

Reading to Listen at the Cocktail Party: Multi-Modal Speech Separation

Akam Rahimi, Triantafyllos Afouras, Andrew Zisserman
VGG, Department of Engineering Science, University of Oxford, UK

{akam, afourast, az}@robots.ox.ac.uk

<https://www.robots.ox.ac.uk/~vgg/research/voiceformer>

Abstract

The goal of this paper is speech separation and enhancement in multi-speaker and noisy environments using a combination of different modalities. Previous works have shown good performance when conditioning on temporal or static visual evidence such as synchronised lip movements or face identity. In this paper, we present a unified framework for multi-modal speech separation and enhancement based on synchronous or asynchronous cues. To that end we make the following contributions: (i) we design a modern Transformer-based architecture tailored to fuse different modalities to solve the speech separation task in the raw waveform domain; (ii) we propose conditioning on the textual content of a sentence alone or in combination with visual information; (iii) we demonstrate the robustness of our model to audio-visual synchronisation offsets; and, (iv) we obtain state-of-the-art performance on the well-established benchmark datasets LRS2 and LRS3.

1. Introduction

Humans have the remarkable ability to focus on conversations even in a room full of talking people, a phenomenon known as the “cocktail party effect” [29]. Our brains carry out this feat by concentrating their attention to a specific speaker while filtering out the rest of the stimuli originating from interfering voices and other environmental noises. Although this ability manifests to a limited extent through hearing alone, it is greatly enhanced when simultaneous information from other modalities is available. For example watching a speaker’s lips can significantly help disambiguate speech in noise [41], while understanding the natural language context of a sentence enables the listener to anticipate the potential next words of the speaker.

In recent years, progress in audio-visual learning has made it possible for machines to also achieve this ability and very effectively isolate individual voices out of multi-




Figure 1. We propose VoiceFormer, a framework for multi-modal speech separation and enhancement, that isolates speech according to either the text content of the target speaker’s utterance, their lip movements or both. Our framework allows conditioning on cues from multiple modalities, without requiring them to be temporally synchronised or have a common temporal rate. This gives it multiple advantages, such as robustness to temporal misalignments between the inputs.

speaker mixtures of speech or noisy audio [1, 14, 31]. Solving this problem enables a great range of practical applications, such as improving subtitle generation in videos with noisy audio, developing smart audio-visual hearing aids to enhance speech conditioned on visual input, or facilitating teleconferencing in noisy settings such as airports or cars.

Previous works have principally taken two approaches: either using synchronous cues, most commonly the lip movements of the target speaker [1, 14]; or using static (fixed embedding) cues such as the voice [49] or face characteristics [11, 21]. The former has the advantage of using dynamic evidence that is very strongly correlated to the desired speech output. However, relying on lip movements has several disadvantages. First they may be momentarily disrupted – e.g. from visual occlusions – therefore strong reliance will make the model sensitive to this form of visual noise; and second, they require synchronisation between the audio and visual streams. On the other hand, static cues arising from biometric characteristics of the speaker are more robust to temporary disruptions, however they are not dynamically correlated to the speech (so are a weaker signal) and may be common among different people. For example it may be increasingly harder to separate speech

among individuals with similar voice timbre or appearance. Recent works have attempted to deal with the inadequacy of conditioning on a single source by either jointly conditioning on more than one modality using naive fusion of static cues with the lip movement [3] or by learning the separation task jointly with a cross-modal prior [21].

However to-date there is no unified framework for: (i) conditioning on *asynchronous* information (such as a delay between the audio and visual streams); or (ii) for seamlessly conditioning (and fusing) on multiple sources of information or on different types of modalities; or (iii) for using a large temporal context so that predictions can be made using a language model.

Our first contribution is to enable conditioning on *asynchronous* visual (lip) streams. Most previous work relies on costly pre-processing steps to synchronise the audio and video streams, and their performance deteriorates in real world situations where out of sync data is a regular occurrence due to transmission delays, jitter or technical issues. We show that in our work there are no detrimental effects with timing delays of 5 frames (200 ms) or more. Furthermore, the audio and visual streams do not even have to have the same temporal sampling rate.

Our second contribution is to enable enhancement by conditioning on *textual* input. This new functionality allows speech to be enhanced without requiring biometric information or even a visual stream. It is applicable where the textual content of the speech is known in advance, e.g. from a prepared speech or lyrics of a song, or where Automatic Speech Recognition (ASR) or lip reading [5, 10] can be used to transcribe what is said, even imprecisely, and then subsequently used to isolate the speaker from background noise. Textual conditioning is asynchronous, as only the order of words (or more precisely the phonemes) is required, but not their precise temporal alignment.

Both of these contributions are facilitated by a new Transformer-based speech separation and enhancement network, where we use the positional encoding of the Transformer to record the timestamp (of the audio and visual samples) or the ordering (of the words in the text) of the conditioning signal. The network operates directly on the waveform level, without requiring spectrograms as an intermediate step of the audio processing. It uses a U-Net [12, 38] architecture to encode noisy audio and then decode it into clean speech, with the Transformer as the network's bottleneck, where the conditioning information can be visual and/or text. The Transformer also enables modelling a longer temporal context (*e.g.* compared to an LSTM) allowing the network to explicitly model structure in natural language. By having the ability to anticipate what follows a certain sequence of words the model can then better approximate the target speech output.

In summary, we make the following contributions: First

we design a modern multi-modal speech enhancement architecture, VoiceFormer, that uses a Transformer-based bottleneck to fuse heterogeneous modality streams, meaning that it can simultaneously condition on multiple non-aligned modalities. Second, we introduce text-conditioned speech enhancement as a novel multi-modal task and show that our proposed architecture is well designed to handle it. Third, we demonstrate that our trained models are robust to audio-visual synchronisation offsets. Fourth, we exhibit state-of-the-art performance, surpassing other audio-only and audio-visual baselines in the tasks of speaker separation and speech enhancement.

2. Related work

Our work is related to a large body of previous works which ranges from traditional audio-based speech enhancement to methods for multi-modal speech and sound-source separation.

