

Hyperbolic Interaction Model for Hierarchical Multi-Label Classification

Boli Chen, Xin Huang, Lin Xiao, Zixin Cai, Liping Jing

Beijing Key Lab of Traffic Data Analysis and Mining

Beijing Jiaotong University, Beijing, China

{18120345, 18120367, 17112079, 18120340, lpjing}@bjtu.edu.cn

Abstract

Different from the traditional classification tasks which assume mutual exclusion of labels, hierarchical multi-label classification (*HMLC*) aims to assign multiple labels to every instance with the labels organized under hierarchical relations. Besides the labels, since linguistic ontologies are intrinsic hierarchies, the conceptual relations between words can also form hierarchical structures. Thus it can be a challenge to learn mappings from word hierarchies to label hierarchies. We propose to model the word and label hierarchies by embedding them jointly in the hyperbolic space. The main reason is that the tree-likeness of the hyperbolic space matches the complexity of symbolic data with hierarchical structures. A new Hyperbolic Interaction Model (*HyperIM*) is designed to learn the label-aware document representations and make predictions for *HMLC*. Extensive experiments are conducted on three benchmark datasets. The results have demonstrated that the new model can realistically capture the complex data structures and further improve the performance for *HMLC* comparing with the state-of-the-art methods. To facilitate future research, our code is publicly available.

Introduction

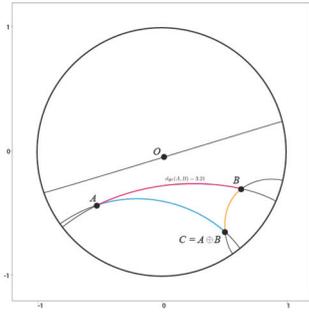
Traditional classification methods suppose the labels are mutually exclusive, whereas for hierarchical classification, labels are not disjointed but organized under a hierarchical structure. Such structure can be a tree or a *Directed Acyclic Graph*, which indicates the parent-child relations between labels. Typical hierarchical classification tasks include protein function prediction in bioinformatics tasks (Wehrmann et al. 2017), image annotation (Dimitrovski et al. 2011) and text classification (Meng et al. 2019). In this paper, we focus on hierarchical multi-label text classification, which aims to assign multiple labels to every document instance with the labels hierarchically structured.

In multi-label classification (*MLC*), there usually exist a lot of infrequently occurring *tail labels* (Bhatia et al. 2015), especially when the label sets are large. The fact that *tail labels* lack of training instances makes it hard to train an efficacious classifier. Fortunately, the effectiveness of utilizing label correlations to address this problem has lately been

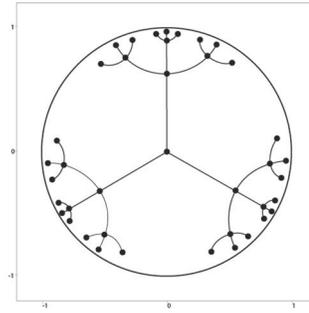
demonstrated. In literatures, label correlations can be determined from label matrix (Zhang et al. 2018) or label content (Wang et al. 2018). The main idea is to project the labels into a latent vectorial space, where each label is represented as a dense low-dimensional vector, so that the label correlations can be characterized in this latent space. For hierarchical multi-label classification (*HMLC*), labels are organized into a hierarchy and located at different hierarchical levels accordingly. Since a parent label generally has several child labels, the number of labels grows exponentially in child levels. In some special cases, most labels are located at the lower levels, and few training instances belong to each of them. In other words, *tail labels* also exist in *HMLC*. Different from the traditional *MLC*, the label structure, which is intuitively useful to detect label correlations, is well provided in *HMLC*.

Inspired by recent works on learning hierarchical representations (Nickel and Kiela 2017), we propose to embed the label hierarchy in the hyperbolic space. Taking advantage of the hyperbolic representation capability, we design a Hyperbolic Interaction Model (*HyperIM*) to classify hierarchically structured labels. *HyperIM* embeds both document words and labels jointly in the hyperbolic space to preserve their latent structures (e.g. structures of conceptual relations between words and parent-child relations between labels). Semantic connections between words and labels can be furthermore explicitly measured according to the word and label embeddings, which benefits extracting the most related components from documents and constructing the label-aware document representations. The prediction is directly optimized by minimizing the *cross-entropy* loss. Our contributions are summarized as follows:

- We adopt hyperbolic space to improve *HMLC*. A novel model *HyperIM* is designed to embed the label hierarchy and the document text in the same hyperbolic space. For the classification, semantic connections between words and labels are explicitly measured to construct the label-aware document representations.
- We present *partial interaction* to improve the scalability of the interaction model. For large label spaces, negative sampling is used to reduce the memory usage during interaction.



(a) Visualization of geodesics and Möbius addition



(b) A tree embedded in the Poincaré disk

Figure 1: (a) Point C represents the Möbius addition of point A and B . In the Poincaré disk model, geodesics between points are arcs and perpendicular to its boundary due to its negative curvature. (b) The line segments indicate the geodesics between each pair of connected nodes in a tree.

- Extensive experiments on three benchmark datasets show the effectiveness of *HyperIM*. An ablation test is performed to demonstrate the superiority of the hyperbolic space over the Euclidean space for *HMLC*. In addition, our code is publicly available.

Preliminaries

Let \mathcal{X} denote the document instance space, and let $\mathcal{L} = \{l_i\}_{i=1}^C$ denote the finite set of C labels. Labels are organized under a hierarchical structure in *HMLC*, $\mathcal{T} = \{(l_p, l_q) \mid l_p \succeq l_q, l_p, l_q \in \mathcal{L}\}$ denotes their parent-child relations, where l_p is the parent of l_q . Given the text sequence of a document instance $x \in \mathcal{X}$ and its one-hot ground truth label vector $y \in \{0, 1\}^C$, the classification model learns the document-label similarities, *i.e.* the probabilities for all the labels given the document. Let $p \in [0, 1]^C$ denote the label probability vector predicted by the model for x , where $p_{[i]} = P(l_i \mid x)$ for $l_i \in \mathcal{L}$ ($i = 1, \dots, c$) (the subscript $[i]$ is used to denote the i -th element in a vector). The model can be trained by optimizing certain loss function that compares y and p .

