

Hyperbolic embedding of multilayer networks

Martin Guillemaud ^{*}

*Paris Brain Institute (ICM), CNRS UMR7225, Inserm U1127, Sorbonne University UM75,
Inria-Paris, Pitié Salpêtrière University Hospital, Paris, France*

Vera Dinkelacker 

Department of Neurology, Strasbourg University Hospital, Strasbourg, France

Mario Chavez 

CNRS, Pitié Salpêtrière University Hospital, Paris, France



(Received 28 May 2025; accepted 1 October 2025; published 1 December 2025)

Multilayer networks offer a powerful framework for modeling complex systems across diverse domains, effectively capturing multiple types of connections and interdependent subsystems commonly found in real-world scenarios. To analyze these networks, embedding techniques that project nodes into a lower-dimensional geometric space are essential. This paper introduces a novel hyperbolic embedding framework that advances the state of the art in multilayer network analysis. Our method, which supports heterogeneous node sets across networks and interlayer connections, generates layer-specific hyperbolic embeddings, enabling detailed intralayer analysis and interlayer comparisons, while simultaneously preserving the global multilayer structure within hyperbolic space—a capability that sets it apart from existing approaches, which typically rely on independent embedding of layers. Through experiments on synthetic multilayer stochastic block models, we demonstrate that our approach effectively preserves community structure, even when layers consist of different node sets. When applied to real brain networks, the method successfully clusters disease-related brain regions from different patients, outperforming layer-independent approaches and highlighting its relevance for comparative analysis. Overall, this work provides a robust tool for multilayer network analysis, enhancing interpretability and offering new insights into the structure and function of complex systems.

DOI: [10.1103/7wd9-dwlr](https://doi.org/10.1103/7wd9-dwlr)

I. INTRODUCTION

The network paradigm has proven to be a powerful and versatile framework for representing and analyzing complex systems across a wide range of domains, from social dynamics and economic systems to neuroscience and biological interactions [1]. In many real-world cases, however, interactions within a system are not adequately captured by a single-layer network, as multiple types of connections or interdependent subsystems often coexist. Multilayer network models were introduced to overcome these limitations, offering a comprehensive and structured way to represent complex systems with multiple dimensions of interaction [2,3]. In this framework, each layer typically corresponds to a specific mode of interaction (or a temporal snapshot in a time-varying model), and nodes can participate in one or several layers. Among the various types of multilayer networks, multiplex networks (where the same set of nodes interact through different types

of links across layers) constitute a particularly common subclass. Multilayer models have shown great utility in diverse applications, including the study of brain connectivity, community detection, and the dynamics of spreading processes [2,4,5].

In recent years, graph embedding techniques have emerged as an essential tool for simplifying and analyzing large-scale networked data. The core idea of graph embedding is to project nodes into a lower-dimensional geometric space such that their proximity reflects structural similarity in the original network [6,7]. By representing networks as vectors or functions, a wide variety of machine learning algorithms can be used to provide a solution for various downstream tasks such as network visualization, link prediction, node classification, and node clustering [8]. A broad array of embedding techniques has been developed for single-layer networks, relying on matrix factorization, random walks, or machine learning approaches [6,7].

Extending these embedding frameworks to multilayer networks has attracted growing interest, leading to a variety of methods that seek to preserve both intra- and interlayer structural information. These include spectral methods [9], random walk-based approaches [10,11], and information-theory based techniques such as mutual information maximization across layers [12,13]. Other strategies use deep learning frameworks (e.g., autoencoders) to generate embeddings useful for

^{*}Contact author: martin.guillemaud@gmail.com

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

multilayer clustering [14], or apply tensor factorization techniques for network decomposition and analysis, particularly in neuroscience [15]. However, most of these methods ultimately produce a single unified embedding for the entire multilayer network. While such summarization is often effective, it may overlook important layer-specific features. In scenarios where layer-level differences are meaningful (such as in interlayer influence quantification, or comparative studies of multiple networks) having one embedding per layer, while still accounting for multilayer dependencies, is crucial. To our knowledge, few methods explicitly address this need. One such exception is the optimization-based approach of [16], which constructs layer-specific embeddings that reflect both within-layer and cross-layer dynamics.

Hyperbolic geometry has emerged as a promising framework for graph representation, as it naturally aligns with the structural characteristics of many real-world networks, including hierarchical organization, scale-free degree distributions, and high clustering coefficients [17–19]. In treelike structures, for example, the number of nodes increases exponentially with distance from high-hierarchy regions, necessitating dimensionalities that exceed the representational limits of Euclidean space. Moreover, increasing the dimensionality in Euclidean embeddings results in greater computational complexity and significant distortion [20]. These challenges have driven the development of models that embed networks in hyperbolic space, such as PSO [21], HyperMap [22], Mercator [23], and its high-dimensional extension, D-Mercator [24]. In parallel, machine learning-based approaches like coalescent embedding [25] and its extensions [26] have provided scalable, data-driven alternatives for hyperbolic network embedding.

Despite these developments, the application of hyperbolic embedding to multilayer networks remains largely unexplored. A common strategy involves aggregating all layers into a single connectivity matrix (typically by summing or averaging the adjacency matrices) before applying conventional hyperbolic embedding methods. However, this simplification obscures the inherent multilayer structure and fails to capture the complex interactions across different types of connections. Temporal networks, often represented as multilayer graphs with each layer corresponding to a distinct time step, have received comparatively greater attention. For instance, in Ref. [27], a method was proposed for embedding temporal graphs into the hyperboloid model using unsupervised random walks. In Ref. [28], an alternative approach was introduced, embedding temporal graphs into the Poincaré ball through a memory mechanism and tangent space optimization. Notably, in Ref. [29], it was demonstrated that embedding each layer of a multilayer network independently into the hyperbolic disk, while analyzing them collectively, improves performance in tasks such as community detection, link prediction, and navigability. These findings underscore the importance of preserving layer-specific information within hyperbolic representations to fully exploit the structural complexity of multilayer networks.

