

Sound Localization by Self-Supervised Time Delay Estimation

Ziyang Chen, David F. Fouhey, and Andrew Owens

University of Michigan

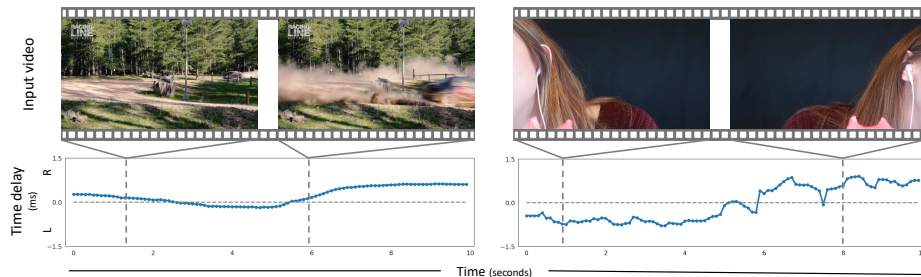


Fig. 1: Given a stereo audio recording, we estimate a sound’s *interaural time delay*. Our model learns through self-supervision to find correspondences between the signals in each channel, from which the time delay can be estimated. We show time delay predictions for two scenes, along with their corresponding video frames (not used by the model). In both cases, the sound source changes its position in a scene, resulting in a corresponding change in time delay.

Abstract. Sounds reach one microphone in a stereo pair sooner than the other, resulting in an *interaural time delay* that conveys their directions. Estimating a sound’s time delay requires finding correspondences between the signals recorded by each microphone. We propose to learn these correspondences through self-supervision, drawing on recent techniques from visual tracking. We adapt the contrastive random walk of Jabri et al. to learn a cycle-consistent representation from unlabeled stereo sounds, resulting in a model that performs on par with supervised methods on “in the wild” internet recordings. We also propose a multimodal contrastive learning model that solves a *visually-guided* localization task: estimating the time delay for a particular person in a multi-speaker mixture, given a visual representation of their face. Project site: <https://ificl.github.io/stereocrw>.

1 Introduction

Sounds in the world arrive at one of our two ears slightly sooner than the other. This *interaural time delay*, which generally lasts only a few hundred microseconds, indicates a sound’s direction and thus provides an important cue for multimodal perception. In humans, for example, time delays convey the positions of objects that move out of sight, and are integrated with visual cues when localizing events [53]. Visual information can also guide the sound localization process, allowing us to find a particular event of interest through binaural cues, while ignoring the others.

While high-quality stereo sound recordings are now abundant, such as in the audio tracks of videos recorded by consumer phones, existing methods often struggle to localize sound sources within them, particularly when they contain correlated noise or multiple sound sources. The localization problem has typically been addressed by matching hand-crafted features [51,67] and, recently, by supervised learning [44,62,21]. However, the difficulty in acquiring natural labeled data has limited their effectiveness. Many approaches, consequently, resort to using simulated training data that may not be fully representative of the world.

We propose to address these problems by learning time delay estimation from real, unlabeled recordings. We take inspiration from work in self-supervised visual tracking that learns space-time correspondences from videos, such as through cycle consistency [78,46,5] and contrastive learning [79]. Analogously, our approach is based on learning audio embeddings that can be used to find *interaural correspondences*: pairs of sounds from different stereo channels that correspond to the same underlying events.

We introduce a model, inspired by the *contrastive random walk* of Jabri et al. [46], that learns cycle consistent features from unlabeled stereo sound. This model maximizes the return probability of a random walk on a graph whose nodes correspond to the audio samples in each channel. In this graph, edges connect samples between channels, and the walk’s transition probabilities are defined by learned embeddings. We show examples of time delay estimates for two real-world videos in Figure 1.

We also propose a model inspired by instance discrimination [23,81,40,12] that can perform a novel *visually-guided* time delay estimation task: localizing a speaker in a multi-speaker audio recording, given only their visual appearance. The resulting model is simple and can accurately localize speakers, without the need for explicitly separating sounds in the mixture. We also use this approach to train audio-based localization models solely from mono audio, which in some domains may be more readily available than stereo sound. This model uses data augmentation to incorporate knowledge about invariances to important sources of variation.

Through experiments on simulated environments with metrically accurate ground truth, and on internet videos with directional judgments annotated by human listeners, we show:

- Interaural time delays can be accurately estimated through self-supervised learning, using either unlabeled stereo and mono training data.
- Our models provide robustness to distracting sounds within a mixture, and perform well on real-world recordings, obtaining competitive performance with state-of-the-art supervised methods.
- Visual signals allow our models to localize specific speakers within mixtures.

2 Related Work

Human binaural localization. Humans use two main cues for estimating the azimuth of a sound: interaural time differences (ITD) and interaural intensity differences (IID), i.e., the difference in the loudness of the sounds entering both

ears [64,75]. In practice, IID is primarily useful for high-frequency sounds that are close to the observer, while ITD is useful for low-frequency sounds and is relatively unaffected by distance [8]. Our work is thus complementary to methods that use IID cues. Humans can accurately estimate azimuth from multi-source mixtures [86,39], and integrate vision with binaural cues [53], motivating our work on multi-speaker time delay estimation.

Time delay estimation. Time delay estimation is a classic signal processing problem. In early work, Carter et al. [9] estimated time delays using generalized cross-correlation with phase transform (GCC-PHAT), which corresponds to a maximum likelihood estimate under low noise [51,7,87]. Other work uses beamforming [22] or subspace methods [67]. Comanducci et al. [18] trained a convolutional network (CNN) to denoise GCC-PHAT features. Other work trains a multi-layer perceptron to predict a time delay from a raw waveform [44], trains recurrent networks on hand-crafted features [62,1], and uses 3D CNNs [21]. Christensen et al. [14] used GCC-PHAT and echolocation to estimate depth maps from audio. In concurrent work, Chen et al. [13] localized multiple sounds by jointly solving source separation and time delay estimation problems. Time delay estimation also has a wide range of applications and modalities, such as oceanography [6], wireless networking [85,61], sonar [10], and possibly directional olfaction [63]. In contrast, we pose time delay estimation as a self-supervised learning problem, and we do not require hand-crafted features or labels.

Supervised binaural localization. Vecchiotti et al. [73] estimated sound direction directly from raw waveforms. Other work uses Short-time Fourier Transform [83,11,2] or beamforming features [65]. Due to the challenge in obtaining labeled data, these methods have largely been trained on synthetic or lab-collected data. In contrast to these approaches, we learn a specific (but widely useful) cue—the time delay—through self-supervision on natural data.

Audio-visual binaural learning. Yang et al. [84] distinguished between audio-visual examples in which the stereo channels have (or have not been) swapped, resulting in a representation that can be finetuned to solve localization tasks. In contrast, our model can optionally be trained and deployed solely with audio, and produces an output—the time delay—that is directly correlated with sound direction, without the need for finetuning. Gan et al. [30] used a car detector to provide pseudo ground truth for a sound-based localization method. Since the training data comes from a supervised car detector, the model relies on labeled training data, whereas ours is self-supervised. Later work [19,72] extends this approach by distilling supervision from multiple visual classifiers and modalities. Other work generates stereo sound from mono audio using images [32,34,82], largely by adjusting the relative volume of the channels to simulate IID cues.

