

Hyperbolic Audio-visual Zero-shot Learning

Jie Hong^{1,2}, Zeeshan Hayder², Junlin Han^{1,2}, Pengfei Fang^{3*},
Mehrtash Harandi⁴, Lars Petersson²

¹Australian National University, ²Data61-CSIRO, ³Southeast University, ⁴Monash University

jie.hong@anu.edu.au, zeeshan.hayder@data61.csiro.au, junlinhcv@gmail.com,

fangpengfei@seu.edu.cn, mehrtash.harandi@monash.edu, lars.petersson@data61.csiro.au

Abstract

Audio-visual zero-shot learning aims to classify samples consisting of a pair of corresponding audio and video sequences from classes that are not present during training. An analysis of the audio-visual data reveals a large degree of hyperbolicity, indicating the potential benefit of using a hyperbolic transformation to achieve curvature-aware geometric learning, with the aim of exploring more complex hierarchical data structures for this task. The proposed approach employs a novel loss function that incorporates cross-modality alignment between video and audio features in the hyperbolic space. Additionally, we explore the use of multiple adaptive curvatures for hyperbolic projections. The experimental results on this very challenging task demonstrate that our proposed hyperbolic approach for zero-shot learning outperforms the SOTA method on three datasets: VGGSound-GZSL, UCF-GZSL, and ActivityNet-GZSL achieving a harmonic mean (HM) improvement of around 3.0%, 7.0%, and 5.3%, respectively.

1. Introduction

Visual and audio signals frequently co-occur. For example, movies combine visual and auditory signals, providing an immersive experience to viewers. Humans perceive multiple sensory inputs jointly and make decisions accordingly. When a vehicle approaches from behind and honks, the driver sees the vehicle in the rear-view mirror, hears the sound, and decides to give way. The integration of joint visual and audio signals benefits numerous applications. For instance, acoustic features can help localize objects that emit sound in videos [30, 58, 52, 45, 13, 46, 44]. Furthermore, the natural correlations between visual and audio signals provide strong supervision for learning video representations. As such, there has been growing interest in learning informative representations from audio sig-

nals for video classification [38, 39, 3, 2, 11, 29, 41]. The benefits of audio-visual multi-modality learning have also been demonstrated in other tasks, such as robotic navigation [9, 18, 12], action recognition [27, 20], highlight detection [54, 5], violence detection [51], aerial scene recognition [26] and speech recognition [53, 1].

Collecting vast amounts of audio-visual data needed for training deep neural networks is time-intensive, expensive, and in some applications, impractical. Additionally, deep models may face difficulties in making accurate predictions when presented with objects from unseen classes in real-world scenarios. This is because they lack the necessary knowledge and context to make well-informed decisions about unfamiliar objects. Low-shot audio-visual tasks [42] have emerged to address the problem of insufficient data, with an aim to generalize models effectively from seen to unseen data. These tasks include one-shot learning [50], few-shot learning [34], and audio-visual zero-shot learning [37, 36, 35, 42]. One-shot and few-shot audio-visual learning deals with classification from a few examples of unseen classes. However, audio-visual zero-shot classification poses a more challenging scenario as the model has no access to audio-visual data from unseen classes during training [42].

In this paper, we focus on audio-visual zero-shot learning. We investigate curvature-aware geometric learning for the audio-visual feature alignment. Our inspiration for using hyperbolic geometry comes from the following observations:

- **Data hierarchy.** Audio-visual datasets exhibit a hierarchy. For instance, in VGGSound [10], all 309 classes can be categorized into 9 parent classes (or super-classes): “people”, “animals”, “music”, “sports”, “nature”, “vehicle”, “home”, “tools” and “others”. Similarly, the dataset ActivityNet [6] provides a rich hierarchy with at least four levels. For example, the class “Hand washing clothes” belongs to “Laundry” (the 2nd level), “Housework” (the 3rd level), and “Household Activities” (the 4th level). Most existing audio-

*Corresponding author

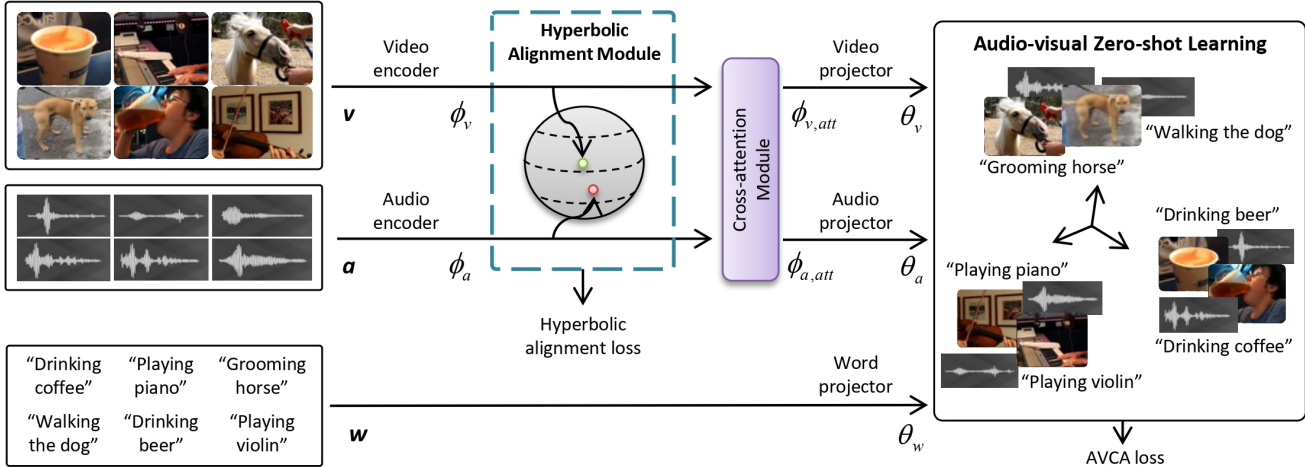


Figure 1. We introduce the *Hyperbolic Alignment Module*, represented by the block with the blue line. The input data features from each modality are encoded by two consecutive networks (encoder and projector). A cross-attention module in between is used to explore the natural correspondence between visual and audio features. Before the cross-attention module processes the features, our *Hyperbolic Alignment Module* computes a hyperbolic alignment loss, which aims to explore more hierarchy in audio-visual data. For example, the model may discriminate embeddings of “Playing piano” and “Walking the dog” when it finds that these embeddings belong to different superclasses: “Playing musical instruments” and “Walking/exercising/playing with animals”.

visual works have not adequately leveraged the hierarchical structure present in the audio-visual data.

- **Hyperbolic geometric properties.** Hyperbolic methods have been shown to be effective in addressing low-shot visual problems [28, 31, 16, 14, 33]. The learned hyperbolic feature embeddings have the ability to capture the hierarchical structure within the data. This is attributed to the tree-like nature of the underlying space, as shown in [24, 7]. One of the benefits of using a hyperbolic space is that the hyperbolic space facilitates the distribution of embeddings in a tree-shaped structure since its volume expands exponentially.

