

Do We Need Sound for Sound Source Localization?

Takashi Oya*, Shohei Iwase*, Ryota Natsume, Takahiro Itazuri, Shugo Yamaguchi, and Shigeo Morishima

Waseda Research Institute for Science and Engineering
`oya.takashi@ruri.waseda.jp`, `sh.iwase@fuji.waseda.jp`,
`nano.poteto@toki.waseda.jp`, `s132800732@fuji.waseda.jp`,
`wasedayshugo@suou.waseda.jp`, `shigeo@waseda.jp`

Abstract. During the performance of sound source localization which uses both visual and aural information, it presently remains unclear how much either image or sound modalities contribute to the result, i.e. do we need both image and sound for sound source localization? To address this question, we develop an unsupervised learning system that solves sound source localization by decomposing this task into two steps: (i) “potential sound source localization”, a step that localizes possible sound sources using only visual information (ii) “object selection”, a step that identifies which objects are actually sounding using aural information. Our overall system achieves state-of-the-art performance in sound source localization, and more importantly, we find that despite the constraint on available information, the results of (i) achieve similar performance. From this observation and further experiments, we show that visual information is dominant in “sound” source localization when evaluated with the currently adopted benchmark dataset. Moreover, we show that the majority of sound-producing objects within the samples in this dataset can be inherently identified using only visual information, and thus that the dataset is inadequate to evaluate a system’s capability to leverage aural information. As an alternative, we present an evaluation protocol that enforces both visual and aural information to be leveraged, and verify this property through several experiments.

Keywords: cross-modal learning, sound source localization, unsupervised learning

1 Introduction

In many scientific areas, it is common to make a prediction of an objective variable using explanatory variables. Nevertheless, the purpose is not always to generate an accurate prediction, but to investigate which variable is important for a prediction. For example, when using learning methods such as Random Forests [1], or their variants like Gradient Boosted Trees [2–4], quantities such

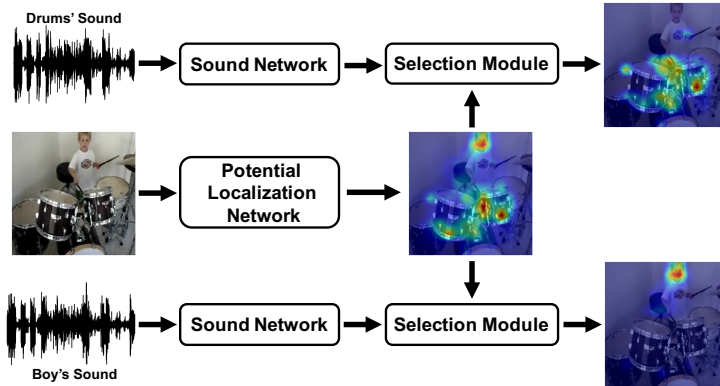


Fig. 1. Overview of our system. The system decomposes sound source localization into two steps, and each step is carried out by a specific module: “potential localization network” and “selection module”. More specifically, the potential localization network obtains a heat map of all objects that are capable of producing sound, and given sound features from a sound network, the selection module identifies objects that are actually producing sound. In the example above, as the boy and drums are capable of producing sound, the potential localization network obtains a heat map that responds to their corresponding regions. Then, when human voice is given as input audio, the pixels that correspond to humans are selected by the selection module. On the other hand, when the sound of drums is given, the pixels that correspond to drums are selected.

as “variable importance” are widely used to determine the predictive strength of specific variables. In a similar manner, linear regression coefficients are known to be very useful to evaluate the contribution of specific variables in a linear regression analysis. However, there are only a few studies that investigate the contributions of each modality in cross-modal perception tasks that leverage both visual and aural information. Our study focuses on the investigation of this issue in sound source localization, which is a cross-modal perception task aiming to determine image pixels that are associated with a given sound source, e.g. selecting the region of trumpet in an image when given sound of trumpet.

In this research, we develop an unsupervised architecture which solves sound source localization by decomposing the problem into two tasks. First, given only an image, our system initially suggests a potential localization map, which is a heat map of all objects that are capable of producing sound. We call this step “potential sound source localization”. Then, given a sound, our system identifies which objects within the potential localization map are actually producing sound, and returns a localization map. By comparing the performance of the potential localization map, which is derived only from visual information, and the localization map, which leverages both visual and aural information, we are able to evaluate the contributions of each modality in the benchmark dataset.

As can be seen from Fig. 1, the architecture has 3 components: the potential localization network, the sound network, and the selection module. The potential localization network is an image network that takes an image as an input

and returns the potential localization map. Notably, the potential localization network operates independently from the sound network, which returns latent features derived from the sound input. The selection module updates the potential localization map to the localization map using the latent features, by selecting objects from the potential localization map. During the training phase, our model can be trained in an unsupervised manner as we utilize audio-visual correspondence [5], i.e. whether image and sound correspond to each other or not.

As a result, when using both visual and aural information, our system achieves state of the art performance. More significantly, although utilizing only visual information, the potential localization map performs similarly to the localization map, indicating that in the current benchmark dataset, high performance can be achieved without the usage of sound as input. Furthermore, from additional quantitative and qualitative analysis, we have found that for the majority of the samples within the current dataset, the sound-producing object can be inherently identified using only image, as all objects that are capable of producing sound are actually producing sound. This is problematic as this is not the case in the real world, and proper evaluation on a system’s capability of leveraging aural information can not be conducted. To overcome this issue, we design an evaluation protocol called *concat-and-localize*, where evaluation is conducted on artificially created samples that enforce the usage of aural information. We verify this property through evaluation of our system using this evaluation protocol.

2 Related Work

Below we provide a list of cross-modal perception tasks that are closely related to the problem of sound source localization. Each task is accompanied with a description of how they are related.

