

# Recursive Visual Sound Separation Using Minus-Plus Net

Xudong Xu

Bo Dai

Dahua Lin

CUHK-SenseTime Joint Lab, The Chinese University of Hong Kong

xx018@ie.cuhk.edu.hk bdai@ie.cuhk.edu.hk dhl@ie.cuhk.edu.hk

## Abstract

*Sounds provide rich semantics, complementary to visual data, for many tasks. However, in practice, sounds from multiple sources are often mixed together. In this paper we propose a novel framework, referred to as MinusPlus Network (MP-Net), for the task of visual sound separation. MP-Net separates sounds recursively in the order of average energy<sup>1</sup>, removing the separated sound from the mixture at the end of each prediction, until the mixture becomes empty or contains only noise. In this way, MP-Net could be applied to sound mixtures with arbitrary numbers and types of sounds. Moreover, while MP-Net keeps removing sounds with large energy from the mixture, sounds with small energy could emerge and become clearer, so that the separation is more accurate. Compared to previous methods, MP-Net obtains state-of-the-art results on two large scale datasets, across mixtures with different types and numbers of sounds.*

## 1. Introduction

Besides visual cues, the sound that comes along with what we see often provides complementary information, which could be used for object detection [12, 13, 17, 18] to clarify ambiguous visual cues, and description generation [6, 7, 26, 5] to enrich semantics. On the other hand, as what we hear in most cases is the mixture of different sounds, coming from different sources, it is necessary to separate sounds and associate them to sources in the visual scene, before utilizing sound data.

The difficulties of visual sound separation lie in several aspects. 1) First, possible sound sources in the corresponding videos may not make any sound, which causes ambiguity. 2) Second, the mixture usually contains a large variance in terms of numbers and types. 3) More importantly, sounds in the mixture often affect each other in multiple ways. For example, sounds with large energy often dominate the mixture, making other sounds less distinguishable or even sound like noise in some cases.

<sup>1</sup>In this paper, average energy of sound stands for the average energy of its spectrogram.

Existing works [10, 28] on visual sound separation mainly separate each sound independently. They assume either fixed types or fixed numbers of sounds, separating sounds independently. Since strong assumptions in [10, 28] have limited their applicability in generalized scenarios, separating sounds independently could lead to inconsistency between the actual mixture and the mixture of separated sounds, *e.g.* some data in the actual mixture does not appear in any sounds. Moreover, the separation of sounds with small energy may be affected by sounds with large energy in such independent processes.

Facing these challenges, we propose a novel solution, referred to as MinusPlus Network (MP-Net), which identifies each sound in the mixture recursively, in descending order of average energy. It can be divided into two stages, namely a *minus* stage and a *plus* stage. At each step of the minus stage, MP-Net identifies the most salient sound from the current mixture, then removes the sound therefrom. This process repeats until the current mixture becomes empty or contains only noise. Due to the removal of preceding separations, only one sound could obtain the component that is shared by multiple sounds. Consequently, to compensate such cases, MP-Net refines each sound in the plus stage, which computes a residual based on the sound itself and the mixture of preceding separated sounds. The final sound is obtained by mixing the outputs of both stages.

MP-Net efficiently overcomes the challenges of visual sound separation. By recursively separating sounds, it adaptively decides the number of sounds in the mixture, without knowing a priori the number and the types of sounds. Moreover, in MP-Net, sounds with large energy will be removed from the mixture after they are separated. In this way, sounds with relatively smaller energy naturally emerge and become clearer, diminishing the effect of imbalanced sound energy.

Overall, our contributions can be briefly summarized as follows: (1) We propose a novel framework, referred to as MinusPlus Network (MP-Net), to separate independent sounds from the recorded mixture based on a corresponding video. Unlike previous works which assume a fixed number of sounds in the mixture, the proposed framework could dy-

namically determine the number of sounds, leading to better generalization ability. (2) MP-Net utilizes a novel way to alleviate the issue of imbalanced energy of sounds in the mixture, by subtracting salient sounds from the mixture after they are separated, so that sounds with less energy could emerge. (3) On two large scale datasets, MP-Net obtains more accurate results, and generalizes better compared to the state-of-the-art method.

## 2. Related Work

Works connecting visual and audio data can be roughly divided into several categories.

The first category is jointly embedding audio-visual data. Aytaç *et al.* [4] transfer discriminative knowledge in visual content to audio data by minimizing the KL divergence of their representations. Arandjelovic *et al.* [2] associate representations of visual and audio data by learning their correspondence (*i.e.* whether they belong to the same video), and authors in [21, 20, 16] further extend such correspondence to temporal alignment, resulting in better representations. Different from these works, visual sound separation requires to separate each independent sound from the mixture, relying on the corresponding video.

The task of sound localization also requires jointly processing visual and audio data, which identifies the region that generates the sound. To solve this task, Hershey *et al.* [15] locate sound sources in video frames by measuring audio-visual synchrony. Both Tian *et al.* [24] and Parascandolo *et al.* [3] apply sound event detection to find sound sources. Finally, Senocak *et al.* [23] and Arandjelovic *et al.* [3] find sound sources by analyzing the activation of feature maps. Although visual sound separation could also locate separated sounds in the corresponding video, it requires separating the sounds at first, making it more challenging.

