



Objects that Sound

Relja Arandjelović¹(✉) and Andrew Zisserman^{1,2}

¹ DeepMind, London, UK
relja@google.com

² VGG, Department of Engineering Science, University of Oxford, Oxford, UK

Abstract. In this paper our objectives are, first, networks that can embed audio and visual inputs into a common space that is suitable for cross-modal retrieval; and second, a network that can localize the object that sounds in an image, given the audio signal. We achieve both these objectives by training from unlabelled video using only *audio-visual correspondence* (AVC) as the objective function. This is a form of cross-modal self-supervision from video.

To this end, we design new network architectures that can be trained for cross-modal retrieval and localizing the sound source in an image, by using the AVC task. We make the following contributions: (i) show that audio and visual embeddings can be learnt that enable both within-mode (*e.g.* audio-to-audio) and between-mode retrieval; (ii) explore various architectures for the AVC task, including those for the visual stream that ingest a single image, or multiple images, or a single image and multi-frame optical flow; (iii) show that the semantic object that sounds within an image can be localized (using only the sound, no motion or flow information); and (iv) give a cautionary tale on how to avoid undesirable shortcuts in the data preparation.

1 Introduction

There has been a recent surge of interest in cross-modal learning from images and audio [1–4]. One reason for this surge is the availability of virtually unlimited training material in the form of videos (*e.g.* from YouTube) that can provide both an image stream and a (synchronized) audio stream, and this cross-modal information can be used to train deep networks. Cross-modal learning itself has a long history in computer vision, principally in the form of images and text [5–7]. Although audio and text share the fact that they are both sequential in nature, the challenges of using audio to partner images are significantly different to those of using text. Text is much closer to a semantic annotation than audio. With text, *e.g.* in the form of a provided caption of an image, the concepts (such as ‘a dog’) are directly available and the problem is then to provide a correspondence between the noun ‘dog’ and a spatial region in the image [5, 8]. Whereas, for audio, obtaining the semantics is less direct, and has more in common with image classification, in that the concept dog is not directly available from the signal but requires something like a ConvNet to obtain it (think of classifying

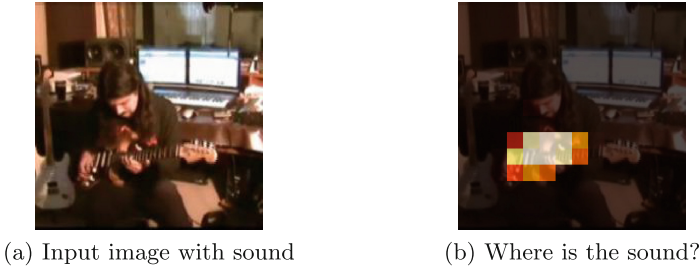


Fig. 1. Where is the sound? Given an input image and sound clip, our method learns, without a single labelled example, to localize the object that makes the sound.

an image as to whether it contains a dog or not, and classifying an audio clip as to whether it contains the sound of a dog or not).

In this paper our interest is in cross-modal learning from images and audio [1–4, 9–12]. In particular, we use unlabelled video as our source material, and employ *audio-visual correspondence* (AVC) as the training objective [4]. In brief, given an input pair of a video frame and 1 s of audio, the AVC task requires the network to decide whether they are in correspondence or not. The labels for the positives (matching) and negatives (mismatched) pairs are obtained directly, as videos provide an automatic alignment between the visual and the audio streams – frame and audio coming from the same time in a video are positives, while frame and audio coming from different videos are negatives. As the labels are constructed directly from the data itself, this is an example of “self-supervision” [13–22], a subclass of unsupervised methods.

The AVC task stimulates the learnt visual and audio representations to be both discriminative, to distinguish between matched and mismatched pairs, and semantically meaningful. The latter is the case because the only way for a network to solve the task is if it learns to classify semantic concepts in both modalities, and then judge whether the two concepts correspond. Recall that the visual network only sees a single frame of video and therefore it cannot learn to cheat by exploiting motion information.

In this paper we propose two networks that enable new functionalities: in Sect. 3 we propose a network architecture that produces *embeddings* directly suitable for cross-modal retrieval; in Sect. 4 we design a network and a learning procedure capable of *localizing* the sound source, *i.e.* answering the basic question – “Which object in an image is making the sound?”. An example is shown in Fig. 1. Both of these are trained from scratch with no labels whatsoever, using the same unsupervised audio-visual correspondence task (AVC).

2 Dataset

Throughout the paper we use the publicly available AudioSet dataset [23]. It consists of 10 s clips from YouTube with an emphasis on audio events, and video-level audio class labels (potentially more than 1 per video) are available, but

are noisy; the labels are organized in an ontology. To make the dataset more manageable and interesting for our purposes, we filter it for sounds of musical instruments, singing and tools, yielding 110 audio classes (the full list is given in the appendix [24], removing uninteresting classes like breathing, sine wave, sound effect, infrasound, silence, *etc.* The videos are challenging as many are of poor quality, the audio source is not always visible, and the audio stream can be artificially inserted on top of the video, *e.g.* it is often the case that a video is compiled of a musical piece and an album cover, text naming the song, still frame of the musician, or even completely unrelated visual motifs like a landscape, *etc.* The dataset already comes with a public train-test split, and we randomly split the public training set into training and validation sets in 90%–10% proportions. The final *AudioSet-Instruments* dataset contains 263k, 30k and 4.3k 10 s clips in the train, val and test splits, respectively.

