

Visualization of High-Dimensional Matrix Manifolds

Anonymous authors

Paper under double-blind review

Abstract

Matrix manifolds play a fundamental role in machine learning, underpinning data representations (*e.g.*, linear subspaces and covariance matrices) and optimization procedures. These manifolds adhere to Riemannian geometry, where intrinsic curvature significantly impacts the performance of geometric learning algorithms. However, traditional visualization methods based on Euclidean assumptions disregard curvature information, leading to distortions of the underlying non-Euclidean structure. To address this limitation, we generalize the popular t-SNE paradigm to the context of Riemannian manifolds and apply it to three types of matrix manifolds, which are the Grassmann manifolds, Correlation manifolds, and Symmetric Positive Semi-Definite (SPSD) manifolds, respectively. By constructing a probability distribution mapping between the original and target spaces, our method transforms high-dimensional manifold-valued data points into low-dimensional ones, preserving curvature information and avoiding distortion caused by Euclidean flattening. This work provides a foundation for general-purpose dimensionality reduction of high-dimensional matrix manifolds. Extensive experimental comparisons with existing visualization methods across synthetic and benchmarking datasets demonstrate the efficacy of our proposal in preserving geometric properties of the data.

1 Introduction

Matrix-valued data, such as linear subspaces, correlation matrices, and covariance matrices, naturally reside on non-Euclidean manifolds. While these descriptors can capture rich geometric and statistical information, their curved geometry fundamentally challenges standard visualization tools built upon Euclidean assumptions. In consequence, there is a growing need for geometry-aware visualization methods capable of faithfully preserving the intrinsic structure of manifold-valued data, thereby facilitating the development of non-Euclidean representation learning.

Grassmann manifolds. The Grassmann manifold $\text{Gr}(d, q)$ is the set of all q -dimensional linear subspaces in \mathbb{R}^d , which can be represented by $d \times q$ orthonormal basis matrix. As a fundamental subspace descriptor, it frequently appears in algorithms such as linear regression, principal component analysis Knudsen (2001); Jansson & Wahlberg (1996), low-rank matrix completion Vidal & Favaro (2014), and image set classification Wang et al. (2020), image fusion Kang et al. (2025). In addition, it demonstrates good performance in several downstream vision applications, such as action recognition Nguyen & Yang (2023); Chen et al. (2024), EEG decoding Ingolfsson et al. (2020); Wang et al. (2024), facial recognition Ingolfsson et al. (2020); Wang et al. (2021), and recommender systems Cao et al. (2016). Nevertheless, visualizing high-dimensional Grassmannian points remains a key challenge.

Correlation matrices. A correlation matrix is a symmetric matrix that encodes linear dependencies among variables, with each entry typically representing a Pearson correlation coefficient. Its key property is scale invariance, making it suitable for settings where absolute magnitudes are irrelevant Thanwerdas & Pennec (2022); Thanwerdas (2024). Taking the Electroencephalogram (EEG) analysis as an example, two electrodes may show a strong correlation despite differing signal amplitudes. Since the correlation matrix has the ability to capture intrinsic geometric relationships independent of scale, it has been widely used in fields such as brain connectivity analysis Varoquaux et al. (2010), finance Rebonato & Jäckel (2011); Marti et al.

```

from grass_imp import Grass_SNE
from cor_imp import Cor_SNE
from spsd_imp import SPSPD_SNE
# Visualization: Correlation matrices, Grassmann suspaces, and SPSPD matrices
embedding_h, labels, size = get_data()
# Correlation matrices implemented metrics: ECM, LEM, PHC, OLM
# Grassmannian implemented metrics: Geodesic, Asimov, Binet Cauchy, Chordal, Fubini
  Study, Martin, Procrustes, Projection, Spectral
# SPSPD matrices implemented metrics: Wasserstein
sne_embedding = run_SNE(embedding_h, manifold='Grassmannian', metric='Asimov', shape=
  size, t_gamma=1, learning_rate=0.05, max_iter=10000)
plot_low_dim(sne_embedding, labels)

```

Figure 1: A concise demo of using our proposed method

(2021), and Gaussian graphical models Lauritzen (1996); Epskamp & Fried (2018). Currently, research on the visualization of correlation matrices is still in its infancy.

SPSD matrices. An SPSPD matrix is a natural superset of SPD matrices, offering greater flexibility when dealing with low-rank or degenerate structures. SPSPD matrices are prevalent in machine learning and statistics, where they commonly appear as kernel matrices and covariance matrices Lanckriet et al. (2004); Huber & Ronchetti (1981). They are also widely encountered in practical domains such as medical imaging Bonnabel et al. (2013) and image set classification Faraki et al. (2016). By encoding both geometric and statistical information, SPSPD matrices provide expressive yet compact representations. This applicability makes SPSPD matrices a powerful modeling tool in scenarios where strict positive definiteness of high-dimensional data cannot be guaranteed. However, this also raises the question of how to visualize the SPSPD data points.

It is well known that visualization plays a central role in analyzing learned representations, evaluating model performance, and characterizing the structure of complex data. Although the aforementioned structured matrices are widely used across artificial intelligence and data science, their inherently non-Euclidean geometry poses substantial challenges for conventional analytical and visualization techniques. Classical methods like t-SNE van der Maaten & Hinton (2008), while highly effective for Euclidean data, assume a flat geometric structure and therefore introduce distortions when applied to curved manifolds. In particular, a point cloud sampled from a Riemannian manifold cannot be faithfully embedded into Euclidean space without geometric distortion. This phenomenon is analogous to the well-known impossibility of projecting the Earth, a curved 2D surface, onto a flat map without altering its geometry Mulcahy & Clarke (2001).

To this end, several geometry-aware extensions of t-SNE have been proposed. In hyperbolic settings, CO-SNE Guo et al. (2022) generalizes t-SNE to the Poincaré disk model to account for negative curvature. Recently, GrassCaré Li & Pimentel-Alarcón (2024) embeds high-dimensional subspace data into a 2D Poincaré disk, enabling curvature-aware visualization of Grassmannian representations. However, GrassCaré suffers from a fundamental curvature mismatch: the Poincaré disk has constant negative curvature, whereas the Grassmann manifold exhibits non-negative or mixed curvature. As a result, hyperbolic embeddings cannot faithfully preserve the intrinsic geometry or relational structure of the Grassmannian data.

To address these limitations, our method extends t-SNE into the Riemannian setting, enabling curvature-aware dimensionality reduction that preserves intrinsic manifold geometry. Specifically, our method produces interpretable low-dimensional embeddings that respect the structural properties of the Grassmann, Correlation, and SPSPD manifolds. For the low-dimensional case of the Grassmann manifold $\text{Gr}(3, 1)$, each point represents a one-dimensional subspace through the origin in \mathbb{R}^3 , which admits a natural visualization as a unit direction vector on the Grassmann hemisphere. On the Correlation manifold, the fixed-unit diagonal and degrees of freedom concentrated in the upper triangular part allow meaningful low-dimensional embeddings with 3×3 correlation matrices. Similarly, for SPSPD matrices, we use 2×2 matrices, whose three degrees of freedom yield compact yet informative representations. Figure 1 briefly illustrates how our approach is used.

Our main contributions are summarized as follows:

- **Riemannian t-SNE framework:** We propose a principled Riemannian extension of t-SNE tailored for matrix-valued data residing on three representative matrix manifolds. This method enables geometry-aware dimensionality reduction and faithful visualization of non-Euclidean data.
- **Visualization on the Grassmann manifold:** We incorporate four different metrics to capture diverse geometric properties of subspaces. Our method yields more interpretable and informative visualizations compared to existing methods that often overlook manifold curvature.
- **Visualization on the Correlation and SPSD manifolds:** For the correlation and SPSD matrices (two important yet underexplored subspaces of symmetric matrices), we design multiple metric-based strategies within our manifold-aware visualization framework. To the best of our knowledge, no prior work has explicitly explored the visualization problem on these two manifolds.
- **Comprehensive Empirical Evaluation:** Our proposed method is evaluated across synthetic data, image sets, EEG signals, and network features. Extensive experimental results demonstrate its effectiveness in preserving the intrinsic geometric structure of different data manifolds.

2 Related Work

Dimensionality Reduction (DR) is a core task in machine learning, aiming to reveal low-dimensional structures in high-dimensional data. DR methods are broadly categorized as linear, nonlinear, or Riemannian.

Linear DR methods. These methods project data onto subspaces that preserve key variance or discriminative information. Principal Component Analysis (PCA) Hastie (2009) identifies directions of maximum variance, while Independent Component Analysis (ICA) Hastie (2009) extracts statistically independent components. Additionally, Linear Discriminant Analysis (LDA) Izenman (2008), a supervised approach, seeks projections that maximize class separability.

Nonlinear DR methods. These approaches assume that data lies on a nonlinear manifold and aim to preserve its intrinsic geometry in the embedding. Representative methods include Multidimensional Scaling (MDS) Borg & Groenen (2007), Kernel PCA Schölkopf et al. (1997), Locally Linear Embedding (LLE) Roweis & Saul (2000), Isometric Feature Mapping (ISOMAP) Tenenbaum et al. (2000), Diffusion Maps (DMaps) Coifman & Lafon (2006), and Uniform Manifold Approximation and Projection (UMAP) McInnes et al. (2018). These methods seek to preserve the neighborhood or topological relationships during the embedding.

Riemannian DR methods. These methods generalize DR to data residing on curved manifolds by incorporating Riemannian geometry. Principal Geodesic Analysis (PGA) Fletcher et al. (2004) extends PCA by modeling variance along geodesics. Hessian LLE Goh & Vidal (2008) has also been applied to curved spaces. In hyperbolic space, HoroPCA Chami et al. (2021) captures hierarchical structure by identifying informative directions in negatively curved spaces, while CO-SNE Guo et al. (2022) develops t-SNE to the Poincaré disk for visualizing hierarchical data. For SPD matrices, de Surrel et al. (2025) extends t-SNE and MDS by introducing Riemannian geometry, such as the Affine-Invariant Riemannian Metric (AIRM) and the Log-Euclidean Metric (LEM) associated with the SPD manifolds. On the Grassmann manifolds, GrassCaré Li & Pimentel-Alarcón (2024) embeds subspace data into a 2D Poincaré disk, while GDMaps dos Santos et al. (2022) generalizes DMaps to the Grassmannian setting by defining appropriate distance measures and projection kernels.