Audio-based speech separation and enhancement. Speech enhancement is a well studied problem with long history in audio processing. Audio-only methods are by design person-agnostic; they often work well for enhancement, retaining speech and filtering out background noise, but struggle with speaker separation. Recent methods aim to tackle this by solving the label permutation problem – assigning audio sources to their corresponding speakers in the audio [23, 48, 54]. Wang *et al.* [50] proposed a method to localize individual speakers and train an enhancement network on spatial as well as spectral features. Lou *et al.* [28] introduced a deep learning framework for speech separation that addresses label permutation and does not require knowledge of the number of speakers. Yu *et al.* [54] devised a deep learning training criterion for solving the label permutation problem. Chen *et al.* [8] perform separation based on clustering of the audio sources in the embedding space. Defossez *et al.* [12] recently proposed the Denoiser, a real-time speech enhancement network that is trained end-to-end on raw waveforms.

Multi-modal methods based on static cues. Various recent methods attempt to solve the audio separation problem based on external cues that contain information about the sound source. Examples of such methods are works that perform identity-specific speech separation using voice [49], or face [11] identity embeddings. Related works that also fall in this category are various audio-visual methods for separating the sounds of musical instruments based on stationary appearance cues [15, 19, 20, 33, 39, 45, 46, 53, 57].

Multi-modal methods based on dynamic cues. Another family of methods solve the source separation task based on dynamic cues that exhibit some variation over time. Such cues are more commonly contained in a synchronised visual stream, therefore these methods are in their

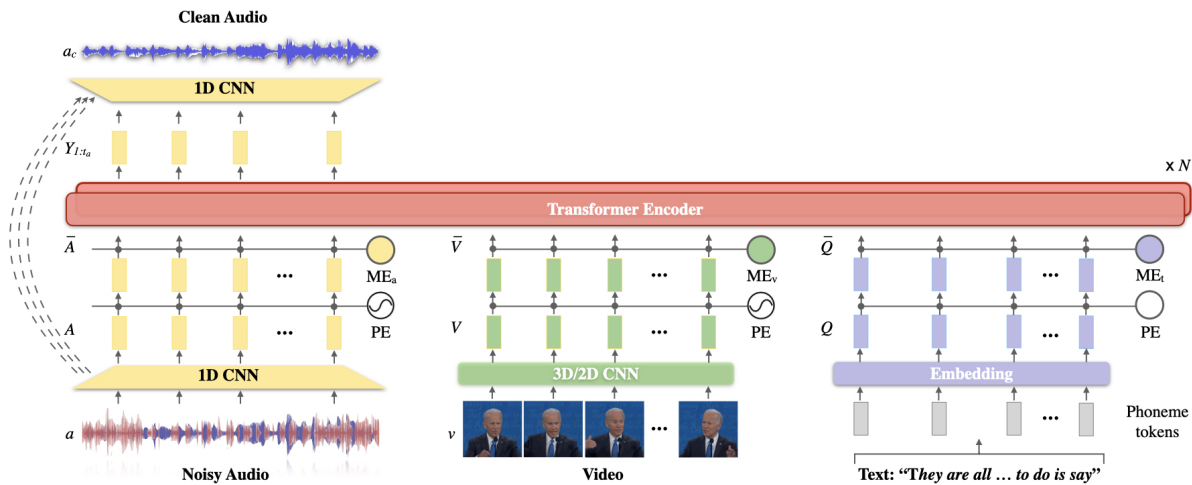


Figure 2. Overview of the proposed multi-modal speech enhancement with transformers (mmset) architecture. it consists of a u-net style encoder-decoder for the audio stream, with the bottleneck layers conditioned on a transformer that can ingest textual and visual modalities. the u-net encoder ingests the raw audio waveform of the target speaker with noise (background or other speakers), and produces a sequence of audio embeddings. the multi-layer transformer conditions on the audio embeddings, the phoneme sequence extracted from the text being spoken, and/or the visual embeddings from the video of the target speaker. the u-net decoder inputs the sequence of refined audio embeddings from the output of the transformer, and produces the clean audio waveform of the target speaker (with the noise removed). in both training and inference the conditioning can include video or text or both.

majority audio-visual. For example, utilising visual features has proven to be very beneficial in separating speakers in audio clips where the corresponding video is accessible. Indeed, recent works [1, 4, 14, 16, 17] conditioned deep learning frameworks on the lip movements of target speakers in order to isolate their voice among multiple other speech signals. Wu *et al.* [51] presented an audio-visual speech separation network which operates in the time-domain (raw waveforms) instead of frequency-domain, while Sadeghi *et al.* solved the task by employing audio-visual VAEs [40]. Owens *et al.* [31] use spatio-temporal video features for separating on-screen from off-screen speech. In a similar line of work, recent methods have proposed using motion cues in videos [18, 32, 56] in order to separate musical instruments belonging to the same class, thus overcoming the limitations of static appearance features. A few recent works investigated combining static and dynamic cues to improve speech separation. For example Gao *et al.* [21] investigated training an audio-visual speaker separation network jointly with voice-face embeddings that provide a prior to aid the separation process. Another proposed direction [3, 22] is conditioning a separation network on both lip-movements and an embedding of the target speaker’s voice to improve robustness to visual occlusions. However none of those works propose a unified framework for conditioning on multiple non-aligned dynamic sources of information.

Multi-modal fusion. Our method is more broadly related to works that fuse different modalities to solve multi-modal tasks, such as audio-visual fusion using Transformers [24, 30] for audio-visual event detection [26, 27, 44] or audio-visual synchronization [7], video-text fusion for vi-

sual grounding [52, 55] and visual keyword spotting [36].

3. Method

In this section we describe our proposed method for multi-modal speech enhancement which we call the VoiceFormer. Given a noisy speech signal, the goal is to separate a target speech component corresponding to other input modalities (text or video), and to filter out the rest of the signal (other speakers or background noise). An overview of the architecture is shown in Figure 2. We use a U-Net style audio encoder-decoder (similar to [12]), with a multi-modal Transformer in its bottleneck, where the noisy audio is fused with the conditioning inputs (video and text). The rest of this section describes the individual components and the training of the model. We refer the reader to the project webpage for full architectural details.

3.1. Architecture

Audio, Visual, and Text representations.