We re-emphasise that no labels whatsoever are used for any of our methods since we treat the dataset purely as a collection of label-less videos. Labels are only used for quantitative evaluation purposes, *e.g.* to evaluate the quality of our unsupervised cross-modal retrieval (Sect. 3.1).

3 Cross-Modal Retrieval

In this section we describe a network architecture capable of learning good visual and audio embeddings from scratch and without labels. Furthermore, the two embeddings are aligned in order to enable querying across modalities, *e.g.* using an image to search for related sounds.

The *Audio-Visual Embedding Network (AVE-Net)* is designed explicitly to facilitate cross-modal retrieval. The input image and 1 s of audio (represented as a log-spectrogram) are processed by vision and audio subnetworks (Figs. 2a and b), respectively, followed by feature fusion whose goal is to determine whether the image and the audio correspond under the AVC task. The architecture is shown in full detail in Fig. 2c. To enforce feature alignment, the AVE-Net computes the correspondence score as a function of the Euclidean distance between the normalized visual and audio embeddings. This information bottleneck, the single scalar value that summarizes whether the image and the audio correspond, forces the two embeddings to be aligned. Furthermore, the use of the Euclidean distance during training is crucial as it makes the features “aware” of the distance metric, therefore making them amenable to retrieval [26].

The two subnetworks produce a 128-D L2 normalized embedding for each of the modalities. The Euclidean distance between the two 128-D features is computed, and this single scalar is passed through a tiny FC, which scales and shifts the distance to calibrate it for the subsequent softmax. The bias of the FC essentially learns the threshold on the distance above which the two features are deemed not to correspond.

Relation to Previous Works. The L^3 -Net introduced in [4] and shown in Fig. 2d, was also trained using the AVC task. However, the L^3 -Net audio and visual

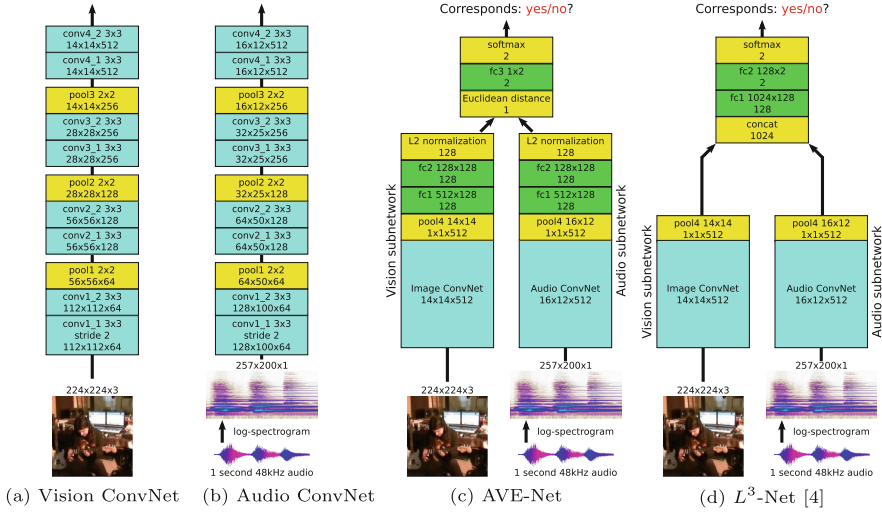


Fig. 2. ConvNet architectures. Each blocks represents a single layer with text providing more information – first row: layer name and optional kernel size, second row: output feature map size. Each convolutional layer is followed by batch normalization [25] and a ReLU nonlinearity, and the first fully connected layer (**fc1**) is followed by ReLU. All pool layers perform max pooling and their strides are equal to the kernel sizes. (a) and (b) show the vision and audio ConvNets which perform initial feature extraction from the image and audio inputs, respectively. (c) Our AVE-Net is designed to produce aligned vision and audio embeddings as the only information, a single scalar, used to decide whether the two inputs correspond is the Euclidean distance between the embeddings. (d) In contrast, the L^3 -Net [4] architecture combines the two modalities by concatenation and a couple of fully connected layers which produce the corresponds or not classification scores.

features are inadequate for cross-modal retrieval (as will be shown in the results of Sect. 3.1) as they are not aligned in any way – the fusion is performed by concatenating the features and the correspondence score is computed only *after* the fully connected layers. In contrast, the AVE-Net moves the fully connected layers into the vision and audio subnetworks and directly optimizes the features for cross-modal retrieval.