However, these methods typically map curved manifold-valued data into Euclidean or fixed non-Euclidean spaces, without explicitly enforcing the preservation of the source manifold’s intrinsic Riemannian structure. Consequently, geometric properties specific to different curved manifolds may not be fully retained in the resulting embeddings. Motivated by these limitations, the goal of this paper is to develop a geometry-aware visualization framework for curved manifolds, enabling geometry-preserved manifold-to-manifold embedding mapping.

3 Preliminary

In this section, we briefly introduce these three manifolds.

Grassmann Manifolds. The Grassmann manifold $\text{Gr}(d, q)$ is the set of all q -dimensional subspaces of \mathbb{R}^d . It is a smooth, compact manifold of dimension $q(d - q)$ Bendokat et al. (2024). Each subspace on $\text{Gr}(d, q)$ can be represented by a matrix $U \in \mathbb{R}^{d \times q}$ with orthonormal columns, *i.e.*, $U^\top U = I_q$, where I_q is the $q \times q$ identity matrix. Different bases of the same subspace are associated by an orthogonal matrix, forming an equivalence class shown below:

$$[U] = \{\tilde{U} \mid \tilde{U} = UO, O \in \mathcal{O}(q)\}. \quad (1)$$

We refer to this as the orthonormal basis (ONB) representation and use U and $[U]$ interchangeably by abuse of notation.

Correlation Manifolds. A correlation matrix is obtained by normalizing a covariance matrix by its variances. Given covariance matrix $\Sigma = (\text{Cov}(X_i, X_j))_{1 \leq i, j \leq d}$, the corresponding correlation matrix has entries

$$C_{ij} = \frac{\Sigma_{ij}}{\sqrt{\Sigma_{ii}}\sqrt{\Sigma_{jj}}} = (\Upsilon^{-1/2})_{ii}\Sigma_{ij}(\Upsilon^{-1/2})_{jj}, \quad (2)$$

where $\Upsilon = \text{Diag}(\Sigma)$. This defines a mapping from SPD matrices to full-rank correlation matrices: $\Sigma \in \text{Sym}^+(d) \mapsto \Upsilon^{-1/2}\Sigma\Upsilon^{-1/2} \in \text{Cor}^+(d)$.

SPSD Manifolds. The SPSD manifold $\mathcal{S}_+(p, d)$ is the set of SPSD matrices of size $d \times d$ and rank $p \leq d$:

$$\mathcal{S}_+(p, d) = \{S \in \mathbb{R}^{d \times d} \mid S = S^\top, S \succeq 0, \text{rank}(S) = p\}. \quad (3)$$

When $p = d$, it reduces to the SPD manifold. This space admits a quotient structure: $\mathcal{S}_+(p, d) \cong \mathbb{R}_*^{d \times p} / \mathcal{O}_p$, where $\mathbb{R}_*^{d \times p}$ denotes full-rank matrices. Given any $F \in \mathbb{R}_*^{d \times p}$, the mapping $\Phi : F\mathcal{O}_p \mapsto FF^\top$ is well-defined and induces a smooth bijection between the quotient space and $\mathcal{S}_+(p, d)$.

4 Riemannian Dimensionality Reduction

t-SNE is a widely used method for nonlinear dimensionality reduction in Euclidean space. It transforms pairwise distances into conditional probabilities that reflect pointwise similarities. Given a high-dimensional point x_i , the probability of choosing x_j as its neighbor is computed by:

$$p_{j|i} = \frac{\exp(-\|x_i - x_j\|^2 / 2\sigma_i^2)}{\sum_{k \neq i} \exp(-\|x_i - x_k\|^2 / 2\sigma_i^2)}, \quad (4)$$

where σ_i is the Gaussian variance centered at x_i . The symmetric joint probability is then defined as $p_{ij} = \frac{p_{i|j} + p_{j|i}}{2N}$, where N signifies the number of data points.

To mitigate the ‘‘crowding problem’’ inherent in previous visualization methods such as SNE Hinton & Roweis (2002), t-SNE introduces a heavy-tailed t-distribution in the low-dimensional data space. This modification enhances repulsion between dissimilar points and prevents overcrowding in the embedding process. The low-dimensional similarity between y_i and y_j is quantified as:

$$q_{ij} = \frac{\exp(-\|y_i - y_j\|^2)}{\sum_{k \neq i} \exp(-\|y_i - y_k\|^2)}, \quad (5)$$

where y_i denotes the low-dimensional representation of x_i .

Finally, t-SNE minimizes the Kullback-Leibler (KL) divergence between the high-dimensional joint distribution and the low-dimensional distribution:

Considering that t-SNE is unable to be directly used in non-Euclidean space, we propose a Riemannian extension of t-SNE tailored for structured matrices residing on curved manifolds. Specifically, our proposed

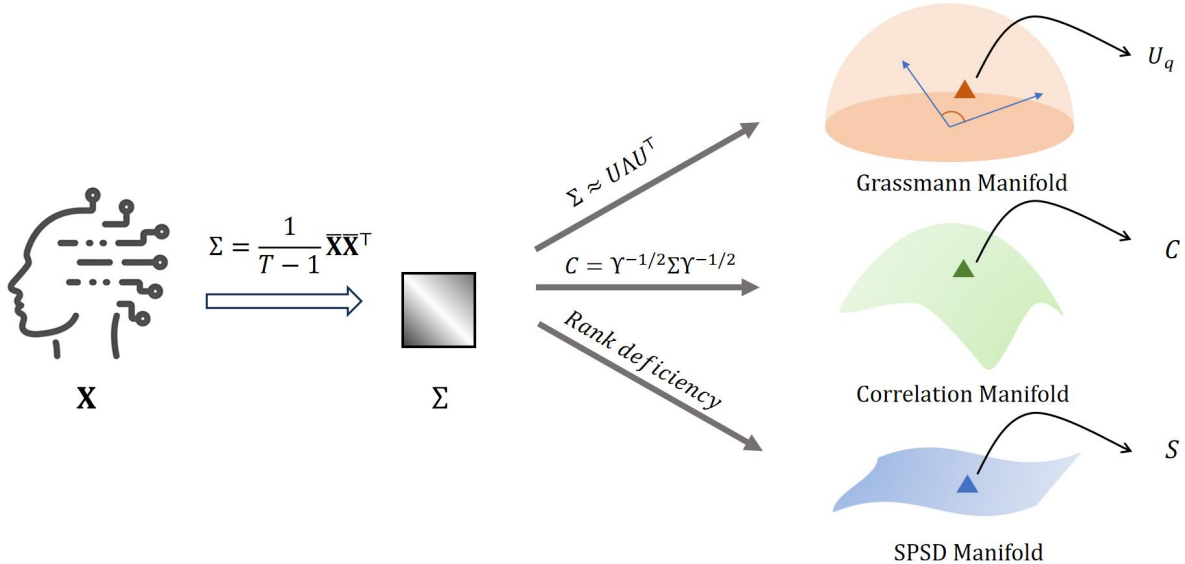


Figure 2: An example of modeling EEG data onto different matrix manifolds. For simplicity, we omit some of the preprocessing steps required for EEG data in this figure.

method comprises four key components: (1) **manifold-aware data modeling**, which embeds the input data onto their corresponding Riemannian manifolds (*e.g.*, Grassmann, Correlation, and SPSD manifolds); (2) **intrinsic similarity measurement**, where Riemannian distances are employed to faithfully quantify pairwise relationships; (3) **data distribution construction** on both the high-dimensional and low-dimensional manifolds; and (4) **geometry-aware optimization** for low-dimensional embedding. The following sections elaborate on each component and describe how they are integrated to form our proposed Riemannian t-SNE framework.

4.1 Manifold-aware Data Modeling

To leverage our proposed method, the data need to be transformed onto their corresponding matrix manifolds, firstly. A key observation is that many matrix manifolds used for modeling time-series data or image sets naturally arise from second-order statistics. In particular, the covariance representation provides a unified, canonical, and widely adopted way to capture the spatiotemporal dependencies among different data regions, serving as a common bridge for constructing Grassmannian, Correlation, and SPSD elements Wang et al. (2024); Hu et al. (2025); Wang et al. For clarity, we illustrate this modeling process using EEG signals as a typical example. Let $X \in \mathbb{R}^{d \times T}$ be an EEG trial, in which d and T signify the number of channels (*e.g.*, electrodes) and the length of time dimension per trial, respectively. Then, the corresponding sample covariance matrix can be computed by $\Sigma = \frac{1}{T-1} \bar{X} \bar{X}^\top$, where \bar{X} denotes the matrix obtained by centering X along the time dimension.

Accordingly, different manifold embeddings can be utilized to obtain the corresponding manifold-valued data:

- **Grassmann Manifold:** The column space spanned by the dominant q eigenvectors of Σ defines a point on $\text{Gr}(d, q)$. This is realized by Eigenvalue Decomposition, *i.e.*, $\Sigma = U\Lambda U^\top$, where $\Lambda = \text{diag}(\sigma_1, \dots, \sigma_d)$ with $\sigma_1 \geq \dots \geq \sigma_d \geq 0$ and $U = [u_1, \dots, u_d] \in \mathbb{R}^{d \times d}$ signify the eigenvalue and the corresponding eigenvector matrices, respectively. The truncated matrix $U_q = [u_1, \dots, u_q] \in \mathbb{R}^{d \times q}$ spans a q -dimensional linear subspace, represented as $[U_q] \in \text{Gr}(d, q)$.




Figure 3: Comparison of geodesic and Euclidean distances on different manifolds. (a) Different distances of a Grassmann manifold. (b) Different distances of a Correlation manifold. (c) Different distances of an SPSD manifold. Wherein, the red curve denotes the intrinsic geodesic, while the green line indicates the ambient Euclidean straight line.

- **Correlation Manifold:** Normalizing Σ by its diagonal yields $C = \Upsilon^{-1/2}\Sigma\Upsilon^{-1/2}$, where $\Upsilon = \text{Diag}(\Sigma)$ is a diagonal matrix consisting of the main diagonal elements of Σ . At this time, C resides on the Correlation manifold, a submanifold of SPD matrices with unit diagonal entries.
- **SPSD Manifold:** The covariance matrix Σ is inherently SPSD, encoding either full-rank or low-rank structure depending on the relationship between d and T . Its rank satisfies $\text{rank}(\Sigma) \leq \min(d, T-1)$. When $T-1 \geq d$, Σ is typically full-rank and lies in the interior of the SPSD cone \mathcal{S}_{++}^d , *i.e.*, the SPD manifold \mathcal{S}_{++}^d . When $T-1 < d$, Σ is rank-deficient and resides on the boundary of \mathcal{S}_{++}^d . In the latter case, Σ lies on a fixed-rank SPSD manifold.

