

HYPERBOLIC DISTANCE-BASED SPEECH SEPARATION

Darius Petermann¹ and Minje Kim² †

¹Indiana University, Department of Intelligent Systems Engineering, Bloomington, IN, USA 47408

²University of Illinois at Urbana-Champaign, Department of Computer Science, IL, USA 61801

ABSTRACT

In this work, we explore the task of hierarchical distance-based speech separation defined on a hyperbolic manifold. Based on the recent advent of audio-related tasks performed in non-Euclidean spaces, we propose to make use of the Poincaré ball to effectively unveil the inherent hierarchical structure found in complex speaker mixtures. We design two sets of experiments in which the distance-based parent sound classes, namely “near” and “far”, can contain up to two or three speakers (i.e., children) each. We show that our hyperbolic approach is suitable for unveiling hierarchical structure from the problem definition, resulting in improved child-level separation. We further show that a clear correlation emerges between the notion of hyperbolic certainty (i.e., the distance to the ball’s origin) and acoustic semantics such as speaker density, inter-source location, and microphone-to-speaker distance.

Index Terms— distance-based source separation, hyperbolic space, sound hierarchy, speech separation

1. INTRODUCTION

Humans can focus on a sound source of interest in complex acoustic scenes with remarkable ease [1]. Such “selective attention” [2] could be based on contextual and subjective motivations, but other more established ones can improve source separation systems. By using the deep neural network (DNN)-based supervised learning approaches [3] as the framework, recent research has attempted to encompass the selective attention concept by conditioning the model with the auxiliary information, such as text queries [4] or language, gender, and spatial information about the source [5]. In this paper we focus on the spatial cues [6], e.g., speaker positions or directional information, which have been proven to be useful [7, 8, 9].

More recently, the task of distance-based source separation was introduced [10], whose goal is to separate a monoaural mixture x into its distance-based constitutive components, the sum of near sources $x^{(\mathcal{N})} = \sum_{k \in \mathcal{N}} s^{(k)}$ and far sources $x^{(\mathcal{F})} = \sum_{k \in \mathcal{F}} s^{(k)}$, where \mathcal{N} and \mathcal{F} stand for the sets of near and far sources, respectively. The grouping is defined by a distance threshold τ . In [10], the input is restricted to single-channel signals while proposing the speaker’s distances to the microphone as a determining factor in the source separation task, differently from the previous work, such as SpeakerBeam [11], where the spatial information mainly comes from a multi-channel context.

In this work, we propose to formulate a hierarchical version of the distance-based source separation problem. We present a two-level hierarchy, where the system separates the mixture into two par-

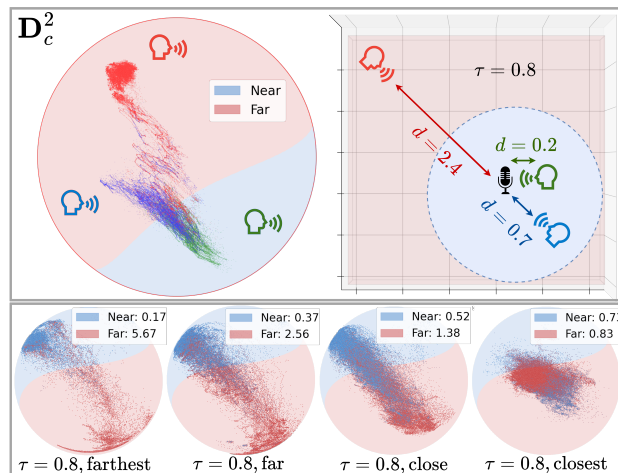


Fig. 1: Illustration of distance-based source separation performed on the Poincaré ball. Top-left is a projection of a three-speaker mixture, where two speakers (green and blue) belong to the near field ($\leq \tau = 0.8$ meter) while one (red) to the far field ($> \tau$). Top-right shows the room configuration of the same mixture. Note that the source at the near-far field boundary tends to be projected at around the center of the Poincaré ball, reflecting the model’s uncertainty. The bottom plots denote projections of two speakers that are progressively placed closer to τ (thus no hierarchy).

ent nodes, $x^{(\mathcal{N})}$ whose source components are closer to the microphone than τ , i.e., $\|\mathbf{m} - \text{loc}(s^{(k)})\|_2 < \tau$, where \mathbf{m} and $\text{loc}(s^{(k)})$ respectively denote the microphone and source locations in the 3D room, and vice versa for the farther source group $x^{(\mathcal{F})}$. In addition, the child-level task aims at separating individual sources $s^{(k)}$.

