
Directed Graph Embeddings in Pseudo-Riemannian Manifolds

Aaron Sim¹ Maciej Wiatrak¹ Angus Brayne¹ Páidí Creed¹ Saeed Paliwal¹

Abstract

The inductive biases of graph representation learning algorithms are often encoded in the background geometry of their embedding space. In this paper, we show that general directed graphs can be effectively represented by an embedding model that combines three components: a pseudo-Riemannian metric structure, a non-trivial global topology, and a unique likelihood function that explicitly incorporates a preferred direction in embedding space. We demonstrate the representational capabilities of this method by applying it to the task of link prediction on a series of synthetic and real directed graphs from natural language applications and biology. In particular, we show that low-dimensional cylindrical Minkowski and anti-de Sitter spacetimes can produce equal or better graph representations than curved Riemannian manifolds of higher dimensions.

1. Introduction

Representation learning of symbolic objects is a central area of focus in machine learning. Alongside the design of deep learning architectures and general learning algorithms, incorporating the right level of inductive biases is key to efficiently building faithful and generalisable entity and relational embeddings (Battaglia et al., 2018).

In graph representation learning, the embedding space geometry itself encodes many such inductive biases, even in the simplest of spaces. For instance, vertices embedded as points in Euclidean manifolds, with inter-node distance guiding graph traversal and link probabilities (Grover & Leskovec, 2016; Perozzi et al., 2014), carry the underlying assumption of homophily with node similarity as a metric function.

The growing recognition that Euclidean geometry lacks the flexibility to encode complex relationships on large graphs

¹BenevolentAI, London, United Kingdom. Correspondence to: Aaron Sim <aaron.sim@benevolent.ai>.

at tractable ranks, without loss of information (Nickel et al., 2011; Bouchard et al., 2015), has spawned numerous embedding models with non-Euclidean geometries. Examples range from complex manifolds for simultaneously encoding symmetric and anti-symmetric relations (Trouillon et al., 2016), to statistical manifolds for representing uncertainty (Vilnis & McCallum, 2015).

One key development was the introduction of hyperbolic embeddings for representation learning (Nickel & Kiela, 2017; 2018). The ability to uncover latent graph hierarchies was applied to directed acyclic graph (DAG) structures in methods like Hyperbolic Entailment Cones (Ganea et al., 2018), building upon Order-Embeddings (Vendrov et al., 2016), and Hyperbolic Disk embeddings (Suzuki et al., 2019), with the latter achieving good performance on complex DAGs with exponentially growing numbers of ancestors and descendants. Further extensions include hyperbolic generalisations of manifold learning algorithms (Sala et al., 2018), product manifolds (Gu et al., 2018), and the inclusion of hyperbolic isometries (Chami et al., 2020).

While these methods continue to capture more complex graph topologies they are largely limited to DAGs with transitive relations, thus failing to represent many naturally occurring graphs, where cycles and non-transitive relations are common features.

In this paper we introduce pseudo-Riemannian embeddings of both DAGs and graphs with cycles. Together with a novel likelihood function with explicitly broken isometries, we are able to represent a wider suite of graph structures. In summary, the model makes the following contributions to graph representation learning, which we will expand in more detail below:

- The ability to disentangle semantic and edge-based similarities using the distinction of space- and timelike separation of nodes in pseudo-Riemannian manifolds.
- The ability to capture directed cycles by introducing a compact timelike embedding dimension. Here we consider Minkowski spacetime with a circle time dimension and anti-de Sitter spacetime.
- The ability to represent chains in a directed graph that flexibly violate local transitivity. We achieve this by

way of a novel edge probability function that decays, asymmetrically, into the past and future timelike directions.

We illustrate the aforementioned features of our model by conducting a series of experiments on several small, simulated toy networks. Because of our emphasis on graph topology, we focus on the standard graph embedding challenge of link prediction. *Link prediction* is the task of inferring the missing edges of, and often solely from, a partially observed graph (Nickel et al., 2015). Premised on the assumption that the structures of real world graphs emerge from underlying mechanistic latent models (e.g. a biological evolutionary process responsible for the growth of a protein-protein interaction network, linguistic rules informing a language graph, etc), performance on the link prediction task hinges critically on one’s ability to render expressive graph representations, which pseudo-Riemannian embeddings allow for beyond existing embedding methods.

With this in mind, we highlight the quality of pseudo-Riemannian embeddings over Euclidean and hyperbolic embeddings in link prediction experiments using both synthetic protein-protein interaction (PPI) networks and the DREAM5 gold standard emulations of causal gene regulatory networks. Additionally, we show that our method has comparable performance to DAG-specific methods such as Disk Embeddings on the WordNet link prediction benchmark. Finally, we explore the ability of anti-de Sitter embeddings to further capture unique graph structures by exploiting critical features of the manifold, such as its intrinsic $S^1 \times \mathbb{R}^N$ topology for representing directed cycles of different lengths.

1.1. Related Work

The disadvantages of Euclidean geometry compared to Minkowski spacetime for graph representation learning was first highlighted in Sun et al. (2015). It was followed by Clough & Evans (2017) who explore DAG representations in Minkowski space, borrowing ideas from the theory of Causal Sets (Bombelli et al., 1987). More broadly, the asymptotic equivalence between complex networks and large-scale causal structure of de Sitter spacetime was proposed and studied in Krioukov et al. (2012). Our work is notably conceptually similar to the hyperbolic disk embedding approach (Suzuki et al., 2019) that embeds a set of symbolic objects with a partial order relation \preceq as *generalized formal disks* in a quasi-metric space (X, d) . A formal disk $(\mathbf{x}, r) \in X \times \mathbb{R}$ is defined by a *center* $\mathbf{x} \in X$ and a *radius* $r \in \mathbb{R}$.¹ Inclusion of disks defines a partial order on formal disks, which enables a natural representation of par-

tially ordered sets as sets of formal disks. These approaches all retain the partial-order transitivity assumption where squared distances decrease monotonically into the future and past. We relax that assumption in our work, alongside considering graphs with cycles and manifolds other than Minkowski spacetime.