Audio-visual sound localization and separation. A variety of methods have been proposed for using vision to localize and separate sounds. Classic work searches for cross-modal similarity in statistical models [42,28,49]. Later work uses contrastive learning to find image regions that are highly correlated with sound [70,4,59,58,88], and separates sounds from synthetic mixtures [26,3,58,31,29]. Recent work has applied the contrastive random walk to

localize multiple sounds within images [45]. This method learns correspondences between image patches and (mono) audio, whereas our model learns correspondences between the signals in each stereo channel.

Audio self-supervision. A variety of methods have been proposed for learning audio representations through self-supervision, typically for semantic recognition tasks, such as music or speech understanding. These include contrastive learning [57,68,77,47,36,76], autoencoding [25], multi-task learning with pretext tasks [60], and generative autoregressive models [17]. In contrast, we learn a representation for learning interaural correspondence in binaural audio.

Learning visual correspondences. We take inspiration from methods that learn space-time correspondences from video. These include methods that colorize grayscale video [74], cycle-consistent feature representations [78,46,38] and slow features [80,37]. Other work [79] has shown that features learned through instance discrimination [12,81,24,40] are effective for tracking. However, these methods have not been applied to learning stereo audio correspondences. In our models, by contrast, vision is used to aid the audio matching process. We adapt several of these methods [46,37,12,5] to learn correspondences between temporal samples of audio for binaural matching.

3 Method

The goal of the time delay estimation problem is to determine how much sooner a sound reaches one microphone than another¹. Given the two channels of a stereo recording, $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^n$, represented as waveforms, and a function $h : \mathbb{R}^n \mapsto \mathbb{R}^{n \times d}$ that computes features for each temporal sample, a common solution is to choose a time delay τ that maximizes the generalized cross-correlation [51]:

$$R_{\mathbf{x}_1, \mathbf{x}_2}(\tau) = \mathbb{E}_t [\mathbf{h}_1(t) \cdot \mathbf{h}_2(t - \tau)], \quad (1)$$

where $\mathbf{h}_i = h(\mathbf{x}_i)$ are the features for \mathbf{x}_i , and $\mathbf{h}_i(t)$ is the d -dimensional feature embedding for time t .

Traditionally, the audio features, h , are defined using hand-crafted features. For example, the widely-used Generalized Cross Correlation with Phase Transform (GCC-PHAT) [51] whitens the audio by dividing by the magnitude of the cross-power spectral density. This approach provides the maximum likelihood solution under certain ideal, low-noise conditions [51,7,87].

We propose, instead, to learn h through self-supervision from unlabeled data. These features ought to capture interaural correspondences: observations in both waveforms that were generated by the same underlying events should be close in embedding space. We consider models that can be trained solely from unlabeled stereo or mono sound (Sec. 3.1), or that learn to perform visually-guided estimation from audio-visual data (Sec. 3.2).

3.1 Learning interaural correspondence

We propose models that learn interaural correspondence from unlabeled data.

¹ This quantity is also known as the *time difference of arrival* (TDOA) or alternatively as the *interaural time difference* or *delay* (ITD).

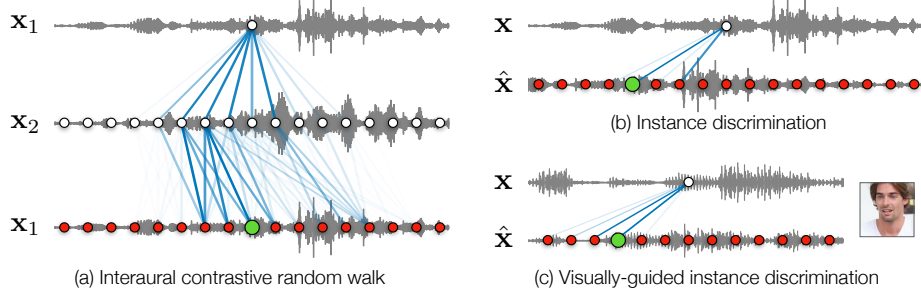


Fig. 2: **Learning interaural correspondence.** We consider several self-supervised models. (a) A random walk moves from one stereo channel to another, then back, with transition probabilities determined by our learned embeddings. We learn to maximize the probability that it returns to the node where it started (marked in green). (b) We apply data augmentation to mono audio, and learn embeddings that maximize the similarity of corresponding timesteps. (c) We learn to match audio for a single speaker from a multi-speaker mixture, given a visual input.

Contrastive random walks. Our embeddings should provide *cycle consistent* matches: the process of matching features from \mathbf{x}_1 to those in \mathbf{x}_2 should yield the same correspondences as matching in the opposite direction, from \mathbf{x}_2 to \mathbf{x}_1 . We use this idea to learn a representation from unlabeled stereo sounds.

We adapt the contrastive random walk model of Jabri et al. [46] to binaural audio (Fig. 2a). We create a graph that contains nodes for each of the temporal sample $\mathbf{x}_i(t)$ from both channels, with edges connecting the nodes that come from different channels.² We then perform a random walk that transitions from nodes in \mathbf{x}_1 to those in \mathbf{x}_2 , then back to \mathbf{x}_1 , with transition probabilities that are defined by dot products between embedding vectors:

$$A_{ij}(s, t) = \frac{\exp(\mathbf{h}_i(s) \cdot \mathbf{h}_j(t)/c)}{\sum_{k=1}^n \exp(\mathbf{h}_i(s) \cdot \mathbf{h}_j(k)/c)}, \quad (2)$$

where $A_{ij}(s, t)$ is the probability of transitioning from sample s in \mathbf{x}_i to sample t in \mathbf{x}_j , and a temperature constant c . The features $\mathbf{h}_i = h(\mathbf{x}_i; \theta)$ are parameterized with network weights θ and are represented using a CNN (Sec. 3.3). We maximize the log return probability of a walk that moves between nodes in the two channels:

$$\mathcal{L}_{\text{crw}} = -\frac{1}{n} \text{tr}(\log(A_{12}A_{21})), \quad (3)$$

where the log is computed element-wise. We also found it helpful to incorporate knowledge about invariances to important sources of variation, such as to noise. To do this, we also apply data augmentation to two audio channels during the walk similar to Hu et al. [45] (see supp. for details).

Slow features. We also train a variation of the model that learns to associate embeddings that temporally co-occur, taking inspiration from methods that learn slow features in video [80, 37] and audio-visual synchronization [16, 58, 52]. These

² Following visual tracking work [46], one could potentially extend this approach to microphone arrays with 3 or more channels by performing the walk over all channels.

pairs of embeddings are more likely (than misaligned timestamps) to correspond to the same events. We minimize:

$$\mathcal{L}_{\text{zero}} = -\frac{1}{n} \text{tr}(\log(A_{12})), \quad (4)$$

where A_{12} is defined as in Eq. 2.

Instance discrimination. We also consider models that can be trained solely with mono audio using *instance discrimination* [81]. In lieu of a second audio channel, we create synthetic views of mono audio, using data augmentation that encourages invariances that are likely to be useful for interaural matching. We minimize:

$$\mathcal{L}_{\text{dis}} = -\log \frac{\exp(\mathbf{h}(t) \cdot \hat{\mathbf{h}}(t)/c)}{\sum_{k=1}^n \exp(\mathbf{h}(t) \cdot \hat{\mathbf{h}}(k)/c)}, \quad (5)$$

over all timesteps t , where $\mathbf{h} = h(\mathbf{x})$ are the features for a mono audio \mathbf{x} , and $\hat{\mathbf{h}} = h(\hat{\mathbf{x}})$ are features computed from an augmented version of \mathbf{x} .