In this work, we introduce a novel multilayer embedding framework in hyperbolic space that addresses these limitations. Our method produces a separate embedding for each layer while accounting for the global multilayer structure,

enabling both intralayer analysis and interlayer comparison. This approach opens new avenues for understanding complex systems through the lens of hyperbolic geometry, with applications ranging from brain networks to social systems.

II. METHODS

Hyperbolic geometry concerns spaces characterized by constant negative curvature K , which deviate fundamentally from the principles of Euclidean geometry. In hyperbolic embedding, a graph is first mapped onto a hyperboloid and can subsequently be projected onto a two-dimensional model of hyperbolic space, such as the Poincaré disk or the Klein disk model. The Poincaré disk is the unit disk $\mathbb{D} = \{x \in \mathbb{R}^2 : \|x\| < 1\}$ equipped with the Riemannian metric

$$ds^2 = \frac{4\|dx\|^2}{(1 - \|x\|^2)^2}, \quad (1)$$

which induces the following distance function between two points $x, y \in \mathbb{D}$:

$$d_{\mathbb{D}}(x, y) = \operatorname{arcosh}\left(1 + \frac{2\|x - y\|^2}{(1 - \|x\|^2)(1 - \|y\|^2)}\right). \quad (2)$$

Our methodology extends the coalescent embedding framework, initially developed for mapping single-layer networks onto the hyperbolic disk [25]. In this study, we introduce two principal advancements to the original approach. First, we generalize the method to support multilayer network structures, enabling more comprehensive modeling of complex systems. Second, we adapt the embedding scheme to project networks into the Poincaré disk $\mathbb{D}$, thereby exploiting its geometric characteristics for enhanced representational accuracy.

A. Prewighting

The first step of our multilayer embedding is edge preweighting. If edge weights already reflect importance, this step can be skipped. Prewighting assigns new weights based on local topological importance using a “repulsion-attraction rule”. For an edge $e_{i,j}$ between nodes i and j , the new weight is defined as follows:

$$W_{\text{RA}}(e_{i,j}) = \frac{d_i + d_j + d_i d_j}{1 + CN_{ij}}, \quad (3)$$

where d_i is the degree of the node i and CN_{ij} the number of common neighbors between nodes i and j .

As in the original methodology [25], edge betweenness centrality can also be used as an alternative weighting strategy during the preweighting step. Nevertheless, nodes that were not previously connected remain unconnected. The framework is applicable to all edges in a multilayer network, including interlayers edges when present.

B. Global connectivity matrix construction – Case of multilayer network with a constant node set across layers

Once the preweighted step has been completed or skipped, a new global connectivity matrix is constructed, incorporating the data from all layers. For clarity, we will denote n_l the number of nodes present in the layer l and let L represents the total




FIG. 1. Global pipeline for multiple network embedding with identical nodes set across layers. The process begins with the extraction of the connectivity matrices for each layer. These matrices are then combined to form a global connectivity matrix G . The dimension reduction algorithm is applied to G to obtain a two-dimensional representation of the dataset. From this embedding, the angular coordinates of the nodes in each layer are extracted, and their radii are initially normalized to one. In a final step, the radial coordinates are reassigned based on the centrality of each node.

number of layers in the network. We begin by considering the case where all layers share the same set of nodes, i.e., $n_l = n$ for all l . We then discuss the adaptations required when layers contain different node sets.

As illustrated in Fig. 1, the matrix G is constructed as follows. Let us denote the connectivity matrix of layer i as $M_{i,j}$ with dimensions of $n \times n$. The global connectivity matrix of the whole network will be designated as G , with dimensions of $n \times L$ for rows and columns [2]. The matrix is divided into $L \times L$ blocks of size $n \times n$ each, noted $g_{i,j}$. The blocks on the diagonal will be occupied with the connectivity matrices of the layers, that is to say,

$$g_{i,i} = M_{i,i} \quad \text{for } i \in \{1, \dots, L\}. \quad (4)$$

The off-diagonal matrices represent the coupling terms between the layers. In the case of disconnected layers, each off-diagonal matrix is composed of an identity matrix of size $n \times n$, scaled by a coupling coefficient μ , which represents the

coupling interaction between layers i and j :

$$g_{i,j} = I_n \cdot \mu, \quad \text{for } i, j \in \{1, \dots, L\}, \quad i \neq j, \quad (5)$$

where $g_{i,j}$ represents the block matrix of size $n \times n$ in G located in row i and column j . Here, I_n denotes the identity matrix of size $n \times n$.

In the case where edges exist between layers (e.g., $M_{i,j}[l, m] = 1$ if an edge between node l of layer i is connected to node m of layer j , preweighting step can also be applied to this connections according to Sec. II A), the off-diagonal block matrices representing the coupling between the layers will take the following form:

$$g_{i,j} = I_n \cdot \mu + M_{i,j}, \quad \text{for } i, j \in \{1, \dots, L\}, \quad i \neq j, \quad (6)$$

where $M_{i,j}$ is the interaction matrix between nodes of layers i and j .