Pseudo-Riemannian manifold optimization was formalized in Gao et al. (2018) and specialized in Law & Stam (2020) to general quadric surfaces in Lorentzian manifolds, which includes anti-de Sitter spacetime as a special case.

For the remainder of the paper, nodes are points on a manifold \mathcal{M} , the probability of edges are functions of the node coordinates, and the challenge is to infer the optimal embeddings via pseudo-Riemannian SGD on the node coordinates.

2. Background

In this section we will provide a very brief overview of the relevant topics in differential geometry.²

2.1. Riemannian Manifold Optimization

The key difference between gradient-based optimization of smooth functions f on Euclidean vs. non-Euclidean manifolds is that for the latter, the trivial isomorphism, for any $p \in \mathcal{M}$, between a manifold \mathcal{M} and the tangent space $T_p\mathcal{M}$ no longer holds in general. In particular, the stochastic gradient descent (SGD) update step $p' \leftarrow p - \lambda \nabla f|_p$ for learning rate λ and gradient ∇f is generalized in two areas (Bonnabel, 2013):

First, ∇f is replaced with the Riemannian gradient vector field

$$\nabla f \rightarrow \text{grad } f := g^{-1}df, \quad (1)$$

where $g^{-1} : T_p^*\mathcal{M} \rightarrow T_p\mathcal{M}$ is the inverse of the positive definite *metric* g , and df the differential one-form. Second, the *exponential map* $\exp_p : T_p\mathcal{M} \rightarrow \mathcal{M}$ generalizes the vector space addition in the update equation. For any $v_p \in T_p\mathcal{M}$ the first-order Taylor expansion is

$$f(\exp_p(v_p)) \approx f(p) + g(\text{grad } f|_p, v_p), \quad (2)$$

from which we infer that $\text{grad } f$ defines the direction of steepest descent, i.e. the Riemannian-SGD (RSGD) update step is simply

$$p' \leftarrow \exp_p(-\lambda \text{grad } f|_p). \quad (3)$$

The classes of manifolds considered here all have analytic expressions for the exponential map (see below). The curves traced out by $\exp_p(tv_p)$ for $t \in \mathbb{R}$ are called *geodesics* - the generalisation of straight lines to curved manifolds.

¹Suzuki et al. (2019) generalize the standard definition of a formal disk/ball to allow for negative radii.

²For a longer introduction, see Robbin & Salamon (2013) and Isham (1999).

2.2. Pseudo-Riemannian Extension

A *pseudo-Riemannian* (or, equivalently, *semi-Riemannian*) manifold is a manifold where g is non-degenerate but no longer positive definite. If g is diagonal with ± 1 entries, it is a *Lorentzian* manifold. If g has just one negative eigenvalue, it is commonly called a *spacetime* manifold. v_p is labelled *timelike* if $g(v_p, v_p)$ is negative, *spacelike* if positive, and *lightlike* or *null* if zero.

It was first noted in Gao et al. (2018) that $\text{grad } f$ is not a guaranteed ascent direction when optimizing f on pseudo-Riemannian manifolds, because its squared norm is no longer strictly positive (see eq. (2)). For diagonal coordinate charts one can simply perform a Wick-rotation (Visser, 2017; Gao et al., 2018) to their Riemannian counterpart and apply eq. (3) legitimately; in all other cases, additional steps are required to reintroduce the guarantee (see Section 2.3.2 eq. (11)).

2.3. Example Manifolds

2.3.1. MINKOWSKI SPACETIME

Minkowski spacetime is the simplest Lorentzian manifold with metric $g = \text{diag}(-1, 1, \dots, 1)$. The $N + 1$ coordinate functions are $(x_0, x_1, \dots, x_N) \equiv (x_0, \mathbf{x})$ with x_0 the time coordinate. The squared distance s^2 between two points with coordinates x and y is

$$s^2 = -(x_0 - y_0)^2 + \sum_{i=1}^N (x_i - y_i)^2. \quad (4)$$

Because the metric is diagonal and flat, the RSGD update is made with the simple map

$$\exp_p(v_p) = p + v_p, \quad (5)$$

where v_p is the vector with components $(v_p)^i = \delta^{ij} (df_p)_j$, with δ^{ij} is the trivial Euclidean metric.

2.3.2. ANTI-DE SITTER SPACETIME

The $(N + 1)$ -dimensional anti-de Sitter spacetime (Bengtsson, 2018) AdS_N can be defined as an embedding in $(N + 2)$ -dimensional Lorentzian manifold (L, g_L) with two ‘time’ coordinates. Given the canonical coordinates $(x_{-1}, x_0, x_1, \dots, x_N) \equiv (x_{-1}, x_0, \mathbf{x})$, AdS_N is the quadric surface defined by $g_L(x, x) \equiv \langle x, x \rangle_L = -1$, or in full,

$$-x_{-1}^2 - x_0^2 + x_1^2 + \dots + x_N^2 = -1. \quad (6)$$

Another useful set of coordinates involves the polar reparameterization $(x_{-1}, x_0) \rightarrow (r, \theta)$:

$$x_{-1} = r \sin \theta, \quad x_0 = r \cos \theta, \quad (7)$$

where by simple substitution

$$r \equiv r(\mathbf{x}) = \left(1 + \sum_{i=1}^N x_i^2 \right)^{\frac{1}{2}}. \quad (8)$$

We define a circle time coordinate to be the arc length

$$t := r\theta = \begin{cases} r \sin^{-1}(x_{-1}/r), & x_0 \geq 0, \\ r(\pi - \sin^{-1}(x_{-1}/r)), & x_0 < 0, \end{cases} \quad (9)$$

with \mathbf{x} -dependent periodicity $t \sim t + 2\pi r(\mathbf{x})$.