Unless otherwise specified, we perform two types of augmentation: time shifting and volume adjustment. To model the challenges in time delay estimation, we choose negative examples exclusively from \mathbf{x} , rather than other examples in the batch [16,58,52]. We ensure that augmented positive views are always taken from the corresponding timestep, i.e., we undo any time-shifting augmentation when indexing $\hat{\mathbf{h}}(t)$.

3.2 Visually-guided time delay estimation

We also apply our model to the novel problem of estimating the time delay for a single sound within a mixture using visual information. Given a sound mixture containing multiple simultaneous speakers, we estimate the time delay for one object, given a visual representation of its appearance (e.g., localizing a speaker using a visual representing their face). The visual input need not co-occur with the audio. For example, the object may be off-screen, or its visual features may have been extracted at an earlier time.

We adapt the instance discrimination variation of our model, with a training procedure that resembles the “mix-and-separate” [88] paradigm used in audio-visual source separation [26,3,58,31,29]. We create a mixture from two sounds, each with its own delay, and ask the model to estimate the delay from only the desired source. Given two audio tracks \mathbf{u} and \mathbf{v} , we create a synthetic binaural sound mixture $\mathbf{x}_1 = \mathbf{u} + \mathbf{v}$ and $\mathbf{x}_2 = \text{shift}(\mathbf{u}, \tau_u) + \text{shift}(\mathbf{v}, \tau_v)$ for randomly sampled values τ_u and τ_v , where $\text{shift}(\mathbf{x}, \tau)$ shifts \mathbf{x} by τ . The model is also provided with I_u , an image depicting \mathbf{u} . We learn audio-visual features by minimizing:

$$\mathcal{L}_{\text{av}} = -\log \frac{\exp(\mathbf{g}_1(t) \cdot \mathbf{g}_2(t - \tau_u)/c)}{\sum_{k=1}^n \exp(\mathbf{g}_1(t) \cdot \mathbf{g}_2(k)/c)}, \quad (6)$$

over all timesteps t , where $\mathbf{g}_i = g(\mathbf{x}_i, I_u)$ are the learned audio-visual features for channel \mathbf{x}_i . Here, g obtains its embedding by fusing audio from one channel with the input image. As in the instance discrimination model, we apply augmentation to \mathbf{g}_2 . Note that this task cannot be solved without I_u : from audio alone, the model would be unable to determine whether the true delay is τ_u or τ_v .

3.3 Learning a time delay estimation model

We now describe how these self-supervised learning models can be trained, and how they can be used to estimate time delays.

Network architectures. We implement the audio embedding h using a CNN that operates on spectrograms [43]. To compute the embedding for sample s , we extract a waveform of length T centered on s . We create a spectrogram representation of size $128 \times 128 \times 2$ using a Short-time Fourier transform (STFT). We keep both magnitude and phase and provide them as input to a ResNet [41], which extracts a $d = 128$ dimensional, ℓ_2 -normalized embedding.

For our audio-visual model, we represent the visual information using (pre-trained) FaceNet [69]. This allows our model to estimate attributes of speakers from face crops, similar to the work in source separation that uses face embeddings [26]. We fuse the audio and visual features after the second convolution block of the audio subnetwork by concatenating the 128-dimensional visual features at each time-frequency position.

Datasets. We train our audio models on datasets of stereo sound: **FAIR-Play** [32], which has 1,871 videos (5.2 hours) of lab-collected music performances from a small number of rooms and **Free-Music-Archive** (FMA) [20], a dataset of 101K (841 hours) music recordings created by a large number of artists.

For the visually-guided model, we train our model on VoxCeleb2 [15] with 500 randomly selected identities. We randomly create training mixtures from mono sound, without the speaker identity labels and without using a simulator.

Training. We use the AdamW optimizer [55,50] with a learning rate $= 10^{-4}$, a cosine decay learning rate scheduler, a batch size of 48, a temperature $c = 0.05$ following [81], and early stopping. Please see supp. for more training details.

Self-supervised learning formulation. For all models, we extract our examples from a 1220-sample waveform, sampled at 16 Khz. We obtain our embeddings using a sliding window of size 0.064s (1024 samples), with a step of 4 samples, yielding 49 audio clips. We apply random stereo channel swapping and channel-wise waveform rescaling for augmentation in all models. In some experiments, we also add random noise, add reverberation, and mix in other sounds as additional augmentation. At test time, we obtain a denser audio graph by using the step of 1 sample for the sliding window. Our model can use input sounds with a variety of durations without retraining, due to the fully-convolutional network architecture [54]. For the audio-visual task, we use a window length of 0.96s or 2.55s to obtain more temporal context, since this problem involves jointly solving a separation task.

Estimating delays from features. After learning our representation h , we can use it to estimate the time delay, such as by maximizing $R_{\mathbf{x}_1, \mathbf{x}_2}$ (Eq. 1). We have found that this procedure can affect the quality of the prediction for both learned and hand-crafted methods (e.g., due to outliers), so we evaluate a number of different variations in our experiments. In our approach, each embedding votes on a value for τ . We then choose a single time delay for the audio from these votes, either by taking the mean or by using a RANSAC-like [27] mode estimation

method. In the latter, we first select the delay with the most votes, then average the inliers (those within a small threshold of the chosen value). This vote can be performed by the nearest neighbor search, or by treating the learned similarities as probabilities (Eq. 2) and taking the expectation, *i.e.*, $\frac{1}{n} \sum_{s,\tau} \tau A_{12}(s, \tau)$ [5].

4 Experiments

We evaluate our methods using both simulated audio, where time delays can be measured exactly, and real-world binaural audio from unknown microphone geometry, where quantized sound direction categories are labeled by humans.

4.1 Evaluation with Simulated Sounds

Before considering real-world audio, we evaluate each model’s performance on time delay estimation task using simulated environments, following [1]. While the resulting sounds are considerably simpler than real-world recordings, they allow us to obtain metrically accurate ground-truth time delays, and to systematically vary different experimental conditions, such as the amount of background noise.

Simulation. Following previous work [1], we simulate stereo sounds using Pyroomacoustics [66]. We create three simulated environments with rooms of different sizes and microphone positions. For our sound sources, we take speech sounds from TIMIT [35] (recorded in anechoic conditions) and place them at random angles sampled uniformly from $(-90^\circ, 90^\circ)$ and distances (0.5m, 3.0m) with respect to the microphone. We add independent Gaussian noise to create conditions with different signal-noise ratio (SNR) levels, and consider a variety of reverberation times (RT_{60}). The ground-truth time delay can straightforwardly be calculated from the sound source and the microphone pose. This simulated test set, which we call **TDE-Simulation**, contains approximately 6K audio samples total (please see the supplement for more details).

Models. We evaluated our audio-based learning methods: 1) **StereoCRW**, contrastive random walks trained on stereo sounds, 2) **ZeroNCE**, slow features trained on stereo sounds (named after VINCE [37]), and 3) **MonoCLR**, and instance discrimination trained on mono sounds (named after SimCLR [12]).