For an intuitive understanding, please refer to the modeling process in Fig. 2.

4.2 Intrinsic Similarity Measurement

Unlike Euclidean spaces, where distances are computed along straight lines, Riemannian manifolds require intrinsic geodesic distances that reflect their underlying curvature. As a consequence, Euclidean metrics cannot be used in the proposed Riemannian t-SNE framework, as they disregard manifold geometry and will produce distorted similarity measurements (see Fig. 3). To ensure that pairwise relationships are faithfully reflected, we employ intrinsic Riemannian geodesic distances as the basis for similarity computation.

For the Grassmann manifold, the geodesic distance between two subspaces $U_1 \in \mathbb{R}^{d \times q}$ and $U_2 \in \mathbb{R}^{d \times q}$ is defined as Ye & Lim (2016):

$$d_{grass}(U_1, U_2) = \left(\sum_{i=1}^q \theta_i^2 \right)^{1/2}, \quad (6)$$

where θ_i ($i = 1 \rightarrow q$) are the principal angles between two subspaces. These angles are obtained from the singular value decomposition (SVD) of $U_1^\top U_2$. Let σ_i denote the i -th singular value, the corresponding principal angle is then given by $\theta_i = \arccos(\sigma_i)$. This distance captures the minimal rotational discrepancy between subspaces. As it is computed from the principal angles, it respects the underlying Riemannian geometry of the Grassmann manifold and provides a natural measure of subspace similarity.

For the Correlation manifolds, we adopt the Poly-Hyperbolic-Cholesky (PHC) metric Thanwerdas & Pennec (2022) to characterize the similarity between any two full-rank correlation matrices:

$$d_{cor}(C_1, C_2) = \left(\sum_{i=2}^n \left[\arccos \left(-Q \left(L_{1i}^\top, L_{2i}^\top \right) \right) \right]^2 \right)^{1/2}, \quad (7)$$

where $L_1 = \text{Chol}(C_1)$ and $L_2 = \text{Chol}(C_2)$ denote the Cholesky factors of C_1 and C_2 , respectively. Here, L_{1i} and L_{2i} represent the i -th rows of the lower triangular matrices, and $Q(x, y) = \sum_{j=1}^k x_j y_j - x_{k+1} y_{k+1}$ signifies the Lorentzian inner product.

The PHC metric is derived from the geodesic structure of hyperbolic geometry. To be specific, each correlation matrix is mapped (via its Cholesky decomposition) into a product space of hyperbolic manifolds \mathbb{H}^{i-1} , where L_{1i}^\top and $L_{2i}^\top \in \mathbb{R}^i$ are interpreted as points in hyperbolic space. Geodesic distances are computed in this hyperbolic product space and subsequently pulled back to the Correlation manifold via a smooth diffeomorphism, yielding an intrinsic and geometry-consistent similarity measure.

For the SPSD manifolds, we follow Massart & Absil (2020) and exploit the quotient-geometric distance for similarity measurement:

$$d_{psd}(S_1, S_2) = [\text{tr}(S_1) + \text{tr}(S_2) - 2 \text{tr}(S^*)]^{1/2}, \quad (8)$$

where $S^* = \left(S_1^{1/2} S_2 S_1^{1/2} \right)^{1/2}$ and $\text{tr}(\cdot)$ represents the matrix trace. This metric arises from the quotient geometry of the SPSD matrices, obtained by projecting the horizontal components of the equivalence-class representatives onto the quotient space. Importantly, the resulting geodesic distance preserves their intrinsic low-rank structure. In other words, the geodesic curve connecting S_1 and S_2 remains within the same fixed-rank stratum, ensuring that the manifold geometry does not artificially increase the rank or alter the underlying subspace structure.

In addition, several alternative Riemannian metrics are available on both the Grassmann and Correlation manifolds. Detailed descriptions of them are provided in Appendix A.1 and Appendix A.2. The metrics introduced above are the primary choices in the proposed framework. In the experiments, we further compare multiple metrics to assess their influence and to validate the rationale behind our selected benchmark.

4.3 Data Distribution Construction.

A core component of the proposed Riemannian t-SNE framework is the construction of probability distributions that characterize pairwise similarities on curved manifolds. In Euclidean t-SNE, these probabilities are derived from Gaussian and Student t-distributions. However, they are defined with respect to Euclidean distance and therefore cannot be directly applied to manifold-valued data. To tackle this limitation, we replace Euclidean distances with intrinsic geodesic distances and generalize the corresponding probability distributions accordingly.

We first introduce the Riemannian normal distribution, a natural analogue of the Euclidean Gaussian that maximizes entropy subject to prescribed expectation and variance Pennec (2006). Given a Fréchet mean μ and dispersion parameter σ , the distribution takes the following form

$$\mathcal{N}(X|\mu, \sigma^2) = \frac{1}{Z} \exp\left(-\frac{\delta(X, \mu)^2}{2\sigma^2}\right), \quad (9)$$

where $\delta(X, \mu)$ denotes the geodesic distance on the manifold and Z is the normalization constant. This construction ensures that probability mass is assigned according to the intrinsic geometry of the manifold, rather than a Euclidean approximation.

To define similarities in the low-dimensional manifold, we extend the Student t-distribution with one degree of freedom using the same geodesic principle. The resulting distribution is given by

$$q_{ij} = \frac{(1 + \delta(Y_i, Y_j)^2)^{-1}}{\sum_{k \neq l} (1 + \delta(Y_k, Y_l)^2)^{-1}}, \quad (10)$$

where Y_i represents the learned low-dimensional embedding of the i -th data point. This heavy-tailed distribution preserves the desirable separation properties of the original t-SNE, while respecting the manifold geometry through $\delta(\cdot, \cdot)$.

Finally, to align similarities between the high-dimensional and low-dimensional manifolds, we minimize the KL divergence between the high-dimensional distribution P and the low-dimensional distribution Q .

$$\mathcal{L} = \sum_i \text{KL}(P \parallel Q) = \sum_i \sum_j p_{ij} \log \frac{p_{ij}}{q_{ij}}. \quad (11)$$

This objective encourages points that are close in the original manifold to remain close after embedding, while allowing distant points to be placed more flexibly. In other words, it retains the key behaviors of Euclidean t-SNE, but reformulated in a manner that is consistent with Riemannian geometry.

4.4 Geometry-aware Optimization.

To obtain the low-dimensional manifold embeddings, we minimize the KL divergence loss \mathcal{L} through Riemannian gradient descent. The gradient of \mathcal{L} with respect to an embedding Y_i is given by

$$\nabla_{Y_i} \mathcal{L} = 2 \sum_j \frac{\partial \mathcal{L}}{\partial \delta_{ij}} \nabla_{Y_i} \delta_{ij} = 4 \sum_j \delta_{ij} (p_{ij} - q_{ij}) (1 + \delta_{ij}^2)^{-1} \nabla_{Y_i} \delta_{ij}, \quad (12)$$

where δ_{ij} denotes the Riemannian distance between Y_i and Y_j , and $\nabla_{Y_i} \delta_{ij}$ is the Riemannian gradient of the distance function.

To illustrate the computation of $\nabla_{Y_i} \delta_{ij}$, we consider the Grassmann manifold case using the geodesic distance in eq. (6). For two subspaces Y_i and Y_j , let $\{\sigma_\ell\}$ and $\{u_\ell, v_\ell\}$ denote the singular values and singular vectors of $Y_i^\top Y_j$. The Riemannian gradient takes the following form

$$\nabla_{Y_i} \delta_{ij} = (I - Y_i Y_i^\top) \left(\sum_{k=1}^q \frac{-\theta_k}{\sqrt{1 - \sigma_k^2}} u_k v_k^\top Y_j^\top \right), \quad (13)$$

where $\theta_k = \arccos(\sigma_k)$ are the principal angles, and $(I - Y_i Y_i^\top)$ is the orthogonal projection onto the tangent space $T_{Y_i} \text{Gr}(d, q)$.

Once the Riemannian gradient is computed, the parameters can be updated along the geodesic on the Grassmann manifold under the *Retraction* operation. This can be expressed as:

$$R(Y_i - \eta \nabla_{Y_i} \mathcal{L}) = \text{Exp}_{Y_i}(-\eta \nabla_{Y_i} \mathcal{L}), \quad (14)$$

where η is the learning rate and $\text{Exp}(\cdot)$ signifies the exponential map. Given a tangent vector Δ , the exponential map is formulated as:

$$\text{Exp}_{Y_i}(\Delta) = (Y_i V \cos \Lambda + U \sin \Lambda) V^\top, \quad (15)$$

where $U \Lambda V^\top$ is the SVD of Δ . This update ensures that the embedding remains on the Grassmann manifold throughout optimization. Additional implementation details are provided in Appendix A.3.

5 Experiments and Results

For each manifold, three types of experiments are conducted to verify the effectiveness of our method: 1) synthetic datasets sampled from a normal distribution, 2) benchmarking datasets, as well as intermediate representations (network features) extracted from designated layers, and 3) trustworthiness-based evaluation to quantify the extent to which global and local structures are preserved during dimensionality reduction.

5.1 Experiments on the Grassmann manifold

Setup. We compare our method against three baselines: GDMaps dos Santos et al. (2022), UMAP McInnes et al. (2018), and t-SNE van der Maaten & Hinton (2008). Both t-SNE and UMAP are widely used nonlinear




Figure 4: Visualization results of synthesized data points at different Grassmannian distances.

dimensionality reduction methods that rely on Euclidean distances to preserve data distributions in a low-dimensional space. GDMaps extends diffusion maps to the Grassmann manifold by incorporating manifold-specific distances and kernels. We adopt the implementations of t-SNE and UMAP in Pedregosa et al. (2011) with default parameters. For GDMaps, we take the implementation from Li & Pimentel-Alarcón (2024). Perplexity is an important parameter in t-SNE. In our method, we set it to $\frac{3}{4}N$, where N denotes the number of data points, consistent with the settings in de Surrel et al. (2025).