The canonical coordinates and metric g_L are not intrinsic to the manifold, so we must treat the pseudo-Riemannian gradient from eq. (1) with the projection operator $\Pi_p : T_p L \rightarrow T_p \text{AdS}_N$, defined for any $v_p \in T_p L$ to be (Robbin & Salamon, 2013)

$$\Pi_p v_p = v_p + g_L(v_p, p)p. \quad (10)$$

Furthermore, as Law & Stam (2020) proved recently for general quadric surfaces, one can sidestep the need to perform the computationally expensive Gram-Schmidt orthogonalization by implementing a double projection, i.e. we have

$$\zeta_p = \Pi_p \left(g_L^{-1} (\Pi_p (g_L^{-1} df)) \right) \quad (11)$$

as our guaranteed descent tangent vector.

The squared distance s^2 between $p, q \in \text{AdS}_N$ is given by

$$s^2 = \begin{cases} -|\cos^{-1}(-\langle p, q \rangle_L)|^2, & -1 < \langle p, q \rangle_L \leq 1, \\ (\cosh^{-1}(-\langle p, q \rangle_L))^2, & \langle p, q \rangle_L < -1, \\ 0, & \langle p, q \rangle_L = -1, \\ -\pi^2, & \langle p, q \rangle_L > 1, \end{cases} \quad (12)$$

where the first three cases are the squared geodesic distances between time-, space-, and lightlike separated points respectively. For $\langle p, q \rangle_L > 1$, there are no geodesics connecting the points. However, we require a smooth loss function in s^2 with complete coverage for all p, q pairs, so we set s^2 to the value at the timelike limit $\langle p, q \rangle_L = 1$.

The exponential map is given by

$$\exp_p(\zeta_p) = \begin{cases} \cos(\|\zeta_p\|)p + \sin(\|\zeta_p\|) \frac{\zeta_p}{\|\zeta_p\|}, \\ \cosh(\|\zeta_p\|)p + \sinh(\|\zeta_p\|) \frac{\zeta_p}{\|\zeta_p\|}, \\ p + \zeta_p, \end{cases} \quad (13)$$

again for the time-, space-, and lightlike ζ_p respectively, and where $\|\zeta_p\| \equiv \sqrt{|g_L(\zeta_p, \zeta_p)|}$.

3. Pseudo-Riemannian Embedding Models

In this section we describe the key components of our embedding model, starting with defining a probability function for a directed edge between any two nodes in a graph.

3.1. Triple Fermi-Dirac Function

The Fermi-Dirac (FD) distribution function³

$$F_{(\tau,r,\alpha)}(x) := \frac{1}{e^{(\alpha x - r)/\tau} + 1}, \quad (14)$$

with $x \in \mathbb{R}$ and parameters $\tau, r \geq 0$ and $0 \leq \alpha \leq 1$, is used to calculate the probabilities of undirected graph edges as functions of node embedding distances (Krioukov et al., 2010; Nickel & Kiela, 2017). For general directed graphs one needs to specify a preferred direction in the embedding space.⁴ This was the approach taken in Suzuki et al. (2019), using the radius coordinate, and our method follows this same principle.

For spacetime manifolds, the time dimension is the natural coordinate, indicating the preferred direction. For two points $p, q \in \mathcal{M}$, we propose the following distribution \mathcal{F} which we refer to as the *Triple Fermi-Dirac* (TFD) function:

$$\mathcal{F}_{(\tau_1, \tau_2, \alpha, r, k)}(p, q) := k(F_1 F_2 F_3)^{1/3}, \quad (15)$$

where $k > 0$ is a tunable scaling factor and

$$F_1 := F_{(\tau_1, r, 1)}(s^2), \quad (16)$$

$$F_2 := F_{(\tau_2, 0, 1)}(-\Delta t), \quad (17)$$

$$F_3 := F_{(\tau_2, 0, \alpha)}(\Delta t), \quad (18)$$

are three FD distribution terms. τ_1, τ_2, r and α are the parameters from the component FD terms (14), s^2 the squared geodesic distance between p and q , and Δt the time coordinate difference, which in the case of Minkowski spacetime, is simply $x_0(q) - x_0(p)$. It is easy to see that for $k \leq 1$, $\mathcal{F}(p, q)$ is a valid probability value between 0 and 1.

The motivations behind eq. (15) can be understood by way of a series of simple plots of \mathcal{F} and its component FD terms over a range of spacetime intervals on 2D Minkowski spacetime. In Figure 1 we fix p at the origin and plot \mathcal{F} as a function of q . Concretely we demonstrate in the remainder of this section that the TFD function combines with the pseudo-Riemannian metric structure to encode three specific inductive biases inherent in general directed graphs: 1. the co-existence of semantic and edge-based similarities, 2. the varying levels of transitivity, and 3. the presence of graph cycles.

³The case of $\alpha = 1$ is the usual parameterization.

⁴There are situations where the probability function itself is isotropic but that the features of the graph displays something analogous to spontaneous symmetry breaking where the optimal node embeddings identifies a natural direction. This is indeed the case for tree-like graphs and hyperbolic embeddings where the direction of exponential volume growth lines up with the exponential growth in the number of nodes as one goes up the tree. However for general directed graphs, an explicit symmetry breaking term in the function is needed in our case.