To capture the fine-grained semantic connections between a document instance and the labels, the document-label similarities are obtained by aggregating the word-label similarities. More specifically, for the text sequence with T word tokens, *i.e.* $x = [x_1, \dots, x_T]$, the i -th label-aware document representation $s_i = [score(x_1, l_i); \dots; score(x_T, l_i)]$ can be calculate via certain score function, $p_{[i]}$ is then deduced from s_i . This process is adapted from the interaction mechanism (Du et al. 2019), which is usually used in tasks like natural language inference (Wang and Jiang 2016). Based on the idea that labels can be considered as abstraction from their word descriptions, sometimes a label is even a word itself, the word-label similarities can be derived from their embeddings in the latent space by the same way as the word similarity, which is widely studied in word embedding methods such as *GloVe* (Pennington, Socher, and Manning 2014).

Note that word embeddings are insufficient to fully represent the meanings of words, especially in the case of *word-sense disambiguation* (Navigli 2009). Take the word "bank" as an example, it has significantly different meanings in the

text sequences "go to the bank and change some money" and "flowers generally grow on the river bank", which will cause a variance when matching with labels "economy" and "environment". In order to capture the real semantics of each word, we use *RNN*-based word encoder which can take the contextual information of text sequences into consideration.

The Poincaré Ball

In *HyperIM*, both document text and labels are embedded in the hyperbolic space. The hyperbolic space is a homogeneous space that has a constant negative sectional curvature, while the Euclidean space has zero curvature. The hyperbolic space can be described via *Riemannian geometry* (Hopper and Andrews 2011). Following previous works (Nickel and Kiela 2017; Ganea, Becigneul, and Hofmann 2018; Tifrea, Bécigneul, and Ganea 2019), we adopt the *Poincaré ball*.

An n -dimensional Poincaré ball ($\mathcal{B}^n, g^{\mathcal{B}^n}$) is a subset of \mathbb{R}^n defined by the *Riemannian manifold* $\mathcal{B}^n = \{x \in \mathbb{R}^n \mid \|x\| < 1\}$ equipped with the *Riemannian metric* $g^{\mathcal{B}^n}$, where $\|\cdot\|$ denotes the Euclidean L^2 norm. As the Poincaré ball is conformal to the Euclidean space (Cannon et al. 1997), the Riemannian metric can be written as $g_p^{\mathcal{B}^n} = \lambda_p^2 g_p^{\mathbb{R}^n}$ with the conformal factor $\lambda_p := \frac{2}{1-\|p\|^2}$ for all $p \in \mathcal{B}^n$, where $g_p^{\mathbb{R}^n} = I_n$ is the Euclidean metric tensor. It is known that the *geodesic distance* between two points $u, v \in \mathcal{B}^n$ can be induced using the ambient Euclidean geometry as $d_{\mathcal{B}^n}(u, v) = \cosh^{-1}(1 + \frac{1}{2}\lambda_u \lambda_v \|u - v\|^2)$. This formula demonstrates that the distance changes smoothly *w.r.t.* $\|u\|$ and $\|v\|$, which is key to learn continuous embeddings for hierarchical structures.

With the purpose of generalizing operations for neural networks in the Poincaré ball, the formalism of the Möbius gyrovector space is used (Ganea, Becigneul, and Hofmann 2018). The Möbius addition for $u, v \in \mathcal{B}^n$ is defined as $u \oplus v = \frac{(1+2\langle u, v \rangle + \|v\|^2)u + (1-\|u\|^2)v}{1+2\langle u, v \rangle + \|u\|^2 + \|v\|^2}$, where $\langle \cdot, \cdot \rangle$ denotes the Euclidean inner product. The Möbius addition operation in the Poincaré disk \mathcal{B}^2 (2-dimensional Poincaré ball) can be visualized in Figure 1a. Then the Poincaré distance can be

rewritten as

$$d_{\mathcal{B}^n}(\mathbf{u}, \mathbf{v}) = 2 \tanh^{-1}(\|\mathbf{u} \oplus \mathbf{v}\|). \quad (1)$$

The Möbius matrix-vector multiplication for $M \in \mathbb{R}^{m \times n}$ and $\mathbf{p} \in \mathcal{B}^n$ when $M\mathbf{p} \neq \mathbf{0}$ is defined as $M \otimes \mathbf{p} = \tanh(\frac{\|M\mathbf{p}\|}{\|\mathbf{p}\|} \tanh^{-1}(\|\mathbf{p}\|)) \frac{M\mathbf{p}}{\|M\mathbf{p}\|}$, and $M \otimes \mathbf{p} = \mathbf{0}$ when $M\mathbf{p} = \mathbf{0}$. Moreover, the closed-form derivations of the exponential map $\exp_{\mathbf{p}} : T_{\mathbf{p}}\mathcal{B}^n \rightarrow \mathcal{B}^n$ and the logarithmic map $\log_{\mathbf{p}} : \mathcal{B}^n \rightarrow T_{\mathbf{p}}\mathcal{B}^n$ for $\mathbf{p} \in \mathcal{B}^n$, $\mathbf{w} \in T_{\mathbf{p}}\mathcal{B}^n \setminus \{\mathbf{0}\}$, $\mathbf{u} \in \mathcal{B}^n \setminus \{\mathbf{p}\}$ are given as $\exp_{\mathbf{p}}(\mathbf{w}) = \mathbf{p} \oplus (\tanh(\frac{\lambda_{\mathbf{p}}}{2} \|\mathbf{w}\|) \frac{\mathbf{w}}{\|\mathbf{w}\|})$ and $\log_{\mathbf{p}}(\mathbf{u}) = \frac{2}{\lambda_{\mathbf{p}}} \tanh^{-1}(\|\mathbf{u} \oplus \mathbf{p}\|) \frac{-\mathbf{p} \oplus \mathbf{u}}{\|\mathbf{u} \oplus \mathbf{p}\|}$.