Based on the observations above, we conjecture that the unique properties of hyperbolic spaces can be leveraged to capture the hierarchical structures in audio-visual data, leading to learning more discriminative embeddings for audio-visual samples. The current SOTA audio-visual zero-shot learning methods [37, 36, 35, 42, 57] operate in the non-curved Euclidean space without considering the data hierarchy. Therefore, there is a need for curvature-aware geometric solutions that can embed the data hierarchy to improve the performance of audio-visual zero-shot learning. The contributions of this work can be summarized as:

- Our work provides a new perspective on using curved geometries for cross-modality, as shown in Figure 1. We propose a hyperbolic alignment loss that learns features in curved space to improve audio-visual zero-shot learning. Specifically, we use alignment between visual and audio features in the hyperbolic space as

an auxiliary learning method for feature fusion across modalities. To the best of our knowledge, we are the first to apply curvature-aware geometric solutions to this task.

- Furthermore, we introduce various frameworks for using the hyperbolic embeddings: 1) Hyper-alignment, 2) Hyper-single, and 3) Hyper-multiple. The Hyper-alignment module maps audio and visual features from Euclidean to hyperbolic space with a fixed negative curvature and compares them using intra-modal similarities. Based on Hyper-alignment, Hyper-single adapts the curvature to the model for flexible data structure exploration. Hyper-multiple generates a set of adaptive curvatures for alignments, enabling more generic embeddings.
- Extensive experiments demonstrate that, in most cases, the proposed modules outperform existing models. Moreover, using the δ_{rel} metric [28, 17], we show that hyperbolic alignment enables the learned features to exhibit stronger hierarchical properties. Additionally, we observe from t-SNE [49] visualizations that the audio-visual feature embeddings from different superclasses become more distinct.
- Ablation studies are provided further to investigate the properties of our hyperbolic alignment module. In addition to the curvature-negative hyperbolic projection approach, we also test Euclidean and Spherical approaches, which have curvature-zero and curvature-positive properties, respectively.

2. Related Works

2.1. Audio-visual Zero-shot Learning

The audio-visual zero-shot learning settings are first proposed in [42], where the Coordinated Joint Multimodal Embedding (CJME) model is introduced to map video, audio, and text into the same feature space and compare them. The triplet loss was used to push video or audio features closer to their corresponding class features. Subsequently, Mazumder *et al.* [35] develop the Audio-Visual Generalized Zero-shot Learning Network (AVGZSLNet) to address audio-visual zero-shot learning, which includes a module that reconstructs text features from visual and audio features. The Audio-Visual Cross-Attention (AVCA) framework [37] is also designed to exchange information between video and audio representations, enabling informative representations help achieve state-of-the-art performance in audio-visual zero-shot classification. A similar design to AVCA for handling temporal data features is presented in [36].

2.2. Hyperbolic Geometry

Hyperbolic geometry has gained considerable interest in the machine learning community since its property of the negative curvature can encode the inherent hierarchical structure of the data. It benefits from a family of learning scenarios [43], including sentiment analysis, recommendation systems, social network studies, etc. Ganea *et al.* first study the necessity to integrate hyperbolic geometry in the learning community due to its intriguing property to encode the hierarchical structure of the data and develop some necessary neural operators in the Poincaré ball [19]. Its superior property is also leveraged to analyze and understand the irregular graph data [8]. Joint efforts are made to develop hyperbolic learning algorithms in the NLP and graph neural network field, *e.g.*, hyperbolic attention network (HAN) [23], hyperbolic graph attention network (HGAN) [56] *etc.* Khrlukov *et al.* analyze that the hierarchical structure also exists in the visual data and successfully integrates the hyperbolic space in the visual domain [28]. The concurrent work in [31] also shows that hyperbolic space is a good alternative for aligning the visual and text domains. The following works prove that hyperbolic geometry can benefit a series of visual tasks [15], like semantic segmentation [4], medical image recognition [55], action recognition [32], anomaly recognition [25] *etc.*

In this paper, we first learn hyperbolic embeddings for both visual and audio modalities. We then introduce a novel hyperbolic alignment loss that minimizes the differences between different modalities. This is different from existing approaches [37, 36, 35, 42, 57] that focus solely on cross-modal correspondence between visual and audio streams. In contrast, our module considers cross-modal alignment in

curvature-negative space (hyperbolic space) for mining hierarchical structures in audio-visual data.

3. Preliminary

To construct the hyperbolic alignment loss, we project the embedding point from Euclidean onto the hyperbolic tangent space. This involves two procedures: the hyperbolic projection and the logarithmic map.

Hyperbolic projection. In this paper, we utilize the Poincaré ball [40, 28] to model the hyperbolic space. The Poincaré ball can be visualized as a ball space with a radius of $r = \sqrt{1/|c|}$, where $c < 0$ represents the curvature. To project a point $\mathbf{x} \in \mathbb{R}^n$ from the Euclidean space onto the n -dimensional Poincaré ball \mathbb{H}_c^n with curvature c , we perform the hyperbolic projection as follows:

$$\mathbf{x}_H = \Gamma_{\mathbb{H}}(\mathbf{x}) = \begin{cases} \mathbf{x} & \text{if } \|\mathbf{x}\| \leq \frac{1}{|c|} \\ \frac{1-\xi}{|c|} \frac{\mathbf{x}}{\|\mathbf{x}\|} & \text{else} \end{cases} \quad (1)$$

where the projected point in the Poincaré ball is denoted by $\mathbf{x}_H \in \mathbb{H}_c^n$. To ensure numerical stability, a small value $\xi > 0$ is introduced. The addition of two points, \mathbf{x}_H and $\mathbf{y}_H \in \mathbb{H}_c^n$, in the Poincaré ball can be obtained via the Möbius addition [28] as follows:

$$\mathbf{x}_H \oplus_c \mathbf{y}_H = \frac{(1+2|c|\langle \mathbf{x}_H, \mathbf{y}_H \rangle + |c|\|\mathbf{y}_H\|^2)\mathbf{x}_H + (1-|c|\|\mathbf{x}_H\|^2)\mathbf{y}_H}{1+2|c|\langle \mathbf{x}_H, \mathbf{y}_H \rangle + |c|^2\|\mathbf{x}_H\|^2\|\mathbf{y}_H\|^2} \quad (2)$$

where $\langle \cdot, \cdot \rangle$ is the inner product.

Logarithmic map. The hyperbolic tangent space is an Euclidean space that locally approximates the hyperbolic space. By selecting a tangent point $\mathbf{z}_H \in \mathbb{H}_c^n$, we can generate a tangent plane $T_{\mathbf{z}_H}\mathbb{H}_c^n$ at \mathbf{z}_H . The process of taking the logarithm is used to map a point $\mathbf{x}_H \in \mathbb{H}_c^n$ onto the tangent space $T_{\mathbf{z}_H}\mathbb{H}_c^n$ of \mathbf{z}_H , which is given as follows:

$$\mathbf{x}_{Tg} = \frac{2}{\sqrt{|c|\lambda_c(\mathbf{z}_H)}} \tanh^{-1}(\sqrt{|c|}\|\mathbf{x}_H - \mathbf{z}_H \oplus_c \mathbf{x}_H\|) \frac{-\mathbf{z}_H \oplus_c \mathbf{x}_H}{\|-\mathbf{z}_H \oplus_c \mathbf{x}_H\|} \quad (3)$$

where $\mathbf{x}_{Tg} \in T_{\mathbf{z}_H}\mathbb{H}_c^n$ is the transformed point in the tangent space. In this paper, we consider the identity tangent space of the Poincaré ball at the origin $\mathbf{0}_H \in \mathbb{H}_c^n$, so we have $\mathbf{z}_H = \mathbf{0}_H$.