Synthetic experiments. To make an intuitive comparison between different distances on the Grassmann manifold (see Appendix A.1 for their definitions), we visualize synthesized data points using the proposed method under varying distance choices. Specifically, we simulate 5 point clusters in a 10-dimensional space. Each cluster contains 20 samples drawn from a Riemannian normal distribution on the Grassmann manifold $\text{Gr}(7, 2)$, where each point is represented as a 7×2 column-orthonormal matrix. The cluster centers are first sampled from a Riemannian normal distribution with unit variance. Around each center, 20 samples

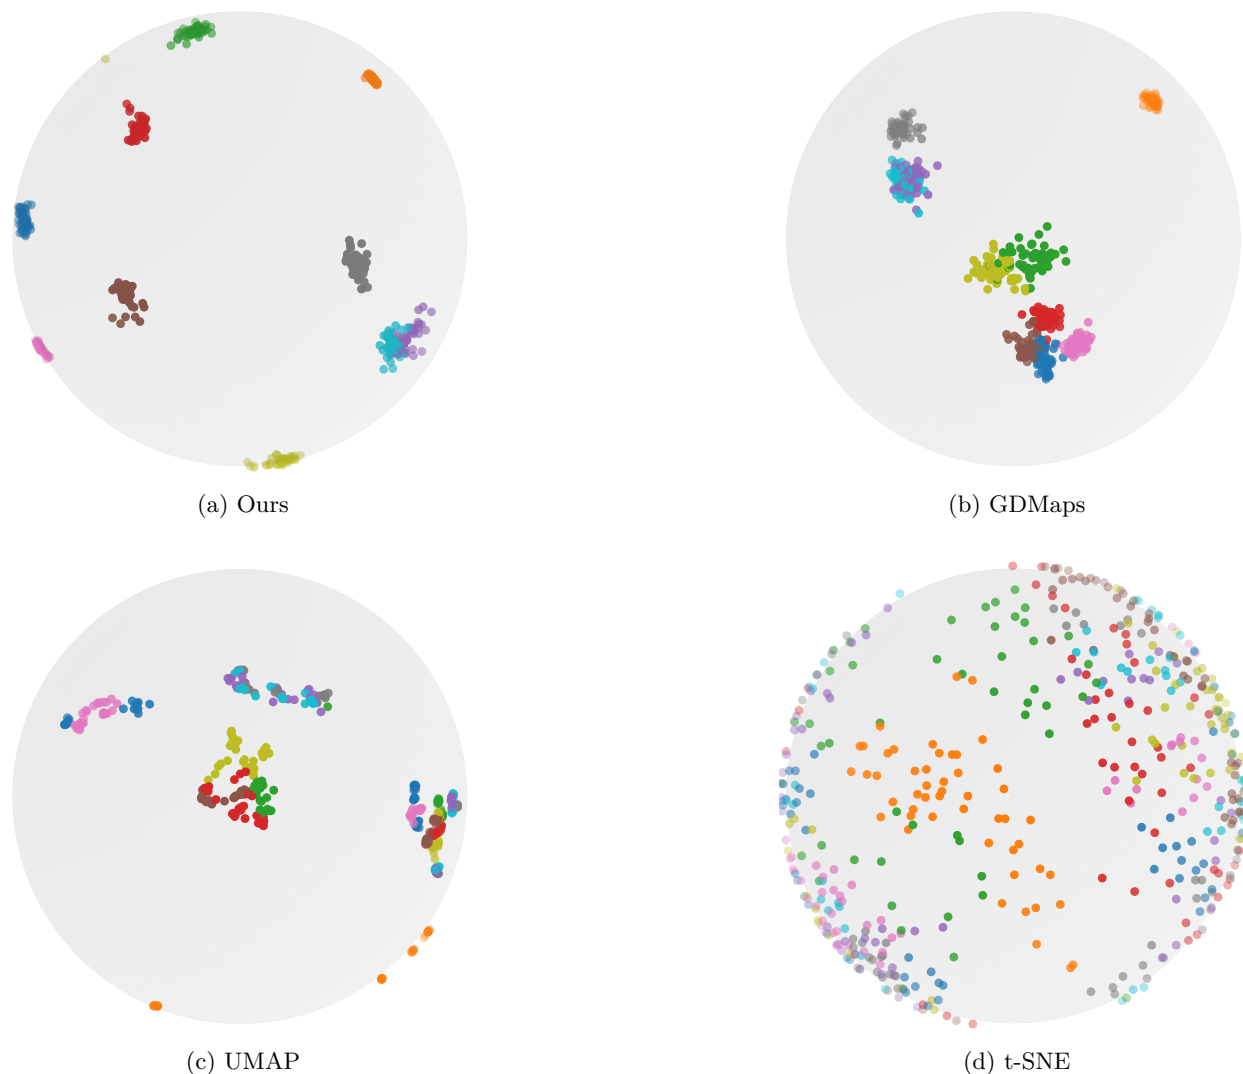


Figure 5: Visualization results of different methods on subspaces of the MNIST dataset.

are generated by adding small Gaussian noise (variance 0.25), followed by orthogonalization to ensure that these points reside on the Grassmann manifold. The results are shown in Fig. 4.

It can be observed that different distances lead to distinct embedding patterns across clusters. Specifically, the results in Fig. 4c and Fig. 4d lead to a severe collapse of high-dimensional data points belonging to the same category into nearly single points in the low-dimensional space. This phenomenon suggests that the projection and chordal distances fail to preserve the intra-cluster geometric structure, which is unfavorable for faithful manifold visualization. In contrast, Fig. 4b excessively stretches the intra-cluster structure, resulting in an over-dispersed embedding that blurs the boundaries between different categories and hinders discriminability. As expected, Fig. 4a provides a more balanced and structure-preserving visualization. This is because the geodesic distance corresponds to the length of geodesic curves induced by the canonical Riemannian metric on the Grassmann manifold. Unlike the chordal and projection distances, which rely on extrinsic Euclidean embeddings, the geodesic distance is defined intrinsically on the manifold itself. Consequently, it more faithfully captures the underlying curvature and geometric relationships of Grassmannian data points. Based on these observations, the geodesic distance is adopted in all subsequent experiments.

Benchmarking datasets. We further evaluate the proposed method on the MNIST dataset by visualizing the image groups represented as linear subspaces. Following the preprocessing procedures in Li & Pimentel-Alarcón (2024), we randomly sample 60 images from each digit class in the training set and form multiple groups. Each group is processed via Singular Value Decomposition, and the top 3 singular vectors are retained to construct a rank-3 subspace representation. Given the original image size of 28×28 , each group is thus modeled as a point on the Grassmann manifold $\text{Gr}(784, 3)$. We sample 50 such subspaces per class.

For ease of observation, the learned low-dimensional embeddings are displayed from a top-down view of the Grassmann hemisphere. As shown in Fig. 5, both our method and GDMaps produce well-separated clusters corresponding to different digit classes, whereas UMAP and t-SNE fail to yield clear class separation. This is mainly because UMAP and t-SNE operate in a flat Euclidean embedding space, which is not well suited for preserving the intrinsic Riemannian geometry of Grassmannian data, often leading to geometric distortions during dimensionality reduction.

Although GDMaps preserves global manifold structure by emphasizing diffusion distances, it exhibits limited discriminability when separating locally similar clusters, resulting in partial overlap among certain digit classes. This limitation is mainly attributed to the subspace construction stage rather than the diffusion process itself. Specifically, GDMaps relies on a truncated q -dimensional subspace representation for each sample prior to diffusion dos Santos et al. (2022). When q is small, geometrically informative directions may be irreversibly discarded, leading to a distorted Grassmannian geometry and consequently degenerate diffusion embeddings. In contrast, our method avoids hard spectral truncation and preserves local neighborhood relationships in a geometrically aligned manner, yielding embeddings that are less sensitive to the choice of subspace dimensionality and more robust for visualization.

Overall, these results demonstrate the effectiveness of our method in producing compact and well-separated embeddings, particularly for visually similar digit classes.

5.2 Experiments on the Correlation manifold

Setup. For experiments on the Correlation manifolds, we adopt a similar set of baselines as those used in the Grassmannian setting, including t-SNE and UMAP. Additionally, MDS Borg & Groenen (2007) is incorporated as a classical baseline. The implementations and hyperparameter settings of all the involved comparative methods are kept consistent with those used in the previous experiments.

Synthetic experiments. Following the experimental protocol on the Grassmann manifold, we evaluate the impact of different Riemannian distances on the visualization of full-rank correlation matrices using a synthetic dataset. It consists of three clusters, each containing 50 samples represented as 5×5 correlation matrices, which reside on a 10-dimensional Correlation manifold. Specifically, these matrices are generated by sampling from a Riemannian normal distribution with variance 0.2, followed by normalizing the diagonal entries to enforce unit variance.

The visualization results are illustrated in Fig. 6. It can be observed that the PHC distance yields clearer separation among different clusters compared to other Riemannian distances. We argue that it is mainly attributed to the favorable geometric and computational properties of the PHC metric on the Correlation manifold. In particular, the PHC distance defines a true geodesic metric and induces a non-positively curved Riemannian structure. Moreover, it admits unique Riemannian means and well-defined logarithmic mappings, while remaining numerically stable and computationally efficient Thanwerdas & Pennec (2022); Thanwerdas (2024). Based on these observations, we designate PHC as the benchmark distance on the Correlation manifold in all subsequent experiments.

Features from a manifold-valued EEG Network. We further evaluate the proposed method on the MAMEM-SSVEP-II dataset Georgiadis et al. (2016), which contains EEG recordings from 11 subjects. Direct visualization of raw EEG signals is often unreliable due to their low signal-to-noise ratio and susceptibility to various nuisance factors, such as sensor noise, power-line interference, and inter-trial as well as inter-subject variability. These factors can distort the correlation structure and degrade the fidelity of manifold embeddings. Therefore, on the correlation manifold, we first employ CorAtt Hu et al. (2025), a correlation-manifold attention network, to extract more stable and discriminative representations, and then perform

manifold-aware dimensionality reduction on the resulting intermediate geometric features from Subject 11. Specifically, we collect 100 feature samples evenly from 5 stimulus classes (20 samples per class), where each sample is a 15×15 correlation matrix.

The visualization results obtained by different dimensionality reduction methods are presented in Fig. 7. It can be seen that our method yields compact and well-separated clusters across the five classes, indicating that the class-discriminative structure of the manifold-valued features is effectively preserved. Although UMAP and MDS also exhibit a certain degree of class separation, they typically treat each correlation matrix as an ordinary data point in Euclidean space, either by directly vectorizing the matrix or by constructing distances in the vectorized space. In consequence, they do not explicitly account for intrinsic matrix constraints and geometry (*e.g.*, symmetry and the manifold properties induced by correlation matrices), which may lead to distortions of the underlying data structure in the low-dimensional embedding.