Visual sound separation belongs to the third category, a special type of which is visual speech separation, where sounds in the mixture are all human speeches. For example, Afouras *et al.* [1] and Ephrat *et al.* [8] obtain a speaker-independent model by leveraging a large amount of news and TV videos, and Xu *et al.* [27] propose an auditory selection framework which uses attention and memory to capture speech characteristics. Unlike these works, we target the general task of separating sounds with different types, which have more diverse sound characteristics.

The most related works are [28] and [10]. In [10], a convolutional network is used to predict the type of objects appeared in the video, and Non-negative Matrix Factorization [9] is used to extract a set of basic components. The association between each object and each basic component will be estimated via a Multi-Instance Multi-Label objective. Consequently, sounds will be separated using the associations between basic components and each predicted object. [28] follows a similar framework, replacing Non-negative Ma-

trix Factorization with a U-Net [22]. In addition, instead of predicting object-base associations, it directly predicts weights conditioned on visual semantics. While the former predicts the existence of different objects in the video, assuming fixed types of sounds, the latter assumes a fixed number of sounds. Such strong assumptions have limited their generalization ability, as the mixture of sounds often has large variance across sound types and numbers. More importantly, each prediction in [28] and [10] is conducted independently. As a result, 1) there may be an inconsistency between the mixture of all predicted sounds and the actual mixture. *e.g.* some data appeared in the actual mixture may not appear in any predicted sounds, or some data has appeared too many times in predicted sounds, exceeding its frequency in the actual mixture. 2) As sounds in the mixture have different average energy, sounds with large energy may affect the prediction accuracy of sounds with less energy. Different from them, our proposed method recursively predicts each sound in the mixture, following the order of average energy. The predicted sound with a large energy will be removed from the mixture after its prediction. In this way, our proposed method requires no assumptions on the type and number of sounds and ensures consistent predictions with the input mixture. Moreover, when sounds with large energy are removed from the mixture continually, sounds with less energy could emerge and become clearer, resulting in more accurate predictions.

## 3. Visual Sound Separation

In the task of visual sound separation, we are given a context video  $V$  and a recorded mixture of sounds  $\mathbf{S}^{\text{mix}}$ , which is the mixture of a set of independent sounds  $\{\mathbf{S}_1^{\text{solo}}, \mathbf{S}_2^{\text{solo}}, \dots, \mathbf{S}_n^{\text{solo}}\}$ . The objective is to separate each sound from the mixture based on the visual context in  $V$ .

We propose a new framework for visual sound separation, referred to as MinusPlus Network (MP-Net), which learns to separate each independent sound from the recorded mixture, without knowing a priori the number of sounds in the mixture (*i.e.*  $n$ ). In addition, MP-Net could also associate each independent sound with a plausible source in the corresponding visual content, providing a way to link data in two different modalities.

### 3.1. Overview

In MP-Net, sound data is represented as spectrograms, and the overall structure of MP-Net has been demonstrated in Figure 1. It has two stages, namely the minus stage and the plus stage.

**Minus Stage.** In the minus stage, MP-Net recursively separates each independent sound from the mixture  $\mathbf{S}^{\text{mix}}$ , where at every recursive step it will focus on the sound that is the most salient one in the remaining sounds. The process

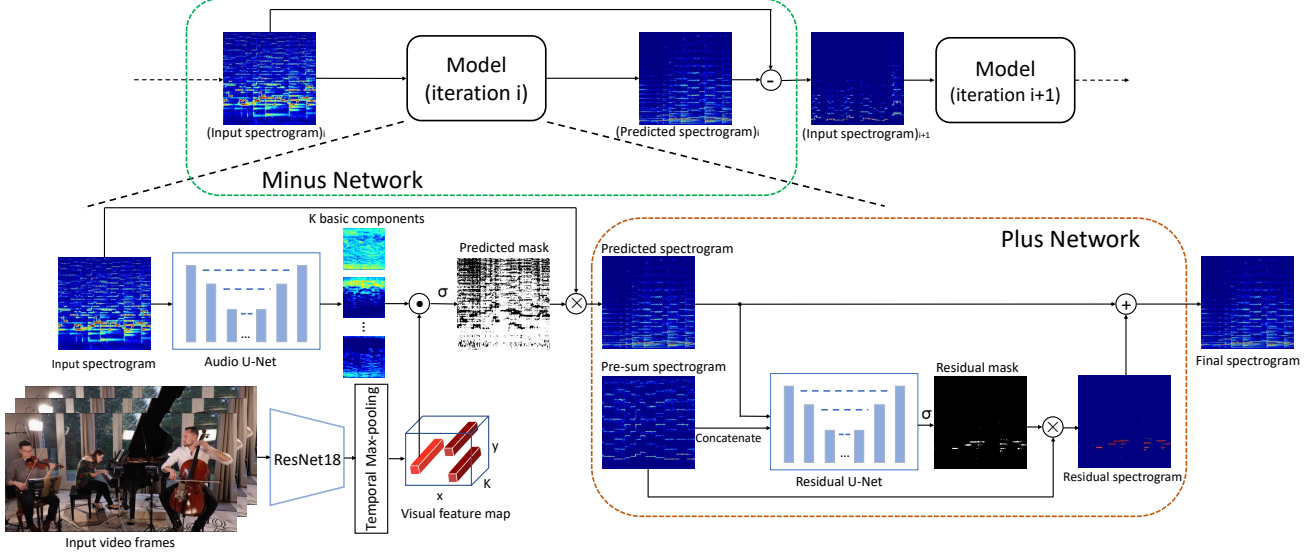


Figure 1: The proposed MinusPlus Network (MP-Net) for visual sound separation. It consists of two sub-networks, namely the Minus Network (M-Net) and the Plus Network (P-Net). In M-Net, sounds are recursively separated based on the input video. At  $i$ -th recursive step, a U-Net [22] is used to predict  $k$  basic components for current mixture, which are then used to estimate a mask  $\mathbf{M}$ , as well as determine a source in the video for the sound to be separated. Based on the mask and visual cues at the sound source, a sound is separated, which will be removed from the mixture. M-Net repeats these operations until the mixture contains only noise. All separated sounds will be refined by P-Net, which computes a residual from the mixture of preceding separated sounds. The final output of MP-Net for each sound is obtained by mixing the outputs of M-Net and P-Net.