Audio-Visual Correspondence (AVC) Audio and visual events tend to occur concurrently within our daily lives. Additionally, humans learn to localize a sound source without any supervision by experiencing these natural co-occurrences a sufficient number of times since they were born. Like humans, in order for a machine to learn from these natural co-occurrences, Arandjelovic & Zisserman [5] proposed an AVC task that predicts whether a pair of sound clip and an image correspond with one another. The concept of AVC facilitates solutions for other cross-modal perception tasks in an unsupervised fashion, such as audio-visual representation learning [5–7], audio-visual source separation [8–10], and sound source localization [6, 9, 11–13]. Although our work does not directly map to the AVC task, it indirectly leverages this idea in an unsupervised setting.

Audio-Visual representation learning This task aims to generate representations of image/sound; these representations are useful for image/sound classification, zero-shot detection, cross-modal retrieval and action recognition. Audio-visual representation learning models can be grouped within 3 categories: the use

of image features to train a sound network in a supervised manner [14–17], the use of sound features to train an image network in a supervised manner [5, 14, 18], and via unsupervised methods [6, 7, 10, 19, 20] including networks to classify motion information. Though the aim of representation learning is different from that of sound source localization, some existing works have reported that the internal representation i.e., image features and sound features of the models developed for sound source localization is useful for image/sound classification and action recognition tasks [7, 10, 20]. Our efforts are different from this task, as we do not focus on internal representations.

Sound generation from video. The purpose of this task is sound synthesis based on visual information. Some works in this domain have generated monaural audio from a video [21, 22], while other works have converted monaural audio into either binaural audio [23] or spatial audio [24]. Gao & Grauman [23] showed that it is possible to localize sound sources by visualizing regions important in conversion and Morgado et al. [24] showed that it is possible to localize a sound source present in a 360 degree video. Our work differs from these works as we neither focus on sound synthesis nor do we use binaural audio or spatial audio.

Audio-visual source separation. Audio-visual source separation is a task that utilizes visual information to guide sound source separation, in contrast to blind sound source separation that does not use visual information [25–28]. There are various methods employed for audio-visual source separation: NMF [29–31], subspace methods [32, 33], the mix-and-separate method [8–10, 12, 13, 34], and the use of facial information for speech separation [35–38]. Many works have attempted simultaneous sound source separation and sound source localization [8–10, 12, 13], as it is important to identify which objects are producing sound to perform sound source separation. For example, Zhao et al. [9] proposed a method that assigns sounds to each image pixel and can localize a sound source by visualizing the volume of each image pixel. However, these previous works applied their methods to datasets containing a limited number of categories, such as a music instrument or speech. In contrast, our work does not focus on sound source separation and can be applied to a dataset containing an unconstrained number of categories.

Sound source localization Sound source localization is a classic problem in both science and engineering. In robotics and signal processing research, sound source localization mainly indicates a task that identifies the spatial locations of the sound sources using only aural information from several microphones [39–42]. On the contrary, in computer vision research, sound source localization indicates a task seeking to determine image pixels that are associated with sound sources, usually employing both visual and aural information from a single microphone, i.e., monaural hearing. There exist various studies concerning sound source localization in the computer vision domain, including works that are not based on neural network techniques: mutual information and CCA [43, 44], subspace methods [45], and keypoints [46]. After the great success of neural networks in

the year 2012 [47], a growing number of works have been conducted that leverage neural network, including the following methods: those that leverage motion information [12, 13, 19], CAM-based methods [10, 20], methods based on attention mechanism [6, 11], and methods that can perform concurrent sound source localization and sound source separation [8–10, 12, 13]. A work that is most similar to the one described in this study is that of Senocak et al. [11], which can be trained in an unsupervised manner and can be applied to a noisy, unlabelled dataset. However, the focus of their work is different from ours, because our study’s aim is to investigate the contributions of each modality in sound source localization. Furthermore, our network differs from theirs in that we introduce potential sound source localization network and object selection module.

3 Proposed Framework

3.1 Dataset and Data Preprocessing

As a benchmark, we employ a dataset created by Senocak et al. [11]; this dataset contains 5k random subset selections of Flickr-SoundNet [15], a large unlabelled dataset containing more than 200M matched pairs of images and sounds. In addition, this dataset contains sound source annotations using bounding-boxes, which were assigned by 3 annotators. The image sizes within the dataset are fixed at 256×256 pixels, the sampling rate of the sound is fixed at 22.05 kHz, and the length of the sound clips is not fixed.

For image data, we performed a simple preprocessing and augmentation scheme: mean subtraction and random cropping (224×224). For audio data, we performed STFT with a window-length of 1022 and a hop-length of 511.

3.2 Potential localization network

The overall network architecture is given in Fig. 2. The potential localization network takes image vectors $I_{i,j} \in R^{H \times W}$ with height H and width W as inputs. Then, a VGG-11 network [48] pretrained on ImageNet [49] and two 1×1 convolution layers followed by ReLU [50] were used to extract a positive vector $v_{k,i,j} \in R^{K \times \frac{H}{16} \times \frac{W}{16}}$ that represents image features with K channels. Finally, a potential localization map $P_{i,j} \in R^{\frac{H}{16} \times \frac{W}{16}}$ is computed by applying a 1×1 convolution layer followed by a softmax function, as shown below.

$$P_{i,j} = \frac{\exp[\sum_k (W_k v_{k,i,j}) + b]}{\sum_{i,j} \exp[\sum_k (W_k v_{k,i,j}) + b]}. \quad (1)$$

Here W_k, b are weight and bias term of the 1×1 convolution layer, respectively.