4. Hyperbolic Audio-visual Learning

Audio-visual zero-shot learning aims to recognize audio-visual samples $\mathbf{u} = (\mathbf{v}_u, \mathbf{a}_u)$ from unseen classes during the evaluation stage [42]. In this task, the training data is restricted to the seen classes. In line with [37], we adopt the generalized setting where both seen and unseen class samples, $\mathbf{s} = (\mathbf{v}_s, \mathbf{a}_s)$ and \mathbf{u} , appear at the test stage. As illustrated in Figure 1, the corresponding semantic words \mathbf{w} are provided along with the audio-visual sample (\mathbf{v}, \mathbf{a}) to serve as textual labels.

In this section, we first introduce the baseline AVCA [37] and then provide details of our approach using a hyperbolic alignment loss for audio-visual zero-shot learning. As depicted in Figure 1, we embed the network with a *Hyperbolic Alignment Module* for computing the hyperbolic alignment loss. We propose three designs of this module: *Hyper-alignment*, *Hyper-single*, and *Hyper-multiple*, as shown in Figure 2 (a), (b) and (c). The subsequent subsections provide more details on each of the proposed designs.

4.1. Baseline

Our method is based on AVCA [37], which serves as the baseline. AVCA is a two-branch network that operates on synchronized visual and audio inputs, as shown in Figure 1. These inputs, denoted by $v, a \in \mathbb{R}^{512}$ respectively, are represented by feature vectors extracted from pre-trained feature extractors, with both vectors having a dimension of 512. In addition to the video clip feature v and corresponding audio feature a , the baseline model takes the text feature $w \in \mathbb{R}^{512}$ as input. The features $\phi_v, \phi_a \in \mathbb{R}^{300}$ are obtained after passing v and a through the video and audio encoder, respectively. With ϕ_v and ϕ_a , a cross-attention module calculates the cross-modality representations $\phi_{v,att}, \phi_{a,att} \in \mathbb{R}^{300}$. These features are then passed through the video and audio projectors to compute the latent embeddings $\theta_v, \theta_a \in \mathbb{R}^{64}$ in the common space for classifying audio-visual inputs (v, a) . The text feature w is processed by the word projector to give $\theta_w \in \mathbb{R}^{64}$ in the common space. The loss of AVCA is represented by \mathcal{L}_{avca} .

4.2. Hyper-alignment

We propose a minimalist design called *Hyper-alignment*, which learns feature alignment in a hyperbolic manifold using a fixed curvature, as illustrated in Figure 2 (a). To construct the loss, both modality features ϕ_v and ϕ_a are first mapped to points $\phi_{v,H}$ and $\phi_{a,H}$ in the Poincaré ball space using Eq. (1). Then, we obtain $\phi_{v,Tg}$ and $\phi_{a,Tg}$ in $T_z\mathbb{H}_c^{300}$ by projecting $\phi_{v,H}$ and $\phi_{a,H}$ from the hyperbolic to tangent space, using Eq. (3).

The alignment loss plays a crucial role in guiding the network to learn informative modal representations. By using this loss, we hope to encourage the audio-visual features ϕ_v and ϕ_a to reflect more hierarchical structures. To achieve this, we transform these features into a curvature-negative space before performing the alignment operation. Additionally, since it is unclear which of the modalities contributes more to hierarchical exploration, our model ensures alignment between features by minimizing the difference in similarities within each modality. Motivated by the approach of [48], we design the loss using pairwise feature similarities within each of the mini-batch. The hyperbolic

alignment loss is written as follows:

$$\mathcal{L}_{align} = \frac{1}{N^2} \|\mathbf{M}_{v,norm} - \mathbf{M}_{a,norm}\|_F^2 \quad (4)$$

where $\mathbf{M}_{v,norm}, \mathbf{M}_{a,norm} \in \mathbb{R}^{N \times N}$ are the normalized matrices of intra-modal similarity matrices, $\mathbf{M}_v, \mathbf{M}_a \in \mathbb{R}^{N \times N}$, for a batch of video or audio features, $\phi_{v,Tg}$ or $\phi_{a,Tg}$, in the hyperbolic tangent space. Hence, we have $\mathbf{M}_{norm} = \frac{\mathbf{M}}{\|\mathbf{M}\|_F}$. The batch size is N and $\|\cdot\|_F$ calculates the Frobenius norm.

Each element in \mathbf{M} is computed via the cosine similarity between the i^{th} and j^{th} feature on the tangent space in a batch: $m_{i,j} = \cos(\phi_{Tg,i}, \phi_{Tg,j}), i, j \in \{1, 2, \dots, N\}$. Essentially, the loss \mathcal{L}_{align} in Eq. (4) computes the mean element-wise squared difference between two similarity matrices. It aligns features from two modalities by minimizing the distance between two within-modal similarities. Finally, the overall loss of Hyper-alignment is comprised of two losses:

$$\mathcal{L}_{total} = \mathcal{L}_{avca} + \mathcal{L}_{align} \quad (5)$$

where \mathcal{L}_{avca} is the loss of the baseline. It is noted Hyper-alignment does not bring any extra weight to the network.

4.3. Hyper-single

We propose Hyper-alignment to align features in a constant-curvature space. However, the representation ability of audio-visual features might suffer as the fixed curvature may not be suitable for the complex audio-visual data structure [21]. To address this, we devise *Hyper-single*, which utilizes an adaptive curvature. Unlike Hyper-alignment, which conforms to a fixed hyperbolic structure, Hyper-single produces a learnable curvature that enables flexible exploration of intrinsic data structures.

Particularly, as illustrated in Figure 2 (b), the modality features, ϕ_v and ϕ_a are first concatenated into one vector $\phi_{va} \in \mathbb{R}^{600}$, and ϕ_{va} passes through one multiple-layer perceptron (MLP) to compute the curvature. Then, we have the curvature adaptively learned as follows:

$$c = c_0 \cdot \text{sigmoid}(\text{MLP}(\phi_{va})) \quad (6)$$

where $\text{MLP}(\cdot) : \mathbb{R}^{600} \rightarrow \mathbb{R}$ is one fully-connected layer and $c_0 < 0$ is the initial curvature. Hyper-single shares the same loss function in Eq. (5) as Hyper-alignment. Compared to Hyper-alignment, Hyper-single introduces the extra weights in MLP to the network.

