



# The Sound of Pixels

Hang Zhao<sup>1</sup>(✉), Chuang Gan<sup>1,2</sup>, Andrew Rouditchenko<sup>1</sup>, Carl Vondrick<sup>1,3</sup>,  
Josh McDermott<sup>1</sup>, and Antonio Torralba<sup>1</sup>

<sup>1</sup> Massachusetts Institute of Technology, Cambridge, USA

{hangzhao,roudi,jhm,torralba}@mit.edu

<sup>2</sup> MIT-IBM Watson AI Lab, Cambridge, USA

ganchuang1990@gmail.com

<sup>3</sup> Columbia University, New York City, USA

cvondrick@gmail.com

**Abstract.** We introduce PixelPlayer, a system that, by leveraging large amounts of unlabeled videos, learns to locate image regions which produce sounds and separate the input sounds into a set of components that represents the sound from each pixel. Our approach capitalizes on the natural synchronization of the visual and audio modalities to learn models that jointly parse sounds and images, without requiring additional manual supervision. Experimental results on a newly collected MUSIC dataset show that our proposed Mix-and-Separate framework outperforms several baselines on source separation. Qualitative results suggest our model learns to ground sounds in vision, enabling applications such as independently adjusting the volume of sound sources.

**Keywords:** Cross-modal learning · Sound separation and localization

## 1 Introduction

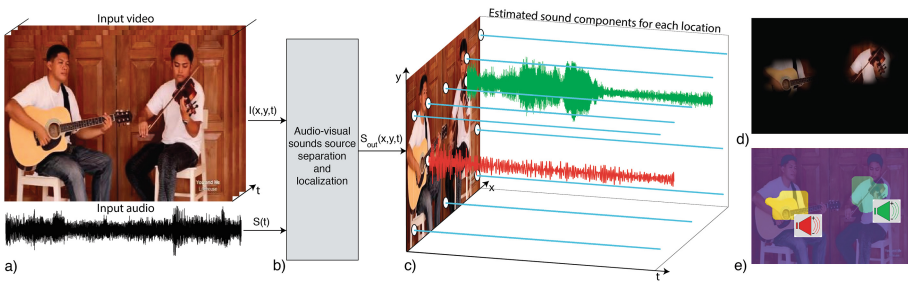
The world generates a rich source of visual and auditory signals. Our visual and auditory systems are able to recognize objects in the world, segment image regions covered by the objects, and isolate sounds produced by objects. While auditory scene analysis [5] is widely studied in the fields of environmental sound recognition [18, 26] and source separation [4, 6, 9, 41, 42, 52], the natural synchronization between vision and sound can provide a rich supervisory signal for grounding sounds in vision [17, 21, 28]. Training systems to recognize objects from vision or sound typically requires large amounts of supervision. In this paper, however, we leverage joint audio-visual learning to discover objects that produce sound in the world without manual supervision [1, 30, 36].

We show that by working with both auditory and visual information, we can learn in an unsupervised way to recognize objects from their visual appearance or the sound they make, to localize objects in images, and to separate the audio component coming from each object. We introduce a new system called PixelPlayer. Given an input video, PixelPlayer jointly separates the accompanying

audio into components and spatially localizes them in the video. PixelPlayer enables us to listen to the sound originating from each pixel in the video.

Figure 1 shows a working example of PixelPlayer (check the project website<sup>1</sup> for sample videos and interactive demos). In this example, the system has been trained with a large number of videos containing people playing instruments in different combinations, including solos and duets. No label is provided on what instruments are present in each video, where they are located, and how they sound. During test time, the input (Fig. 1a) is a video of several instruments played together containing the visual frames  $I(x, y, t)$ , and the mono audio  $S(t)$ . PixelPlayer performs audio-visual source separation and localization, splitting the input sound signal to estimate output sound components  $S_{out}(x, y, t)$ , each one corresponding to the sound coming from a spatial location  $(x, y)$  in the video frame. As an illustration, Fig. 1c shows the recovered audio signals for 11 example pixels. The flat blue lines correspond to pixels that are considered as silent by the system. The non-silent signals correspond to the sounds coming from each individual instrument. Figure 1d shows the estimated sound energy, or volume of the audio signal from each pixel. Note that the system correctly detects that the sounds are coming from the two instruments and not from the background. Figure 1e shows how pixels are clustered according to their component sound signals. The same color is assigned to pixels that generate very similar sounds.

The capability to incorporate sound into vision will have a large impact on a range of applications involving the recognition and manipulation of video. PixelPlayer's ability to separate and locate sounds sources will allow more isolated



**Fig. 1.** PixelPlayer localizes sound sources in a video and separates the audio into its components without supervision. The figure shows: (a) The input video frames  $I(x, y, t)$ , and the video mono sound signal  $S(t)$ . (b) The system estimates the output sound signals  $S_{out}(x, y, t)$  by separating the input sound. Each output component corresponds to the sound coming from a spatial location  $(x, y)$  in the video. (c) Component audio waveforms at 11 example locations; straight lines indicate silence. (d) The system's estimation of the sound energy (or volume) of each pixel. (e) Clustering of sound components in the pixel space. The same color is assigned to pixels with similar sounds. As an example application of clustering, PixelPlayer would enable the independent volume control of different sound sources in videos.

<sup>1</sup> <http://sound-of-pixels.csail.mit.edu>.

processing of the sound coming from each object and will aid auditory recognition. Our system could also facilitate sound editing in videos, enabling, for instance, volume adjustments for specific objects or removal of the audio from particular sources.