3.3 Sound network

The sound network takes the amplitudes and phases of a spectrogram as an input. This information is then fed into a VGG-like Network. Here, the sound

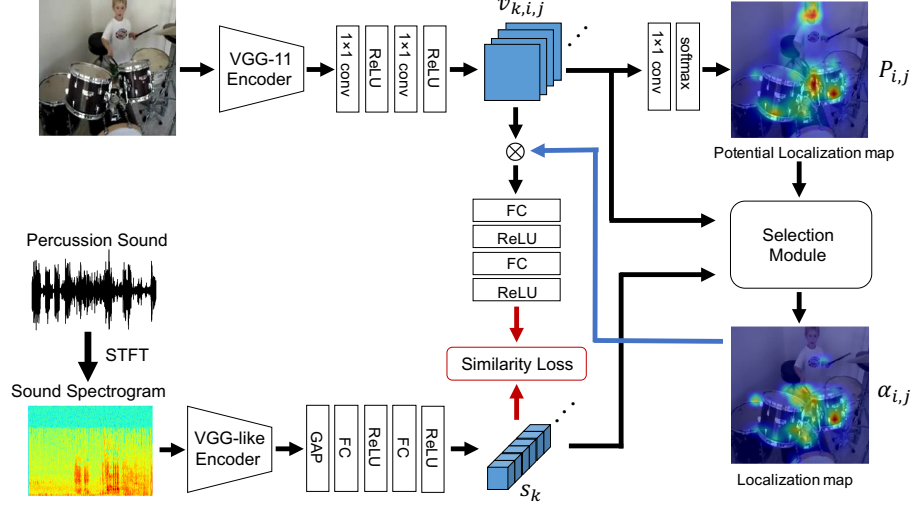


Fig. 2. The detailed network architecture of our unsupervised model. The potential localization map is obtained using image features extracted from the input image. Then from the potential localization map, the image features, and the sound features extracted from the input sound, the selection module produces a localization map. The similarity loss is calculated using sound features and image features of the sound-producing object obtained from the pixel-wise product between the localization map and the image features.

length is arbitrary, caused by the application of a global average pooling (GAP) layer [51] along the time axis after the VGG-like network has acted on the data. Then, the positive vector $s_k \in R^K$ is obtained which represents the sound features after having been fed through two fully-connected (FC) layers followed by a ReLU.

3.4 Selection module

The selection module takes 3 inputs: image features $v_{k,i,j}$, sound features s_k , and the potential localization map $P_{i,j}$. An attention map $A_{i,j}$ is obtained by calculating the cosine similarity between image features $v_{k,i,j}$, and sound features s_k , as shown below:

$$A_{i,j} = \frac{\sum_k v_{k,i,j} s_k}{\sqrt{\sum_k s_k^2} \sqrt{\sum_k v_{k,i,j}^2}}. \quad (2)$$

It should be noted that $A_{i,j}$ ranges between 0 and 1, as both $v_{k,i,j}$ and s_k are positive. Finally, the localization map $\alpha_{i,j}$ is calculated by L1 normalization along pixels after multiplying $A_{i,j}$ and $P_{i,j}$ together as shown by the following expression:

$$\alpha_{i,j} = \frac{A_{i,j} P_{i,j}}{\sum_{i,j} A_{i,j} P_{i,j}}. \quad (3)$$

3.5 Training phase

For training, we leverage concepts from an AVC task [5] in which networks are trained to determine whether data pairs (image and sound) correspond to each other. Our loss function which we call “similarity loss” is calculated as follows. We first calculated $\hat{s}_k \in R^K$ by

$$\hat{s}_k = \text{ReLU}(\text{FC}(\text{ReLU}(\text{FC}(\sum_{i,j} v_{k,i,j} \alpha_{i,j}))))). \quad (4)$$

Intuitively, \hat{s}_k indicates image features $v_{k,i,j}$ that were filtered by the localization map $\alpha_{i,j}$.

Then as in a Siamese network [52], the loss function is defined differently for the positive pairs (pairs corresponding to each other) and the negative pairs (pairs not corresponding to each other).

For both positive and negative pairs, the cosine similarity between \hat{s}_k and s_k is used. For the positive pairs, the loss function is given as follows:

$$L = 1 - \frac{\sum_k s_k \hat{s}_k}{\sqrt{\sum_k s_k^2} \sqrt{\sum_k \hat{s}_k^2}}. \quad (5)$$

For the negative pairs, the loss function is given by the following expression:

$$L = \frac{\sum_k s_k \hat{s}_k}{\sqrt{\sum_k s_k^2} \sqrt{\sum_k \hat{s}_k^2}}. \quad (6)$$

The loss function forces \hat{s}_k to be similar to sound features s_k for positive pairs and vice versa for negative pairs. Additional details such as hyperparameters are provided in the supplementary material.

3.6 Inference phase

As we have shown, the potential localization network does not depend on aural information and works independently from other components. Therefore, a potential localization map $P_{i,j}$ can be obtained using only visual information. On the other hand, aural information is needed to obtain a localization map $\alpha_{i,j}$, since the attention map $A_{i,j}$ necessary to derive $\alpha_{i,j}$ depends on aural information.

4 Experiments

We conducted 3 experiments. First, to verify the claim that potential localization maps respond to objects that potentially produce sound and not to objects that can not produce sound such as tables or chairs, we compared the results of the potential localization map to a saliency map. Then, to investigate the contributions of each modality, we trained our unsupervised method and compared the performance of (i) our potential localization map (ii) our localization map (iii) an existing unsupervised sound source localization method. Finally, we also conducted a similar evaluation in a supervised setting.