We compared our methods with the widely-used **GCC-PHAT** [51], a hand-crafted audio feature. We also compared with the recent *supervised* method **Salvati et al.** [1], which trains a CNN on parameterized GCC-PHAT features to regress time delay. We trained this model on simulated stereo sounds, based on audio clips from VoxCeleb2 [15] to obtain human speech signals. To improve this baseline’s performance, we make a modification: in addition to the noise and reverberation augmentations from [1], we train with synthetic sound mixtures, in which a background sound is added to the input waveform. We regard this supervised method as an approximate upper bound for the simulation-based experiments. We provide all methods with the same duration audio as input, and evaluate different post-processing methods.

Evaluation with moderate noise. We first evaluate our models on TDE-Simulation with $SNR = 10$ and $RT_{60} = 0.5s$, a condition with moderate amounts of noise and reverberation. This simulated setup is also well-suited for analyzing

traditional techniques [51,87], which are designed to deal with unstructured, independent noise and reverberation. We evaluate three audio-based variations of our model, with and without augmentation, and on the different unlabeled training sets. For the instance discrimination model, we always include time-shifting, since the model cannot be trained without some form of augmentation. To measure prediction accuracy, we use mean absolute error (MAE) and root mean square error (RMSE) in milliseconds (ms). For all the methods, we provide 1024 (0.064s) audio samples (at 16Khz) as input and perform 128 time delay prediction votes (Sec. 3.3). For GCC-PHAT and Salvati et al., we combine the votes into a single prediction by computing the mode

or mean. For our method, we use the mode. We provide an ablation study about post-process, input duration and data distribution gap in the supplement.

As shown in Tab. 1, the StereoCRW model substantially outperforms GCC-PHAT when it is trained on a large stereo dataset, FreeMusic-Archive, obtaining performance comparable with supervised models trained on synthetic data. While ZeroNCE is trained with real stereo sound, its loss implicitly assumes the true time delay is zero, which is violated in real scenes. The cycle consistency loss in StereoCRW does not make this assumption, allowing it to learn from more complex data and learn better representations. The ZeroNCE model outperforms MonoCLR, suggesting that stereo sounds are useful training signals. Augmentations are important for all the models. We also note that the data distribution between training and test cases is quite different (*i.e.*, training with music signals and testing on the human speech), suggesting that our approach is capable of generalization.

Table 1: **Delay estimation on TDE-Simulation data.** We use SNR=10 and $RT_{60}=0.5$ s. FAIR is FAIR-Play [32], FMA is FreeMusic-Archive [20]. *Vox-Sim* is the simulator [66] with VoxCeleb2 [15] clips and *FMA-Sim* is the simulator with FreeMusic-Archive clips. Errors in ms. *Sup* refers to supervision, and *Aug* refers to augmentation.

Model	Variation	Data	Sup	Aug	MAE	RMSE
Salvati et al. [1]	Mean	Vox-Sim	✓		0.126	0.254
	Mean	Vox-Sim	✓	✓	0.169	0.294
	Mean	FMA-Sim	✓		0.135	0.256
	Mean	FMA-Sim	✓	✓	0.146	0.267
GCC-PHAT [51]	Mode	—			0.179	0.396
	Mean	—			0.160	0.318
Ours	Random	—			0.448	0.505
	MonoCLR	FAIR			0.395	0.566
	MonoCLR	FAIR	✓		0.202	0.340
	ZeroNCE	FAIR			0.241	0.362
	ZeroNCE	FAIR	✓		0.196	0.366
	StereoCRW	FAIR			0.241	0.364
	StereoCRW	FAIR	✓		0.174	0.322
	MonoCLR	FMA			0.430	0.648
	MonoCLR	FMA	✓		0.187	0.335
	ZeroNCE	FMA			0.227	0.347
	ZeroNCE	FMA	✓		0.174	0.319
	StereoCRW	FMA			0.434	0.654
	StereoCRW	FMA	✓		0.133	0.259

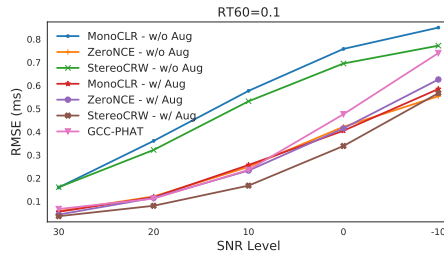


Fig. 3: Robustness to random noise.

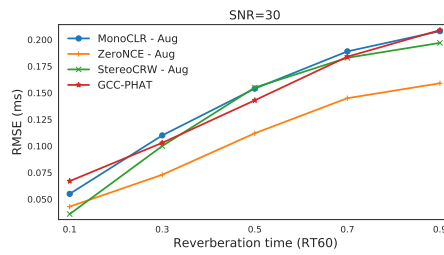


Fig. 4: Robustness to reverberation.

Table 2: **Sound mixtures on TDE-Simulation data.** We evaluate time delay estimation with sound mixtures under simulator settings $\text{SNR} = 30$ and $\text{RT}_{60} = 0.1\text{s}$. We report RMSE (ms).

Model	Variation	Aug	Dataset	Intensity level of distracting sound				
				0.1	0.3	0.5	0.7	0.9
Salvati et al. [1]	Mean		Vox-Sim	0.030	0.068	0.171	0.298	0.415
	Mean	✓	Vox-Sim	0.027	0.055	0.092	0.200	0.370
	Mean		FMA-Sim	0.047	0.091	0.196	0.326	0.423
	Mean	✓	FMA-Sim	0.051	0.073	0.102	0.217	0.370
GCC-PHAT [51]	Mean		–	0.024	0.177	0.304	0.395	0.459
Ours	MonoCLR		FMA	0.078	0.204	0.302	0.389	0.452
	MonoCLR	✓	FMA	0.055	0.103	0.185	0.306	0.427
	ZeroNCE		FMA	0.065	0.129	0.212	0.294	0.369
	ZeroNCE	✓	FMA	0.091	0.164	0.249	0.332	0.405
	StereoCRW		FMA	0.052	0.138	0.254	0.358	0.438
	StereoCRW	✓	FMA	0.041	0.079	0.144	0.273	0.417

Robustness to noise and reverberation. Following [1], we evaluate our model’s robustness to noise. We simulate sounds with the fixed reverberation time $\text{RT}_{60} = 0.1\text{s}$. We use the models trained on FreeMusic-Archive dataset and evaluate them with different SNR levels. In Fig. 3, we see that our methods trained with augmentation outperform GCC-PHAT, with a gap that widens as the amount of noise increases. This suggests that augmentation allows us to build in useful invariances that may not be captured by hand-crafted features.

We also evaluated robustness to reverberation. We fixed the SNR level to 30dB and used reverberation conditions RT_{60} in the range of $[0.1, 0.9]$ via the simulation. As shown in Fig. 4, our slow feature method outperforms the baseline under each reverberation condition while the other two approaches show similar overall performances as GCC-PHAT. The results suggest our approaches are robust to some amount of reverberation.

Robustness to mixed-in sounds. In the previous experiments, the noise in two channels is designed to be random and uncorrelated. However, in real-world audio, a major source of error comes from other sound sources (*e.g.*, background sounds). These sound sources are also present in the scene at some spatial position, thus generating a correlated error (and possibly time delay) in both channels.

We design experiments to investigate the model’s ability to ignore distracting background sounds within a mixture. We create synthetic mixtures from TDE-Simulation by mixing two sounds with different angles and distances, using $\text{SNR} = 30$ and $\text{RT}_{60} = 0.1\text{s}$. We set one sound source to be the “dominant” signal and rescale the distracting sound to be 10% – 90% loudness of the dominant source (such that the delay of the louder sound is considered to be the correct answer). We used the models from FMA and evaluate them with 0.5s audio to ensure that there is sufficient signal to identify the dominant sound.