C. Extension to multilayer graphs with heterogeneous node sets

The method introduced in the previous subsection is formulated for multilayer graphs in which all layers share an identical set of nodes. Nevertheless, in practical applications, it is common to encounter layers comprising distinct node sets. The proposed approach remains applicable in such scenarios, with only minor modifications, as described below.

Let n_i denote the number of nodes in layer i . To construct the global connectivity matrix G , the off-diagonal blocks (originally identity matrices scaled by coupling parameters μ) are replaced with rectangular matrices when node counts differ across layers. Specifically, the off-diagonal block $g_{i,j}$ (for $i \neq j$) is of size $n_i \times n_j$ and contains entries equal to μ at positions corresponding to interlayer node associations. That is, if node a in layer i corresponds to node b in layer j and there is no direct link between them, then the entry (a, b) in $g_{i,j}$ is simply assigned the value μ , otherwise it is set to zero. In the absence of correspondence between layers nodes, the corresponding values are assigned a value of zero. When interlayer edges (regardless of whether they are weighted) are explicitly included in the multilayer graph, they are likewise represented in the off-diagonal blocks of G , as specified in Eq. (6), irrespective of whether the node sets across layers are identical. The construction of the diagonal matrices follows the same methodology employed in the standard multiplex framework.

D. Dimension reduction

Following the original approach proposed in Ref. [25], once the global connectivity matrix G has been constructed, a nonlinear dimensionality reduction algorithm, such as Isomap or Laplacian eigenmaps [30], is applied to embed the network into a two-dimensional space. The matrix G serves as input to the algorithm, which produces 2D coordinates for all $\sum_{l=1}^L n_l$ nodes, where, for instance, the positions of the first n_1 nodes correspond to those in Layer 1. For illustration, see Fig. 1 in the case where $n_l = n$ for all l . In the present study, we adopt the Isomap algorithm, as it performs a global optimization that preserves the overall network structure by incorporating geodesic (i.e., shortest path) distances between nodes [31].

E. Layers extraction and radii attribution

Once the nonlinear dimensionality reduction has been performed, nodes retain their angular coordinates in the two-dimensional embedding, while their radial positions are reassigned based on their centrality within the network. As illustrated in Fig. 1, the initial step involves extracting the positions of the nodes for each layer and setting all radial coordinates to 1. Radial positions are then reassigned according to each node's centrality, measured by its degree. To obtain a more informative estimate of centrality, we use the weighted degree w_i , which accounts for the strength of connections rather than merely their number. The weighted degree of a node is defined as the sum of the weights of its incident edges, providing a measure of the node's overall connectivity strength within the weighted network. Notably, this centrality measure reflects any preweighting applied to the edges; in the absence of such preprocessing, the original edge weights are used. Although weighted degree is employed here, alternative centrality metrics may also be used, depending on the specific application or network characteristics [25,32]. Given a node i with (weighted) degree w_i , its radial coordinate is updated as follows:

$$r_i = 1 - \tanh\left(\frac{w_i}{\beta}\right), \quad (7)$$

where r_i denotes the radial coordinate of node i , $\tanh$ is the hyperbolic tangent function, and β is a scaling parameter. This formulation ensures that all nodes are embedded within the Poincaré disk, with nodes of degree zero mapped at the boundary ($r = 1$) and the most central nodes positioned closer to the origin. The parameter β controls the rate at which the radius decreases with increasing centrality and plays a role analogous to the popularity fading parameter in the original coalescent embedding method. Unlike the original coalescent embedding, which assigns radii based on the degree ranking within a layer, here we directly use the (weighted) degree itself. This choice ensures that nodes with the same degree across different layers are mapped at consistent radial distances, while also naturally placing disconnected nodes ($w_i = 0$) at the boundary ($r = 1$). This adaptation makes the method more suitable for multilayer networks, where degree distributions can differ significantly between layers. A practical strategy for setting β consists in prescribing the radius r_i of a node i with maximal (weighted) degree w_i , such that it lies within a desired distance from the origin (for instance, $r_i = 0.1/0.2$) and solving $\beta = w_i/\text{arctanh}(r_i)$. This ensures that node radii span a broad portion of the disk while maintaining an interpretable mapping between centrality and geometry.

F. Gaussian distributions in the Poincaré disk

The proposed embedding method exhibits the ability to consistently cluster critical nodes across multiple layers. To quantitatively assess this property, we introduce a Gaussian model in hyperbolic space to characterize the distribution of a given node's embedded positions across layers. The underlying rationale is that the covariance among the projections of the same node from different layers decreases as these projections converge toward similar coordinates. Consequently,

a Gaussian model offers an appropriate statistical framework for assessing the localization a node's position across layers.

To define the parameters of the hyperbolic Gaussian distribution (specifically the barycenter and covariance) we employ the Klein model of hyperbolic space. Like the Poincaré model, the Klein model is defined over the open unit disk $\mathbb{D}$, that is, $\mathbb{K} = \{x \in \mathbb{R}^2 : \|x\| < 1\}$, but it is endowed with a different metric ($ds^2 = \frac{\|dx\|^2}{1-\|x\|^2} + \frac{(x \cdot dx)^2}{(1-\|x\|^2)^2}$). Consequently, a given point will generally have different coordinate representations in the Poincaré and Klein models, due to the distinct metrics underlying each.