4.1 Evaluation metrics

To evaluate our system, we employ the consensus intersection over union (cIoU) [11], which is similar to intersection over union (IoU), but can be applied when ground truth takes continuous values. The ground truth $G_{i,j}$ is given by

$$G_{i,j} = \sum_{k=1}^M \frac{G_{k,i,j}}{M}. \quad (7)$$

Here, M indicates the number of annotators and $G_{k,i,j}$ indicates ground truth by k -th annotator, which takes 1 if the bounding box includes (i,j) -th pixel, and takes 0 otherwise. The definition of cIoU is as follows.

$$\text{cIoU} = \frac{\sum_{i,j} G_{i,j} B_{i,j}}{\sum_{i,j} G_{i,j} + \sum_{(i,j) \in \{(i,j) | G_{i,j}=0\}} B_{i,j}}, \quad (8)$$

where $B_{i,j}$ is obtained by binarizing the prediction $P_{i,j}$ using a threshold τ .

$$B_{i,j} = \begin{cases} 1 & (P_{i,j} > \tau) \\ 0 & (\text{otherwise}) \end{cases}. \quad (9)$$

Intuitively, the numerator of cIoU indicates the intersection of the ground truth and the prediction, and the denominator indicates the union. Therefore, a higher cIoU score indicates better performance. We also report AUC score, which indicates the area under the curve of cIoU plotted by varying the threshold from 0 to 1. For evaluation, we use the same test set as Senocak et al. [11], which contains 250 pairs of image and sound.

4.2 Configurations

Saliency Map We obtain the saliency map for the top-1 predicted class and the top-100 predicted classes of the pretrained VGG network using Grad-CAM [53]. We conducted this experiment to address the question: is a potential localization map just a saliency map, thus visualizing the important region for image classification? This experiment allows us to differentiate whether potential localization maps respond only to objects that potentially produce sound or to any given object within the image, including those that can not produce sound. Additional details of this experiment are given in the supplementary material.

Unsupervised learning We trained our unsupervised method with various numbers of training samples (1k, 2.5k, 10k, 144k) using pretrained weights for the VGG network. As a means of comparison, both the potential localization map and the localization map were evaluated to investigate the degree to which aural information contributes to the overall performance in the benchmark dataset. We also trained our unsupervised method without pretrained weights, in order to make a fair comparison with results reported in the existing study [11], which

Table 1. Evaluation of localization map and potential localization map in an unsupervised setting. cIoU score with threshold 0.5 and AUC score are calculated for different sized training samples. We also report the results of random prediction. “de novo” indicates training without pretrained weights, and localization map is abbreviated as loc. map. Additional comparison with other methods are listed in the supplementary material.

# of training samples	loc. map by [11] (reported)		loc. map (pretrain.)		potential loc. map (pretrain.)		loc. map (de novo)		potential loc. map (de novo)	
	cIoU	AUC	cIoU	AUC	cIoU	AUC	cIoU	AUC	cIoU	AUC
1k	—	—	48.7	46.4	45.9	43.8	36.5	34.1	35.3	33.4
2.5k	—	—	50.3	47.7	48.1	46.0	40.7	37.3	38.5	36.8
10k	43.6	44.9	56.8	50.7	53.9	48.6	48.4	45.3	47.4	44.6
144k	66.0	55.8	68.4	57.0	66.8	56.2	66.7	56.3	65.5	55.5
Random			cIoU		AUC					
			34.1		32.3					

Table 2. Evaluation of potential localization map in a supervised setting, i.e. prediction of U-Net which only uses visual information. We use 2.5k samples for training as in the supervised setting of Senocak et al. [11].

Method	cIoU	AUC
loc. map by [11] (reported)	80.4	60.3
potential loc. map	79.3	60.9

does not use pretrained weights. In addition, based on this experiment, we can verify whether our model can be trained without pretrained weights.

Supervised learning To also investigate the contribution of visual information in a supervised setting, we trained a U-Net [54] with the ground truth $G_{i,j}$ as a target and use Dice-coefficient as a loss function. Dice-coefficient is defined as

$$DC = \frac{\sum_{i,j} G_{i,j} U_{i,j}}{\sum_{i,j} G_{i,j} + \sum_{i,j} U_{i,j}}, \quad (10)$$

where $U_{i,j}$ indicates the prediction of U-Net. Additional details including hyperparameters are given in the supplementary material.

4.3 Results and Analysis

Table 1 and Table 2 show the results reported by Senocak et al. [11] and the results of our potential localization map and localization map in an unsupervised setting and a supervised setting respectively. Table 3 shows results for the saliency maps. Insights obtained from these results are as follows.

Table 3. Evaluation of saliency maps for the top 1 class and the top 100 classes.

Method	cIoU	AUC
saliency map for the top1 class	45.7	41.1
saliency map for the top100 classes	52.7	46.3

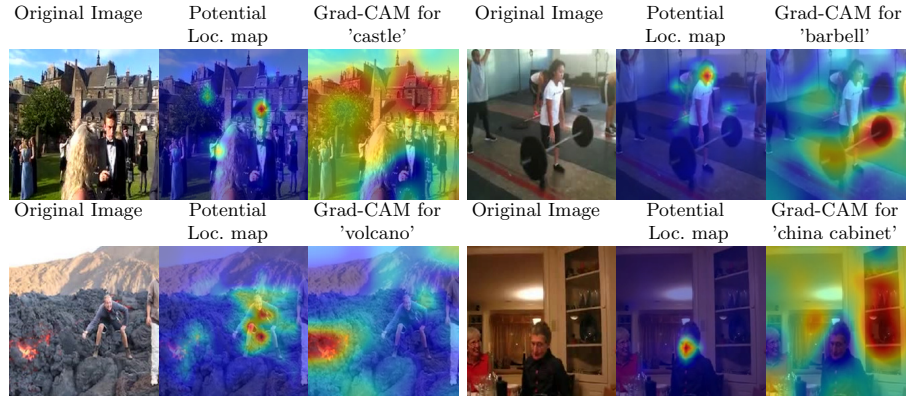


Fig. 3. Visualization results of potential localization map and saliency map (Grad-CAM) for the top 1 class. As the results show, the saliency maps react to objects that can not produce sound like “china cabinet”, while our potential localization map does not. For the visualization, we use the model trained with 144k samples and pretrained weights.