3.1.1. SEMANTIC SIMILARITY VS. DIRECT NEIGHBORS

A pair of graph vertices x, y can be similar in two ways. They could define an edge, in which case they are simply neighbors, or they could have overlapping neighbor sets N_x and N_y . However, in Riemannian manifold graph embedding models, where node distances determine edge probabilities, a high degree of this latter *semantic similarity* suggests that x and y are in close proximity, especially if both N_x and N_y are themselves sparsely connected (Sun et al., 2015). This can be in conflict with the absence of an edge between the vertex pair. A schematic of this structure is shown in Figure 2A. For example, in sporting events this can occur when two members of a team may compete against a largely overlapping set of competitors, but would never be pitted against one another.

Embedding graphs in a spacetime resolves this inconsistency – for a pair of vertices that do not share an edge there is no constraint on the Jaccard index $(N_x \cap N_y)/(N_x \cup N_y)$. This claim can be verified by examining the contribution from the first FD term F_1 (16) in the TFD function. As shown in Figure 1D, F_1 is low when x and y are spacelike separated and high when within each other’s past and future lightcones. Our claim can then be rephrased in geometric terms by stating that two spacelike separated points (hence low edge probability) can have an arbitrarily large overlap in their lightcones (and hence in their neighbor sets).

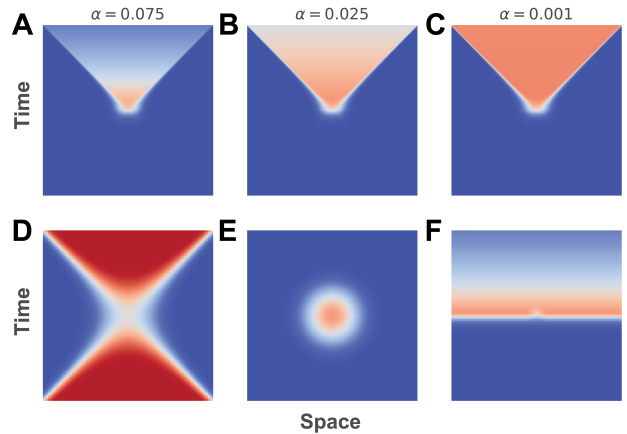


Figure 1. **A-C:** The TFD function \mathcal{F} (eq. (15)) in 2D Minkowski spacetime for varying α . **D:** The FD function F_1 (eq. (16)) in Minkowski space, **E:** F_1 in Euclidean space, **F:** \mathcal{F} in Euclidean space. Blue to red corresponds probabilities from 0 to 1.

3.1.2. DIRECTED NON-TRANSITIVE RELATIONS

Simply using F_1 for edge probabilities is problematic in two ways. First, the past and future are indistinguishable (Figure 1D), and hence too the probabilities for both edge directions of any node pair. Second, lightcones, being cones,

impose a strict partial ordering on the graph. Specifically, any univariate function of s^2 that assigns, by way of future lightcone inclusion, a high probability to the directed chain

$$p_1 \rightarrow p_2 \rightarrow p_3 \quad (19)$$

will always predict the transitive relation $p_1 \rightarrow p_3$ with a higher probability than either $p_1 \rightarrow p_2$ or $p_2 \rightarrow p_3$ (see Figure 3A). This is a highly restrictive condition imposed by a naive graph embedding on a pseudo-Riemannian manifold. In real-world networks such as social networks, these friend-of-friend links are often missing, violating this constraint.

We tackle both the need for edge anti-symmetries and for flexible transitive closure constraints via the other two FD terms F_2 and F_3 (Eqs. (17) and (18)). Both terms introduce exponential decays into the respective past and future timelike directions, thereby breaking the strict partial order. Crucially, they are asymmetric due to the parameter α in F_3 , which introduces different decay rates. As shown in Figure 1A, the TFD function has a local maximum in the near timelike future. α controls the level of transitivity in the inferred graph embeddings, as we see in Figures 1A-C. Euclidean disk embeddings (Suzuki et al., 2019) can be viewed as approximately equivalent to the $\alpha = 0$ case in flat Minkowski spacetime, where the region of high edge probability extends far out into the future lightcone (see Supplementary Material for a precise statement of this equivalence).

We note that the TFD function is well-defined on Riemannian spaces, as long as one designates a coordinate to take the role of the time dimension (see Figure 1E-F).

3.1.3. GRAPH CYCLES

Another feature of the TFD probabilistic model in (15) is that the lightcone boundaries are soft transitions that can be modified by adjusting the temperature hyperparameters τ_1 and τ_2 . Specifically, short (in a coordinate-value sense) directed links into the past or in spacelike directions have probabilities close to $1/2$, as can be verified by a Taylor expansion of eq. (15) around $p = q = 0$ (see Figure 1D).

This feature alone does not constitute a sufficient basis to promote pseudo-Riemannian embeddings with the TFD distribution as a model for cyclic graph representations (e.g. the toy example in Figure 4A). Embedding long directed cycles in this way is not efficient. For example, a length- n chain with $O(n^2)$ number of possible (and missing) transitive edges would have a similar embedding pattern to a fully-connected clique. This distinction is important when modeling real world systems, such as gene regulation, in which important information is encoded in sparse networks with cycles (Leclerc, 2008). For this we turn our attention to the global topology of the manifold.

3.2. Cylinder Topology: $S^1 \times \mathbb{R}^N$

The problem with embedding cycles via a model that favors future-timelike directed links is the unavoidable occurrence of at least one low-probability past-directed edge. Here we propose a circular time dimension as a global topological solution. We consider two such pseudo-Riemannian manifolds with a $S^1 \times \mathbb{R}^N$ topology - a modified cylindrical Minkowski spacetime and anti-de Sitter spacetime.

