

Visual Sound Localization in the Wild by Cross-Modal Interference Erasing

Xian Liu^{1,2*}, Rui Qian^{1,3*}, Hang Zhou^{1*}, Di Hu⁴, Weiyao Lin³, Ziwei Liu⁵,
Bolei Zhou¹, Xiaowei Zhou^{2‡}

¹ The Chinese University of Hong Kong,

² Zhejiang University,

³ Shanghai Jiao Tong University,

⁴ Gaoling School of Artificial Intelligence, Renmin University of China,

⁵ S-Lab, Nanyang Technological University

alvinliu@ie.cuhk.edu.hk, qr021@ie.cuhk.edu.hk, zhouhang@link.cuhk.edu.hk

Abstract

The task of audiovisual sound source localization has been well studied under constrained scenes, where the audio recordings are clean. However, in real world scenarios, audios are usually contaminated by off screen sound and background noise. They will interfere with the procedure of identifying desired sources and building visual sound connections, making previous studies nonapplicable. In this work, we propose the **Interference Eraser (IEr)** framework, which tackles the problem of audiovisual sound source localization in the wild. The key idea is to eliminate the interference by redefining and carving discriminative audio representations. Specifically, we observe that the previous practice of learning only a single audio representation is insufficient due to the additive nature of audio signals. We thus extend the audio representation with our Audio Instance Identifier module, which clearly distinguishes sounding instances when audio signals of different volumes are unevenly mixed. Then we erase the influence of the audible but off screen sounds and the silent but visible objects by a Cross modal Referrer module with cross modality distillation. Quantitative and qualitative evaluations demonstrate that our framework achieves superior results on sound localization tasks, especially under real world scenarios.

Introduction

Humans can grasp the relevance between audio and visual information by leveraging their natural correspondences (Proulx et al. 2014; Stein and Meredith 1993). Even for an in-the-wild scenario as shown in Fig. 1, humans are able to distinguish the sounding objects despite the background sound and noises. In order for machines to achieve human-like multi-modality perception, researchers have explored the task of visual sound source localization (Zhao et al. 2018; Senocak et al. 2018; Hu et al. 2020; Qian et al. 2020), which aims at localizing the objects that produce sound in an image given its corresponding audio clip.

With the recent development of neural networks, studies have firstly been made on learning audio-visual correspondences towards associating sound and visual objects. They



Figure 1: An example of an in-the-wild scene¹. It contains sounding clarinet, drum, human, etc. (red); visible but silent trumpet (green); background sound of water and wind (blue). Given an image and its corresponding audio as input, we aim to localize objects that produce sound in the image.

propose to optimize the similarity between audio and visual features according to spatial-level correspondences or temporal synchronization. Such pipeline effectively detects sound sources in scenarios with single sound source (Arandjelovic and Zisserman 2017, 2018; Owens and Efros 2018; Senocak et al. 2018). However, models learned under such settings cannot generalize well to mixed sound scenes. To further improve the model’s capacity, fine-grained audio-visual components are leveraged to establish sound-object associations for multiple sound source localization. Qian et al. (2020) propose to solve the problem in a coarse-to-fine manner. But it relies on audio class labels, which are sometimes inaccessible. Hu et al. (2020) design a two-stage pipeline to transfer single-sound prior to multi-sound cases and suppress the influence of silent objects. Nevertheless, they care less about the interference in audios, thus confining the applicable scenarios to clean environments.

The interference in audio is caused by the additive nature of audio signals. When the audio clip is a mixture of different sources and noise, several challenges occur: 1) In the localization task, sound sources of different volumes should be evenly identified. However, machines tend to distinguish only a few dominant sounds from a mixture as demonstrated

*Equal contribution. ‡ Corresponding author.

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://www.youtube.com/watch?v=3dXbFE0tIbs>.

by Phan et al. (2017); Adavanne et al. (2018). Thus the normal practice of previous studies (Arandjelovic and Zisserman 2017, 2018; Qian et al. 2020; Hu et al. 2020) to learn a single audio embedding is not enough for such complicated scenarios. More distinguishable representations are needed. 2) The undesired off-screen sound and background noise will be evenly taken into account during audio encoding, which might compromise the procedure of audio-visual modality matching.

In this paper, we propose a label-free method targeting in-the-wild visual sound localization called **Interference Eraser (IEr)**. The key idea is to *eliminate the interference by redefining and carving discriminative audio representations*. We regard volume differences between mixed sound sources, audible but off-screen sounds, and visible but silent objects as *interference* in the task of visual sound localization. In order to erase them, two questions need to be addressed: 1) How to identify audio instances regardless of their volume? 2) How to match the visible and sounding instances from the two modalities?

Specifically, we first build audio and visual *prototypes* in the feature space through learning the traditional single-source audio-visual correspondences (Arandjelovic and Zisserman 2018). The prototypes are the centroid features of each clustered category, which are responsible for recognizing individual instances in the absence of label information. For the case of mixed sound sources, we design an **Audio-Instance-Identifier** module to evenly identify all sounding instances by learning a latent step for each prototype to compensate for the mixture’s differences in volumes. In this way, we expand the audio representation from a single feature to a set of features, each taking care of one sounding category. At the second stage, we propose the **Cross-modal Referrer** that particularly targets the off-screen interference. The confidence score of all cluster-classes are transformed to class distributions for both domains, and we particularly focus on erasing the interference on the distributions by cross-referring the two modalities. Finally, the predicted distributions on all clustered categories from both modalities serve as soft targets for each other to perform cross distillation.

Our contributions are summarized as follows:

- We design the Audio-Instance-Identifier module which erases the influence of sound volume in uneven audio mixtures for discriminative audio perception.
- We devise the Cross-modal Referrer module to symmetrically erase the interference of silent but visible objects and audible but off-screen sounds.
- Through cross-modal distillation learning, our proposed **Interference Eraser** remarkably outperforms the state-of-the-art methods in terms of sound localization for both synthetic and in-the-wild audio-visual scenes.

