

Hypformer: Exploring Efficient Hyperbolic Transformer Fully in Hyperbolic Space

Menglin Yang

Yale University

New Haven, United States

menglin.yang@outlook.com

Harshit Verma

Birla Institute of Technology and

Science

Hyderabad, India

verma08harshit@gmail.com

Delvin Ce Zhang

Yale University

New Haven, United States

delvincezhang@gmail.com

Jiahong Liu

The Chinese University of Hong Kong

Hong Kong, China

jiahong.liu21@gmail.com

Irwin King

The Chinese University of Hong Kong

Hong Kong, China

king@cse.cuhk.edu.hk

Rex Ying

Yale University

New Haven, United States

rex.ying@yale.edu

Abstract

Hyperbolic geometry have shown significant potential in modeling complex structured data, particularly those with underlying tree-like and hierarchical structures. Despite the impressive performance of various hyperbolic neural networks across numerous domains, research on adapting the Transformer to hyperbolic space remains limited. Previous attempts have mainly focused on modifying self-attention modules in the Transformer. However, these efforts have fallen short of developing a complete hyperbolic Transformer. This stems primarily from: (i) the absence of well-defined modules in hyperbolic space, including linear transformation layers, Layer-Norm layers, activation functions, dropout operations, etc. (ii) the quadratic time complexity of the existing hyperbolic self-attention module w.r.t the number of input tokens, which hinders its scalability. To address these challenges, we propose, **Hypformer**, a novel hyperbolic Transformer based on the Lorentz model of hyperbolic geometry. In Hypformer, we introduce two foundational blocks that define the essential modules of the Transformer in hyperbolic space. Furthermore, we develop a *linear* self-attention mechanism in hyperbolic space, enabling hyperbolic Transformer to process billion-scale graph data and long-sequence inputs for the first time. Our experimental results confirm the effectiveness and efficiency of Hypformer across various datasets, demonstrating its potential as an effective and scalable solution for large-scale data representation and large models.

CCS Concepts

- Computing methodologies → Machine learning; Knowledge representation and reasoning; • Mathematics of computing → Geometric topology.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '24, August 25–29, 2024, Barcelona, Spain

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0490-1/24/08

<https://doi.org/10.1145/3637528.3672039>

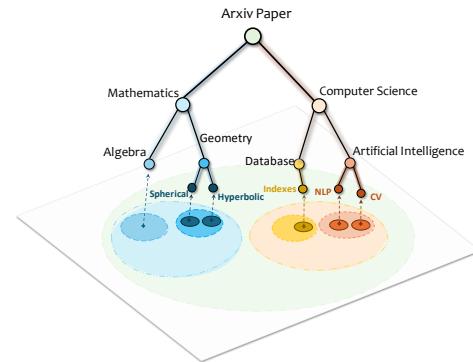


Figure 1: In a variety of real-world scenes, when we classify instances in a dataset (e.g., classifying the node in an Arxiv paper), we can group them into larger groups (e.g., Physics, Computer Science) that contain smaller subgroups (e.g., {Quantum Computing, Geometry}, {Database, Artificial Intelligence}), which may also contain even smaller sub-subgroups. The relationships between these various levels of groups and subgroups can be represented by dendograms, which are tree-like structures that reveal the underlying hierarchies in the data.

Keywords

Transformer; Hyperbolic geometry; Linear self-attention; Foundation model

ACM Reference Format:

Menglin Yang, Harshit Verma, Delvin Ce Zhang, Jiahong Liu, Irwin King, and Rex Ying. 2024. Hypformer: Exploring Efficient Hyperbolic Transformer Fully in Hyperbolic Space. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24)*, August 25–29, 2024, Barcelona, Spain. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3637528.3672039>

1 Introduction

In many real-world scenarios, data frequently exhibit a hierarchical or tree-like structure, either implicitly or explicitly [32, 49, 95]. This is evident in complex networks [8, 38, 78, 94], the semantics of words in natural language processing [50, 51, 63], and conceptual hierarchies in vision tasks [15, 32]. As illustrated in Figure 1, such

data can be organized into large and abstract groups that encompass small and specific subgroups, which can further be subdivided into even smaller and more specific sub-subgroups, and so on. The relationships between these groups and subgroups can be effectively approximated by tree-like structures [13]. This hierarchical representation mirrors human cognitive processes [14, 28], making it an intuitive approach to data representation.

Recent initiatives have explored the use of hyperbolic learning spaces to encode complex non-Euclidean data, achieving impressive performance in representing tree-like data [8, 24, 43, 50, 51, 62, 76, 79, 80, 83]. This success is attributed to the unique property of hyperbolic space, which expands exponentially compared to the polynomial expansion of Euclidean spaces. This property aligns hyperbolic space with the metric of trees, making it particularly suitable for representing tree-like or hierarchically structured data [38]. Despite the growing interest in hyperbolic representation and deep learning, the Transformer [1, 65, 70], a cornerstone model in the various domains, was seldom explored within the realm of hyperbolic space. Despite preliminary attempts in hyperbolic Transformers [9, 26, 61], numerous challenges remain to be addressed.

Challenge (1): Insufficient definitions for operations in the hyperbolic Transformer. Prior works of HAN [26] and HNN++ [61] primarily concentrated on the self-attention module, yet they fell short of constructing a comprehensive Transformer architecture, lacking basic components such as LayerNorm layer and positional encoding layer. This is primarily due to the inadequate definition of fundamental operations in previous studies.

Challenge (2): Inefficient and ineffective definitions for linear transformation in the hyperbolic Transformer. While some techniques [8, 24] employ the tangent space to achieve the linear transformation, they often necessitate frequent logarithmic and exponential mappings, heavily dependent on the tangent space at the origin. This leads to an increased computational load, accumulation of mapping errors, and unstable training procedures. Although Chen et al. [9] introduced a fully Lorentz linear transformation in hyperbolic space, it is constrained by its immutable curvature and normalization term.

Challenge (3): Absence of a linear attention mechanism in hyperbolic Transformer. The hyperbolic self-attention mechanisms proposed by Gulcehre et al. [26], Shimizu et al. [61], and Chen et al. [9] exhibit quadratic time complexity, posing a significant challenge when handling long-sequence input and large-scale graph data.

Proposed work: In this work, we propose an efficient hyperbolic Transformer, referred to as Hypformer. In particular, to address Challenges (1) and (2), we propose two foundational blocks, Hyperbolic Transformation with Curvatures (HTC) and Hyperbolic Readjustment and Refinement with Curvatures (HRC), to build all essential modules in the hyperbolic Transformer. HTC and HRC are built on the Lorentz model of hyperbolic geometry, working directly on the hyperbolic space without frequently mapping. HTC defines the linear transformation and facilitates mapping from a hyperbolic space with one curvature to another different curvature while preserving the relative distance. HRC further enables the definition of basic operations commonly used in the Transformer,

such as LayerNorm layer, activation function, dropout, and concatenation, within a hyperbolic context. To tackle Challenge (3), we introduce a self-attention mechanism in Hypformer with linear complexity, enabling efficient large-scale data processing.

To validate the effectiveness of the proposed methodology, we have undertaken extensive experiments across a diverse range of tasks. These include graph analysis [40, 42, 77, 85, 89], text classification [90, 91], and image classification [17, 72]. The empirical evidence gathered from these experiments indicates that the proposed method significantly reduces the GPU computation cost by a factor of 10 and concurrently halves the training time compared with the existing hyperbolic softmax attention. Furthermore, the proposed method consistently surpasses the performance of competitive baselines, yielding substantial improvements on both tree-like and non-tree-like datasets.

Contributions. In summary, this study offers the following contributions: *First*, we introduce two fundamental hyperbolic blocks, HTC and HRC. Building upon these, we have formulated fundamental modules for linear transformation, LayerNorm, activation function, dropout, and concatenation operations within a hyperbolic context. *Second*, we propose the first hyperbolic linear attention mechanism, which enables the hyperbolic Transformer to be scalable and efficient. Based on the above efforts, we construct a Hypformer¹, the first comprehensive and efficient hyperbolic Transformer model fully designed to operate within hyperbolic space. *Last*, we extend the hyperbolic model to handle billion-level graph data for the first time, laying a crucial foundation for the application of big data and large-scale models.

2 Related Work

2.1 Hyperbolic Neural Networks

Recent studies have demonstrated that hyperbolic space is particularly adept at capturing the hierarchical and tree-like structures [23, 25, 39, 41, 50, 51, 53, 63, 83]. Building on hyperbolic space, a variety of hyperbolic neural networks, HNN [24], HAN [26], HNN++ [61], HGNC [8], HGNN [43], F-HNN [9], Poincaré Resnet [64], HGTM [91] have been developed to leverage the advantages of the hyperbolic geometry. These neural networks have obtained an impressive performance in domains like computer vision [2, 29, 32], natural language processing [4, 7, 37, 47], recommender systems [11, 62, 68, 69, 76, 80], graph learning [3, 8, 41, 43, 78, 81, 92] and so on [41, 75].

2.2 Transformer and Hyperbolic Transformer

Introduced by Vaswani et al. [65], Transformer models have brought about a paradigm shift in the field of artificial intelligence. Transformer [5, 16, 18, 65] has made a tremendous impact in many fields, such as language understanding [5, 16, 54], image processing [6, 52] and graph learning [33, 55]. A well-known concern with self-attention is the quadratic time complexity, which can hinder model scalability in many settings. Efficient self-attention models are crucial in applications that model long sequences [27, 31, 36, 59].

Despite these advancements, existing Transformer architectures predominantly operate within the Euclidean domain. There have been limited attempts to extend these models to hyperbolic and

¹Code is available at <https://github.com/Graph-and-Geometric-Learning/hyperbolic-transformer>

other non-Euclidean spaces. Gulcehre et al. [26] proposed hyperbolic attention networks, which replace the dot-product between the *query* and *key* in self-attention with a function of negative hyperbolic distance. They then utilize the Einstein midpoint to compute the attentive output with *value*. Similarly, Chen et al. [9] and Shimizu et al. [61] adopt similar strategies that result in the attentive output with *key* being based on the Lorentzian midpoint and gyromidpoint, respectively.² However, these methods exhibit quadratic time complexity, limiting their scalability. Besides, they focused more on the self-attention module and did not define the essential modules, like LayerNorm in Transformer. Recently, Cho et al. [12] proposed a fully Product-Stereographic Transformer, presenting a kernelized approach to non-Euclidean attention, which is linear time complexity. However, this method heavily relies on the tangent space, necessitating frequent mappings between the tangent space and manifolds. Ermolov et al. [19] proposed mapping the last layer features obtained from a Euclidean Transformer to hyperbolic space, which essentially does not establish a true Hyperbolic Transformer. Our work aims to address these challenges and further the development of hyperbolic Transformers.

3 Preliminaries

In this section, we introduce concepts related to Lorentz model of hyperbolic geometry and self-attention module briefly.

3.1 Lorentz Model of Hyperbolic Geometry

There are several isometric models [23, 26, 50, 51, 56, 63] of hyperbolic geometry that have been employed in prior research. In this study, we choose the Lorentz model as the foundational framework due to the numerical stability it offers [46, 51]. Also, the proposed Hypformer can be easily adapted to other hyperbolic models, as they are isometrically equivalent.