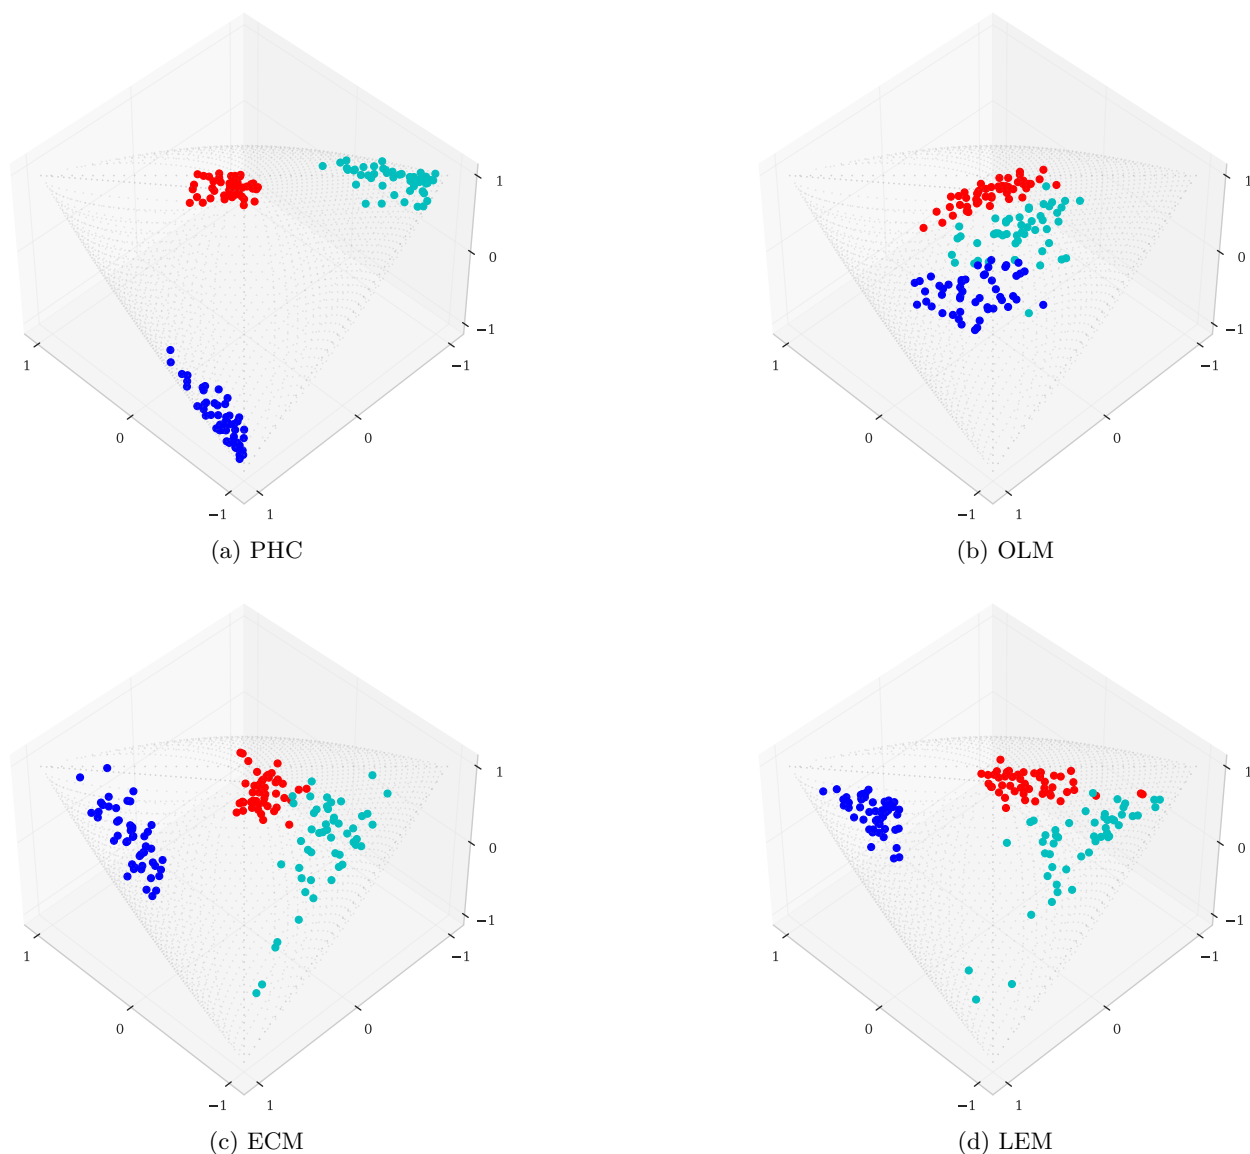


Figure 6: The impact of different Riemannian distances on the visualization of correlation matrices on a synthetic dataset.

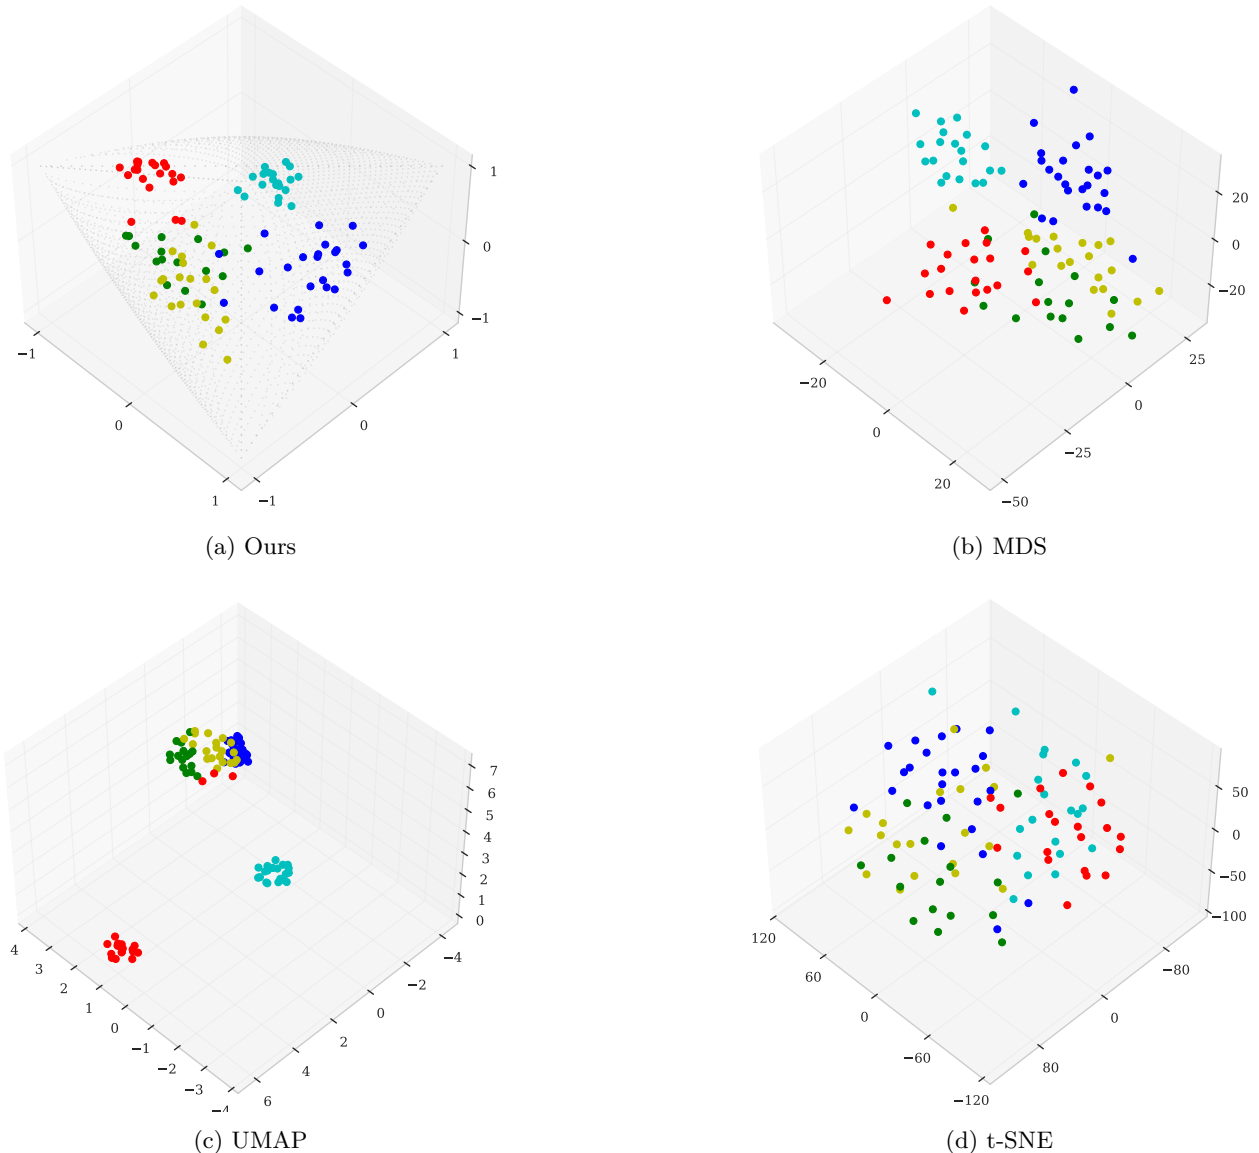


Figure 7: Visualization results of the manifold-valued features from an EEG network under different methods.

5.3 Experiments on the SPSP manifold

Setup. To the best of our knowledge, there currently exist no dedicated visualization methods specifically designed for the SPSP matrices. Since SPSP matrices and SPD matrices share closely related geometric structures, the SPD-SNE de Surrel et al. (2025) is adopted as a reasonable surrogate baseline. In this case, we impose a widely used regularization method on the aforementioned covariance matrix Σ , *i.e.*, $\Sigma + \epsilon I$, to produce the required SPD data points. In parallel, we generate the corresponding SPSP samples via Eigenvalue Decomposition by retaining the leading components that preserve 95% of the spectral energy. In addition, we use the official implementation provided in de Surrel et al. (2025) with default parameter settings. To ensure a fair comparison, our method is configured identically to that in the previous experiments.