In Tab. 2, we see that our proposed approaches significantly outperform GCC-PHAT when the distracting sounds became louder, suggesting that our model has obtained robustness to sound mixtures. For very quiet mixtures (10% the volume of the dominant source), GCC-PHAT outperforms our model, which

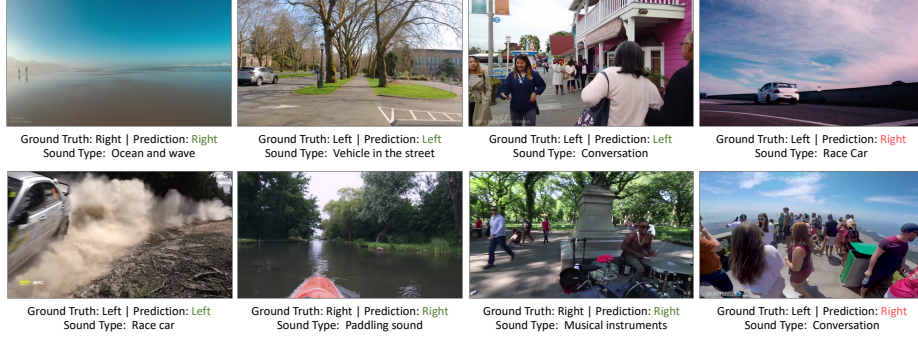


Fig. 5: **Qualitative results for in-the-wild audio.** We provide an overview of our In-the-wild Binaural dataset along with predictions from our StereoCRW model. **Green** denotes the correct predictions and **red** means the wrong predictions. We show our failure cases in the last samples on each row. For clarity, we chose examples in which the sound source is evident from the video frame. Please see supp. for more videos.

is understandable, given that this domain resembles noise-free audio, which it is well-suited to [51]. The supervised model of Salvati et al. [1] performs worse than our method when distracting sounds have a high intensity level (above 50%) unless we explicitly add the same mixture augmentation used in our methods.

4.2 Evaluation with In-the-Wild Audio Recordings

Next, we ask how well our time delay estimation methods can localize sound directions in challenging real-world scenes, using audio collected from the internet.

In-the-wild data. We collected 30 internet binaural videos, extracted 1K samples from them, and use human judgments to label sound directions. These videos contain a variety of sounds, including engine noise and human speech, which are often far from the viewer. Many also contain multiple sound sources and background noise. We provide examples of these videos in the supplement.

Since it is difficult for humans to describe sound directions in terms of time delay, we asked listeners to annotate the direction of the loudest sound. The annotator (one of the authors) listened to the audio with headphones and labeled 5 directions: *left/right*, *center left/right*, and *center*. From these, we created binary *left/right* labels, which can be objectively evaluated by thresholding the delay: we discard the *center* label and merge the remaining directional labels (resulting in 885 examples). We measure the accuracy of the thresholded time delay using these labels. We balance the dataset by swapping stereo channels, such that chance is 50%. As in the mixture experiments, we provide models with 0.5s audio to ensure that they have sufficient context. We also

Table 3: **In-the-wild evaluation.** We evaluate our models’ ability of localizing sound-ing objects on **in-the-wild** test cases.

Model	Variation	Aug.	Dataset	Acc (%) ↑
Salvati et al. [1]	Mean		Vox-Sim	87.0
	Mean	✓	Vox-Sim	87.3
	Mean		FMA-Sim	87.9
	Mean	✓	FMA-Sim	89.1
Chance	–		–	50.0
IID	–		–	75.5
GCC-PHAT [51]	Mean		–	81.3
	Random	–	–	72.5
Ours	MonoCLR		FMA	84.2
	MonoCLR	✓	FMA	87.9
	ZeroNCE		FMA	86.0
	ZeroNCE	✓	FMA	86.0
	StereoCRW		FMA	82.5
	StereoCRW	✓	FMA	88.7



Fig. 6: **Motion and time delays.** We show the time delays for both our method and for GCC-PHAT, along with the x coordinate for a tracked vehicle. We also show their correlation. We continue showing the time delay when the car moves off-screen.

compared with a method that uses interaural intensity difference (IID) cues, by comparing the root mean square (RMS) of each audio channel to determine which channel is louder to predict its left/right direction (equivalent to thresholding based on $\|\mathbf{x}\|$).

As shown in Tab. 3, our proposed approaches all substantially outperform GCC-PHAT. Interestingly, our top-performing model shows comparable results to a state-of-the-art *supervised* method, Salvati et al. We also found that augmentation generally improved results, perhaps due to the complexity of the scenes. The IID-based model performed poorly, which may be due to the fact that the sounds were often distant from the camera, and due to the presence of multiple sound sources. We show qualitative results in Fig. 5.

Correlation between visual motion and time delay. To help understand how our predicted time delays vary with motion and change over time, we correlated the visual motions in our dataset with the predicted time delays. We tracked race cars in the subset of our in-the-wild dataset that contains them using CenterTrack [89], and manually removed erroneous tracks (obtaining 49 trajectories). We applied our StereoCRW model (with 1024 samples and 128 votes). In qualitative examples (Fig. 6), we see that the car’s on-screen position is closely correlated with the time delay. Interestingly, our model continues to convey the car’s position when it moves off-screen. We computed the Spearman rank correlation coefficient [71] between the x position of each track (using the center of the bounding box) and the time delay predictions (averaged over all video clips), yielding $\rho = 0.57$ for our model and $\rho = 0.48$ for GCC-PHAT. More video results are shown in the supplement.

Application to phone recordings. We also found that our model worked successfully on sounds from ordinary video recordings from recent iPhones. For qualitative results, please see the supplement.

4.3 Visually-guided Time Delay Estimation

Instead of attending to a loud, dominant sound in a mixture (Sec. 4.1), we use vision to specify which, of several, sound sources to localize. The model solves this task by first learning to associate a voice with the given visual attribute of the speakers. Unlike audio spatialization and active speaker detection tasks, which provide temporal and spatial cues from the video, the audio-visual streams

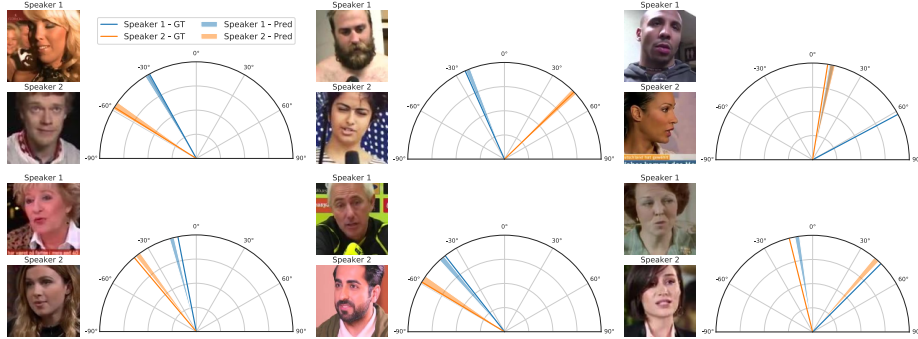


Fig. 7: **Qualitative results for visually-guided localization.** We estimate the azimuth of speakers by converting time delay predictions to angles. To convey uncertainty, we show the angular cone corresponding to ± 0.25 samples of error. We show failure cases in the last example in each row. Please see supp. for more video results.

in our case are not necessarily temporally aligned (*e.g.*, modeling the challenges of tracking a speaker when they move out of sight).