Lorentz Model. An n -dimensional Lorentz model with negative constant curvature κ ($\kappa < 0$) is a Riemannian manifold denoted by $\mathbb{L}^{n,\kappa}$. The corresponding Riemannian metric is given by $g^\kappa = \text{diag}(1/k, 1, \dots, 1)$. Each point in $\mathbb{L}^{n,\kappa}$ can be represented as $\mathbf{x} = \begin{bmatrix} x_t \\ \mathbf{x}_s \end{bmatrix}$ where $\mathbf{x} \in \mathbb{R}^{n+1}$, $x_t \in \mathbb{R}$ and $\mathbf{x}_s \in \mathbb{R}^n$. The set of points, $\mathbb{L}^{n,\kappa}$, that constitute the manifold are defined as

$$\mathbb{L}^{n,\kappa} := \{ \mathbf{x} \in \mathbb{R}^{n+1} \mid \langle \mathbf{x}, \mathbf{x} \rangle_{\mathcal{L}} = 1/\kappa, x_t > 0 \}. \quad (1)$$

Here, $\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}} = -x_t y_t + \mathbf{x}_s^\top \mathbf{y}_s = \mathbf{x}^\top g^\kappa \mathbf{y}$ represents the Lorentzian inner product. Lorentz model, also known as the hyperboloid model, is an upper hyper-surface in an $(n+1)$ dimensional Minkowski space with the origin point $(\sqrt{-1/\kappa}, 0, \dots, 0)$. Lorentz model has its roots in the theory of special relativity [57] and employs terminology borrowed from this field. The hyperboloid's axis of symmetry, represented by the 0-th element x_t , is referred to as the time-like dimension, while all other axes \mathbf{x}_s are called space-like dimensions.

Tangent Space of Lorentz Model. Given $\mathbf{x} \in \mathbb{L}^{n,\kappa}$, the tangent space $\mathcal{T}_{\mathbf{x}} \mathbb{L}^{n,\kappa} := \{ \mathbf{u} \in \mathbb{R}^{n+1} \mid \langle \mathbf{u}, \mathbf{x} \rangle_{\mathcal{L}} = 0 \}$ is the orthogonal space of $\mathbb{L}^{n,\kappa}$ at \mathbf{x} with respect to the Lorentzian inner product.³ To

²In theory, the Einstein midpoint, Lorentzian centroid, and gyromidpoint are equivalent midpoint operations projected onto each manifold [61].

³The orthogonality condition $\langle \mathbf{u}, \mathbf{x} \rangle_{\mathcal{L}} = 0$ ensures that \mathbf{u} lies in the tangent space, preserving the manifold's geometry.

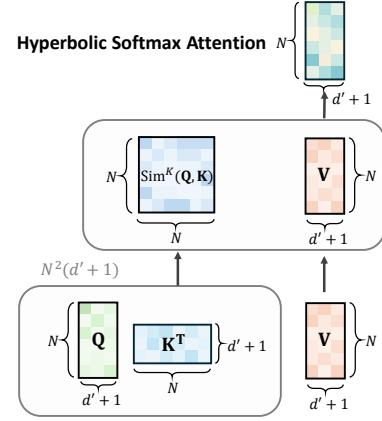


Figure 2: Illustration of hyperbolic softmax attention defined on Lorentz model. Unlike the attention mechanism in Euclidean space, this hyperbolic attention obtains the similarity between Q and K by a negative hyperbolic distance defined in Equation (5). The time complexity is quadratic w.r.t the number of input tokens.

achieve the mapping from the Lorentz model to the tangent space at \mathbf{x} , we can use the logarithmic map, $\log_{\mathbf{x}}^{\kappa} : \mathbb{L}^{n,\kappa} \rightarrow \mathcal{T}_{\mathbf{x}} \mathbb{L}^{n,\kappa}$. The exponential map defines the inverse process, $\exp_{\mathbf{x}}^{\kappa} : \mathcal{T}_{\mathbf{x}} \mathbb{L}^{n,\kappa} \rightarrow \mathbb{L}^{n,\kappa}$. For the details about exponential, logarithmic maps and the relevant distance functions, please refer to Appendix A.

3.2 Self-Attention Module

We first examine the general form of self-attention in Euclidean Transformers. Given the input of N tokens $\mathbf{X} \in \mathbb{R}^{N \times d}$, within each head, self-attention can be expressed as:

$$\begin{aligned} \mathbf{Q} &= \mathbf{X} \mathbf{W}^Q, \mathbf{K} = \mathbf{X} \mathbf{W}^K, \mathbf{V} = \mathbf{X} \mathbf{W}^V, \\ \mathbf{Z}_i &= \sum_{j=1}^N \frac{\text{Sim}(\mathbf{Q}_i, \mathbf{K}_j)}{\sum_{j=1}^N \text{Sim}(\mathbf{Q}_i, \mathbf{K}_j)} \mathbf{V}_j, \end{aligned} \quad (2)$$

where $\mathbf{W}^Q, \mathbf{W}^K, \mathbf{W}^V \in \mathbb{R}^{d \times d'}$ are projection matrices and $\text{Sim}(\cdot, \cdot)$ denotes the similarity function. Modern Euclidean Transformers primarily use Softmax attention [65] where similarity is calculated as $\text{Sim}(\mathbf{Q}_i, \mathbf{K}_j) = \exp(\mathbf{Q}_i \mathbf{K}_j^\top / \sqrt{d'})$. In this scenario, the attention map is derived by computing the similarity between all query-key pairs, which results in a computational complexity of $O(N^2)$.

The concept of hyperbolic self-attention, as defined by previous works [9, 26, 61], bears a similar idea to Equation (2). Figure 2 presents an illustration for this hyperbolic operation on Lorentz model. It can be expressed as follows:

$$\begin{aligned} \mathbf{Q} &= \mathbf{X} \otimes^{\kappa} \mathbf{W}^Q, \mathbf{K} = \mathbf{X} \otimes^{\kappa} \mathbf{W}^K, \mathbf{V} = \mathbf{X} \otimes^{\kappa} \mathbf{W}^V, \\ \mathbf{Z}_i &= \sum_{j=1}^N \frac{\text{Sim}^{\kappa}(\mathbf{Q}_i, \mathbf{K}_j)}{\sum_{j=1}^N \text{Sim}^{\kappa}(\mathbf{Q}_i, \mathbf{K}_j)} \odot^{\kappa} \mathbf{V}_j. \end{aligned} \quad (3)$$

In this equation, \otimes^{κ} denotes the hyperbolic linear transformation, which can be computed using Equations (6) and (7) given in the following section. The symbol \odot^{κ} represents the weighted sum

in hyperbolic space. Let Att_i denotes the i -th row of the attention matrix in the Lorentz model, it can be computed by Lorentzian midpoint [39]:

$$\text{Att}_i \odot^\kappa \mathbf{v}_j := \frac{\sum_{j=1}^N \alpha_{ij} \mathbf{v}_j}{\sqrt{|\kappa| \left\| \sum_{k=1}^N \alpha_{ik} \mathbf{v}_k \right\|_{\mathcal{L}}}}. \quad (4)$$

The function $\text{Sim}^\kappa(\cdot, \cdot)$ denotes the similarity function defined by the hyperbolic distance $d_{\mathcal{H}}$ [9, 26, 61]⁴ or the tangent inner product [44]. Specifically, Chen et al. [9] defined the similarity function as:

$$\text{Sim}^\kappa(\mathbf{Q}_i, \mathbf{K}_j) = \exp\left(-d_{\mathcal{H}}^2(\mathbf{Q}_i, \mathbf{K}_j)/\sqrt{d'}\right). \quad (5)$$

Both Gulcehre et al. [26] and Shimizu et al. [61] utilized similar forms of this function. They all employ negative distance to define similarity, and each has a computational complexity of $\mathcal{O}(N^2)$.

3.3 Lorentz Transformation

Lorentz Tangent Space Transformation. Previous works [8, 24, 43, 82, 93] mainly define hyperbolic linear transformations by the tangent space method, termed as $\text{LT}_{\mathcal{T}}$. Given the Lorentz embedding vector \mathbf{x} and operation function f , the tangent space method maps \mathbf{x} to the tangent space at a local reference point by the logarithmic map. Then, the transformation operation f is applied in this tangent space. Finally, the resulting vector is mapped back to the Lorentz model using the exponential mapping, that is⁵,

$$\text{LT}_{\mathcal{T}}(\mathbf{x}; f, \kappa_1, \kappa_2) := \exp_{\mathbf{o}}^{\kappa_1}(f(\log_{\mathbf{o}}^{\kappa_2}(\mathbf{x}))), \quad (6)$$

where \mathbf{o} is the local reference point (generally the origin point), and the curvatures κ_1 and κ_2 could be different since they share the same tangent space. Using this method, previous works define the linear transformation, neighbor's aggregation, dropout, and non-linear activation [8, 43].

Limitations. While this method is intuitive, it has notable limitations. *First*, parallel computation is feasible if the same reference point is used for the entire embedding. However, this approach can lead to significant mapping errors for distant points due to the point-specific nature of the tangent space. Conversely, using local-specific points enhances accuracy but increases computational load by requiring separate mappings. *Second*, frequent use of hyperbolic functions like \cosh or \cosh^{-1} can destabilize learning. While the `clamp` function can mitigate this issue, its use may compromise computational precision.

Fully Lorentz Transformation. To overcome the above limitations, Chen et al. [9] defined an alternative Lorentz transformation without using tangent space, termed as $\text{LT}_{\mathcal{F}}$:

$$\text{LT}_{\mathcal{F}}(\mathbf{x}; f, \mathbf{W}, \kappa) := \left(\sqrt{\|f(\mathbf{W}\mathbf{x}, \mathbf{v})\|^2 - 1/\kappa}, f(\mathbf{W}\mathbf{x}, \mathbf{v}) \right)^T, \quad (7)$$

which involves a function f that operates on vectors $\mathbf{v} \in \mathbb{R}^{n+1}$ and $\mathbf{W} \in \mathbb{R}^{m \times (n+1)}$, $f(\mathbf{W}\mathbf{x}, \mathbf{v}) = \frac{\lambda \sigma(\mathbf{v}^\top \mathbf{x} + \mathbf{b}')}{\|\mathbf{W}\mathbf{x} + \mathbf{b}\|} (\mathbf{W}\mathbf{x} + \mathbf{b})$. Here, σ is the sigmoid function, \mathbf{b} and \mathbf{b}' are bias terms, $\lambda > 0$ controls the

⁴Here we use subscript \mathcal{H} other than \mathcal{L} since it is not limited to Lorentz model.

⁵Some studies [82, 93] proposed an improved version of tangential linear transformations only on the space-like dimension and then incorporated a zero value to the transformed results, in order to respect the constraints of the tangent space at the origin. They have a similar formula as Equation (6), which we omit for brevity.

scaling range, and \mathbf{h} is the activation function. Depending on the type of function, it can perform different operations. For instance, for dropout, the operation function is $f(\mathbf{W}\mathbf{x}, \mathbf{v}) = \mathbf{W}$ dropout (\mathbf{x}) .