Synthetic experiments. On the SPSP manifold, we also conduct a synthetic experiment in a 10-dimensional setting, where each data point is represented as a 4×4 SPSP matrix. Three clusters are generated, each containing 50 samples drawn from a Riemannian normal distribution with an intra-class variance of 0.4. The visualization results are shown in Fig. 8.




Figure 8: Visualizing a set of synthesized data points using the proposed method and SPD-SNE.

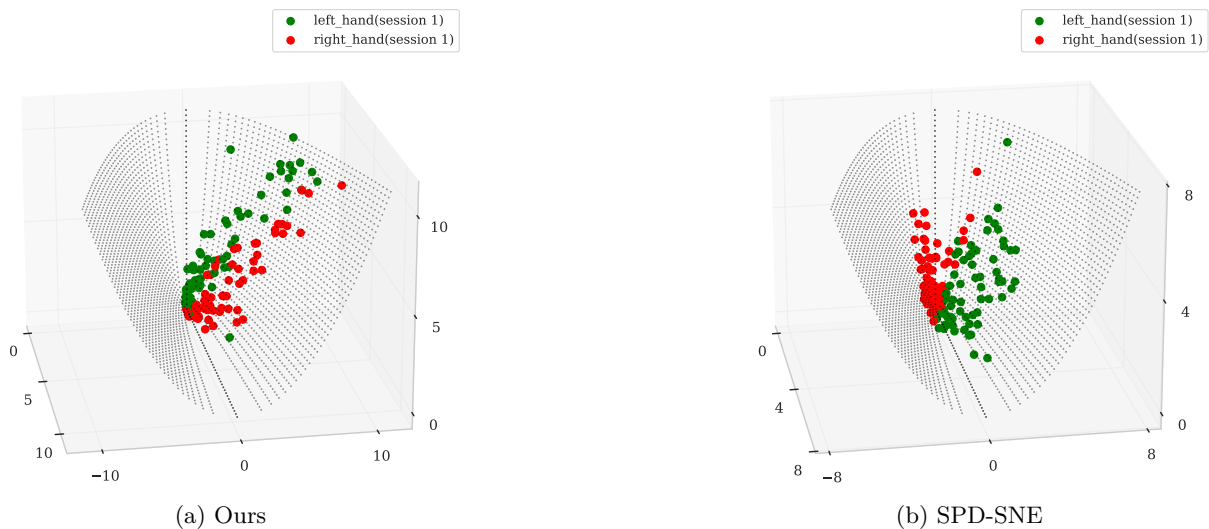


Figure 9: Visualizing EEG signal data using our method and SPD-SNE.

Both methods produce visually meaningful embeddings. However, they exhibit distinct distribution patterns due to the different geodesic structures they employ. To be specific, SPD-SNE tends to embed data points toward the interior of the positive definite cone, enforcing strict SPD constraints and resulting in clusters concentrated away from the boundary. In contrast, our proposed method distributes data points close to the cone’s boundary, thereby respecting the intrinsic geometry of SPSD matrices and better preserving the structural characteristics associated with rank-deficient or nearly singular samples.

EEG signal data. To the best of our knowledge, there are currently no established SPSD manifold neural architectures for EEG representation learning that can serve as a valid baseline. As a result, we directly apply the proposed visualization method to the raw EEG data for fair and reproducible evaluation. Specifically, we visualized the raw EEG data of Subject 8 in the BNCI2014001 dataset Tangermann et al. (2012), which is consistent with de Surrel et al. (2025). The experimental results are shown in Fig. 9.

It can be seen that both methods achieve a clear separation between the two categories. Note that the two visualizations are not directly comparable in a strict sense, since SPD-SNE assumes full-rank SPD inputs, whereas our method operates on low-rank SPSD matrices. This difference in matrix constraints

and underlying geometry naturally leads to different embeddings. Importantly, in the low-rank regime, our method preserves local neighborhoods more faithfully, yielding a higher trustworthiness coefficient (Fig. 12).

This improvement is expected because the low-rank SPSD construction suppresses eigen directions dominated by noise, which otherwise introduce instability into distance-based neighborhood relations. As a result, the embeddings better preserve local neighborhood structure, leading to higher trustworthiness scores.

5.4 Trustworthiness coefficient evaluation




Figure 10: Trustworthiness evaluation on the Grassmann manifold.




Figure 11: Trustworthiness evaluation on the Correlation manifold.




Figure 12: Trustworthiness evaluation on the SPSD manifold.

The trustworthiness coefficient quantifies how well neighborhood relationships are preserved during dimensionality reduction. Given a hyperparameter $n \in [1, k/2]$, where k is the number of data points, it compares the rankings of nearest neighbors in the original and embedded spaces, penalizing discrepancies. The score ranges from 0 to 1, with higher values indicating better local or global structure preservation.

We evaluate the visualization performance by measuring the trustworthiness in previous experiments. Specifically, we assess the benchmarking datasets on the Grassmann manifold, manifold-valued EEG network features on the Correlation manifold, and raw EEG signals on the SPSD manifold. Results are shown in Fig. 10, Fig. 11, and Fig. 12. It can be found that t-SNE and UMAP are not as good as their competitors. The

basic reason is that both methods are primarily designed for Euclidean data and typically rely on Euclidean distances. Hence, this will induce a metric mismatch and does not explicitly respect the intrinsic manifold geometry and constraints when applied to manifold data. In consequence, the original local geometric structure will be distorted during data embedding, leading to reduced trustworthiness.

As discussed in Sec. 5.1, GDMaps relies on a hard truncation of the fixed subspace dimension prior to the diffusion process, whereas our method preserves local subspace neighborhoods through geometric probability model alignment. As a result, our approach is less sensitive to the choice of subspace dimensionality and is more faithful to fine-scale geometric variations. This makes the proposed method more reliable than GDMaps in evaluating the trustworthiness across different neighborhood sizes.

Moreover, by minimizing global pairwise distance distortion Borg & Groenen (2007), MDS inherently prioritizes the preservation of global structure over local neighborhood relationships. This design enables MDS to achieve relatively high trustworthiness when the neighborhood size n is large. In contrast, our approach can be regarded as a manifold-valued t-SNE, which explicitly emphasizes the local neighborhood structure on the underlying manifold. This objective leads to superior trustworthiness performance in the small n regime, where faithful local neighborhood preservation is most critical. Notably, even for larger neighborhood sizes (*e.g.*, n approaching $k/2$), our method remains competitive with MDS, demonstrating a favorable balance between local fidelity and global consistency.

6 Conclusion

In this paper, we introduce a visualization framework for high-dimensional matrix-valued data defined on Riemannian manifolds, with particular emphasis on the Grassmannian subspaces, correlation matrices, and SPSD matrices. By explicitly incorporating intrinsic geodesic distances and curvature-aware probability distributions, the proposed method effectively preserves both local and global geometric structures during dimensionality reduction. Extensive experimental results on synthetic data points, benchmarking datasets, and manifold network features verify that our method consistently outperforms Euclidean geometry-based baselines as well as existing manifold-specific methods in capturing geometric relationships of manifold data.

References

- Thomas Bendokat, Ralf Zimmermann, and P-A Absil. A grassmann manifold handbook: Basic geometry and computational aspects. *Advances in Computational Mathematics*, 50(1):6, 2024.
- Silvère Bonnabel, Anne Collard, and Rodolphe Sepulchre. Rank-preserving geometric means of positive semi-definite matrices. *Linear Algebra and its Applications*, 438(8):3202–3216, 2013. ISSN 0024-3795. doi: <https://doi.org/10.1016/j.laa.2012.12.009>. URL <https://www.sciencedirect.com/science/article/pii/S0024379512008646>.
- Ingwer Borg and Patrick JF Groenen. *Modern multidimensional scaling: Theory and applications*. Springer Science & Business Media, 2007.
- Jiuwen Cao, Kai Zhang, Minxia Luo, Chun Yin, and Xiaoping Lai. Extreme learning machine and adaptive sparse representation for image classification. *Neural Networks*, 81:91–102, 2016. ISSN 0893-6080. doi: <https://doi.org/10.1016/j.neunet.2016.06.001>. URL <https://www.sciencedirect.com/science/article/pii/S0893608016300673>.
- Ines Chami, Albert Gu, Dat P Nguyen, and Christopher Re. Horopca: Hyperbolic dimensionality reduction via horospherical projections. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 1419–1429. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/chami21a.html>.
- Ziheng Chen, Yue Song, Yunmei Liu, and Nicu Sebe. A lie group approach to riemannian batch normalization. *arXiv preprint arXiv:2403.11261*, 2024.
- Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006. ISSN 1063-5203. doi: <https://doi.org/10.1016/j.acha.2006.04.006>. URL <https://www.sciencedirect.com/science/article/pii/S1063520306000546>. Special Issue: Diffusion Maps and Wavelets.
- Thibault de Surrél, Sylvain Chevallier, Fabien Lotte, and Florian Yger. Geometry-aware visualization of high dimensional symmetric positive definite matrices. *Transactions on Machine Learning Research Journal*, 2025.
- Ketson R. dos Santos, Dimitrios G. Giovanis, and Michael D. Shields. Grassmannian diffusion maps–based dimension reduction and classification for high-dimensional data. *SIAM Journal on Scientific Computing*, 44(2):B250–B274, 2022. doi: 10.1137/20M137001X. URL <https://doi.org/10.1137/20M137001X>.
- Sacha Epskamp and Eiko I Fried. A tutorial on regularized partial correlation networks. *Psychological methods*, 23(4):617, 2018.
- Masoud Faraki, Mehrtash T. Harandi, and Fatih Porikli. Image set classification by symmetric positive semi-definite matrices. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1–8, 2016. doi: 10.1109/WACV.2016.7477621.
- P.T. Fletcher, Conglin Lu, S.M. Pizer, and Sarang Joshi. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Transactions on Medical Imaging*, 23(8):995–1005, 2004. doi: 10.1109/TMI.2004.831793.
- K Georgiadis, G Liaros, VP Oikonomou, E Chatzilari, K Adam, S Nikolopoulos, and I Kompatsiaris. Mamem eeg ssvp dataset i (256 channels, 11 subjects, 5 frequencies), 2016.
- Alvina Goh and Rene Vidal. Clustering and dimensionality reduction on riemannian manifolds. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–7, 2008. doi: 10.1109/CVPR.2008.4587422.
- Yunhui Guo, Haoran Guo, and Stella X. Yu. Co-sne: Dimensionality reduction and visualization for hyperbolic data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 21–30, June 2022.