4.4. Hyper-multiple

For the design of *Hyper-multiple*, as depicted in Figure 2 (c), we first attain several curvatures. Then, multiple Hyper-alignment modules are used for learning the alignment losses via the obtained curvatures. The multiple-curvature method leads to more generic spaces since diverse manifold structures may coexist in audio-visual data [22].

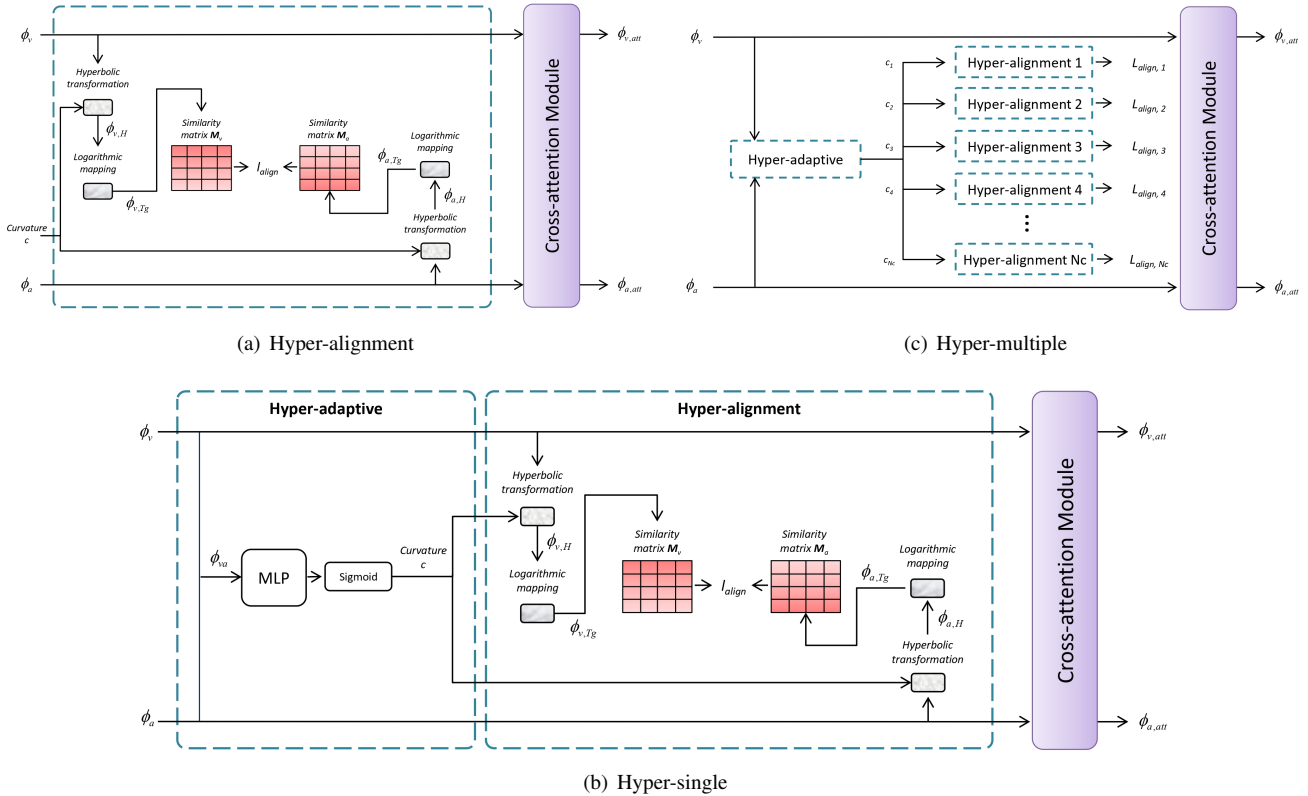


Figure 2. The frameworks of the proposed *Hyperbolic Alignment Module*: (a) *Hyper-alignment*, which utilizes the hyperbolic space to learn feature alignment, (b) *Hyper-single*, which automatically computes the curvature using the Hyper-adaptive module after concatenating visual and audio features as input, and (c) *Hyper-multiple*, which performs alignment using multiple learnable curvatures. While Hyper-alignment uses a fixed curvature, Hyper-single and Hyper-multiple enable flexible exploration of intrinsic data structures by adapting the curvature. Hyper-multiple is particularly suited when diverse manifold structures coexist, improving the generality of the learned representations.

Hyper-multiple generates a set of adaptive curvatures $\mathbf{c} = \{c_1, c_2, \dots, c_{N_c}\}$ and synchronously map ϕ_v and ϕ_a into N_c hyperbolic tangent spaces with these curvatures. Similar to Hyper-single, we make use of $\text{MLP}(\cdot) : \mathbb{R}^{600} \rightarrow \mathbb{R}^{N_c}$ to output curvatures \mathbf{c} . The total loss in Eq. (5) applies Hyper-multiple, where the alignment loss \mathcal{L}_{align} becomes:

$$\mathcal{L}_{align} = \frac{1}{N_c} \sum_i^{N_c} \mathcal{L}_{align,i} \quad (7)$$

where $\mathcal{L}_{align,i}$ is the hyperbolic alignment loss given in Eq. (4) with the i^{th} curvature, $i \in \{1, 2, \dots, N_c\}$. Instead of restricting to only single curvature, Hyper-multiple computes multiple appropriate curvatures to learn more generic embeddings.

5. Experiments

In this section, we provide extensive experiments to verify the effectiveness of our method. The proposed *Hyper-*

bolistic Alignment Module is evaluated in audio-visual zero-shot classification on three datasets: VGGSound-GZSL, UCF-GZSL, and ActivityNet-GZSL. The details of each dataset are as follows.

VGGSound-GZSL [37] is a modified version of the audio-visual dataset VGGSound [10]. **UCF-GZSL** [37] is a subset of the action video recognition dataset UCF101 [47] that includes audio information. **ActivityNet-GZSL** [37] is based on the action recognition dataset ActivityNet [6]. The statistics of the class split for three datasets are provided in Table 1 of [37] and Table 2 of the supplementary material of [37].

We use classical metrics S and U for evaluating the performances of seen and unseen classes, respectively. We also use the metric ZSL on unseen classes. In addition to ZSL, which solely evaluates the performance for unseen classes, the metric harmonic mean (HM) is used. It considers both seen and unseen classes' performances, S and U, in the test stage: $\text{HM} = \frac{2US}{U+S}$.