Limitations. There are several limitations to this method. *First*, the curvature is unchangeable. Although it appears that $\text{LT}_{\mathcal{F}}$ provides a way to directly modify κ in Equation (7), this modification results in a loss of previously learned information, introducing distortions. *Direct alteration of curvature cannot guarantee the preservation of relative distance relationships within the learned embedding.* The derivation is shown as follows:

Let $\mathbf{x}' = f(\mathbf{W}\mathbf{x}, \mathbf{v})$, and $g(\mathbf{x}') = \left(\sqrt{\|\mathbf{x}'\|^2 - 1/\kappa'}, \mathbf{x}' \right)$. Then:

$$\begin{aligned} d_{\mathcal{L}}^{\kappa'}(g(\mathbf{x}'), g(\mathbf{y}')) &= \sqrt{1/|\kappa'|} \operatorname{arcosh}(\kappa' \langle g(\mathbf{x}'), g(\mathbf{y}') \rangle_{\mathcal{L}}) \\ &= \sqrt{1/|\kappa'|} \operatorname{arcosh}\left(\kappa' \left(\alpha_{\text{time}} + \mathbf{x}'^T \mathbf{y}' \right)\right), \end{aligned} \quad (8)$$

where $\alpha_{\text{time}} = \sqrt{(\|\mathbf{x}'\|^2 - 1/\kappa') (\|\mathbf{y}'\|^2 - 1/\kappa')}$.

It can be observed that changing κ results in a non-linear transformation of the Lorentz distance $d_{\mathcal{L}}^{\kappa'}$. Consequently, the relative distances between data points may not be preserved as they were in the original κ Lorentz space. Even small changes in the parameter κ can significantly affect the resulting distances, potentially distorting the previously learned hierarchical structure.

Second, the requirement for the \mathbf{W} matrix and normalization term pose another challenge. In [9], \mathbf{W} is applied to both time-like and space-like dimensions, in order to achieve Lorentz boosts and rotations simultaneously. However, its introduction constrains the usage of certain functions. For instance, dropout, activation operation do not necessarily interact with the matrix \mathbf{W} . Taking the ReLU activation function as an example, it only requires filtering out negative values without needing matrix multiplication in Euclidean space. Additionally, Chen et al. [9] introduced a normalization term that constrains the value within a limited range, thereby limiting the expressiveness of the transformation.

Lastly, some basic operations, such as LayerNorm and Concatenation, cannot be achieved within this definition.

4 Method

The proposed method is designed to overcome the limitations of the existing attempts in hyperbolic Transformer, as outlined in the Section 1. To address Challenges (1) and (2), we designed two foundational blocks, namely HTC and HRC in Section 4.1, and 4.2, respectively. To overcome Challenge (3), we developed a hyperbolic linear attention module in Section 4.3, which equips the Transformer with linear time complexity.

4.1 Hyperbolic Transformation with Curvatures (HTC)

Novelty. Unlike the design of linear transformation using the tangent space method in Equation (6), we build the transformation fully in hyperbolic space. Besides, compared with Lorentz transformation defined by Equation (7), we have two improvements:



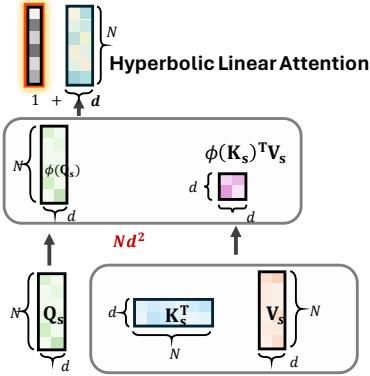


Figure 3: Illustration of Hyperbolic Linear attention. This attention operates in the space-like dimension (Q_s, K_s, V_s) and reduces the time complexity by changing the computation order.

(1) making the curvature changeable with preserving the relative ordering; (2) being disentangled with normalization term.

Given a point \mathbf{x} in Lorentz model, $\mathbf{x} \in \mathbb{L}^{d, \kappa_1}$ (implies $\mathbf{x} \in \mathbb{R}^{d+1}$), and transformation matrix $\mathbf{W} \in \mathbb{R}^{(d+1) \times d'}$ and bias $b \in \mathbb{R}^{d'}$, the **HTC** is given as the following equation:

$$\text{HTC}(\mathbf{x}; f_t, \mathbf{W}, \kappa_1, \kappa_2) := \begin{pmatrix} \sqrt{\frac{\kappa_1}{\kappa_2} \|f_t(\mathbf{x}; \mathbf{W})\|_2^2 - 1/\kappa_2} & \sqrt{\frac{\kappa_1}{\kappa_2}} f_t(\mathbf{x}; \mathbf{W}) \\ \text{time-like dimension} & \text{space-like dimension} \end{pmatrix}^T, \quad (9)$$

where the $f_t(\mathbf{x}; \mathbf{W}) = \mathbf{W}^T \mathbf{x} + b$ denotes the linear transformation with bias addition and κ_1, κ_2 represent the curvatures before and after the transformation. Note that **HTC** does not entangle a normalization term with this linear transformation. Besides, It is easy to prove that the defined transformation also satisfies the **Lorentz rotation and boost operations**, described in [9]. The proof is similar in [9], we omit for brevity.

The proposed HTC avoids the use of tangent space and minimizes the usage of logarithmic and exponential mappings in comparison to Equation (6). When contrasted with Equation (7), the variable curvature of the HTC enhances the flexibility of the transformation. This is because linear transformations generally alter the feature dimension, and varying curvatures can express more than a fixed one.

Next, we study the theoretical aspects of the proposed HTC. First and foremost, we prove that HTC is closed in hyperbolic space in Proposition (4.1) with different dimensions and curvatures so that the mapping is done correctly. Next, in Proposition (4.2), we show that the curvature-changing strategy of the proposed HTC, along with the subsequent HRC, maintains the relative distance among any points between pre and post-curvature changing.

PROPOSITION 4.1. Let $\mathbf{x} \in \mathbb{L}^{d_a, \kappa_a}$ and $\mathbf{W} \in \mathbb{R}^{(d_a+1) \times d_b}$. The LTC operation, defined as $\text{LTC}(\mathbf{x}; f_t; \mathbf{W}, \kappa_a, \kappa_b)$, correctly transforms \mathbf{x} from the Lorentz model with curvature κ_a to the Lorentz model with curvature κ_b , such that

$$\text{LTC}(\mathbf{x}; f_t; \mathbf{W}, \kappa_a, \kappa_b) \in \mathbb{L}^{d_b, \kappa_b}. \quad (10)$$

PROPOSITION 4.2. Let $\mathbf{z}_i, \mathbf{z}_j, \mathbf{z}_k \in \mathbb{L}^{\kappa_a}$ be points in the Lorentz model with curvature κ_a . Consider the curvature changing transformations defined in HTC (Equation (7)) and HRC (Equation 13). Let $\mathbf{z}'_i, \mathbf{z}'_j, \mathbf{z}'_k \in \mathbb{L}^{\kappa_b}$ denote the transformed points in the Lorentz model with curvature κ_b . The relative distances within $(\mathbf{z}_i, \mathbf{z}_j, \mathbf{z}_k)$ are preserved after the curvature alteration. Specifically, if

$$d_{\mathcal{L}}^{\kappa_a}(\mathbf{z}_i, \mathbf{z}_j) \geq d_{\mathcal{L}}^{\kappa_a}(\mathbf{z}_i, \mathbf{z}_k), \quad (11)$$

then

$$d_{\mathcal{L}}^{\kappa_b}(\mathbf{z}'_i, \mathbf{z}'_j) \geq d_{\mathcal{L}}^{\kappa_b}(\mathbf{z}'_i, \mathbf{z}'_k). \quad (12)$$

4.2 Hyperbolic Readjustment and Refinement with Curvatures (HRC)

Novelty. Within the Transformer, we have several basic operations beyond linear transformation, which include *Dropout* and *Concatenation*, *Activation function* (e.g., ReLU), and *LayerNorm*. We interpret these operations within the hyperbolic space as a readjustment or refinement process, referred to as HRC. Similarly, given a point \mathbf{x} in Lorentz model, the proposed operation HRC is defined as:

$$\text{HRC}(\mathbf{x}; f_r, \kappa_1, \kappa_2) := \begin{pmatrix} \sqrt{\frac{\kappa_1}{\kappa_2} \|f_r(\mathbf{x}_{[1:]})\|_2^2 - 1/\kappa_2} & \sqrt{\frac{\kappa_1}{\kappa_2}} f_r(\mathbf{x}_{[1:]}) \\ \text{time-like dimension} & \text{space-like dimension} \end{pmatrix}^T. \quad (13)$$

Here, f_r represents a function applied to the space-like dimensions.

It is evident that HRC shares similar advantages with HTC, which we will not repeat for the sake of brevity. However, unlike HTC, HRC performs the transformation only in space-like dimensions. The primary motivation is as follows: HTC involves a Lorentz boost, essential for mapping between **different** inertial reference frames, tied to causality and affecting the observed sequence of events in relativity. However, operations such as LayerNorm, activation functions, dropout, and concatenation serve as readjustments or refinements within the **same** frame of reference, acting on space-like features to standardize, activate, or regularize them. Applying these to the space-like dimension ensures the causal structure remains intact. In practical, it ensures dimensional consistency, improves interpretability, and allows for more efficient computation. Nonetheless, it is important to note that HRC does not completely discard the time-like information. According to the definition in Equation (1), the time-like dimension is determined by the space-like dimensions. By operating on the space-like dimensions, HRC implicitly utilizes the time-like information.

4.3 Hyperbolic Linear Attention

In hyperbolic space, the traditional way of calculating self-attention is quadratic time complexity, which hinders scalability. Therefore, we defined a linear attention through HTC and HRC modules.

Specifically, given the N input token feature with dimension d , $\mathbf{X} \in \mathbb{L}^{N \times d, \kappa_1}$ in the Lorentz model with transformation matrix

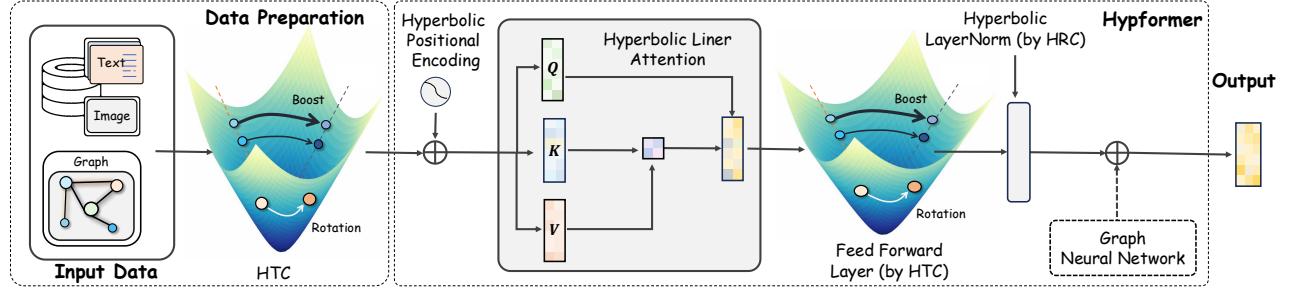


Figure 4: Framework of Hypformer. Input data (text, images, graphs) are projected onto the Lorentz model, then transformed via Hyperbolic Linear Transformation (HTC). The result passes through the hyperbolic linear attention block with positional encoding, followed by a Feedforward layer (built by HTC) and LayerNorm (built by HRC). This serves as an encoder which can optionally incorporate a GNN. For classification tasks in this study, the decoder is the fully connected layer. Dropout, activation, and residual connections are omitted for brevity.