Concurrent to this work, there are papers [11, 29] at the same conference that also show the power of combining vision and audio to decompose sounds into components. [11] shows how person appearance could help solving the cocktail party problem in speech domain. [29] demonstrates an audio-visual system that separates on-screen sound *vs.* background sounds not visible in the video.

This paper is presented as follows. In Sect. 2, we first review related work in both the vision and sound communities. In Sect. 3, we present our system that leverages cross-modal context as a supervisory signal. In Sect. 4, we describe a new dataset for visual-audio grounding. In Sect. 5, we present several experiments to analyze our model. Subjective evaluations are presented in Sect. 6.

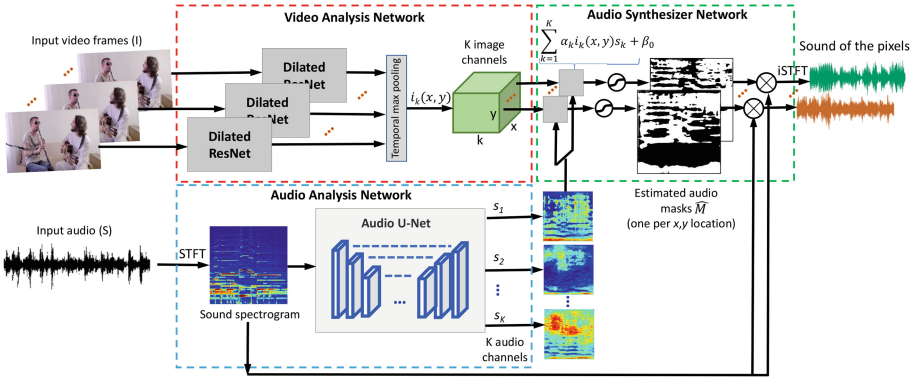
## 2 Related Work

Our work relates mainly to the fields of sound source separation, visual-audio cross-modal learning, and self-supervised learning, which will be briefly discussed in this section.

**Sound Source Separation.** Sound source separation, also known as the “cocktail party problem” [14, 25], is a classic problem in engineering and perception. Classical approaches include signal processing methods such as Non-negative Matrix Factorization (NMF) [8, 40, 42]. More recently, deep learning methods have gained popularity [7, 45]. Sound source separation methods enable applications ranging from music/vocal separation [39], to speech separation and enhancement [12, 16, 27]. Our problem differs from classic sound source separation problems because we want to separate sounds into visually and spatially grounded components.

**Learning Visual-Audio Correspondence.** Recent work in computer vision has explored the relationship between vision and sound. One line of work has developed models for generating sound from silent videos [30, 51]. The correspondence between vision and sound has also been leveraged for learning representations. For example, [31] used audio to supervise visual representations, [3, 18] used vision to supervise audio representations, and [1] used sound and vision to jointly supervise each other. In work related to our paper, people studied how to localize sounds in vision according to motion [19] or semantic cues [2, 37], however they do not separate multiple sounds from a mixed signal.

**Self-Supervised Learning.** Our work builds off efforts to learn perceptual models that are “self-supervised” by leveraging natural contextual signals in images [10, 22, 24, 33, 38], videos [13, 20, 32, 43, 44, 46], and even radio signals [48]. These approaches utilize the power of supervised learning while not requiring manual annotations, instead deriving supervisory signals from the structure in natural data. Our model is similarly self-supervised, but uses self-supervision to learn to separate and ground sound in vision.



**Fig. 2.** Procedure to generate the sound of a pixel: pixel-level visual features are extracted by temporal max-pooling over the output of a dilated ResNet applied to T frames. The input audio spectrogram is passed through a U-Net whose output is K audio channels. The sound of each pixel is computed by an audio synthesizer network. The audio synthesizer network outputs a mask to be applied to the input spectrogram that will select the spectral components associated with the pixel. Finally, inverse STFT is applied to the spectrogram computed for each pixel to produce the final sound.

### 3 Audio-Visual Source Separation and Localization

In this section, we introduce the model architectures of PixelPlayer, and the proposed Mix-and-Separate training framework that learns to separate sound according to vision.

#### 3.1 Model Architectures

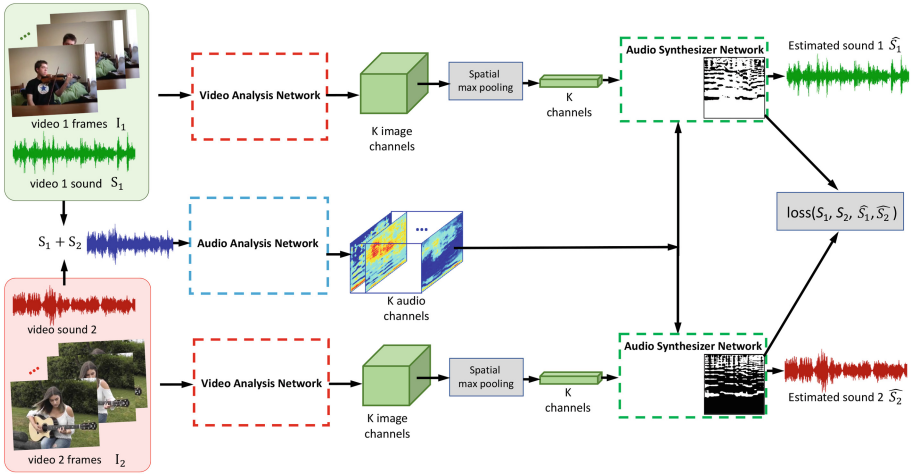
Our model is composed of a video analysis network, an audio analysis network, and an audio synthesizer network, as shown in Fig. 2.