Behavioral analysis of potential localization map As you can see from Table 1 and Table 3, when more than 10k samples are used to conduct unsupervised training with pretrained weights, and when more than 144k samples are used to conduct training without pretrained weights, our potential localization map outperforms the baselines of saliency maps. This implies our potential localization map technique, trained with a sufficient number of samples can successfully identify objects that potentially produce sound and not just visualize a saliency map. As can be seen in Figure 3, it can be confirmed that our potential localization map focuses on objects that potentially produce sound, while a saliency map does not.

Assessment of individual modalities’ importance When the quantity of training data is within the range 10k to 144k, our localization maps achieve state-of-the-art performance (68.4%, 56.8% in cIoU). More notably, our potential localization map performs comparably to the localization map which uses both visual and aural information (1.6% gap in cIoU for 144k training samples with pretrained weights). The difference in performance is consistent in the supervised setting (1.1% gap in cIoU). This suggests that although sound source localization is a cross-modal task, when evaluated using the current benchmark dataset, visual information is dominant and the combined advantage of having both visual and aural information is not fully assessed.

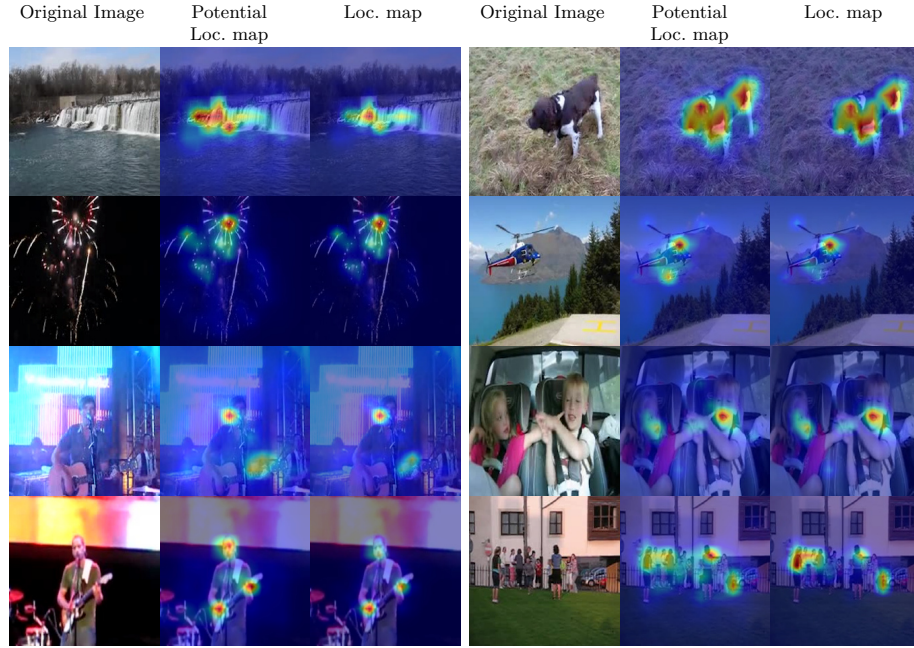


Fig. 4. Visualization results of potential localization maps and localization maps of various objects when all possible objects are producing sound (Type A). Samples in the two top rows only have one object capable of producing sound. Samples in the bottom two rows have multiple objects but the sound of both objects are given. In these cases the results show that the potential localization map and the localization map are very similar. For the visualization, we use the model trained with 144k samples and pretrained weights.

Analysis of the performance gap In order to investigate the difference in behavior between the potential localization map and the localization map, we make qualitative observations as shown in Figure 4 and Figure 5. As seen from Figure 4, it can be observed that in an unsupervised setting, when all objects capable of producing sound are actually producing sound, the potential localization map performs similarly to the localization map. We call this type of sample Type A. On the other hand, when only some of the objects capable of producing sound are actually producing sound, such as in Figures 5, the localization map performs better than the potential localization map. We call this type of sample Type B.

To quantitatively evaluate this difference in behavior between samples that are Type A and Type B, we manually picked out 30 samples for each type from Flickr-SoundNet [15] excluding the samples contained in the current benchmark dataset [11], and then annotated sound sources using bounding-boxes. Further details on the annotation process are given in the supplementary material. IoU and AUC scores for both types were calculated to check which type of samples leads to the difference in performance between potential localization map and localization map. The results are shown in Table 4. In accordance with our qual-

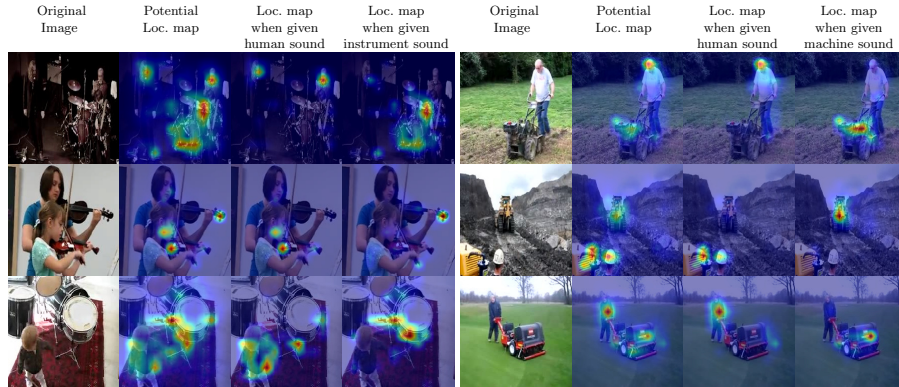


Fig. 5. Visualization results of potential localization maps and localization maps of samples with several objects capable of producing sound in the image, but of which only one object is actually producing sound (Type B). The localization map successfully responds to an object that is actually producing sound, while the potential localization map does not. For the visualization, we use the model trained with 144k samples and pretrained weights.