Related Work

Audio-Visual Correspondence Learning. The correspondence between co-occurring modalities, *e.g.*, audio and vision, provides natural and expressive self-supervision for representation learning in multi-modal scenarios (Korbar, Tran, and Torresani 2018; Aytar, Vondrick, and Torralba

2016; Tian et al. 2018; Zhou et al. 2019b,a, 2021; Liu et al. 2022b,a). Recent works (Alwassel et al. 2019; Patrick et al. 2020) use temporal consistency as a self-supervisory signal to facilitate audio-visual learning. Typically, Arandjelovic and Zisserman (2018) calculate the similarity between global audio embedding and visual features of each spatial grid to generate audio-visual correlation maps. Zhao et al. (2018, 2019) employ audio-visual consistency to align channel-wise cross-modal distributions, then leverage the learned representation for sound source separation. However, most previous works are performed in single-source settings. In real-world scenarios with multiple sources, as proved in previous works on sound separation (Tzinis et al. 2021; Gao and Grauman 2019; Gan et al. 2020; Pu et al. 2017), the audio prediction could be quite unbalanced, which affects the mapping between audio and visual modality. In this work, we focus on in-the-wild scenes.

Visual Sound Source Localization. Visually localizing sound source aims to figure out specific sounding area in the visual scene given a clip of audio, and can be achieved in various ways (Arandjelovic and Zisserman 2018; Chen et al. 2021; Lin et al. 2021; Zhou et al. 2020; Xu et al. 2021; Tian, Hu, and Xu 2021; Tian and Xu 2021; Afouras et al. 2020). For example, Arandjelovic and Zisserman (2018); Owens and Efros (2018) perform sound localization by measuring similarity between audio embedding and visual features on spatial area. Senocak et al. (2018) develop audio-guided attention mechanisms to refine audio-related visual features and improve cross-modal alignment. In complex scenes, Zhao et al. (2018) introduce sound separation as a pretext task to perform visual grounding through channel correlation. Particularly, Qian et al. (2020) target multi-sound scenes by adopting category prior to associate one-to-one sound-object pairs. One most recent work (Hu et al. 2020) proposes to discriminatively localize sounding objects via category-level audio-visual distribution alignment. Specifically, it first trains audio-visual correspondence in single-sound scenes, then *directly* transfers to mixed sound cases for audio-visual matching. However, the quite unbalanced audio mixture and off-screen sound in multi-source scenarios impose great challenges to previous works. Alternatively, we develop the Interference Eraser to achieve robust mixed sound perception and filter out interference.

Method

In this section, we introduce the details of our **Interference Eraser (IEr)** for label-free visual sound localization. The whole pipeline is illustrated in Fig. 3. Below we first illustrate the first-stage learning: audio-visual prototype extraction by single-source audio-visual learning (Sec.), and the Audio-Instance-Identifier with a curriculum learning scheme for mixed-audio instance discrimination (Sec.). Then, we introduce the second-stage learning, *i.e.*, Cross-modal Referrer (Sec.) to perform in-the-wild scene training and distribution matching for robust sound localization.

Notably, the training videos are denoted as $\mathcal{X} = \{(a_i, v_i) | i = 1, 2, \dots, N\}$, where N is the number of videos and (a_i, v_i) is the i -th audio-visual pair. \mathcal{X} can be divided into two disjoint subsets: $\mathcal{X}^s = \{(a_i^s, v_i^s) | i =$

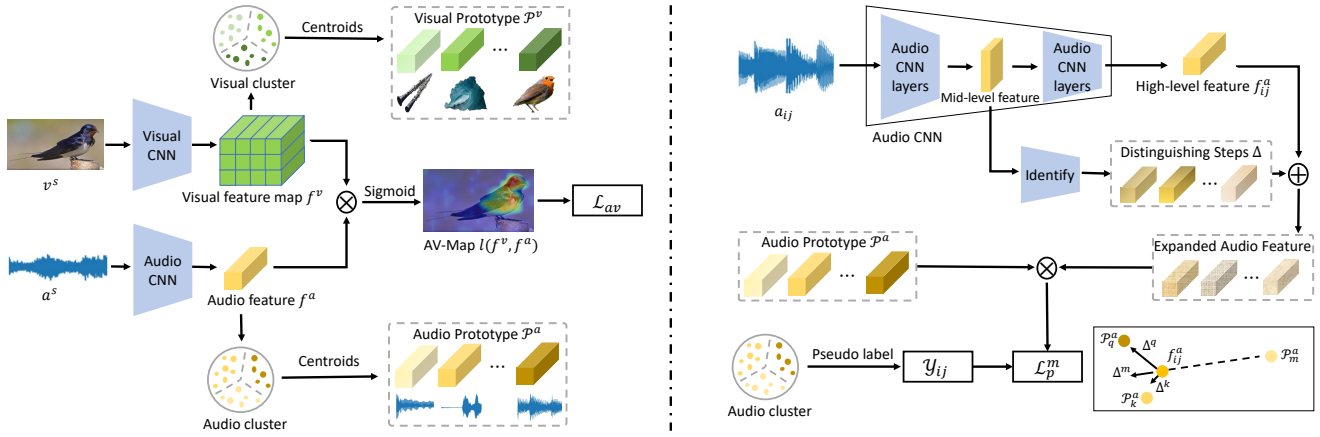


Figure 2: An illustration of the first-stage learning. The left presents audio-visual prototype extraction. We calculate the cosine similarity between the audio embedding and the visual feature at each position to generate the localization map, and cluster audio-visual features into centroids as prototypes. The right shows Audio-Instance-Identifier. Given mixed audio, we predict distinguishing-steps to generate expanded audio features. As shown in the corner, it predicts a small/big step for dominant/non-dominant sound to approach all sounding prototypes, and a reverse step to repel from non-existent sounds’ prototypes.