$\mathbf{W}^Q, \mathbf{W}^K, \mathbf{W}^V \in \mathbb{R}^{(d+1) \times d'}$, we first transform it to Q, \mathcal{K} and \mathcal{V} , that is

$$\begin{aligned} Q &= \text{HTC}(\mathbf{X}; f_t, \mathbf{W}^Q, \kappa_1, \kappa_2), \\ \mathcal{K} &= \text{HTC}(\mathbf{X}; f_t, \mathbf{W}^K, \kappa_1, \kappa_2), \\ \mathcal{V} &= \text{HTC}(\mathbf{X}; f_t, \mathbf{W}^V, \kappa_1, \kappa_2), \end{aligned} \quad (14)$$

where Q, \mathcal{K} and $\mathcal{V} \in \mathbb{L}^{N \times d', \kappa_2}$. Given that the subsequent pairwise similarity computation and aggregation essentially constitute a weighted sum, and their calculation does not involve transformations on the time-like dimension, we adopted the idea of HRC to achieve this. Specifically, we first slice the values of the space-like dimension,

$$\mathcal{Q}_s, \mathcal{K}_s, \mathcal{V}_s = \phi(Q_{[1:]}) \cdot \phi(\mathcal{K}_{[1:]}) \cdot \phi(\mathcal{V}_{[1:]}) \quad (15)$$

To achieve linearity, we alter the computation sequence, i.e., transitioning from $(Q^T \mathcal{K}) \mathcal{V}$ to $Q(\mathcal{K}^T \mathcal{V})$, inspired by [27]. Our innovation lies in defining this operation in space-like dimensions and recalibrating the time-like value to respect the Lorentz constraint,

$$\mathcal{Z}_s = \frac{\mathcal{Q}_s (\mathcal{K}_s^T \mathcal{V}_s)}{\mathcal{Q}_s (\mathcal{K}_s^T \mathbf{1})} \quad (16)$$

Before recalibrating, we incorporate the following residual connection:

$$\tilde{\mathcal{Z}}_s = \mathcal{Z}_s + \psi(\mathcal{V}_s), \quad (17)$$

and then do the time-like calibration and concatenation,

$$\begin{aligned} \mathcal{Z}_t &= \sqrt{\frac{\kappa_2}{\kappa_3}} \|\tilde{\mathcal{Z}}_s\|^2 - 1/\kappa_3, & (\text{Time-like Calibration}) \\ \mathcal{Z} &= \left(\mathcal{Z}_t, \sqrt{\frac{\kappa_2}{\kappa_3}} \tilde{\mathcal{Z}}_s \right). & (\text{Re-concatenation}) \end{aligned} \quad (18)$$

In Equation (15, 16, 17), $\mathbf{1}$ denotes an all "1" vector, ψ is a linear layer and ϕ signifies the functions employed to enhance the focus of the linear attention, i.e.,

$$\phi(\tilde{\mathbf{e}}) = \frac{\|\tilde{\mathbf{e}}\|}{\|\tilde{\mathbf{e}}^p\|} \tilde{\mathbf{e}}^p, \text{ where } \tilde{\mathbf{e}}_i = \text{ReLU}(\mathbf{e}_i)/t, \quad (19)$$

where $\mathbf{e} \in \mathbb{L}^{d', \kappa_2}$ represents the transpose of row in Q, \mathcal{K} and \mathcal{V} . In this case, t represents a scaling factor, which we set as a trainable parameter in the experiments. The focused strategy is inspired

by the work in [27]. A $p > 1$ sharpens the paired points, i.e., it enhances the similarity within each group while diminishing the similarity between the groups. Conversely, a $p < 1$ has the opposite effect.

This linear attention approach allows us to handle large datasets and long sequences more efficiently while respecting the properties of the Lorentz model.

4.4 Hyperbolic Positional Encoding

Positional encoding in a Transformer model is instrumental in preserving the sequence of input tokens. In what follows, we introduce a relative positional encoding with a trainable model inspired by [39, 65].

$$\tilde{\mathbf{x}} = \frac{\mathbf{x} + \epsilon \cdot \mathbf{p}}{\sqrt{\kappa \|\mathbf{x} + \epsilon \cdot \mathbf{p}\| \|\mathbf{L}\|}}. \quad (20)$$

Here, $\mathbf{p} := \text{HTC}(\mathbf{x})$ functions as a Lorentz position vector, and ϵ specifies the magnitude of \mathbf{p} and we use 1 in our experiments. This definition calculates the midpoint between \mathbf{x} and $\epsilon \cdot \mathbf{p}$, with respect to the Lorentz constraint. We add the positional encoding before the linear transformation in the self-attention block. We reserve the exploration of more advanced positional encoding for future works.

4.5 Hyperbolic LayerNorm, Dropout, Activation, and Concatenation

LayerNorm, Dropout, Activation, and Concatenation are fundamental components of the Transformer architecture. For these operations, we employ HRC in our definitions. This choice is motivated by the fact that these functions are performed within the same reference system and do not involve a time-like dimension. Consequently, we define our operations as follows⁶:

$$\begin{aligned} \text{HypLayerNorm}(\mathbf{X}) &= \text{HRC}(\mathbf{X}, f_{\text{LayerNorm}}), \\ \text{HypBatchNorm}(\mathbf{X}) &= \text{HRC}(\mathbf{X}, f_{\text{BatchNorm}}), \\ \text{HypDropout}(\mathbf{X}) &= \text{HRC}(\mathbf{X}, f_{\text{Dropout}}), \\ \text{HypActivation}(\mathbf{X}) &= \text{HRC}(\mathbf{X}, f_{\text{Activation}}), \\ \text{HypConcatenation}(\mathbf{X}) &= \text{HRC}((\mathbf{X}_i, \mathbf{X}_j), f_{\text{concatenation}}), \end{aligned} \quad (21)$$

⁶We also include BatchNorm for reference.

Table 1: Testing results (ROC-AUC for ogbn-proteins and Accuracy for other datasets) on large-scale node property prediction benchmarks. OOM denotes out of memory during training or testing, and OOT indicates the model could not complete within the allocated time budget. The best and second-best results are highlighted in red bold and underlined, respectively.

Method	ogbn-proteins	Amazon2m	ogbn-arxiv	Papers100M
#Nodes	132,534	2,449,029	169,343	111,059,956
#Edges	39,561,252	61,859,140	1,166,243	1,615,685,872
MLP	72.0 ± 0.5	63.5 ± 0.1	55.5 ± 0.2	47.2 ± 0.3
GCN [34]	72.5 ± 0.4	83.9 ± 0.1	71.7 ± 0.3	OOM
SGC [71]	70.3 ± 0.2	81.2 ± 0.1	67.8 ± 0.3	63.3 ± 0.2
GCN-NSampler	73.5 ± 1.3	83.8 ± 0.4	68.5 ± 0.2	62.0 ± 0.3
GAT-NSampler	74.6 ± 1.2	85.2 ± 0.3	67.6 ± 0.2	63.5 ± 0.4
SIGN [22]	71.2 ± 0.5	81.0 ± 0.3	70.3 ± 0.3	65.1 ± 0.1
GraphFormer [84]	OOM	OOM	OOM	OOM
GraphTrans [74]	OOM	OOM	OOM	OOM
GraphGPS [55]	OOM	OOM	OOM	OOM
HAN [26]	OOM	OOM	OOM	OOM
HNN++ [61]	OOM	OOM	OOM	OOM
F-HNN [9]	OOM	OOM	OOM	OOM
NodeFormer [72]	77.5 ± 1.2	87.9 ± 0.2	59.9 ± 0.4	OOT
SGFormer [73]	79.5 ± 0.3	89.1 ± 0.1	72.4 ± 0.3	65.8 ± 0.5
Hypformer	80.4 ± 0.5	89.4 ± 0.3	73.2 ± 0.2	66.1 ± 0.4

where $f_{Dropout}$, $f_{LayerNorm}$, $f_{BatchNorm}$, and f_{σ} as well as $f_{concatenation}$ represent traditional Euclidean *Dropout*, *LayerNorm*, and *Activation* functions, respectively. In general, we define κ as unchanged before and after the HRC. In the actual implementation process, for two operations that appear consecutively, such as $f_1 = f_{Dropout}$ and $f_2 = f_{ReLU}$, we merge them into $f = f_1 \circ f_2$ for computational efficiency.

4.6 Overall Architecture

The framework of Hypformer is shown in Figure 4, it can accept a variety of data types, such as text, images, and graphs. During the data preparation phase, the input data is mapped to the Lorentz model using an exponential map⁷. This mapped embedding is then transformed using a HTC layer. In the encoder part of Hypformer, the transformed data is processed through a hyperbolic linear attention block with hyperbolic position encoding. This is followed by the Feedforward layer implemented by HTC, and LayerNorm layer built by HRC. For graph-based inputs, we incorporate the graph neural networks and adopt the parallel paradigm [45] for Transformer and GNN encoder to form a graph Transformer model. The processed data is then forwarded to the decoder. The decoder can either be the similar structure of encoder, hyperbolic multinomial logistic regression (HypMLR) [24, 61] or a tailored design, we leave it in future exploration. In this research, the decoder is a fully connected layer used for classification tasks.

Time complexity. In the proposed Hypformer, the linear attention module is the main computational bottleneck. The complexity comes from two key operations. In Equation (16), we perform a space-like inner product computation of \mathcal{K}^T and \mathcal{V} within the Lorentz model, which incurs a complexity of $O(d'^2N)$. Following this, we calculate the inner product of these results with Q , which also has a complexity of $O(d'^2N)$. Given that $d' \ll N$, the total

⁷This step is necessary since most data are built from Euclidean space.

computational complexity of our method is $O(N)$. When dealing with graph inputs, the computational complexity of a GNN model is typical $O(N+E)$, where E represents the number of edges. Owing to the typical sparsity of graphs (i.e., $E \ll N^2$), the proposed method can scale linearly with respect to the number of nodes in a graph. This design make Hypformer operate on graphs with billion-level nodes.

5 Experiments

In this work, we propose a novel hyperbolic Transformer with linear complexity, which is especially well-suited for processing graph-structured data. Graphs often exhibit intricate topological and hierarchical relationships, making them an ideal testbed for evaluating the effectiveness of our proposed hyperbolic Transformer. As such, we primarily focus on comparing our model’s performance with other state-of-the-art graph models.

5.1 Experiments on Large Graphs

Experimental Settings. We first evaluate Hypformer on diverse large-scale graphs for node classification, with node counts ranging from millions to billions, including ogbn-arxiv, ogbn-protein, and Papers100M (for dataset details, see Appendix C.1). To our knowledge, this represents the first application of hyperbolic or non-Euclidean transformations to graphs of this scale. Our comparative analysis focuses on state-of-the-art Euclidean GNNs and graph Transformers. We evaluate Hypformer against a spectrum of baselines, including MLP, GCN [35], SGC [71], advanced GNN variants (SIGN [22], GCN-NSampler, GAT-NSampler), recent graph Transformer architectures (GraphFormer [84], GraphTrans [74], GraphGPS [55], NodeFormer [72], SGFormer [73]) and hyperbolic models HAN [26], HNN++ [61] and F-HNN [9].