All these metrics are evaluated on both the “main fea-

Method	VGGSound-GZSL				UCF-GZSL				ActivityNet-GZSL			
	S \uparrow	U \uparrow	HM \uparrow	ZSL \uparrow	S \uparrow	U \uparrow	HM \uparrow	ZSL \uparrow	S \uparrow	U \uparrow	HM \uparrow	ZSL \uparrow
CJME [42]	8.69	4.78	6.17	5.16	26.04	8.21	12.48	8.29	5.55	4.75	5.12	5.84
AVGZSLNet [35]	18.05	3.48	5.83	5.28	52.52	10.90	18.05	13.65	8.93	5.04	6.44	5.40
AVCA [37]	14.90	4.00	6.31	6.00	51.53	18.43	27.15	20.01	24.86	8.02	12.13	9.13
Hyper-alignment	13.22	5.01	7.27	6.14	57.28	17.83	27.19	19.02	23.50	8.47	12.46	9.83
Hyper-single	9.79	6.23	7.62	6.46	52.67	19.04	27.97	22.09	23.60	10.13	14.18	10.80
Hyper-multiple	15.02	6.75	9.32	7.97	63.08	19.10	29.32	22.24	23.38	8.67	12.65	9.50

Table 1. Experimental results of audio-visual zero-shot learning on three datasets (main feature). AVCA [37] is adopted as the baseline for the proposed Hyper-alignment, Hyper-single, and Hyper-multiple modules. The best results in HM and ZSL are in **bold**. The curvatures of Hyper-alignment on VGGSound-GZSL, UCF-GZSL, and ActivityNet-GZSL are set as -0.5 , -0.2 , and -0.2 . The numbers of adaptive curvatures N_c for Hyper-multiple are 2, 3 and 2.

Method	VGGSound-GZSL ^{cls}				UCF-GZSL ^{cls}				ActivityNet-GZSL ^{cls}			
	S \uparrow	U \uparrow	HM \uparrow	ZSL \uparrow	S \uparrow	U \uparrow	HM \uparrow	ZSL \uparrow	S \uparrow	U \uparrow	HM \uparrow	ZSL \uparrow
CJME [42]	10.86	2.22	3.68	3.72	33.89	24.82	28.65	29.01	10.75	5.55	7.32	6.29
AVGZSLNet [35]	15.02	3.19	5.26	4.81	74.79	24.15	36.51	31.51	13.70	5.96	8.30	6.39
AVCA [37]	12.63	6.19	8.31	6.91	63.15	30.72	41.34	37.72	16.77	7.04	9.92	7.58
Hyper-alignment	12.50	6.44	8.50	7.25	57.13	33.86	42.52	39.80	29.77	8.77	13.55	9.13
Hyper-single	12.56	5.03	7.18	5.47	63.47	34.85	44.99	39.86	24.61	10.10	14.32	10.37
Hyper-multiple	15.62	6.00	8.67	7.31	74.26	35.79	48.30	52.11	36.98	9.60	15.25	10.39

Table 2. Experimental results of audio-visual zero-shot learning on three datasets (cls feature). AVCA [37] is adopted as the baseline for the proposed Hyper-alignment, Hyper-single, and Hyper-multiple modules. The best results in HM and ZSL are in **bold**. The curvatures of Hyper-alignment on VGGSound-GZSL^{cls}, UCF-GZSL^{cls}, and ActivityNet-GZSL^{cls} are set as -0.1 , -0.2 and -0.2 . The numbers of adaptive curvatures N_c for Hyper-multiple are 3, 2 and 3.

ture” and “cls feature,” which indicate the pre-trained models used to extract data features v , a , and w [37]. To fairly compare the performances of the Hyperbolic Alignment Module and the baseline, we keep all settings of hyperparameters in AVCA [37], including the learning rate, training epoch, *etc.* In this case, c_0 is set as -0.4 . In the ablation study, we analyze the effects of different parameters: curvature c of Hyper-alignment and the number of curvatures N_c of Hyper-multiple.

5.1. Result Analysis

The main results of the proposed modules for audio-visual zero-shot learning are presented in Table 1 and 2. In general, Hyper-alignment outperforms the baseline in most cases. For instance, on UCF-GZSL^{cls}, Hyper-alignment achieves an accuracy of 42.52%/39.80% in HM/ZSL, which is higher than the baseline’s accuracy of 41.34%/37.72%. Hyper-single, which computes an adaptive curvature for flexible exploration of audio-visual data structures, performs better than Hyper-alignment, except in the case of VGGSound-GZSL^{cls}. For example, on ActivityNet-GZSL, Hyper-single surpasses Hyper-alignment in HM/ZSL by 2.05%/0.97%. Meanwhile, Hyper-multiple, which learns multiple adaptive curvatures, clearly outperforms Hyper-single, except in the case of ActivityNet-GZSL.

As evidenced by our results on UCF-GZSL, Hyper-multiple outperforms all other methods, achieving a performance of 29.32%/22.24%. Moreover, the classification accuracy for seen classes (S) using Hyper-multiple is typically higher than those of other approaches. This suggests that the Hyper-multiple module generates more generic representations that are better suited for audio-visual zero-shot learning. In light of these results, we conclude that Hyper-multiple is the preferred design.

In conclusion, our proposed Hyperbolic Alignment Module improves audio-visual zero-shot learning performance over the baseline, as demonstrated in Table 1 and 2, with effective geometric alignments.

5.2. Ablation Study

Curvature. Experiments of Hyper-alignment are conducted under varying curvatures to understand how the degree to which the Euclidean space is distorted by hyperbolic geometry affects the model’s performance. In addition to the hyperbolic space ($c < 0$), feature alignments are also performed in Euclidean space ($c = 0$) and spherical space ($c > 0$). The results of these experiments are shown in Table 3, which indicate that the performance of Hyper-alignment on UCF-GZSL^{cls} and ActivityNet-GZSL^{cls} peaks at a curvature of -0.2 . This suggests that there is an optimal curvature for Hyper-alignment that

Method	c	UCF-GZSL ^{cls}		ActivityNet-GZSL ^{cls}	
		HM \uparrow	ZSL \uparrow	HM \uparrow	ZSL \uparrow
AVCA [37]	-	41.34	37.72	9.92	7.58
Sphere-alignment	1.0	38.62	31.64	9.26	6.23
	0.8	35.33	31.58	10.24	7.01
	0.6	38.80	38.62	10.28	7.01
	0.4	38.69	38.22	10.55	7.58
	0.2	38.56	34.45	10.56	6.95
	0.1	34.14	31.72	10.13	6.74
Euclidean-alignment	0.0	38.82	36.01	10.47	8.00
Hyper-alignment	-0.1	40.64	38.09	12.51	8.67
	-0.2	42.52	39.80	13.55	9.13
	-0.4	41.62	36.83	13.15	9.08
	-0.6	37.69	36.56	12.86	9.77
	-0.8	35.48	36.16	12.63	9.62
	-1.0	38.25	34.11	10.95	7.30

Table 3. Ablation study: curvature. Different curvature values of Hyper-alignment are tested on UCF-GZSL^{cls} and ActivityNet-GZSL^{cls}. Besides Hyper-alignment, Euclidean-alignment and Sphere-alignment are also considered. From the table, Hyper-alignment shows advantages over Euclidean-alignment and Sphere-alignment.

Method	c	UCF-GZSL ^{cls}		ActivityNet-GZSL ^{cls}	
		HM \uparrow	ZSL \uparrow	HM \uparrow	ZSL \uparrow
AVCA [37]	-	41.34	37.72	9.92	7.58
Mixed-curvature learning	E + E	39.26	37.05	9.83	7.35
	E + S	38.31	36.71	11.93	8.03
	E + H	42.73	37.74	7.49	5.70
	S + S	23.97	33.02	10.02	6.32
	S + H	40.37	38.16	12.25	7.78
	H + H	48.30	52.11	13.70	8.95

Table 4. Ablation study: mixed curvature learning. Different variants of mixed-curvature geometries are tested on UCF-GZSL^{cls} and ActivityNet-GZSL^{cls}. H, S and E indicate hyperbolic, spherical and Euclidean alignments. From the table, Geometry-multiple, which incorporates two hyperbolic geometries, gains the most performance improvements.

achieves the best performance. Moreover, we observe that the model is robust against the curvature of the hyperbolic space, as performances with different curvatures are better than the baseline in most cases.