could be described as:

$$\mathbf{S}_0^{\text{mix}} = \mathbf{S}^{\text{mix}}, \quad (1)$$

$$\mathbf{S}_i^{\text{solo}} = \text{M-Net}(\mathbf{V}, \mathbf{S}_{i-1}^{\text{mix}}), \quad (2)$$

$$\mathbf{S}_i^{\text{mix}} = \mathbf{S}_{i-1}^{\text{mix}} \ominus \mathbf{S}_i^{\text{solo}}, \quad (3)$$

where  $\mathbf{S}_i^{\text{solo}}$  is  $i$ -th predicted sound, M-Net stands for the sub-net used in the minus stage, and  $\ominus$  is the element-wise subtraction on spectrograms. As shown in Eq.(3), MP-Net keeps removing  $\mathbf{S}_i^{\text{solo}}$  from previous mixture  $\mathbf{S}_{i-1}^{\text{mix}}$ , until current mixture  $\mathbf{S}_i^{\text{mix}}$  is empty or contains only noise with significantly low energy.

**Plus Stage.** While in the minus stage we remove preceding predictions from the mixture by subtraction, a prediction  $\mathbf{S}_i^{\text{solo}}$  may miss some content that is shared by it and preceding predictions  $\{\mathbf{S}_1^{\text{solo}}, \dots, \mathbf{S}_{i-1}^{\text{solo}}\}$ . Inspired by this, MP-Net contains a plus stage, which further refines each separated sound following:

$$\mathbf{S}_i^{\text{remix}} = \mathbf{S}_1^{\text{solo}} \oplus \dots \oplus \mathbf{S}_{i-1}^{\text{solo}}, \quad (4)$$

$$\mathbf{S}_i^{\text{residual}} = \text{P-Net}(\mathbf{S}_i^{\text{remix}}, \mathbf{S}_i^{\text{solo}}), \quad (5)$$

$$\mathbf{S}_i^{\text{solo, final}} = \mathbf{S}_i^{\text{solo}} \oplus \mathbf{S}_i^{\text{residual}}, \quad (6)$$

where P-Net stands for the sub-network used in the plus stage, and  $\oplus$  is the element-wise addition on spectrograms.

As shown in Eq.(6), MP-Net computes a residual  $\mathbf{S}_i^{\text{residual}}$  for  $i$ -th prediction, based on  $\mathbf{S}_i^{\text{solo}}$  and the mixture of all preceding predictions, and finally refines  $i$ -th prediction by mixing  $\mathbf{S}_i^{\text{solo}}$  and  $\mathbf{S}_i^{\text{residual}}$ . Subsequently, in practice, we use  $\mathbf{S}_i^{\text{solo, final}}$  instead of  $\mathbf{S}_i^{\text{solo}}$  in Eq.(3) and Eq.(4).

The benefits of using two stages lie in several aspects.

1) The minus stage could effectively determine the number of independent sounds in the mixture, without knowing it a priori. 2) Removing preceding predictions from the mixture could diminish their disruption on the remaining sounds, so that remaining sounds continue emerging as the recursion goes. 3) Removing preceding predictions from the mixture potentially helps M-Net focus on the distinct characteristics of remaining sounds, enhancing the accuracy of its predictions. 3) The plus stage could compensate for the loss of shared information between each prediction and all its preceding predictions, potentially smoothing the final prediction of each sound. Subsequently, we will introduce the two subnets, namely M-Net and P-Net, respectively.

### 3.2. M-Net

M-Net is the sub-net responsible for separating each independent sound from the mixture, following a recursive procedure. Specifically, To separate the most salient sound  $\mathbf{S}_i^{\text{solo}}$  at  $i$ -th recursive step, M-Net will predict  $k$  sub-spectrograms  $\{\mathbf{S}_1^{\text{sub}}, \dots, \mathbf{S}_k^{\text{sub}}\}$  using a U-Net [22], which

NSDR: 31.03 SIR: 49.20 SAR: 31.10

NSDR: 20.89 SIR: 31.14 SAR: 21.32




Figure 2: As shown in this figure, when mismatch appears at different locations of spectrograms, scores in terms of SDR, SIR and SAR can vary significantly.

