

Visually Informed Binaural Audio Generation without Binaural Audios

Xudong Xu^{1*} Hang Zhou^{1*} Ziwei Liu² Bo Dai² Xiaogang Wang¹ Dahua Lin¹
¹CUHK - SenseTime Joint Lab, The Chinese University of Hong Kong
²S-Lab, Nanyang Technological University

{xx018@ie, zhouhang@link, xgwang@ee, dhlin@ie}.cuhk.edu.hk, {ziwei.liu, bo.dai}@ntu.edu.sg

Abstract

*Stereophonic audio, especially binaural audio, plays an essential role in immersive viewing environments. Recent research has explored generating visually guided stereophonic audios supervised by multi-channel audio collections. However, due to the requirement of professional recording devices, existing datasets are limited in scale and variety, which impedes the generalization of supervised methods in real-world scenarios. In this work, we propose **PseudoBinaural**, an effective pipeline that is free of binaural recordings. The key insight is to carefully build pseudo visual-stereo pairs with mono data for training. Specifically, we leverage spherical harmonic decomposition and head-related impulse response (HRIR) to identify the relationship between spatial locations and received binaural audios. Then in the visual modality, corresponding visual cues of the mono data are manually placed at sound source positions to form the pairs. Compared to fully-supervised paradigms, our binaural-recording-free pipeline shows great stability in cross-dataset evaluation and achieves comparable performance under subjective preference. Moreover, combined with binaural recordings, our method is able to further boost the performance of binaural audio generation under supervised settings¹.*

1. Introduction

Auditory and visual experiences are implicitly but strongly connected. In immersive environments, the perception of sound is impacted by visual scenes [44]. Therefore, researchers have explored ways to generate stereophonic audios with visual guidance, in order to improve the user experience in multimedia products. Specifically, supervised learning methods [22, 14, 20, 45] have been considered for this purpose.

However, it is noteworthy that fully-supervised learning methods, despite the positive results that they achieve under constrained settings, would face significant difficulties in real-world applications. 1) They rely on videos associated with stereophonic recordings, which we refer to as “visual-stereo” pairs [22, 14]. Obtaining a high-quality collection of real stereo data requires complicated and professional recording systems (e.g. microphone arrays or dummy heads), thus is both resource-demanding and time-consuming. 2) The models trained on datasets collected under controlled environments may overfit to the layout of the rooms, rather than capturing the general associations between sound effects and the visual locations of the sound sources. The resultant models would also have poor generalization capability.

The privilege of learning representations from unlabeled data has been well discussed in different fields of deep learning [40, 18, 3, 23, 25]. This inspires us to explore an alternative approach, namely, to use only mono audios which can be acquired much more easily compared to binaural audios. We note that mono audios have been successfully applied in learning visually informed sound separation [9, 1, 13, 43, 42, 15]. Zhou *et al.* [45] recently leverage mono audios for stereo generation. However, their stereophonic learning procedure still depends on stereo data.

In this work, we propose **PseudoBinaural**, a novel pipeline that generates visually coherent binaural audios without accessing any recorded binaural data. Our key insight is to carefully build *pseudo visual-stereo* pairs from mono data. Two questions need to be identified in order to achieve our goal. Given a spatial location, 1) *what is the relationship between a mono audio and its binaural counterpart sourcing from that location?* 2) *How should visual cues be organized to represent the source visually?* Our solution is to utilize two mappings. A *Mono-Binaural-Mapping* to reproduce binaural audios of a single source positioned at any spatial location, and a *Visual-Coordinate-Mapping* that associates visual modality with spatial locations. Specifically, the Mono-Binaural-Mapping is achieved by adopting spherical harmonic decomposition [8]. A head-related impulse response (HRIR) [6] is then used to render binaural

*Equal contribution.

¹Code, models and demo videos are available at <https://sheldontsui.github.io/projects/PseudoBinaural>.

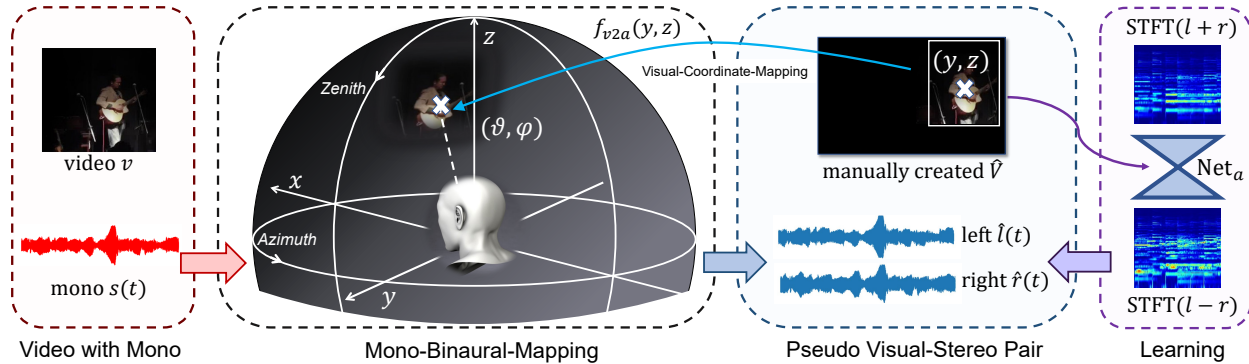


Figure 1: The pipeline of our method. Given one mono source, we create a *pseudo visual-stereo* pair $\{\hat{V}, (\hat{l}, \hat{r})\}$ by assigning the source direction $\vartheta = (\vartheta, \varphi)$ in the spherical coordinates according to our manually created \hat{V} . Then mono source $s(t)$ is converted to binaural channels $(\hat{l}(t, \vartheta), \hat{r}(t, \vartheta))$ through our Mono-Binaural-Mapping procedure by leveraging spherical harmonics decomposition. Within this pipeline, multiple sources can be linearly blended together to build training pairs. Then mono-to-binaural networks can be trained on the created pseudo data.

audios from the zero- and first-order terms of the decomposition. As for the Visual-Coordinate-Mapping, we pre-define a correspondence between pixel coordinates and spherical coordinates, so that we can easily manipulate visual content to meet the designation of the corresponding source direction.