Inspired by the recent success of hyperbolic manifolds to model hierarchical structures in computer vision [12, 13], natural language processing [14], and more recently audio [15, 16, 17], we propose to utilize a popular hyperbolic model, the Poincaré ball whose embeddings efficiently preserves the intrinsic hierarchical structure of the data in low dimensional space with minimal distortion [18].

The Poincaré ball has also proven to provide reliable and natural measures of certainty estimation [12, 13, 19], which we can interpret as the model’s source-specific confidence level put in the context of audio source separation [15]. In our hierarchical setup, we assume that the geometry between the sources and the microphone introduces uncertainty to the parent-level separation task, i.e., the closer individual sources are to the threshold τ , the more challenging the parent-level separation will be. In addition, the child-level separation faces the challenges that a separation system typically needs to deal with, e.g., same-gender sources being more difficult to separate and therefore inducing additional uncertainty to the separation task. Yet,

*This material is based upon work supported by the National Science Foundation under Grant No. 2046963.

†Work done at Indiana University.

the child-level sources’ geometry also contributes to the uncertainty mapping. Fig. 1 exemplifies this paradigm; the close proximity of the blue source ($d=0.7$) to the threshold ($\tau=0.8$) leads its projections on the Poincaré ball to be closer to the origin, while the green source’s proximity to the microphone ($d=0.2$) mitigates the uncertainty, thus pushing its projections to the edge of the ball. In this case the main cause of the performance degradation in the parent-level separation task is the blue source’s location which puts many of its embeddings on the other side of the decision boundary. Meanwhile, the red source’s farthest location pushes its embeddings away from the origin and consequently the near-far decision boundary as well, signifying the relevance of the geometry to both levels.

Our paper empirically demonstrates two conjectures. First, the hyperbolic model does a better job at hierarchically separating the parent and children sources in terms of scale-invariant signal-to-distortion ratio (SI-SDR) [20] than its Euclidean counterpart. Second, we analytically show that the network can intrinsically map out the hyperbolic certainty on the Poincaré ball to a few acoustic semantics. We envision that the interpretability could benefit geometry-sensitive real-world applications, e.g., disambiguating meeting participants from other interferences based on their physical locations (i.e., inside the same meeting room or not), discerning close and important sound events from too-far-to-be-relevant ones, etc. We further claim that the simplicity of the projection-based conversion to the Poincaré ball makes it an efficient alternative to the sampling or ensemble-based uncertainty quantification method [21].

2. HYPERBOLIC DISTANCE-BASED SPEECH SEPARATION

2.1. Problem Formulation

Following [15], we formulate our source separation task as a mask-inference problem. Given a complex mixture spectrogram $X \in \mathbb{C}^{T \times F}$ with K underlying sources and a separator model $f(\cdot)$, we estimate a K -dimensional mask $M \in \{0, 1\}^{T \times F \times K}$, which is a one-hot vector and whose k -th element of a time-frequency (TF) bin $M_{t,f,k} = 1$ only if the k -th source is the most dominant one in that bin. Likewise, the problem boils down to a K -class classification, where softmax applied to $f(\cdot)$ infers the posterior probability of each TF bin (t, f) belonging to the sources, i.e., $\hat{M}_{t,f,k} = P(k|X_{t,f})$. The mask estimates \hat{M} are then element-wise multiplied with the original mixture X in order to retrieve the k -th spectrogram estimate. The inverse short-time Fourier transform (iSTFT) converts it back to the time domain: $\hat{s}^{(k)} = \text{iSTFT}(X \odot \hat{M}_{:, :, k})$. An established method to obtain these masks is to compute embeddings $Z_{t,f,:} \in \mathbb{R}^L$ of certain dimension L for each of the TF bins, using a neural network-based transformation function, i.e., $Z \leftarrow f(X)$.