Table 4. IoU score with threshold 0.5 and AUC score for Type A and Type B samples. We use the model trained with 144k samples and pretrained weights.

	loc. map		potential loc. map		perform. gap		Random	
	IoU	AUC	IoU	AUC	IoU	AUC	IoU	AUC
Type A	58.5	49.6	55.7	48.3	2.8	1.3	29.9	28.4
Type B	44.5	38.0	31.9	27.2	12.6	10.8	17.1	16.6

itative observation, the results show that the performance gap is smaller in Type A and larger in Type B.

Based on the differences seen between samples of Type A and Type B, it can be implied that the lack of performance gap when evaluated with the current benchmark dataset is caused by the majority of samples within the dataset being Type A. This indicates that, inherently, the majority of sound sources can be inferred based on visual information only, thus confirming that the contribution of aural information is minimal in the current benchmark dataset. This is problematic as a proper evaluation of whether a system is capable of leveraging aural information can not be conducted. It is important to realize that in the real world, there are many examples where objects that are capable of producing sound are silent, and thus the usage of sound would be inevitable. Hence, datasets used for evaluation of sound source localization should necessitate the usage of sound.

5 Alternative Evaluation Protocol

From our observations drawn in Sec. 4.3, it can be deduced that a dataset which consists of abundant samples of Type B is required to properly assess a sys-

Table 5. Evaluation of localization map and potential localization map in an unsupervised setting following our evaluation protocol. cIoU score with threshold 0.5 and AUC score are calculated for different sized training samples.

# of training samples	loc. map (pretrain.)		potential loc. map (pretrain.)		perform. gap (pretrain.)		loc. map (de novo)		potential loc. map (de novo)		perform. gap (de novo)	
	cIoU	AUC	cIoU	AUC	cIoU	AUC	cIoU	AUC	cIoU	AUC	cIoU	AUC
1k	24.8	23.5	23.4	21.8	1.4	1.7	21.8	20.4	21.1	20.2	0.7	0.2
2.5k	27.3	24.4	23.5	21.3	3.8	3.1	23.2	21.5	22.9	20.8	0.3	0.7
10k	35.4	29.5	28.2	25.1	7.2	4.4	30.1	26.6	24.0	22.1	6.1	4.5
144k	41.6	35.6	30.2	26.4	11.4	9.2	37.4	31.6	28.9	25.4	8.5	6.2
Random					cIoU		AUC					
					20.0		18.5					

tem’s capability of leveraging aural information. However, as the retrieval of such samples in sufficient amounts is implausible, we introduce a *concat-and-localize* method, which is an alternative evaluation protocol that utilizes artificially created samples which structurally have the same properties as Type B.

Inspired by the mix-and-separate method used in audio-visual source separation [8–10, 12, 13, 34], the process of artificially creating a sample in our *concat-and-localize* method is as follows. First, two pairs of image and sound are sampled from the test set of the current dataset [11]. The two images are concatenated side-by-side and then rescaled to the size of the original image. Only one of the two sounds is given as input, and the new ground truth is defined as the rescaled original ground truth of the sample corresponding to the input sound as shown in Fig. 6. This process is conducted for both sounds, i.e. two synthetic samples are obtained per process. By design, the newly created sample always consists of multiple objects that are capable of producing sound as each original image contains such objects. In addition, as only one of the sounds is given, we can reproduce a situation where only some of the objects are producing sound, meeting the requirements for the sample to be Type B. Hence, in order to determine whether the sound originates from the objects on the left side of the concatenated image or the ones on the right side, aural information must be leveraged.

We created a new dataset of 1000 samples using the above process and evaluated our system using this dataset. The same metrics (cIoU and AUC) as Sec. 4 were utilized. In Table 5, we show the results of our potential localization map and localization map following the same configuration as Sec. 4. As expected, the results show that the performance gap is significantly greater than the currently adopted evaluation method, especially when training with a large number of samples (11.4% gap in cIoU with 144k training samples). This can be attributed to the design of our evaluation protocol, where the problem cannot be solved using only visual information. In this respect, our *concat-and-localize* method is more suitable for evaluating cross-modal perception task than simply using the current benchmark dataset.

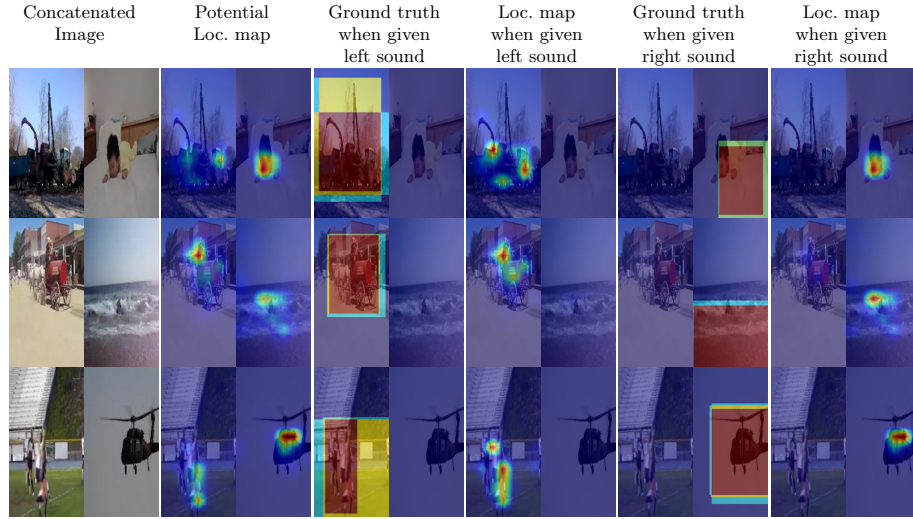


Fig. 6. Visualization results of potential localization maps and localization maps for artificially created samples. When given the concatenated image, the potential localization map shows response to all objects capable of producing sound in both images. Then when given only one of the sounds, the localization map correctly responds to the corresponding object. For the visualization, we use the model trained with 144k samples and pretrained weights.