Experimental Findings. Table 1 summarizes the results of our experiments. Hypformer consistently outperforms other models across various large-scale graph datasets, demonstrating substantial improvements. It is worth noting that models, such as GraphFormer [84], GraphTrans [74], and GraphGPS [55], HAN [26], HNN++ [61] and F-HNN [9], have difficulty operating effectively on large-scale graph data. In addition, our method significantly outperforms the recent approaches such as, SGFormer and NodeFormer accross all tested scenarios, highlighting its superior effectiveness. Importantly, Hypformer exhibits robust scalability, maintaining its performance advantage even on the largest dataset, ogbn-papers100M, where previous Transformer-based models have encountered limitations.

5.2 Experiments on Small/Medium Graphs

To complement our large-scale evaluations, we assessed Hypformer on small- and medium-scale graph datasets. This additional testing allows for a more comprehensive comparison against current state-of-the-art models, including GNNs, graph transformers, and hyperbolic approaches that may not scale effectively to larger datasets. By expanding our evaluation scope, we aim to isolate Hypformer’s effectiveness in graph learning from its scalability advantages.

Experimental Settings. We conducted experiments on five small/medium-scale graph datasets, adhering closely to the settings used in HGCRN works [8]. These datasets included three low-degree

Table 2: Testing results (F1-score for DISEASE and Accuracy for other datasets) on small and medium-sized graph benchmarks. The best and second-best results are highlighted in red bold and underlined, respectively.

Models	DISEASE	AIRPORT	CORA	CITESEER	PUBMED
#Nodes	1,044	2,665	2,708	3,327	19,717
#Edges	1,043	2,664	5,429	4,732	88,651
GCN [35]	69.7 ± 0.4	81.4 ± 0.6	81.3 ± 0.3	71.6 ± 0.4	78.1 ± 0.2
GAT [66]	70.4 ± 0.4	81.5 ± 0.3	83.0 ± 0.7	72.5 ± 1.1	79.0 ± 0.3
SGC [71]	69.1 ± 0.6	82.1 ± 0.5	80.1 ± 0.2	71.9 ± 0.1	78.7 ± 0.1
HGNN [43]	81.3 ± 3.5	84.7 ± 1.0	77.1 ± 0.8	70.0 ± 1.0	78.3 ± 1.2
HGCN [8]	88.2 ± 0.7	89.3 ± 1.2	76.5 ± 0.6	68.0 ± 0.6	78.0 ± 1.0
HGAT [8]	90.3 ± 0.6	89.6 ± 1.0	77.4 ± 0.7	68.6 ± 0.3	78.3 ± 1.4
GraphFormer [84]	75.2 ± 0.0	88.1 ± 1.2	60.0 ± 0.5	61.4 ± 0.6	73.3 ± 0.7
GraphTrans [74]	89.3 ± 3.2	94.3 ± 0.6	77.6 ± 0.8	65.1 ± 1.4	77.5 ± 0.7
GraphGPS [55]	92.8 ± 2.7	94.5 ± 0.9	73.0 ± 1.4	62.0 ± 1.5	72.8 ± 1.4
FPS-T [12]	88.6 ± 0.9	96.0 ± 0.6	82.3 ± 0.7	70.0 ± 0.7	78.5 ± 0.6
HAN [26]	85.1 ± 0.8	92.9 ± 0.6	83.1 ± 0.5	72.4 ± 0.5	79.0 ± 0.6
HNN++ [61]	89.5 ± 0.2	92.3 ± 0.3	82.8 ± 0.6	71.5 ± 1.3	79.9 ± 0.4
F-HNN [9]	92.3 ± 1.1	93.0 ± 0.7	81.0 ± 0.7	71.2 ± 0.4	77.5 ± 0.8
NodeFormer [72]	75.9 ± 0.9	80.2 ± 0.6	82.2 ± 0.9	72.5 ± 1.1	79.9 ± 1.0
SGFormer [73]	89.0 ± 3.9	92.9 ± 0.5	83.2 ± 0.9	72.2 ± 0.3	80.0 ± 0.8
Hypformer	93.0 ± 0.7	95.0 ± 0.5	85.0 ± 0.3	73.3 ± 0.4	81.3 ± 0.3

hyperbolicity datasets: CITESEER, CORA [60], and PUBMED [48], as well as two high-degree hyperbolicity datasets: AIRPORT and DISEASE. The number of nodes and edges are shown in Table 2. For data split and processing, please refer to Appendix C.2.

Experimental Findings. Table 2 showcases all the experimental results⁸. Our findings suggest that the proposed method significantly surpasses both standard GNNs and hyperbolic GNN models by a substantial margin. Importantly, the method exhibits effectiveness not only in scenarios with hyperbolic datasets (like DISEASE, AIRPORT) but also in situations with non-hyperbolic dataset (like CORA, CITESEER and PUBMED). The existing hyperbolic GNN model [8] had a notable deficiency in this non-hyperbolic datasets. However, by introducing a hyperbolic Transformer, we have successfully overcome this problem. This thanks to that Transformers possess long-distance learning capabilities. However, on datasets such as CORA, CITESEER, and PUBMED, the existing graph Transformers cannot perform well. The primary reason might be that the Transformer equals a fully linked aggregation, which will introduce substantial noise. Nevertheless, our method employs linear-focused attention to solve this issue effectively.

5.3 Comparisons on Text and Vision Datasets

Additionally, we apply our model to semi-supervised image and text classification tasks on the Mini-ImageNet and 20News-Groups datasets. We also construct a graph using k-NN (based on input node features) to utilize graph model. These experiments are conducted closely in NodeFormer. More comprehensive details are provided in Appendix C.3. Table 3 presents the comparative results for varying k values. Notably, our method outperforms in seven out of eight cases. In contrast, the performance of competing baselines models vary significantly with different k values, while our method demonstrates greater stability.

⁸Missing values indicate that there were no previous experiments conducted and the results could not be reproduced.

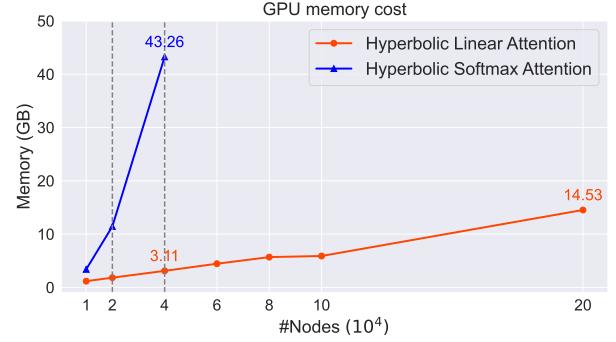


Figure 5: Scalability test on an A100 with the proposed linear attention and softmax attention of training time per epoch and GPU memory usage w.r.t. the number of input tokens/nodes.

6 Analysis

Scalability of Hypformer. We conducted additional tests on the model's scalability regarding the number of nodes in a single batch. The Amazon2M dataset was used, and we randomly selected a subset of nodes, with the number of nodes varying from 10K to 200K. We made a comparison between softmax attention defined by Equation (3) and linear attention defined by Equation (16), keeping all other parameters the same. As depicted in Figure 5, the memory usage of the proposed method exhibits a linear increase with the size of the graph. When the node count exceeds 40K, the softmax attention experiences an out-of-memory (OOM) issue. However, the proposed method continues to function effectively, resulting in a 10X reduction in GPU cost.

Efficiency and Effectiveness of Hypformer. The linear attention designed for Hypformer enhances its efficiency significantly. Table 4 presents the efficiency of both softmax attention and linear attention within Hypformer.⁹ As indicated in Table 4, the proposed linear attention mechanism significantly reduces the training time by half compared to the softmax attention in Hypformer. Furthermore, The left subfigure in Figure 6 presents the performance comparison between Hypformer equipped with Softmax attention (Hypformer(S)) and Linear attention (Hypformer(L)). The results demonstrate that both models perform well, with the linear attention exhibiting better accuracy.

Effectiveness of Curvature κ . In this work, we propose that both the HTC and HRC basic blocks involve two variable curvatures. In our experiment, we set these as trainable parameters. In the right Figure 6, we compare the impact of varying κ and fixed curvature on the Hypformer. Experiments show that varying κ can always perform better than the unified one.

Ablation Study. To gain a deeper understanding of the proposed Hyperbolic Transformer's effectiveness, we conducted an ablation study on three diverse datasets. We compared the performance of the original Hyperbolic Transformer with two variants: one without the graph component (W/o Graph) and another without

⁹To ensure a fair comparison, we have maintained the batch size at a smaller value (40K for Amazon2M and 10K for others) across all tests.

Table 3: Experimental results on semi-supervised classification on Mini-ImageNet and 20News-Groups where we use k-NN (with different k’s) for artificially constructing an input graph. The best and second-best results are highlighted in red bold and underlined, respectively, where models with and without graph are compared separately.

Method	Mini-ImageNet				20News-Group			
	$k = 5$	$k = 10$	$k = 15$	$k = 20$	$k = 5$	$k = 10$	$k = 15$	$k = 20$
GCN [35]	84.86 ± 0.42	85.61 ± 0.40	85.93 ± 0.59	85.96 ± 0.66	65.98 ± 0.68	64.13 ± 0.88	62.95 ± 0.70	62.59 ± 0.62
GAT [67]	84.70 ± 0.48	85.24 ± 0.42	85.41 ± 0.43	85.37 ± 0.51	64.06 ± 0.44	62.51 ± 0.71	61.38 ± 0.88	60.80 ± 0.59
DropEdge [58]	83.91 ± 0.24	85.35 ± 0.44	85.25 ± 0.63	85.81 ± 0.65	64.46 ± 0.43	64.01 ± 0.42	62.46 ± 0.51	62.68 ± 0.71
IDGL [10]	83.63 ± 0.32	84.41 ± 0.35	85.50 ± 0.24	85.66 ± 0.42	65.09 ± 1.23	63.41 ± 1.26	61.57 ± 0.52	62.21 ± 0.79
LDS [21]	OOM	OOM	OOM	OOM	66.15 ± 0.36	64.70 ± 1.07	63.51 ± 0.64	63.51 ± 1.75
NodeFormer [72]	86.77 ± 0.45	86.74 ± 0.23	86.87 ± 0.41	86.64 ± 0.42	66.01 ± 1.18	65.21 ± 1.14	64.69 ± 1.31	64.55 ± 0.97
SGFormer	86.21 ± 0.66	86.46 ± 0.61	86.73 ± 0.84	86.76 ± 0.72	68.55 ± 0.54	67.96 ± 0.68	66.44 ± 0.87	65.46 ± 0.59
NodeFormer w/o graph [73]			87.46 ± 0.36				64.71 ± 1.33	
SGFormer w/o graph [73]			87.25 ± 0.38				67.53 ± 0.43	
Hypformer	87.36 ± 0.73	87.30 ± 0.65	87.41 ± 0.59	87.48 ± 0.61	68.21 ± 0.78	68.01 ± 0.34	66.87 ± 0.30	66.74 ± 0.19
Hypformer w/o graph			87.73 ± 0.63				67.73 ± 0.23	

Table 4: Efficiency comparison by running time (ms) per epoch between the softmax full and the proposed linear attention in Hypformer on an A100 GPU.