In the proposed system, the separator network projects the input to a hyperbolic manifold, i.e., the Poincaré ball, to perform classification on it [15]. To this end, $f(\cdot)$ is repurposed: it originally consists of a bidirectional long short-term memory (BLSTM) stack followed by a linear layer that converts the BLSTM output to Z . As for the hyperbolic counterpart, Z goes through an additional deterministic projection function $H = \exp_0^c(Z_{t,f,:})$ as in Eq. (2), followed by the hyperbolic multinomial logistic regression (HMLR) to substitute softmax. The whole process can be summarized as follows:

$$P(k|X_{t,f}) \approx \begin{cases} \text{softmax}(Z_{t,f,:}) & \text{Euclidean,} \\ \text{HMLR}(H_{t,f,:}) & \text{Hyperbolic.} \end{cases} \quad (1)$$

Given two parent groups $x^{(\mathcal{N})} = \sum_{k \in \mathcal{N}} s^{(k)}$ and $x^{(\mathcal{F})} = \sum_{k \in \mathcal{F}} s^{(k)}$ as well as their respective children $s^{(k \in \mathcal{N})}$ and $s^{(k \in \mathcal{F})}$,

source separation aims to estimate two sets of masks for parent and children levels, separately. The parent-level masks are $\hat{M}_{t,f,0}^{\text{parent}} = P(x^{(\mathcal{N})}|X_{t,f})$ and $\hat{M}_{t,f,1}^{\text{parent}} = P(x^{(\mathcal{F})}|X_{t,f})$, thus $\hat{M}_{t,f,0}^{\text{parent}} + \hat{M}_{t,f,1}^{\text{parent}} = 1$. For the child level, an individual separation task per parent, i.e., $\sum_{k \in \mathcal{N}} \hat{M}_{t,f,k}^{\text{near}} = 1$ and $\sum_{k \in \mathcal{F}} \hat{M}_{t,f,k}^{\text{far}} = 1$, respectively, is defined. We assume independence between the two hierarchical levels, so the *hierarchical softmax* produces level-specific mask estimates for the parents and children independently.

2.2. Hierarchy in the Hyperbolic Space

Hyperbolic spaces have the inherent ability to efficiently represent tree-like structures in a continuous way [22]. Because of the unique geometry induced by these spaces, the distance from the origin grows exponentially (i.e., as a function of the curvature) with the radius, reaching infinity at the ball’s edge.

Let \mathbb{D}_c^n be the n -dimensional hyperbolic space with a Riemannian geometry and negative curvature $c < 0$, which defines a cone-shaped hyperbolic manifold, while the other two isotropic model spaces being Euclidean ($c = 0$) and spherical ($c > 0$). In this paper, we use the Poincaré ball, one of the most popular hyperbolic models for gradient-based learning [12, 19, 13, 15], defined by the manifold and Riemannian metric pair \mathcal{M} and $g_c^{\mathbb{D}}(x) = (\lambda_x^c)^2 g^E$, respectively, where $\lambda_x^c = 2/(1 - c\|x\|^2)$ is a so-called *conformal factor* and g^E the Euclidean metric. The model defines its *induced distance metric* d in relation to the center of the ball: $d_{\mathbb{D}}(a, b) = \cosh^{-1} \left(1 + 2 \frac{\|a-b\|^2}{(1-\|a\|^2)(1-\|b\|^2)} \right)$, where $d_{\mathbb{D}}(a, b)$ defines the shortest distance path, or *geodesic*, between points a and b on the Poincaré ball. In the context of the hierarchical trees, the mean between two leaves in this space would result in a parent rather than another intermediate leaf, and would be located closer to the ball’s origin (i.e., analogous to a continuous hierarchy).

2.3. Hyperbolic Learning

Our architecture operates in both Euclidean and hyperbolic spaces through exponential and logarithmic maps, allowing us to project the embeddings to and from the hyperbolic space, respectively. For embeddings $v \in \mathbb{R}^n \setminus \{0\}$ and $y \in \mathbb{D}_c^n \setminus \{0\}$,

$$\exp_0^c(v) = \frac{\tanh(\sqrt{c}\|v\|)}{\sqrt{c}\|v\|}v, \quad \log_0^c(y) = \frac{\tanh^{-1}(\sqrt{c}\|y\|)}{\sqrt{c}\|y\|}y, \quad (2)$$

where the former represents the deterministic conversion applied to the output of the separator network, i.e., $H_{t,f} = \exp_0^c(Z_{t,f})$, as a mapping from the tangent space to the hyperbolic manifold, while the latter the inverse (i.e., \mathbb{R}^n to \mathbb{D}_c^n and inversely). Note that both projections are centered around their subscript 0. Once the embeddings are projected into the hyperbolic space, the key operations such as vector additions and multiplication follow the rules derived by the Riemannian metric in the Poincaré space [14].