The training bears resemblance to metric learning via the contrastive loss [27], but (i) unlike contrastive loss which requires tuning of the margin hyper-parameter, ours is parameter-free, and (ii) it explicitly computes the corresponds-or-not output, thus making it directly comparable to the L^3 -Net while contrastive loss would require another hyper-parameter for the distance threshold. Wang *et al.* [28] also train a network for cross-modal retrieval but use a triplet loss which also contains the margin hyper-parameter, they use pretrained networks, and consider different modalities (image-text) with fully supervised correspondence labels. In concurrent work, Hong *et al.* [29] use a similar technique with pretrained networks and triplet loss for joint embedding of music and

video. Recent work of [12] also trains networks for cross-modal retrieval, but uses an ImageNet pretrained network as a teacher. In our case, we train the entire network from scratch.

3.1 Evaluation and Results

The architectures are trained on the AudioSet-Instruments train-val set, and evaluated on the AudioSet-Instruments test set described in Sect. 2. Implementation details are given below in Sect. 3.3.

On the audio-visual correspondence task, AVE-Net achieves an accuracy of 81.9%, beating slightly the L^3 -Net which gets 80.8%. However, AVC performance is not the ultimate goal since the task is only used as a proxy for learning good embeddings, so the real test of interest here is the retrieval performance.

To evaluate the intra-modal (*e.g.* image-to-image) and cross-modal retrieval, we use the AudioSet-Instruments test dataset. A single frame and surrounding 1s of audio are sampled randomly from each test video to form the retrieval database. All combinations of image/audio as query and image/audio as database are tested, *e.g.* audio-to-image uses the *audio* embedding as the query vector to search the database of visual embeddings, answering the question “Which image could make this sound?”; and image-to-image uses the *visual* embedding as the query vector to search the same database.

Evaluation Metric. The performance of a retrieval system is assessed using a standard measure – the normalized discounted cumulative gain (nDCG). It measures the quality of the ranked list of the top k retrieved items (we use $k = 30$ throughout) normalized to the $[0, 1]$ range, where 1 signifies a perfect ranking in which items are sorted in a non-increasing relevance-to-query order. For details on the definition of the relevance, refer to the appendix [24]. Each item in the test dataset is used as a query and the average nDCG@30 is reported as the final retrieval performance. Recall that the labels are noisy, and note that we only extract a single frame/1s audio per video and can therefore miss the relevant event, so the ideal nDCG of 1 is highly unlikely to be achievable.

Baselines. We compare to the L^3 -Net as it is also trained in an unsupervised manner, and we train it using an identical procedure and training data to our method. As the L^3 -Net is expected not to work for cross-modal retrieval since the representation are not aligned in any way, we also test the L^3 -Net representations aligned with CCA as a baseline. In addition, vision features extracted from the last hidden layer of the VGG-16 network trained in a fully-supervised manner on ImageNet [30] are evaluated as well. For cross-modal retrieval, the VGG16-ImageNet visual features are aligned with the L^3 -Net audio features using CCA, which is a strong baseline as the vision features are fully-supervised while the audio features are state-of-the-art [4]. Note that the vanilla L^3 -Net produces 512-D representations, while VGG16 yields a 4096-D visual descriptor. For computational reasons, and for fair comparison with our AVE-Net which produces 128-D embeddings, all CCA-based methods use 128 components. For

all cases the representations are L2-normalized as we found this to significantly improve the performance; note that AVE-Net includes L2-normalization in the architecture and therefore the re-normalization is redundant.

Table 1. Cross-modal and intra-modal retrieval. Comparison of our method with unsupervised and supervised baselines in terms of the average nDCG@30 on the AudioSet-Instruments test set. The columns headers denote the modalities of the query and the database, respectively, where *im* stands for *image* and *aud* for *audio*. Our AVE-Net beats all baselines convincingly.

Method	im-im	im-aud	aud-im	aud-aud
Random chance	.407	.407	.407	.407
L^3 -Net [4]	.567	.418	.385	.653
L^3 -Net with CCA	.578	.531	.560	.649
VGG16-ImageNet [30]	.600	–	–	–
VGG16-ImageNet + L^3 -Audio CCA	.493	.458	.464	.618
AVE-Net	.604	.561	.587	.665

Results. The nDCG@30 for all combinations of query-database modalities is shown in Table 1. For intra-modal retrieval (image-image, audio-audio) our AVE-Net is better than all baselines including slightly beating VGG16-ImageNet for image-image, which was trained in a fully supervised manner on another task. It is interesting to note that our network has never seen same-modality pairs during training, so it has not been trained explicitly for image-image and audio-audio retrieval. However, intra-modal retrieval works because of transitivity – an image of a violin is close in feature space to the sound of a violin, which is in turn close to other images of violins. Note that despite learning essentially the same information on the same task and training data as the L^3 -Net, our AVE-Net outperforms the L^3 -Net because it is Euclidean distance “aware”, *i.e.* it has been designed and trained with retrieval in mind.

For cross-modal retrieval (image-audio, audio-image), AVE-Net beats all baselines, verifying that our unsupervised training is effective. The L^3 -Net representations are clearly not aligned across modalities as their cross-modal retrieval performance is on the level of random chance. The L^3 -Net features aligned with CCA form a strong baseline, but the benefits of directly training our network for alignment are apparent. It is interesting that aligning vision features trained on ImageNet with state-of-the-art L^3 -Net audio features using CCA performs worse than other methods, demonstrating a case for unsupervised learning from a more varied dataset, as it is not sufficient to just use ImageNet-pretrained networks as black-box feature extractors.