These operations make hyperbolic neural networks available (Ganea, Becigneul, and Hofmann 2018) and gradient-based optimizations can be performed to estimate the model parameters in the Poincaré ball (Bécigneul and Ganea 2019).

Hyperbolic Interaction Model

We design a Hyperbolic Interaction Model (*HyperIM*) for hierarchical multi-label text classification. Given the text sequence of a document, *HyperIM* measures the word-label similarities by calculating the geodesic distance between the jointly embedded words and labels in the Poincaré ball. The word-label similarity scores are then aggregated to estimate the label-aware document representations and further predict the probability for each label. Figure 2 illustrates the framework of *HyperIM*.

Hyperbolic Label Embedding

The tree-likeness of the hyperbolic space (Hamann 2018) makes it natural to embed hierarchical structures. For instance, Figure 1b presents a tree embedded in the Poincaré disk, where the root is placed at the origin and the leaves are close to the boundary. It has been shown that any finite tree can be embedded with arbitrary low distortion into the Poincaré ball while the distances are approximately preserved (Sarkar 2011). Conversely, it is difficult to perform such embedding in the Euclidean space even with unbounded dimensionality (Sala et al. 2018). Since the label hierarchy is defined in the set $\mathcal{T} = \{(l_p, l_q) \mid l_p \succeq l_q, l_p, l_q \in \mathcal{L}\}$, the goal is to maximize the distance between labels without parent-child relation (Nickel and Kiela 2017). Let $\Theta^L = \{\theta_i^l\}_{i=1}^C$, $\theta_i^l \in \mathcal{B}^k$ be the label embedding set, using Riemannian adaptive optimization methods (Bécigneul and Ganea 2019), Θ^L can be efficiently estimated by minimizing the loss function

$$\mathcal{L}_{loss}^h = - \sum_{(l_p, l_q) \in \mathcal{T}} \log \frac{\exp(-d_{\mathcal{B}^k}(\theta_p^l, \theta_q^l))}{\sum_{l_{q'} \in \mathcal{N}(l_p)} \exp(-d_{\mathcal{B}^k}(\theta_p^l, \theta_{q'}^l))}, \quad (2)$$

where $\mathcal{N}(l_p) = \{l_{q'} \mid (l_p, l_{q'}) \notin \mathcal{T}\} \cup \{l_p\}$ is the set of negative samples. The obtained Θ^L can capture the hierarchical structure among labels.

Hyperbolic Word Embedding

For natural language processing, word embeddings are essential in neural networks as intermediate features. Given

the statistics of word co-occurrences in the corpus, we adopt the *Poincaré GloVe* (Tifrea, Bécigneul, and Ganea 2019) to capture the elementary relations between words by embedding them in the hyperbolic space. Let X_{ij} indicate the times that word i and word j co-occur in the same context window, $\theta_i^e \in \mathcal{B}^k$ be the target embedding vector in the k -dimensional Poincaré ball for word i , and $\tilde{\theta}_j^e \in \mathcal{B}^k$ be the context embedding vector for word j . With the aid of Riemannian adaptive optimization methods, the embeddings $\Theta^E = \{\theta_i^e\}_{i=1}^V$ and $\tilde{\Theta}^E = \{\tilde{\theta}_j^e\}_{j=1}^V$ for the corpus with vocabulary size V are estimated by minimizing the loss function

$$\mathcal{L}_{loss}^e = \sum_{i,j=1}^V f(X_{ij}) (-h(d_{\mathcal{B}^k}(\theta_i^e, \tilde{\theta}_j^e)) + b_i + \tilde{b}_j - \log(X_{ij}))^2, \quad (3)$$

where b_i, \tilde{b}_j are the biases, and the two suggested weight functions are defined as $f(x) = \min(1, (x/100)^{3/4})$, $h(x) = \cosh^2(x)$.

Moreover, some words in the corpus are also in the *WordNet* (Miller 1995) hypernym set $\mathcal{T}^w = \{(x_p, x_q) \mid x_p \succeq x_q\}$, where word x_p is the hypernym of word x_q . As it is similar to the label hierarchy, providing the hypernym information to the word embeddings, latent correlations between the two hierarchies can be later captured via interaction. Therefore, a post-processing step is further conducted to minimize a loss similar to Eq. (2) on top of the learned *Poincaré GloVe* embeddings Θ^E , which do not explicitly capture the conceptual relations among words.

Hyperbolic Word Encoder

Considering the *word-sense disambiguation* (Navigli 2009), meanings of polysemous words are difficult to distinguish if the word and label embeddings interact with each other directly, since all the meanings of a word are embedded on the same position. However, polysemous words can usually be inferred from the context.

Given the text sequence of a document with T word tokens $\mathbf{x} = [x_1, \dots, x_T]$, pre-trained hyperbolic word embeddings Θ^E can be used to learn the final word representations according to the text sequence. To consider the sequentiality of the text sequence, we take advantage of the hyperbolic space adaptive *RNN*-based architectures (Ganea, Becigneul, and Hofmann 2018). More specifically, given $\Theta^e = [\theta_1^e, \dots, \theta_T^e]$ where $\theta_t^e \in \Theta^E (t = 1, \dots, T)$, the hyperbolic word encoder based on the *GRU* architecture adjusts the embedding for each word to fit its context via

$$\begin{aligned} \mathbf{r}_t &= \sigma(\log_0(\mathbf{W}^r \otimes \theta_{t-1}^w \oplus \mathbf{U}^r \otimes \theta_t^e \oplus \mathbf{b}^r)), \\ \mathbf{z}_t &= \sigma(\log_0(\mathbf{W}^z \otimes \theta_{t-1}^w \oplus \mathbf{U}^z \otimes \theta_t^e \oplus \mathbf{b}^z)), \\ \tilde{\theta}_t^w &= \varphi((\mathbf{W}^g \text{diag}(\mathbf{r}_t)) \otimes \theta_{t-1}^w \oplus \mathbf{U}^g \otimes \theta_t^e \oplus \mathbf{b}^g), \\ \theta_t^w &= \theta_{t-1}^w \oplus \text{diag}(\mathbf{z}_t) \otimes (-\theta_{t-1}^w \oplus \tilde{\theta}_t^w), \end{aligned} \quad (4)$$

where $\Theta^w = [\theta_1^w, \dots, \theta_T^w]$ denotes the encoded embeddings for the text sequence, the initial hidden state $\theta_0^w := \mathbf{0}$, \mathbf{r}_t is the reset gate, \mathbf{z}_t is the update gate, $\text{diag}(\cdot)$ denotes the diagonal matrix with each element in the vector on its