In Euclidean geometry, MLR works with the assumption that logits are the representation of the embedding’s distance to the class hyperplanes. In the hyperbolic space, this assumption can be reformulated from the perspective of margin hyperplanes. For a given class $k \in \{1, \dots, K\}$, one can define a geodesic in \mathbb{D}_c^n orthogonal to a_k and containing p_k . Considering an embedding $Z_{t,f}$ projected onto the ball, the final hyperbolic MLR formula is: $P(k|H_{t,f}) \propto \exp \left(\frac{\lambda_{p_k}^c \|a_k\|}{\sqrt{c}} \sinh^{-1} \left(\frac{2\sqrt{c} | \langle -p_k \oplus_c H_{t,f}, a_k \rangle |}{(1-c\| -p_k \oplus_c H_{t,f} \|^2) \|a_k\|} \right) \right)$, where \oplus_c denotes the Möbius addition in \mathbb{D}_c^n [23]. Note that both p_k and a_k are learnable parameters.

| Density (2 children) | {2,0} | {2,1} | {2,2} | {1,2} | {0,2} |
|----------------------|-------|-------|--------|--------|--------|
| Density (3 children) | {3,0} | {3,1} | {2,2} | {1,3} | {0,3} |
| Training | 3156 | 8217 | 28 059 | 71 152 | 89 416 |
| Validation | 45 | 101 | 321 | 916 | 1117 |
| Testing | 400 | 400 | 400 | 400 | 400 |

Table 1: Distribution of the various speaker density configurations across our two hierarchical use-cases.

3. DATA CURATION AND SIMULATION

To validate our method with various acoustics configurations, we use similar methods as presented in [10]. We first generate random room impulse responses (RIRs) using the Pyroomacoustics toolbox [24]. The room dimensions are uniformly randomized between $3.0 \times 4.0 \times 2.13$ and $7.0 \times 8.0 \times 3.03$ meters. Microphone locations are randomly sampled within the resulting room dimension. In each room, four source locations are independently and randomly sampled from a beta distribution based on their distance from the microphone. The target RT_{60} (in seconds) for each room simulation is randomly sampled in a similar fashion from the range $[0.1, 0.5]$. For the speech sources, we opt to use Libri-light [25] to generate the training set, and Librispeech-clean [26] for validation and testing.

During training, a room is first randomly sampled without replacement for each training example. Six-second-long speech chunks are then randomly sampled from the pool of speech utterances with replacement by zero-padding short utterances. For each training mixture, all four sources are expected to be exploited unless their associated parent class has already reached the maximum number of allowed children. The validation and testing sets are prepared using the same procedure.

4. EXPERIMENTAL SETTINGS

To validate our approach, we design a set of three experiments; the first one solely aims at separating the parent sources $x^{(N)}$ and $x^{(F)}$ (i.e., one-level model) as in [10]. The latter two exploit the two-level models with different hierarchical configurations: one assumes a maximum of two child sources per parent, while the other allows three child sources per parent. Table 1 denotes the five possible speaker setups found in our dataset for each configuration.

Model architecture: Our model follows the same architecture as described in [15], which consists of four BLSTM layers with 600 units in each direction; a dense layer is then used to obtain an Euclidean embedding for each TF bin where $L = 2$. A dropout of rate 0.3 is applied to the output of each BLSTM layer, except the last one. For the hyperbolic model with negative curvature ($c < 0$), an additional exponential projection layer maps the Euclidean embeddings onto the Poincaré ball. As discussed in Section 2.1, either softmax with Euclidean logits or the hyperbolic MLR version then computes the masks for each of the source classes. In practice, we follow the hierarchical softmax approach from [12], and therefore have two MLR layers: one for the binary classification of the parent sources $x^{(N)}$ and $x^{(F)}$, and the second one with $K = 4$ or $K = 6$ for the leaf classes, as dictated by our two possible two-level experiment configurations. We use the mixture phase for resynthesis.