capture different patterns in  $\mathbf{S}^{\text{mix}}$ . At the same time, we will obtain a feature map  $\mathbf{V}$  of size  $H/16 \times W/16 \times k$  from the input video  $V$ , which estimates the association score between each sub-spectrogram and the visual content at different spatial location. With  $\mathbf{V}$  and  $\{\mathbf{S}_1^{\text{sub}}, \mathbf{S}_2^{\text{sub}}, \dots, \mathbf{S}_k^{\text{sub}}\}$ , we could then identify the associated visual content for  $\mathbf{S}_i^{\text{solo}}$ :

$$(x^*, y^*) = \underset{(x, y)}{\operatorname{argmax}} E \left[ \sigma \left( \sum_{j=1}^k \mathbf{V}(x, y, j) * \mathbf{S}_j^{\text{sub}} \right) * \mathbf{S}^{\text{mix}} \right], \quad (7)$$

where  $\sigma \left( \sum_{j=1}^k \mathbf{V}(x, y, j) * \mathbf{S}_j^{\text{sub}} \right)$  computes a location-specific mask.  $E[\cdot]$  computes the average energy of a spectrogram, and  $\sigma$  stands for the sigmoid function. We regard  $(x^*, y^*)$  as the source location of  $\mathbf{S}_i^{\text{solo}}$ , and the feature vector  $\mathbf{v}$  in  $\mathbf{V}$  at that location as the visual feature of  $\mathbf{S}_i^{\text{solo}}$ . To separate  $\mathbf{S}_i^{\text{solo}}$ , we reuse the vector  $\mathbf{v}$  as the attention weights on sub-spectrograms and get the actual mask  $\mathbf{M}$  by

$$\mathbf{M} = \sigma \left( \sum_{j=1}^k \mathbf{v}_j \mathbf{S}_j^{\text{sub}} \right), \quad (8)$$

where  $\sigma$  stands for the sigmoid function. Following [28], we refer to  $\mathbf{M}$  as a ratio mask, and an alternative choice is to further binarize  $\mathbf{M}$  to get a binary mask. Finally,  $\mathbf{S}_i^{\text{solo}}$  is obtained by  $\mathbf{S}_i^{\text{solo}} = \mathbf{M} \otimes \mathbf{S}^{\text{mix}}$ . It is worth noting that we could also directly predict  $\mathbf{S}_i^{\text{solo}}$  following  $\mathbf{S}_i^{\text{solo}} = \sum_{j=1}^k \mathbf{v}_j \mathbf{S}_j^{\text{sub}}$ . However, it is reported that an intermediate mask leads to better results [28]. At the end of  $i$ -th recursive step, MP-Net will remove the predicted  $\mathbf{S}_i^{\text{solo}}$  from previous mixture  $\mathbf{S}_{i-1}^{\text{mix}}$  via  $\mathbf{S}_i^{\text{mix}} = \mathbf{S}_{i-1}^{\text{mix}} \ominus \mathbf{S}_i^{\text{solo}}$ , so that less salient sounds could emerge in later recursive steps. When the average energy in  $\mathbf{S}_i^{\text{mix}}$  less than a threshold  $\epsilon$ , M-Net stops the recursive process, assuming all sounds have been separated.

### 3.3. P-Net

While M-Net makes succeeding predictions more accurate by removing preceding predictions from the mixture, succeeding predictions may miss some content shared by preceding predictions, leading to incomplete spectrograms. To overcome this issue, MP-Net further applies a P-Net to refine sounds separated by the M-Net. Specifically, for  $\mathbf{S}_i^{\text{solo}}$ , P-Net applies a U-Net [22] to get a residual mask  $\mathbf{M}_r$ , based on two inputs, namely  $\mathbf{S}_i^{\text{solo}}$  and  $\mathbf{S}_i^{\text{remix}} = \mathbf{S}_1^{\text{solo}} \oplus \dots \oplus \mathbf{S}_{i-1}^{\text{solo}}$ , which is the re-mixture of preceding predictions, as the missing content of  $\mathbf{S}_i^{\text{solo}}$  could only appear in them. The final sound spectrogram for  $i$ -th sound is obtained by:

$$\mathbf{S}_i^{\text{residual}} = \mathbf{S}_i^{\text{remix}} \otimes \mathbf{M}_r, \quad (9)$$

$$\mathbf{S}_i^{\text{solo, final}} = \mathbf{S}_i^{\text{solo}} \oplus \mathbf{S}_i^{\text{residual}}. \quad (10)$$

### 3.4. Mutual Distortion Measurement

To evaluate models for visual sound separation, previous approaches [10, 28] utilize Normalized Signal-to-Distortion Ratio (NSDR), Signal-to-Interference Ratio (SIR), and Signal-to-Artifact Ratio (SAR). While these traditional metrics could reflect the separation performance to some extent, they are sensitive to frequencies, so that scores of different separated sounds are not only affected by their similarities with ground-truths, but also the locations of mismatches. Consequently, as shown in Figure 2, scores in terms of SDR, SIR and SAR vary significantly when the mismatch appears at different locations. To compensate such cases, we propose to measure the quality of visual sound separation under the criterion that two pairs of spectrograms need to obtain approximately the same score if they have the same level of similarities. This metric, referred to as Average Mutual Information Distortion (AMID), computes the average similarity between a separated sound and a ground-truth of another sound, where the similarity is estimated via the Structural Similarity (SSIM) [25] over spectrograms. Specifically, for a set of separated sounds  $\{\mathbf{S}_1^{\text{solo}}, \dots, \mathbf{S}_m^{\text{solo}}\}$  and its corresponding annotations  $\{\mathbf{S}_1^{\text{gt}}, \dots, \mathbf{S}_m^{\text{gt}}\}$ , AMID is computed as:

$$\text{AMID}(\{\mathbf{S}_i^{\text{solo}}\}, \{\mathbf{S}_j^{\text{gt}}\}) = \frac{1}{m(m-1)} \sum_{i \neq j} \text{SSIM}(\mathbf{S}_i^{\text{solo}}, \mathbf{S}_j^{\text{gt}}). \quad (11)$$

As AMID relies on SSIM over spectrograms, it is insensitive to frequencies. Moreover, a low AMID score indicates the model can distinctly separate sounds in a mixture, which meets the evaluation requirements of visual sound separation.