$1, 2, \dots, N^s\}$ consisting of N^s single source videos and $\mathcal{X}^u = \{(a_i^u, v_i^u) | i = 1, 2, \dots, N^u\}$ consisting of N^u unconstrained in-the-wild videos. The first stage training is conducted on \mathcal{X}^s , and the second stage is conducted on \mathcal{X}^u .

Audio-Visual Prototype Extraction

Towards localizing sound-making objects, the first step is to associate the visual appearance of objects with their sound. While visual instances can be divided spatially, audio signals are naturally mixed together in real-world scenarios. Moreover, no label information can be leveraged in our setting. In order to discriminate different instances from a mixture, representations of each possible sounding category need to be identified.

To this end, our first step is to perform the traditional instance-level feature correspondence learning on the single source subset \mathcal{X}^s (Arandjelovic and Zisserman 2018; Zhao et al. 2018) and find the latent representatives for each class in the feature space, namely the *prototypes* in the same way as (Hu et al. 2020).

Single Source Audio-Visual Learning. As illustrated in Fig. 2 left, we employ the variants of ResNet-18 (He et al. 2016) as visual and audio backbones to extract visual feature map $f^v \in \mathbb{R}^{H \times W \times C}$ and the audio feature $f^a \in \mathbb{R}^C$ from \mathcal{X}^s . The cosine similarity between f^a and each positional feature $f^{v(x,y)}$ of f^v are used to present an audio-visual association map (localization map) $l(f^v, f^a) \in \mathbb{R}^{H \times W}$. The training is through contrastively positive and negative pair sampling:

$$\mathcal{L}_{av} = \mathcal{L}_{bce}(\text{GMP}(l(f_i^a, f_j^v)), \delta), \quad (1)$$

where GMP denotes Global Max Pooling, \mathcal{L}_{bce} is binary cross-entropy loss function and δ indicates whether f_i^a and f_j^v are extracted from the same video, *i.e.*, $\delta = 1$ when $i = j$, otherwise $\delta = 0$.

Embedding Prototypes. As visual appearances of different objects are comparatively easier to learn, the categories

can be roughly defined through deep clustering (Caron et al. 2018; Alwassel et al. 2019) the visual object features extracted from \mathcal{X}^s . We manually set a large cluster number and finally settle a total number of K categories. The visual *prototype* $\mathcal{P}_j^v \in \mathbb{R}^C$ is defined as the centroid for the j -th class in the feature space. All visual prototypes can be represented as $\mathcal{P}^v \in \mathbb{R}^{K \times C}$.

On the other hand, the pseudo-class labels can be automatically assigned to the video’s accompany audios. Thus the audio prototypes $\mathcal{P}^a \in \mathbb{R}^{K \times C}$ can be identified in the same way as visual ones. In this way, a pseudo-class k can be assigned to each single source audio-visual data pair $(a_i^{s_k}, v_i^{s_k})$ in \mathcal{X}^s ,

Audio-Instance-Identifier

As stated before, localizing sounding objects in multi-source scenes requires identifying each object. Even with the prototypes learned, one embedded feature from a mixture of audios can not be similar to all prototypes, thus limits the feature’s capacity to identify all instances. Particularly in an audio mixture with different volumes, the dominant sound might severely interfere with the audio discrimination.

To accurately identify different sounding components, we propose the Audio-Instance-Identifier module which expands the audio representation. The key is to learn a *Distinguishing-step* for each prototype in a curriculum manner as shown in the right part of Fig. 2.

Learning the Distinguishing-step. Inspired by research in sound source separation (Zhao et al. 2018; Gao and Grauman 2019), we propose to manually mix single sources together as input and aim at recognizing them individually. We denote the pseudo label for an single audio clip $a_i^{s_k}$ as $\mathcal{Y}_i \in \mathbb{R}^K$ whose k -th element $\mathcal{Y}_i^k = 1$ and others $\mathcal{Y}_i^n = 0$ when $n \neq k$. By mixing the i -th audio clip $a_i^{s_k}$ with pseudo-class k and the j -th clip $a_j^{s_q}$ with pseudo-class $q (q \neq k)$ together, a mixture a_{ij} with pseudo label $\mathcal{Y}_{ij} = \mathcal{Y}_i + \mathcal{Y}_j$ can be rendered.

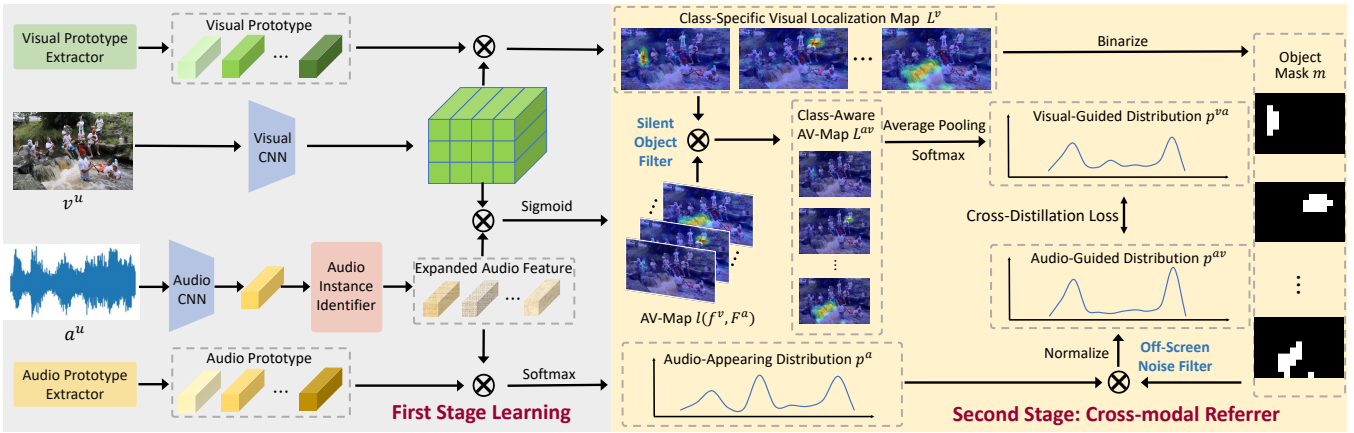


Figure 3: An overview of the Interference Eraser (IEr) framework. The left presents the first stage, where we train on single-sound data to establish audio-visual prototypes, and then utilize Audio-Instance-Identifier to facilitate class-specific audio features. The right presents the second stage, where the silent object filter suppresses visual interference by referring to balanced audio distribution, and the off-screen noise filter eliminates audio interference by referring to binary visual masks. A cross-distillation loss helps to achieve robust sound localization in-the-wild. The class-aware AV-Maps L^{av} are the final output in inference.