Method	ogbn-proteins		Amazon2M		ogbn-arxiv	
	Train (ms)	Test (ms)	Train (ms)	Test (ms)	Train (ms)	Test (ms)
Hypformer (Softmax)	11.9	OOM	37.38	OOM	7.8	OOM
Hypformer (Linear)	5.3	2.4	16.32	2.5	3	2.5

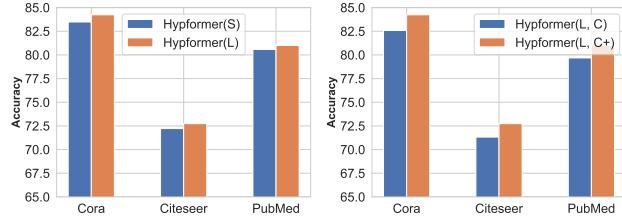


Figure 6: Left: Comparison of the proposed linear attention and softmax attention on small/medium datasets. Right: Comparison of unified curvature (Hypformer(L, C)) and varying curvature (Hypformer(L, C+))

the Transformer component (W/o Transformer). The results of this study are presented in Table 5.

For the CORA dataset, a citation network, removing the graph component leads to a substantial performance drop. This indicates the crucial role of the graph structure in capturing the relationships between nodes in this context. The Transformer component alone (W/o Graph) is insufficient for effectively modeling node interactions. Conversely, removing the Transformer component (W/o Transformer) still yields reasonable performance, highlighting the importance of the graph component for this dataset. In the case of the ogbn-proteins dataset, which represents a protein-protein interaction network, both the graph and Transformer components contribute significantly to the model’s performance. This suggests that the interplay between the graph structure and the Transformer’s ability to capture long-range dependencies is essential for accurately modeling the complex interactions in this biological network. For the 20news dataset, which comprises textual data, the graph is constructed from the original features and may not accurately reflect the true relationships between documents. In this case, the

Table 5: Ablation study

Dataset	W/o Graph	W/o Transformer	Hypformer
Cora	64.6 ± 0.5	82.8 ± 0.3	85.0 ± 0.3
ogbn-proteins	70.0 ± 0.2	75.9 ± 0.4	80.4 ± 0.5
Mini-ImageNet	87.7 ± 0.6	85.8 ± 0.5	87.4 ± 0.7

model performs best when the graph component is removed (W/o Graph), indicating that the graph structure might not be as informative for this particular dataset. The Hyperbolic Transformer component alone is sufficient to capture the semantic relationships between documents. These findings underscore the adaptability of the Hyperbolic Transformer to various datasets and its ability to leverage both graph structure and long-range dependencies when appropriate.

7 Conclusion

In this work, we introduce a efficient hyperbolic Transformer, Hypformer. This method operates directly and fully on hyperbolic representations and employs a linear attention mechanism, enabling it to be both scalable and effective. Furthermore, this study introduces two basic blocks, HTC and HRC, which are foundational in constructing hyperbolic models. Nonetheless, the research presented is an initial exploration and numerous challenges warrant further investigation. These include the initial determination of a curvature that better reflects the data geometry, the setting of curvature at different levels for Hypformer, and the design of effective decoders for different downstream tasks. We plan to address these issues in our future work.

Acknowledgements

We express gratitude to the anonymous reviewers and area chairs for their valuable comments and suggestions. In this study, Menglin Yang was partly supported by Tony Massini Postdoctoral Fellowship in Data Science from Yale University. Jiahong Liu and Irwin King were partly supported by the Research Grants Council of the Hong Kong Special Administrative Region, China (CUHK14222922, RGC GRF 2151185).

References

- [1] Xavier Amatriain. 2023. Transformer models: an introduction and catalog. *arXiv preprint arXiv:2302.07730* (2023).
- [2] Mina Ghadimi Atigh, Julian Schoep, Erman Acar, Nanne van Noord, and Pascal Mettes. 2022. Hyperbolic Image Segmentation. In *CVPR*. 4453–4462.
- [3] Qijie Bai, Changli Nie, Haiwei Zhang, Dongming Zhao, and Xiaojie Yuan. 2023. Hgwavenet: A hyperbolic graph neural network for temporal link prediction. In *WWW*. 523–532.
- [4] Yushi Bai, Zhitao Ying, Hongyu Ren, and Jure Leskovec. 2021. Modeling heterogeneous hierarchies with relation-specific hyperbolic cones. In *NeurIPS*, Vol. 34. 12316–12327.
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. In *NeurIPS*, Vol. 33. 1877–1901.
- [6] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *ECCV*. Springer, 213–229.
- [7] Ines Chami, Adva Wolf, Da-Cheng Juan, Frederic Sala, Sujith Ravi, and Christopher Ré. 2020. Low-Dimensional Hyperbolic Knowledge Graph Embeddings. In *ACL*. 6901–6914.
- [8] Ines Chami, Zhitao Ying, Christopher Ré, and Jure Leskovec. 2019. Hyperbolic graph convolutional neural networks. In *NeurIPS*, Vol. 32. 4868–4879.
- [9] Weize Chen, Xu Han, Yankai Lin, Hexu Zhao, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. 2022. Fully Hyperbolic Neural Networks. In *ACL*. 5672–5686.
- [10] Yu Chen, Lingfei Wu, and Mohammed Zaki. 2020. Iterative deep graph learning for graph neural networks: Better and robust node embeddings. In *NeurIPS*, Vol. 33. 19314–19326.
- [11] Yankai Chen, Menglin Yang, Yingxue Zhang, Mengchen Zhao, Ziqiao Meng, Jianye Hao, and Irwin King. 2022. Modeling scale-free graphs with hyperbolic geometry for knowledge-aware recommendation. In *WSDM*. 94–102.
- [12] Sungjun Cho, Seunghyun Cho, Sungwoo Park, Hankook Lee, Honglak Lee, and Moontae Lee. 2023. Curve Your Attention: Mixed-Curvature Transformers for Graph Representation Learning. *arXiv preprint arXiv:2309.04082* (2023).
- [13] Aaron Clauset, Christopher Moore, and Mark EJ Newman. 2008. Hierarchical structure and the prediction of missing links in networks. *Nature* 453, 7191 (2008), 98–101.
- [14] Allan M Collins and M Ross Quillian. 1969. Retrieval time from semantic memory. *Journal of verbal learning and verbal behavior* 8, 2 (1969), 240–247.
- [15] Karan Desai, Maximilian Nickel, Tanmay Rajpurohit, Justin Johnson, and Shambhukha Ramakrishna Vedantam. 2023. Hyperbolic image-text representations. In *ICML*. PMLR, 7694–7731.
- [16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [17] Ankit Dhall, Anastasia Makarova, Octavian Ganea, Dario Pavlo, Michael Grefe, and Andreas Krause. 2020. Hierarchical image classification using entailment cone embeddings. In *Proceedings of the CVPR workshops*. 836–837.
- [18] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
- [19] Aleksandr Ermolov, Leyla Mirvakhabova, Valentin Khrulkov, Nicu Sebe, and Ivan Oseledets. 2022. Hyperbolic vision transformers: Combining improvements in metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7409–7419.
- [20] Yifan Feng, Haoxuan You, Zizhao Zhang, Rongrong Ji, and Yue Gao. 2019. Hypergraph neural networks. In *AAAI*, Vol. 33. 3558–3565.
- [21] Luca Franceschi, Matthias Niepert, Massimiliano Pontil, and Xiao He. 2019. Learning discrete structures for graph neural networks. In *ICML*. PMLR, 1972–1982.
- [22] Fabrizio Frasca, Emanuele Rossi, Davide Eynard, Ben Chamberlain, Michael Bronstein, and Federico Monti. 2020. Sign: Scalable inception graph neural networks. *arXiv preprint arXiv:2004.11198* (2020).
- [23] Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. 2018. Hyperbolic entailment cones for learning hierarchical embeddings. In *ICML*. PMLR, 1646–1655.
- [24] Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. 2018. Hyperbolic neural networks. In *NeurIPS*. 5345–5355.
- [25] Albert Gu, Frederic Sala, Beliz Gunel, and Christopher Ré. 2019. Learning mixed-curvature representations in product spaces. In *ICLR*.
- [26] Caglar Gulcehre, Misha Denil, Mateusz Malinowski, Ali Razavi, Razvan Pascanu, Karl Moritz Hermann, Peter Battaglia, Victor Bapst, David Raposo, Adam Santoro, et al. 2019. Hyperbolic attention networks. In *ICLR*.
- [27] Dongchen Han, Xuran Pan, Yizeng Han, Shiji Song, and Gao Huang. 2023. Flatten transformer: Vision transformer using focused linear attention. In *ICCV*. 5961–5971.
- [28] Stephen C Hirtle and John Jonides. 1985. Evidence of hierarchies in cognitive maps. *Memory & cognition* 13, 3 (1985), 208–217.
- [29] Joy Hsu, Jeffrey Gu, Gong Wu, Wah Chiu, and Serena Yeung. 2021. Capturing implicit hierarchical structure in 3D biomedical images with self-supervised hyperbolic representations. In *NeurIPS*, Vol. 34. 5112–5123.
- [30] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. 2020. Open graph benchmark: Datasets for machine learning on graphs. In *NeurIPS*, Vol. 33. 22118–22133.
- [31] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. 2020. Transformers are rnns: Fast autoregressive transformers with linear attention. In *ICML*. PMLR, 5156–5165.
- [32] Valentin Khrulkov, Leyla Mirvakhabova, Evgeniya Ustinova, Ivan Oseledets, and Victor Lempitsky. 2020. Hyperbolic image embeddings. In *CVPR*. 6418–6428.
- [33] Jinwoo Kim, Dat Nguyen, Seonwoo Min, Sungjun Cho, Moontae Lee, Honglak Lee, and Seunghoon Hong. 2022. Pure transformers are powerful graph learners. In *NeurIPS*, Vol. 35. 14582–14595.
- [34] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [35] Thomas N Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*.
- [36] Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. 2020. Reformer: The efficient transformer. *arXiv preprint arXiv:2001.04451* (2020).
- [37] Prodromos Kolyvakis, Alexandros Kalousis, and Dimitris Kirlitsis. 2019. HyperKG: Hyperbolic knowledge graph embeddings for knowledge base completion. *arXiv preprint arXiv:1908.04895* (2019).
- [38] Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguñá. 2010. Hyperbolic geometry of complex networks. *Physical Review E* 82, 3 (2010), 036106.
- [39] Marc Law, Renjie Liao, Jake Snell, and Richard Zemel. 2019. Lorentzian distance learning for hyperbolic representations. In *ICML*. PMLR, 3672–3681.
- [40] Bisheng Li, Min Zhou, Shengzhong Zhang, Menglin Yang, Defu Lian, and Zengfeng Huang. 2022. Bsa!: A framework of bi-component structure and attribute learning for link prediction. In *SIGIR*. 2053–2058.
- [41] Jiahong Liu, Menglin Yang, Min Zhou, Shanshan Feng, and Philippe Fournier-Viger. 2022. Enhancing Hyperbolic Graph Embeddings via Contrastive Learning. In *NeurIPS 2022 workshop on 2nd SSL*.
- [42] Jiahong Liu, Min Zhou, Philippe Fournier-Viger, Menglin Yang, Lujia Pan, and Mourad Nouioua. 2022. Discovering representative attribute-stars via minimum description length. In *ICDE*. IEEE, 68–80.
- [43] Qi Liu, Maximilian Nickel, and Douwe Kiela. 2019. Hyperbolic graph neural networks. In *NeurIPS*. 8230–8241.
- [44] Marko Valentin Micic and Hugo Chu. 2018. Hyperbolic deep learning for chinese natural language understanding. *arXiv preprint arXiv:1812.10408* (2018).
- [45] Erxue Min, Runfa Chen, Yatao Bian, Tingyang Xu, Kangfei Zhao, Wenbing Huang, Peilin Zhao, Junzhou Huang, Sophia Ananiadou, and Yu Rong. 2022. Transformer for graphs: An overview from architecture perspective. *arXiv preprint arXiv:2202.08455* (2022).
- [46] Gal Mishne, Zhengchan Wan, Yusu Wang, and Sheng Yang. 2023. The numerical stability of hyperbolic representation learning. In *International Conference on Machine Learning*. PMLR, 24925–24949.
- [47] Sébastien Montella, Lina M Rojas Barahona, and Johannes Heinecke. 2021. Hyperbolic Temporal Knowledge Graph Embeddings with Relational and Time Curvatures. In *Findings of the ACL: ACL-IJCNLP 2021*. 3296–3308.
- [48] Galileo Namata, Ben London, Lise Getoor, Bert Huang, and U Edu. 2012. Query-driven active surveying for collective classification. In *10th international workshop on mining and learning with graphs*, Vol. 8. 1.
- [49] Mark EJ Newman. 2005. Power laws, Pareto distributions and Zipf's law. *Contemporary physics* 46, 5 (2005), 323–351.
- [50] Maximilian Nickel and Douwe Kiela. 2017. Poincaré embeddings for learning hierarchical representations. In *NeurIPS*. 6338–6347.
- [51] Maximilian Nickel and Douwe Kiela. 2018. Learning Continuous Hierarchies in the Lorentz Model of Hyperbolic Geometry. In *ICML*. 3779–3788.
- [52] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Łukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. 2018. Image transformer. In *ICML*. PMLR, 4055–4064.
- [53] Zexuan Qiu, Jiahong Liu, Yankai Chen, and Irwin King. 2024. HiHPQ: Hierarchical Hyperbolic Product Quantization for Unsupervised Image Retrieval. In *AAAI*. 4614–4622.
- [54] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR* 21, 1 (2020), 5485–5551.
- [55] Ladislav Rampášek, Michael Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and Dominique Beaini. 2022. Recipe for a general, powerful, scalable graph transformer. In *NeurIPS*, Vol. 35. 14501–14515.
- [56] Arlan Ramsay and Robert D Richtmyer. 2013. *Introduction to hyperbolic geometry*. Springer Science & Business Media.
- [57] Robert Resnick. 1991. *Introduction to special relativity*. John Wiley & Sons.