3.2.1. TFD FUNCTION ON CYLINDRICAL MINKOWSKI SPACETIME

For an $(N+1)$ -dimensional spacetime with time coordinates (x_0, \mathbf{x}) , we construct our cylinder by identifying

$$x_0 \sim x_0 + nC, \quad (20)$$

for some circumference $C > 0$ and $n \in \mathbb{Z}$. To ensure a smooth TFD function on the cylinder we define a *wrapped TFD* function as

$$\tilde{\mathcal{F}}(p, q) := \sum_{n=-\infty}^{\infty} \mathcal{F}(p, q^{(n)}), \quad (21)$$

where $x_0(q^{(n)}) \equiv x_0(q) + nC$, with all other coordinates equal. Unlike the case of the wrapped normal distribution, there is no (known) closed-form expression for $\tilde{\mathcal{F}}$. However, provided $\alpha > 0$, the exponentially decreasing probabilities from F_2 and F_3 into the far past and future time directions respectively enables one to approximate the infinite sum in eq. (21) by truncating to a finite sum over integers from $-m$ to m . The scaling factor k in eq. (15) can be used to ensure $\max \tilde{\mathcal{F}} \leq 1$.

In this topology, the concept of space- and timelike separated points is purely a local feature, as there will be multiple timelike paths between *any* two points on the manifold.

3.2.2. TFD FUNCTION ON ANTI-DE SITTER SPACETIME

Unlike for Minkowski spacetime, AdS already has an intrinsic $S^1 \times \mathbb{R}^N$ topology.

To adapt the TFD distribution to AdS space, we adopt polar coordinates and let the angular difference between p and q be

$$\theta_{pq} := \theta_q - \theta_p, \quad (22)$$

where θ_p and θ_q correspond to the polar angle θ from eq. (7). Then we define

$$\Delta t_{pq} := r_q \theta_{pq}. \quad (23)$$

Similar to the cylindrical Minkowski case in eq. (21) we modify \mathcal{F} such that for two points p and q , we have

$$\mathcal{F}(p, q) = \mathcal{F}(p, q^{(n)}), \quad (24)$$

where $q^{(n)} \in \text{AdS}_N$ have identical coordinates with q except that

$$\Delta t_{pq}^{(n)} := \Delta t_{pq} + 2\pi n r q. \quad (25)$$

The wrapped TFD distribution is identical to eq. (21) with Δt_{pq} as the time difference variable Δt in the F_2 and F_3 components (eqs. (17) and (18) respectively).

4. Experiments

4.1. Datasets and Training Details

In this section we evaluate the quality of pseudo-Riemannian embeddings via a series of graph reconstruction and link prediction experiments. Using small synthetic graphs, we begin with a demonstration of the model’s ability to encode the particular set of graph-specific features outlined in Section 3. We then run the model on a series of benchmarks and an ablation study to characterize the utility of these embeddings in downstream applications. For this second group of experiments, we rely on the following three classes of directed graph datasets.

Duplication Divergence Model: A two-parameter model that simulates the growth and evolution of large protein-protein interaction networks (Ispolatov et al., 2005). Depending on the topology of the initial seed graphs, the final graph is either a DAG or a directed graph with cycles.

DREAM5: Gold standard edges from a genome-scale network inference challenge, comprising of a set of gene regulatory networks across organisms and an *in silico* example (Marbach et al., 2012). These networks contain a relatively small number of cycles.

WordNet: An acyclic, hierarchical, tree-like network of nouns, each with relatively few ancestors and many descendants. We consider networks with different proportions of transitive closure. We use the same train / validation / test split as in Suzuki et al. (2019) and Ganea et al. (2018).

In all our experiments, we seek to minimize the negative log-likelihood (NLL) loss based on the probabilities (14) or (15). We fix a negative sampling ratio of 4 throughout. Similar to Nickel & Kiela (2017), we initialize our embeddings in a small random patch near the origin ($x = (1, 0, \dots, 0)$) for AdS and perform a burn-in phase of several epochs with the learning rate scaled by a factor of 0.01. All of our models and baselines are run across various dimensions to ensure our methods do not display higher performance only at certain dimensionality level. We provide full details of each experiment, including the hyperparameter sweep and average runtimes, alongside a more full description of the datasets in the Supplementary Material. The models in this paper were implemented in JAX (Bradbury et al., 2018).

4.2. Graph Feature Encoding




Figure 2. **A** Tri-partite graph of an unconnected node pair with large number of common predecessors and successors. **B** Euclidean embeddings with inferred edge probability between the node pair as 0.50. **C** Minkowski embeddings with inferred edge probability of 0.01.

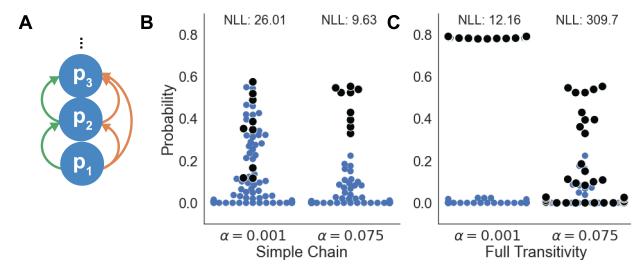


Figure 3. Performance of Minkowski embeddings for varying α values in the TFD function on two example graphs. **A:** Schematic of a three-node graph, with green edges indicating a simple chain, and orange edges indicating a full transitive closure. **B:** 10-node chain. **C:** 10-node fully transitive DAG. Black and blue markers indicate positive and negative directed edges respectively. Also annotated is the NLL of the graph/model pair.

Node pairs with arbitrarily large overlapping neighbor sets can remain unconnected.

We construct a graph consisting of an unconnected node pair, and two large sets of unconnected, common predecessors and successors (Figure 2A). The Euclidean embedding model trained with the FD edge probability function cannot reconcile the high semantic similarity of the node pair with the absence of an edge; we see that the pair of nodes are unavoidably closer to each other than the majority of their neighbor connections. (Figure 2B). On the contrary, the Minkowski model effectively encodes the tri-partite graph as three sets of spacelike separated points with high degree of overlap in the lightcones of the nodes in the pair (Figure 2C).