Loss functions: Among the possible spectrogram, waveform, and mask-based loss functions proposed in [15], we opt to use the cross-entropy (CE) loss to compare the ground-truth binary masks and the prediction as in classification setups [3, 27]. In other words, we assume each TF bin embedding to belong to one source only and use a one-hot vector encoding (or IBM in the two-source mixture case) as the target mask on CE. Equations (3) to (5) denote the loss functions

for three cases: the parent level, near sources and far sources in the children level, respectively:

$$\text{Parents: } \mathcal{L}^{\text{parent}} = \frac{1}{TF} \sum_{t,f} \text{CE} \left(M_{t,f,:}^{\text{parent}} \parallel \hat{M}_{t,f,:}^{\text{parent}} \right) \quad (3)$$

$$\text{Near children: } \mathcal{L}^{\text{near}} = \frac{1}{TF} \sum_{t,f} \text{CE} \left(M_{t,f,k \in \mathcal{N}}^{\text{near}} \parallel \hat{M}_{t,f,k \in \mathcal{N}}^{\text{near}} \right) \quad (4)$$

$$\text{Far children: } \mathcal{L}^{\text{far}} = \frac{1}{TF} \sum_{t,f} \text{CE} \left(M_{t,f,k \in \mathcal{F}}^{\text{far}} \parallel \hat{M}_{t,f,k \in \mathcal{F}}^{\text{far}} \right), \quad (5)$$

where the parent-level loss effectively compares the source estimate to the sum of near sources group, e.g., $x^{(N)} \approx \text{iSTFT}(\hat{M}_{t,f,:}^{\text{parent}} \odot X)$, and ditto for the far sources. In contrast, the child-level loss functions are parent-specific, meaning the source separation results are compared to the child sources that belong to the same parent, either near or far. Note that while the order of the parent source groups is assumed, permutation invariant training (PIT) [28] for the child-level is required as each parents’ child sources are in no particular order. Hence, the $\mathcal{L}^{\text{parent}}$ works as the sole loss function for parent-only source separation task as in [10], i.e., the one-level model. Meanwhile, our two-level hierarchical model is trained from the compound of all loss terms: $\mathcal{L} = \mathcal{L}^{\text{parent}} + \mathcal{L}^{\text{near}} + \mathcal{L}^{\text{far}}$.

Training: All experiments use the Adam optimizer for the Euclidean parameters, and the Riemannian Adam [29] implementation from *geoopt* [30] for the hyperbolic parameters. All models are trained using chunks of 6.0 s, batch size of 96 over a total of 200 and 300 epochs for the one-level and two-level models, respectively. We use an initial learning rate of 10^{-3} , which is halved if the validation loss does not improve for 10 epochs. We set the STFT size to 32-ms with a 50% overlap and the square-root Hann window.

Evaluation: Our main focus is to show the predictive uncertainty naturally conveyed by the hyperbolic space and how this information could potentially be exploited on downstream systems. We report the validity of the proposed method in terms of SI-SDR improvement (SI-SDRi) [20] and use the best permutation at the child-level. If one of the parents is silent (i.e., all speakers belong to either the “near” or “far” field), we evaluate the predicted silent sources by how much signal was bled into it. For that reason, we report the loudness ratio between the input mixture and the predicted silent source, called *noise reduction*, as in [10] (i.e., the larger the better).

5. EXPERIMENTAL RESULTS

5.1. Uncertainty Analysis

Our experiments explore how well the Poincaré ball can model the task of hierarchical distance-based source separation. Precisely, we investigate to what extent the hyperbolic space can tie the intrinsic notion of certainty to the acoustic configurations, such as source-to-microphone distance, the relative distance among speakers, and speaker density, with a fixed distance threshold $\tau = 0.8$ meter.

Speaker density: We refer to speaker density as the relation between the number of speakers present in the parent classes “near” and “far”. For example, our two-level hierarchy (see Table 1), can have a “dense” setup with up to two child sources on each parent side ($\{2, 2\}$), while its “moderate” ($\{2, 1\}$, $\{1, 2\}$), and “sparse” ($\{2, 0\}$, $\{0, 2\}$) setups are also considered. Fig 2a shows the embeddings’ L_2 norm distribution across all three configurations. We note that when all children belong to one parent, the model showcases high level of certainty, with most of the projected embeddings located far from the origin.