Existing models for visually informed binaural audio generation can be readily adapted to train on our *pseudo visual-stereo* pairs. In order to make the best use of mono data, we further propose a new way of leveraging the task of audio-visual source separation [43, 15] to assist the training. The inference procedure is to simply apply the trained models to videos with mono audios and generate corresponding binaural audios. Our framework renders stable performances on two datasets and in-the-wild scenarios. Moreover, we can mix our *pseudo* data with real stereophonic recordings to further boost the performance of binaural audio generation under the supervised setting.

Our contributions can be summarized as follows: **1)** We identify the mapping between source directions and binaural audios with theoretical analysis. **2)** By manipulating the visual modality, pseudo visual-stereo pairs can be generated for model training without relying on any recorded binaural data. **3)** Extensive experiments validate the effectiveness and stability of our method on a variety of scenes. Moreover, our pseudo visual-stereo data can serve as a strong augmentation under the supervised setting.

2. Related Work

Visually Informed Stereophonic Audio Generation.

While stereo is strongly correlated with visual information, only few papers have proposed to guide the generation of stereo with vision. Li *et al.* [19] combine a synthesized early reverberation and a measured late reverberation tail for the generation of stereo sound in the desired room. However, the usage of such method is restricted to specific rooms

and serves for 360° videos. Morgado *et al.* [22] propose to recover ambisonics based on the datasets collected from YouTube. They assume that their end-to-end network is able to separate sound sources and reformulate them with learnable weights. Lu *et al.* [20] leverage flow with corresponding classifier for stereo generation. Specifically, Gao *et al.* [14] collect the FAIR-Play dataset using professional binaural audio collecting mics in a music room. Then they propose the Mono2binaural pipeline for converting mono audios to binaural ones in a U-Net framework. Their data is precious yet limited, models trained on their lab-collected data are difficult to generalize well on in-the-wild scenarios. Very recently, Zhou *et al.* [45] leverage mono data and propose to tackle stereophonic audio generation and source separation at the same time. Nevertheless, their method uses mono data to train separation only. All the above methods rely on recorded stereophonic data and visual-stereo pairs for training. We target to generate visually guided binaural audios without any binaural data.

Visually Indicated Sound Source Separation. The task of visually guided sound source separation aims at separating a mixed audio into independent ones, according to their sound source’s visual appearances. It has long been an interest of research for both human speech [10, 21, 26, 9, 1] and music [27, 13, 43, 42, 15, 11]. Recent learning-based methods [1, 43, 42, 15, 39] all leverage the **Mix-and-Separate** training pipeline that creates training pairs using collected solo data. Our work also exploits the same type of data to build training samples for binaural generation. We also adopt the setting of separating two sources [45] to boost the final performance.

Sound Source Localization. One of the most important features for human auditory system is to localize sound by the subtle differences of intensity, spectral and time cues between ears [7]. In the audio domain, previous research

mostly relies on microphone arrays to perform direction of arrival estimation [41]. Multi-modality works learn audio-visual associations [4, 5, 35, 34, 38, 46] and propose to localize the responses of sound in the visual domains [43, 31, 28, 17] a different type of “source localization”. Recent works [12, 36] propose to detect the position of vehicles with stereo audio, which deals with the opposite of our task. Normally, the visualization of activation is used to show auditorily associated visual information [4, 26, 25, 5, 42, 22, 14]. Our model also shows the ability of source localization by training only on our pseudo visual-stereo pairs.

3. Methodology

Different from previous completely learning-based and data-driven methods that rely on the ground-truth stereo, we train networks on self-created *pseudo visual-stereo* pairs. Our method is thus called **PseudoBinaural**. The overall pipeline is illustrated in Fig. 1.

3.1. Mapping Mono to Binaural

The key of our method relies on identifying the relationship between mono and stereo. This whole procedure, as illustrated in Fig. 2, is called *Mono-Binaural-Mapping*. Given a mono audio with an arbitrarily assigned source position $\vartheta = (\vartheta, \varphi)$ (Fig. 2 a), our goal is to first convert it to binaural channels with correct auditory sense of location. We empirically choose spherical harmonic decomposition for its expressive ability and its substantial connection with ambisonics (Fig. 2 b). Finally, the decomposed coefficients are transformed to virtual array and rendered to audios with HRIR (Fig. 2 c).

Spherical Harmonic Decomposition. The Laplace spherical harmonics represent a complete set of orthonormal basis defined on sphere surface [8]. The normalized form of spherical harmonics defined at azimuth angle ϑ and zenith angle φ in spherical coordinates can be represented as:

$$Y_l^m(\varphi, \vartheta) = N_l^{|m|} P_l^{|m|}(\cos \varphi) e^{jm\vartheta}, \quad (1)$$

where $P_l^{|m|}(\cos \varphi)$ is the associated Legendre polynomials, integer l is its order and m is the degree, limited to $[-l, l]$. $N_l^{|m|}$ is a normalization factor. Real spherical harmonics can serve as a type of generalized Fourier series, to decompose any function f :

$$f(\varphi, \vartheta) = \sum_{l=0}^{\infty} \sum_{m=-l}^l \Psi_l^m Y_l^m(\varphi, \vartheta). \quad (2)$$

The coefficients Ψ_l^m are the analogs of Fourier coefficients which can be represented as (* denotes the conjunction):

$$\Psi_l^m = \int_0^{2\pi} \int_0^\pi f(\varphi, \vartheta) Y_l^m(\varphi, \vartheta)^* \sin \varphi \, d\varphi d\vartheta. \quad (3)$$

Decomposed Coefficients for Mono Source. Here we follow the simplest assumption that only the impulse response from the direction $\vartheta = (\vartheta, \varphi)$ of a single sound source $s(t)$ is received, the Fourier coefficients can be derived from Eq. (3) as:

$$\Psi_l^m(\vartheta) = s(t) Y_l^m(\varphi, \vartheta). \quad (4)$$

This is the same as the encoding of ambisonics, where the $\Psi_l^m(\vartheta)$ can also be regarded as ambisonics’ components. For simplicity, ϑ is omitted in the following representations associated with this pre-defined direction.