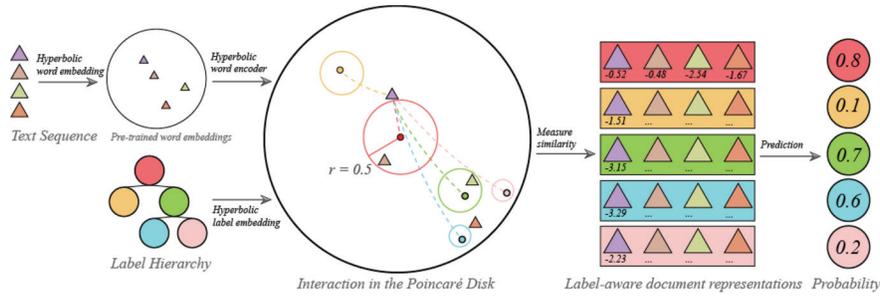


Figure 2: Framework of the Hyperbolic Interaction Model (*HyperIM*). Word-label similarities are measured in the Poincaré Disk. The label nodes are the centres of the hyperbolic circles, which have the same radius. The dash lines are the geodesics from the label nodes to a word node. Note that the hyperbolic centers of the circles in general do not correspond to the Euclidean ones. Labels have the same similarity scores for words embedded on the boundary of their circles.

diagonal, σ is the *sigmoid* function, φ is a pointwise non-linearity, typically *sigmoid*, *tanh* or *ReLU*. Since the hyperbolic space naturally has non-linearity, φ can be *identity* (no non-linearity) here.

The formula of the *hyperbolic GRU* is derived by connecting the Möbius gyrovector space with the Poincaré ball (Ganea, Becigneul, and Hofmann 2018). The six weights $\mathbf{W}, \mathbf{U} \in \mathbb{R}^{k \times k}$ are trainable parameters in the Euclidean space and the three biases $\mathbf{b} \in \mathcal{B}^k$ are trainable parameters in the hyperbolic space (the superscripts are omitted for simplicity). Thus the weights \mathbf{W} and \mathbf{U} are updated via vanilla optimization methods, and the biases \mathbf{b} are updated with Riemannian adaptive optimization methods. Θ^w will be used for measuring the word-label similarities during the following interaction process.

Interaction in the Hyperbolic Space

The major objective of text classification is to build connections from the word space to the label space. In order to capture the fine-grained semantic information, we first construct the label-aware document representations, and then learn the mappings between the document instance and the labels.

Label-Aware Document Representations Once the encoded word embeddings Θ^w and label embeddings Θ^L are obtained, it is expected that every pair of word and label embedded close to each other based on their geodesic distance if they are semantically similar. Note that cosine similarity (Wang, Hamza, and Florian 2017) is not appropriate to be the metric since there does not exist a clear hyperbolic inner-product for the the Poincaré ball (Tifrea, Bécigneul, and Ganea 2019), so the geodesic distance is more intuitively suitable. The similarity between the t -th word $x_t (t = 1, \dots, T)$ and the i -th label $l_i (i = 1, \dots, C)$ is calculated as $score(x_t, l_i) = -d_{\mathcal{B}^k}(\theta_t^w, \theta_i^l)$, where θ_t^w and θ_i^l are their corresponding embeddings, $d_{\mathcal{B}^k}(\cdot, \cdot)$ is the Poincaré distance function defined in Eq. (1). The i -th label-aware document representation can be formed as the concatenation of all the similarities along the text sequence, *i.e.* $\mathbf{s}_i = [score(x_1, l_i); \dots; score(x_T, l_i)]$. The set $\mathcal{S} = \{\mathbf{s}_i\}_{i=1}^C$ acquired along the labels can be taken as the label-aware document representations under the hyperbolic word and label embeddings.

Prediction Given the document representations in \mathcal{S} , predictions can be made by a fully-connected layer and an output layer. The probability of each label for the document instance can be obtained by

$$p_i = \sigma(\mathbf{W}^e \varphi(\mathbf{W}^f \mathbf{s}_i)), \forall \mathbf{s}_i \in \mathcal{S}, i = 1, \dots, C, \quad (5)$$

where σ is the *sigmoid* function, φ is a non-linearity. The weights $\mathbf{W}^e \in \mathbb{R}^{1 \times (T/2)}$ and $\mathbf{W}^f \in \mathbb{R}^{(T/2) \times T}$ are trainable parameters.

Partial Interaction

During the above interaction process, the amount of computation increases with the number of labels. When the output label space is large, it is a burden to calculate the label-aware document representations. On account of the fact that only a few labels are assigned to one document instance, we propose to use a negative sampling method to improve the scalability during training. Let \mathcal{L}^+ denote the set of true labels and \mathcal{L}^- denote the set of randomly selected negative labels, the model is trained by minimizing the loss function which is derived from the *binary cross-entropy* loss as it is commonly used for *MLC* (Liu et al. 2017), *i.e.*

$$\mathcal{L}_{loss}^b = -\left(\sum_{i \in \mathcal{L}^+} \log(p_i) + \sum_{j \in \mathcal{L}^-} \log(1 - p_j)\right). \quad (6)$$

The hyperbolic parameters, *i.e.* Θ^E , Θ^L and \mathbf{b} in the hyperbolic word encoder, are updated via Riemannian adaptive optimization methods. The Euclidean parameters, *i.e.* \mathbf{W} , \mathbf{U} in the hyperbolic word encoder and \mathbf{W} in the prediction layers, are updated via vanilla optimization methods. *Partial interaction* can significantly reduce the memory usage during training especially when the label set is large.