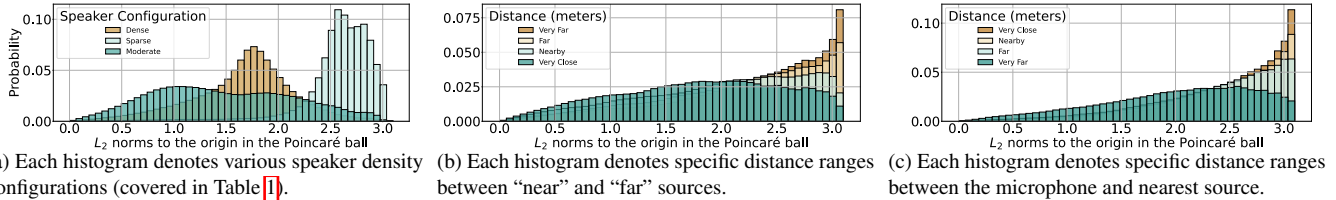


Fig. 2: Distributions of L_2 norms to the origin of the Poincaré ball for all embeddings in their dedicated testing-set, as a function of various acoustics paradigms, such as speakers configuration (left), sources (center), and microphone (right) distances.

| c | SI-SDRi | | Noise Reduction |
|---------|------------|------------|-----------------|
| | Near | Far | Near |
| No Proc | 3.1 | -3.1 | — |
| 0.0 | 6.5 | 9.1 | 27.1 |
| -0.1 | 6.3 | 9.0 | 27.1 |
| -1.0 | 6.3 | 9.2 | 25.8 |

Table 2: SI-SDR improvement (dB) comparisons on models involving different curvatures (i.e., $c = 0.0$ for Euclidean and $c < 0$ for hyperbolic) on a parents-only training setup. Power reduction is reported when “near” is silent (the larger the better).

Distance among sources: To observe how the relative distance among the sources can affect the hyperbolic certainty in the speakers’ embeddings, we design a small test set with one child per parent, i.e., $x^{(N)} = s^{(1)}$ and $x^{(F)} = s^{(2)}$. These sources are equidistant from the threshold τ , meaning $d(s^{(2)}) - \tau = \tau - d(s^{(1)})$, where $d(\cdot)$ denotes the distance from the microphone. We then vary $d(s^{(1)})$ and $d(s^{(2)})$ such that the two sources’ relative distance $d(s^{(2)}) - d(s^{(1)})$ increases from 0, i.e., $d(s^{(2)}) = d(s^{(1)}) = \tau$, to 1.4 meters, i.e., $d(s^{(2)}) = 1.5$ and $d(s^{(1)}) = 0.1$. Fig. 2b shows the observed distribution of speakers embeddings’ L_2 norms as a function of their relative distance. We note an overall decrease in certainty as their relative distance decreases, and vice versa.

Distance from the microphone: We also consider the sources’ distance from the microphone by designing another one-child-per-parent test set with the far source away from the microphone with a fixed distance $d(s^{(2)}) = 2.9$, while the near source $s^{(1)}$ varies its distance to the microphone in the range of $[0.2, 0.8]$. From Fig 2c, we observe that the closer $s^{(1)}$ is to the microphone, the higher the certainty is observed from the hyperbolic embeddings.

5.2. Separation performances

Table 2 reports the results in terms of SI-SDR improvement from our one-level model (i.e., the parent-only case), where no performance gain is induced by the hyperbolic setup. We speculate that it is due to the lack of hierarchy in the problem setup, negating the benefit from the hyperbolic manifold. However, we emphasize that as shown in Fig. 1 the notion of certainty in the proposed hyperbolic approach can provide an additional user interface, identifying borderline cases.

Table 3 shows the results from the two-level setups. We report the performance of both Euclidean ($c = 0.0$) and hyperbolic models ($c \in [-0.1, -1.0]$). In the top rows (two children sources per parent), the parental separation results are not affected by our hyperbolic setup. However, we observe clearer improvement in the child-level results using the hyperbolic method, with an average improvement of 0.2 dB SI-SDR over their Euclidean counterpart. The more complex two-level setup with up to three child sources per parent (bottom rows) shows a solid 0.5 dB SI-SDR improvement with