**Video Analysis Network.** The video analysis network extracts visual features from video frames. Its choice can be an arbitrary architecture used for visual classification tasks. Here we use a dilated variation of the ResNet-18 model [15] which will be described in detail in the experiment section. For an input video of size  $T \times H \times W \times 3$ , the ResNet model extracts per-frame features with size  $T \times (H/16) \times (W/16) \times K$ . After temporal pooling and sigmoid activation, we obtain a visual feature  $i_k(x, y)$  for each pixel with size K.

**Audio Analysis Network.** The audio analysis network takes the form of a U-Net [35] architecture, which splits the input sound into K components  $s_k$ ,  $k = (1, \dots, K)$ . We empirically found that working with audio spectrograms gives better performance than using raw waveforms, so the network described in this

paper uses the Time-Frequency (T-F) representation of sound. First, a Short-Time Fourier Transform (STFT) is applied on the input mixture sound to obtain its spectrogram. Then the magnitude of spectrogram is transformed into log-frequency scale (analyzed in Sect. 5), and fed into the U-Net which yields  $K$  feature maps containing features of different components of the input sound.

**Audio Synthesizer Network.** The synthesizer network finally predicts the predicted sound by taking pixel-level visual feature  $i_k(x, y)$  and audio feature  $s_k$ . The output sound spectrogram is generated by vision-based spectrogram masking technique. Specifically, a mask  $M(x, y)$  that could separate the sound of the pixel from the input is estimated, and multiplied with the input spectrogram. Finally, to get the waveform of the prediction, we combine the predicted magnitude of spectrogram with the phase of input spectrogram, and use inverse STFT for recovery.



**Fig. 3.** Training pipeline of our proposed Mix-and-Separate framework in the case of mixing two videos ( $N = 2$ ). The dashed boxes represent the modules detailed in Fig. 2. The audio signals from the two videos are added together to generate an input mixture with known constituent source signals. The network is trained to separate the audio source signals conditioned on corresponding video frames; its output is an estimate of both sound signals. Note that we do not assume that each video contains a single source of sound. Moreover, no annotations are provided. The system thus learns to separate individual sources without traditional supervision.

### 3.2 Mix-and-Separate Framework for Self-supervised Training

The idea of the Mix-and-Separate training procedure is to artificially create a complex auditory scene and then solve the auditory scene analysis problem of

separating and grounding sounds. Leveraging the fact that audio signals are approximately additive, we mix sounds from different videos to generate a complex audio input signal. The learning objective of the model is to separate a sound source of interest conditioned on the visual input associated with it.

Concretely, to generate a complex audio input, we randomly sample  $N$  videos  $\{I_n, S_n\}$  from the training dataset, where  $n = (1, \dots, N)$ .  $I_n$  and  $S_n$  represent the visual frames and audio of the  $n$ -th video, respectively. The input sound mixture is created through linear combinations of the audio inputs as  $S_{mix} = \sum_{n=1}^N S_n$ . The model  $f$  learns to estimate the sounds in each video  $\hat{S}_n$  given the audio mixture and the visual of the corresponding video  $\hat{S}_n = f(S_{mix}, I_n)$ .

Figure 3 shows the training framework in the case of  $N = 2$ . The training phase differs from the testing phase in that (1) we sample multiple videos randomly from the training set, mix the sample audios and target to recover each of them given their corresponding visual input; (2) video-level visual features are obtained by **spatial-temporal max pooling** instead of pixel-level features. Note that although we have clear targets to learn in the training process, it is still unsupervised as we do not use the data labels and do not make assumptions about the sampled data.

The learning target in our system are the spectrogram masks, they can be binary or ratios. In the case of binary masks, the value of the ground truth mask of the  $n$ -th video is calculated by observing whether the target sound is the dominant component in the mixed sound in each T-F unit,

$$M_n(u, v) = \llbracket S_n(u, v) \geq S_m(u, v) \rrbracket, \quad \forall m = (1, \dots, N), \quad (1)$$

where  $(u, v)$  represents the coordinates in the T-F representation and  $S$  represents the spectrogram. Per-pixel sigmoid cross entropy loss is used for learning. For ratio masks, the ground truth mask of a video is calculated as the ratio of the magnitudes of the target sound and the mixed sound,

$$M_n(u, v) = \frac{S_n(u, v)}{S_{mix}(u, v)}. \quad (2)$$


In this case, per-pixel  $L1$  loss [47] is used for training. Note that the values of the ground truth mask do not necessarily stay within  $[0, 1]$  because of interference.

## 4 MUSIC Dataset


The most commonly used videos with audio-visual correspondence are musical recordings, so we introduce a musical instrument video dataset for the proposed task, called MUSIC (Multimodal Sources of Instrument Combinations) dataset.

We retrieved the MUSIC videos from YouTube by keyword query. During the search, we added keywords such as “cover” to find more videos that were not post-processed or edited.

MUSIC dataset has 714 untrimmed videos of musical solos and duets, some sample videos are shown in Fig. 4. The dataset spans 11 instrument categories:



**Fig. 4.** Example frames and associated sounds from our video dataset. The top row shows videos of solos and the bottom row shows videos of duets. The sounds are displayed in the time-frequency domain as spectrograms, with frequency on a log scale.