Let $X_{\mathbb{D}}$ and $X_{\mathbb{K}}$ denote the coordinates of the same point in the Poincaré and Klein disk models, respectively. The bijective mapping between these representations is defined by

$$X_{\mathbb{D}} = \frac{X_{\mathbb{K}}}{1 + \sqrt{1 - \|X_{\mathbb{K}}\|^2}}, \quad X_{\mathbb{K}} = \frac{2X_{\mathbb{D}}}{1 + \|X_{\mathbb{D}}\|^2}. \quad (8)$$

Unlike other models of hyperbolic geometry, the Klein model admits a particularly simple expression for the hyperbolic barycenter (a generalization of the Euclidean mean) in Klein coordinates [33]:

$$\text{HypBary}(X_1, \dots, X_n) = \frac{\sum_{i=1}^n \gamma_i X_i}{\sum_{i=1}^n \gamma_i}, \quad (9)$$

where $\gamma_i = \frac{1}{\sqrt{1 - \|X_i\|^2}}$ is the Lorentz factor associated with X_i .

Once the barycenter is obtained in the Klein disk, it is mapped back to the Poincaré disk using Eq. (8).

To estimate the hyperbolic Gaussian distributions in the Poincaré disk representation, we project the node positions onto the tangent space at the barycenter—an Euclidean space where Gaussian statistics are well defined. Specifically, the logarithmic map $\text{Log}_z(X)$ projects a point $X \in \mathbb{D}$ onto the tangent space $T_z\mathbb{D}$ centered at $z \in \mathbb{D}$, while the exponential map $\text{Exp}_z(V)$ performs the inverse operation, mapping a tangent vector $V \in T_z\mathbb{D}$ back to the hyperbolic disk. These maps are defined as follows:

$$\text{Log}_z(X) = \frac{2}{\lambda_z} \text{arctanh}(\| -z \oplus X \|) \frac{-z \oplus X}{\| -z \oplus X \|}, \quad (10)$$

$$\text{Exp}_z(V) = z \oplus \left(\tanh\left(\frac{\lambda_z \|V\|}{2}\right) \frac{V}{\|V\|} \right). \quad (11)$$

Here, $\oplus$ denotes the Möbius addition, defined by [34]

$$X \oplus Y = \frac{(1 + 2\langle X, Y \rangle + \|Y\|^2)X + (1 - \|X\|^2)Y}{1 + 2\langle X, Y \rangle + \|X\|^2\|Y\|^2}. \quad (12)$$

To estimate the covariance matrix of a set of points X_i in the hyperbolic disk, we first compute their barycenter z in the Klein disk and subsequently map it back to the Poincaré disk. The set of points is then projected onto the tangent space $T_z\mathbb{D}$, centered on the estimated barycenter z using the logarithmic map $\text{Log}_z(\cdot)$. Within this Euclidean subspace, the barycenter is the point with coordinates $(0, 0)$ and the covariance matrix can be calculated as $V_z = \text{Covariance}([\text{Log}_z(X_i)]^T)$.

For a given position $\tilde{X}$ in the Poincaré disk, the hyperbolic Gaussian pdf of the nodes X_i estimation is

$$\text{Pdf}_z(\tilde{X}) = \frac{1}{\sqrt{(2\pi)^2 \|V_z\|}} \text{Exp}\left[-\frac{1}{2} \tilde{W}^T V_z^{-1} \tilde{W}\right], \quad (13)$$

where

$$\tilde{W} = \text{Log}_z(\tilde{X}) \quad (14)$$

represents the projection of the position $\tilde{X}$ from the Poincaré disk into the tangent space $T_z\mathbb{D}$ centered on z , the hyperbolic barycenter of all occurrences for the set of points X_i . $\|V_z\|$ denotes the determinant of covariance matrix V_z computed from the tangent space projection, and j denotes the imaginary unit $j = \sqrt{-1}$.

III. RESULTS

This section illustrates the effectiveness of the proposed hyperbolic embedding method for multilayer graphs through a series of case studies. We first apply the method to a multilayer stochastic block model (SBM), assessing both its capacity to recover underlying community structures and its robustness in scenarios where layers contain different sets of nodes. We then evaluate the approach on brain network data from patients with temporal lobe epilepsy, demonstrating its ability to yield consistent node localization across layers. Finally, we investigate the influence of the coupling parameter μ on the quality of the resulting embeddings. Together, these examples showcase the versatility and strengths of our method across diverse network analysis settings.

A. Multilayer community structure – Multilayer SBM example

To evaluate the capacity of our method to recover and emphasize community structures through hyperbolic embedding, we constructed a multilayer Stochastic Block Model (SBM). The model consists of n_l layers, each containing n nodes partitioned into n_c communities. Community assignments are kept consistent across layers. As in the classical single-layer SBM, nodes within the same community are more likely to be connected than nodes belonging to different communities. This intralayer community structure is extended across layers to model interlayer organization.

Let us consider two nodes i and j , assigned to communities c_i and c_j , and belonging to layers l_i and l_j , respectively. The probability $p_{i,j}$ of an edge between these two nodes is defined as

$$p_{i,j} = \begin{cases} p_{\neq} & \text{if } c_i \neq c_j \text{ and } l_i = l_j, \\ p_{=} & \text{if } c_i = c_j \text{ and } l_i = l_j, \\ \frac{p_{\neq}}{\alpha} & \text{if } c_i \neq c_j \text{ and } l_i \neq l_j, \\ \frac{p_{=}}{\alpha} & \text{if } c_i = c_j \text{ and } l_i \neq l_j. \end{cases} \quad (15)$$

Here, $p_{=}$ denotes the probability of a connection between two nodes belonging to the same community within the same layer, while $p_{\neq}$ corresponds to the probability of a connection for nodes in different communities within the same layer. To ensure a detectable community structure, we assume $p_{=} \geq p_{\neq}$. The parameter $\alpha \geq 1$ controls the strength of interlayer coupling: $\alpha = 1$ corresponds to maximal coupling (no distinction between intra- and interlayer edges), while larger values of α reduce interlayer connectivity.