- Trevor Hastie. The elements of statistical learning: data mining, inference, and prediction, 2009.
- Geoffrey E Hinton and Sam Roweis. Stochastic neighbor embedding. In S. Becker, S. Thrun, and K. Obermayer (eds.), *Advances in Neural Information Processing Systems*, volume 15. MIT Press, 2002. URL https://proceedings.neurips.cc/paper_files/paper/2002/file/6150ccc6069bea6b5716254057a194ef-Paper.pdf.
- Chen Hu, Rui Wang, Xiaoning Song, Tao Zhou, Xiao-Jun Wu, Nicu Sebe, and Ziheng Chen. A correlation manifold self-attention network for eeg decoding. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, 2025.
- Peter J Huber and Elvezio M Ronchetti. Robust statistics john wiley & sons. *New York*, 1(1), 1981.
- Thorir Mar Ingolfsson, Michael Hersche, Xiaying Wang, Nobuaki Kobayashi, Lukas Cavigelli, and Luca Benini. Eeg-tcnet: An accurate temporal convolutional network for embedded motor-imagery brain-machine interfaces. In *2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 2958–2965. IEEE, 2020.
- Alan Julian Izenman. *Linear Discriminant Analysis*, pp. 237–280. Springer New York, New York, NY, 2008. doi: 10.1007/978-0-387-78189-1_8. URL https://doi.org/10.1007/978-0-387-78189-1_8.
- Magnus Jansson and Bo Wahlberg. A linear regression approach to state-space subspace system identification. *Signal Processing*, 52(2):103–129, 1996. ISSN 0165-1684. doi: [https://doi.org/10.1016/0165-1684\(96\)00048-5](https://doi.org/10.1016/0165-1684(96)00048-5). URL <https://www.sciencedirect.com/science/article/pii/0165168496000485>. Subspace Methods, Part II: System Identification.
- Huan Kang, Hui Li, Xiao-Jun Wu, Tianyang Xu, Rui Wang, Chunyang Cheng, and Josef Kittler. Grformer: A novel transformer on grassmann manifold for infrared and visible image fusion. *Information Fusion*, pp. 103402, 2025.
- Torben Knudsen. Consistency analysis of subspace identification methods based on a linear regression approach. *Automatica*, 37(1):81–89, 2001. ISSN 0005-1098. doi: [https://doi.org/10.1016/S0005-1098\(00\)00125-4](https://doi.org/10.1016/S0005-1098(00)00125-4). URL <https://www.sciencedirect.com/science/article/pii/S0005109800001254>.
- Gert R. G. Lanckriet, Nello Cristianini, Peter Bartlett, Laurent El Ghaoui, and Michael I. Jordan. Learning the kernel matrix with semidefinite programming. *J. Mach. Learn. Res.*, 5:27–72, December 2004. ISSN 1532-4435.
- S.L. Lauritzen. *Graphical Models*. Oxford Statistical Science Series. Clarendon Press, 1996. ISBN 9780191591228. URL <https://books.google.com.hk/books?id=mGQWkx4guhAC>.
- Huanran Li and Daniel Pimentel-Alarcón. Grasscaré: Visualizing the grassmannian on the poincaré disk. *SN Computer Science*, 5(3):298, 2024.
- Gautier Marti, Victor Goubet, and Frank Nielsen. ccorrigan: Conditional correlation gan for learning empirical conditional distributions in the ellipsope. In *International Conference on Geometric Science of Information*, pp. 613–620. Springer, 2021.
- Estelle Massart and P.-A. Absil. Quotient geometry with simple geodesics for the manifold of fixed-rank positive-semidefinite matrices. *SIAM Journal on Matrix Analysis and Applications*, 41(1):171–198, 2020. doi: 10.1137/18M1231389. URL <https://doi.org/10.1137/18M1231389>.
- Leland McInnes, John Healy, and James Melville. Umap: uniform manifold approximation and projection for dimension reduction. arxiv. *arXiv preprint arXiv:1802.03426*, 10, 2018.
- Karen A. Mulcahy and Keith C. Clarke. Symbolization of map projection distortion: A review. *Cartography and Geographic Information Science*, 28(3):167–182, 2001. doi: 10.1559/152304001782153044. URL <https://doi.org/10.1559/152304001782153044>.

- Xuan Son Nguyen and Shuo Yang. Building neural networks on matrix manifolds: A gyrovector space approach. In *International Conference on Machine Learning*, pp. 26031–26062. PMLR, 2023.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
- Xavier Pennec. Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127–154, 2006.
- Riccardo Rebonato and Peter Jäckel. The most general methodology to create a valid correlation matrix for risk management and option pricing purposes. *Available at SSRN 1969689*, 2011.
- Sam T. Roweis and Lawrence K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000. doi: 10.1126/science.290.5500.2323. URL <https://www.science.org/doi/abs/10.1126/science.290.5500.2323>.
- Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. Kernel principal component analysis. In Wulfram Gerstner, Alain Germond, Martin Hasler, and Jean-Daniel Nicoud (eds.), *Artificial Neural Networks — ICANN’97*, pp. 583–588, Berlin, Heidelberg, 1997. Springer Berlin Heidelberg. ISBN 978-3-540-69620-9.
- Michael Tangermann, Klaus-Robert Müller, Ad Aertsen, Niels Birbaumer, Christoph Braun, Clemens Brunner, Robert Leeb, Carsten Mehring, Kai J Miller, Gernot R Müller-Putz, et al. Review of the bci competition iv. *Frontiers in neuroscience*, 6:55, 2012.
- Joshua B. Tenenbaum, Vin de Silva, and John C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000. doi: 10.1126/science.290.5500.2319. URL <https://www.science.org/doi/abs/10.1126/science.290.5500.2319>.
- Yann Thanwerdas. Permutation-invariant log-euclidean geometries on full-rank correlation matrices. *SIAM Journal on Matrix Analysis and Applications*, 45(2):930–953, 2024. doi: 10.1137/22M1538144. URL <https://doi.org/10.1137/22M1538144>.
- Yann Thanwerdas and Xavier Pennec. Theoretically and computationally convenient geometries on full-rank correlation matrices. *SIAM Journal on Matrix Analysis and Applications*, 43(4):1851–1872, 2022. doi: 10.1137/22M1471729. URL <https://doi.org/10.1137/22M1471729>.
- Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008. URL <http://jmlr.org/papers/v9/vandermaaten08a.html>.
- Gaël Varoquaux, Flore Baronnet, Andreas Kleinschmidt, Pierre Fillard, and Bertrand Thirion. Detection of brain functional-connectivity difference in post-stroke patients using group-level covariance modeling. In *International conference on medical image computing and computer-assisted intervention*, pp. 200–208. Springer, 2010.
- René Vidal and Paolo Favaro. Low rank subspace clustering (lsrc). *Pattern Recognition Letters*, 43: 47–61, 2014. ISSN 0167-8655. doi: <https://doi.org/10.1016/j.patrec.2013.08.006>. URL <https://www.sciencedirect.com/science/article/pii/S0167865513003012>. ICPR2012 Awarded Papers.
- Rui Wang, Chen Hu, Xiaojun Wu, Xiaoning Song, Nicu Sebe, and Ziheng Chen. Gyroatt: A gyro attention framework for matrix manifolds.
- Rui Wang, Xiao-Jun Wu, and Josef Kittler. Graph embedding multi-kernel metric learning for image set classification with grassmannian manifold-valued features. *IEEE Transactions on Multimedia*, 23:228–242, 2020.
- Rui Wang, Xiao-Jun Wu, Zhen Liu, and Josef Kittler. Geometry-aware graph embedding projection metric learning for image set classification. *IEEE Transactions on Cognitive and Developmental Systems*, 14(3): 957–970, 2021.

Rui Wang, Chen Hu, Ziheng Chen, Xiao-Jun Wu, and Xiaoning Song. A grassmannian manifold self-attention network for signal classification. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pp. 5099–5107, 2024.

Ke Ye and Lek-Heng Lim. Schubert varieties and distances between subspaces of different dimensions. *SIAM Journal on Matrix Analysis and Applications*, 37(3):1176–1197, 2016. doi: 10.1137/15M1054201. URL <https://doi.org/10.1137/15M1054201>.

A Appendix

A.1 Distance metrics on Grassmannian manifolds

Metric	Riemannian distance
Chordal	$\left(\sum_{i=1}^k \sin^2 \theta_i\right)^{1/2}$
Martin	$\left(\log \prod_{i=1}^k 1/\cos^2 \theta_i\right)^{1/2}$
Projection	$\sin \theta_{max}$

Table 1: Distances on the Grassmann manifold Ye & Lim (2016).

Principal angles are a set of angles that measure the angle between two subspaces. They are calculated starting from the “closest” direction and can fully describe the relative position of the two subspaces. It can be defined as follows. Let $U_1, U_2 \subset \mathbb{R}^d$ be two subspaces with dimensions $\dim(U_1) = k$ and $\dim(U_2) = l$, respectively. Define $r = \min(k, l)$. The principal angles between U_1 and U_2 are denoted by

$$0 \leq \theta_1 \leq \theta_2 \leq \dots \leq \theta_r \leq \frac{\pi}{2},$$

and are defined recursively as follows.

For each $i = 1, \dots, r$, the i -th principal angle θ_i is defined via the optimization problem:

$$\cos \theta_i = \max_{\substack{p \in U_1, \|p\|=1, p \perp p_1, \dots, p_{i-1} \\ q \in U_2, \|q\|=1, q \perp q_1, \dots, q_{i-1}}} p^\top q, \quad (16)$$

where (p_j, q_j) are the principal vectors associated with the j -th principal angle θ_j for $j < i$.

Geometrically, each θ_i represents the smallest possible angle between directions in U_1 and U_2 that are orthogonal to the previously identified principal vectors. Collectively, the principal angles fully characterize the relative positions of the two subspaces up to rotation, and they can be computed via SVD.

Let $U_1 \in \mathbb{R}^{d \times k}$ and $U_2 \in \mathbb{R}^{d \times l}$ be matrices whose columns form orthonormal bases for U_1 and U_2 , respectively. The principal angles $\theta_1, \dots, \theta_r$ can be efficiently computed from SVD of the matrix $U_1^\top U_2 \in \mathbb{R}^{k \times l}$. Specifically, let

$$U_1^\top U_2 = W \Sigma Z^\top,$$

where $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_r)$, with singular values $\sigma_1 \geq \dots \geq \sigma_r \in [0, 1]$. Then the principal angles are given by:

$$\theta_i = \cos^{-1}(\sigma_i), \quad i = 1, \dots, r. \quad (17)$$

The associated principal vectors are given by the columns of $U_1 W$ and $U_2 Z$, corresponding to the left and right singular vectors of $U_1^\top U_2$, respectively.