**Fig. 5.** Dataset Statistics: (a) Shows the distribution of video categories. There are 565 videos of solos and 149 videos of duets. (b) Shows the distribution of video durations. The average duration is about 2 min.

accordion, acoustic guitar, cello, clarinet, erhu, flute, saxophone, trumpet, tuba, violin and xylophone. Figure 5 shows the dataset statistics.

Statistics reveal that due to the natural distribution of videos, duet performances are less balanced than the solo performances. For example, there are almost no videos of tuba and violin duets, while there are many videos of guitar and violin duets.

## 5 Experiments

### 5.1 Audio Data Processing

There are several steps we take before feeding the audio data into our model. To speed up computation, we sub-sampled the audio signals to 11 kHz, such that the highest signal frequency preserved is 5.5 kHz. This preserves the most perceptually important frequencies of instruments and only slightly degrades the overall audio quality. Each audio sample is approximately 6s, randomly cropped from the untrimmed videos during training. An STFT with a window size of 1022 and a hop length of 256 is computed on the audio samples, resulting in a  $512 \times 256$  Time-Frequency (T-F) representation of the sound. We further re-sample this signal on a log-frequency scale to obtain a  $256 \times 256$  T-F representation. This step is similar to the common practice of using a Mel-Frequency scale, *e.g.* in speech



recognition [23]. The log-frequency scale has the dual advantages of (1) similarity to the frequency decomposition of the human auditory system (frequency discrimination is better in absolute terms at low frequencies) and (2) translation invariance for harmonic sounds such as musical instruments (whose fundamental frequency and higher order harmonics translate on the log-frequency scale as the pitch changes), fitting well to a ConvNet framework. The log magnitude values of T-F units are used as the input to the audio analysis network. After obtaining the output mask from our model, we use an inverse sampling step to convert our mask back to linear frequency scale with size  $512 \times 256$ , which can be applied on the input spectrogram. We finally perform an inverse STFT to obtain the recovered signal.

## 5.2 Model Configurations

In all the experiments, we use a variant of the ResNet-18 model for the video analysis network, with the following modifications made: (1) removing the last **average pooling** layer and **fc** layer; (2) removing the stride of the last residual block, and making the convolution layers in this block to have a dilation of 2; (3) adding a last  $3 \times 3$  convolution layer with  $K$  output channels. For each video sample, it takes  $T$  frames with size  $224 \times 224 \times 3$  as input, and outputs a feature of size  $K$  after **spatiotemporal max pooling**.

The audio analysis network is modified from U-Net. It has 7 convolutions (or down-convolutions) and 7 de-convolutions (or up-convolution) with skip connections in between. It takes an audio spectrogram with size  $256 \times 256 \times 1$ , and outputs  $K$  feature maps of size  $256 \times 256 \times K$ .

The audio synthesizer takes the outputs from video and audio analysis networks, fuses them with a weighted summation, and outputs a mask that will be applied on the spectrogram. The audio synthesizer is a linear layer which has very few trainable parameters ( $K$  weights +1 bias). It could be designed to have more complex computations, but we choose the simple operation in this work to show interpretable intermediate representations, which will be shown in Sect. 5.6.

Our best model takes 3 frames as visual input, and uses the number of feature channels  $K = 16$ .

## 5.3 Implementation Details

Our goal in the model training is to learn on natural videos (with both solos and duets), evaluate quantitatively on the validation set, and finally solve the source separation and localization problem on the natural videos with mixtures. Therefore, we split our MUSIC dataset into 500 videos for training, 130 videos for validation, and 84 videos for testing. Among them, 500 training videos contain both solos and duets, the validation set only contains solos, and the test set only contains duets.

During training, we randomly sample  $N = 2$  videos from our MUSIC dataset, which can be solos, duets, or silent background. Silent videos are made by pairing



silent audio waveforms randomly with images from the ADE dataset [50] which contains images of natural environments. This technique regularizes the model better in localizing objects that sound by introducing more silent videos. To recap, the input audio mixture could contain 0 to 4 instruments. We also experimented with combining more sounds, but that made the task more challenging and the model did not learn better.

In the optimization process, we use a SGD optimizer with momentum 0.9. We set the learning rate of the audio analysis network and the audio synthesizer both as 0.001, and the learning rate of the video analysis network as 0.0001 since we adopt a pre-trained CNN model on ImageNet.

#### 5.4 Sound Separation Performance

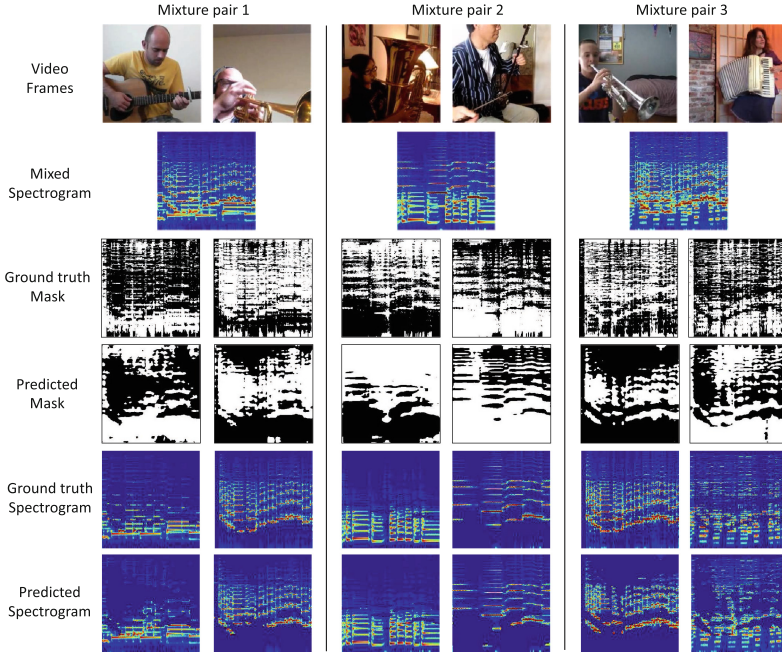
To evaluate the performance of our model, we also use the Mix-and-Separate process to make a validation set of synthetic mixture audios and the separation is evaluated.