To demonstrate the effectiveness of our approach in revealing community structure, we compared the results of the hyperbolic embedding in the Poincaré disk under two




FIG. 2. Hyperbolic embeddings of multilayer SBM models. Parameters: three layers, three communities; $p_{=} = 0.16$, $p_{\neq} = 0.06$, $\mu = 20$, $\alpha = 10$. (a) Comparison between multilayer and independent embeddings for identical node sets across layers. Each layer contains 100 nodes. Top row: embeddings computed using the proposed multilayer approach. Bottom row: embeddings computed independently for each layer. (b) Multilayer embeddings with varying node sets across layers. The multilayer model comprises three layers with different numbers of nodes (Layer 0: 100 nodes; Layer 1: 90 nodes; Layer 2: 85 nodes).

settings: (i) independent embedding of each layer without cross-layer information, and (ii) joint embedding of all layers using our multilayer method. Remarkably, even under weak interlayer coupling ($\alpha = 10$), the multilayer approach reveals well-separated community clusters, unlike the independent embeddings where the structure appears more dispersed. These results, shown in Fig. 2, highlight the advantage of the multilayer embedding in networks with latent community structure. When using our multilayer framework, the communities are clearly delineated in the hyperbolic space, whereas the independent approach requires stronger intra-community and lower intercommunity connectivity to achieve comparable separation. This is due to the fact that a lower intercommunity connectivity enhances the contrast between intra- and intercommunity edge probabilities, which is essential for the independent embedding to produce a clear separation of communities. When this contrast diminishes and the connection probabilities become similar, the network approaches a regime akin to an Erdős–Rényi graph, where communities are less distinguishable without supplementary interlayer information. When replacing Isomap with Laplacian eigenmaps on the same dataset, the multilayer embedding still preserves community separation, while the performance of the independent embedding further deteriorates with the spectral method.

Another key advantage of our method lies in its capacity to accommodate multilayer networks comprising different sets of nodes across layers, provided that a correspondence matrix between node-layer pairs is available. To assess this capability, we generated a three-layer stochastic block model (SBM) using our framework and randomly removed nodes from certain layers prior to embedding the network into the Poincaré disk. The results, depicted in Fig. 2, illustrate that Layer 0 retains all 100 nodes, whereas Layers 1 and 2 contain only 90 and 85 nodes, respectively, due to the random removal of 10 and 15 nodes. Despite these discrepancies, our method effectively embeds all layers while preserving the latent community structure. These findings indicate that our approach robustly captures the multilayer community organization even when the individual layers consist of nonidentical node sets.

Other multilayer/multiplex network models exist, including the Geometric Multiplex Model (GMM) [29,35]. We applied our embedding method to networks generated using the GMM model and compared the pairwise hyperbolic distances between nodes in the original configuration and the embedded version. Using a four-layer network with 100 nodes, we observed a mean Spearman correlation of 0.79 between the original and embedded hyperbolic distances, indicating that the embedding effectively captures the network geometry of the network.”

B. Nodes localization – An example on brain epilepsy networks

A potential application of our method lies in the comparison of network populations, such as comparing the brain networks of patients with those of healthy controls. The spatial localization of nodes in the hyperbolic disk could serve as a valuable biomarker. To illustrate this use case, we analyzed diffusion MRI data from patients with a left temporal lobe epilepsy ($n = 19$) and healthy controls ($n = 28$). Each brain network consists of 164 nodes, corresponding to regions defined by a standard anatomical atlas (Destrieux atlas). The dataset is described in detail in [36].

To illustrate the practical utility of our approach, we computed the hyperbolic Gaussian distribution of node positions (specifically from the left temporal lobes) within the Poincaré disk (see Sec. II F). We evaluated two embedding strategies: our proposed multilayer method, in which each layer represents a different patient, and an independent embedding approach, wherein each patient is embedded separately and the resulting configurations are subsequently aligned via rotational transformation. For the rotational transformation, reflection symmetry (axis inversion) is also tested. The alignment error is computed both with and without reflection, and the configuration with the lowest error is retained. The alignment is performed by rotating each embedding (beyond the reference layer) to minimize the sum of hyperbolic distances between corresponding nodes, effectively minimizing the G_{score} defined Eq. (16). This procedure, similar to the one described in [37], ensures comparability across layers while preserving the geometric invariance of the embedding space.

As shown in Fig. 3, the multilayer embedding produces a more consistent and spatially coherent localization of the temporal lobe regions compared to the independent embedding strategy. From this result we observe that, compared with




FIG. 3. Hyperbolic embedding of patients with left temporal lobe epilepsy (red points) and controls (blue points). Visualization of the Gaussian distributions of node positions corresponding to the left temporal lobes (two nodes per patients) within the Poincaré disk, using two embedding strategies: (a) the proposed multilayer embedding approach, where each layer represents an individual patient $\beta = 50$, $\mu = 20$; and (b) an independent embedding of each patient followed by *post hoc* rotational alignment.

independently embedding each layer, the nodes within the left temporal lobe are more tightly clustered under our method, yielding a sharper Gaussian distribution. This enhanced localization highlights the capacity of the multilayer embedding to integrate information across patients, thereby generating more stable and reliable individual embeddings. Moreover, the method clearly uncovers distributional differences between patients with left temporal epilepsy and healthy controls, underscoring its potential utility in clinical settings. Permutation tests, where group labels were randomly reassigned, resulted in entropy distributions for both groups converging towards a common intermediate value, distinct from the lower entropy observed in patients (a more concentrated node positions) and higher entropy in controls (more dispersed configurations). This indicates that the tighter clustering in patients is not an artifact of the embedding method but reflects genuine group-specific network organization.