The zero- and first-order components ($l = 0, 1$) that contribute most to 3D audio effect are leveraged in our model. Based on Eq. (1) and Eq. (4), the coefficients $\{\Psi_0^0, \Psi_1^1, \Psi_1^{-1}, \Psi_1^0\}$ (ϑ omitted) can be written as:

$$\begin{aligned} \Psi_0^0 &= s(t) N_0^0, \\ \Psi_1^1 &= s(t) N_1^1 \cos \varphi \cos \vartheta, \\ \Psi_1^{-1} &= s(t) N_1^1 \cos \varphi \sin \vartheta, \\ \Psi_1^0 &= s(t) N_1^0 \sin \varphi, \end{aligned} \quad (5)$$

which correspond to the W, X, Y and Z channels of ambisonics, respectively. W is the omnidirectional base channel, X, Y and Z are the orthogonal channels lie along 3D Cartesian axes as illustrated in Fig. 2 (b).

We adopt the Schmidt semi-normalization (SN3D) [37] to Eq. (1), which can be written as $N_l^m = \sqrt{(2 - \delta_m) \frac{(l-|m|)!}{(l+|m|)!}}$, where $\delta_m = 1$ if $m = 0$ else 0.

Binaural Decoding. Regarding the decomposed coefficients as ambisonic channels, we can roughly predict the left and right binaural channels $\hat{l}(t)$ and $\hat{r}(t)$ using simple transformation: $\hat{l}(t) = W + Y$ and $\hat{r}(t) = W - Y$. However, this paradigm is unable to recover real binaural.

On the other hand, binaural sound can be directly synthesized given a source position with the head-related impulse response (HRIR). One set of HRIR data can serve as filters $h_r(\vartheta)$ and $h_l(\vartheta)$ with respect to the direction ϑ of the sound source. The transferred binaural sound can be represented as $\hat{l}(t) = h_l(\vartheta) \otimes s(t)$ and $\hat{r}(t) = h_r(\vartheta) \otimes s(t)$, where \otimes is the convolution operation. However, open-sourced HRIR [2] are recorded in the free-field, thus cannot recover binaurals in a normal scene owing to reverberations.

Our solution is to leverage the binaural rendering technique that combines ambisonics with HRIR [24]. A virtual speaker array is pre-defined to make up for the reverberations as shown in Fig. 2 (c). Denoting the Fourier coefficients (ambisonic channels) as vectors $\Psi(\vartheta) = s(t) \mathbf{Y}(\vartheta) = [\Psi_0^0, \Psi_1^1, \Psi_1^{-1}, \Psi_1^0]^T$ (Refer to Eq. 5), we can further decompose $\Psi(\vartheta)$ into M virtual speakers at directions $\Theta = [\vartheta'_1, \dots, \vartheta'_M]$ to analog the multi-source effect caused by room reverberations. We define the matrix $D(\Theta) = [\mathbf{Y}(\vartheta'_1), \dots, \mathbf{Y}(\vartheta'_M)]$, each column representing

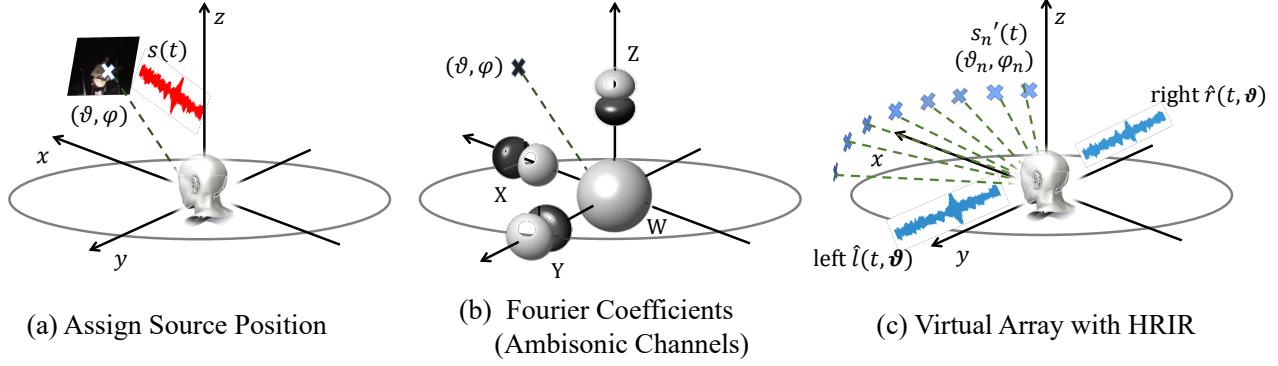


Figure 2: The steps for the Mono-Binaural-Mapping procedure. (a) Firstly the mono sound source $s(t)$ is assigned at direction $\vartheta = (\vartheta, \varphi)$. (b) Then through spherical harmonics decomposition of the source, we can derive the zero- and first-order spherical-based Fourier coefficients (which are also ambisonic channels). The figure represents the directions of the channels. (c) Finally, the Fourier coefficients can be transferred to a set of speaker array with fixed positions, and generate binaural with HRIR.

the harmonics for each virtual source. The virtual audio signals $\mathbf{s}'(t) = [s'_1(t), \dots, s'_M(t)]^T$ can be constraint as:

$$D(\Theta)\mathbf{s}'(t) = \Psi(\vartheta), \quad (6)$$

As the matrix $D(\Theta)^T D(\Theta)$ is of full-column rank, the virtual signals can be computed:

$$\mathbf{s}'(t) = (D(\Theta)^T D(\Theta))^{-1} D(\Theta)^T \Psi(\vartheta). \quad (7)$$

Finally, we take advantage of HRIR filters to acquire the desired left and right channels:

$$\begin{aligned} \hat{l}(t, \vartheta) &= \sum_{m=1}^M h_l(\vartheta'_m) \otimes s'_m(t), \\ \hat{r}(t, \vartheta) &= \sum_{m=1}^M h_r(\vartheta'_m) \otimes s'_m(t). \end{aligned} \quad (8)$$

The above is the binaural audio generated from a single mono source $s(t)$ from direction ϑ .

3.2. Creating Pseudo Visual-Stereo Pairs

With the Mono-Binaural-Mapping, there are two questions left to achieve visually informed binaural generation: 1) How to leverage the visual information, and 2) how to connect the direction of the sound source with visual information. To this end, we create *pseudo* visual information and define a *Visual-Coordinate-Mapping*.