Mixed curvature learning. In this experiment, variants of multiple geometric curvatures are evaluated. Hyper-multiple combines two hyperbolic alignments when the number of curvatures N_c is set to 2. Other combinations, such as Euclidean and spherical alignments, spherical and hyperbolic alignments, *etc.*, are also tested. The results in Table 4 show that the combination of two hyperbolic alignments (H + H) achieves the best performance among all the combinations. This observation is consistent with the results in Table 3, which demonstrate the advantage of hyperbolic geometry over other geometries in audio-visual zero-shot learning.

Effectiveness of multiple curvatures. Here, we test the performance of Hyper-multiple with different numbers of adaptive curvatures, N_c , on VGGSound-GZSL and

Method	N_c	VGGSound-GZSL		VGGSound-GZSL ^{cls}	
		HM \uparrow	ZSL \uparrow	HM \uparrow	ZSL \uparrow
AVCA [37]	-	6.31	6.00	8.31	6.91
Hyper-multiple	1	7.62	6.46	7.18	5.47
	2	9.32	7.97	8.97	6.72
	3	8.91	7.12	8.67	7.31

Table 5. Ablation study: effectiveness of multiple curvatures. Different curvature numbers of Hyper-multiple are evaluated on VGGSound-GZSL and VGGSound-GZSL^{cls}.

Method	aligned feature	UCF-GZSL		ActivityNet-GZSL	
		HM \uparrow	ZSL \uparrow	HM \uparrow	ZSL \uparrow
AVCA [37]	-	27.15	20.01	12.13	9.13
Hyper-single	ϕ_v/ϕ_a	27.97	22.09	14.18	10.80
	$\phi_{v,att}/\phi_{a,att}$	13.78	11.89	8.71	6.59
	θ_v/θ_a	25.58	18.77	12.96	10.00

Table 6. Ablation study: audio-visual feature alignment. Different locations of aligned features by Hyper-single are tested on datasets UCF-GZSL and ActivityNet-GZSL.

VGGSound-GZSL^{cls}. The results in Table 5 demonstrate that Hyper-multiple with 2 or 3 adaptive curvatures achieves the best performance. This suggests that, for Hyper-multiple, having more curvatures does not necessarily lead to better performance in audio-visual zero-shot learning.

Audio-visual feature alignment. In the main paper, we present the hyperbolic alignment modules that align the visual and audio features ϕ_v and ϕ_a before the cross-attention module of the baseline. In this experiment, we investigate aligning the features after the cross-attention and compare the results in Table 6. As seen from the table, aligning ϕ_v and ϕ_a before the cross-attention achieves the best result. On the other hand, aligning $\phi_{v,att}$ and $\phi_{a,att}$ harms the performance. These results suggest that the cross-modal fusion performed by the cross-attention module may alter the hierarchy of the data’s original features from each modality.

5.3. Discussion

δ -Hyperbolicity. The underlying structures of audio-visual datasets may exhibit high non-Euclidean latent geometry since there is a rich hierarchy in audio-visual data. In this part, we attempt to demonstrate that the proposed hyperbolic alignment modules are capable of uncovering some hierarchical knowledge. To measure the degree of the hier-

Method	VGGSound-GZSL	ActivityNet-GZSL ^{cls}
	δ_{rel}	δ_{rel}
AVCA [37]	0.30	0.37
Hyper-alignment	0.27	0.32

Table 7. Discussion: δ -Hyperbolicity. We compute δ_{rel} on the feature θ_v . The smaller value of δ_{rel} indicates the stronger hyperbolicity inside θ_v . Hyper-alignment gets a lower value of δ_{rel} than AVCA, which indicates the representations via the hyperbolic alignment loss learn more hierarchical structures.

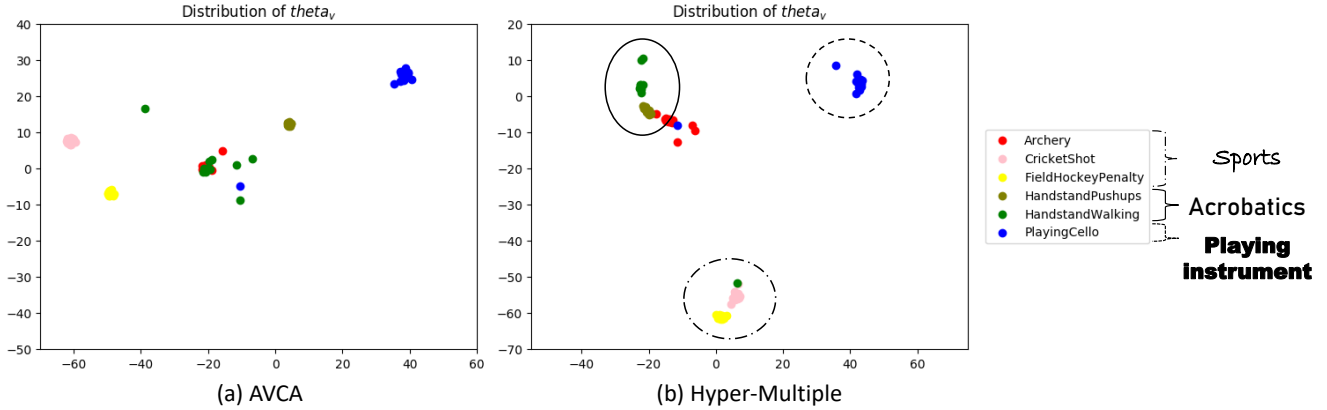


Figure 3. Visualization examples on UCF-GZSL^{cls}. We give t-SNE visualizations of θ_v from six seen classes: “Archery”, “CricketShot”, “FieldHockeyPenalty”, “HandstandPushups”, “HandstandWalking” and “PlayingCello”. They can be categorized into three superclasses: “Sports”, “Acrobatics” and “Playinginstruments”. Hyper-multiple, which learns a hyperbolic alignment loss, pushes away features from different superclasses.