6 Conclusion

In this paper, we develop an unsupervised architecture that solves sound source localization by decomposing this task into two steps: potential sound source localization and object selection. We provide sufficient evidence to show that our system can localize possible sound sources from only an image, and achieve state-of-the-art in sound source localization when leveraging both image and sound inputs. More importantly, our system, even when using only visual information, performs comparably to our state-of-the-art output which leverages both image and sound, suggesting sound sources can be localized without the usage of sound in the current dataset. From this observation and both qualitative and quantitative analysis, we pointed out the problem of the currently adopted evaluation method, and introduced the *concat-and-localize* method as an alternative. Our evaluation protocol is more suitable for evaluating the cross-modal aspect of sound source localization as the design enforces the usage of both modalities.

For future works, we believe it to be valuable to investigate the importance and contribution of different modalities in other cross-modal perception tasks too, as it may lead to a deeper understanding of the tasks, a dataset with better quality, and a better performance of future models in the real world.

Acknowledgements

This research is supported by the JST ACCEL (JPMJAC1602), JST-Mirai Program (JPMJMI19B2), JSPS KAKENHI (JP17H06101, JP19H01129 and JP19H04137).

References

1. Breiman, L.: Random forests. *Machine Learning* **45** (2001) 5–32 [1](#)
2. Chen, T., Guestrin, C.: Xgboost: A scalable tree boosting system. In: *International Conference on Knowledge Discovery and Data Mining (KDD)*. (2016) [1](#)
3. Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., Liu, T.Y.: Lightgbm: A highly efficient gradient boosting decision tree. In: *Neural Information Processing Systems (NIPS)*. (2017) [1](#)
4. Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A.V., Gulin, A.: Catboost: Unbiased boosting with categorical features. In: *Neural Information Processing Systems (NIPS)*. (2018) [1](#)
5. Arandjelovic, R., Zisserman, A.: Look, listen and learn. In: *International Conference on Computer Vision (ICCV)*. (2017) [3, 4, 7](#)
6. Arandjelović, R., Zisserman, A.: Objects that sound. In: *European Conference on Computer Vision (ECCV)*. (2018) [3, 4, 5](#)
7. Hu, D., Nie, F., Li, X.: Deep multimodal clustering for unsupervised audiovisual learning. In: *Computer Vision and Pattern Recognition (CVPR)*. (2019) [3, 4](#)
8. Rouditchenko, A., Zhao, H., Gan, C., McDermott, J., Torralba, A.: Self-supervised audio-visual co-segmentation. In: *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. (2019) [3, 4, 5, 13](#)
9. Zhao, H., Gan, C., Rouditchenko, A., Vondrick, C., McDermott, J., Torralba, A.: The sound of pixels. In: *European Conference on Computer Vision (ECCV)*. (2018) [3, 4, 5, 13](#)
10. Owens, A., Efros, A.A.: Audio-visual scene analysis with self-supervised multisensory features. In: *European Conference on Computer Vision (ECCV)*. (2018) [3, 4, 5, 13](#)
11. Senocak, A., Oh, T.H., Kim, J., Yang, M.H., So Kweon, I.: Learning to localize sound source in visual scenes. In: *Computer Vision and Pattern Recognition (CVPR)*. (2018) [3, 5, 8, 9, 11, 13](#)
12. Zhao, H., Gan, C., Ma, W.C., Torralba, A.: The sound of motions. In: *International Conference on Computer Vision (ICCV)*. (2019) [3, 4, 5, 13](#)
13. Gan, C., Huang, D., Zhao, H., Tenenbaum, J.B., Torralba, A.: Music gesture for visual sound separation. In: *Computer Vision and Pattern Recognition (CVPR)*. (2020) [3, 4, 5, 13](#)
14. Aytar, Y., Vondrick, C., Torralba, A.: See, hear, and read: Deep aligned representations. *arXiv preprint arXiv:1706.00932* (2017) [4](#)
15. Aytar, Y., Vondrick, C., Torralba, A.: Soundnet: Learning sound representations from unlabeled video. In: *Neural Information Processing Systems (NIPS)*. (2016) [4, 5, 11](#)
16. Harwath, D., Glass, J.R.: Learning word-like units from joint audio-visual analysis. In: *Association for Computational Linguistics (ACL)*. (2017) [4](#)
17. Harwath, D., Torralba, A., Glass, J.: Unsupervised learning of spoken language with visual context. In: *Neural Information Processing Systems (NIPS)*. (2016) [4](#)
18. Owens, A., Wu, J., McDermott, J.H., Freeman, W.T., Torralba, A.: Ambient sound provides supervision for visual learning. In: *European Conference on Computer Vision (ECCV)*. (2016) [4](#)
19. Tian, Y., Shi, J., Li, B., Duan, Z., Xu, C.: Audio-visual event localization in unconstrained videos. In: *European Conference on Computer Vision (ECCV)*. (2018) [4, 5](#)