Pseudo Visual Information Creation. As illustrated in Fig. 1, assuming that the listener is facing towards the x axis in 3D Cartesian coordinates, frontal-view scenes can thus be projected to the $y-z$ image plane. Given a video v_k with a single sound source, we can place the center of the sound source v'_k to a random position in the image plane. More specifically, the cropped frames v'_k are placed to a background image \hat{V} according to $\hat{V}(y, z) = v'_k$.

Visual-Coordinate-Mapping. We then define the mapping f_{v2a} from pixel position (y, z) on frontal-view images to spherical angles $\vartheta = f_{v2a}(y, z)$. In the spherical coordinate, the frontal-view image plane is defined as part of a cylinder centered at the coordinate origin as shown in Fig. 3 (a).

Based on the fact that the effective visual field of humans is approximately 120 degrees [33], we define the border azimuth angle as $\vartheta_{v0} = \pi/3$. So that objects within \hat{V} are distributed within $\vartheta \in [-\vartheta_{v0}, \vartheta_{v0}]$. The ratio between the height and width of the background image is set to $H/W = 1/2$, thus the top edge of the image is corresponding to $\varphi_{v0} = \pi/2 - \arctan(\pi/3)$. The range is $\varphi \in [\pi/2 - \arctan(\pi/3), \pi/2 + \arctan(\pi/3)]$. In this way, for each point in the image plane, we can find an angle in the spherical coordinate. We have also explored other field of view settings and find subtle differences.

Create Pseudo Visual-Stereo Pairs. By calculating the corresponding angle ϑ_k in spherical coordinates, a pair of binaural audios $\{\hat{r}_k(t, \vartheta_k), \hat{l}_k(t, \vartheta_k)\}$ can be recovered from the mono audio $s_k(t)$ accompanying $v_k(t)$ through Eq. (8).

Audio recordings collected in real-world scenarios are mostly mixed, therefore, we propose to mix multiple solo videos together in one scene to create *pseudo visual-stereo* pairs $\{\hat{V}, (\hat{l}, \hat{r})\}$. Each time, we assemble a random number of K independent mono videos $\mathbf{s}(t) = \{s_1(t), \dots, s_K(t)\}$ together to form a pseudo visual-stereo pair. The self-created binaural can be written as $\hat{l}(t) = \sum_{k=1}^K \hat{l}_k(t, \vartheta_k)$ and $\hat{r}(t) = \sum_{k=1}^K \hat{r}_k(t, \vartheta_k)$. The manually built visual information is \hat{V} , where $\hat{V}(\vartheta_k) = v_k, k \in [1, K]$.

Note that the patch size and audio amplitude are both directly proportional to the reciprocal of the depth. Therefore, the mono audio is firstly normalized according to its wave amplitude, and the corresponding cropped patch v'_k is normalized in the same scale as the audio. When the pseudo scene is assembled, v'_k is randomly resized and placed on \hat{V}

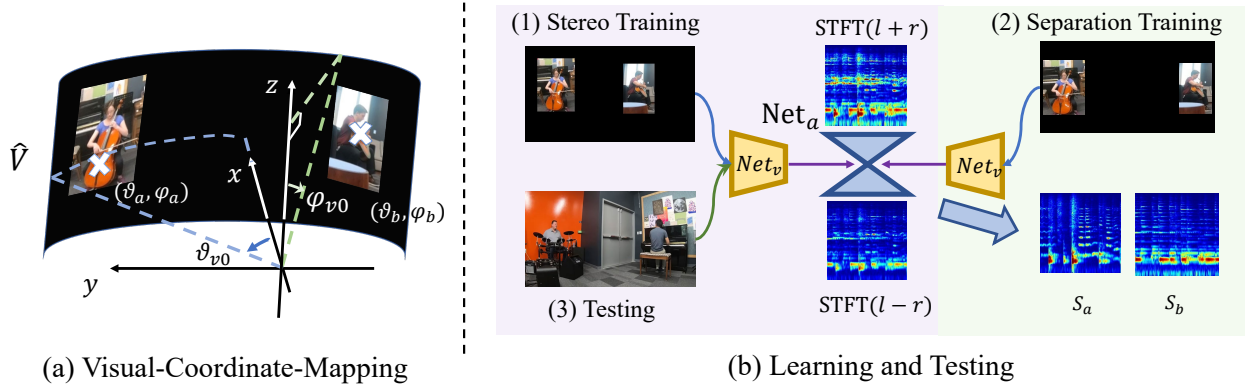


Figure 3: (a) Mapping from visual positions in the image domain to spherical angles. Normally we place the image \hat{V} at the frontal-view to be part of a cylinder. The borders of the image are corresponding to the angles φ_{v0} and ϑ_{v0} . (b) Network details. The input to the audio UNet Net_a is the STFT of a mono audio. The output is the prediction of the STFT difference between the left and right channels. During training, Net_v extract the visual feature f_v from our self-created image \hat{V} , and concatenate it to the audio UNet. During testing, this network can be applied to normal frames.

to represent objects at different depths. Only mono audios are used in building the data.

3.3. Learning

We leverage neural networks for learning from mono and visual guidance to pseudo binaural output. Previous networks and training paradigms from Mono2Binaural [14] and Sep-Stereo [45] can be readily adapted to train on our data. The learning procedure is depicted in Fig. 3 (b).

Stereo Training. The main part of our learning procedure is to directly train networks using our pseudo visual-stereo pairs. The whole training procedure is basically following [14, 45]. It consists of a backbone U-Net [30] audio network Net_a , and a ResNet18 [16] visual network Net_v . The audios are all processed in the complex Time-Frequency domain in form of Short-Time Fourier Transformation (STFT). Mono is created from the left and right channels $s_m(t) = \hat{l}(t) + \hat{r}(t)$ and the input to Net_a is the transformed mono spectrum $S_m = \text{STFT}(s_m(t))$. Net_a returns the complex mask \mathcal{M} for final predictions. The training objective is the difference of the left and right spectrums $S_D = \text{STFT}(\hat{l}(t) - \hat{r}(t))$, which can be written as:

$$\mathcal{L}_{\text{stereo}} = \|S_D - \mathcal{M} * S_m\|_2 \quad (9)$$

Then the predicted difference spectrum is transferred back to the difference audio $\tilde{s}_D = \text{ISTFT}(\mathcal{M} * S_m)$. The predicted left and right can be computed as $\tilde{l}(t) = (s_m(t) + \tilde{s}_D(t))/2$ and $\tilde{r}(t) = (s_m(t) - \tilde{s}_D(t))/2$.