Figure 4. Visualization examples on ActivityNet-GZSL^{cls}. We give t-SNE visualizations of θ_v . The features θ_v of three unseen classes, which are from different parent classes, are visualized in the upper figure. The lower figure depicts the feature distribution of three unseen classes from the same parent class. Comparing two figures, Hyper-multiple pushes away features more from different superclasses than the same superclass.

archy inside the features, we use the δ -Hyperbolicity metric, which has been shown to be effective in detecting hierarchical structure [28, 17]. In our case, we compute δ_{rel} [28, 17], a metric overall testing sample features that indicate the degree to which the feature distribution resembles a tree-like structure. Smaller values of δ_{rel} reflect stronger hyperbolicity within the feature space. As we can see from Table 7, Hyper-alignment achieves lower δ_{rel} values than AVCA on VGGSound-GZSL and ActivityNet-GZSL^{cls}. The results support our hypothesis that aligned representations using hyperbolic alignment loss convey more hierarchical information.

t-SNE Visualizations. In addition to measuring δ -Hyperbolicity, we provide t-SNE visualizations that reveal the capability of hyperbolic spaces in exploring data hierarchy. As shown in Figure 3, the hyperbolic alignment loss facilitates pulling together features from the same parent class while pushing away features belonging to different parent classes on UCF-GZSL^{cls}. For instance, features of “CricketShot” and “FieldHockeyPenalty”, which belong to the same parent class “Sports,” are pulled together, while features of “CricketShot” and “HandstandWalking”, which belong to different parent classes, are pushed away. To some extent, this explains how hyperbolic alignment between visual and audio features in hyperbolic spaces enables more accurate audio-visual zero-shot classification. Similarly, the visualization in Figure 4 suggests that on ActivityNet^{cls}, features belonging to different superclasses are more discriminately distributed than features from the same superclass. These visualizations somewhat describe how our approach uncovers the hierarchical knowledge among classes.

6. Conclusion

We introduce a novel hyperbolic alignment module that improves cross-modality representations in audio-visual zero-shot learning. This is achieved by leveraging the properties of hyperbolic spaces to explore more hierarchical structures within the audio-visual data and generate more distinctive features. Our approach includes three modules, Hyper-alignment, Hyper-single, and Hyper-multiple, for computing the loss, and we demonstrate their effectiveness through empirical evaluation. Our use of curvature-aware geometric learning to leverage data hierarchy may inspire further research in the field of audio-visual learning.

Acknowledgements

We acknowledge the support from the Australian Research Council (ARC) for M. Harandi's project DP230101176 and the Southeast University Start-Up Grant for New Faculty (No. 4009002309). This research work is also supported by the Big Data Computing Center of Southeast University.

References

- [1] Triantafyllos Afouras, Joon Son Chung, Andrew Senior, Oriol Vinyals, and Andrew Zisserman. Deep audio-visual speech recognition. *IEEE transactions on pattern analysis and machine intelligence*, 44(12):8717–8727, 2018. 1
- [2] Triantafyllos Afouras, Andrew Owens, Joon Son Chung, and Andrew Zisserman. Self-supervised learning of audio-visual objects from video. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, pages 208–224. Springer, 2020. 1
- [3] Yuki Asano, Mandela Patrick, Christian Rupprecht, and Andrea Vedaldi. Labelling unlabelled videos from scratch with multi-modal self-supervision. *Advances in Neural Information Processing Systems*, 33:4660–4671, 2020. 1
- [4] Mina Ghadimi Atigh, Julian Schoep, Erman Acar, Nanne van Noord, and Pascal Mettes. Hyperbolic image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4453–4462, June 2022. 3
- [5] Taivanbat Badamdorj, Mrigank Rochan, Yang Wang, and Li Cheng. Joint visual and audio learning for video highlight detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8127–8137, 2021. 1
- [6] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–970, 2015. 1, 5
- [7] Ines Chami, Albert Gu, Dat P Nguyen, and Christopher Ré. Horopca: Hyperbolic dimensionality reduction via horospherical projections. In *International Conference on Machine Learning*, pages 1419–1429. PMLR, 2021. 2
- [8] Ines Chami, Zhitao Ying, Christopher Ré, and Jure Leskovec. Hyperbolic graph convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 4869–4880, 2019. 3
- [9] Changan Chen, Sagnik Majumder, Ziad Al-Halah, Ruohan Gao, Santhosh Kumar Ramakrishnan, and Kristen Grauman. Learning to set waypoints for audio-visual navigation. *arXiv preprint arXiv:2008.09622*, 2020. 1
- [10] Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. Vggsound: A large-scale audio-visual dataset. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 721–725. IEEE, 2020. 1, 5
- [11] Jason Cramer, Ho-Hsiang Wu, Justin Salamon, and Juan Pablo Bello. Look, listen, and learn more: Design choices for deep audio embeddings. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3852–3856. IEEE, 2019. 1
- [12] Victoria Dean, Shubham Tulsiani, and Abhinav Gupta. See, hear, explore: Curiosity via audio-visual association. *Advances in Neural Information Processing Systems*, 33:14961–14972, 2020. 1
- [13] Yifan Ding, Yong Xu, Shi-Xiong Zhang, Yahuan Cong, and Liqiang Wang. Self-supervised learning for audio-visual speaker diarization. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4367–4371. IEEE, 2020. 1
- [14] Pengfei Fang, Mehrtash Harandi, Zhenzhong Lan, and Lars Petersson. Poincaré kernels for hyperbolic representations. *International Journal of Computer Vision*, 2023. 2
- [15] Pengfei Fang, Mehrtash Harandi, Trung Le, and Dinh Phung. Hyperbolic geometry in computer vision: A survey. *arXiv preprint arXiv:2304.10764*, 2023. 3
- [16] Pengfei Fang, Mehrtash Harandi, and Lars Petersson. Kernel methods in hyperbolic spaces. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10665–10674, 2021. 2
- [17] Hervé Fournier, Anas Ismail, and Antoine Vigneron. Computing the gromov hyperbolicity of a discrete metric space. *Information Processing Letters*, 115(6-8):576–579, 2015. 2, 8
- [18] Chuang Gan, Yiwei Zhang, Jiajun Wu, Boqing Gong, and Joshua B Tenenbaum. Look, listen, and act: Towards audio-visual embodied navigation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9701–9707. IEEE, 2020. 1
- [19] Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. Hyperbolic neural networks. In *Advances in neural information processing systems*, pages 5345–5355, 2018. 3
- [20] Ruohan Gao, Tae-Hyun Oh, Kristen Grauman, and Lorenzo Torresani. Listen to look: Action recognition by previewing audio. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10457–10467, 2020. 1
- [21] Zhi Gao, Yuwei Wu, Mehrtash Harandi, and Yunde Jia. Curvature-adaptive meta-learning for fast adaptation to manifold data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):1545–1562, 2022. 4
- [22] Zhi Gao, Yuwei Wu, Yunde Jia, and Mehrtash Harandi. Curvature generation in curved spaces for few-shot learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8691–8700, 2021. 4
- [23] Caglar Gulcehre, Misha Denil, Mateusz Malinowski, Ali Razavi, Razvan Pascanu, Karl Moritz Hermann, Peter Battaglia, Victor Bapst, David Raposo, Adam Santoro, and Nando de Freitas. Hyperbolic attention networks. In *International Conference on Learning Representations*, 2019. 3
- [24] Yunhui Guo, Haoran Guo, and Stella X Yu. Co-sne: Dimensionality reduction and visualization for hyperbolic data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21–30, 2022. 2