The model has three input streams: one ingesting an audio waveform $a \in \mathbb{R}^{T_a}$, one the corresponding video input $v \in \mathbb{R}^{3 \times T_v \times H \times W}$, and one a textual representation $s = (s_1, s_2 \dots, s_{n_s})$ of the sentence being uttered. Similar to [13], we extract a representation of the noisy audio, $A \in \mathbb{R}^{t_a \times c}$, directly from the input waveform a using the encoder part of a U-Net, which consists of 1D convolution layers. We also use a spatiotemporal CNN [1] to obtain visual representations, $V \in \mathbb{R}^{t_v \times c}$. The textual representation is obtained using the phonimzer library [6] with espeakng as its backend; the words in the input sentence are first mapped to a phonetic sequence of length t_q based on the International Phonetic Alphabet, and are then mapped to a

sequence of learnable embedding vectors $Q \in \mathbb{R}^{t_q \times c}$.

Transformer bottleneck. In order to inform the model of the temporal order of its inputs, *i.e.* signal timestamps for the video/audio features and phoneme ordering for the text, we add positional encodings, $PE_{\{a,v,q\}} \in \mathbb{R}^{t_{\{a,v,q\}} \times c}$. PE_a and PE_v are implemented as sinusoidal vectors and PE_q as learnable embedding vectors. Moreover, in order to allow the model to distinguish which signal comes from which modality, we also add modality encodings, $ME_{\{a,v,q\}} \in \mathbb{R}^c$, which are three learnable vectors, one for each modality type. In summary, the order and modality aware uni-modal representations are calculated as

$$\bar{A} = A + PE_a + ME_a, \quad (1)$$

$$\bar{V} = V + PE_v + ME_v, \quad (2)$$

$$\bar{Q} = Q + PE_q + ME_q, \quad (3)$$

These are concatenated along the temporal dimension

$$Z = (\bar{A}; \bar{V}; \bar{Q}) \in \mathbb{R}^{(t_a+t_v+t_q) \times c} \quad (4)$$

and the resulting feature vector is processed with a Transformer encoder with N layers and h heads:

$$Y = \text{TRANSFORMER-ENCODER}(Z).$$

The Transformer bottleneck fuses the three inputs together, allowing for full cross-attention between all modality combinations. In particular, the textual and video evidence is attended upon and used to extract only the relevant parts of the embedded audio. We note that neither explicit alignment nor common frame-rates between the different signals are required. The output of the Transformer corresponding to its audio input, $Y_{1:t_a}$, contains the representation of the separated/enhanced audio (the outputs corresponding to video and text are discarded).

U-Net Decoder. The enhanced audio representation is decoded into a waveform \hat{a}_c with the decoder part of the U-Net, which is comprised of a stack of transposed 1D convolutions, including shortcut connections from the audio encoder. The resulting output \hat{a}_c is an enhanced waveform containing only the speech corresponding to the visual and text input.

3.2. Training objective

Given a training dataset \mathcal{D} of tuples (a, v, s, a_c) of noisy audio waveforms and corresponding video, text and clean target waveforms, we train the model using an L1 loss between the predicted enhanced and the target clean waveforms:

$$\mathcal{L} = \mathbb{E}_{(a,v,s,a_c) \in \mathcal{D}} \|a_c - \hat{a}_c\|_1 \quad (5)$$

4. Experiments

4.1. Synthetic sequences

Following previous work [1, 14] we train and evaluate our models by creating synthetic noisy samples by adding

the waveforms of two separate clips and requiring the model to reconstruct the individual waveforms in the output, after conditioning on the corresponding video/text input. In particular, one of the clips always contains clean speech of a single speaker, while the interfering audio might be either speech from another speaker in the speaker separation experiments, or a noise audio clip, simulating background noise for the speech enhancement experiments. Note that although we train and evaluate the model on synthetic audio mixtures it is applicable to real noisy sequences as the domain gap between synthetic and real samples is small.

4.2. Datasets, training & evaluation protocol

Data. We obtain audio-visual speech samples from the LRS2 [9] and LRS3 [2] lip reading datasets. LRS2 contains broadcast footage from British television while LRS3 has been created from TED and TEDx clips downloaded from YouTube. Both datasets contain audio-visual tracks of tightly cropped talking heads, engaging in continuous speech. All tracks are accompanied by text transcriptions of the utterances that are well aligned to the video and audio, which have also been automatically synchronised. On examining the datasets we determined that some of the samples include two speakers while others included background noise such as clapping, bird chiming, music or crowd laughter. These samples were removed from the dataset such that every sample contains speech from only one speaker without other background noise. A combinations of diarization and background noise detection methods were employed to detect and remove the noisy samples. As a result, 57 hours out of 197 hours was retained from the LRS2, and 439 hours out of 440 hours was kept from LRS3.

Moreover, for obtaining noise for our denoising experiments we follow [12] and use a subset of the DNS [35] dataset, which contains approximately 181 hours of noise audio from a wide variety of events. These samples were used as background noise to construct synthetic noisy audio waveforms during training and evaluation.

Training sequences. Approximately 23,000 samples from the LRS2 dataset and over 100,000 samples from the LRS3 dataset were used at training time. The samples were mixed together randomly at each training pass leading to the generation of numerous new and unseen examples. The speech signals are mixed together in sequences of 4 seconds for the version of the model trained on audio and video input. The starting point of each sequence was randomly chosen as a data augmentation method. The length of the sequences for the models that have text as their input are dictated by the number of words in the text that correspond to the associated audio and videos, ranging between 1 to 6 seconds. The samples with similar length were batched together at training time to avoid padding the sequences with zeros whenever possible. Each audio track was inde-

pendently normalised before mixing them together to create synthetic mixtures or feeding them into the network.

Evaluation sequences. We evaluate on synthetic mixtures of two speakers, or a mixture of one speaker and a noise sample. To distinguish between these two related tasks we refer to the first one as speaker separation and the second as denoising.

We created separate test sets from LRS2 and LRS3 with 2515 and 3229 samples from each dataset respectively. These test sets were used to evaluate the aforementioned tasks, comparing our model with baselines and for performing model ablations and robustness tests. The samples in all the test sets contain noisy audio, video and text. The duration of the samples vary as they are cropped based on the length of the corresponding text. The text samples are 9 words long and do not necessarily form a sentence. We include qualitative examples on real sequences on the project webpage.