Separation Training. Specifically, we leverage the task of separating two sources inspired by Sep-Stereo [45]. We care less about the performance on separation, but more about its benefits on distinguishing sound sources. Thus different from their visual feature rearrangement, we directly place two sources at separate edges when creating the pseudo

visual input \hat{V} (as shown in Fig. 3 (b)). The network input would be the pseudo pair $\{\hat{V}, (s_a, s_b)\}$, where s_a and s_b are the individual mono audio signals. Then we leverage one APNet branch from [45] to predict the original STFTs S_a and S_b . In this way, the backbone network can learn better the association between the sources' visual and audio information. Please refer to the supplementary materials for details.

4. Experiments

4.1. Datasets

We emphasize creating binaurals for music, which is an important scenario for stereo production. We will at first show our analysis on the FAIR-Play dataset, then introduce other datasets we use.

Revisiting FAIR-Play. Collected in a music room, FAIR-Play [14] is one of the most influential datasets in this field. However, by carefully examining the dataset, we find that the original train-test splits are somewhat problematic. The whole dataset contains 1,871 clips cut from several different long camera recordings with approximately the same camera view and scene layouts. The clips are randomly divided into 10 different train-test splits. As a result, the scenes within train and test splits are overlapped, probably originate from the same recording. This would lead to serious overfitting problems. The models might learn layouts of the room instead of visual-stereo association that we desired.

In order to evaluate the true generalization ability of different models on this dataset, we take efforts to re-organize the FAIR-Play dataset through reconstructing the original videos and re-splitting them. Specifically, we first run a clustering algorithm on all the clips to roughly group them according to the scenes. Then by matching the first and last

Table 1: Quantitative results of binaural audio generation on FAIR-Play and MUSIC-Stereo dataset. Except for SNR, the lower the score, the better the results. The upper half shows the results of standard benchmarks and our PseudoBinaural method. The lower half shows the augmentation results and our ablation studies on different binaural decoding schemes (Sec. 3.1). Our method outperforms previous methods when augmented with binaural recordings. Moreover, our chosen decoding scheme achieves the best performance among three decoding methods.

Method	FAIR-Play					MUSIC-Stereo				
	STFT	ENV	Mag	\mathcal{D}_{phase}	SNR \uparrow	STFT	ENV	Mag	\mathcal{D}_{phase}	SNR \uparrow
Mono-Mono	1.024	0.145	2.049	1.571	4.968	1.014	0.144	2.027	1.568	7.858
Mono2Binaural [14]	0.917	0.137	1.835	1.504	5.203	0.942	0.138	1.885	1.550	8.255
PseudoBinaural (w/o sep.)	0.951	0.140	1.914	1.539	5.037	0.953	0.139	1.902	1.564	8.129
PseudoBinaural (Ours)	0.944	0.139	1.901	1.522	5.124	0.943	0.139	1.886	1.562	8.198
Sep-Stereo [45]	0.906	0.136	1.811	1.495	5.221	0.929	0.135	1.803	1.544	8.306
Augment-HRIR	0.896	0.137	1.791	1.472	5.255	0.940	0.138	1.866	1.550	8.259
Augment-ambisonic	0.912	0.139	1.823	1.477	5.220	0.909	0.137	1.817	1.546	8.277
Augment-PseudoBinaural	0.878	0.134	1.768	1.467	5.316	0.891	0.132	1.762	1.539	8.419

frame of each clip within groups, we find the original order of the clips and concatenate them to recover the recorded videos. Finally, we select the videos whose scenes are completely absent in other videos as the validate and test sets. In this way, we create 5 different splits in which train and test sets are not overlapped. In our experiments, we re-train all supervised models (including Mono2Binaural [14] and Sep-Stereo [45]) and report the average results on the five splits. Please be noted that our model is also trained on this dataset, using only the solo part and *mono audios*.

MUSIC-Stereo [43, 42]. Containing 21 different types of instruments, MUSIC(21) is originally collected for visually guided sound separation. We select all the videos with binaural audio from MUSIC(21) and MUSIC-duet [43] to form a new dataset MUSIC-Stereo. Composed of solo and duet parts, it includes 1,120 unique videos of different musical performances. MUSIC-Stereo lasts 49.7 hours in total, which is 10 times larger than the FAIR-Play dataset. Following the post-processing steps in [14], we cut these videos into 17,940 10s clips and split them into training, validation, and test sets in an 8:1:1 ratio. Similar to FAIR-Play, only the solo part and *mono audios* are exploited for our model’s training.

YT-Music [22]. This dataset is collected from 360° videos on YouTube in the ambisonic format. The audios are transferred to binaural in the same way as our decoding scheme. With distinct vision configurations and stereo audio characteristics, YT-MUSIC is the most challenging dataset.

4.2. Evaluation Metrics

Previous Metrics. The evaluation protocol within this field is basically the STFT distance and the envelope distance (ENV) between recovered audios and recorded ones [22, 14]. The STFT distance represents the mean square error com-

puted on predicted spectrums, and the ENV distance is performed on raw audio waves through Hilbert transform [32]. To evaluate the predicted binaural audios more comprehensively, we also adopt two widely-used metrics *Magnitude Distance (Mag)* and *Signal-to-noise Ratio (SNR)* from [22]. The Mag distance reflects the L_2 deviation on the magnitude of spectrums and SNR is operated on the waveform directly. **Newly Proposed Metric.** In 3D audio sensation, audiences care more about sensing the source direction, where the phase of binaural audio is the key. As illustrated in [29], phase errors will introduce perceivable distortions but are always neglected during the optimization. Inspired by this, we further propose a new metric named *Difference Phase Distance (\mathcal{D}_{phase})*, which is performed on the Time-Frequency domain. Note that, the binaural audio is completely determined by the difference between left and right spectrums 3.3. Hence, \mathcal{D}_{phase} is to evaluate the phase distortion between the ground-truth difference S_D and the predicted one $\tilde{S}_D = \mathcal{M} * S_m$:

$$\mathcal{D}_{phase} = \|\text{phase}(S_D) - \text{phase}(\tilde{S}_D)\|_2, \quad (10)$$

where the phase is represented by the angle values, thus $\mathcal{D}_{phase} \in [0, 2\pi]$. It’s worth emphasizing that \mathcal{D}_{phase} is sensitive to the audio directions, *i.e.*, switching left and right channels would bring a significant change on this metric.