C. Impact of the μ coupling parameter

An important parameter in the method is the coupling strength μ , which governs the degree of alignment between embeddings across layers. The effectiveness of this alignment can be quantitatively assessed using the global error score, defined as

$$G_{\text{score}}(E_1^h, E_2^h) = \frac{1}{N} \sum_{i=1}^N d_{\mathbb{D}}(X_i^1, X_i^2), \quad (16)$$

where X_i^1 and X_i^2 denote the positions of node i in the first and second layer embeddings, respectively. $d_{\mathbb{D}}(\cdot, \cdot)$ is the hyperbolic distance function in the Poincaré’s disk model [see Eq. (2)]. E_1^h and E_2^h denote the embeddings of the first and second layer, respectively, in the Poincaré disk. In other words, the score corresponds to the mean hyperbolic displacement of nodes between the two layers. A lower G_{score} indicates a better alignment between the embeddings, and thus a more coherent multilayer representation.

To assess the impact of the coupling parameter μ , we performed the embedding over a range of μ values for a two-layer graph (a multilayer SBM model), computing the




FIG. 4. G_{score} [see Eq. (16)] as a function of the coupling parameter μ and the mean edge weight of the graph. Calculations were performed on a two-layer SBM model with $n = 100$ nodes per layer.

corresponding G_{score} in each case. Interestingly, we observed a sharp transition in the behavior of G_{score} as μ increases. Beyond a critical value, denoted μ^* , the score drops significantly and stabilizes at a plateau value, referred to as G_{score}^* (see Fig. 4). Plateau is stable even when μ takes large values such as $\mu \rightarrow 10^4 \mu^*$ (result not shown). This indicates that, to reach a stable regime on the plateau, the parameter μ must satisfy $\mu \gg \mu^* \simeq$ the mean edge weight on the graph. Since the alignment remains stable once the plateau is reached, choosing a sufficiently high value of μ ensures convergence without adverse effects.

To further interpret the physical meaning of μ^* , we replicated this analysis using a multilayer stochastic block model (SBM) with varying the mean edge weights of the graphs. This was implemented by scaling the edge weights of the original graph to achieve the desired mean. Remarkably, we found that μ^* scales with the mean edge weight of the network, and the plateau value G_{score}^* is consistently reached when μ matches this mean value. This suggests that optimal embedding performance occurs when the coupling strength μ reflects the average internal connectivity strength of the layers.

The global plateau value G_{score}^* also appears to carry physical significance. While the critical coupling μ^* is governed by the average edge weight, the value of G_{score}^* seems to reflect the similarity between the connectivity distributions of the layers. To assess this hypothesis, we perturbed the connection probabilities in the second layer (i.e., modifying $p_{=}$ and $p_{\neq}$) while keeping the first layer unchanged. As

shown in Fig. 4(b), the relationship between μ^* and the mean edge weight remains robust to these perturbations, indicating that the phase transition is not disrupted. However, the plateau value G_{score}^* increases when the two layers become less similar, compared to the case with identical probability distributions [Fig. 4(a)]. These results suggest that G_{score}^* can serve as a quantitative indicator of interlayer similarity: the more similar the layers, the lower the score. To further validate this interpretation, we compared G_{score}^* with the graph edit distance [38] in a controlled perturbation experiment, where edges were randomly added and removed from an initial network. The two metrics showed a strong linear relationship, with a Pearson correlation coefficient of 0.969, confirming that G_{score}^* reliably captures structural dissimilarities between layers.

IV. DISCUSSION

In this work, we have presented a novel framework for embedding multilayer networks into hyperbolic space by extending the coalescent embedding approach to incorporate interlayer interactions via a global connectivity matrix, with the Poincaré disk serving as the embedding space. In contrast to prior methods such as [29], which embed each layer independently and perform comparative analyses *post hoc*, our approach integrates information from all layers simultaneously. This joint embedding strategy enhances the coherence and interpretability of the resulting representations, offering a significant advantage in contexts where structural patterns are shared across layers.

Our method is entirely data driven and does not rely on explicit generative assumptions, in contrast to model-based approaches such as Mercator [23], which infer latent geometric parameters through a probabilistic framework grounded in network generative models. While Mercator provides robust parameter estimation and interpretable embeddings even in weakly geometric regimes, our approach offers complementary advantages by allowing flexible and direct embedding without requiring prior model parameter estimation. This makes our framework particularly suitable for empirical settings where model assumptions may not hold or where parameters are challenging to infer, such as in heterogeneous or partially observed multilayer networks.

The embedding is performed in the Poincaré disk, a bounded model of hyperbolic space. This choice brings multiple benefits. First, it offers a well-defined geometric boundary at radius $r = 1$, which naturally encodes disconnected or peripheral nodes addressing a common issue in the native hyperbolic disk where disconnected nodes tend toward infinity. Second, it allows for direct and intuitive comparison of node positions across layers, as all layers are mapped into a common finite geometric domain. Finally, it provides a visual interpretability, which is critical for downstream analyses such as classification or clustering.