In order to erase the interference of volume difference and identify all sounding categories, we propose to learn a set of *distinguishing-steps* $\Delta^n(a) \in \mathbb{R}^{K \times C}$ towards each prototype \mathcal{P}_n^a from the input audio. The motivation is to compensate for the amplitude information in the feature space. For example, as shown in Fig. 2, in the two-mix case listed above, when the learned feature f_{ij}^a is too close to the dominant \mathcal{P}_k^a but far away from \mathcal{P}_q^a , the desired $\Delta^q(a_{ij})$ should be a large step towards \mathcal{P}_q^a while $\Delta^k(a_{ij})$ is a smaller step towards \mathcal{P}_k^a . We can formulate a binary classification problem for each pseudo-class:

$$\mathcal{L}_p^m = \frac{1}{K} \sum_{n=1}^K \mathcal{L}_{bce}(\text{sim}(\mathcal{P}_n^a, f_{ij}^a + \Delta^n(a)), \mathcal{Y}_{ij}^n). \quad (2)$$

where $\text{sim}(f_1, f_2)$ is the cosine similarity function defined as $\text{sim}(f_1, f_2) = \frac{f_1 \cdot f_2}{\|f_1\| \|f_2\|}$. It can be seen that this distinguishing-step also helps to repel from the non-existent audio instances' prototypes by predicting a step to the reverse direction of $\mathcal{P}_m^a (m \neq q, k)$ as shown in Fig. 2.

Particularly, distinguishing-step is not directly learned from high-level feature f_{ij}^a . It is verified that mid-level features can capture more fine-grained information (Yosinski et al. 2014; Lin et al. 2017; Zou et al. 2020), thus we employ the weighted combination of linearly transformed mid-level features (shown in Fig. 2) to make prediction.

Curriculum Learning Strategy. We develop a curriculum learning strategy for the training process of the Audio-Instance-Identifier. Detailedly, we choose to alternatively train on single-source data without the discriminating-step and mixed sound with \mathcal{L}_p^m (Eq. 2). In each epoch, a proportion of $p < 1$ samples are mixed ones. p is initialized as 0.5 and gradually increases to 0.9 at the end of training. While Eq. 2 shows only a two-mix case, it can be expanded to cases with 3 or more mixtures. The number of mixtures is set as 2

at first and gradually increases to 4 during the training process. As the audio backbone is keeping updated, the audio prototypes are re-computed at the beginning of each epoch.

Cross-modal Referrer

At the second stage, we move to real-world multi-source scenarios, where the training and inference are carried out on \mathcal{X}^u . Normally, the localization map on general scenes can be calculated in the same way as learning audio-visual correspondences, that is, to directly encode the audio feature f^a from a mixed audio a^u , and compute cosine similarities $l(f^v, f^a)$ on the visual feature map f^v (Sec.). However, this operation will suffer from the same uneven audio feature encoding problem as stated in Sec. .

Thus, we propose the Cross-modal Referrer as shown in Fig. 3. It first identifies the audio-visual instances, and then predicts the class-aware audio-visual localization maps where visual interference is erased by referring to our updated audio representation. Afterwards, we leverage the acquired localization map to erase the audio interference and compute the category distributions of both audio and visual domains. Finally, a cross-modal distillation loss is used to facilitate better cross-modality matching.

Class-Aware Audio-Visual Localization Maps. Given an input image with the visual prototypes \mathcal{P}^v , it is easy for us to find all potential sound-making objects by computing the visual localization map for the category k as $L_k^v = l(f^v, \mathcal{P}_k^v)$, $L_k^v \in \mathbb{R}^{H \times W}$. However, within the L^v s, there might be responses on silent but visible object's map L_m^v , thus we refer to the audio domain.

We suppose the audio clip a^u accompanies v^u is a mixture of sources including an off-screen sound a_j . Its feature f^a together with its *distinguishing-steps* $\Delta(a) = [\Delta^1(a), \dots, \Delta^K(a)]$ for each pseudo class can be computed as illustrated in Sec. . The audio feature can be expanded to a set of *class-specific* features $F_k^a = f^a + \Delta^k(a)$.

Scenario	Single Sound Scene (a)				General in-the-Wild Scene (b)							
Dataset	MUSIC		VGGSound		MUSIC-Syn.		MUSIC-Duet		MUSIC-Un.		VGG-Un.	
Methods\Metrics	IoU	AUC	IoU	AUC	CIoU	AUC	CIoU	AUC	CIoU	AUC	CIoU	AUC
Object-that-sound	26.1	35.8	48.4	46.1	3.7	10.2	13.2	18.3	0.1	6.8	7.8	15.1
Sound-of-pixel	40.5	43.3	42.5	45.1	8.1	11.8	16.8	16.8	7.5	11.6	7.9	14.4
DSOL	51.4	43.6	49.3	45.8	32.3	23.5	30.2	22.1	3.2	7.3	8.1	12.2
Interference Eraser	53.9	50.7	51.3	46.9	47.6	29.8	52.9	33.8	15.6	15.3	12.8	17.6

Table 1: Visual Sound Localization results in (a) Single Sound Scene: MUSIC, VGGSound; (b) General in-the-Wild Scene: MUSIC-Syn., MUSIC-Duet, MUSIC-Un., VGG-Un. We compare IEr against recent SOTA methods. The IoU reported in Table(a) is IoU@0.5, and the CIoU reported in Table(b) is CIoU@0.3. The results of (Arandjelovic and Zisserman 2018; Zhao et al. 2018; Hu et al. 2020) on MUSIC, MUSIC-Syn. and MUSIC-Duet are reported from (Hu et al. 2020).