Figure 3 shows some qualitative retrieval results, illustrating the efficacy of our approach. The system generally does retrieve relevant items from the database, while making reasonable mistakes such as confusing the sound of a zither with an acoustic guitar.

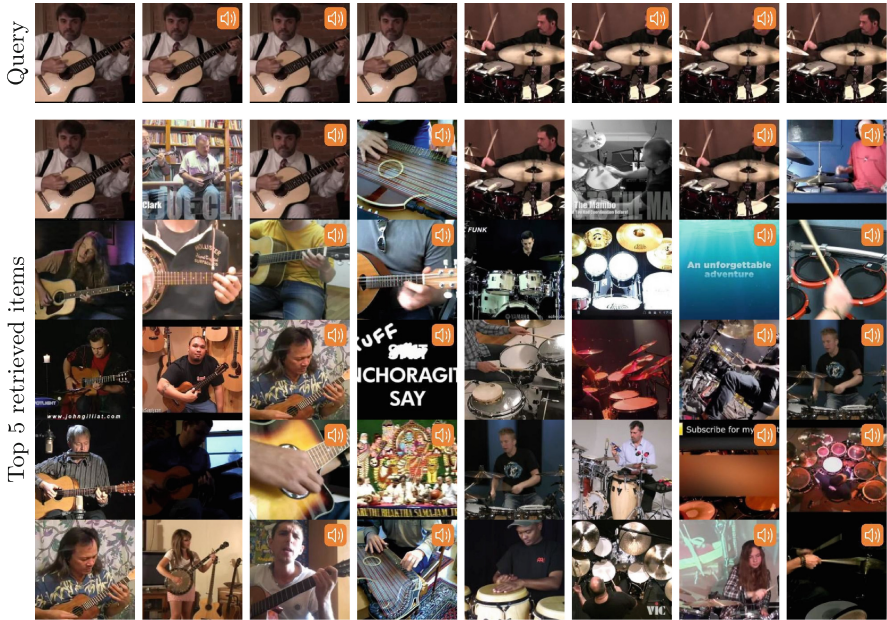


Fig. 3. Cross-modal and intra-modal retrieval. Each column shows one query and retrieved results. Purely for visualization purposes, as it is hard to display sound, the frame of the video that is aligned with the sound is shown instead of the actual sound form. The sound icon or lack of it indicates the audio or vision modality, respectively. For example, the last column illustrates query by image into an audio database, thus answering the question “Which sounds are the most plausible for this query image?” Note that many audio retrieval items are indeed correct despite the fact that their corresponding frames are unrelated – *e.g.* the audio of the blue image with white text does contain drums – this is just an artefact of how noisy real-world YouTube videos are.

3.2 Extending the AVE-Net to Multiple Frames

It is also interesting to investigate whether using information from multiple frames can help solving the AVC task. For these results only, we evaluate two modifications to the architecture from Fig. 2a to handle a different visual input – multiple frames (AVE+MF) and optical flow (AVE+OF). For conciseness, the details of the architectures are explained in the appendix [24], but the overall idea is that for AVE+MF we input 25 frames and convert convolution layers from 2D to 3D, while for AVE+OF we combine information from a single frame and 10 frames of optical flow using a two-stream network in the style of [31].

The performance of the AVE+MF and AVE+OF networks on the AVC task are 84.7% and 84.9%, respectively, compared to our single input image network’s 81.9%. However, when evaluated on retrieval, they fail to provide a boost, *e.g.* the AVE+OF network achieves 0.608, 0.558, 0.588, and 0.665 for im-im, im-aud, aud-im and aud-aud, respectively; this is comparable to the performance

of the vanilla AVE-Net that uses a single frame as input (Table 1). One explanation of this underwhelming result is that, as is the case with most unsupervised approaches, the performance on the training objective is not necessarily in perfect correlation with the quality of learnt features and their performance on the task of interest. More specifically, the AVE+MF and AVE+OF could be using the motion information available at input to solve the AVC task more easily by exploiting some lower-level information (*e.g.* changes in the motion could be correlated with changes in sound, such as when seeing the fingers playing a guitar or flute), which in turn provides less incentive for the network to learn good semantic embeddings. For this reason, a single frame input is used for all other experiments.

3.3 Preventing Shortcuts and Implementation

Preventing Shortcuts. Deep neural networks are notorious for finding subtle data shortcuts to exploit in order to “cheat” and thus not learn to solve the task in the desired manner; an example is the misuse of chromatic aberration in [14] to solve the relative-position task. To prevent such behaviour, we found it important to carefully implement the sampling of AVC negative pairs to be as similar as possible to the sampling of positive pairs. In detail, a positive pair is generated by sampling a random video, picking a random frame in that video, and then picking a 1 s audio with the frame at its mid-point. It is tempting to generate a negative pair by randomly sampling two different videos and picking a random frame from one and a random 1 s audio clip from the other. However, this produces a slight statistical difference between positive and negative audio samples, in that the mid-point of the positives is always aligned with a frame and is thus at a multiple of 0.04 s (the video frame rate is 25 fps), while negatives have no such restrictions. This allows a shortcut as it appears the network is able to learn to recognize audio samples taken at multiples of 0.04 s, therefore distinguishing positives from negatives. It probably does so by exploiting low-level artefacts of MPEG encoding and/or audio resampling. Therefore, with this naive implementation of negative pair generation the network has less incentive to strongly learn semantically meaningful information.