\mathcal{F} can be tuned to encode networks with varying proportion of transitive relations

We construct a simple, 10-node chain and a fully transitive variant to explore our ability to tune the TFD loss to best capture graph transitive

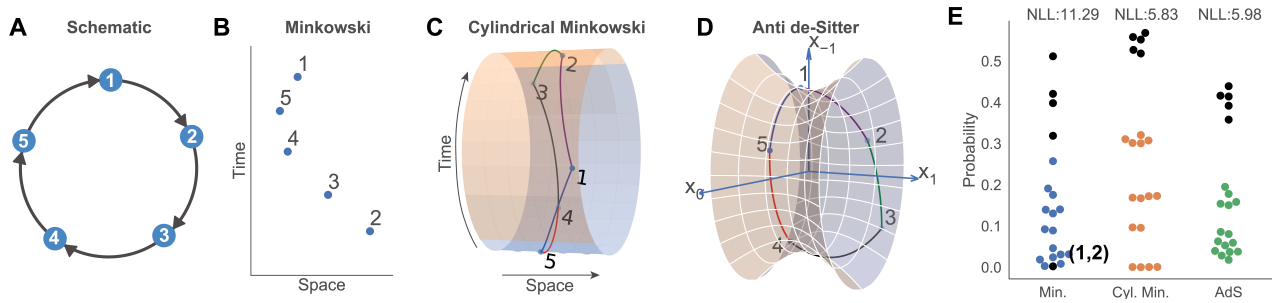


Figure 4. **A:** Illustration of a 5-node chain used as training data. **B-D:** 5-node chain embedding on various manifolds. **E:** Edge probabilities for 5-node cycle using Minkowski (Min.), cylindrical Minkowski (Cyl. Min.) and AdS embeddings. The five positive edges are indicated in black dots while the 15 negative edges are the dots with colors.

structure, using the alpha parameter to vary probability decay in the time direction. We use two α values, $\alpha = 0.001$ and $\alpha = 0.075$, corresponding to the heatmaps shown in Figure 1A and C. In Figure 3, we see that smaller values of α are able to more successfully capture full transitivity. In the case of a simple chain, setting $\alpha = 0.001$ results in a higher probability being assigned to negative edges, compared to the $\alpha = 0.075$ case.

Graph cycles wrap around the S^1 dimension. We embed a simple five-node loop (Figure 4A) in 2D Minkowski (B), cylindrical Minkowski (C), and anti-de Sitter (D) spacetimes. In all three manifold embeddings, we recover the true node ordering as reflected in the ordered time coordinates. However for the Minkowski manifold, there is one positive edge ($1 \rightarrow 2$ in this case) that unavoidably gets incorrectly assigned a very low probability, which we indicate in Figure 4E. On the contrary, the S^1 time dimension for cylindrical Minkowski and AdS spacetime ensures that all edges have high probabilities thereby demonstrating the utility of the circle time dimension. We note that all the pseudo-Riemannian embedding models, including the (non-cylindrical) Minkowski case, outperform the Euclidean (NLL: 17.24) and the hyperboloid models (NLL: 13.89) (not shown in Figure 4).

4.3. Cyclic Graph Inference

We perform link prediction on the Duplication Divergence Model network and the *in silico* DREAM5 datasets.⁵ The results on randomly held-out positive edges for a range of embedding models of various dimensions are presented in Table 1. Here we see that for both datasets pseudo-Riemannian embeddings significantly outperform Euclidean and Hyper-

⁵For full results on DREAM5 networks for *Escherichia coli* and *Saccharomyces Cerevisiae* please refer to the Supplementary Material. The conclusions there are similar to the main text experiment.

boloid embedding baselines across all dimensions, with the cylindrical Minkowski model achieving the best result.

Next we perform an ablation study to examine the relative contributions of 1. the pseudo-Riemannian geometry, as reflected in the presence of negative squared distances, 2. the global cylindrical topology, and 3. TFD vs. FD likelihood model. The results are presented in Figure 5.

We see that all three components combine to produce the superior performances of the anti-de Sitter and cylindrical Minkowski manifold embeddings, with the S^1 topology arguably being the biggest contributor.

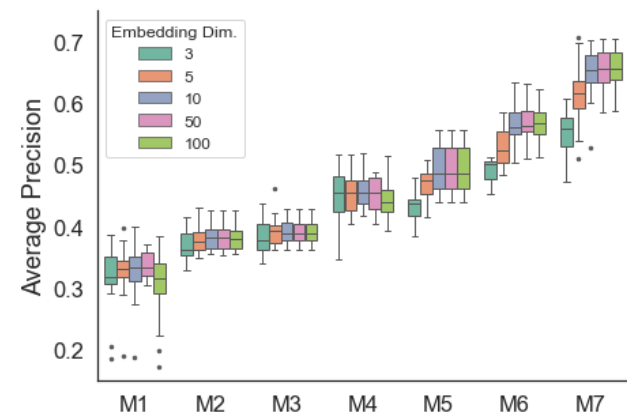


Figure 5. The average precision values of link prediction on the Duplication Divergence Model graph across manifold / likelihood model combinations and embedding dimensions, over multiple random seeds. M1: Euclidean manifold with TFD, M2: Hyperboloid + FD, M3: Euclidean with FD, M4: cylindrical Euclidean + TFD, M5: Minkowski + TFD, M6: Anti-de Sitter + TFD, M7: cylindrical Minkowski + TFD.

Table 1. Link prediction for directed cyclic graphs. We show the median average precision (AP) percentages across 20 random initializations on a held-out test set, calculated separately for different embedding dimensions d . Annotated in **bold** is the top-performing model for the given dimension. For reference, the asymptotic random baseline AP is 20%.