Then we can predict the *class-aware audio-visual localization maps* by associating the class-specific audio features F^a with the visual feature map f^v again, then mask it by the class-specific visual localization maps L^v . It is computed as $L_k^{av} = l(f^v, F_k^a) \cdot L_k^v$ for each class k , where \cdot is the element-wise matrix multiplication. As there would be little responses on the silent object m for $l(f^v, F_m^a) \cdot L_m^v$, the interference of m is thus suppressed. This whole operation is marked as the **Silent Object Filter** on Fig. 3. Note that these class-aware audio-visual localization maps (AVMaps) are also the desired output of **IEr** system during inference.

Acquiring Audio-Visual Distribution. Besides acquiring the audio-visual localization maps (AVMaps), we can also predict the on-screen objects’ distributions on the pseudo-classes. The way is to perform global average pooling (GAP) on each AVMap as a rough accumulation of the class’s response. Then use a softmax function to regularize all responses on each class. This visual-guided distribution can be written as:

$$p^{va} = \text{softmax}([\text{GAP}(L_1^{av}), \dots, \text{GAP}(L_K^{av})]). \quad (3)$$

On the other hand, we can also compute the audio-appearing probability for each pseudo-class as cosine similarity between the audio prototype and class-specific audio features $p_k^a = \text{sim}(\mathcal{P}_k^a, F_k^a)$. However, the off-screen sound a_j might interfere with the final on-screen sound distribution prediction. To erase this interference, we refer to the help of visual localization maps L^v . Specifically, we first binarize each visual localization map L_k^v to a class-specific binary mask m_K by setting a threshold. This is to suppress noisy peaks in L_k^v . Then we leverage the global average-pooled binary mask $\text{GAP}(m_k)$ to accounts for the portion of this instance in the visual domain. Thus the visual-associated audio-appearing probability on class k can be written as:

$$p_k^{av} = \frac{\text{GAP}(m_k) \cdot p_k^a}{\sum_{k=1}^K \text{GAP}(m_k) \cdot p_k^a}. \quad (4)$$

After this operation, the influence of off-screen sound a_j can be erased, as the reference visual contribution of its corresponding class would be subtle. This whole operation is marked as the **Off-screen Noise Filter** on Fig. 3. The audio-guided class distribution can then be computed as:

$$p^{av} = \text{softmax}([p_1^{av}, \dots, p_K^{av}]). \quad (5)$$

Notably, the visual size of objects does not necessarily connect with the volume of their sound. Thus this cross-modal referring operation on the audio distribution not only erase off-screen sounds but also modulate the audio subject’s probability according to its visual size, which would be beneficial to cross-modality distribution matching.

Cross-Modal Distillation. Finally, given the class distributions p^{va} and p^{av} , which are guided by the visual and audio modalities respectively, we can regard one of them as soft targets in the formulation of knowledge distillation (Hinton, Vinyals, and Dean 2015). Substantially, it is to leverage the Kullback–Leibler divergence \mathcal{D}_{KL} for similarity measurement between the two distributions. This additional training objective for in-the-wild scene is:

$$\mathcal{L}_u = \frac{1}{2} \mathcal{D}_{KL}(p^{va} || p^{av}) + \frac{1}{2} \mathcal{D}_{KL}(p^{av} || p^{va}). \quad (6)$$

Experiments

Dataset

MUSIC (Synthetic) MUSIC dataset (Zhao et al. 2018) covers 11 types of instruments. Since some YouTube videos have been removed, we collect 489 solo and 141 duet videos. Following (Hu et al. 2020), we employ half of the solo data to form \mathcal{X}^s , with the other half to synthesize \mathcal{X}^u . Specifically, we use the first five/two videos of each category in solo/duet to evaluate localization performance and employ the rest videos to generate MUSIC-Synthetic (Hu et al. 2020) as well as MUSIC-Unconstrained.

VGGSound VGGSound dataset (Vedaldi et al. 2020) consists of more than 210k single-sound videos, covering 310 categories. For better evaluation, we filter out a subset of 98 categories with specific sounding objects that can be localized. Similarly, we use half of the data for single-sound and the other half for synthesizing VGG-Synthetic. We finally obtain 28,756 videos for training and 2,787 for evaluation. This subset is very suitable for class-aware sound localization and also provides diverse challenging scenarios.

Unconstrained-Synthetic Since most samples in two datasets do not simultaneously contain the aforementioned interference, they fail to fully reflect the difficulty of in-the-wild scenarios. To better validate sound localization in general scenes, we establish two challenging datasets called MUSIC-Unconstrained and VGG-Unconstrained. For each

synthetic audio-visual pair, the image contains two sounding and two silent objects, and the audio is a mixture of on-screen and randomly chosen off-screen audio clips. In this way, the Unconstrained dataset simulates in-the-wild scenes with visible but silent objects and off-screen noise.

Experimental Settings

Comparing Baselines. We compare our model with the following methods: **Object-that-sound** (Arandjelovic and Zisserman 2018) which optimizes similarity between audio embedding and spatial visual features; **Sound-of-pixel** (Zhao et al. 2018) that establishes channel-wise audio-visual correlation by training sound separation. (Qian et al. 2020) requires audio category prior to discriminate multiple sounds, which is different from our setting, so we do not make comparison. Particularly, we compare with the SOTA **DSOL** (Hu et al. 2020), which also learns object representation from single-sound data and then expands to complex cases.

Since we focus on visual sound localization without labels, where all classes are cluster-based pseudo ones. So we do not compare with other mixed audio perception methods.