A key strength of our method lies in its ability to embed layers with distinct node sets. Many real-world systems, particularly in biological and social contexts, exhibit partial or evolving node participation across layers. Our approach accommodates such heterogeneity by adapting the global connectivity matrix with rectangular coupling blocks that reflect

known node correspondences. As demonstrated in the multilayer SBM experiments, our method successfully captures multilayer community structure even when layers do not share the same set of nodes, highlighting its robustness and general applicability.

Compared to conventional approaches, where each layer is embedded independently and subsequently aligned, the proposed multilayer framework achieves more accurate spatial localization of nodes with similar connectivity across layers, as quantified through hyperbolic Gaussian distributions. When applied to real-world brain networks from patients with temporal lobe epilepsy, the method consistently identifies disease-related brain regions in the Poincaré disk across individuals, outperforming the independent layer embedding approach and demonstrating its potential for comparative analysis in clinical neuroscience. These results indicate that our approach provides a principled basis for constructing group-representative networks by associating each node with a Gaussian probability distribution in the hyperbolic disk. The improved consistency of the embeddings suggests broader applicability to tasks such as anomaly detection, network-based phenotyping, and cross-subject analyses.

Another central contribution of this work is the analysis of the coupling parameter μ , which modulates the influence of interlayer interactions in the global embedding. This parameter governs the influence of interlayer interactions in the global embedding. We show that the alignment quality of the embeddings undergoes a clear transition as μ increases: below a critical threshold μ^* , embeddings remain poorly aligned, whereas above μ^* , a plateau is reached, indicating successful integration of layers. Notably, we observe that μ^* scales with the average edge weight of the network, suggesting that optimal coupling reflects the intrinsic connectivity strength within the layers. Furthermore, the plateau value of the alignment score G_{score} serves as a proxy for interlayer similarity, with lower scores corresponding to more structurally homogeneous layers. These findings suggest that the coupling parameter μ acts not only as a tunable parameter for controlling multilayer integration, but also as a diagnostic indicator of the underlying network structure. Specifically, the optimal value μ must be sufficiently large relative to the mean edge weight of the network to ensure stable alignment across layers. Furthermore, the plateau value of the alignment score, G_{score}^* , quantitatively reflects the degree of similarity or dissimilarity between layers, thus providing insight into the latent geometric and structural organization of the multilayer system.

In addition to structural analysis, the proposed framework offers new perspectives for perturbation detection. For instance, one can treat the original network as one layer and a perturbed version as another. The joint embedding then reveals how perturbations affect node positions in hyperbolic

space. This approach provides a principled alternative to methods like [37], where networks are embedded separately and compared postalignment. The inclusion of interlayer coupling allows our method to capture both global and local deviations more effectively, offering a potential refinement to existing perturbation scoring techniques.

While this study focuses on two-dimensional embeddings for clarity and visualization, the framework is readily extendable to higher dimensions. Embedding into the n -dimensional Poincaré ball $\mathbb{D}^n$ would allow for the representation of more complex topological features and larger networks, in a similar spirit to the D-Mercator method [24]. Furthermore, our embedding pipeline, centered on the global connectivity matrix and the coupling parameter, can be generalized to other types of geometries or embedding techniques.

An important limitation of the proposed method concerns its computational scalability. The approach requires constructing and processing a global connectivity matrix of size $M \times M$, where $M = N \times L$, with N the number of nodes per layer and L the number of layers. Consequently, the computational complexity scales approximately as $\mathcal{O}(M^2)$, which can quickly become prohibitive for multilayer networks with a large number of layers or nodes. This limitation restricts the applicability of our method to moderately sized multilayer systems, particularly in the context of temporal networks with high temporal resolution, where the number of layers (snapshots) is very large. Addressing this challenge will require algorithmic optimizations or approximation strategies, representing a promising direction for future work.

In summary, this work provides a unified and flexible approach for embedding multilayer networks into hyperbolic space. It enhances interpretability, accommodates heterogeneity across layers, captures perturbations, and reveals meaningful geometric transitions through the coupling parameter. These contributions open new directions for the study of complex systems in a geometric framework.

ACKNOWLEDGMENTS

M.G. acknowledges doctoral support from the Ecole Normale Supérieure de Lyon, France. We thank Marián Ángeles Serrano for kindly sharing the implementation of the Geometric Multiplex Model (GMM), which was used in our simulations.

DATA AVAILABILITY

Some of the data that support the findings of this article are openly available [39]. Patients Diffusion MRI Data are not publicly available because of legal restrictions preventing unrestricted public distribution. The data are available from the authors upon reasonable request.

- [1] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, and D.-U. Hwang, *Phys. Rep.* **424**, 175 (2006).
- [2] S. Boccaletti, G. Bianconi, R. Criado, C. I. del Genio, J. G.-G. nes, M. Romance, I. S. na Nadal, Z. Wang, and M. Zanin, *Phys. Rep.* **544**, 1 (2014).

- [3] M. Kivelä, A. Arenas, M. Barthélemy, J. P. Gleeson, Y. Moreno, and M. A. Porter, *J. Complex Networks* **2**, 203 (2014).
- [4] T. Xu, K. H. Nanning, E. Schwartz, S. J. Hong, J. T. Vogelstein, A. Goulas, D. A. Fair, C. E. Schroeder, D. S. Margulies, J.