Evaluation metrics. For evaluating our methods and baselines we use standard speech enhancement metrics, including Signal-to-Distortion-Ratio (SDR) [34, 47], a common blind source separation criterion, measuring the ratio between the energy of the target signal and of the errors contained in the separated output, Short-Time Objective Intelligibility (STOI) [43], which measures the intelligibility of the signal, and the Perceptual Evaluation of Speech Quality (PESQ) [37], that rates the overall perception quality of the output signal.

4.3. Implementation details

Our network is implemented and trained in Pytorch. The faces in the video recordings are cropped and resembled into 25 FPS. The audio input is converted into mono by taking the mean of both channels and resampled to have a rate of 16kHz, and the signals are up sampled by 3.2 times before feeding into the network. The audio output of the model is down-sampled by the same ratio. For the audio U-Net we use the Denoiser implementation of [12] without any changes to the architecture. For the visual backbone we follow [36] and use a 3D/2D residual CNN [42], pre-trained on word-level lip reading task. To speed up training, the backbone is frozen and visual features are pre-extracted and saved on a hard drive. For the Transformer encoder we use $N = 3$ layers and $h = 8$ heads, with a model size 532. The embedding dimensions across all the modalities is set to 768 to match the channel dimension of the audio features from the output of the U-Net encoder layer.

We obtain models that condition on a single modality (e.g. only video or only text), by simply not including the corresponding input in Eq. 4.

Training started with a network that includes an LSTM on synchronised audio and corresponding visual sequences as a pre-training stage for the Transformer model. For train-

ing the transformer the learning rate was set to $5 \cdot 10^{-5}$. In all cases the Adam optimizer was used with weight decay of 0.0001, batch size of 64 and the learning rate was reduced linearly after each epoch on plateau. The models are trained on the LRS2 and LRS3 datasets separately starting with the LRS2 mixtures. The training curriculum for the speech enhancement models started with mixtures of 2 speakers before training on one speaker and background noise. When experimenting with more than two transformer layers, all the trainable parameters in the model are frozen apart from the additional encoder layers, which helped to stabilize and accelerate the training.

4.4. Results

In this section we give a detailed evaluation of the proposed method, including robustness analyses, ablations and comparison to baselines. We first compare the performance of our models when they are conditioned on different input modality combinations; we then perform robustness tests in settings where parts of the modalities are missing as well as when there is a misalignment between video and audio; we finally compare to the state-of-the-art on the speaker separation and speech enhancement tasks.

Modalities comparison. In order to assess the effect of using different modalities as conditioning input, we compare models using only video ($A+V$), only text ($A+T$) or both ($A+V+T$) in Table 1. We observe that the text-only model successfully separates the speakers, obtaining reasonable performance. This result is evidence that (a) indeed text can be used to separate speech in cocktail party scenarios, and (b) that our proposed architecture is flexible enough to capture information from different conditioning sources without any changes and can successfully solve the novel text-conditioned speaker separation task without any help from other modalities. Observing the performance of the video-conditioned models shows that video obtains a stronger performance than text. This is an interesting finding, suggesting that lip movements are stronger cues for the separation task, presumably because they carry more information than the language content of the utterance (e.g. speaker mood, accent etc). We refer to the project webpage for qualitative separation results using the different models.

Cross-modal attention. To assess the effectiveness of our design for cross-modal attention through a concatenation of the modalities, we examine the attention in the transformer bottleneck. The attention maps in Figure 3 reveal the correspondence between the audio tokens and the other modalities. attend to the features in the corresponding modalities. This indicates that the model is able to implicitly learn the audio-visual and audio-text alignments and confirms our intuition that it does not require manual alignment or common temporal rates.

| Model | LRS2 | | | LRS3 | | |
|-------|----------------|-----------------|-----------------|----------------|-----------------|-----------------|
| | SDR \uparrow | STOI \uparrow | PESQ \uparrow | SDR \uparrow | STOI \uparrow | PESQ \uparrow |
| A+T | 13.1 | 89.7 | 2.16 | 14.1 | 91.4 | 2.37 |
| A+V | 14.1 | 91.3 | 2.36 | 15.5 | 93.4 | 2.62 |
| A+V+T | 14.2 | 91.7 | 2.41 | 15.5 | 93.5 | 2.63 |

Table 1. **Speaker separation performance using different modalities.** We compare VoiceFormer models conditioning on different modality combinations. We observe that the text-only model (A+T) performs reasonably well for this task, although the performance is lower compared to when conditioning on lip movements from the visual stream. The full A+V+T model that conditions on both text and video obtains only a slight improvement over the A+V model. \uparrow denotes higher is better

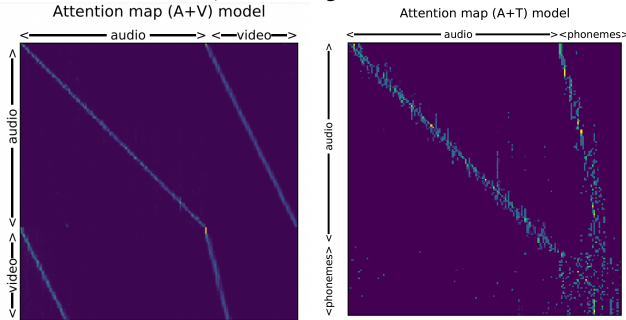


Figure 3. **Attention map visualisations of the first Transformer layer.** The visualisations show the average score of the attention heads in the first multi-head attention layer of the transformer. Brighter colours indicate higher scores and brighter pixels on the same row indicate correspondence between modalities. Left: audio and video correlation; and Right: audio and text correlation. Higher scores are given to the audio token and its corresponding token in the other modality at each timestep. This indicates that the model is able to elegantly fuse the mixed/noisy audio stream with the conditioning vectors from different modalities, without the need for explicit alignment between the signals, or requiring them to be operating at the same temporal rate.