4.3. Quantitative Results

Binaural-Recording-Free Evaluation. Since no binaural-recording-free method has been proposed before, supervised method *Mono2Binaural* [14] whose backbone we borrow from, can be served as our baseline and upper bound. The evaluation is made on our newly-split FAIR-Play [14] and MUSIC-Stereo. For comparison, *Mono2Binaural* is trained with both visual frames and binaural audio, whereas our

Table 2: Cross-dataset evaluation results on five metrics. While the model trained on FAIR-Play is used for testing on the others, the model trained on MUSIC-Stereo is for the evaluation on FAIR-Play. PseudoBinaural presents better generalization ability than the supervised method Mono2Binaural on all datasets.

Dataset	Mono2Binaural [14]					PseudoBinaural (Ours)				
	STFT	ENV	Mag	\mathcal{D}_{phase}	SNR \uparrow	STFT	ENV	Mag	\mathcal{D}_{phase}	SNR \uparrow
FAIR-Play	0.996	0.142	1.993	1.562	5.876	0.959	0.140	1.917	1.496	6.057
MUSIC-Stereo	0.971	0.140	1.942	1.552	7.933	0.952	0.139	1.904	0.574	8.099
YT-MUSIC	0.717	0.118	1.435	1.597	9.214	0.653	0.111	1.306	1.357	9.848

method *PseudoBinaural* only leverages frames and mono audio to do the training. Please be noted that we do not rely on an extra dataset. The result of *Mono-Mono* is also listed, which copies the mono input two times as the stereo channels. This method should have no stereo effect at all, thus outperforming it means the success of generating the sense of directions. As the whole model of ours includes the separation training described in Sec. 3.3, we also evaluation the ablation of this module (**w/o sep.**). Table 1 shows the results of these methods on all metrics.

With no supervision, it is reasonable that our *PseudoBinaural* cannot outperform the supervised setting. However, the fact that our model outperforms *Mono-Mono* on all metrics proves the effectiveness of our proposed method. In line with previous work [45], introducing the separation task in the training framework can further improve the overall performance of generated binaural audio.

Augmentation to Binaural Audio Training. Since our method just relies on the pseudo visual-stereo pairs, a natural idea is to leverage both pseudo data and recorded ones to boost the performance of the traditional fully-supervised approach. As demonstrated in the lower half of Table 1, our method, denoted as *Augment-PseudoBinaural*, can surpass the traditional setting *Mono2Binaural* [14] on all 5 metrics. Moreover, compared to *Sep-Stereo* [45], which incorporates extra data, we create pseudo pairs with the same set of collected data, providing more effective and complementary information to guide the training. Consequently, our method outperforms theirs on both FAIR-Play and MUSIC-Stereo.

Cross-Dataset Evaluation. We specifically show the results of cross-dataset evaluation in Table. 2 to prove that 1) supervised methods can easily overfit to a specific domain and 2) the generalization ability of our method. YT-MUSIC [22] with special 360° videos and ambisonics sounds is also used for evaluation. Here we use the non-augmented version of *PseudoBinaural* for evaluation. For *Mono2Binaural*, the model evaluated on FAIR-Play is trained on MUSIC-Stereo, and the model tested on MUSIC-Stereo and YT-MUSIC is trained on FAIR-Play. During cross-testing on FAIR-Play and MUSIC-Stereo, the visual to angle mapping f_{v2a} is defined in the frontal view. But when cross-testing on YT-MUSIC, the video is defined in the form of 360°.

Table 3: Ablation study on the number K of mono videos to mix based on FAIR-Play dataset. When K is a mixture of the three different numbers, the ratio is empirically set to $1 : 2 : 3 = 0.4 : 0.5 : 0.1$.

K	STFT	ENV	Mag	\mathcal{D}_{phase}	SNR \uparrow
1	0.965	0.143	1.914	1.483	4.976
2	0.935	0.141	1.871	1.480	5.026
3	0.967	0.142	1.936	1.527	5.004
1,2	0.895	0.136	1.793	1.479	5.282
1,2,3	0.878	0.134	1.768	1.467	5.316

We can see from the table that our method stably outperforms *Mono2Binaural* in all cross-dataset evaluations. The supervised method tends to perform badly when testing in a different domain, while our recording-free method generalizes well by training only on mono data.

Ablation Study. The lower half of Table 1 presents our ablation studies on different binaural decoding schemes. As written in Sec. 3.1, binaural audios can be decoded directly from *HRIR* or *ambisonic*. It can be seen that our way of combining both leads to the best results.

When preparing pseudo visual-stereo pairs, the number of mono videos to mix is also another important hyperparameter for consideration. As shown in Table 3, a fixed mixing number K always fails to construct various training samples, introducing inconsistency with those naturally collected datasets. Hence, an empirical ratio of $1 : 2 : 3 = 0.4 : 0.5 : 0.1$ for the number K is applied to ensure the diversity of generated visual-stereo pairs.

Additionally, we evaluate the choice of visual field-of-view (FOV) when building the *Visual-Coordinate-Mapping*. The influential parameter is the border azimuth angle ϑ_{v0} which is set to $\pi/3$ (Sec. 3.2). The results are shown in Table 4, that our choice achieves the best results.

4.4. Qualitative Results

User Study². In total 30 users with normal hearing participated in our study to perform the quality evaluation. There

²Please refer to <https://sheldontsui.github.io/projects/PseudoBinaural> for demo videos.