Evaluation dataset. We evaluate the audio-visual model in a simulated environment. This allows us to control the positions of the speakers, and to remove other localization cues from the images and audio. We use audio clips from VoxCeleb2 [15] with the simulation parameters from Sec. 4.1. We select 500 speakers from the database (same as training) and pair them with their corresponding face images. Following the work in source separation, we remove loudness cues by normalizing the volume of sound sources, and placing two speakers in the simulator at the same distance (but at different angles). Note that, since the position of the speakers is randomized, the visual signal does not provide localization cues (*e.g.*, via perspective). We also convert delay predictions to direction-of-arrival angles, using the known radius and microphone distance.

Comparisons. To provide points of comparison for this novel task, we compare our audio-visual approach with audio-only methods including GCC-PHAT. We also provide a (oracle) baseline which selects one of the two speakers’ ground-truth time delay at random, thus simulating a method that perfectly solves the localization task but which is unable to match faces to

Table 4: **Quantitative results of visual-guided time delay estimation on simulated data.** We evaluate our models’ ability of predicting ITD signals from mixtures with the aid of visual information.

Audio duration	0.96s		2.55s	
Model	RMSE	Err ≤ 0.1 \uparrow	RMSE	Err ≤ 0.1 \uparrow
GCC-PHAT [51]	0.503	56.6	0.504	56.9
Salvati et al. [1]	0.490	52.5	0.483	50.1
Random Oracle	0.502	56.9	0.502	56.9
Ours - Random	0.493	10.0	0.503	9.76
Ours - StereoCRW	0.493	56.8	0.488	55.7
Ours - AV	0.304	72.5	0.295	76.1
Sep [33] + GCC	0.361	77.6	0.323	82.2
Sep [33] + StereoCRW	0.309	82.8	0.281	85.5

voices. We also consider a two-stage method that first separates the speaker’s voice for each channel using VisualVoice [33], a state-of-the-art audio-visual separation model, then applies audio-based time delay estimation methods to the separated sound. To ensure a fair comparison, we retrain the static-image based separation model [33] with input audio duration of 1.27s and 2.55s.

We evaluated the methods using two metrics: RMSE, and the percentage of predictions with less than 0.1 ms (1.6 samples) of error ($\text{Err} \leq 0.1$). We feed models 1.0s or 2.55s audio, resulting in 512 votes. We show results in Tab. 4. Our audio-visual model substantially outperforms the baselines. Interestingly, it outperforms the baseline that chooses one speaker’s delay at random, an upper bound on audio-only performance. This suggests that our model successfully uses visual information. We found that the model that combines audio-visual separation with our learned audio representation performs best. Our audio-visual model (without separation) performs comparably to Sep+GCC-PHAT in regression metrics. We provide qualitative results (for the 0.96s case) in Fig. 7.

Real-world visually-guided localization.

We perform visually-guided time delay estimation on a self-recorded video (Fig. 8). Two speakers talk concurrently while moving off-screen. Our model localizes each speaker in the mixture with a cropped image of their face. We show the mean and standard deviation of delay predictions in 2.0s windows.

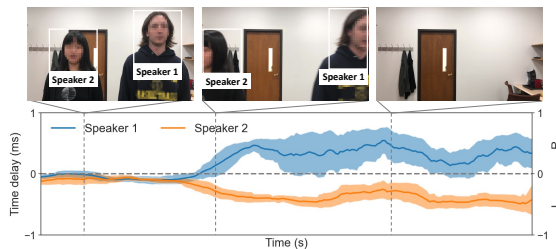


Fig. 8: **Visually-guided localization for a real-world scene.** Images are blurred to preserve privacy.

5 Discussion and Limitations

We have proposed to use *self-supervised time delay estimation* to localize sounds by learning interaural correspondence. We also introduced a novel visually-guided localization task. Our audio models obtain performance on par with supervised methods on real-world sound, while our audio-visual model successfully localizes speakers in mixtures. We see our work opening two directions: first, integrating more visual information for multisensory localization and, second, finding finer-grained delays using recent methods from optical flow [5, 48].

Limitations. Our audio-visual model associates the appearance of speakers with the sound of their voice. We tested on speakers that our model has been trained on, avoiding the need to generalize based solely on a person’s appearance [56]. However, there is still a potential for the model to exhibit bias. Our sound-based models are trained on music, which may not be representative of all downstream tasks. We released code, data, and models on our [project site](#).

Acknowledgments. We would like to thank Xixi Hu for her valuable suggestions about the augmentation idea on random walk graphs. We thank Justin Salamon for helpful discussions and Daniele Salvati for the help on the simulator setup. We thank Zhaoying Pan and Matthew Sticha for the help recording real-world examples. We also thank Daniel Geng for his comments and feedback on the paper. This work was funded in part by DARPA Semafor and Cisco Systems. The views, opinions and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

References

1. Time delay estimation for speaker localization using cnn-based parametrized gcc-phat features [3](#), [8](#), [9](#), [10](#), [11](#), [13](#)
2. Adavanne, S., Politis, A., Virtanen, T.: Direction of arrival estimation for multiple sound sources using convolutional recurrent neural network. In: 2018 26th European Signal Processing Conference (EUSIPCO). pp. 1462–1466. IEEE (2018) [3](#)
3. Afouras, T., Chung, J.S., Zisserman, A.: The conversation: Deep audio-visual speech enhancement. arXiv preprint arXiv:1804.04121 (2018) [3](#), [6](#)
4. Arandjelović, R., Zisserman, A.: Objects that sound. arXiv preprint arXiv:1712.06651 (2017) [3](#)
5. Bian, Z., Jabri, A., Efros, A.A., Owens, A.: Learning pixel trajectories with multi-scale contrastive random walks. arXiv (2022) [2](#), [4](#), [8](#), [14](#)
6. Bianco, M.J., Gerstoft, P., Traer, J., Ozanich, E., Roch, M.A., Gannot, S., Deledalle, C.A.: Machine learning in acoustics: Theory and applications. The Journal of the Acoustical Society of America **146**(5), 3590–3628 (2019) [3](#)
7. Brandstein, M.S., Silverman, H.F.: A practical methodology for speech source localization with microphone arrays. Computer Speech & Language **11**(2), 91–126 (1997) [3](#), [4](#)
8. Brungart, D.S.: Near-field auditory localization. Ph.D. thesis, Massachusetts Institute of Technology (1998) [3](#)
9. Carter, G.C., Nuttall, A.H., Cable, P.G.: The smoothed coherence transform. Proceedings of the IEEE **61**(10), 1497–1498 (1973) [3](#)
10. Carter, G.: Time delay estimation for passive sonar signal processing. IEEE Transactions on Acoustics, Speech, and Signal Processing **29**(3), 463–470 (1981) [3](#)
11. Chakrabarty, S., Habets, E.A.: Broadband doa estimation using convolutional neural networks trained with noise signals. In: 2017 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA). pp. 136–140. IEEE (2017) [3](#)
12. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. arXiv preprint arXiv:2002.05709 (2020) [2](#), [4](#), [8](#)
13. Chen, Y., Liu, B., Zhang, Z., Kim, H.S.: An end-to-end deep learning framework for multiple audio source separation and localization. International Conference on Acoustics, Speech, and Signal Processing (ICASSP) (2022) [3](#)
14. Christensen, J.H., Hornauer, S., Yu, S.: Batvision with gcc-phat features for better sound to vision predictions. arXiv preprint arXiv:2006.07995 (2020) [3](#)
15. Chung, J.S., Nagrani, A., Zisserman, A.: Voxceleb2: Deep speaker recognition. arXiv preprint arXiv:1806.05622 (2018) [7](#), [8](#), [9](#), [13](#)
16. Chung, J.S., Zisserman, A.: Out of time: automated lip sync in the wild. In: Asian conference on computer vision. pp. 251–263. Springer (2016) [5](#), [6](#)
17. Chung, Y.A., Hsu, W.N., Tang, H., Glass, J.: An unsupervised autoregressive model for speech representation learning. arXiv preprint arXiv:1904.03240 (2019) [4](#)
18. Comanducci, L., Cobos, M., Antonacci, F., Sarti, A.: Time difference of arrival estimation from frequency-sliding generalized cross-correlations using convolutional neural networks. In: ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 4945–4949. IEEE (2020) [3](#)
19. Dai, D., Vasudevan, A.B., Matas, J., Van Gool, L.: Binaural soundnet: Predicting semantics, depth and motion with binaural sounds. arXiv preprint arXiv:2109.02763 (2021) [3](#)
20. Defferrard, M., Benzi, K., Vandergheynst, P., Bresson, X.: Fma: A dataset for music analysis. arXiv preprint arXiv:1612.01840 (2016) [7](#), [9](#)