	Mask	mix-2				mix-3			
		NSDR↑	SIR↑	SAR↑	AMID↓	NSDR↑	SIR↑	SAR↑	AMID↓
MIML [10]	-	1.12	3.05	5.26	8.74	0.34	1.32	2.39	8.91
PixelPlayer[28]	Binary	2.20	<b>7.98</b>	9.00	6.99	1.00	4.82	3.82	6.49
	Ratio	<b>2.96</b>	5.91	13.77	10.35	2.99	2.59	10.69	10.55
M-Net	Binary	2.02	7.48	9.22	5.96	1.23	4.76	4.69	5.96
	Ratio	2.66	5.17	14.19	6.80	3.54	2.31	15.92	11.54
MP-Net (M-Net + P-Net)	Binary	2.14	7.66	9.47	<b>5.78</b>	1.48	<b>4.99</b>	4.80	<b>5.76</b>
	Ratio	2.81	5.45	<b>14.49</b>	6.53	<b>3.75</b>	2.52	<b>16.77</b>	10.59

Table 1: This table lists the results of visual sound separation on VEGAS [29], where MP-Net obtains best performance under various metrics and settings.

	Mask	mix-2				mix-3			
		NSDR↑	SIR↑	SAR↑	AMID↓	NSDR↑	SIR↑	SAR↑	AMID↓
MIML [10]	-	2.82	4.94	9.21	16.37	1.76	3.32	4.54	25.32
PixelPlayer[28]	Binary	5.16	10.96	10.60	15.81	3.01	6.38	6.27	24.01
	Ratio	6.09	8.07	14.93	18.81	4.83	4.87	11.19	29.84
M-Net	Binary	5.47	12.63	10.21	11.83	4.01	7.89	6.76	23.76
	Ratio	6.82	10.12	14.98	13.90	5.61	5.03	13.42	24.05
MP-Net (M-Net + P-Net)	Binary	5.73	<b>12.75</b>	10.50	<b>11.22</b>	4.23	<b>8.18</b>	6.95	<b>23.10</b>
	Ratio	<b>7.00</b>	10.39	<b>15.31</b>	13.36	<b>5.75</b>	5.37	<b>13.68</b>	23.51

Table 2: This table lists the results of visual sound separation on MUSIC [28], where MP-Net obtains best performance under various metrics and settings.

## 4. Experiments

### 4.1. Datasets

We test MP-Net on two datasets, namely Multimodal Sources of Instrument Combinations (MUSIC) [28] and (Visually Engaged and Grounded AudioSet (VEGAS) [29].

MUSIC mainly contains untrimmed videos of people playing instruments belonging to 11 categories, namely *accordion*, *acoustic guitar*, *cello*, *clarinet*, *erhu*, *flute*, *saxophone*, *trumpet*, *tuba*, *violin* and *xylophone*. There are respectively 500, 130 and 40 samples in the train, validation and test set of MUSIC. While the test set of MUSIC contains only duets without ground-truths of sounds in mixtures, we use its validation set as test set, and train set for training and validation. While MUSIC focuses on instrumental sounds, VEGAS, another dataset with a larger scale, covers 10 types of natural sounds, including *baby crying*, *chainsaw*, *dog*, *drum*, *fireworks*, *helicopter*, *printer*, *rail transport*, *snoring* and *water flowing*, trimmed from AudioSet [11]. 2,000 samples in VEGAS are used as the test, with remaining samples being used for training and validation.

### 4.2. Training and Testing Details

Due to the lack of ground-truth of real mixed data, *i.e.* those videos that contain multiple sounds. We construct such data from *solo* video clips instead. Each clip contains at most one sound. We denote the collection of solo video clips by  $\{\mathbf{S}_j^{\text{solo}}, V_j\}_{j=1}^N$ , where  $\mathbf{S}_j^{\text{solo}}$  and  $V_j$  respectively represent the sound and visual content. Note that a video clip can be silent, for such case,  $\mathbf{S}_j^{\text{solo}}$  is an empty spectrogram. For each  $V_j$ , we sample  $T = 6$  frames at even intervals, and extract visual features for each frame using ResNet-18 [14]. This would result in a feature tensor of size  $T \times (H/16) \times (W/16) \times k$ . In both training and testing, this feature tensor will be reduced into a vector to represent the visual content by performing max pooling along the first three dimensions. On top of this solo video collection, we then follow the Mix-and-Separate strategy as in [28, 10] to construct the mixed video/sound data, where each sample mixes  $n$  videos, called a mix- $n$  sample.