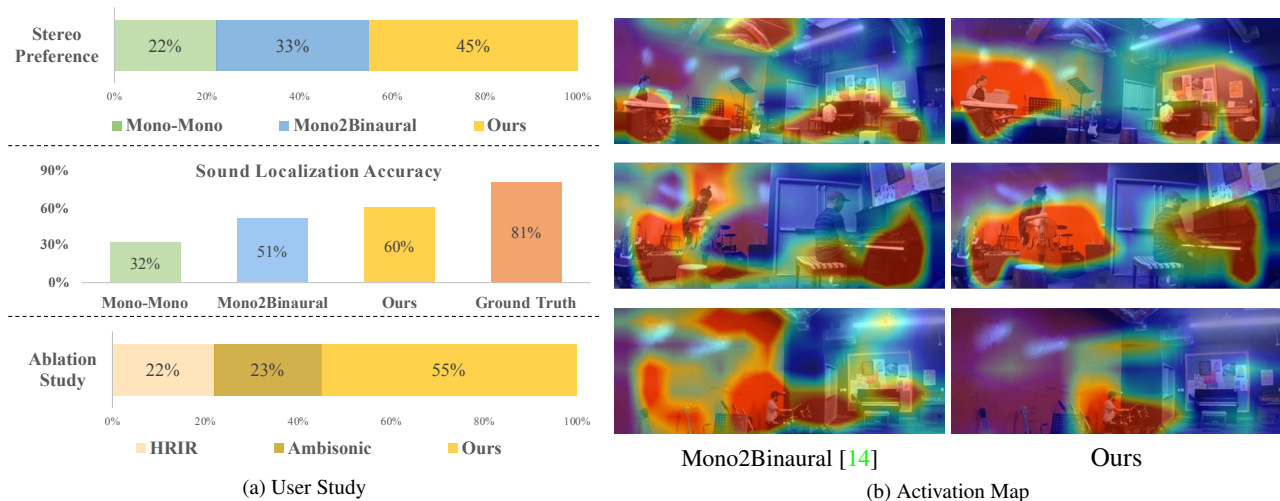


Figure 4: Qualitative results. (a) shows the results of our user studies. It can be seen that the users slightly prefer our approach over supervised Mono2Binaural [14]. (b) is the visualization of the activation maps of ours and Mono2Binaural. While the attention of theirs is messier, the results of ours are more compact. We focus more on sound sources.

are three sets of studies, each with 20 videos selected from the test set of FAIR-Play [14] and MUSIC-Stereo [42], most of which are duets. 1) The users are asked to watch one video and listen to the binaural audios generated by *PseudoBinaural*, *Mono2Binaural* [14] or *Mono-Mono*. The question is which one of the three methods creates the best stereo sensation. The results show the percentage of the users’ **Stereo Preferences**. 2) The users are asked to listen to the audio generated by the above methods without viewing videos, and decide where is the specific instrument (left, right, or center). Ground truth audios are also included for reference. The results show the **Sound Localization Accuracy** of these methods. 3) A subjective **Ablation Study** is conducted to show the influence of different choices of binaural decoding methods. The users are asked to tell which decoding diagram, direct HRIR, ambisonic, or ours, creates the best 3D hearing experience.

The results are shown in Fig 4 (a). From the first and second experiments, we can see that users find our method slightly better than supervised *Mono2Binaural* on both the two measurements. This is enough to validate that our results are highly competitive to supervised methods in subjective evaluations, which is extremely important for auditory tasks. In the sound localization experiment, users can only achieve 81% accuracy even given the ground-truth audio, which demonstrates the misleading caused by the room reverberations. The subjective ablation study shows that our decoding procedure apparently creates the best sense of hearing among all decoding choices.

Visualization. We also visualize the activation map generated by our method and *Mono2Binaural* [14] on the visual domain. In Fig. 4 (b) we can see that *Pseu-*

Table 4: Ablation study on the border azimuth angle ϑ_{v0} . The horizontal visual field-of-view is $2\vartheta_{v0}$.

ϑ_{v0}	$\pi/6$	$\pi/4$	$\pi/3$	$5\pi/12$	$\pi/2$
STFT ↓	0.923	0.896	0.878	0.884	0.886
SNR ↑	5.138	5.181	5.316	5.302	5.271

doBinaural can successfully attend to sound sources while *Mono2Binaural* [14] would focus on less important areas. For example, their approach would attend to the ceiling for all three scenes shown, which is not the sound source.

5. Conclusion

In this work, we propose *PseudoBinaural*, a binaural-recording-free method for generating binaural audios from corresponding mono audios and visual cues. For the first time, the problem of visually informed binaural audio generation is tackled without binaural audio recordings. Based on the theoretical analysis of Mono-Binaural-Mapping, the created pseudo visual-stereo pairs can be capitalized to train models for binaural audio generation. Extensive experiments validate that our framework can be very competitive both quantitatively and qualitatively. More impressively, augmented with real binaural audio recordings, our *PseudoBinaural* could outperform current state-of-the-art methods on various standard benchmarks.

Acknowledgements. This work is supported by the Collaborative Research Grant from SenseTime (CUHK Agreement No. TS1712093), the General Research Fund (GRF) of Hong Kong (No. 14205719, 14202217, 14203118, 14208619), NTU NAP and A*STAR through the Industry Alignment Fund - Industry Collaboration Projects Grant.