Experiments

Datasets Experiments are carried out on three publicly available multi-label text classification datasets, including the small-scale *RCVI* (Lewis et al. 2004), the middle-scale *Zhihu*¹ and the large-scale *WikiLSHTC* (Partalas et al. 2015). All the datasets are equipped with labels that explicitly exhibit a hierarchical structure. Their statistics can be found in Table 1.

¹<https://biendata.com/competition/zhihu/>.

Table 1: Statistics of the datasets: N_{train} and N_{test} are the number of training and test instances, L is the number of labels, \hat{L} is the average number of label per document, \tilde{L} is the average number of documents per label, W_{train} and W_{test} denote the average number of words per document in the training and test set respectively.

Dataset	N_{train}	N_{test}	L	\hat{L}	\tilde{L}	W_{train}	W_{test}
<i>RCV1</i>	23,149	781,265	103	3.18	729.67	259.47	269.23
<i>Zhifu</i>	2,699,969	299,997	1,999	2.32	3513.17	38.14	35.56
<i>WikiLSHTC</i>	456,886	81,262	36,504	1.86	4.33	117.98	118.31

Evaluation Metrics We use the rank-based evaluation metrics which have been widely adopted for multi-label classification tasks, *i.e.* $Precision@k$ ($P@k$ for short) and $nDCG@k$ for $k = 1, 3, 5$ (Bhatia et al. 2015; Liu et al. 2017; Zhang et al. 2018). Let $\mathbf{y} \in \{0, 1\}^C$ be the ground truth label vector for a document instance and $\mathbf{p} \in [0, 1]^C$ be the predicted label probability vector. $P@k$ records the fraction of correct predictions in the top k possible labels. Let the vector $\mathbf{r} \in \{1, \dots, C\}^k$ denote the indices for k most possible labels in descending order, *i.e.* the $r_{[1]}$ -th label has the largest probability to be true, then the metrics are defined as

$$P@k = \frac{1}{k} \sum_{i=1}^k \mathbf{y}_{[r_{[i]}]}, \quad (7)$$

$$nDCG@k = \frac{\sum_{i=1}^k \mathbf{y}_{[r_{[i]}]}/\log(i+1)}{\sum_{i=1}^{\min(k, \|\mathbf{y}\|_0)} 1/\log(i+1)},$$

where $\|\mathbf{y}\|_0$ denotes the number of true labels, *i.e.* the number of 1 in \mathbf{y} . The final results are averaged over all the test document instances. Notice that $nDCG@1$ is omitted in the results since it gives the same value as $P@1$.

Baselines To demonstrate the effectiveness of *HyperIM* on the benchmark datasets, five comparative multi-label classification methods are chosen. *EXAM* (Du et al. 2019) is the state-of-the-art interaction model for text classification. *EXAM* use pre-trained word embeddings in the Euclidean space, its label embeddings are randomly initialized. To calculate the similarity scores, *EXAM* uses the dot-product between word and label embeddings. *SLEEC* (Bhatia et al. 2015) and *DXML* (Zhang et al. 2018) are two label-embedding methods. *SLEEC* projects labels into low-dimensional vectors which can capture label correlations by preserving the pairwise distance between them. *SLEEC* uses the k -nearest neighbors when predicting, and clustering is used to speed up its prediction. Ensemble method is also used to improve the performance of *SLEEC*. *DXML* uses *DeepWalk* (Perozzi, Al-Rfou, and Skiena 2014) to embed the label co-occurrence graph into vectors, and uses neural networks to map the features into the embedding space. *HR-DGCNN* (Peng et al. 2018) and *HMCN-F* (Wehrmann, Cerri, and Barros 2018) are two neural network models specifically designed for hierarchical classification tasks. Taking advantage of the label hierarchy, *HR-DGCNN* adds a regularization term on the weights of the fully-connected layer. The original *HMCN-F* cannot take in the raw text data.

To make *HMCN-F* more competitive, *CNN*-based architecture similar to *XML-CNN* (Liu et al. 2017) is adopted to extract the primary features. *HMCN-F* then fits its neural network layers to the label hierarchy, each layer focuses on predicting the labels in the corresponding hierarchical level.

Results As shown in Table 2, *HyperIM* consistently outperforms all the baselines. *HyperIM* effectively takes advantage of the label hierarchical structure comparing with *EXAM*, *SLEEC* and *DXML*. *EXAM* uses the interaction mechanism to learn word-label similarities, whereas clear connections between the words and the label hierarchy cannot be captured since its label embeddings are randomly initialized. The fact that *HyperIM* achieves better results than *EXAM* further confirms that *HyperIM* benefits from the retention of the hierarchical label relations. Meanwhile, the word embeddings learned by *HyperIM* have strong connections to the label structure, which is helpful to the measurement of word-label similarities and the acquirement of the label-aware document representations. *SLEEC* and *DXML* take the label correlations into account. However, the label correlations they use are captured from the label matrix, *e.g.* embedding the label co-occurrence graph, which may be influenced by *tail labels*. For *HyperIM*, the label relations are determined from the label hierarchy, so the embeddings of labels with parent-child relations are dependable to be correlated.

As expected, *HyperIM* is superior to the existing hierarchical classification methods *HR-DGCNN* and *HMCN-F*, even though they take advantage of the label hierarchy information. By investigating the properties of these three methods, we summarize the main reasons as follows. *HR-DGCNN* adds the regularization terms based on the assumption that labels with parent-child relations should have similar weights in the fully-connected layer, which may not always be true in real applications. *HMCN-F* highly depends on the label hierarchy, it assumes that different paths pass through the same number of hierarchical levels. Unfortunately, in the real data, different paths may have totally different lengths. *HyperIM* models the label relations by embedding the label hierarchy in the hyperbolic space. Any hierarchical structure can be suitable and labels are not required to sit on a specific hierarchical level, which makes *HyperIM* less reliant on the label hierarchy. Furthermore, *HyperIM* can learn the word-label similarities and preserve the label relations simultaneously to acquire label-aware document representations, whereas *HR-DGCNN* and *HMCN-F* treat document words and labels separately.