To prevent this from happening, the audio for the negative pair is also sampled only from multiples of 0.04 s. Without shortcut prevention, the AVE-Net achieves an artificially high accuracy of 87.6% on the AVC task, compared to 81.9% with the proper sampling safety mechanism in place, but the performance of the network without shortcut prevention on the retrieval task is consistently 1–2% worse. Note that, for fairness, we train the L^3 -Net with shortcut prevention as well.

The L^3 -Net training in [4] does not encounter this problem due to performing additional data augmentation by randomly misaligning the audio and the frame by up to 1 s for both positives and negatives. We apply this augmentation as well, but our observation is important to keep in mind for future unsupervised approaches where exact alignment might be required, such as audio-visual synchronization.

Implementation Details. We follow the same setup and implementation details as in [4]. Namely, the input frame is a 224×224 colour image, while the 1 s of audio is resampled at 48 kHz, converted into a log-spectrogram (window length 0.01 s and half-window overlap) and treated as a 257×200 greyscale image. Standard data augmentation is used – random cropping, horizontal flipping and brightness and saturation jittering for vision, and random clip-level amplitude jittering for audio. The network is trained with cross-entropy loss for the binary classification task – whether the image and the audio correspond or not – using the Adam optimizer [32], weight decay 10^{-5} , and learning rate obtained by grid search. Training is done using 16 GPUs in parallel with synchronous updates implemented in TensorFlow, where each worker processes a 128-element batch, thus making the effective batch size 2048.

Note that the only small differences from the setup of [4] are that: (i) We use a stride of 2 pixels in the first convolutional layers as we found it to not affect the performance while yielding a $4\times$ speedup and saving in GPU memory, thus enabling the use of $4\times$ larger batches (the extra factor of $2\times$ is through use of a better GPU); and (ii) We use a learning rate schedule in the style of [33] where the learning rate is decreased by 6% every 16 epochs. With this setup we are able to fully reproduce the L^3 -Net results of [4], achieving even slightly better performance (+0.5% on the ESC-50 classification benchmark [34]), probably due to the improved learning rate schedule and the use of larger batches.

4 Localizing Objects that Sound

A system which understands the audio-visual world should associate appearance of an object with the sound it makes, and thus be able to answer “where is the object that is making the sound?” Here we outline an architecture and a training procedure for learning to localize the sounding object, while still operating in the scenario where there is no supervision, neither on the object location level nor on their identities. We again make use of the AVC task, and show that by designing the network appropriately, it is possible to learn to localize sounding objects in this extremely challenging label-less scenario.

In contrast to the standard AVC task where the goal is to learn a single embedding of the entire image which explains the sound, the goal in sound localization is to find regions of the image which explain the sound, while other regions should not be correlated with it and belong to the background. To operationalize this, we formulate the problem in the Multiple Instance Learning (MIL) framework [35]. Namely, local region-level image descriptors are extracted on a spatial grid and a similarity score is computed between the audio embedding and each of the vision descriptors. For the goal of finding regions which correlate well with the sound, the maximal similarity score is used as the measure of the image-audio agreement. The network is then trained in the same manner as for the AVC task, *i.e.* predicting whether the image and the audio correspond. For corresponding pairs, the method encourages one region to respond highly and therefore localize the object, while for mismatched pairs the maximal score

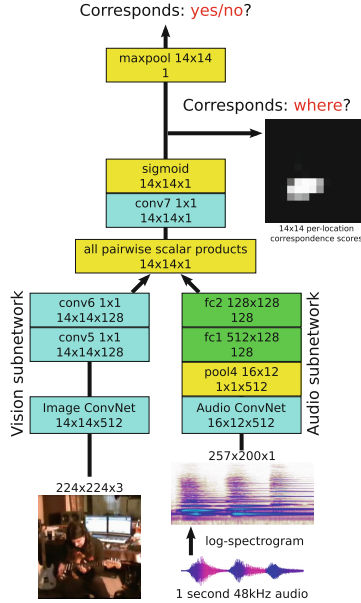


Fig. 4. Audio-Visual Object Localization (AVOL-Net). The notation and some building blocks are shared with Fig. 2. The audio subnetwork is the same as in AVE-Net (Fig. 2c). The vision network, instead of globally pooling the feature tensor, continues to operate at the 14×14 resolution, with relevant FCs (vision-**fc1**, vision-**fc2**, **fc3**) converted into their “fully convolutional” equivalents (*i.e.* 1×1 convolutions **conv5**, **conv6**, **conv7**). The similarities between the audio and all vision embeddings reveal the location of the object that makes the sound, while the maximal similarity is used as the correspondence score.

should be low thus making the entire score map low, indicating, as desired, there is no object which makes the input sound. In essence, the audio representation forms a filter which “looks” for relevant image patches in a similar manner to an attention mechanism.