	Duplication Divergence Model					DREAM5 : <i>in silico</i>				
	$d = 3$	$d = 5$	$d = 10$	$d = 50$	$d = 100$	$d = 3$	$d = 5$	$d = 10$	$d = 50$	$d = 100$
Euclidean + FD	37.8	39.4	39.0	38.9	38.9	29.4	32.9	39.7	39.8	34.8
Hyperboloid + FD	36.3	37.5	38.2	38.2	38.1	28.8	46.8	50.8	50.9	52.5
Minkowski + TFD	43.7	47.5	48.5	48.5	48.5	36.3	43.1	51.2	57.7	58.0
Anti-de Sitter + TFD	50.1	52.4	56.2	56.3	56.8	38.1	45.2	51.9	55.6	56.0
Cylindrical Minkowski + TFD	55.8	61.6	65.3	65.7	65.6	41.0	48.4	56.3	58.9	61.0

4.4. DAG Inference

We test our method on the WordNet dataset, which is a DAG that has been used in prior work on DAG inference (Nickel & Kiela, 2017; Ganea et al., 2018; Suzuki et al., 2019). Table 2 shows the performance of the pseudo-Riemannian approach on the link prediction task, benchmarked against a set of competitive methods (results taken from Suzuki et al. (2019)). We break down the methods into flat- and curved-space embedding models. Our approach outperforms all of the other flat-space as well as some curved-space methods, including Poincare embeddings (Nickel & Kiela, 2017). Furthermore, the Minkowski embedding model achieves competitive performance with hyperbolic and spherical approaches. These methods are well suited to representing hierarchical relationships central to WordNet, thus this result shows that pseudo-Riemannian models are highly capable of representing hierarchies by encoding edge direction. We can show that the representational power of a special case of the TFD probability on flat Minkowski spacetime is similar to that of Euclidean Disk Embeddings (see Supplementary Material). The difference in performance between Euclidean Disk Embeddings and our model on this task could be due to the additional flexibility allowed by the TFD probability function or in differences in the resulting optimization problem, something that could be further explored in future work.

Table 2. F1 percentage score on the test data of WordNet. The best flat-space performance (top-half) for each dataset/embedding dimension combination has been highlighted in gray and the best overall is in **bold**. The benchmark method’s results were taken from Suzuki et al. (2019). The full results table is presented in the Supplementary Material.

Transitive Closure Percentage	$d = 5$		$d = 10$	
	25%	50%	25%	50%
Minkowski + TFD (Ours)	86.2	92.1	89.3	94.4
Order Emb. (Vendrov et al., 2016)	75.9	82.1	79.4	84.1
Euclidean Disk (Suzuki et al., 2019)	42.5	45.1	65.8	72.0
Spherical Disk (Suzuki et al., 2019)	90.5	93.4	91.5	93.9
Hyperbolic EC (Ganea et al., 2018)	87.1	92.8	90.8	93.8
Hyperbolic Disk (Suzuki et al., 2019)	81.3	83.1	90.5	94.2
Poincare Emb. (Nickel & Kiela, 2017)	78.3	83.9	82.1	85.4

5. Discussion and Future Work

While Riemannian manifolds embeddings have made significant advances in representing complex graph topologies, there are still areas in which these methods fall short, due largely to the inherent constraints of the metric spaces used. In this paper, we demonstrated how pseudo-Riemannian embeddings can effectively address some of the open challenges in graph representation learning, summarized in the following key results. Firstly, we are able to successfully represent directed cyclic graphs using the novel Triple Fermi-Dirac probability function and the cylindrical topology of the manifolds, and to disambiguate direct neighbors from functionally similar but unlinked nodes, as demonstrated in a series of experiments on synthetic datasets. Also, Minkowski embeddings strongly outperform Riemannian baselines on a link prediction task in directed cyclic graphs, and achieve results comparable with state-of-the-art methods on DAGs across various dimensions. In this latter case, applying our approach gave us the flexibility to lift the constraints of transitivity due to the temporal decay of the TFD probability function. Using these characteristics, we demonstrate superior performance on a number of directed cyclic graphs: the duplication-divergence model and a set of three DREAM5 gold standard gene regulatory networks.

There are two areas for further work. First, in addition to further validation of the pseudo-Riemannian optimization procedure introduced in Law & Stam (2020), one can experiment with AdS coordinates different to the one proposed in Section 3.2.2, or extend the methods in this work to more general classes of pseudo-Riemannian manifolds. Second, we note that the field closely related to link prediction is network inference – inferring causal networks from observational and interventional data is a critical problem with a wide range of applications. In biology, for example, robust and accurate inference of gene regulatory networks from gene expression data is a long-standing grand challenge in systems biology. In this case, limited ground-truth connectivity is usually already available as a prior – a logical extension of the model presented here would be a hybrid model that combines pseudo Riemannian-based link prediction with metric space network inference.

Acknowledgements

We would like to thank Craig Glastonbury for helpful discussions and suggestions on the manuscript.