Experimental Details In order to evaluate the baselines on the benchmark datasets, hyperparameters recommended by their authors are used. On account of the numeric error issue caused by the constrain $\|\mathbf{p}\| < 1$ for $\mathbf{p} \in \mathcal{B}^k$ when the embedding dimension k is large, the workaround is taken to address this issue, *i.e.* the embedding vector is a concatenation of vectors in the low-dimensional Poincaré ball. Consequently, the embedding dimension for *HyperIM* is $75 \times 2D$ as it generally outperforms the baselines.

Table 2: Results in $P@k$ and $nDCG@k$, bold face indicates the best in each line.

Dataset	Metrics	EXAM	SLEEC	DXML	HR-DGCNN	HMCN-F	HyperIM
RCVI	$P@1$	95.98	94.45	95.27	95.17	95.35	96.78
	$P@3$	80.83	78.60	77.86	80.32	78.95	81.46
	$P@5$	55.80	54.24	53.44	55.38	55.90	56.79
	$nDCG@3$	90.74	90.05	89.69	90.02	90.14	91.52
	$nDCG@5$	91.26	90.32	90.24	90.28	90.82	91.89
Zhihu	$P@1$	51.41	51.34	50.34	50.97	50.24	52.14
	$P@3$	32.81	32.56	31.21	32.41	32.18	33.66
	$P@5$	24.29	24.23	23.36	23.87	24.09	24.99
	$nDCG@3$	49.32	49.27	47.92	49.02	48.36	50.13
	$nDCG@5$	50.74	49.71	48.65	49.91	49.21	51.05
WikiLSHTC	$P@1$	54.90	53.57	52.02	52.67	53.23	55.06
	$P@3$	30.50	31.25	30.57	30.13	29.32	31.73
	$P@5$	22.02	22.46	21.66	22.85	21.79	23.08
	$nDCG@3$	49.50	46.06	47.97	49.24	48.93	50.46
	$nDCG@5$	50.46	47.52	48.14	50.42	49.87	51.36

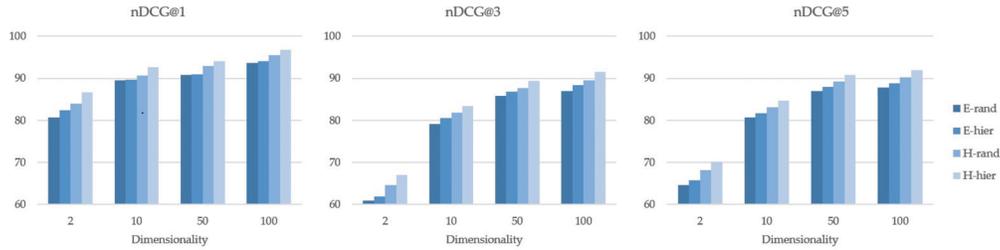


Figure 3: Results in $nDCG@k$ for the ablation test. E -rand and H -rand denote *EuclideanIM* and *HyperIM* take the randomly initialized label embeddings respectively, E -hier and H -hier take the same label embeddings initialized according to the hierarchical label relations.

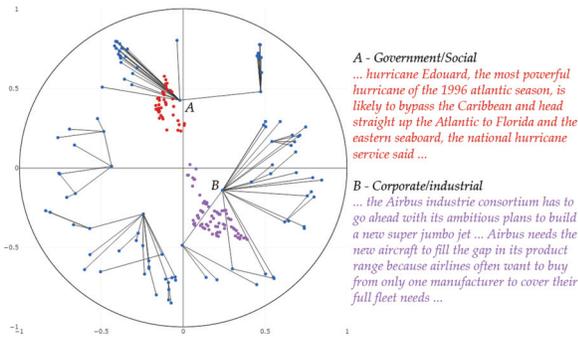


Figure 4: Visualization of labels (blue nodes) and words jointly embedded in the Poincaré disk. The connected labels denote that they are related to each other.

Ablation Test

In order to show the characteristics of *HyperIM* and justify the superiority of the hyperbolic space for hierarchical multi-label text classification, we are interested in comparing it with an analogous model in the Euclidean space.

Euclidean Interaction Model The analogous model in the Euclidean space (*EuclideanIM*) has a similar architec-

ture as *HyperIM*. *EuclideanIM* takes the vanilla pre-trained *GloVe* word embeddings (Pennington, Socher, and Manning 2014) and uses the vanilla *GRU* (Chung et al. 2014) as the word encoder. The label embeddings are randomly initialized for E -rand, while E -hier takes the same label embeddings initialized by the hierarchical label relations as H -hier. The word-label similarities are computed as the negative of the Euclidean distance between their embeddings, i.e. $score(x_t, l_i) = -\|\theta_t^w - \theta_i^l\|$ for $\theta_t^w, \theta_i^l \in \mathbb{R}^k$. The same architecture of the prediction layers is adopted.

Results Figure 3 shows the $nDCG@k$ results for different embedding dimensions on the *RCVI* dataset. The fact that E -hier slightly outperforms E -rand indicates that the label correlations provide useful information for classification. However, H -rand still achieves better results than the Euclidean models even without the hierarchical label structure information, which confirms that the hyperbolic space is more suitable for *HMLC*. For E -hier, the hierarchical label structure is not appropriate to be embedded in the Euclidean space, thus it cannot fully take advantage of such information. *HyperIM* generally outperforms *EuclideanIM* and achieves significant improvement especially in low-dimensional latent space. H -hier takes in the label correlations and outperforms H -rand as expected.

Interaction Visualization The 2-dimensional hyperbolic label embeddings and the encoded word embeddings (not

the pre-trained word embeddings) can be visualize jointly in the Poincaré disk as shown in Figure 4. The hierarchical label structure which can represent the parent-child relations between labels is well preserved by *HyperIM*. Note that the embedded label hierarchy resembles the embedded tree in Figure 1b. The top-level nodes (e.g. the label node *A*) are embedded near the origin of the Poincaré disk, while the leaf nodes are close to the boundary. The hierarchical label relations are well modeled by such tree-like structure. Moreover, in the dataset, the top-level labels are not connected to an abstract "root". The structure of the embedded label hierarchy still suggests that there should be a "root" that connects all the top-level labels to put at the very origin of the Poincaré disk, which indicates that *HyperIM* can really make use of the hierarchical label relations.