Architecture component contribution. To analyse the contribution of our architecture components on the performance of the model, we examine various architecture configurations. The results are shown in Table 2

Robustness to missing information. As can be seen from the experiments presented so far, although text can be sufficient for performing speech separation, it provides a very limited performance boost when strong visual evidence (*i.e.* clear, unoccluded lip movements) is already available. However, we emphasise that adding text in this setting improves the robustness of the model against missing information. We therefore conduct further experiments where we evaluate the same models, but artificially limit the amount of information from one of the sources. For simulating missing video information we mask out (by zeroing) the visual features for a percentage of video frames. We

| Model | input | SDR \uparrow | STOI \uparrow | PESQ \uparrow |
|--------------------------|-------|----------------|-----------------|-----------------|
| LSTM (A+V) [4] | S | 9.25 | 84.0 | 1.91 |
| UNet + LSTM (A+V) | W | 12.8 | 89.9 | 2.17 |
| UNet + Transformer (A+V) | W | 14.1 | 91.3 | 2.36 |

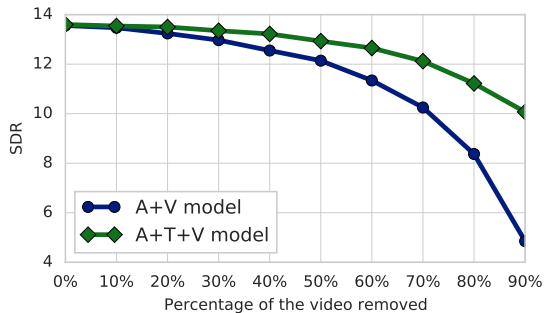
Table 2. **Performance based on various architecture configurations.** We observe that the UNet module, which digests audio as a raw waveform, improves the SDR by 3.55dB in comparison to the spectrograms based baseline model [4]. Replacing the LSTM with a transformer bottleneck further increases the SDR by 1.3dB. Note that the UNet removes the necessity to calculate spectrograms and predict phase, while the transformer offers robustness to audio visual misalignment. \uparrow denotes higher is better. S indicates the audio input is in mel-spectrogram form and W indicates raw waveform is used.

show the results of this experiments in Figure 4a. It is clear that the A+T+V model that conditions on text in addition to video is much more robust to distortions in the video input.

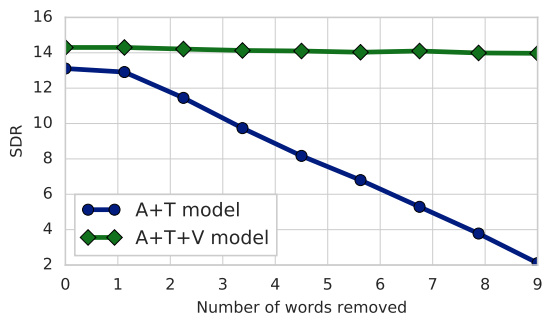
Similarly, for simulating missing text information, we remove a variable number of words from the text input. The performance against the number of removed words is shown in Figure 4b. We make two observations. First, the performance of the A+T model sharply deteriorates when an increasing number of words are removed from the text input. This indicates that the model relies on the text input to perform the separation of the whole utterance, rather than for just disambiguation between two separated streams. Second, that the A+V+T model is, as expected, very robust to the missing text information.

Robustness to inconsistent modalities. To further probe into our models and gain a better understanding of them, we perform a series of experiments where we replace the conditioning video or text input with input from a different, irrelevant video clip. The results are shown in Table 3. We observe that the models that condition on a single source (A+T, A+V) completely fail when presented with wrong inputs. This is expected as the visual/textual evidence in those cases is not consistent with any speech components contained in the audio. From the similar analysis on the A+V+T model we observe that, although the model performs a lot worse when supplied with inconsistent video, but consistent text, input, it does not completely fail (*e.g.* obtains 5.1 SDR). On the other hand, the performance drops only marginally when the model is presented with inconsistent text but consistent video input.

To sum up our analysis of the behaviour of VoiceFormer, we conclude that the A+V+T model provides good robustness to disruptions in the video inputs, which comes with virtually no risk; even if for some reason the textual input provided is missing or inconsistent with the video (*e.g.* if we use noisy ASR approximations), the model still performs on par with the video-only model.



(a) **Robustness to missing video information.** Conditioning on text in combination with video clearly improves the robustness to missing video information (e.g. due to occlusions), compared to only using video.



(b) **Robustness to missing text information.** The performance of the text-only conditioned model quickly deteriorates when words are removed from the text input. This demonstrates that the separation is indeed reliant on the text content throughout the utterance, and not just for disambiguation between two separated signals. As expected, combining video and text provides robustness to missing text input.

Figure 4. Experiments with missing information.

Bottleneck ablation and robustness to AV misalignment. To assess our choice of the Transformer encoder in the U-Net bottleneck we conduct an experiment replacing the Transformer by an 2 layer LSTM (similar to the architecture used by [12]), where the audio and video features are concatenated in the channel dimension.

We argue that using a Transformer to fuse the different modalities offers the advantage of not requiring synchronised input streams (e.g. video and audio). To verify this hypothesis, we experiment with artificially shifting the audio input by a random offset within the range from -200 to 200 ms. We train and evaluate both the LSTM and Transformer models under those conditions. The results of this comparison are shown in Figure 5. It is clear that the performance of the LSTM baseline steeply deteriorates when the audio and video inputs are not properly synchronised. VoiceFormer on the other hand is very robust to synchronisation issues, retaining high SDR scores (> 12), even when the two modalities are offset by up to 200ms.

Comparison to the state-of-the-art. We report our method’s performance on speaker separation and compare

| Modality | Text | Video | SDR \uparrow | STOI \uparrow | PESQ \uparrow |
|----------|------|-------|----------------|-----------------|-----------------|
| A+T | ✓ | ✗ | 13.1 | 89.7 | 2.16 |
| A+T | ↷ | ✗ | 1.18 | 60.5 | 1.53 |
| A+V | ✗ | ✓ | 14.1 | 91.3 | 2.36 |
| A+V | ✗ | ↷ | -1.51 | 51.9 | 1.39 |
| A+V+T | ✓ | ✓ | 14.2 | 91.7 | 2.41 |
| A+V+T | ✓ | ↷ | 5.11 | 70.5 | 1.69 |
| A+V+T | ✓ | ✗ | 10.9 | 83.6 | 2.12 |
| A+V+T | ↷ | ✓ | 14.1 | 91.3 | 2.36 |
| A+V+T | ✗ | ✓ | 14.1 | 91.2 | 2.34 |

Table 3. **Missing or inconsistent modalities.** ✓ indicates that the correct signal for the corresponding modality is in input, ↷ that the input for this modality is supplied from a different sentence or video, and ✗ that this input is not provided. The results are reported for the speaker separation task on the LRS2 test set. We observe that the models that use a single modality completely fail to solve the task when the conditioning input is wrong. On the other hand the A+V+T model that uses both video and text is fully robust to inconsistent text inputs, and can even partially recover from inconsistent video inputs. \uparrow denotes higher is better