References

- [1] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. The conversation: Deep audio-visual speech enhancement. *Proceedings of the Interspeech*, 2018. 1, 2
- [2] V Ralph Algazi, Richard O Duda, Dennis M Thompson, and Carlos Avendano. The cipic hrtf database. In *Proceedings of the 2001 IEEE Workshop on the Applications of Signal Processing to Audio and Acoustics (Cat. No. 01TH8575)*, pages 99–102. IEEE, 2001. 3
- [3] Humam Alwassel, Dhruv Mahajan, Lorenzo Torresani, Bernard Ghanem, and Du Tran. Self-supervised learning by cross-modal audio-video clustering. *arXiv preprint arXiv:1911.12667*, 2019. 1
- [4] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 3
- [5] Relja Arandjelovic and Andrew Zisserman. Objects that sound. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 3
- [6] Durand R Begault and Leonard J Trejo. 3-d sound for virtual reality and multimedia. 2000. 1
- [7] Jens Blauert. *Spatial hearing: the psychophysics of human sound localization*. MIT press, 1997. 2
- [8] R. Courant and D. Hilbert. *Methods of mathematical physics*. Wiley classics library. Interscience Publishers, 1962. 1, 3
- [9] Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinatan Hassidim, William T Freeman, and Michael Rubinstein. Looking to listen at the cocktail party: a speaker-independent audio-visual model for speech separation. *ACM Transactions on Graphics (TOG)*, 37(4), 2018. 1, 2
- [10] John W Fisher III, Trevor Darrell, William T Freeman, and Paul A Viola. Learning joint statistical models for audio-visual fusion and segregation. In *Advances in neural information processing systems (NeurIPS)*, 2001. 2
- [11] Chuang Gan, Deng Huang, Hang Zhao, Joshua B Tenenbaum, and Antonio Torralba. Music gesture for visual sound separation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [12] Chuang Gan, Hang Zhao, Peihao Chen, David Cox, and Antonio Torralba. Self-supervised moving vehicle tracking with stereo sound. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 3
- [13] Ruohan Gao, Rogerio Feris, and Kristen Grauman. Learning to separate object sounds by watching unlabeled video. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 1, 2
- [14] Ruohan Gao and Kristen Grauman. 2.5 d visual sound. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 3, 5, 6, 7, 8
- [15] Ruohan Gao and Kristen Grauman. Co-separating sounds of visual objects. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 1, 2
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 770–778, 2016. 5
- [17] Di Hu, Rui Qian, Minyue Jiang, Xiao Tan, Shilei Wen, Errui Ding, Weiyao Lin, and Dejing Dou. Discriminative sounding objects localization via self-supervised audiovisual matching. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 3
- [18] Bruno Korbar, Du Tran, and Lorenzo Torresani. Cooperative learning of audio and video models from self-supervised synchronization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. 1
- [19] Dingzeyu Li, Timothy R. Langlois, and Changxi Zheng. Scene-aware audio for 360 videos. *ACM Trans. Graph.*, 37(4), 2018. 2
- [20] Yu-Ding Lu, Hsin-Ying Lee, Hung-Yu Tseng, and Ming-Hsuan Yang. Self-supervised audio spatialization with correspondence classifier. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 3347–3351. IEEE, 2019. 1, 2
- [21] Hari Krishna Maganti, Daniel Gatica-Perez, and Iain McCowan. Speech enhancement and recognition in meetings with an audio-visual sensor array. *IEEE Transactions on Audio, Speech, and Language Processing*, 15(8):2257–2269, 2007. 2
- [22] Pedro Morgado, Nuno Nvasconcelos, Timothy Langlois, and Oliver Wang. Self-supervised generation of spatial audio for 360 video. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. 1, 2, 3, 6, 7
- [23] Pedro Morgado, Nuno Vasconcelos, and Ishan Misra. Audio-visual instance discrimination with cross-modal agreement. *arXiv preprint arXiv:2004.12943*, 2020. 1
- [24] Markus Noisternig, Alois Sontacchi, Thomas Musil, and Robert Holdrich. A 3d ambisonic based binaural sound reproduction system. In *Audio Engineering Society Conference: 24th International Conference: Multichannel Audio, The New Reality*. Audio Engineering Society, 2003. 3
- [25] Andrew Owens and Alexei A Efros. Audio-visual scene analysis with self-supervised multisensory features. *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 1, 3
- [26] Andrew Owens, Phillip Isola, Josh McDermott, Antonio Torralba, Edward H Adelson, and William T Freeman. Visually indicated sounds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 3
- [27] Sanjeel Parekh, Slim Essid, Alexey Ozerov, Ngoc QK Duong, Patrick Pérez, and Gaël Richard. Motion informed audio source separation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6–10. IEEE, 2017. 2
- [28] Rui Qian, Di Hu, Heinrich Dinkel, Mengyue Wu, Ning Xu, and Weiyao Lin. Multiple sound sources localization from coarse to fine. *ECCV*, 2020. 3
- [29] Alexander Richard, Dejan Markovic, Israel D. Gebru, Steven Krenn, Gladstone Alexander Butler, Fernando Torre, and Yaser Sheikh. Neural synthesis of binaural speech. In *International Conference on Learning Representations (ICLR)*, 2021. 6
- [30] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation.

- In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 5
- [31] 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 (CVPR)*, 2018. 3
- [32] Julius Orion Smith. *Mathematics of the discrete Fourier transform (DFT): with audio applications*. 2007. 6
- [33] John Smythies. A note on the concept of the visual field in neurology, psychology, and visual neuroscience. *Perception*, 25(3):369–371, 1996. 4
- [34] Yapeng Tian, Dingzeyu Li, and Chenliang Xu. Unified multisensory perception: weakly-supervised audio-visual video parsing. *ECCV*, 2020. 3
- [35] Yapeng Tian, Jing Shi, Bochen Li, Zhiyao Duan, and Chenliang Xu. Audio-visual event localization in unconstrained videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 3
- [36] Arun Balajee Vasudevan, Dengxin Dai, and Luc Van Gool. Semantic object prediction and spatial sound super-resolution with binaural sounds. *ECCV*, 2020. 3
- [37] DE Winch, DJ Ivers, JPR Turner, and RJ Stening. Geomagnetism and schmidt quasi-normalization. *Geophysical Journal International*, 2005. 3
- [38] Yu Wu, Linchao Zhu, Yan Yan, and Yi Yang. Dual attention matching for audio-visual event localization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 3
- [39] Xudong Xu, Bo Dai, and Dahua Lin. Recursive visual sound separation using minus-plus net. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [40] Karren Yang, Bryan Russell, and Justin Salamon. Telling left from right: Learning spatial correspondence of sight and sound. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1
- [41] Qilin Zhang, Habti Abeida, Ming Xue, William Rowe, and Jian Li. Fast implementation of sparse iterative covariance-based estimation for source localization. *The Journal of the Acoustical Society of America*, 2012. 3
- [42] Hang Zhao, Chuang Gan, Wei-Chiu Ma, and Antonio Torralba. The sound of motions. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 1, 2, 3, 6, 8
- [43] Hang Zhao, Chuang Gan, Andrew Rouditchenko, Carl Vondrick, Josh McDermott, and Antonio Torralba. The sound of pixels. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 1, 2, 3, 6
- [44] Hang Zhou, Ziwei Liu, Xudong Xu, Ping Luo, and Xiaogang Wang. Vision-infused deep audio inpainting. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 1
- [45] Hang Zhou, Xudong Xu, Dahua Lin, Xiaogang Wang, and Ziwei Liu. Sep-stereo: Visually guided stereophonic audio generation by associating source separation. *ECCV*, 2020. 1, 2, 5, 6, 7
- [46] Ye Zhu, Yu Wu, Hugo Latapie, Yi Yang, and Yan Yan. Learning audio-visual correlations from variational cross-modal generation. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021. 3