| Speaker Density | Parents | | | Children | | |
|-------------------|--------------------------|-------------------|-------------------|----------|------------|------------|
| | $c = 0.0$ | -0.1 | -1.0 | 0.0 | -0.1 | -1.0 |
| 2 near, 0 far | 31.8 [†] | 29.8 [†] | 28.7 [†] | 8.4 | 8.4 | 8.4 |
| 2 near, 1 far | 1.5 | 1.5 | 1.5 | 5.9 | 6.1 | 6.3 |
| 2 near, 2 far | 2.4 | 2.3 | 2.3 | 4.0 | 4.1 | 4.1 |
| 1 near, 2 far | 6.7 | 6.7 | 6.8 | 7.1 | 7.5 | 7.1 |
| 0 near, 2 far | 63.3 [†] | 61.3 [†] | 63.2 [†] | 8.9 | 9.2 | 9.2 |
| Average (SI-SDRi) | 3.5 | 3.5 | 3.5 | 6.8 | 7.0 | 7.0 |
| 3 near, 0 far | 14.2 [†] | 19.5 [†] | 15.7 [†] | 5.7 | 4.9 | 6.0 |
| 3 near, 1 far | 0.6 | 0.8 | 0.7 | 4.2 | 4.0 | 4.4 |
| 2 near, 2 far | 2.5 | 2.5 | 2.6 | 4.6 | 4.6 | 4.9 |
| 1 near, 3 far | 4.9 | 4.7 | 5.0 | 5.0 | 5.1 | 5.4 |
| 0 near, 3 far | 49.2 [†] | 42.6 [†] | 41.6 [†] | 5.9 | 6.9 | 7.0 |
| Average (SI-SDRi) | 2.6 | 2.7 | 2.8 | 5.1 | 5.1 | 5.6 |

Table 3: SI-SDR improvement (dB) of the models involving different curvatures (i.e., $c = 0.0$ for Euclidean and $c < 0$ for hyperbolic) on our 2-level setup with up to two and three child sources per parent (top and bottom rows, respectively). Noise Reduction, when one of the parent is silent, is indicated by [†] (the larger the better).

$c = -1.0$. We believe that the children benefit the most from the hyperbolic treatment due to the intrinsic nature of the space to efficiently model the hierarchical structure modelled by the data. This could also explain the greater gain when more children are present.

This relatively small performance margin from the hyperbolic method aligns well with the literature [12, 13, 15]. In this work, we rather bring the readers’ attention to the explanatory and intuitive nature of the hyperbolic representation. For example, obtaining a natural measure of uncertainty predictions for these hierarchical tasks could help tackle the typical trade-off between sensitivity and specificity: one can still choose to acquire the closer source group if the system is uncertain about the child-level source separation results.

6. CONCLUSIONS

In this work, we redefined distance-based single-channel speech separation as a hierarchical task performed on an hyperbolic manifold (i.e., the Poincaré ball). We designed a two-level hierarchy, whose parent level distinguishes sources based on their distance (i.e., near and far) to the microphone, while each parent may contain up to two or three speakers (i.e., children). We empirically showed that the hyperbolic space consistently outperformed our Euclidean baseline when the task was challenging enough with more child-level sources. In addition, the network was successful in modeling some key acoustic concepts such as speaker density, the relative distance between sources, or microphones, and translating them to the notion of hyperbolic certainty which could be effectively used towards practical downstream tasks. Code will be made available [1].