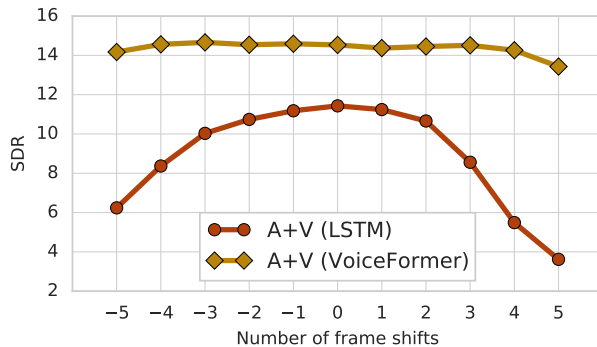


Figure 5. **Robustness to audio-visual misalignment.** We compare our proposed model with a baseline using an LSTM bottleneck in the audio-visual speaker separation setting. It is clear that while the LSTM baseline struggles when the video and audio streams are misaligned, VoiceFormer is robust to synchronisation offsets. A five frame offset corresponds to 200 ms.

it to previous methods in Table 4. As baselines we use the state-of-the-art speech enhancement method of [12] (audio-only), as well as the recently proposed audio-visual methods [1, 4, 21]. It is clear that VoiceFormer outperforms the previous works on all metrics, obtaining state-of-the-art performance in the (comparable) audio-visual setting A+V.

Comparison on speech enhancement (denoising). For completeness, we also assess our model’s performance on the denoising task. We show the results in Table 5. Our proposed models perform on par with the state-of-the-art Denoiser model [12]. We note that this is expected as the speech enhancement task is easier than speaker separation and can be solved by using only the audio modality. In-

| Model | V | T | LRS2 | | | LRS3 | | |
|-------------------------------------|---|---|-------------|-------------|-------------|-------------|-------------|-------------|
| | | | SDR | STOI | PESQ | SDR | STOI | PESQ |
| Noisy input | ✗ | ✗ | 0.2 | 66.6 | 1.17 | 1.3 | 69.7 | 1.3 |
| Denoyer [12] | ✗ | ✗ | 0.2 | 66.6 | 1.17 | 1.3 | 69.7 | 1.3 |
| AVObjects [4] | ✓ | ✗ | 8.86 | 83.9 | 1.94 | 9.72 | 85.1 | 2.02 |
| Conversation [1] | ✓ | ✗ | 9.25 | 84.0 | 1.91 | 10.15 | 86.5 | 2.08 |
| VisualVoice [21] [†] | ✓ | ✗ | 10.8 | 88.4 | 2.16 | 11.7 | 90.0 | 2.41 |
| Ours A+V | ✓ | ✗ | <u>14.1</u> | <u>91.3</u> | <u>2.36</u> | 15.5 | 93.4 | 2.62 |
| Lee <i>et al.</i> [25] [◦] | ✓ | ✗ | 10.01 | 88.0 | 0.94 | 9.78 | 85.0 | 0.710 |
| Ours A+V [◦] | ✓ | ✗ | 12.71 | 92.0 | 2.06 | 14.74 | 94.0 | 2.42 |
| Ours A+V+T | ✓ | ✓ | 14.2 | 91.7 | 2.41 | 15.5 | 93.5 | 2.63 |

Table 4. **Comparison to the state-of-the-art on the speaker separation task.** We evaluate our best models on the synthetic LRS2 and LRS3 test sets. The *V* and *T* columns denote which modalities are used for conditioning by each model. *A+V+T* indicates our full model and *A+V* for a version conditioning only on video and not on text. Our *A+V* model clearly outperforms the previous work on all the measures, obtaining state-of-the-art performance in this setting. We note that the state-of-the-art speech enhancement model of [12] cannot deal with mixture of different speakers and outputs the mixed signal, even after we attempted to fine-tune it on this task. [†]not fine-tuned on the synthetic two-speaker LRS2 training set. [◦] The comparison is on a different test set published by Lee *et al.* [25]. Higher is better for all metrics.

deed the high numbers obtained on all metrics indicate that the performance in this setting is potentially saturated. This experiment was included to demonstrate that denoising can also be well handled by the proposed method. We leave stress-testing of our models on more challenging enhancement settings to future work.

| Model | SDR [↑] | STOI [↑] | PESQ [↑] |
|--------------|------------------|-------------------|-------------------|
| Denoyer [12] | 16.49 | 88.4 | 2.44 |
| Ours A+V | <u>16.13</u> | 88.9 | 2.49 |
| Ours A+T | 15.8 | 88.4 | 2.38 |
| Ours A+V+T | 16.0 | 88.8 | 2.42 |

Table 5. **Model performance and comparisons on the speech enhancement (denoising) task.** Results are reported on the LRS2 test set. All our models obtain good performance, matching the state-of-the-art audio-only [13] denoising method. [↑] denotes higher is better

Qualitative examples. We strongly encourage the reader to refer to the project webpage for multiple examples of speaker separation and denoising on real videos, which cover many of the above scenarios. We demonstrate that our model can perform speech separation in challenging real world scenarios.

5. Discussion

5.1. Limitations

The main limitation of the proposed method is the strong assumption that the textual content of the target spoken utterance is available as input to the model during inference. As we have argued in the introduction there are many practical uses where this is a valid assumption (*e.g.* a prepared conference speech or song lyrics). In our qualitative examples we also show that it is possible to use related technologies such as ASR to obtain approximate transcriptions which can be use in our method for target speaker separation.

5.2. Societal impact

The development of strong multi-modal speech enhancement models opens up exciting opportunities for useful applications, as noted in the introduction. The new method can be used without requiring synchronisation while achieving improved performance. It also provides the opportunity for novel applications, for example it can use subtitles in films to suppress all background music. However there are possible malign uses of providing a new method to isolate a speaker from others in terms of surveillance. In the novel setting that we consider in this paper however, the natural language content is presumed to be already known (or can be obtained beforehand through other means). We therefore argue that this concern does not apply in our setting and that overall the potential benefits from benevolent uses (*e.g.* medical applications and smart hearing aids) outweigh the limited risks.