20. Korbar, B., Tran, D., Torresani, L.: Cooperative learning of audio and video models from self-supervised synchronization. In: Neural Information Processing Systems (NIPS). (2018) 4, 5
21. Zhou, Y., Wang, Z., Fang, C., Bui, T., Berg, T.L.: Visual to sound: Generating natural sound for videos in the wild. In: Computer Vision and Pattern Recognition (CVPR). (2018) 4
22. Owens, A., Isola, P., McDermott, J., Torralba, A., Adelson, E.H., Freeman, W.T.: Visually indicated sounds. In: Computer Vision and Pattern Recognition (CVPR). (2016) 4
23. Gao, R., Grauman, K.: 2.5 d visual sound. In: Computer Vision and Pattern Recognition (CVPR). (2019) 4
24. Morgado, P., Nvasconcelos, N., Langlois, T., Wang, O.: Self-supervised generation of spatial audio for 360° video. In: Neural Information Processing Systems (NIPS). (2018) 4
25. Lyon, R.F.: A computational model of binaural localization and separation. In: International Conference on Acoustics, Speech and Signal Processing (ICASSP). (1983) 4
26. Hershey, J.R., Chen, Z., Roux, J.L., Watanabe, S.: Deep clustering: Discriminative embeddings for segmentation and separation. In: International Conference on Acoustics, Speech and Signal Processing (ICASSP). (2016) 4
27. Chen, Z., Luo, Y., Mesgarani, N.: Deep attractor network for single-microphone speaker separation. In: International Conference on Acoustics, Speech and Signal Processing (ICASSP). (2017) 4
28. Yu, D., Kolbæk, M., Tan, Z.H., Jensen, J.: Permutation invariant training of deep models for speaker-independent multi-talker speech separation. In: International Conference on Acoustics, Speech and Signal Processing (ICASSP). (2017) 4
29. Parekh, S., Essid, S., Ozerov, A., Duong, N., Pérez, P., Richard, G.: Motion informed audio source separation. In: International Conference on Acoustics, Speech and Signal Processing (ICASSP). (2017) 4
30. Sedighin, F., Babaie-Zadeh, M., Rivet, B., Jutten, C.: Two multimodal approaches for single microphone source separation. In: European Signal Processing Conference (EUSIPCO). (2016) 4
31. Gao, R., Feris, R., Grauman, K.: Learning to separate object sounds by watching unlabeled video. In: European Conference on Computer Vision (ECCV). (2018) 4
32. Smaragdis, P., Casey, M.: Audio/visual independent components. In: International Conference on Independent Component Analysis and Signal Separation (ICA). (2003) 4
33. Pu, J., Panagakis, Y., Petridis, S., Pantic, M.: Audio-visual object localization and separation using low-rank and sparsity. In: International Conference on Acoustics, Speech and Signal Processing (ICASSP). (2017) 4
34. Gao, R., Grauman, K.: Co-separating sounds of visual objects. In: International Conference on Computer Vision (ICCV). (2019) 4, 13
35. 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. In: Special Interest Group on Computer GRAPHics and Interactive Techniques (SIGGRAPH). (2018) 4
36. Gabbay, A., Shamir, A., Peleg, S.: Visual speech enhancement. In: International Speech Communication Association (INTERSPEECH). (2018) 4
37. Afouras, T., Chung, J.S., Zisserman, A.: The conversation: Deep audio-visual speech enhancement. In: International Speech Communication Association (INTERSPEECH). (2018) 4

38. Llagostera Casanovas, A., Monaci, G., Vandergheynst, P., Gribonval, R.: Blind audiovisual source separation based on sparse redundant representations. *Transactions on Multimedia* **12** (2010) 358–371 [4](#)
39. Nakadai, K., Okuno, H.G., Kitano, H.: Real-time sound source localization and separation for robot audition. In: *International Conference on Spoken Language Processing (ICSLP)*. (2002) [4](#)
40. Argentieri, S., Danès, P., Souères, P.: A survey on sound source localization in robotics: From binaural to array processing methods. *Comput. Speech Language* **34** (2015) 87–112 [4](#)
41. Nakamura, K., Nakadai, K., Asano, F., Ince, G.: Intelligent sound source localization and its application to multimodal human tracking. In: *International Conference on Intelligent Robots and Systems (IROS)*. (2011) [4](#)
42. Strobel, N., Spors, S., Rabenstein, R.: Joint audio-video object localization and tracking. *Signal Processing Magazine* **18** (2001) 22–31 [4](#)
43. Hershey, J.R., Movellan, J.R.: Audio vision: Using audio-visual synchrony to locate sounds. In: *Neural Information Processing Systems (NIPS)*. (1999) [4](#)
44. Kidron, E., Schechner, Y.Y., Elad, M.: Pixels that sound. In: *Computer Vision and Pattern Recognition (CVPR)*. (2005) [4](#)
45. 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)*. (2001) [4](#)
46. Barzelay, Z., Schechner, Y.: Harmony in motion. In: *Computer Vision and Pattern Recognition (CVPR)*. (2007) [4](#)
47. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Neural Information Processing Systems (NIPS)*. (2012) [5](#)
48. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: *International Conference on Learning Representations (ICLR)*. (2015) [5](#)
49. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *Computer Vision and Pattern Recognition (CVPR)*. (2009) [5](#)
50. Nair, V., Hinton, G.E.: Rectified linear units improve restricted boltzmann machines. In: *International Conference on Machine Learning (ICML)*. (2010) [5](#)
51. Lin, M., Chen, Q., Yan, S.: Network in network. In: *International Conference on Learning Representations (ICLR)*. (2014) [6](#)
52. Bromley, J., Guyon, I., LeCun, Y., Säckinger, E., Shah, R.: Signature verification using a "siamese" time delay neural network. In: *Neural Information Processing Systems (NIPS)*. (1994) [7](#)
53. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Gradcam: Visual explanations from deep networks via gradient-based localization. In: *International Conference on Computer Vision (ICCV)*. (2017) [8](#)
54. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. (2015) [9](#)