21. Diaz-Guerra, D., Miguel, A., Beltran, J.R.: Robust sound source tracking using srp-phat and 3d convolutional neural networks. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **29**, 300–311 (2020) [2](#), [3](#)
22. DiBiase, J.H.: A high-accuracy, low-latency technique for talker localization in reverberant environments using microphone arrays. Brown University (2000) [3](#)
23. Dosovitskiy, A., Fischer, P., Springenberg, J.T., Riedmiller, M., Brox, T.: Discriminative unsupervised feature learning with exemplar convolutional neural networks. *IEEE transactions on pattern analysis and machine intelligence* **38**(9), 1734–1747 (2015) [2](#)
24. Dosovitskiy, A., Springenberg, J.T., Riedmiller, M., Brox, T.: Discriminative unsupervised feature learning with convolutional neural networks. In: *Neural Information Processing Systems (NIPS)* (2014) [4](#)
25. Eloff, R., Nortje, A., van Niekerk, B., Govender, A., Nortje, L., Pretorius, A., Van Biljon, E., van der Westhuizen, E., van Staden, L., Kamper, H.: Unsupervised acoustic unit discovery for speech synthesis using discrete latent-variable neural networks. *arXiv preprint arXiv:1904.07556* (2019) [4](#)
26. Ephrat, A., Mosseri, I., Lang, O., Dekel, T., Wilson, K., Hassidim, A., Freeman, W.T., Rubinstein, M.: Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation. *arXiv preprint arXiv:1804.03619* (2018) [3](#), [6](#), [7](#)
27. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* (1981) [7](#)
28. Fisher III, J.W., Darrell, T., Freeman, W.T., Viola, P.A.: Learning joint statistical models for audio-visual fusion and segregation. In: *Neural Information Processing Systems (NIPS)* (2000) [3](#)
29. Gabbay, A., Ephrat, A., Halperin, T., Peleg, S.: Seeing through noise: Visually driven speaker separation and enhancement. In: *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. pp. 3051–3055. IEEE (2018) [3](#), [6](#)
30. Gan, C., Zhao, H., Chen, P., Cox, D., Torralba, A.: Self-supervised moving vehicle tracking with stereo sound. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 7053–7062 (2019) [3](#)
31. Gao, R., Feris, R., Grauman, K.: Learning to separate object sounds by watching unlabeled video. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 35–53 (2018) [3](#), [6](#)
32. Gao, R., Grauman, K.: 2.5 d visual sound. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 324–333 (2019) [3](#), [7](#), [9](#)
33. Gao, R., Grauman, K.: Visualvoice: Audio-visual speech separation with cross-modal consistency. In: *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 15490–15500. IEEE (2021) [13](#)
34. Garg, R., Gao, R., Grauman, K.: Geometry-aware multi-task learning for binaural audio generation from video. *arXiv preprint arXiv:2111.10882* (2021) [3](#)
35. Garofolo, J.S.: Timit acoustic phonetic continuous speech corpus. *Linguistic Data Consortium, 1993* (1993) [8](#)
36. Gong, Y., Lai, C.I.J., Chung, Y.A., Glass, J.: Ssast: Self-supervised audio spectrogram transformer. *arXiv preprint arXiv:2110.09784* (2021) [4](#)
37. Gordon, D., Ehsani, K., Fox, D., Farhadi, A.: Watching the world go by: Representation learning from unlabeled videos (2020) [4](#), [5](#), [8](#)

38. Hadji, I., Derpanis, K.G., Jepson, A.D.: Representation learning via global temporal alignment and cycle-consistency. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11068–11077 (2021) [4](#)
39. Hawley, M.L., Litovsky, R.Y., Colburn, H.S.: Speech intelligibility and localization in a multi-source environment. *The Journal of the Acoustical Society of America* **105**(6), 3436–3448 (1999) [3](#)
40. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722* (2019) [2](#), [4](#)
41. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Computer Vision and Pattern Recognition (CVPR)* (2016) [7](#)
42. Hershey, J.R., Movellan, J.R.: Audio vision: Using audio-visual synchrony to locate sounds. In: *Neural Information Processing Systems (NIPS)* (1999) [3](#)
43. Hershey, S., Chaudhuri, S., Ellis, D.P.W., Gemmeke, J.F., Jansen, A., Moore, C., Plakal, M., Platt, D., Saurous, R.A., Seybold, B., Slaney, M., Weiss, R., Wilson, K.: Cnn architectures for large-scale audio classification. In: *International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2017), <https://arxiv.org/abs/1609.09430> [7](#)
44. Houegnigan, L., Safari, P., Nadeu, C., van der Schaar, M., Solé, M., Andre, M.: Neural networks for high performance time delay estimation and acoustic source localization. In: *Proceedings of the Second International Conference on Computer Science, Information Technology and Applications*. pp. 137–146 (2017) [2](#), [3](#)
45. Hu, X., Chen, Z., Owens, A.: Mix and localize: Localizing sound sources in mixtures. *Computer Vision and Pattern Recognition (CVPR)* (2022) [4](#), [5](#)
46. Jabri, A., Owens, A., Efros, A.A.: Space-time correspondence as a contrastive random walk. *arXiv* (2020) [2](#), [4](#), [5](#)
47. Jiang, D., Li, W., Cao, M., Zou, W., Li, X.: Speech simclr: Combining contrastive and reconstruction objective for self-supervised speech representation learning. *arXiv preprint arXiv:2010.13991* (2020) [4](#)
48. Jonschkowski, R., Stone, A., Barron, J.T., Gordon, A., Konolige, K., Angelova, A.: What matters in unsupervised optical flow. *arXiv preprint arXiv:2006.04902* (2020) [14](#)
49. Kidron, E., Schechner, Y.Y., Elad, M.: Pixels that sound. In: *Computer Vision and Pattern Recognition (CVPR)* (2005) [3](#)
50. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. *International Conference on Learning Representation* (2015) [7](#)
51. Knapp, C., Carter, G.: The generalized correlation method for estimation of time delay. *IEEE transactions on acoustics, speech, and signal processing* **24**(4), 320–327 (1976) [2](#), [3](#), [4](#), [8](#), [9](#), [10](#), [11](#), [13](#)
52. Korbar, B., Tran, D., Torresani, L.: Cooperative learning of audio and video models from self-supervised synchronization. In: *Advances in Neural Information Processing Systems* (2018) [5](#), [6](#)
53. Kumpik, D.P., Campbell, C., Schnupp, J.W., King, A.J.: Re-weighting of sound localization cues by audiovisual training. *Frontiers in Neuroscience* **13**, 1164 (2019) [1](#), [3](#)
54. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *CVPR* (2015) [7](#)
55. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017) [7](#)
56. Nagrani, A., Albanie, S., Zisserman, A.: Seeing voices and hearing faces: Cross-modal biometric matching. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 8427–8436 (2018) [14](#)