Evaluation Metrics. For the multi-modal prototype extraction, we use *Normalized Mutual Information* (NMI) between cluster assignments and category labels to measure the discrimination of audio-visual backbones. To validate model’s ability in challenging mixed sound perception, following (Vedaldi et al. 2020), *mean Average Precision* (mAP) is adopted to evaluate audio classification performance. Besides, we adopt *Precision* (Prec.) and *Recall* (Rec.) of multi-label classification for more detailed analysis.

For sound localization, we use *Intersection over Union* (IoU) and *Area Under Curve* (AUC) (Senocak et al. 2018) for single-sound scene evaluation and adopt *Class-aware IoU* (CIoU) (Hu et al. 2020) for in-the-wild localization:

$$\text{CIoU} = \frac{\sum_{t=1}^K \mathcal{Y}^m(t) \text{IoU}_t}{\sum_{t=1}^K \mathcal{Y}^m(t)}, \quad (7)$$

where $\mathcal{Y}^m(t)$ is the pseudo-label for the t -th category and IoU_t is calculated based on the predicted sounding object area and annotated bounding box for the t -th class.

Implementation Details. We split videos in each dataset into non-overlapped 1-second clips to form audio-visual pairs. Concretely, we sample audio at 16kHz with window length of 160ms, hop length of 80ms and transform it into log-mel spectrograms with 64 frequency bins. For visual stream, we randomly sample a frame from each clip and resize it into 256×256 , then randomly/center crop into 224×224 for training/evaluation. The model is trained by Adam optimizer with learning rate 10^{-4} . For evaluation, we use Faster RCNN (Ren et al. 2015) to detect bounding box as reference. We aggregate cluster assignments to category labels to class-specifically measure localization performance.

Quantitative Analysis

Single Sound Localization. We first analyze the quantitative result of single-sound localization. As shown in Table 1(a), our **IER** framework outperforms existing methods on both MUSIC and VGGSound datasets. Considering that

there still exists noise in single-sound data and not all audio-visual pairs are clean, such superiority indicates the **Audio-Instance-Identifier** trained on mixed audio improves the robustness and discrimination of extracted audio embedding.

Discriminative Sound Localization in-the-Wild. General in-the-wild scenarios are often accompanied by interference of unbalanced audio mixture, visible but silent objects and off-screen sound, which makes it intractable to localize sounding objects of each class without extra annotations. As shown in Table 1(b), our model outperforms other methods by a large margin on both realistic and synthetic scenes. Such significant improvements majorly benefit from two aspects: 1) The **Audio-Instance-Identifier** leverages the distinguishing-step with curriculum learning to obtain expanded class-specific mixed audio features, while (Arandjelovic and Zisserman 2018; Zhao et al. 2019; Hu et al. 2020) failed to handle mixed-sound perception well. 2) The **Cross-modal Referrer** eliminates the pervasive interference of visible but silent objects and off-screen noise in-the-wild, while previous methods ignored this problem and failed to achieve robust audio-visual matching.

Qualitative Analysis

Besides the quantitative analysis, we compare visualization results for class-specific sound localization and perceived in-the-scene audio distribution of different methods in Fig. 4. Note that Object-that-sound (Arandjelovic and Zisserman 2018) is only applicable to single-sound cases, we do not present their results. For in-the-wild scene in the right part of Fig. 4, Sound-of-pixel (Zhao et al. 2018) cannot identify the exact location of objects and fails to filter out silent objects or suppress off-screen sound. This is because it requires clean data for the pretext task training and has difficulty adapting to noisy cases. The audio perception of DSOL (Hu et al. 2020) is quite unbalanced and dominated by the dog barking, which makes it only manage to localize the dog while treating the talking person as silent. On the contrary, with the help of two key modules, our method achieves balanced audio perception as well as off-screen sound erasing. Besides, the proposed framework can precisely localize objects in a class-aware manner and filter out silent ones.

Ablation Study

Audio-Instance-Identifier. To quantify the efficacy of Audio-Instance-Identifier with curriculum learning strategy, we present the ablation study results in terms of NMI between cluster assignments and category labels as well as multi-label audio classification in Table 2. It indicates that Audio-Instance-Identifier leads to improvement over all metrics. The significant gain in recall reveals that this module helps to avoid neglecting non-dominant sounds and balance audio perception. Besides, curriculum learning strategy further improves performance since it provides the model with effectively arranged audio samples to enhance the robustness to noise and the expressiveness of prototypes.

Cross-modal Referrer. Table 3 shows the ablative results on two sub-modules of Cross-modal Referrer. The results demonstrate that it is necessary to erase interference in both modalities. That is, either off-screen sound or visible but