The explicit label correlations can further help *HyperIM* to learn to encode the word embeddings via interaction. The encoded text of a document instance are generally embedded close to the assigned labels. This clear pattern between the encoded word embeddings and the label hierarchy indicates that *HyperIM* learns the word-label similarities with the label correlations taken into consideration. This is the main reason that *HyperIM* outperforms *EuclideanIM* significantly in low dimensions. Some of the words such as "the", "is" and "to" do not provide much information for classification, putting these words near the origin can make them equally similar to labels in the same hierarchical level. A nice by-product is that the predicted probabilities for labels in the same hierarchical level won't be influenced by these words. Moreover, the variance of word-label distance for labels in different hierarchical levels make parent labels distinguishable from child labels, e.g. top-level labels can be made different from the leaf labels since they are generally closer to the word embeddings. Such difference suggests that *HyperIM* treats the document instances differently along the labels in different hierarchical levels.

Related Work

Hierarchical Multi-Label Classification

The existing methods dedicating to hierarchical classification usually focus on the design of loss functions or neural network architectures (Cerri, Barros, and de Carvalho 2015; Gargiulo et al. 2019). Traditional hierarchical classification methods optimize a loss function locally or globally (Silla and Freitas 2011). Local methods are better at capturing label correlations, whereas global methods are less computationally expensive. Researchers recently try to use a hybrid loss function associated with specifically designed neural networks, e.g. *HR-DGCNN* (Peng et al. 2018). The archetype of *HMCN-F* (Wehrmann, Cerri, and Barros 2018) employs a cascade of neural networks, where each neural networks layer corresponds to one level of the label hierarchy. Such neural network architectures generally require all the paths in the label hierarchy to have the same length, which limits their application. Moreover, on account of the fact that labels in high hierarchical levels usually contain much more instances than labels in low levels, whereas neural network layers for low levels need to classify more la-

bels than layers for high levels, such architectures also lead to imbalance classification.

Hyperbolic Deep Learning

Research on representation learning (Nickel and Kiela 2017; Wang, Zhang, and Shi 2019; Chen and Quirk 2019) indicates that the hyperbolic space is more suitable for embedding symbolic data with hierarchical structures than the Euclidean space, since the tree-likeness properties (Hamann 2018) of the hyperbolic space make it efficient to learn hierarchical representations with low distortion (Sarkar 2011). As linguistic ontologies are innately hierarchies, hierarchies are ubiquitous in natural language, (e.g. *WordNet* (Miller 1995)). Some works lately demonstrate the superiority of the hyperbolic space for natural language processing tasks such as textual entailment (Ganea, Bécigneul, and Hofmann 2018), machine translation (Gulcehre et al. 2019) and word embedding (Tifrea, Bécigneul, and Ganea 2019).

Riemannian Optimization In the same way that gradient-based optimization methods are used for trainable parameters in the Euclidean space, the hyperbolic parameters can be updated via Riemannian adaptive optimization methods (Bécigneul and Ganea 2019). For instance, *Riemannian adaptive SGD* updates the parameters $\theta \in \mathcal{B}^k$ by $\theta_{t+1} = \exp_{\theta_t}(-\eta \nabla_R \mathcal{L}(\theta_t))$, where η is the learning rate, and the Riemannian gradient $\nabla_R \mathcal{L}(\theta_t) \in T_{\theta} \mathcal{B}^k$ is the rescaled Euclidean gradient, i.e. $\nabla_R \mathcal{L}(\theta) = \frac{1}{\lambda_{\theta}^2} \nabla_E \mathcal{L}(\theta)$ (Wilson and Leimeister 2018).

Conclusion

The hierarchical parent-child relations between labels can be well modeled in the hyperbolic space. The proposed *HyperIM* is able to explicitly learn the word-label similarities by embedding the words and labels jointly and preserving the label hierarchy simultaneously. *HyperIM* acquires label-aware document representations to extract the fine-grained text content along each label, which significantly improves the hierarchical multi-label text classification performance. Indeed, *HyperIM* makes use of the label hierarchy, whereas there is usually no such hierarchically organized labels in practice, especially for extreme multi-label classification (*XMLC*). Nevertheless, the labels in *XMLC* usually follow a power-law distribution due to the amount of *tail labels* (Babbar and Schölkopf 2018), which can be traced back to hierarchical structures (Ravasz and Barabási 2003). Thus, it will be interesting to extend *HyperIM* for *XMLC* in the future. Our code is publicly available to facilitate other research.²

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant 61822601, 61773050, and 61632004; the Beijing Natural Science Foundation under Grant Z180006; the Beijing Municipal Science & Technology Commission under Grant

²<https://github.com/bcol23/HyperIM>

Z181100008918012; National Key Research and Development Program (2017YFC1703506); the Fundamental Research Funds for the Central Universities (2019JBZ110).