5.3. Conclusion

We have presented a multi-modal speech enhancement method that can condition on multiple non-aligned modalities. We also introduced text-conditioned speech enhancement as a new task and showed how our proposed architecture can efficiently solve it. Our trained models demonstrated state-of-the art performance in a variety of settings as well as robustness to synchronisation issues and missing information. In future work we will consider extending our framework by adding new types of embeddings to condition on modalities such as (i) a particular person (expanding on previous work in [21])(ii) the language spoken (*e.g.* to pick a French speaker out of English ones).

Acknowledgements. This work is funded by the UK EPSRC AIMS CDT, the EPSRC Programme Grant VisualAI EP/T028572/1, and a Google-DeepMind Graduate Scholarship.

References

- [1] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. The conversation: Deep audio-visual speech enhancement. In *INTERSPEECH*, 2018. 1, 3, 4, 7, 8
- [2] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. LRS3-TED: a large-scale dataset for visual speech recognition. In *arXiv preprint arXiv:1809.00496*, 2018. 4
- [3] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. My lips are concealed: Audio-visual speech enhancement through obstructions. In *INTERSPEECH*, 2019. 2, 3
- [4] Triantafyllos Afouras, Andrew Owens, Joon Son Chung, and Andrew Zisserman. Self-supervised learning of audio-visual objects from video. In *Proc. ECCV*, 2020. 3, 6, 7, 8
- [5] Yannis M. Assael, Brendan Shillingford, Shimon Whiteson, and Nando de Freitas. Lipnet: Sentence-level lipreading. *arXiv:1611.01599*, 2016. 2
- [6] Mathieu Bernard and Hadrien Titeux. Phonemizer: Text to phones transcription for multiple languages in python. *Journal of Open Source Software*, 6:3958, 12 2021. 3
- [7] Honglie Chen, Weidi Xie, Triantafyllos Afouras, Arsha Nagrani, Andrea Vedaldi, and Andrew Zisserman. Audio-visual synchronisation in the wild. In *Proc. BMVC*, 2021. 3
- [8] Zhuo Chen, Yi Luo, and Nima Mesgarani. Deep attractor network for single-microphone speaker separation. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 246–250, 2017. 2
- [9] Joon Son Chung, Andrew Senior, Oriol Vinyals, and Andrew Zisserman. Lip reading sentences in the wild. In *Proc. CVPR*, 2017. 4
- [10] Joon Son Chung and Andrew Zisserman. Lip reading in the wild. In *Proc. ACCV*, 2016. 2
- [11] Soo-Whan Chung, Soyeon Choe, Joon Son Chung, and Hong-Goo Kang. Facefilter: Audio-visual speech separation using still images. *arXiv preprint arXiv:2005.07074*, 2020. 1, 2
- [12] Alexandre Defossez, Gabriel Synnaeve, and Yossi Adi. Real time speech enhancement in the waveform domain. In *Interspeech*, 2020. 2, 3, 4, 5, 7, 8
- [13] Alexandre Défossez, Nicolas Usunier, Léon Bottou, and Francis R. Bach. Music Source Separation in the Waveform Domain. *CoRR*, abs/1911.13254, 2019. 3, 8
- [14] Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinatan Hassidim, William Freeman, and Michael Rubinstein. Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation. *arXiv preprint arXiv:1804.03619*, 2018. 1, 3, 4
- [15] John W Fisher III, Trevor Darrell, William T Freeman, and Paul A Viola. Learning joint statistical models for audio-visual fusion and segregation. In *NeurIPS*, 2000. 2
- [16] Aviv Gabbay, Ariel Ephrat, Tavi Halperin, and Shmuel Peleg. Seeing through noise: Visually driven speaker separation and enhancement. In *Proc. ICASSP*, pages 3051–3055. IEEE, 2018. 3
- [17] Aviv Gabbay, Asaph Shamir, and Shmuel Peleg. Visual Speech Enhancement using Noise-Invariant Training. *arXiv preprint arXiv:1711.08789*, 2017. 3
- [18] Chuang Gan, Deng Huang, Hang Zhao, Joshua B. Tenenbaum, and Antonio Torralba. Music gesture for visual sound separation. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10475–10484, 2020. 3
- [19] Ruohan Gao, Rogério Schmidt Feris, and Kristen Grauman. Learning to separate object sounds by watching unlabeled video. *CoRR*, abs/1804.01665, 2018. 2
- [20] Ruohan Gao and Kristen Grauman. Co-separating sounds of visual objects. *arXiv preprint arXiv:1904.07750*, 2019. 2
- [21] Ruohan Gao and Kristen Grauman. VisualVoice: Audio-Visual Speech Separation with Cross-Modal Consistency. In *Proc. CVPR*, 2021. 1, 2, 3, 7, 8
- [22] Rongzhi Gu, Shi-Xiong Zhang, Yong Xu, Lianwu Chen, Yuexian Zou, and Dong Yu. Multi-modal multi-channel target speech separation. *IEEE Journal of Selected Topics in Signal Processing*, 14(3):530–541, 2020. 3
- [23] John R. Hershey, Zhuo Chen, Jonathan Le Roux, and Shinji Watanabe. Deep clustering: Discriminative embeddings for segmentation and separation. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 31–35, 2016. 2
- [24] Andrew Jaegle, Felix Gimeno, Andrew Brock, Andrew Zisserman, Oriol Vinyals, and Joao Carreira. Perceiver: General perception with iterative attention. *arXiv preprint arXiv:2103.03206*, 2021. 3
- [25] Jiyoung Lee, Soo-Whan Chung, Sunok Kim, Hong-Goo Kang, and Kwanghoon Sohn. Looking into your speech: Learning cross-modal affinity for audio-visual speech separation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 8
- [26] Sangho Lee, Youngjae Yu, Gunhee Kim, Thomas Breuel, Jan Kautz, and Yale Song. Parameter efficient multimodal transformers for video representation learning. In *Proc. ICLR*, 2021. 3
- [27] Yan-Bo Lin and Yu-Chiang Frank Wang. Audiovisual transformer with instance attention for audio-visual event localization. In *Proc. ACCV*, 2020. 3
- [28] Yi Luo, Zhuo Chen, and Nima Mesgarani. Speaker-independent speech separation with deep attractor network. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 26(4):787–796, 2018. 2
- [29] Neville Moray. Attention in dichotic listening: Affective cues and the influence of instructions. *Quarterly Journal of Experimental Psychology*, 11(1):56–60, 1959. 1
- [30] Arsha Nagrani, Shan Yang, Anurag Arnab, Aren Jansen, Cordelia Schmid, and Chen Sun. Attention bottlenecks for multimodal fusion. *NeurIPS*, 2021. 3
- [31] Andrew Owens and Alexei A. Efros. Audio-visual scene analysis with self-supervised multisensory features. *Proc. ECCV*, 2018. 1, 3
- [32] Sanjeel Parekh, Slim Essid, Alexey Ozerov, Ngoc Q. K. Duong, Patrick Pérez, and Gaël Richard. Motion informed audio source separation. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6–10, 2017. 3