Figure 6 shows qualitative results of our best model, which predicts binary masks that apply on the mixture spectrogram. The first row shows one frame per sampled videos that we mix together, the second row shows the spectrogram (in log frequency scale) of the audio mixture, which is the actual input to the audio analysis network. The third and fourth rows show ground truth masks and the predicted masks, which are the targets and output of our model. The fifth and sixth rows show the ground truth spectrogram and predicted spectrogram after applying masks on the input spectrogram. We could observe that even with the complex patterns in the mixed spectrogram, our model can “segment” the target instrument components out successfully.

**Table 1.** Model performances of baselines and different variations of our proposed model, evaluated in NSDR/SIR/SAR. Binary masking in log frequency scale performs best in most metrics.

|      | NMF<br>[42] | DeepConvSep<br>[7] | Spectral<br>regression | Ratio mask      |              | Binary mask     |              |
|------|-------------|--------------------|------------------------|-----------------|--------------|-----------------|--------------|
|      |             |                    |                        | Linear<br>scale | Log<br>scale | Linear<br>scale | Log<br>scale |
| NSDR | 3.14        | 6.12               | 5.12                   | 6.67            | 8.56         | 6.94            | <b>8.87</b>  |
| SIR  | 6.70        | 8.38               | 7.72                   | 12.85           | 13.75        | 12.87           | <b>15.02</b> |
| SAR  | 10.10       | 11.02              | 10.43                  | 13.87           | <b>14.19</b> | 11.12           | 12.28        |

To quantify the performance of the proposed model, we use the following metrics: the Normalized Signal-to-Distortion Ratio (NSDR), Signal-to-Interference Ratio (SIR), and Signal-to-Artifact Ratio (SAR) on the validation set of our synthetic videos. The NSDR is defined as the difference in SDR of the separated signals compared with the ground truth signals and the SDR of the mixture



**Fig. 6.** Qualitative results on vision-guided source separation on synthetic audio mixtures. This experiment is performed only for quantitative model evaluation.

signals compared with the ground truth signals. This represents the improvement of using the separated signal compared with using the mixture as each separated source. The results reported in this paper were obtained by using the open-source `mir_eval` [34] library.


Results are shown in Table 1. Among all the models, baseline approaches NMF [42] and DeepConvSep [7] use audio and ground-truth labels to do source separation. All variants of our model use the same architecture we described, and take both visual and sound input for learning. Spectral Regression refers to the model that directly regresses output spectrogram values given an input mixture spectrogram, instead of outputting spectrogram mask values. From the numbers in the table, we can conclude that (1) masking based approaches are generally better than direct regression; (2) working in the log frequency scale performs better than in the linear frequency scale; (3) binary masking based method achieves similar performance as ratio masking.

Meanwhile, we found that the NSDR/SIR/SAR metrics are not the best metrics for evaluating perceptual separation quality, so in Sect. 6 we further conduct user studies on the audio separation quality.

## 5.5 Visual Grounding of Sounds


As the title of paper indicates, we are fundamentally solving two problems: localization and separation of sounds.

**Sound Localization.** The first problem is related to the spatial grounding question, “which pixels are making sounds?” This is answered in Fig. 7: for natural videos in the dataset, we calculate the sound energy (or volume) of each pixel in the image, and plot their distributions in heatmaps. As can be seen, the model accurately localizes the sounding instruments.




**Fig. 7.** “Which pixels are making sounds?” Energy distribution of sound in pixel space. Overlaid heatmaps show the volumes from each pixel.

**Clustering of Sounds.** The second problem is related to a further question: “what sounds do these pixels make?” In order to answer this, we visualize the sound each pixel makes in images in the following way: for each pixel in a video frame, we take the feature of its sound, namely the vectorized log spectrogram magnitudes, and project them onto 3D RGB space using PCA for visualization purposes. Results are shown in Fig. 8, different instruments and the background in the same video frame have different color embeddings, indicating different sounds that they make.



**Fig. 8.** “What sounds do these pixels make?” Clustering of sound in space. Overlaid colormap shows different audio features with different colors. (Color figure online)

**Discriminative Channel Activations.** Given our model could separate sounds of different instruments, we explore its channel activations for different categories. For validation samples of each category, we find the strongest activated channel, and then sort them to generate a confusion matrix. Figure 9 shows the (a) visual and (b) audio confusion matrices from our best model. If we simply evaluate classification by assigning one category to one channel, the accuracy is 46.2% for vision and 68.9% for audio. Note that no learning is involved here, we expect much higher performance by using a linear classifier. This experiment demonstrates that the model has implicitly learned to discriminate instruments visually and auditorily.



**Fig. 9.** (a) Visual and (b) audio confusion matrices by sorting channel activations with respect to ground truth category labels.

In a similar fashion, we evaluate object localization performance of the video analysis network based on the channel activations. To generate a bounding box from the channel activation map, we follow [49] to threshold the map. We first segment the regions of which the value is above 20% of the max value of the activation map, and then take the bounding box that covers the largest connected component in the segmentation map. Localization accuracy under different intersection over union (IoU) criterion are shown in Table 2.