- Smallwood, M. P. Milham, and G. Langs, *NeuroImage* **223**, 117346 (2020).
- [5] M. G. Puxeddu, M. Petti, and L. Astolfi, *Front. Syst. Neurosci.* **15**, 624183 (2021).
- [6] P. Goyal and E. Ferrara, *Knowledge-Based Syst.* **151**, 78 (2018).
- [7] P. Cui, X. Wang, J. Pei, and W. Zhu, *IEEE Trans. Knowl. Data Eng.* **31**, 833 (2019).
- [8] M. Xu, *SIAM Rev.* **63**, 825 (2021).
- [9] Z. Liu, C. Huang, Y. Yu, B. Fan, and J. Dong, in *International Conference on Information and Knowledge Management, Proceedings* (Association for Computing Machinery, 2020), pp. 995–1004.
- [10] L. Pio-Lopez, A. Valdeolivas, L. Tichit, É. Remy, and A. Baudot, *Sci. Rep.* **11**, 8794 (2021).
- [11] W. Liu, P. Y. Chen, S. Yeung, T. Suzumura, and L. Chen, in *IEEE International Conference on Data Mining Workshops, ICDMW*, Vol. 2017–November (IEEE Computer Society, 2017), pp. 134–141.
- [12] Q. Wang, H. Jiang, Y. Jiang, S. Yi, Q. Nie, and G. Zhang, *Digital Commun. Networks* **9**, 1157 (2023).
- [13] C. Park, J. Han, and H. Yu, *Knowledge-Based Syst.* **197**, 105861 (2020).
- [14] Y. Wang, D. Chang, Z. Fu, and Y. Zhao, *IEEE Trans. Multimedia* **25**, 1008 (2023).
- [15] Y. Liu, L. He, B. Cao, P. S. Yu, A. B. Ragin, and A. D. Leow, in *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32 (AAAI Press, 2018), pp. 117–124.
- [16] J. Li, C. Chen, H. Tong, and H. Liu, in *Proceedings of the 2018 SIAM International Conference on Data Mining* (Society for Industrial and Applied Mathematics Publications, 2018), pp. 684–692.
- [17] M. Boguñá, D. Krioukov, and K. C. Claffy, *Nat. Phys.* **5**, 74 (2009).
- [18] D. Krioukov, F. Papadopoulos, M. Kitsak, A. Vahdat, and M. Boguñá, *Phys. Rev. E: Stat. Nonlinear Soft Matter Phys.* **82**, 036106 (2010).
- [19] M. Boguñá, I. Bonamassa, M. D. Domenico, S. Havlin, D. Krioukov, and M. Á. Serrano, *Nat. Rev. Phys.* **3**, 114 (2021).
- [20] M. Nickel and D. Kiela, in *Advances in Neural Information Processing Systems (NIPS 2017)*, Vol. 30, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran Associates, Inc., Long Beach, CA, USA), pp. 6341–6350.
- [21] F. Papadopoulos, M. Kitsak, M. Á. Serrano, M. Boguñá, and D. Krioukov, *Nature (London)* **489**, 537 (2012).
- [22] F. Papadopoulos, C. Psomas, and D. Krioukov, *IEEE/ACM Trans. Networking* **23**, 198 (2015).
- [23] G. García-Pérez, A. Allard, M. Á. Serrano, and M. Boguñá, *New J. Phys.* **21**, 123033 (2019).
- [24] R. Jankowski, A. Allard, M. Boguñá, and M. Á. Serrano, *Nat. Commun.* **14**, 7585 (2023).
- [25] A. Muscoloni, J. M. Thomas, S. Ciucci, G. Bianconi, and C. V. Cannistraci, *Nat. Commun.* **8**, 1615 (2017).
- [26] B. Kovács and G. Palla, *Sci. Rep.* **11**, 8350 (2021).
- [27] L. Wang, C. Huang, W. Ma, R. Liu, and S. Vosoughi, *Data Min. Knowl. Discovery* **35**, 1906 (2021).
- [28] M. Yang, M. Zhou, H. Xiong, and I. King, *IEEE Trans. Knowl. Data Eng.* **35**, 11489 (2023).
- [29] K. K. Kleineberg, M. Boguñá, M. Á. Serrano, and F. Papadopoulos, *Nat. Phys.* **12**, 1076 (2016).
- [30] U. von Luxburg, *Stat. Comput.* **17**, 395 (2007).
- [31] J. B. Tenenbaum, V. de Silva, and J. C. Langford, *Science* **290**, 2319 (2000).
- [32] F. Battiston, V. Nicosia, and V. Latora, *Phys. Rev. E* **89**, 032804 (2014).
- [33] V. Khrukov, L. Mirvakhabova, E. Ustinova, I. Oseledets, and V. Lempitsky, in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR, 2020)*, pp. 6418–6428.
- [34] A. A. Ungar, *Comput. Math. Appl.* **41**, 135 (2001).
- [35] J. van der Kolk, D. Krioukov, M. Boguñá, and M. Ángeles Serrano, Multiplexity amplifies geometry in networks, [arXiv:2505.17688](https://arxiv.org/abs/2505.17688).
- [36] P. Besson, V. Dinkelacker, R. Valabregue, L. Thivard, X. Leclerc, M. Baulac, D. Sammler, O. Colliot, S. Lehericy, S. Samson, and S. Dupont, *NeuroImage* **100**, 135 (2014).
- [37] A. Longhena, M. Guillemaud, and M. Chavez, *Chaos: An Int. J. Nonlinear Sci.* **34**, 063117 (2024).
- [38] Z. Abu-Aisheh, R. Raveaux, J. Ramel, and P. Martineau, in *Proceedings of the International Conference on Pattern Recognition Applications and Methods - Volume 1: ICPRAM, INSTICC* (SciTePress, 2015), pp. 271–278.
- [39] <https://github.com/MartinG-38/MLNHypEmb>.