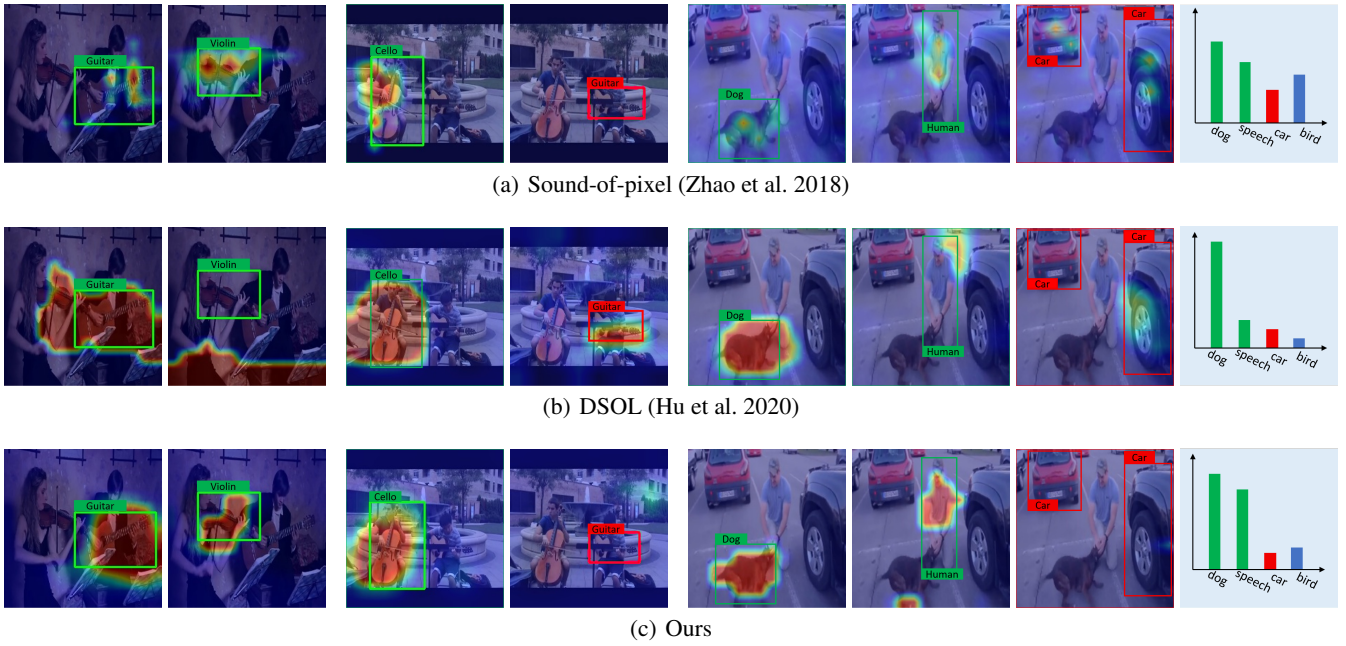


Figure 4: Comparison of sound localization and audio perception between different methods. In class-specific localization maps, the green boxes are sounding objects, the red boxes indicate visible but silent objects. And the histogram represents perceived audio distribution, the green columns indicate in-the-scene sounds, the red column means silent category, and the blue column means off-screen sound that is expected to be suppressed. More qualitative results can be found in demo video.

Dataset	MUSIC				VGGSound			
Method	NMI	Prec.	Rec.	mAP	NMI	Prec.	Rec.	mAP
w/o Audio-Instance-Identifier	0.692	0.372	0.335	0.355	0.410	0.218	0.076	0.189
Audio-Instance-Identifier w/o curriculum	0.758	0.441	0.489	0.403	0.414	0.298	0.160	0.231
Audio-Instance-Identifier w/ curriculum	0.809	0.461	0.715	0.433	0.436	0.346	0.232	0.283

Table 2: Ablation study on Audio-Instance-Identifier with curriculum learning strategy in terms of NMI and multi-label audio classification. We empirically set the threshold as 0.5 when calculating Prec. and Rec.

Dataset		MUSIC-Un.		VGG-Un.	
Silent.	Off-Screen.	CIoU	AUC	CIoU	AUC
✗	✗	2.7	6.9	5.9	10.2
✓	✗	4.6	9.5	7.2	14.1
✗	✓	5.1	10.8	7.8	13.9
✓	✓	15.6	15.3	12.8	17.6

Table 3: Ablation study on Silent Object Filter and Off-Screen Noise Filter for MUSIC-Un. and VGG-Un.

Dataset		MUSIC-Syn.		VGG-Un.	
All.	CMR.	CIoU	AUC	CIoU	AUC
✗	✗	0.2	6.8	3.1	7.8
✓	✗	12.8	11.1	5.9	10.2
✗	✓	29.5	22.3	6.9	13.5
✓	✓	47.6	29.8	12.8	17.6

Table 4: Ablative experiments on Audio-Instance-Identifier (All.) and Cross-modal Referrer (CMR.).

silent objects introduce severe interference for audio-visual correspondence learning and corrupt sound localization.

The overall ablation study of Audio-Instance-Identifier and Cross-modal Referrer is shown in Table 4, verifying that both modules are crucial for in the wild sound localization.

Conclusion

In this work, we introduce a novel framework Interference Eraser (**IEr**) to enhance robust visual sound localization for

in-the-wild scenes. We identify the two strengths of our method: 1) Our proposed Audio-Instance-Identifier learns the distinguishing-step to achieve volume agnostic mixed sound perception, which erases the interference of uneven sound mixtures. 2) We propose the Cross-modal Referrer to eliminate the interference of visible but silent objects and audible but off-screen sounds. Extensive experiments demonstrate the superiority of our proposed method for sound localization, particularly for in-the-wild scenarios.

Acknowledgments

This study is supported by National Natural Science Foundation of China (No. U21B2013), NTU NAP, MOE AcRF Tier 1 (2021-T1-001-088), and under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