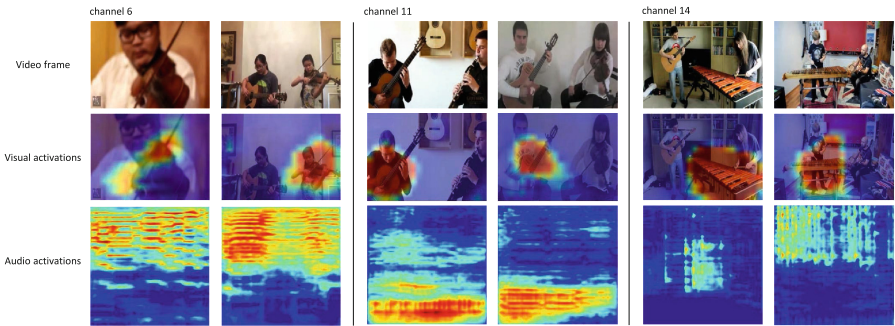
**Table 2.** Object localization performance of the learned video analysis network.

| IoU threshold | 0.3   | 0.4   | 0.5   |
|---------------|-------|-------|-------|
| Accuracy(%)   | 66.10 | 47.92 | 32.43 |

### 5.6 Visual-Audio Corresponding Activations

As our proposed model is a form of self-supervised learning and is designed such that both visual and audio networks learn to activate simultaneously on the same channel, we further explore the representations learned by the model. Specifically, we look at the K channel activations of the video analysis network

before **max pooling**, and their corresponding channel activations of the audio analysis network. The model has learned to detect important features of specific objects across the individual channels. In Fig. 10 we show the top activated videos of channel 6, 11 and 14. These channels have emerged as violin, guitar and xylophone detectors respectively, in both visual and audio domains. Channel 6 responds strongly to the visual appearance of violin and to the higher order harmonics in violin sounds. Channel 11 responds to guitars and the low frequency region in sounds. And channel 14 responds to the visual appearance of xylophone and to the brief, pulse-like patterns in the spectrogram domain. For other channels, some of them also detect specific instruments while others just detect specific features of instruments.



**Fig. 10.** Visualizations of corresponding channel activations. Channel 6 has emerged as a violin detector, responding strongly to the presence of violins in the video frames and to the high order harmonics in the spectrogram, which are colored brighter in the spectrogram of the figure. Likewise, channel 11 and 14 seems to detect the visual and auditory characteristics of guitars and xylophones.

## 6 Subjective Evaluations

The objective and quantitative evaluations in Sect. 5.4 are mainly performed on the synthetic mixture videos, the performance on the natural videos needs to be further investigated. On the other hand, the popular NSDR/SIR/SAR metrics used are not closely related to perceptual quality. Therefore we conducted crowd-sourced subjective evaluations as a complementary evaluation. Two studies are conducted on Amazon Mechanical Turk (AMT) by human raters, a sound separation quality evaluation and a visual-audio correspondence evaluation.

### 6.1 Sound Separation Quality

For the sound separation evaluation, we used a subset of the solos from the dataset as ground truth. We prepared the outputs of the baseline NMF model

and the outputs of our models, including spectral regression, ratio masking and binary masking, all in log frequency scale. For each model, we take 256 audio outputs from the same set for evaluation and each audio is evaluated by 3 independent AMT workers. Audio samples are randomly presented to the workers, and the following question is asked: “Which sound do you hear? 1. A, 2. B, 3. Both, or 4. None of them”. Here A and B are replaced by their mixture sources, *e.g.* A=clarinet, B=flute.

Subjective evaluation results are shown in Table 3. We show the percentages of workers who heard only the correct solo instrument (**Correct**), who heard only the incorrect solo instrument (**Wrong**), who heard both of the instruments (**Both**), and who heard neither of the instruments (**None**). First, we observe that although the NMF baseline did not have good NSDR numbers in the quantitative evaluation, it has competitive results in our human study. Second, among our models, the binary masking model outperforms all other models by a margin, showing its advantage in separation as a classification model. The binary masking model gives the highest correct rate, lowest error rate, and lowest confusion (percentage of **Both**), indicating that the binary model performs source separation perceptively better than the other models. It is worth noticing that even the ground truth solos do not give 100% correct rate, which represents the upper bound of performance.

**Table 3.** Subjective evaluation of sound separation performance. Binary masking-based model outperforms other models in sound separation.

| Model               | Correct(%)   | Wrong(%) | Both(%) | None(%) |
|---------------------|--------------|----------|---------|---------|
| NMF                 | 45.70        | 15.23    | 21.35   | 17.71   |
| Spectral regression | 18.23        | 15.36    | 64.45   | 1.95    |
| Ratio mask          | 39.19        | 19.53    | 27.73   | 13.54   |
| Binary mask         | <b>59.11</b> | 11.59    | 18.10   | 11.20   |
| Ground truth solo   | 70.31        | 16.02    | 7.68    | 5.99    |

## 6.2 Visual-Sound Correspondence Evaluations

The second study focuses on the evaluation of the visual-sound correspondence problem. For a pixel-sound pair, we ask the binary question: “Is the sound coming from this pixel?” For this task, we only evaluate our models for comparison as the task requires visual input, so audio-only baselines are not applicable. We select 256 pixel positions (50% on instruments and 50% on background objects) to generate corresponding sounds with different models, and get the percentage of **Yes** responses from the workers, which tells the percentage of pixels with good source separation and localization, results are shown in Table 4. This evaluation also demonstrates that the binary masking-based model gives the best performance in the vision-related source separation problem.

