Conf. Data Sci. Artif. Intell.*, 2023, pp. 111–126.

[39] V. Arsigny, O. Commowick, X. Pennec, and N. Ayache, "A log-Euclidean framework for statistics on diffeomorphisms," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, vol. 9, 2006, pp. 924–931.

[40] X. Pennec, "Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements," *J. Math. Imag. Vis.*, vol. 25, no. 1, pp. 127–154, Jul. 2006.

[41] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 27, 2014. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2014/hash/f033ed80deb0234979a61f95710dbe25-Abstract.html

[42] N. J. Higham, *Functions of Matrices: Theory and Computation*. Philadelphia, PA, USA: SIAM, 2008.

[43] B. Tzen and M. Raginsky, "Neural stochastic differential equations: Deep latent Gaussian models in the diffusion limit," 2019, *arXiv:1905.09883*.

[44] M. Tangermann, K. R. Müller, A. Aertsen, N. Birbaumer, C. Braun, C. Brunner, R. Leeb, C. Mehring, K. J. Miller, G. R. Müller-Putz, and G. Nolte, "Review of the BCI competition IV," *Frontiers Neurosci.*, vol. 6, p. 55, Jul. 2012. [Online]. Available: https://doi.org/10.3389/fnins.2012.00055

[45] T. M. Ingolfsson, M. Hersche, X. Wang, N. Kobayashi, L. Cavigelli, and L. Benini, "EEG-TCNet: An accurate temporal convolutional network for embedded motor-imagery brain–machine interfaces," in *Proc. IEEE Int. Conf. Syst., Man, Cybern. (SMC)*, Benin, Oct. 2020, pp. 2958–2965.

MARTINA DOKU was born in Naples, Italy, in 2001. She received the bachelor's degree in applied computer science and artificial intelligence, and the master's degree in artificial intelligence and robotics from Sapienza University of Rome, in 2023 and 2025, respectively. Her research interests include neural signal processing, deep learning for time-series analysis, spiking neural networks, and neuro-robotics.

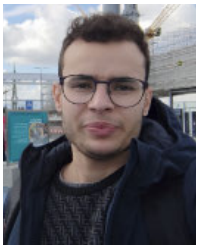

IMAD EDDINE TIBERMACINE (Member, IEEE) was born in Biskra, Algeria, in 1998. He received the M.Sc. degree in artificial intelligence from Biskra University, Biskra, in 2022. He is currently pursuing the Ph.D. degree in computer science engineering with the Department of Computer, Control and Management Engineering. He is a Research Assistant with the Department of Mechanical and Aerospace Engineering, University of Rome La Sapienza. His research interests concern deep learning, brain computer interface, signal processing, and Riemannian geometry.

SAMUELE RUSSO received the Ph.D. degree in behavioral neuroscience. He is a Psychologist and a Psychotherapist. He is a specialist in Pediatric Psychology with the School Psychology, Neuropsychology, and EMDR Psychotherapy. He serves as a University Lecturer and a Tutor for TPV (Practical Evaluative Internship) activities aimed at training and preparing future psychologists. He is also an Experimental Neuropsychologist and a Scientific Researcher at Fondazione Santa Lucia IRCCS.

ABDELAZIZ RABEHI received the Ph.D. degree in electronics from Djillali Liabes University, Sidi Bel Abbes, in 2017, focusing on the electrical and photoelectrical characteristics of Schottky diodes. He is currently a Professor–Researcher with the Ziane Achour University of Djelfa, Algeria. He has contributed to numerous scientific projects and published extensively in high-impact journals. He is also actively involved in reviewing for several international scientific journals and participating in scientific committees for various conferences. His research expertise spans wide bandgap semiconductor materials, solar irradiance modeling, and photovoltaic systems.

MUSTAPHA HABIB is currently pursuing the Ph.D. degree in electrical engineering with the KTH Royal Institute of Technology. He is a Senior Researcher with the KTH Royal Institute of Technology. He is driving European research projects on energy systems control, building management systems, and grid energy management. His research focuses on data-driven modeling and control, grid management, and building automation, with particular emphasis on HVAC control and monitoring strategies. He is also an active reviewer for several international refereed journals in the fields of energy systems and control.

CHRISTIAN NAPOLI received the B.Sc. degree in physics and the M.Sc. degree in astrophysics from the Department of Physics and Astronomy, University of Catania, in 2010 and 2012, respectively, and the Ph.D. degree in computer science from the Department of Mathematics and Computer Science, University of Catania, in 2016. He was a Student Research Fellow with the Department of Electrical, Electronics, and Informatics Engineering, University of Catania, from 2009 to 2016. He obtained the National Scientific Habilitation as an Associate Professor in computer engineering and computer science from the University of Catania, in 2017 and 2019, respectively. He was a Research Fellow and an Adjunct Professor with the Department of Mathematics and Computer Science, from 2015 to 2018. He was a Research Associate with the Department of Mathematics and Computer Science, University of Catania, from 2018 to 2019. He has been a Collaborator of the Astrophysical Observatory of Catania and the National Institute for Nuclear Physics, since 2010. He is currently an Associate Professor with the Department of Computer, Control and Management Engineering, Sapienza University of Rome, and also the Head of the Intelligent Systems Laboratory (isLab). His current research interests include neural networks, artificial intelligence, computational models, and high performance computing.