57. Van den Oord, A., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv e-prints pp. arXiv-1807 (2018) [4](#)
58. Owens, A., Efros, A.A.: Audio-visual scene analysis with self-supervised multisensory features. European Conference on Computer Vision (ECCV) (2018) [3](#), [5](#), [6](#)
59. Owens, A., Wu, J., McDermott, J.H., Freeman, W.T., Torralba, A.: Learning sight from sound: Ambient sound provides supervision for visual learning. In: International Journal of Computer Vision (IJCV) (2018) [3](#)
60. Pascual, S., Ravanelli, M., Serra, J., Bonafonte, A., Bengio, Y.: Learning problem-agnostic speech representations from multiple self-supervised tasks. arXiv preprint arXiv:1904.03416 (2019) [4](#)
61. Patwari, N., Ash, J.N., Kyperountas, S., Hero, A.O., Moses, R.L., Correal, N.S.: Locating the nodes: cooperative localization in wireless sensor networks. IEEE Signal processing magazine **22**(4), 54–69 (2005) [3](#)
62. Pertilä, P., Parviainen, M.: Time difference of arrival estimation of speech signals using deep neural networks with integrated time-frequency masking. In: ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 436–440. IEEE (2019) [2](#), [3](#)
63. Rajan, R., Clement, J.P., Bhalla, U.S.: Rats smell in stereo. Science **311**(5761), 666–670 (2006) [3](#)
64. Rayleigh, L.: Xii. on our perception of sound direction. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science **13**(74), 214–232 (1907) [3](#)
65. Salvati, D., Drioli, C., Foresti, G.L.: Exploiting cnns for improving acoustic source localization in noisy and reverberant conditions. IEEE Transactions on Emerging Topics in Computational Intelligence **2**(2), 103–116 (2018) [3](#)
66. Scheibler, R., Bezzam, E., Dokmanić, I.: Pyroomacoustics: A python package for audio room simulation and array processing algorithms. In: 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 351–355. IEEE (2018) [8](#), [9](#)
67. Schmidt, R.: Multiple emitter location and signal parameter estimation. IEEE transactions on antennas and propagation **34**(3), 276–280 (1986) [2](#), [3](#)
68. Schneider, S., Baevski, A., Collobert, R., Auli, M.: wav2vec: Unsupervised pre-training for speech recognition. arXiv preprint arXiv:1904.05862 (2019) [4](#)
69. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 815–823 (2015) [7](#)
70. Senocak, A., Oh, T.H., Kim, J., Yang, M.H., Kweon, I.S.: Learning to localize sound source in visual scenes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4358–4366 (2018) [3](#)
71. Spearman, C.: The proof and measurement of association between two things. (1961) [12](#)
72. Valverde, F.R., Hurtado, J.V., Valada, A.: There is more than meets the eye: Self-supervised multi-object detection and tracking with sound by distilling multimodal knowledge. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11612–11621 (2021) [3](#)
73. Vecchiotti, P., Ma, N., Squartini, S., Brown, G.J.: End-to-end binaural sound localisation from the raw waveform. In: ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 451–455. IEEE (2019) [3](#)
74. Vondrick, C., Shrivastava, A., Fathi, A., Guadarrama, S., Murphy, K.: Tracking emerges by colorizing videos. In: ECCV (2017) [4](#)

75. Wang, D., Brown, G.J.: Computational auditory scene analysis: Principles, algorithms, and applications. Wiley-IEEE press (2006) [3](#)
76. Wang, L., Luc, P., Wu, Y., Recasens, A., Smaira, L., Brock, A., Jaegle, A., Alayrac, J.B., Dieleman, S., Carreira, J., et al.: Towards learning universal audio representations. arXiv preprint arXiv:2111.12124 (2021) [4](#)
77. Wang, L., Oord, A.v.d.: Multi-format contrastive learning of audio representations. arXiv preprint arXiv:2103.06508 (2021) [4](#)
78. Wang, X., Jabri, A., Efros, A.A.: Learning correspondence from the cycle-consistency of time. In: CVPR (2019) [2](#), [4](#)
79. Wang, Z., Zhao, H., Li, Y.L., Wang, S., Torr, P., Bertinetto, L.: Do different tracking tasks require different appearance models? NeruIPS (2021) [2](#), [4](#)
80. Wiskott, L., Sejnowski, T.J.: Slow feature analysis: Unsupervised learning of invariances. Neural computation **14**(4), 715–770 (2002) [4](#), [5](#)
81. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3733–3742 (2018) [2](#), [4](#), [6](#), [7](#)
82. Xu, X., Zhou, H., Liu, Z., Dai, B., Wang, X., Lin, D.: Visually informed binaural audio generation without binaural audios. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15485–15494 (2021) [3](#)
83. Yalta, N., Nakadai, K., Ogata, T.: Sound source localization using deep learning models. Journal of Robotics and Mechatronics **29**(1), 37–48 (2017) [3](#)
84. Yang, K., Russell, B., Salamon, J.: Telling left from right: Learning spatial correspondence of sight and sound. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9932–9941 (2020) [3](#)
85. Yang, M., Chuo, L.X., Suri, K., Liu, L., Zheng, H., Kim, H.S.: ilps: Local positioning system with simultaneous localization and wireless communication. In: IEEE INFOCOM 2019-IEEE Conference on Computer Communications. pp. 379–387. IEEE (2019) [3](#)
86. Yost, W.A., Dye, R.H., Sheft, S.: A simulated “cocktail party” with up to three sound sources. Perception & psychophysics **58**(7), 1026–1036 (1996) [3](#)
87. Zhang, C., Florêncio, D., Zhang, Z.: Why does phat work well in lownoise, reverberative environments? In: 2008 IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 2565–2568. IEEE (2008) [3](#), [4](#), [9](#)
88. Zhao, H., Gan, C., Rouditchenko, A., Vondrick, C., McDermott, J., Torralba, A.: The sound of pixels. In: Proceedings of the European conference on computer vision (ECCV). pp. 570–586 (2018) [3](#), [6](#)
89. Zhou, X., Koltun, V., Krähenbühl, P.: Tracking objects as points. In: European Conference on Computer Vision. pp. 474–490. Springer (2020) [12](#)