**Table 4.** Subjective evaluation of visual-sound correspondence. Binary masking-based model best relates vision and sound.

| Model               | Yes(%) |
|---------------------|--------|
| Spectral regression | 39.06  |
| Ratio mask          | 54.68  |
| Binary mask         | 67.58  |

## 7 Conclusions

In this paper, we introduced PixelPlayer, a system that learns from unlabeled videos to separate input sounds and also locate them in the visual input. Quantitative results, qualitative results, and subjective user studies demonstrate the effectiveness of our cross-modal learning system. We expect our work can open up new research avenues for understanding the problem of sound source separation using both visual and auditory signals.

**Acknowledgement.** This work was supported by NSF grant IIS-1524817. We thank Adria Recasens, Yu Zhang and Xue Feng for insightful discussions.

## References

1. Arandjelovic, R., Zisserman, A.: Look, listen and learn. In: 2017 IEEE International Conference on Computer Vision (ICCV), pp. 609–617. IEEE (2017)
2. Arandjelović, R., Zisserman, A.: Objects that sound (2017). arXiv preprint [arXiv:1712.06651](https://arxiv.org/abs/1712.06651)
3. Aytar, Y., Vondrick, C., Torralba, A.: Soundnet: learning sound representations from unlabeled video. In: Advances in Neural Information Processing Systems, pp. 892–900 (2016)
4. Belouchrani, A., Abed-Meraim, K., Cardoso, J.F., Moulines, E.: A blind source separation technique using second-order statistics. *IEEE Trans. Sig. Process.* **45**(2), 434–444 (1997)
5. Bregman, A.S.: *Auditory Scene Analysis: The Perceptual Organization of Sound*. MIT Press, Cambridge (1994)
6. Cardoso, J.F.: Infomax and maximum likelihood for blind source separation. *IEEE Sig. Process. Lett.* **4**(4), 112–114 (1997)
7. Chandna, P., Miron, M., Janer, J., Gómez, E.: Monoaural audio source separation using deep convolutional neural networks. In: Tichavský, P., Babaie-Zadeh, M., Michel, O.J.J., Thirion-Moreau, N. (eds.) *LVA/ICA 2017*. LNCS, vol. 10169, pp. 258–266. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-53547-0\\_25](https://doi.org/10.1007/978-3-319-53547-0_25)
8. Cichocki, A., Zdunek, R., Phan, A.H., Amari, S.I.: *Nonnegative Matrix and Tensor Factorizations: Applications to Exploratory Multi-Way Data Analysis and Blind Source Separation*. Wiley, Chichester (2009)
9. Comon, P., Jutten, C.: *Handbook of Blind Source Separation: Independent Component Analysis and Applications*. Academic Press, San Diego (2010)



10. Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1422–1430 (2015)
11. Ephrat, A., et al.: Looking to listen at the cocktail party: a speaker-independent audio-visual model for speech separation (2018). arXiv preprint [arXiv:1804.03619](https://arxiv.org/abs/1804.03619)
12. Gabbay, A., Ephrat, A., Halperin, T., Peleg, S.: Seeing through noise: speaker separation and enhancement using visually-derived speech (2017). arXiv preprint [arXiv:1708.06767](https://arxiv.org/abs/1708.06767)
13. Gan, C., Gong, B., Liu, K., Su, H., Guibas, L.J.: Geometry-guided CNN for self-supervised video representation learning (2018)
14. Haykin, S., Chen, Z.: The cocktail party problem. *Neural Comput.* **17**(9), 1875–1902 (2005)
15. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)
16. Hershey, J.R., Chen, Z., Le Roux, J., Watanabe, S.: Deep clustering: discriminative embeddings for segmentation and separation. In: 2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 31–35. IEEE (2016)
17. Hershey, J.R., Movellan, J.R.: Audio vision: using audio-visual synchrony to locate sounds. In: Solla, S.A., Leen, T.K., Müller, K. (eds.) *Advances in Neural Information Processing Systems*, vol. 12, pp. 813–819. MIT Press (2000). <http://papers.nips.cc/paper/1686-audio-vision-using-audio-visual-synchrony-to-locate-sounds.pdf>
18. Hershey, S., Chaudhuri, S., Ellis, D.P., Gemmeke, J.F., Jansen, A., Moore, R.C., Plakal, M., Platt, D., Saurous, R.A., Seybold, B., et al.: Cnn architectures for large-scale audio classification. In: 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 131–135. IEEE (2017)
19. Izadinia, H., Saleemi, I., Shah, M.: Multimodal analysis for identification and segmentation of moving-sounding objects. *IEEE Trans. Multimed.* **15**(2), 378–390 (2013)
20. Jayaraman, D., Grauman, K.: Learning image representations tied to ego-motion. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1413–1421 (2015)
21. Kidron, E., Schechner, Y.Y., Elad, M.: Pixels that sound. In: Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2005), vol. 1, pp. 88–95. IEEE Computer Society, Washington (2005). <https://doi.org/10.1109/CVPR.2005.274>
22. Larsson, G., Maire, M., Shakhnarovich, G.: Colorization as a proxy task for visual understanding. In: CVPR, vol. 2, p. 8 (2017)
23. Logan, B.: Mel frequency cepstral coefficients for music modeling. *Int. Soc. Music Inf. Retrieval* **270**, 1–11 (2000)
24. Ma, W.C., Chu, H., Zhou, B., Urtasun, R., Torrallba, A.: Single image intrinsic decomposition without a single intrinsic image. In: Ferrari, V., et al. (eds.) *ECCV 2018, Part XIV*. LNCS, vol. 11205, pp. 211–229. Springer, Cham (2018)
25. McDermott, J.H.: The cocktail party problem. *Curr. Biol.* **19**(22), R1024–R1027 (2009)
26. Mesaros, A., Heittola, T., Diment, A., Elizalde, B., Ankit Shah, E.A.: Dcase 2017 challenge setup: tasks, datasets and baseline system. In: *DCASE 2017 - Workshop on Detection and Classification of Acoustic Scenes and Events* (2017)