Our *Audio-Visual Object Localization Network (AVOL-Net)* is depicted in Fig. 4. Compared to the AVE-Net (Fig. 2c), the vision subnetwork does not pool **conv4.2** features but keeps operating on the 14×14 resolution. To enable this, the two fully connected layers **fc1** and **fc2** of the vision subnetwork are converted to 1×1 convolutions **conv5** and **conv6**. Feature normalization is removed to enable features to have a low response on background regions. Similarities between each of the 14×14 128-D visual descriptors and the single 128-D audio descriptor are computed via a scalar product, producing a 14×14 similarity score map. Similarly to the AVE-Net, the scores are calibrated using a tiny 1×1 convolution (**fc3** converted to be “fully convolutional”), followed by a sigmoid which produces the localization output in the form of the image-audio correspondence score for each spatial location. Max pooling over all spatial locations is performed to obtain the final correspondence score, which is then used for training on the AVC task using the logistic loss.

Relation to Previous Works. While usually hinting at object localization, previous cross-modal works fall short from achieving this goal. Harwath *et al.* [2] demonstrate localizing objects in the audio domain of a spoken text, but do not design their network for localization. In [4], the network, trained from scratch, internally learns object detectors, but has never been demonstrated to be able to answer the question “Where is the object that is making the sound?”, nor, unlike our approach, was it trained with this ability in mind. Rather, their heatmaps are produced by examining responses of its various neurons given *only* the input image. The output is computed *completely independently of the sound* and therefore cannot answer “Where is the object that is making the sound?”.

Our approach has similarities with [36,37] who used max and average pooling, respectively, to learn object detectors without bounding box annotations in the single visual modality setting, but use ImageNet pretrained networks and image-level labels. The MIL-based approach also has connections with attention mechanisms as it can be viewed as “infinitely hard” attention [8,38]. Note that we do not use information from multiple audio channels which could aid localization [39] because (i) this setup generally requires known calibration of the multi-microphone rig which is unknown for unconstrained YouTube videos, (ii) the number of channels changes across videos, (iii) quality of audio on YouTube varies significantly while localization methods based on multi-microphone information are prone to noise and reverberation, and (iv) we desire that our system learns to detect semantic concepts rather than localize by “cheating” through accessing multi-microphone information. Finally, a similar technique to ours appears in the concurrent work of [40], while later works of [41,42] are also relevant.

4.1 Evaluation and Results

First, the accuracy of the localization network (AVOL-Net) on the AVC task is the same as that of the AVE-Net embedding network in Sect. 3, which is encouraging as it means that switching to the MIL setup does not cause a loss in accuracy and the ability to detect semantic concepts in the two modalities.

The ability of the network to localize the object(s) that sound is demonstrated in Fig. 5. It is able to detect a wide range of objects in different viewpoints and scales, and under challenging imaging conditions. A more detailed discussion including the analysis of some failure cases is available in the figure caption. As expected from an unsupervised method, it is not necessarily the case that it detects the entire object but can focus only on specific discriminative parts such as the interface between the hands and the piano keyboard. This interacts with the more philosophical question of what is an object and what is it that is making the sound – the body of the piano and its strings, the keyboard, the fingers on the keyboard, the whole human together with the instrument, or the entire orchestra? How should a gramophone or a radio be handled by the system, as they can produce arbitrary sounds?

From the impressive results in Fig. 5, one question that comes to mind is whether the network is simply detecting the salient object in the image, which

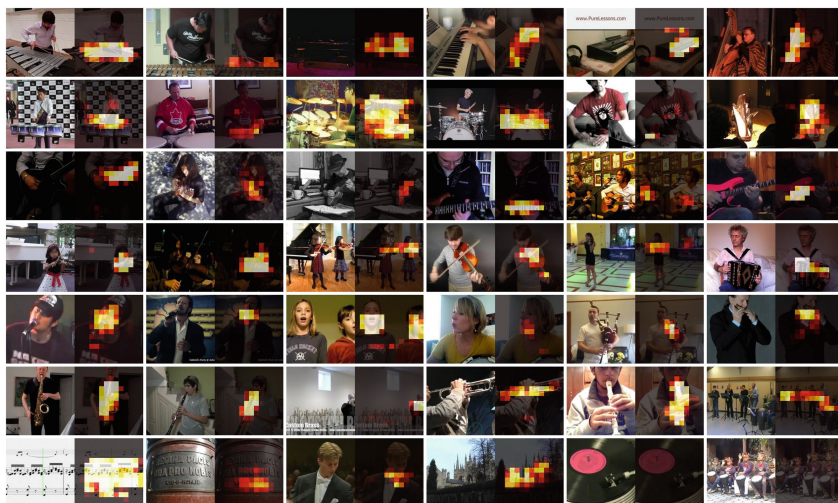


Fig. 5. What is making the sound? Localization output of the AVOL-Net on the unseen test data; see Fig. 1 and <https://goo.gl/JVsJ7P> for more. Recall that the network sees a single frame and therefore cannot “cheat” by using motion information. Each pair of images shows the input frame (left) and the localization output for the input frame and 1s of audio around it, overlaid over the frame (right). Note the wide range of detectable objects, such as keyboards, accordions, drums, harps, guitars, violins, xylophones, people’s mouths, saxophones, etc. Sounding objects are detected despite significant clutter and variations in lighting, scale and viewpoint. It is also possible to detect multiple relevant objects: two violins, two people singing, and an orchestra. The final row shows failure cases, where the first two likely reflects the noise in the training data as many videos contain just music sheets or text overlaid with music playing, in columns 3–4 the network probably just detects the salient parts of the scene, while in columns 5–6 it fails to detect the sounding objects.