[58] Yu Rong, Wenbing Huang, Tingyang Xu, and Junzhou Huang. 2019. Dropedge: Towards deep graph convolutional networks on node classification. *arXiv preprint arXiv:1907.10903* (2019).

[59] Aurko Roy, Mohammad Saffar, Ashish Vaswani, and David Grangier. 2021. Efficient content-based sparse attention with routing transformers. *Transactions of the Association for Computational Linguistics* 9 (2021), 53–68.

[60] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. 2008. Collective classification in network data. *AI magazine* 29, 3 (2008), 93–93.

[61] Ryohei Shimizu, Yusuke Mukuta, and Tatsuya Harada. 2020. Hyperbolic Neural Networks++. In *ICLR*.

[62] Jianing Sun, Zhaoyue Cheng, Saba Zuberi, Felipe Pérez, and Maksims Volkovs. 2021. HGCF: Hyperbolic Graph Convolution Networks for Collaborative Filtering. In *WWW*. 593–601.

[63] Alexandru Tifrea, Gary Bécigneul, and Octavian-Eugen Ganea. 2018. Poincaré glove: Hyperbolic word embeddings. *arXiv preprint arXiv:1810.06546* (2018).

[64] Max van Spengler, Erwin Berkhouit, and Pascal Mettes. 2023. Poincaré ResNet. *arXiv preprint arXiv:2303.14027* (2023).

[65] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*. 5998–6008.

[66] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903* (2017).

[67] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *ICLR*.

[68] Lucas Vinh Tran, Yi Tay, Shuai Zhang, Gao Cong, and Xiaoli Li. 2020. HyperML: A Boosting Metric Learning Approach in Hyperbolic Space for Recommender Systems. In *WSDM*. New York, NY, USA, 609–617.

[69] Hao Wang, Defu Lian, Hanghang Tong, Qi Liu, Zhenya Huang, and Enhong Chen. 2021. HyperSoRec: Exploiting Hyperbolic User and Item Representations with Multiple Aspects for Social-aware Recommendation. *TOIS* (2021), 1–28.

[70] Hongyu Wang, Shuming Ma, Shaohan Huang, Li Dong, Wenhui Wang, Zhihang Peng, Yu Wu, Payal Bajaj, Saksham Singhal, Alon Benhaim, et al. 2022. Foundation transformers. *arXiv preprint arXiv:2210.06423* (2022).

[71] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. 2019. Simplifying graph convolutional networks. In *ICML*. PMLR, 6861–6871.

[72] Qitian Wu, Wentao Zhao, Zenan Li, David P Wipf, and Junchi Yan. 2022. Nodeformer: A scalable graph structure learning transformer for node classification. In *NeurIPS*, Vol. 35. 27387–27401.

[73] Qitian Wu, Wentao Zhao, Chenxiao Yang, Hengrui Zhang, Fan Nie, Haitian Jiang, Yatao Bian, and Junchi Yan. 2023. Simplifying and empowering transformers for large-graph representations. In *NeurIPS*, Vol. 36.

[74] Zhanghao Wu, Paras Jain, Matthew Wright, Azalia Mirhoseini, Joseph E Gonzalez, and Ion Stoica. 2021. Representing long-range context for graph neural networks with global attention. In *NeurIPS*, Vol. 34. 13266–13279.

[75] Bo Xiong, M. Cochez, Mojtaba Nayyeri, and Steffen Staab. 2022. Hyperbolic Embedding Inference for Structured Multi-Label Prediction. In *NeurIPS*.

[76] Menglin Yang, Zhihao Li, Min Zhou, Jiahong Liu, and Irwin King. 2022. HICF: Hyperbolic informative collaborative filtering. In *KDD*. 2212–2221.

[77] Menglin Yang, Ziqiao Meng, and Irwin King. 2020. Featurenorm: L2 feature normalization for dynamic graph embedding. In *ICDM*. IEEE, 731–740.

[78] Menglin Yang, Min Zhou, Marcus Kalander, Zengfeng Huang, and Irwin King. 2021. Discrete-time Temporal Network Embedding via Implicit Hierarchical Learning in Hyperbolic Space. In *KDD*. 1975–1985.

[79] Menglin Yang, Min Zhou, Zhihao Li, Jiahong Liu, Lujia Pan, Hui Xiong, and Irwin King. 2022. Hyperbolic Graph Neural Networks: A Review of Methods and Applications. *arXiv preprint arXiv:2202.13852* (2022).

[80] Menglin Yang, Min Zhou, Jiahong Liu, Defu Lian, and Irwin King. 2022. HRCF: Enhancing Collaborative Filtering via Hyperbolic Geometric Regularization. In *WWW*.

[81] Menglin Yang, Min Zhou, Lujia Pan, and Irwin King. 2023. κ HGCN: Tree-likeness modeling via continuous and discrete curvature learning. In *KDD*. 2965–2977.

[82] Menglin Yang, Min Zhou, Hui Xiong, and Irwin King. 2022. Hyperbolic Temporal Network Embedding. *IEEE Trans. Knowl. Data Eng.* (2022).

[83] Menglin Yang, Min Zhou, Rex Ying, Yankai Chen, and Irwin King. 2023. Hyperbolic representation learning: Revisiting and advancing. In *ICML*. PMLR, 39639–39659.

[84] Chengxuan Ying, Tianhe Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. 2021. Do Transformers Really Perform Badly for Graph Representation? In *NeurIPS*.

[85] Ce Zhang and Hady W Lauw. 2020. Topic modeling on document networks with adjacent-encoder. In *AAAI*, Vol. 34. 6737–6745.

[86] Delvin Ce Zhang and Hady Lauw. 2022. Dynamic topic models for temporal document networks. In *ICML*. PMLR, 26281–26292.

[87] Delvin Ce Zhang and Hady W Lauw. 2021. Semi-supervised semantic visualization for networked documents. In *Machine Learning and Knowledge Discovery in Databases. Research Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021. Proceedings, Part III* 21. Springer, 762–778.

[88] Delvin Ce Zhang and Hady W Lauw. 2022. Variational graph author topic modeling. In *KDD*. 2429–2438.

[89] Delvin Ce Zhang and Hady W Lauw. 2023. Topic Modeling on Document Networks with Dirichlet Optimal Transport Barycenter. *IEEE Trans. Knowl. Data Eng.* (2023).

[90] Delvin Ce Zhang, Menglin Yang, Rex Ying, and Hady W Lauw. 2024. Text-Attributed Graph Representation Learning: Methods, Applications, and Challenges. In *WWW*. 1298–1301.

[91] Delvin Ce Zhang, Rex Ying, and Hady W Lauw. 2023. Hyperbolic graph topic modeling network with continuously updated topic tree. In *KDD*. 3206–3216.

[92] Yiding Zhang, Xiao Wang, Chuan Shi, Xunqiang Jiang, and Yanfang Fanny Ye. 2021. Hyperbolic graph attention network. *IEEE Trans. Big Data* (2021).

[93] Yiding Zhang, Xiao Wang, Chuan Shi, Nian Liu, and Guojie Song. 2021. Lorentzian Graph Convolutional Networks. In *WWW 2021*. 1249–1261.

[94] Min Zhou, Bisheng Li, Menglin Yang, and Lujia Pan. 2022. TeleGraph: A benchmark dataset for hierarchical link prediction. *arXiv preprint arXiv:2204.07703* (2022).

[95] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A Efros. 2016. Generative visual manipulation on the natural image manifold. In *ECCV*. Springer, 597–613.