27. Nagrani, A., Albanie, S., Zisserman, A.: Seeing voices and hearing faces: cross-modal biometric matching (2018). arXiv preprint [arXiv:1804.00326](https://arxiv.org/abs/1804.00326)
28. Ngiam, J., Khosla, A., Kim, M., Nam, J., Lee, H., Ng, A.Y.: Multimodal deep learning. In: Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML 2011, pp. 689–696 (2011)
29. Owens, A., Efros, A.A.: Audio-visual scene analysis with self-supervised multisensory features (2018). arXiv preprint [arXiv:1804.03641](https://arxiv.org/abs/1804.03641)
30. Owens, A., Isola, P., McDermott, J., Torralba, A., Adelson, E.H., Freeman, W.T.: Visually indicated sounds. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2405–2413 (2016)
31. Owens, A., Wu, J., McDermott, J.H., Freeman, W.T., Torralba, A.: Ambient sound provides supervision for visual learning. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9905, pp. 801–816. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46448-0\\_48](https://doi.org/10.1007/978-3-319-46448-0_48)
32. Pathak, D., Girshick, R., Dollár, P., Darrell, T., Hariharan, B.: Learning features by watching objects move. In: Proceedings of CVPR, vol. 2 (2017)
33. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: feature learning by inpainting. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2536–2544 (2016)
34. Raffel, C., et al.: mir\_eval: a transparent implementation of common mir metrics. In: Proceedings of the 15th International Society for Music Information Retrieval Conference, ISMIR. Citeseer (2014)
35. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-24574-4\\_28](https://doi.org/10.1007/978-3-319-24574-4_28)
36. de Sa, V.R.: Learning classification with unlabeled data. In: Advances in Neural Information Processing Systems, pp. 112–119 (1993)
37. Senocak, A., Oh, T.H., Kim, J., Yang, M.H., Kweon, I.S.: Learning to localize sound source in visual scenes (2018). arXiv preprint [arXiv:1803.03849](https://arxiv.org/abs/1803.03849)
38. Shu, Z., Yumer, E., Hadap, S., Sunkavalli, K., Shechtman, E., Samaras, D.: Neural face editing with intrinsic image disentangling (2017). arXiv preprint [arXiv:1704.04131](https://arxiv.org/abs/1704.04131)
39. Simpson, A.J.R., Roma, G., Plumbley, M.D.: Deep karaoke: extracting vocals from musical mixtures using a convolutional deep neural network. In: Vincent, E., Yeredor, A., Koldovský, Z., Tichavský, P. (eds.) LVA/ICA 2015. LNCS, vol. 9237, pp. 429–436. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-22482-4\\_50](https://doi.org/10.1007/978-3-319-22482-4_50)
40. Smaragdis, P., Brown, J.C.: Non-negative matrix factorization for polyphonic music transcription. In: 2003 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics, pp. 177–180. IEEE (2003)
41. Vincent, E., Gribonval, R., Févotte, C.: Performance measurement in blind audio source separation. *IEEE Trans. Audio Speech Lang. Process.* **14**(4), 1462–1469 (2006)
42. Virtanen, T.: Monaural sound source separation by nonnegative matrix factorization with temporal continuity and sparseness criteria. *IEEE Trans. Audio Speech Lang. Process.* **15**(3), 1066–1074 (2007)
43. Vondrick, C., Pirsiavash, H., Torralba, A.: Generating videos with scene dynamics. In: Advances in Neural Information Processing Systems, pp. 613–621 (2016)
44. Vondrick, C., Shrivastava, A., Fathi, A., Guadarrama, S., Murphy, K.: Tracking emerges by colorizing videos (2018). arXiv preprint [arXiv:1806.09594](https://arxiv.org/abs/1806.09594)

45. Wang, D., Chen, J.: Supervised speech separation based on deep learning: an overview (2017). arXiv preprint [arXiv:1708.07524](https://arxiv.org/abs/1708.07524)
46. Wang, X., Gupta, A.: Unsupervised learning of visual representations using videos. In: ICCV, pp. 2794–2802 (2015)
47. Zhao, H., Gallo, O., Frosio, I., Kautz, J.: Loss functions for image restoration with neural networks. *IEEE Trans. Comput. Imaging* **3**(1), 47–57 (2017)
48. Zhao, M., et al.: Through-wall human pose estimation using radio signals. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 7356–7365 (2018)
49. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2921–2929 (2016)
50. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ADE20K dataset. In: Proceedings of CVPR (2017)
51. Zhou, Y., Wang, Z., Fang, C., Bui, T., Berg, T.L.: Visual to sound: Generating natural sound for videos in the wild (2017). arXiv preprint [arXiv:1712.01393](https://arxiv.org/abs/1712.01393)
52. Zibulevsky, M., Pearlmutter, B.A.: Blind source separation by sparse decomposition in a signal dictionary. *Neural Comput.* **13**(4), 863–882 (2001)