References

- Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., Gulcehre, C., Song, F., Ballard, A., Gilmer, J., Dahl, G., Vaswani, A., Allen, K., Nash, C., Langston, V., Dyer, C., Heess, N., Wierstra, D., Kohli, P., Botvinick, M., Vinyals, O., Li, Y., and Pascanu, R. Relational inductive biases, deep learning, and graph networks. *arXiv*, 2018.
- Bengtsson, I. Anti-de Sitter Space, 2018. URL <http://3dhouse.se/ingemar/relteori/Kurs.pdf>.
- Bombelli, L., Lee, J., Meyer, D., and Sorkin, R. D. Spacetime as a causal set. *Physical review letters*, 59(5):521, 1987.
- Bonnabel, S. Stochastic gradient descent on Riemannian manifolds. *IEEE Transactions on Automatic Control*, 58(9):2217–2229, 2013.
- Bouchard, G., Singh, S., and Trouillon, T. On Approximate Reasoning Capabilities of Low-Rank Vector Spaces. *AAAI spring symposia*, 2015.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Chami, I., Wolf, A., Juan, D.-C., Sala, F., Ravi, S., and Ré, C. Low-dimensional hyperbolic knowledge graph embeddings. *arXiv preprint arXiv:2005.00545*, 2020.
- Clough, J. R. and Evans, T. S. Embedding graphs in Lorentzian spacetime. *PLOS ONE*, 12(11):e0187301, 2017.
- Ganea, O., Becigneul, G., and Hofmann, T. Hyperbolic entailment cones for learning hierarchical embeddings. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, pp. 1646–1655, Stockholm, Sweden, 2018.
- Gao, T., Lim, L.-H., and Ye, K. Semi-Riemannian Manifold Optimization. *arXiv*, 2018.
- Grover, A. and Leskovec, J. Node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 855–864, 2016.
- Gu, A., Sala, F., Gunel, B., and Ré, C. Learning mixed-curvature representations in product spaces. In *International Conference on Learning Representations*, 2018.
- Isham, C. J. *Modern differential geometry for physicists*, volume 61. World Scientific Publishing Company, 1999.
- Ispolatov, I., Krapivsky, P. L., and Yuryev, A. Duplication-divergence model of protein interaction network. *Phys. Rev. E*, 71:061911, Jun 2005.
- Krioukov, D., Papadopoulos, F., Kitsak, M., Vahdat, A., and Boguñá, M. Hyperbolic geometry of complex networks. *Phys. Rev. E*, 82:036106, Sep 2010.
- Krioukov, D., Kitsak, M., Sinkovits, R. S., Rideout, D., Meyer, D., and Boguñá, M. Network cosmology. *Scientific reports*, 2(1):1–6, 2012.
- Law, M. T. and Stam, J. Ultrahyperbolic representation learning. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33*, 2020.
- Leclerc, R. D. Survival of the sparsest: robust gene networks are parsimonious. *Molecular systems biology*, 4(1):213, 2008.
- Marbach, D., Costello, J. C., Küffner, R., Vega, N. M., Prill, R. J., Camacho, D. M., Allison, K. R., Aderhold, A., Allison, K. R., Bonneau, R., Camacho, D. M., Chen, Y., Collins, J. J., Cordero, F., Costello, J. C., Crane, M., Dondelinger, F., Drton, M., Esposito, R., Foygel, R., Fuente, A. d. I., Gertheiss, J., Geurts, P., Greenfield, A., Grzegorzczak, M., Haury, A.-C., Holmes, B., Hothorn, T., Husmeier, D., Huynh-Thu, V. A., Irrthum, A., Kellis, M., Karlebach, G., Küffner, R., Lèbre, S., Leo, V. D., Madar, A., Mani, S., Marbach, D., Mordelet, F., Ostrer, H., Ouyang, Z., Pandya, R., Petri, T., Pinna, A., Poultney, C. S., Prill, R. J., Rezny, S., Ruskin, H. J., Saeys, Y., Shamir, R., Sîrbu, A., Song, M., Soranzo, N., Statnikov, A., Stolovitzky, G., Vega, N., Vera-Licona, P., Vert, J.-P., Visconti, A., Wang, H., Wehenkel, L., Windhager, L., Zhang, Y., Zimmer, R., Kellis, M., Collins, J. J., and Stolovitzky, G. Wisdom of crowds for robust gene network inference. *Nature Methods*, 9(8):796–804, 2012. ISSN 1548-7091.
- Nickel, M. and Kiela, D. Poincaré embeddings for learning hierarchical representations. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 6338–6347, 2017.
- Nickel, M. and Kiela, D. Learning continuous hierarchies in the Lorentz model of hyperbolic geometry. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International*

Conference on Machine Learning, volume 80, pp. 3779–3788, 2018.

Nickel, M., Tresp, V., and Kriegel, H.-P. A three-way model for collective learning on multi-relational data. In Getoor, L. and Scheffer, T. (eds.), *Proceedings of the 28th International Conference on Machine Learning*, pp. 809–816, 2011.

Nickel, M., Murphy, K., Tresp, V., and Gabrilovich, E. A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE*, 104(1):11–33, 2015.

Perozzi, B., Al-Rfou, R., and Skiena, S. Deepwalk: On-line learning of social representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 701–710, 2014.

Robbin, J. W. and Salamon, D. A. Introduction to differential geometry, March 2013.

Sala, F., De Sa, C., Gu, A., and Re, C. Representation trade-offs for hyperbolic embeddings. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4460–4469, 2018.

Sun, K., Wang, J., Kalousis, A., and Marchand-Maillet, S. Space-time local embeddings. In Cortes, C., Lawrence, N., Lee, D., Sugiyama, M., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 28*, pp. 100–108, 2015.

Suzuki, R., Takahama, R., and Onoda, S. Hyperbolic disk embeddings for directed acyclic graphs. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, pp. 6066–6075, 2019.

Trouillon, T., Welbl, J., Riedel, S., Gaussier, E., and Bouchard, G. Complex embeddings for simple link prediction. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of the 33rd International Conference on Machine Learning*, pp. 2071–2080, 2016.

Vendrov, I., Kiros, R., Fidler, S., and Urtasun, R. Order-embeddings of images and language. In Bengio, Y. and LeCun, Y. (eds.), *4th International Conference on Learning Representations*, 2016.

Vilnis, L. and McCallum, A. Word representations via Gaussian embedding. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations*, 2015.

Visser, M. How to Wick rotate generic curved spacetime. *arXiv*, 2017.