Audios are preprocessed before training and testing. Specifically, we sample audios at 16kHz, and use the open-sourced package *librosa* [19] to transform the sound clips of around 6 seconds into STFT spectrograms of size  $750 \times 256$ , where the window size and the hop length are respectively

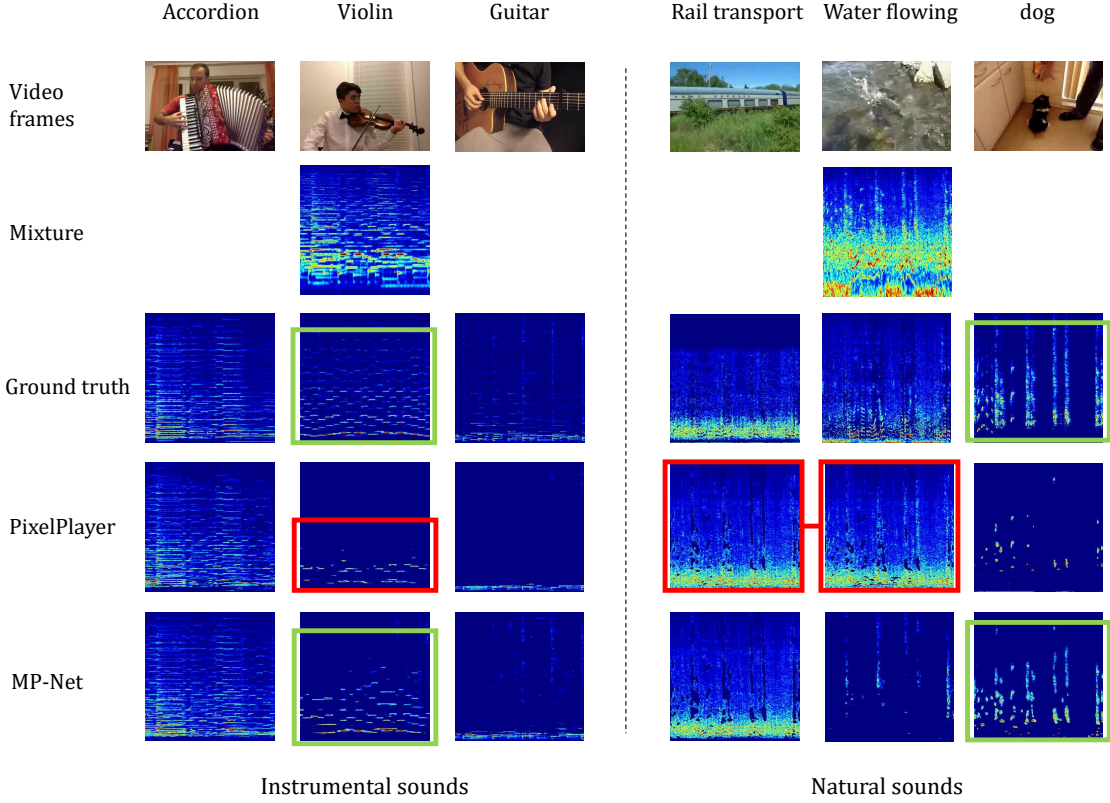


Figure 3: Qualitative results for visual sound separation, from MP-Net and PixelPlayer [28]. On the left, the mixture of instrumental sounds is demonstrated, where MP-Net successfully separates violin’s sound, unlike its baseline. And on the right, natural sounds are separated from the mixture. As the sound of *rail transport* and *water flowing* share a high similarity, MP-Net separates dog sounds but predicts silence for *water flowing*, while its baseline reuses the sound of *rail transport*.




Figure 4: Curves of varying the number of sounds in the testing mixture on MUSIC, obtained by models respectively trained with mixtures of 2 sounds (first row), and 3 sounds (second row). green, red, and blue lines respectively stand for MP-Net, PixelPlayer [28] and MIML [10].

set as 1, 500 and 375. We down-sample it on mel scale and obtain the spectrogram with size  $256 \times 256$ .

We use  $k = 16$  for the M-Net. We adopt a three-round training strategy, where in the first round, we train the M-

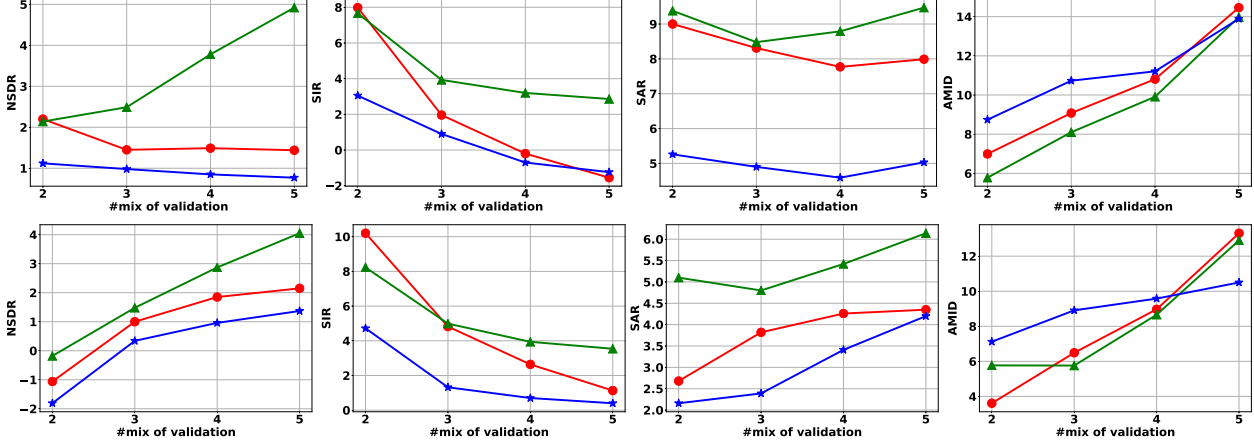


Figure 5: Curves of varying the number of sounds in the testing mixture on VEGAS, obtained by models respectively trained with mixtures of 2 sounds (first row), and 3 sounds (second row). green, red, and blue lines respectively stand for MP-Net, PixelPlayer [28] and MIML [10].