References

- Babbar, R., and Schölkopf, B. 2018. Adversarial extreme multi-label classification. *arXiv preprint arXiv:1803.01570*.
- Bécigneul, G., and Ganea, O.-E. 2019. Riemannian adaptive optimization methods. In *Proceedings of the Seventh International Conference on Learning Representations*.
- Bhatia, K.; Jain, H.; Kar, P.; Varma, M.; and Jain, P. 2015. Sparse local embeddings for extreme multi-label classification. In *Advances in neural information processing systems* 28, 730–738.
- Cannon, J. W.; Floyd, W. J.; Kenyon, R.; Parry, W. R.; et al. 1997. Hyperbolic geometry. *Flavors of geometry* 31:59–115.
- Cerri, R.; Barros, R. C.; and de Carvalho, A. C. 2015. Hierarchical classification of gene ontology-based protein functions with neural networks. In *2015 International Joint Conference on Neural Networks*, 1–8. IEEE.
- Chen, M., and Quirk, C. 2019. Embedding edge-attributed relational hierarchies. In *Proceedings of the Forty-second International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM.
- Chung, J.; Gulcehre, C.; Cho, K.; and Bengio, Y. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *NIPS 2014 Workshop on Deep Learning*.
- Dimitrovski, I.; Kocev, D.; Loskovska, S.; and Džeroski, S. 2011. Hierarchical annotation of medical images. *Pattern Recognition* 44(10-11):2436–2449.
- Du, C.; Chin, Z.; Feng, F.; Zhu, L.; Gan, T.; and Nie, L. 2019. Explicit interaction model towards text classification. In *Proceedings of the Thirty-third AAAI Conference on Artificial Intelligence*.
- Ganea, O.; Bécigneul, G.; and Hofmann, T. 2018. Hyperbolic neural networks. In *Advances in neural information processing systems* 31, 5345–5355.
- Gargiulo, F.; Silvestri, S.; Ciampi, M.; and De Pietro, G. 2019. Deep neural network for hierarchical extreme multi-label text classification. *Applied Soft Computing* 79:125–138.
- Gulcehre, C.; Denil, M.; Malinowski, M.; Razavi, A.; Pascanu, R.; Hermann, K. M.; Battaglia, P.; Bapst, V.; Raposo, D.; Santoro, A.; et al. 2019. Hyperbolic attention networks. In *Proceedings of the Seventh International Conference on Learning Representations*.
- Hamann, M. 2018. On the tree-likeness of hyperbolic spaces. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 164, 345–361. Cambridge University Press.
- Hopper, C., and Andrews, B. 2011. *The Ricci flow in Riemannian geometry*. Springer.
- Lewis, D. D.; Yang, Y.; Rose, T. G.; and Li, F. 2004. Rcv1: A new benchmark collection for text categorization research. *Journal of machine learning research* 5(Apr):361–397.
- Liu, J.; Chang, W.-C.; Wu, Y.; and Yang, Y. 2017. Deep learning for extreme multi-label text classification. In *Proceedings of the Forty International ACM SIGIR Conference on Research and Development in Information Retrieval*, 115–124. ACM.
- Meng, Y.; Shen, J.; Zhang, C.; and Han, J. 2019. Weakly-supervised hierarchical text classification. In *Proceedings of the Thirty-third AAAI Conference on Artificial Intelligence*.
- Miller, G. A. 1995. Wordnet: a lexical database for english. *Communications of the ACM* 38(11):39–41.
- Navigli, R. 2009. Word sense disambiguation: A survey. *ACM computing surveys (CSUR)* 41(2):10.
- Nickel, M., and Kiela, D. 2017. Poincaré embeddings for learning hierarchical representations. In *Advances in neural information processing systems* 30, 6338–6347.
- Partalas, I.; Kosmopoulos, A.; Baskiotis, N.; Artières, T.; Paliouras, G.; Gaussier, É.; Androutsopoulos, I.; Amini, M.; and Gallinari, P. 2015. LSHTC: A benchmark for large-scale text classification. *CoRR* abs/1503.08581.
- Peng, H.; Li, J.; He, Y.; Liu, Y.; Bao, M.; Song, Y.; and Yang, Q. 2018. Large-scale hierarchical text classification with recursively regularized deep graph-cnn. In *WWW*.
- Pennington, J.; Socher, R.; and Manning, C. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, 1532–1543.
- Perozzi, B.; Al-Rfou, R.; and Skiena, S. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the Twenty ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 701–710. ACM.
- Ravasz, E., and Barabási, A.-L. 2003. Hierarchical organization in complex networks. *Physical review E* 67(2):026112.
- Sala, F.; De Sa, C.; Gu, A.; and Re, C. 2018. Representation trade-offs for hyperbolic embeddings. In *Proceedings of the Thirty-fifth International Conference on Machine Learning*, 4457–4466.
- Sarkar, R. 2011. Low distortion delaunay embedding of trees in hyperbolic plane. In *International Symposium on Graph Drawing*, 355–366. Springer.
- Silla, C. N., and Freitas, A. A. 2011. A survey of hierarchical classification across different application domains. *Data Mining and Knowledge Discovery* 22(1-2):31–72.
- Tifrea, A.; Bécigneul, G.; and Ganea, O.-E. 2019. Poincaré glove: Hyperbolic word embeddings. In *Proceedings of the Seventh International Conference on Learning Representations*.
- Wang, S., and Jiang, J. 2016. Learning natural language inference with lstm. In *Proceedings of NAACL-HLT*, 1442–1451.
- Wang, G.; Li, C.; Wang, W.; Zhang, Y.; Shen, D.; Zhang, X.; Heno, R.; and Carin, L. 2018. Joint embedding of words and labels for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2321–2331.
- Wang, Z.; Hamza, W.; and Florian, R. 2017. Bilateral multi-perspective matching for natural language sentences. In *Proceedings of the Twenty-sixth International Joint Conference on Artificial Intelligence*, 4144–4150. AAAI Press.
- Wang, X.; Zhang, Y.; and Shi, C. 2019. Hyperbolic heterogeneous information network embedding. In *Proceedings of the Thirty-third AAAI Conference on Artificial Intelligence*.
- Wehrmann, J.; Barros, R. C.; Dôres, S. N. d.; and Cerri, R. 2017. Hierarchical multi-label classification with chained neural networks. In *Proceedings of the Symposium on Applied Computing*, 790–795. ACM.
- Wehrmann, J.; Cerri, R.; and Barros, R. 2018. Hierarchical multi-label classification networks. In *Proceedings of the Thirty-fifth International Conference on Machine Learning*, 5075–5084.
- Wilson, B., and Leimeister, M. 2018. Gradient descent in hyperbolic space. *arXiv preprint arXiv:1805.08207*.
- Zhang, W.; Yan, J.; Wang, X.; and Zha, H. 2018. Deep extreme multi-label learning. In *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval*, 100–107. ACM.