A Exponential and Logarithmic Map

Exponential Map. The exponential map, denoted as $\exp_x^\kappa : \mathcal{T}_x \mathbb{L}^{n,\kappa} \rightarrow \mathbb{L}^{n,\kappa}$, is a function that project any tangent vector \mathbf{u} from the tangent space at point \mathbf{x} , $\mathcal{T}_x \mathbb{L}^{n,\kappa}$, to the manifold $\mathbb{L}^{n,\kappa}$, which is given as

$$\exp_x^\kappa(\mathbf{u}) = \cosh\left(\sqrt{|\kappa|}\|\mathbf{u}\|_{\mathcal{L}}\right)\mathbf{x} + \frac{\sinh\left(\sqrt{|\kappa|}\|\mathbf{u}\|_{\mathcal{L}}\right)}{\sqrt{|\kappa|}\|\mathbf{u}\|_{\mathcal{L}}}\mathbf{u}. \quad (22)$$

Logarithmic Map. The logarithmic map $\log_u^\kappa : \mathbb{L}^{n,\kappa} \rightarrow \mathcal{T}_u \mathbb{L}^{n,\kappa}$ plays an opposite role, more specifically,

$$\log_u^\kappa(\mathbf{x}) = \frac{\cosh^{-1}(\kappa\langle \mathbf{u}, \mathbf{x} \rangle_{\mathcal{L}})}{\sinh(\cosh^{-1}(\kappa\langle \mathbf{u}, \mathbf{x} \rangle_{\mathcal{L}}))}(\mathbf{x} - \kappa\langle \mathbf{u}, \mathbf{x} \rangle_{\mathcal{L}}\mathbf{u}). \quad (23)$$

Lorentz Distance. The Lorentz distance between two points ($\mathbf{x} \in \mathbb{L}^{n,\kappa}$, $\mathbf{y} \in \mathbb{L}^{n,\kappa}$) is given as:

$$d_{\mathcal{L}}^\kappa(\mathbf{x}, \mathbf{y}) = \frac{1}{\sqrt{|\kappa|}} \cosh^{-1}(\kappa\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}}) \quad (24)$$

B Proof

Proof of Proposition 4.1

PROOF. Let $\mathbf{L}_x = \text{LTC}(\mathbf{x}; f_t; \mathbf{W}, \kappa_a, \kappa_b)$ and $\langle \mathbf{L}_x, \mathbf{L}_x \rangle_{\mathcal{L}} = 1/\kappa_b$ holds. Besides, $f_t(\mathbf{x}; \mathbf{W}) : \mathbb{R}^{d_a+1} \rightarrow \mathbb{R}^{d_b}$. With the time-like dimension re-calibration and concatenation, $\mathbf{L}_x \in \mathbb{R}^{d_b+1}$. Therefore $\mathbf{L}_x \in \mathbb{L}^{d_b, \kappa_b}$. \square

Proof of Proposition 4.2

PROOF. First, let

$$\mathbf{z} = \left(\sqrt{\|f(\mathbf{x}; \mathbf{W})\|_2^2 - 1/\kappa_1}, f(\mathbf{x}; \mathbf{W}) \right)$$

and then

$$\text{HTC}(\mathbf{x}, \mathbf{W}, \kappa_1, \kappa_2) = \sqrt{\frac{\kappa_1}{\kappa_2}} \cdot \mathbf{z}.$$

We know $\mathbf{z} \in \mathbb{L}^{\kappa_1}$ and $\mathbf{z}' = \text{HTC}(\mathbf{x}, \mathbf{W}, \kappa_1, \kappa_2) \in \mathbb{L}^{\kappa_2}$. Consider the distance between any pair of points in \mathbb{L}^{κ_1} and \mathbb{L}^{κ_2} :

$$d_{\mathcal{L}}^{\kappa_1}(\mathbf{z}_i, \mathbf{z}_j) = \sqrt{1/|\kappa_1|} \text{arcosh}(\kappa_1 \langle \mathbf{z}_i, \mathbf{z}_j \rangle_{\mathcal{L}}), \quad (25)$$

$$d_{\mathcal{L}}^{\kappa_1}(\mathbf{z}_i, \mathbf{z}_k) = \sqrt{1/|\kappa_1|} \operatorname{arccosh}(\kappa_1 \langle \mathbf{z}_i, \mathbf{z}_k \rangle_{\mathcal{L}}), \quad (26)$$

$$d_{\mathcal{L}}^{\kappa_2}(\mathbf{z}'_i, \mathbf{z}'_j) = \sqrt{1/|\kappa_2|} \operatorname{arccosh}(\kappa_2 \langle \mathbf{z}'_i, \mathbf{z}'_j \rangle_{\mathcal{L}}), \quad (27)$$

$$d_{\mathcal{L}}^{\kappa_2}(\mathbf{z}'_i, \mathbf{z}'_k) = \sqrt{1/|\kappa_2|} \operatorname{arccosh}(\kappa_2 \langle \mathbf{z}'_i, \mathbf{z}'_k \rangle_{\mathcal{L}}). \quad (28)$$

We aim to prove that if

$$d_{\mathcal{L}}^{\kappa_1}(\mathbf{z}_i, \mathbf{z}_j) \geq d_{\mathcal{L}}^{\kappa_1}(\mathbf{z}_i, \mathbf{z}_k), \quad (29)$$

then

$$d_{\mathcal{L}}^{\kappa_2}(\mathbf{z}'_i, \mathbf{z}'_j) \geq d_{\mathcal{L}}^{\kappa_2}(\mathbf{z}'_i, \mathbf{z}'_k). \quad (30)$$

Let us expand $d_{\mathcal{L}}^{\kappa_2}(\mathbf{z}'_i, \mathbf{z}'_j)$:

$$\begin{aligned} d_{\mathcal{L}}^{\kappa_2}(\mathbf{z}'_i, \mathbf{z}'_j) &= \sqrt{1/|\kappa_2|} \operatorname{arccosh}(\kappa_2 \langle \mathbf{z}'_i, \mathbf{z}'_j \rangle_{\mathcal{L}}), \\ &= \sqrt{1/|\kappa_2|} \operatorname{arccosh}\left(\kappa_2 \sqrt{\frac{\kappa_1}{\kappa_2}} \cdot \sqrt{\frac{\kappa_1}{\kappa_2}} \langle \mathbf{z}_i, \mathbf{z}_j \rangle_{\mathcal{L}}\right) \\ &= \sqrt{1/|\kappa_2|} \operatorname{arccosh}(\kappa_1 \langle \mathbf{z}_i, \mathbf{z}_j \rangle_{\mathcal{L}}) \\ &= \sqrt{\frac{\kappa_1}{\kappa_2}} (d_{\mathcal{L}}^{\kappa_1}(\mathbf{z}_i, \mathbf{z}_j)) \end{aligned} \quad (31)$$

Similarly, we can show that:

$$d_{\mathcal{L}}^{\kappa_2}(\mathbf{z}'_i, \mathbf{z}'_k) = \sqrt{\frac{\kappa_1}{\kappa_2}} d_{\mathcal{L}}^{\kappa_1}(\mathbf{z}_i, \mathbf{z}_k) \quad (32)$$

Given the inequality in (29), we can multiply both sides by $\sqrt{\frac{\kappa_1}{\kappa_2}}$ (which is positive):

$$\sqrt{\frac{\kappa_1}{\kappa_2}} d_{\mathcal{L}}^{\kappa_1}(\mathbf{z}_i, \mathbf{z}_j) \geq \sqrt{\frac{\kappa_1}{\kappa_2}} d_{\mathcal{L}}^{\kappa_1}(\mathbf{z}_i, \mathbf{z}_k) \quad (33)$$

Substituting from (31) and (32), we obtain:

$$d_{\mathcal{L}}^{\kappa_2}(\mathbf{z}'_i, \mathbf{z}'_j) \geq d_{\mathcal{L}}^{\kappa_2}(\mathbf{z}'_i, \mathbf{z}'_k). \quad (34)$$

This proves the desired inequality (30). \square

C Data Processing and Experimental Details

C.1 Data Processing for Large-graph Data

We employ the public splits offered by OGB [30] for ogbn-proteins and ogbn-arxiv datasets. Additionally, we assess our approach using models on the Amazon2M item co-occurrence network, which comprises 2.45 million nodes and 61.86 million edges. For Amazon2M, we follow the same splits used in recent studies [72, 73]. The largest dataset we employ is ogbn-papers100M, boasting an impressive 0.11 billion nodes and 1.61 billion edges. We also adhere to the publicly available OGB splits for this dataset.

C.2 Data Processing for Medium-graph Data

We used standard splits [34] for the citation networks. For the Airport and Disease datasets, the train/val/test splits were 70%/15%/15% and 30%/10%/60%, respectively, which is the same as [8]. We report the results of five runs on the node classification task. For DISEASE and AIRPORT, which are imbalanced, we report the F1-score. For the other datasets, we report the accuracy. **Baselines.** For the baselines, we compare Hypformer against the basic GNN models, including

GCN [34], GAT [66] and SGC [71]. For Hyperbolic GNN models, we utilized HGNCN [8], LGCN [93] and HGNN [20] as the competitors. We also compared with state-of-the-art Euclidean graph Transformers models viz. Graphomer [84], GraphTrans [74], GraphGPS [55], NodeFormer [72] and SGFormer [73]. Graphomer suggested the incorporation of edge connectivities into the model by employing shortest-path distances to bias the attention mechanism. GraphTrans introduced a permutation-invariant Transformer module combined with a GNN module. While, Nodeformer, GraphGPS and SGFormer each introduced linear attention mechanisms. Specifically, Nodeformer employed a kernelized Gumbel-Softmax, while GraphGPS separated the local real-edge aggregation and the fully-connected Transformer to achieve this complexity. Besides, we also compared with non-Euclidean transformers, i.e., the HAN [26] and FPS-T [12].

C.3 Data Processing for Text and Image Data

We tested our model on two datasets without a graph structure: 20News-Groups and Mini-ImageNet. For our experiment, we selected 30 classes from the dataset, each with 600 images with 128 features extracted by a CNN. These settings closely follow the Nodeformer [72]. For each dataset, we randomly allocate instances into training, validation, and testing sets, comprising 50%, 25%, and 25% of the data, respectively. Following existing works [86–88], we also construct a graph using k-NN (based on input node features) to facilitate the message passing of GNN and the graph transformer. All the datasets we used in the experiment were directly sourced, except for Mini-ImageNet, for which we extracted the features ourselves. Following the approach of [72], we computed node embeddings using a CNN model with four convolutional layers followed by a fully connected layer, resulting in a 128-dimensional embedding. These 128-dimensional outputs are then used as the features of the nodes (images) for subsequent tasks based on Graph Neural Networks (GNNs).

C.4 Implementation Details

In the reported results, we mainly refer to findings from several relevant works for the baseline comparisons [8, 71, 81]. For the most relevant studies, such as SGFormer, other Graph Transformers, we reproduce results using identical experimental settings to ensure a fair comparison. It is important to note that the results of SGFormer cannot be fully reproduced due to errors in its official code implementation. To maintain the integrity of our analysis, we report the performance of SGFormer based on the available information while acknowledging the discrepancy caused by the implementation issues. Our experimental setup for Hypformer mainly follows the configurations used in SGformer. Additionally, we performed parameter tuning for the input curvature and output curvature, exploring values within [1.0, 2.0, 3.0]. This is grounded in our hypothesis that the input attributes and hidden states belong to different curvature spaces. While a more detailed curvature setting could be employed, we leave this for future exploration. Furthermore, we conducted a parameter search for p in Equation (19) within [1.0, 2.0, 3.0]. Regarding the decoder, we created Euclidean and hyperbolic classifiers for experiments, with the Euclidean classifier performing better in most cases.