Figure 6: Qualitative samples on found associations between visual sound sources (light regions) and different types of sounds, using MP-Net and PixelPlayer [28].

Net in isolation. And in the second round, we train the P-Net while fixing parameters of the M-Net. Finally, in the third round, the M-Net and the P-Net are jointly finetuned.

During training, for each mix- $n$  sample, we first perform data augmentation, randomly scaling the energy of the spectrograms. Then, MP-Net makes  $n$  predictions in the descending order of the average energy of the ground-truth

sounds. Particularly, for the  $t$ -th prediction, MP-Net predicts  $\mathbf{M}$  and  $\mathbf{M}_r$  for the sound with  $t$ -th largest average energy and computes the BCE loss between  $\mathbf{M} + \mathbf{M}_r$  and the ground-truth mask if binary masks are used, or  $\mathcal{L}_1$  loss if ratio masks are used. After all  $n$  predictions are done, we add an extra loss between the remaining mixture and an empty spectrogram – ideally if all  $n$  predictions are precise, there

should be no sound left.

During evaluation, we determine the prediction order by Eq.(7). Since all baselines need to know the number of sounds in the mixture, to compare fairly, we also provide the number of sounds to MP-Net. It is, however, noteworthy that MP-Net could work without this information, relying only on the termination criterion to determine the number. On *MUSIC*, MP-Net predicts the correct number of sounds with over 90% of accuracy.

### 4.3. Experimental Results

**Results on Effectiveness** To study the effectiveness of our model, we compared our model to state-of-the-art methods, namely PixelPlayer [28] and MIML [10], across datasets and settings, providing a comprehensive comparison. Specifically, on both *MUSIC* and *VEGAS*, we train and evaluate all methods twice, respectively using mix-2 and mix-3 samples, which contain 2 and 3 sounds in the mixture. For PixelPlayer and MP-Net, we further alter the form of masks to switch between ratio masks and binary masks. The results in terms of NSDR, SIR, SAR and AMID are listed in Table 1 for *VEGAS* and Table 2 for *MUSIC*. We observe that 1) our proposed MP-Net obtains best results in most settings, outperforming PixelPlayer and MIML by large margins, which indicate the effectiveness of separating sounds in the order of average energy. 2) Using ratio masks is better in terms of NSDR and SAR, while using binary masks is better in terms of SIR and AMID. 3) Our proposed metric AMID correlates well with other metrics, which intuitively verifies its effectiveness. 4) Scores of all methods on mix-2 samples are much higher than scores on mix-3 samples, which add only one more sound in the mixture. Such differences in scores have shown the challenges of visual sound separation. 5) In general, methods obtain higher scores on *MUSIC*, meaning natural sounds are more complicated than instrumental sounds, as instrumental sounds often contain regular patterns.

**Results on Ablation Study** While the proposed MP-Net contains two sub-nets, we have compared MP-Net with and without P-Net. As shown in Table 1 and Table 2, on all metrics, MP-Net with P-Net outperforms MP-Net without P-Net by large margins, indicating 1) different sounds have shared patterns, a good model needs to take this into consideration, so that the mixture of separated sounds is consistent with the actual mixture. 2) P-Net could effectively compensate the loss of shared patterns caused by sound removals, filling blanks in the spectrograms.

**Results on Robustness** A benefit of recursively separating sounds from the mixture is that MP-Net is robust when the number of sounds in the mixture varies, although trained with a fixed number of sounds. To verify the generalization ability of MP-Net, we have tested all methods that trained

with mix-2 or mix-3 samples, on samples with an increasing number of sounds in the mixtures. The resulting curves on *MUSIC* are shown in Figure 4, and Figure 5 includes the curves on *VEGAS*. In Figure 4 and Figure 5, up to mixtures consisting of 5 sounds, MP-Net trained with a fixed number of sounds in the mixture outperforms baselines steadily as the number of sounds in the mixture increases.

**Qualitative Results** In Figure 3, we show qualitative samples with sounds separated by respectively MP-Net and PixelPlayer, in the form of spectrograms. In the sample with a mixture of instrumental sounds, PixelPlayer fails to separate sounds belonging to violin and guitar, as their sounds are overwhelmed by the sound of accordion. On the contrary, MP-Net successfully separates sounds of violin and guitar, alleviating the effect of the accordion’s sound. Unlike PixelPlayer that separates sounds independently, MP-Net recursively separates the dominant sound in current mixture, and removes it from the mixture, leading to accurate separation results. A similar phenomenon can also be observed in the sample with a mixture of natural sounds, PixelPlayer predicts the same sound for *rail\_transport* and *water\_flow*ing, and fails to separate the sound of a dog.

**Localization Results** MP-Net could also be used to associate sound sources in the video with separated sounds, using Eq.(7). We show some samples in Figure 6, where compared to PixelPlayer, MP-Net produces more precise associations between separated sounds and their possible sound sources.