- [33] Jie Pu, Yannis Panagakis, Stavros Petridis, and Maja Pantic. Audio-visual object localization and separation using low-rank and sparsity. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2901–2905, 2017. 2
- [34] Colin Raffel, Brian Mcfee, Eric Humphrey, Justin Salamon, Oriol Nieto, Dawen Liang, and Daniel Ellis. mir_eval: A transparent implementation of common mir metrics. In *Proceedings - 15th International Society for Music Information Retrieval Conference (ISMIR 2014)*, 10 2014. 5
- [35] Chandan KA Reddy, Harishchandra Dubey, Kazuhito Koishida, Arun Nair, Vishak Gopal, Ross Cutler, Sebastian Braun, Hannes Gamper, Robert Aichner, and Sriram Srinivasan. Interspeech 2021 deep noise suppression challenge. In *INTERSPEECH*, 2021. 4
- [36] Prajwal K Renukanand, Liliane Momeni, Triantafyllos Afouras, and Andrew Zisserman. Visual keyword spotting with attention. In *Proc. BMVC*, 2021. 3, 5
- [37] Antony W Rix, John G Beerends, Michael P Hollier, and Andries P Hekstra. Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs. In *Proc. ICASSP*, volume 2, pages 749–752. IEEE, 2001. 5
- [38] 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
- [39] Andrew Rouditchenko, Hang Zhao, Chuang Gan, Josh McDermott, and Antonio Torralba. Self-supervised audio-visual co-segmentation. In *Proc. ICASSP*, pages 2357–2361. IEEE, 2019. 2
- [40] Mostafa Sadeghi, Simon Leglaive, Xavier Alameda-Pineda, Laurent Girin, and Radu Horaud. Audio-visual speech enhancement using conditional variational auto-encoders. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:1788–1800, 2020. 3
- [41] Jean-Luc Schwartz, Frédéric Berthommier, and Christophe Savariaux. Seeing to hear better: evidence for early audio-visual interactions in speech identification. *Cognition*, 93(2):B69–B78, 2004. 1
- [42] Themis Stafylakis and Georgios Tzimiropoulos. Combining residual networks with lstms for lipreading. In *Interspeech*, 2017. 5
- [43] Cees Taal, Richard Hendriks, Richard Heusdens, and Jesper Jensen. An algorithm for intelligibility prediction of time-frequency weighted noisy speech. *IEEE Transactions on Audio, Speech and Language Processing*, 2011. 5
- [44] Yapeng Tian, Dingzeyu Li, and Chenliang Xu. Unified multisensory perception: Weakly-supervised audio-visual video parsing. In *Proc. ECCV*, 2020. 3
- [45] Efthymios Tzinis, Scott Wisdom, Aren Jansen, Shawn Hershey, Tal Remez, Daniel PW Ellis, and John R Hershey. Into the wild with audioscope: Unsupervised audio-visual separation of on-screen sounds. *arXiv preprint arXiv:2011.01143*, 2020. 2
- [46] Efthymios Tzinis, Scott Wisdom, Tal Remez, and John R. Hershey. Improving on-screen sound separation for open domain videos with audio-visual self-attention. *ArXiv*, abs/2106.09669, 2021. 2
- [47] Emmanuel Vincent, Rémi Gribonval, and Cédric Févotte. BSS EVAL toolbox user guide. *IRISA Technical Report 1706*. http://www.irisa.fr/metiss/bss_eval/, 2005. 5
- [48] DeLiang Wang and Jitong Chen. Supervised speech separation based on deep learning: An overview. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 26(10):1702–1726, 2018. 2
- [49] Quan Wang, Hannah Muckenhirn, Kevin Wilson, Prashant Sridhar, Zelin Wu, John Hershey, Rif A Saurous, Ron J Weiss, Ye Jia, and Ignacio Lopez Moreno. Voicefilter: Targeted voice separation by speaker-conditioned spectrogram masking. In *Interspeech*, 2018. 1, 2
- [50] Zhong-Qiu Wang and DeLiang Wang. Combining spectral and spatial features for deep learning based blind speaker separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(2):457–468, 2019. 2
- [51] Jian Wu, Yong Xu, Shi-Xiong Zhang, Lian-Wu Chen, Meng Yu, Lei Xie, and Dong Yu. Time domain audio visual speech separation. In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 667–673, 2019. 3
- [52] Huijuan Xu, Kun He, Bryan A Plummer, Leonid Sigal, Stan Sclaroff, and Kate Saenko. Multilevel language and vision integration for text-to-clip retrieval. In *AAAI*, 2019. 3
- [53] Xudong Xu, Bo Dai, and Dahua Lin. Recursive visual sound separation using minus-plus net. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 882–891, 2019. 2
- [54] Dong Yu, Morten Kolbæk, Zheng-Hua Tan, and Jensen Jensen. Permutation invariant training of deep models for speaker-independent multi-talker speech separation. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 241–245, 2017. 2
- [55] Yitian Yuan, Tao Mei, and Wenwu Zhu. To find where you talk: Temporal sentence localization in video with attention based location regression. In *AAAI*, 2019. 3
- [56] Hang Zhao, Chuang Gan, Wei-Chiu Ma, and Antonio Torralba. The sound of motions. *Proc. ICCV*, 2019. 3
- [57] Hang Zhao, Chuang Gan, Andrew Rouditchenko, Carl Vondrick, Josh McDermott, and Antonio Torralba. The sound of pixels. *arXiv preprint arXiv:1804.03160*, 2018. 2