Table 1 shows different distance metrics based on principal angles on the Grassmann manifold.

On Grassmann manifolds, common subspace distances include Chordal distance, Martin distance, and Projection distance. All three can characterize the relative positions between subspaces using principal angles, but they differ in the geometric features they focus on. Chordal distance considers all principal angles, equivalent to comparing the overall difference between the projection matrices of two subspaces; it is numerically

stable and insensitive to noise. Projection distance is determined only by the largest principal angle, characterizing the degree of deviation between two subspaces in the direction of greatest inconsistency, suitable for worst-case scenarios or scenes requiring strong discriminative power. In contrast, Martin distance, through a nonlinear transformation of the principal angles, emphasizes the differences when subspaces are nearly orthogonal, is more sensitive to geometric changes, but is relatively complex in numerical computation.

For more details, please refer to Ye & Lim (2016).

A.2 Distance metrics on Correlation manifolds

Metric	Riemannian Distance
ECM	$\ \phi(C_2) - \phi(C_1)\ _F$
LEM	$\ \log(\phi(C_2)) - \log(\phi(C_1))\ _F$
OLM	$\ \text{Off}(\log(C_2) - \log(C_1))\ _F$

Table 2: Distances on the correlation manifold Thanwerdas & Pennec (2022); Thanwerdas (2024).

The Euclidean-Cholesky metric (ECM) and the log-Euclidean-Cholesky metric (LEM) are both constructed by mapping correlation matrices to a lower triangular matrix space via the Cholesky decomposition, equipping this space with a Euclidean structure, and pulling the metric back to the original manifold. Both metrics are flat and geodesically complete, ensuring the uniqueness of the Fréchet mean. ECM directly uses a Euclidean metric on the Cholesky factors, leading to simple and efficient computations, whereas LEM applies a logarithmic transform in the Cholesky domain, yielding a scale-adaptive geometry at a slightly higher computational cost. In contrast, off-log metrics (OLM) rely on a log-Euclidean-type parametrization that separates diagonal and off-diagonal components without using a triangular factorization. OLMs preserve permutation invariance and also induce a flat geometry, but differ from ECM and LEM in both interpretation and numerical behavior. The definitions of these three distances are as follows:

The ECM distance. Let $C \in \text{Cor}^+(n)$ be a full-rank correlation matrix, and denote its normalized Cholesky factor by

$$\phi(C) = \text{Diag}(\text{Chol}(C))^{-1}\text{Chol}(C) \in \text{LT}_1(n), \quad (18)$$

where $\text{Chol}(C) \in \text{LT}_+(n)$ is the unique lower triangular matrix with positive diagonal such that $C = \text{Chol}(C) \cdot \text{Chol}(C)^\top$, and $\text{LT}_1(n)$ denotes the set of unit-diagonal lower triangular matrices.

The ECM distance between two correlation matrices $C_1, C_2 \in \text{Cor}^+(n)$ is defined as

$$d_{\text{EC}}(C_1, C_2) = \|\phi(C_1) - \phi(C_2)\|_F, \quad (19)$$

where $\|\cdot\|_F$ denotes the Frobenius norm. This distance corresponds to the Euclidean norm in the space $\text{LT}_1(n)$.

The LEM distance. Let $\log : \text{LT}_1(n) \rightarrow \text{LT}_0(n)$ denote the matrix logarithm, which is a diffeomorphism from the Lie group of unit-diagonal lower triangular matrices $\text{LT}_1(n)$ to its Lie algebra $\text{LT}_0(n)$, the space of strictly lower triangular matrices. The LEM distance between two correlation matrices $C_1, C_2 \in \text{Cor}^+(n)$ is defined by

$$d_{\text{LE}}(C_1, C_2) = \|\log(\phi(C_1)) - \log(\phi(C_2))\|_F, \quad (20)$$

where the logarithm is computed in the matrix sense. This metric corresponds to the Euclidean distance in the logarithmic coordinates of the Lie algebra $\text{LT}_0(n)$.

Next, we define the off-log distance as the geodesic distance induced by the off-log metric (OLM) on the manifold of full-rank correlation matrices $\text{Cor}^+(n)$.

The OLM distance. Let $\text{Log} : \text{Cor}^+(n) \rightarrow \text{Hol}(n)$ be the off-log diffeomorphism, given by

$$\text{Log}(C) := \text{Off}(\log(C)), \quad (21)$$

where $\text{Hol}(n)$ denotes the space of symmetric matrices with vanishing diagonal:

$$\text{Hol}(n) := \{X \in \text{Sym}(n) \mid \text{Diag}(X) = 0\}. \quad (22)$$

Given two correlation matrices $C_1, C_2 \in \text{Cor}^+(n)$, the off-log distance between them is defined by

$$d_{\text{OL}}(C_1, C_2) := \|\text{Log}(C_2) - \text{Log}(C_1)\|_{\text{Hol}}, \quad (23)$$

where $\|\cdot\|_{\text{Hol}}$ is the norm induced by a permutation-invariant inner product on $\text{Hol}(n)$. A commonly used instance is the Frobenius norm:

$$\|X\|_{\text{Hol}}^2 = \text{tr}(X^2), \quad \forall X \in \text{Hol}(n),$$

which leads to the following closed-form expression:

$$d_{\text{OL}}(C_1, C_2) = \|\text{Off}(\log(C_2) - \log(C_1))\|_F. \quad (24)$$

This distance is smooth, flat, and invariant under permutations of variables.

These three distance metrics are shown in Table 2. For more details, please refer to Thanwerdas & Pennec (2022); Thanwerdas (2024)

A.3 Derivation of gradients

We now compute the gradient of the loss function with respect to each embedded point on the Grassmann manifold. Recall that the objective function is the Kullback–Leibler divergence between the pairwise similarities in the high-dimensional and low-dimensional spaces:

$$C = \sum_i \text{KL}(P \parallel Q) = \sum_i \sum_{j \neq i} p_{ij} \log \frac{p_{ij}}{q_{ij}},$$

To compute the gradient with respect to the embedding Y_i , we first apply the chain rule:

$$\nabla_{Y_i} C = 2 \sum_{j \neq i} \frac{\partial C}{\partial \delta_{ij}} \nabla_{Y_i} \delta_{ij},$$

Computing $\frac{\partial C}{\partial \delta_{ij}}$

Let the normalization factor Z be defined as:

$$Z = \sum_{k \neq l} (1 + \delta_{kl}^2)^{-1},$$

so that the low-dimensional similarity q_{ij} becomes:

$$q_{ij} = \frac{(1 + \delta_{ij}^2)^{-1}}{Z}.$$

Rewriting the cost function C , we have:

$$C = \sum_i \sum_{j \neq i} p_{ij} \log p_{ij} - \sum_i \sum_{j \neq i} p_{ij} \log \left(\frac{(1 + \delta_{ij}^2)^{-1}}{Z} \right).$$

The first term is constant with respect to Y_i and thus can be ignored in the gradient. Taking derivatives yields:

$$\frac{\partial C}{\partial \delta_{ij}} = \frac{\partial}{\partial \delta_{ij}} [p_{ij} \log(1 + \delta_{ij}^2)] + \frac{\partial}{\partial \delta_{ij}} \left[\sum_{k \neq l} p_{kl} \log Z \right].$$

The first term evaluates to:

$$\frac{\partial}{\partial \delta_{ij}} [p_{ij} \log(1 + \delta_{ij}^2)] = p_{ij} \cdot \frac{2\delta_{ij}}{1 + \delta_{ij}^2}.$$

For the second term, since $\sum_{k \neq l} p_{kl} = 1$, and noting:

$$\frac{\partial Z}{\partial \delta_{ij}} = \frac{\partial}{\partial \delta_{ij}} \left[\sum_{k \neq l} (1 + \delta_{kl}^2)^{-1} \right] = -\frac{2\delta_{ij}}{(1 + \delta_{ij}^2)^2},$$

we obtain:

$$\frac{\partial}{\partial \delta_{ij}} \left[\sum_{k \neq l} p_{kl} \log Z \right] = \frac{1}{Z} \cdot \left(-\frac{2\delta_{ij}}{(1 + \delta_{ij}^2)^2} \right) = -\frac{2\delta_{ij} q_{ij}}{1 + \delta_{ij}^2}.$$

Therefore, the full derivative becomes:

$$\frac{\partial C}{\partial \delta_{ij}} = 2\delta_{ij} \cdot \frac{p_{ij} - q_{ij}}{1 + \delta_{ij}^2}.$$

Computing $\nabla_{Y_i} \delta_{ij}$

Recall that the geodesic distance on the Grassmann manifold is:

$$\delta_{ij} = d(Y_i, Y_j) = \left(\sum_{k=1}^q \theta_k^2 \right)^{1/2},$$

Then:

$$\nabla_{Y_i} \delta_{ij}^2 = 2 \sum_{k=1}^q \theta_k \nabla_{Y_i} \theta_k.$$

Since $\theta_k = \arccos(\sigma_k)$, where σ_k is the k -th singular value of $Y_i^\top Y_j$, we obtain:

$$\nabla_{Y_i} \theta_k = -\frac{1}{\sqrt{1 - \sigma_k^2}} \cdot \nabla_{Y_i} \sigma_k,$$

and from matrix calculus:

$$\nabla_{Y_i} \sigma_k = u_k v_k^\top Y_j^\top,$$

where u_k and v_k are the k -th left and right singular vectors of $Y_i^\top Y_j$, respectively.

Combining these results and projecting onto the tangent space of the Grassmann manifold, we obtain:

$$\nabla_{Y_i} \delta_{ij} = (I - Y_i Y_i^\top) \left(\sum_{k=1}^q \frac{-\theta_k}{\sqrt{1 - \sigma_k^2}} u_k v_k^\top Y_j^\top \right).$$

Final Expression

Substituting the expressions of $\frac{\partial C}{\partial \delta_{ij}}$ and $\nabla_{Y_i} \delta_{ij}$, we obtain the final gradient of the cost function:

$$\nabla_{Y_i} C = 4 \sum_{j \neq i} \delta_{ij} \cdot \frac{p_{ij} - q_{ij}}{1 + \delta_{ij}^2} \cdot \nabla_{Y_i} \delta_{ij}.$$