## 5. Conclusion

We propose MinusPlus Network (MP-Net), a novel framework for visual sound separation. Unlike previous methods that separate each sound independently, MP-Net jointly considers all sounds, where sounds with larger energy are separated firstly, followed by them being removed from the mixture, so that sounds with smaller energy keep emerging. In this way, once trained, MP-Net could deal with mixtures made of an arbitrary number of sounds. On two datasets, MP-Net is shown to consistently outperform state-of-the-arts, and maintains steady performance as the number of sounds in mixtures increases. Besides, MP-Net could also associate separated sounds to possible sound sources in the corresponding video, potentially linking data from two modalities.

**Acknowledgement** This work is partially supported by the Collaborative Research Grant from SenseTime (CUHK Agreement No.TS1610626 & No.TS1712093), and the General Research Fund (GRF) of Hong Kong (No.14236516 & No.14203518).



## References

- [1] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. The conversation: Deep audio-visual speech enhancement. *Proc. Interspeech 2018*, pages 3244–3248, 2018. 2
- [2] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 609–617, 2017. 2
- [3] Relja Arandjelovic and Andrew Zisserman. Objects that sound. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 435–451, 2018. 2
- [4] Yusuf Aytar, Carl Vondrick, and Antonio Torralba. Soundnet: Learning sound representations from unlabeled video. In *Advances in Neural Information Processing Systems*, 2016. 2
- [5] Bo Dai, Sanja Fidler, and Dahua Lin. A neural compositional paradigm for image captioning. In *Advances in Neural Information Processing Systems*, pages 658–668, 2018. 1
- [6] Bo Dai, Sanja Fidler, Raquel Urtasun, and Dahua Lin. Towards diverse and natural image descriptions via a conditional gan. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2970–2979, 2017. 1
- [7] Bo Dai, Deming Ye, and Dahua Lin. Rethinking the form of latent states in image captioning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 282–298, 2018. 1
- [8] Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinatan Hassidim, William T Freeman, and Michael Rubinstein. Looking to listen at the cocktail party: a speaker-independent audio-visual model for speech separation. *ACM Transactions on Graphics (TOG)*, 37(4):112, 2018. 2
- [9] Cédric Févotte, Nancy Bertin, and Jean-Louis Durrieu. Non-negative matrix factorization with the itakura-saito divergence: With application to music analysis. *Neural computation*, 21(3):793–830, 2009. 2
- [10] Ruohan Gao, Rogerio Feris, and Kristen Grauman. Learning to separate object sounds by watching unlabeled video. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 35–53, 2018. 1, 2, 4, 5, 6, 7, 8
- [11] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 776–780. IEEE, 2017. 5
- [12] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 1
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 1
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [15] John R Hershey and Javier R Movellan. Audio vision: Using audio-visual synchrony to locate sounds. In *Advances in neural information processing systems*, pages 813–819, 2000. 2
- [16] Bruno Korbar, Du Tran, and Lorenzo Torresani. Cooperative learning of audio and video models from self-supervised synchronization. In *Advances in Neural Information Processing Systems*, pages 7774–7785, 2018. 2
- [17] Hongyang Li, Bo Dai, Shaoshuai Shi, Wanli Ouyang, and Xiaogang Wang. Feature intertwiner for object detection. *arXiv preprint arXiv:1903.11851*, 2019. 1
- [18] Hongyang Li, Xiaoyang Guo, Bo Dai, Wanli Ouyang, and Xiaogang Wang. Neural network encapsulation. In *European Conference on Computer Vision*, pages 266–282. Springer, 2018. 1
- [19] Brian McFee, Colin Raffel, Dawen Liang, Daniel PW Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto. librosa: Audio and music signal analysis in python. In *Proceedings of the 14th python in science conference*, pages 18–25, 2015. 5
- [20] Andrew Owens and Alexei A Efros. Audio-visual scene analysis with self-supervised multisensory features. *European Conference on Computer Vision (ECCV)*, 2018. 2
- [21] Andrew Owens, Jiajun Wu, Josh H McDermott, William T Freeman, and Antonio Torralba. Ambient sound provides supervision for visual learning. In *European Conference on Computer Vision (ECCV)*, 2016. 2
- [22] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 2, 3, 4
- [23] Arda Senocak, Tae-Hyun Oh, Junsik Kim, Ming-Hsuan Yang, and In So Kweon. Learning to localize sound source in visual scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4358–4366, 2018. 2
- [24] Yapeng Tian, Jing Shi, Bochen Li, Zhiyao Duan, and Chenliang Xu. Audio-visual event localization in unconstrained videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 247–263, 2018. 2
- [25] Zhou Wang, Alan C Bovik, Hamid R Sheikh, Eero P Simoncelli, et al. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 4
- [26] Yilei Xiong, Bo Dai, and Dahua Lin. Move forward and tell: A progressive generator of video descriptions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 468–483, 2018. 1
- [27] Jiaming Xu, Jing Shi, Guangcan Liu, Xiuyi Chen, and Bo Xu. Modeling attention and memory for auditory selection in a cocktail party environment. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018. 2
- [28] Hang Zhao, Chuang Gan, Andrew Rouditchenko, Carl Vondrick, Josh McDermott, and Antonio Torralba. The sound of pixels. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 570–586, 2018. 1, 2, 4, 5, 6, 7, 8

- [29] Yipin Zhou, Zhaowen Wang, Chen Fang, Trung Bui, and Tamara L Berg. Visual to sound: Generating natural sound for videos in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3550–3558, 2018. 5