is not the desired behaviour. To test this hypothesis we can provide mismatched frame and audio pairs as inputs to interrogate the network to answer “what would make this sound?”, and check if salient objects are still highlighted regardless of the irrelevant sound. Figure 6 shows that this is indeed not the case, as when, for example, drums are played on top of an image of a violin, the localization map is empty. In contrast, when another violin is played, the network highlights the violin. Furthermore, to completely reject the saliency hypothesis – in the case of an image depicting a piano and a flute, it is possible to play a flute sound and the network will pick the flute, while if a piano is played, the piano is highlighted in the image. Therefore, the network has truly learnt to disentangle multiple objects in an image and maintain a discriminative embedding for each of them.

To evaluate the localization performance quantitatively, 500 clips are sampled randomly from the validation data and the middle frame annotated with the localization of the instrument producing the sound. We then compare two



Fig. 6. What would make this sound? Similarly to Fig. 5, the AVOL-Net localization output is shown given an input image frame and 1 s of audio. However, here the frame and audio are mismatched. Each triplet of images shows the (left) input audio, (middle) input frame, and (right) localization output overlaid over the frame. Purely for visualization purposes, as it is hard to display sound, the frame of the video that is aligned with the sound is shown instead of the actual sound form (left). On the example of the first triplet: (left) flute sound illustrated by an image of a flute, (middle) image of a piano and a flute, (right) the flute from the middle image is highlighted as our network successfully answers the question “What in the piano-flute image would make a flute sound?” In each row the input frame is fixed while the input audio varies, showing that object localization does depend on the sound and therefore our system is not just detecting salient objects in the scene but is achieving the original goal – localizing the object that sounds.

methods of predicting the localization (as in [36]): first, a baseline method that always predicts the center of the image; second, the mode of the AVOL-Net heatmap produced by inputting the sound of the clip. The baseline achieves 57.2%, whilst AVOL-Net achieves 81.7%. This demonstrates that the AVOL-NET is not simply highlighting the salient object at the center of the image. Failure cases are mainly due to the problems with the AudioSet dataset described in Sect. 2. Note, it is necessary to annotate the data, rather than using a standard benchmark, since datasets such as PASCAL VOC, COCO, DAVIS, KITTI, do not contain musical instruments. This also means that off-the-shelf object detectors for instruments are not available, so could not be used to annotate AudioSet frames with bounding boxes.

Finally, Fig. 7 shows the localization results on videos. Note that each video frame and surrounding audio are processed completely independently, so no motion information is used, nor is there any temporal smoothing. The results reiterate the ability of the system to detect an object under a variety of poses, and to highlight different objects depending on the varying audio context. Please see YouTube playlist <https://goo.gl/JVsJ7P> for more video results.



Fig. 7. What is making the sound? The visualization is the same as for Fig. 5 but here each column contains frames from a single video, taken 1 s apart. The frames are processed completely independently, motion information is not used, nor there is any temporal smoothing. Our method reliably detects the sounding object across varying poses (columns 1–2), and shots (column 3). Furthermore, it is able to switch between objects that are making the sound such as interleaved speech and guitar during a guitar lesson (column 4).

5 Conclusions and Future Work

We have demonstrated that the unsupervised audio-visual correspondence task enables, with appropriate network design, two entirely new functionalities to be learnt: cross-modal retrieval, and semantic based localization of objects that sound. The AVE-Net was shown to perform cross-modal retrieval even better than supervised baselines, while the AVOL-Net exhibits impressive object localization capabilities. Potential improvements could include modifying the AVOL-Net to have an explicit soft attention mechanism, rather than the max-pooling used currently.

Acknowledgements. We thank Carl Doersch for useful insights regarding preventing shortcuts.