References

- Adavanne, S.; Politis, A.; Nikunen, J.; and Virtanen, T. 2018. Sound event localization and detection of overlapping sources using convolutional recurrent neural networks. *IEEE Journal of Selected Topics in Signal Processing*, 13(1): 34–48.
- Afouras, T.; Owens, A.; Chung, J. S.; and Zisserman, A. 2020. 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*, 208–224. Springer.
- Alwassel, H.; Mahajan, D.; Torresani, L.; Ghanem, B.; and Tran, D. 2019. Self-supervised learning by cross-modal audio-video clustering. *arXiv preprint arXiv:1911.12667*.
- Arandjelovic, R.; and Zisserman, A. 2017. Look, listen and learn. In *Proceedings of the IEEE International Conference on Computer Vision*, 609–617.
- Arandjelovic, R.; and Zisserman, A. 2018. Objects that sound. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 435–451.
- Aytar, Y.; Vondrick, C.; and Torralba, A. 2016. Soundnet: Learning sound representations from unlabeled video. In *Advances in neural information processing systems*, 892–900.
- Caron, M.; Bojanowski, P.; Joulin, A.; and Douze, M. 2018. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 132–149.
- Chen, H.; Xie, W.; Afouras, T.; Nagrani, A.; Vedaldi, A.; and Zisserman, A. 2021. Localizing Visual Sounds the Hard Way. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16867–16876.
- Gan, C.; Huang, D.; Zhao, H.; Tenenbaum, J. B.; and Torralba, A. 2020. Music Gesture for Visual Sound Separation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10478–10487.
- Gao, R.; and Grauman, K. 2019. Co-separating sounds of visual objects. In *Proceedings of the IEEE International Conference on Computer Vision*, 3879–3888.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the Knowledge in a Neural Network. In *NIPS Deep Learning and Representation Learning Workshop*.
- Hu, D.; Qian, R.; Jiang, M.; Tan, X.; Wen, S.; Ding, E.; Lin, W.; and Dou, D. 2020. Discriminative Sounding Objects Localization via Self-supervised Audiovisual Matching. *arXiv preprint arXiv:2010.05466*.
- Korbar, B.; Tran, D.; and Torresani, L. 2018. Cooperative learning of audio and video models from self-supervised synchronization. In *Advances in Neural Information Processing Systems*, 7763–7774.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125.
- Lin, Y.-B.; Tseng, H.-Y.; Lee, H.-Y.; Lin, Y.-Y.; and Yang, M.-H. 2021. Unsupervised Sound Localization via Iterative Contrastive Learning. *arXiv preprint arXiv:2104.00315*.
- Liu, X.; Wu, Q.; Zhou, H.; Xu, Y.; Qian, R.; Lin, X.; Zhou, X.; Wu, W.; Dai, B.; and Zhou, B. 2022a. Learning Hierarchical Cross-Modal Association for Co-Speech Gesture Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Liu, X.; Xu, Y.; Wu, Q.; Zhou, H.; Wu, W.; and Zhou, B. 2022b. Semantic-Aware Implicit Neural Audio-Driven Video Portrait Generation. *arXiv preprint arXiv:2201.07786*.
- Owens, A.; and Efros, A. A. 2018. Audio-visual scene analysis with self-supervised multisensory features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 631–648.
- Patrick, M.; Asano, Y. M.; Fong, R.; Henriques, J. F.; Zweig, G.; and Vedaldi, A. 2020. Multi-modal self-supervision from generalized data transformations. *arXiv preprint arXiv:2003.04298*.
- Phan, H.; Koch, P.; Katzberg, F.; Maass, M.; Mazur, R.; McLoughlin, I.; and Mertins, A. 2017. What makes audio event detection harder than classification? In *2017 25th European signal processing conference (EUSIPCO)*, 2739–2743. IEEE.
- Proulx, M. J.; Brown, D. J.; Pasqualotto, A.; and Meijer, P. 2014. Multisensory perceptual learning and sensory substitution. *Neuroscience & Biobehavioral Reviews*, 41: 16–25.
- Pu, J.; Panagakis, Y.; Petridis, S.; and Pantic, M. 2017. Audio-visual object localization and separation using low-rank and sparsity. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2901–2905. IEEE.
- Qian, R.; Hu, D.; Dinkel, H.; Wu, M.; Xu, N.; and Lin, W. 2020. Multiple Sound Sources Localization from Coarse to Fine. *arXiv preprint arXiv:2007.06355*.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, 91–99.
- Senocak, A.; Oh, T.-H.; Kim, J.; Yang, M.-H.; and So Kweon, I. 2018. Learning to localize sound source in visual scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4358–4366.

- Stein, B. E.; and Meredith, M. A. 1993. *The merging of the senses*. The MIT Press.
- Tian, Y.; Hu, D.; and Xu, C. 2021. Cyclic Co-Learning of Sounding Object Visual Grounding and Sound Separation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2745–2754.
- Tian, Y.; Shi, J.; Li, B.; Duan, Z.; and Xu, C. 2018. Audio-visual event localization in unconstrained videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 247–263.
- Tian, Y.; and Xu, C. 2021. Can audio-visual integration strengthen robustness under multimodal attacks? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5601–5611.
- Tzinis, E.; Wisdom, S.; Remez, T.; and Hershey, J. R. 2021. Improving On-Screen Sound Separation for Open Domain Videos with Audio-Visual Self-attention. *arXiv preprint arXiv:2106.09669*.
- Vedaldi, A.; Zisserman, A.; Chen, H.; and Xie, W. 2020. Vggsound: a large-scale audio-visual dataset. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*. IEEE.
- Xu, X.; Zhou, H.; Liu, Z.; Dai, B.; Wang, X.; and Lin, D. 2021. Visually Informed Binaural Audio Generation without Binaural Audios. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*.
- Yosinski, J.; Clune, J.; Bengio, Y.; and Lipson, H. 2014. How transferable are features in deep neural networks? In *Advances in neural information processing systems*, 3320–3328.
- Zhao, H.; Gan, C.; Ma, W.-C.; and Torralba, A. 2019. The sound of motions. In *Proceedings of the IEEE International Conference on Computer Vision*, 1735–1744.
- Zhao, H.; Gan, C.; Rouditchenko, A.; Vondrick, C.; McDermott, J.; and Torralba, A. 2018. The sound of pixels. In *Proceedings of the European conference on computer vision (ECCV)*, 570–586.
- Zhou, H.; Liu, Y.; Liu, Z.; Luo, P.; and Wang, X. 2019a. Talking Face Generation by Adversarially Disentangled Audio-Visual Representation. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- Zhou, H.; Liu, Z.; Xu, X.; Luo, P.; and Wang, X. 2019b. Vision-Infused Deep Audio Inpainting. In *The IEEE International Conference on Computer Vision (ICCV)*.
- Zhou, H.; Sun, Y.; Wu, W.; Loy, C. C.; Wang, X.; and Liu, Z. 2021. Pose-controllable talking face generation by implicitly modularized audio-visual representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4176–4186.
- Zhou, H.; Xu, X.; Lin, D.; Wang, X.; and Liu, Z. 2020. Sep-Stereo: Visually Guided Stereophonic Audio Generation by Associating Source Separation. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Zou, Y.; Zhang, S.; Moura, J. M.; Yu, J.; and Tian, Y. 2020. Revisiting Mid-Level Patterns for Distant-Domain Few-Shot Recognition. *arXiv preprint arXiv:2008.03128*.