¹<https://minjekim.com/research-projects/hdss>

7. REFERENCES

- [1] J. H. McDermott, “The cocktail party problem,” *Current Biology*, vol. 19, no. 22, pp. R1024–R1027, 2009.
- [2] C. M. Karns, E. Isbell, R. J. Giuliano, and H. J. Neville, “Auditory attention in childhood and adolescence: An event-related potential study of spatial selective attention to one of two simultaneous stories,” *Developmental Cognitive Neuroscience*, vol. 13, pp. 53 – 67, 2015.
- [3] D. L. Wang and J. Chen, “Supervised Speech Separation Based on Deep Learning: An Overview,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 10, pp. 1702–1726, 2018.
- [4] X. Liu *et al.*, “Separate What You Describe: Language-Queried Audio Source Separation,” in *Proc. Interspeech*, 2022.
- [5] E. Tzinis *et al.*, “Heterogeneous Target Speech Separation,” in *Proc. Interspeech*, 2022, pp. 1796–1800.
- [6] C. J. Darwin and R. W. Hukin, “Effectiveness of spatial cues, prosody, and talker characteristics in selective attention,” *The Journal of the Acoustical Society of America*, vol. 107, no. 2, pp. 970–977, 2000.
- [7] G. Rongzhi *et al.*, “Neural spatial filter: Target speaker speech separation assisted with directional information,” in *Proc. Interspeech*, 09 2019, pp. 4290–4294.
- [8] J. Heitkaemper, T. Fehér, M. Freitag, and R. Haeb-Umbach, “A study on online source extraction in the presence of changing speaker positions,” in *Statistical Language and Speech Processing*, C. Martín-Vide, M. Purver, and S. Pollak, Eds. Springer International Publishing, 2019, pp. 198–209.
- [9] S. E. Eskimez *et al.*, “Personalized Speech Enhancement: New Models and Comprehensive Evaluation,” in *Proc. of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2022, pp. 356–360.
- [10] K. Patterson, K. W. Wilson, S. Wisdom, and J. R. Hershey, “Distance-based sound separation,” in *Interspeech 2022, 23rd Annual Conference of the International Speech Communication Association, Incheon, Korea, 18-22 September 2022*, H. Ko and J. H. L. Hansen, Eds. ISCA, 2022, pp. 901–905.
- [11] M. Delcroix *et al.*, “Single Channel Target Speaker Extraction and Recognition with Speaker Beam,” in *Proc. of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2018, pp. 5554–5558.
- [12] M. G. Atigh *et al.*, “Hyperbolic image segmentation,” in *Proc. of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2022.
- [13] V. Khrukov *et al.*, “Hyperbolic image embeddings,” in *Proc. of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [14] R. Shimizu, Y. Mukuta, and T. Harada, “Hyperbolic neural networks++,” in *International Conference on Learning Representations*, 2021.
- [15] D. Petermann, G. Wichern, A. Subramanian, and J. Le Roux, “Hyperbolic audio source separation,” in *Proc. of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Jun. 2023.
- [16] F. Nakashima *et al.*, “Hyperbolic timbre embedding for musical instrument sound synthesis based on variational autoencoders,” in *Proc. APSIPA ASC*, 2022.
- [17] F. Germain, G. Wichern, and J. Le Roux, “Hyperbolic unsupervised anomalous sound detection,” in *Proc. of the IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2023.
- [18] R. Sarkar, “Low distortion delaunay embedding of trees in hyperbolic plane,” in *Proc. International Symposium On Graph Drawing*, 2012, pp. 355–366.
- [19] D. Surís, R. Liu, and C. Vondrick, “Learning the predictability of the future,” in *Proc. of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [20] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, “SDR – half-baked or well done?” in *Proc. of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2019.
- [21] H. Fang, D. Becker, S. Wermter, and T. Gerkmann, “Integrating uncertainty into neural network-based speech enhancement,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 1587–1600, 2023.
- [22] M. Nickel and D. Kiela, “Poincaré embeddings for learning hierarchical representations,” in *Advances in Neural Information Processing Systems*, I. Guyon *et al.*, Eds., vol. 30. Curran Associates, Inc., 2017.
- [23] W. Chen *et al.*, “Fully hyperbolic neural networks,” in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Dublin, Ireland: Association for Computational Linguistics, May 2022, pp. 5672–5686.
- [24] R. Scheibler, E. Bezzam, and I. Dokmanić, “Pyroomacoustics: A python package for audio room simulation and array processing algorithms,” in *Proc. of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2018, p. 351–355.
- [25] J. Kahn *et al.*, “Libri-light: A benchmark for asr with limited or no supervision,” in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7669–7673.
- [26] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An ASR corpus based on public domain audio books,” in *Proc. of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.
- [27] A. Subramanian, S. J. Chen, and S. Watanabe, “Student-teacher learning for blstm mask-based speech enhancement,” in *Proc. Interspeech*, 09 2018, pp. 3249–3253.
- [28] M. Kolbæk, D. Yu, Z. H. Tan, and J. Jensen, “Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 10, pp. 1901–1913, 2017.
- [29] H. Kasai, P. Jawanpuria, and B. Mishra, “Riemannian adaptive stochastic gradient algorithms on matrix manifolds,” in *Proc. of the International Conference on Machine Learning (ICML)*, Jun. 2019.
- [30] M. Kochurov, R. Karimov, and S. Kozlukov, “Geopt: Riemannian optimization in PyTorch,” *arXiv preprint arXiv:2005.02819*, 2020.