- [25] Jie Hong, Pengfei Fang, Weihao Li, Junlin Han, Lars Petersson, and Mehrtaash Harandi. Curved geometric networks for visual anomaly recognition. *arXiv preprint arXiv:2208.01188*, 2022. [3](#)
- [26] Di Hu, Xuhong Li, Lichao Mou, Pu Jin, Dong Chen, Liping Jing, Xiaoxiang Zhu, and Dejing Dou. Cross-task transfer for geotagged audiovisual aerial scene recognition. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIV 16*, pages 68–84. Springer, 2020. [1](#)
- [27] Evangelos Kazakos, Arsha Nagrani, Andrew Zisserman, and Dima Damen. Epic-fusion: Audio-visual temporal binding for egocentric action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5492–5501, 2019. [1](#)
- [28] Valentin Khruklov, Leyla Mirvakhabova, Evgeniya Ustinova, Ivan Oseledets, and Victor Lempitsky. Hyperbolic image embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6418–6428, 2020. [2](#), [3](#), [8](#)
- [29] Bruno Korbar, Du Tran, and Lorenzo Torresani. Cooperative learning of audio and video models from self-supervised synchronization. *Advances in Neural Information Processing Systems*, 31, 2018. [1](#)
- [30] Yan-Bo Lin, Hung-Yu Tseng, Hsin-Ying Lee, Yen-Yu Lin, and Ming-Hsuan Yang. Exploring cross-video and cross-modality signals for weakly-supervised audio-visual video parsing. *Advances in Neural Information Processing Systems*, 34:11449–11461, 2021. [1](#)
- [31] Shaoteng Liu, Jingjing Chen, Liangming Pan, Chong-Wah Ngo, Tat-Seng Chua, and Yu-Gang Jiang. Hyperbolic visual embedding learning for zero-shot recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9273–9281, 2020. [2](#), [3](#)
- [32] Teng Long, Pascal Mettes, Heng Tao Shen, and Cees G. M. Snoek. Searching for actions on the hyperbole. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2020. [3](#)
- [33] Rongkai Ma, Pengfei Fang, Tom Drummond, and Mehrtaash Harandi. Adaptive poincaré point to set distance for few-shot classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 1926–1934, 2022. [2](#)
- [34] Sagnik Majumder, Changan Chen, Ziad Al-Halah, and Kristen Grauman. Few-shot audio-visual learning of environment acoustics. *arXiv preprint arXiv:2206.04006*, 2022. [1](#)
- [35] Pratik Mazumder, Pravendra Singh, Kranti Kumar Parida, and Vinay P Nambodiri. Avgzslnet: Audio-visual generalized zero-shot learning by reconstructing label features from multi-modal embeddings. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3090–3099, 2021. [1](#), [2](#), [3](#), [6](#)
- [36] Otniel-Bogdan Mercea, Thomas Hummel, A Sophia Koepke, and Zeynep Akata. Temporal and cross-modal attention for audio-visual zero-shot learning. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XX*, pages 488–505. Springer, 2022. [1](#), [2](#), [3](#)
- [37] Otniel-Bogdan Mercea, Lukas Riesch, A Koepke, and Zeynep Akata. Audio-visual generalised zero-shot learning with cross-modal attention and language. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10553–10563, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [38] Himangi Mittal, Pedro Morgado, Unnat Jain, and Abhinav Gupta. Learning state-aware visual representations from audible interactions. *arXiv preprint arXiv:2209.13583*, 2022. [1](#)
- [39] Pedro Morgado, Ishan Misra, and Nuno Vasconcelos. Robust audio-visual instance discrimination. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12934–12945, 2021. [1](#)
- [40] Maximillian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. *Advances in neural information processing systems*, 30, 2017. [3](#)
- [41] Andrew Owens, Jiajun Wu, Josh H McDermott, William T Freeman, and Antonio Torralba. Ambient sound provides supervision for visual learning. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 801–816. Springer, 2016. [1](#)
- [42] Kranti Parida, Neeraj Matiyali, Tanaya Guha, and Gaurav Sharma. Coordinated joint multimodal embeddings for generalized audio-visual zero-shot classification and retrieval of videos. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3251–3260, 2020. [1](#), [2](#), [3](#), [6](#)
- [43] Wei Peng, Tuomas Varanka, Abdelrahman Mostafa, Henglin Shi, and Guoying Zhao. Hyperbolic deep neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. [3](#)
- [44] Jie Pu, Yannis Panagakis, Stavros Petridis, and Maja Pantic. Audio-visual object localization and separation using low-rank and sparsity. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2901–2905. IEEE, 2017. [1](#)
- [45] Rui Qian, Di Hu, Heinrich Dinkel, Mengyue Wu, Ning Xu, and Weiyao Lin. Multiple sound sources localization from coarse to fine. In *European Conference on Computer Vision*, pages 292–308. Springer, 2020. [1](#)
- [46] Arda Senocak, Tae-Hyun Oh, Junsik Kim, Ming-Hsuan Yang, and In So Kweon. Learning to localize sound source in visual scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4358–4366, 2018. [1](#)
- [47] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. [5](#)
- [48] Frederick Tung and Greg Mori. Similarity-preserving knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1365–1374, 2019. [4](#)
- [49] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. [2](#)

- [50] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10039–10049, 2021. [1](#)
- [51] Peng Wu, Jing Liu, Yujia Shi, Yujia Sun, Fangtao Shao, Zhaoyang Wu, and Zhiwei Yang. Not only look, but also listen: Learning multimodal violence detection under weak supervision. In *European conference on computer vision*, pages 322–339. Springer, 2020. [1](#)
- [52] Xinyi Wu, Zhenyao Wu, Lili Ju, and Song Wang. Binaural audio-visual localization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 2961–2968, 2021. [1](#)
- [53] Bo Xu, Cheng Lu, Yandong Guo, and Jacob Wang. Discriminative multi-modality speech recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14433–14442, 2020. [1](#)
- [54] Qinghao Ye, Xiyue Shen, Yuan Gao, Zirui Wang, Qi Bi, Ping Li, and Guang Yang. Temporal cue guided video highlight detection with low-rank audio-visual fusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7950–7959, 2021. [1](#)
- [55] Zhen Yu, Toan Nguyen, Yaniv Gal, Lie Ju, Shekhar S. Chandra, Lei Zhang, Paul Bonnington, Victoria Mar, Zhiyong Wang, and Zongyuan Ge. Skin lesion recognition with class-hierarchy regularized hyperbolic embeddings. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 594–603, 2022. [3](#)
- [56] Yiding Zhang, Xiao Wang, Chuan Shi, Xunqiang Jiang, and Yanfang Fanny Ye. Hyperbolic graph attention network. *IEEE Transactions on Big Data*, pages 1690–1701, 2021. [3](#)
- [57] Qichen Zheng, Jie Hong, and Moshir Farazi. A generative approach to audio-visual generalized zero-shot learning: Combining contrastive and discriminative techniques. In *2023 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2023. [2](#), [3](#)
- [58] Jinxing Zhou, Liang Zheng, Yiran Zhong, Shijie Hao, and Meng Wang. Positive sample propagation along the audio-visual event line. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8436–8444, 2021. [1](#)