References

1. Aytar, Y., Vondrick, C., Torralba, A.: SoundNet: learning sound representations from unlabeled video. In: NIPS (2016)
2. Harwath, D., Torralba, A., Glass, J.R.: Unsupervised learning of spoken language with visual context. In: NIPS (2016)
3. 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
4. Arandjelović, R., Zisserman, A.: Look, listen and learn. In: Proceedings of ICCV (2017)
5. Barnard, K., Duygulu, P., de Freitas, N., Forsyth, D., Blei, D., Jordan, M.: Matching words and pictures. *JMLR* **3**, 1107–1135 (2003)
6. Duygulu, P., Barnard, K., de Freitas, J.F.G., Forsyth, D.A.: Object recognition as machine translation: learning a lexicon for a fixed image vocabulary. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) ECCV 2002. LNCS, vol. 2353, pp. 97–112. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-47979-1_7
7. Frome, A., et al.: Devise: a deep visual-semantic embedding model. In: NIPS (2013)
8. Xu, K., et al.: Show, attend and tell: neural image caption generation with visual attention. arXiv preprint [arXiv:1502.03044](https://arxiv.org/abs/1502.03044) (2015)
9. de Sa, V.R.: Learning classification from unlabelled data. In: NIPS (1994)
10. Kidron, E., Schechner, Y.Y., Elad, M.: Pixels that sound. In: Proceedings of CVPR (2005)
11. Owens, A., Isola, P., McDermott, J.H., Torralba, A., Adelson, E.H., Freeman, W.T.: Visually indicated sounds. In: Proceedings of CVPR, pp. 2405–2413 (2016)
12. Aytar, Y., Vondrick, C., Torralba, A.: See, hear, and read: deep aligned representations. CoRR abs/1706.00932 (2017)
13. Dosovitskiy, A., Springenberg, J.T., Riedmiller, M., Brox, T.: Discriminative unsupervised feature learning with convolutional neural networks. In: NIPS (2014)
14. Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: Proceedings of CVPR (2015)
15. Agrawal, P., Carreira, J., Malik, J.: Learning to see by moving. In: Proceedings of ICCV (2015)
16. Wang, X., Gupta, A.: Unsupervised learning of visual representations using videos. In: Proceedings of ICCV, pp. 2794–2802 (2015)
17. Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9907, pp. 649–666. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46487-9_40
18. Misra, I., Zitnick, C.L., Hebert, M.: Shuffle and learn: unsupervised learning using temporal order verification. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9905, pp. 527–544. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46448-0_32
19. Pathak, D., Krähenbühl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: feature learning by inpainting. In: Proceedings of CVPR, pp. 2536–2544 (2016)
20. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9910, pp. 69–84. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46466-4_5

21. Fernando, B., Bilen, H., Gavves, E., Gould, S.: Self-supervised video representation learning with odd-one-out networks. In: Proceedings of ICCV (2017)
22. Doersch, C., Zisserman, A.: Multi-task self-supervised visual learning. In: Proceedings of ICCV (2017)
23. Gemmeke, J.F., et al.: Audio set: an ontology and human-labeled dataset for audio events. In: ICASSP (2017)
24. Arandjelović, R., Zisserman, A.: Objects that sound. CoRR abs/1712.06651 (2017)
25. Ioffe, S., Szegedy, C.: Batch normalization: accelerating deep network training by reducing internal covariate shift. In: Proceedings of ICML (2015)
26. Arandjelović, R., Gronat, P., Torii, A., Pajdla, T., Sivic, J.: NetVLAD: CNN architecture for weakly supervised place recognition. In: IEEE PAMI (2017)
27. Chopra, S., Hadsell, R., LeCun, Y.: Learning a similarity metric discriminatively, with application to face verification. In: Proceedings of CVPR, vol. 1, pp. 539–546. IEEE (2005)
28. Wang, L., Li, Y., Lazebnik, S.: Learning deep structure-preserving image-text embeddings. In: Proceedings of CVPR (2016)
29. Hong, S., Im, W., S. Yang, H.: CBVMR: content-based video-music retrieval using soft intra-modal structure constraint. In: ACM ICMR (2018)
30. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: International Conference on Learning Representations (2015)
31. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: NIPS (2014)
32. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. In: Proceedings of ICLR (2015)
33. Szegedy, C., et al.: Going deeper with convolutions. In: Proceedings of CVPR (2015)
34. Piczak, K.J.: ESC: dataset for environmental sound classification. In: Proceedings of ACMM (2015)
35. Dietterich, T.G., Lathrop, R.H., Lozano-Perez, T.: Solving the multiple instance problem with axis-parallel rectangles. *Artif. Intell.* **89**(1–2), 31–71 (1997)
36. Oquab, M., Bottou, L., Laptev, I., Sivic, J.: Is object localization for free? - Weakly-supervised learning with convolutional neural networks. In: Proceedings of CVPR (2015)
37. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: Proceedings of CVPR (2016)
38. Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. In: Proceedings of ICLR (2015)
39. Shivappa, S.T., Rao, B.D., Trivedi, M.M.: Audio-visual fusion and tracking with multilevel iterative decoding: framework and experimental evaluation. *IEEE J. Sel. Top. Signal Process.* **4**(5), 882–894 (2010)
40. Senocak, A., Oh, T.H., Kim, J., Yang, M.H., Kweon, I.S.: On learning association of sound source and visual scenes. In: Proceedings of CVPR (2018)
41. Zhao, H., Gan, C., Rouditchenko, A., Vondrick, C., McDermott, J., Torralba, A.: The sound of pixels. In: Ferrari, (eds.) ECCV 2018, Part I. LNCS, vol. 11205, pp. 587–604. Springer, Cham (2018)
42. Owens, A., Efros, A.A.: Audio-visual scene analysis with self-supervised multisensory features. In: Proceedings of ECCV (2018, to appear)