Looking to Listen at the Cocktail Party: A Speaker-Independent Audio-Visual Model for Speech Separation

ARIEL EPHRAT, Google Research and The Hebrew University of Jerusalem, Israel

INBAR MOSSERI, Google Research

ORAN LANG, Google Research

TALI DEKEL, Google Research

KEVIN WILSON, Google Research

AVINATAN HASSIDIM, Google Research

WILLIAM T. FREEMAN, Google Research

MICHAEL RUBINSTEIN, Google Research

Fig. 1. We present a model for isolating and enhancing the speech of desired speakers in a video. (a) The input is a video (frames + audio track) with one or more people speaking, where the speech of interest is interfered by other speakers and/or background noise. (b) Both audio and visual features are extracted and fed into a joint audio-visual speech separation model. The output is a decomposition of the input audio track into clean speech tracks, one for each person detected in the video (c). This allows us to then compose videos where speech of specific people is enhanced while all other sound is suppressed. Our model was trained using thousands of hours of video segments from our new dataset, AVSPEECH. The “Stand-Up” video (a) is courtesy of Team Coco.

We present a joint audio-visual model for isolating a single speech signal from a mixture of sounds such as other speakers and background noise. Solving this task using only audio as input is extremely challenging and does not provide an association of the separated speech signals with speakers in the video. In this paper, we present a deep network-based model that incorporates both visual and auditory signals to solve this task. The visual features are used to “focus” the audio on desired speakers in a scene and to improve the speech separation quality. To train our joint audio-visual model, we introduce AVSPEECH, a new dataset comprised of thousands of hours of video segments from the Web. We demonstrate the applicability of our method to classic speech separation tasks, as well as real-world scenarios involving heated interviews, noisy bars, and screaming children, only requiring the user to specify the face of the person in the video whose speech they want to isolate. Our method shows clear advantage over state-of-the-art audio-only speech separation in cases of mixed speech. In addition, our model, which is speaker-independent (trained once, applicable to any speaker), produces better results than recent audio-visual speech separation

methods that are speaker-dependent (require training a separate model for each speaker of interest).

Additional Key Words and Phrases: Audio-Visual, Source Separation, Speech Enhancement, Deep Learning, CNN, BLSTM

ACM Reference Format:

Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinatan Hassidim, William T. Freeman, and Michael Rubinstein. 2018. Looking to Listen at the Cocktail Party: A Speaker-Independent Audio-Visual Model for Speech Separation. *ACM Trans. Graph.* 37, 4, Article 109 (August 2018), 11 pages. <https://doi.org/10.1145/3197517.3201357>

1 INTRODUCTION

Humans are remarkably capable of focusing their auditory attention on a single sound source within a noisy environment, while de-emphasizing (“muting”) all other voices and sounds. The way neural systems achieve this feat, which is known as the *cocktail party effect* [Cherry 1953], remains unclear. However, research has shown that viewing a speaker’s face enhances a person’s capacity to resolve perceptual ambiguity in a noisy environment [Golumbic et al. 2013; Ma et al. 2009]. In this paper we achieve this ability computationally.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2018 Copyright held by the owner/author(s).

0730-0301/2018/8-ART109

<https://doi.org/10.1145/3197517.3201357>

The first author performed this work as an intern at Google.

Automatic speech separation—separating an input audio signal into its individual speech sources—is well-studied in the audio processing literature. Since this problem is inherently ill-posed, it requires prior knowledge or special microphone configurations in order to obtain a reasonable solution [McDermott 2009]. In addition, a fundamental problem with audio-only speech separation is the *label permutation problem* [Hershey et al. 2016]: there is no easy way to associate each separated audio source with its corresponding speaker in the video [Hershey et al. 2016; Yu et al. 2017].

In this work, we present a joint audio-visual method for “focusing” audio on a desired speaker in a video. The input video can then be recomposed such that the audio corresponding to specific people is enhanced while all other sound is suppressed (Fig. 1). More specifically, we design and train a neural network-based model that takes the recorded sound mixture, along with tight crops of detected faces in each frame in the video as input, and splits the mixture into separate audio streams for each detected speaker. The model uses visual information both as a means to improve the source separation quality (compared to audio-only results), as well as to associate the separated speech tracks with visible speakers in the video. All that is required from the user is to specify which faces of the people in the video they want to hear the speech from.

To train our model, we collected 100,000 high-quality lectures, TED talks and how-to videos from YouTube, then automatically extracted from these videos roughly 1500 hours of video clips with visible speakers and clean speech with no interfering sounds (Fig. 2). We call our new dataset AVSPEECH. With this dataset in hand, we then generated a training set of “*synthetic cocktail parties*”—mixtures of face videos with clean speech, and other speech audio tracks and background noise.

We demonstrate the benefits of our approach over recent speech separation methods in two ways. First, we show superior results compared to a state-of-the-art audio-only method on pure speech mixtures. Second, we demonstrate our model’s capability of producing enhanced sound streams from mixtures containing both overlapping speech and background noise in real-world scenarios.

To summarize, our paper makes two main contributions: (a) An audio-visual speech separation model that outperforms audio-only and audio-visual models on classic speech separation tasks, and is applicable in challenging, natural scenes. To our knowledge, our paper is the first to propose a speaker-independent audio-visual model for speech separation. (b) A new, large-scale audio-visual dataset, AVSPEECH, carefully collected and processed, comprised of video segments where the audible sound belongs to a single person, visible in the video, and no audio background interference. This dataset allows us to achieve state-of-the-art results on speech separation and may be useful for the research community for further studies. Our videos, results and supplementary material are available on the project web page: <http://looking-to-listen.github.io/>.

2 RELATED WORK

We briefly review related work in the areas of speech separation and audio-visual signal processing.

Speech separation. Speech separation is one of the fundamental problems in audio processing and has been the subject of extensive study over the last decades. Wang and Chen [2017] give a comprehensive overview of recent audio-only methods based on deep

learning that tackle both speech denoising [Erdogan et al. 2015; Wenginger et al. 2015] and speech separation tasks.

Two recent works have emerged which solve the aforementioned label permutation problem to perform speaker-independent, multi-speaker separation in the single-channel case. Hershey et al. [2016] propose a method called *deep clustering* in which discriminatively-trained speech embeddings are used to cluster and separate the different sources. Hershey et al. [2016] also introduced the idea of a permutation-free or *permutation invariant* loss function, but they did not find that it worked well. Isik et al. [2016] and Yu et al. [2017] subsequently introduced methods which successfully use a permutation invariant loss function to train a DNN.

The advantages of our approach over such audio-only methods are threefold: First, we show that the separation results of our audio-visual model are of higher quality than those of a state-of-the-art-inspired audio-only model. Second, our approach performs well in the setting of multiple speakers mixed with background noise, which, to our knowledge, no audio-only method has satisfactorily solved. Third, we jointly solve two speech processing problems: speech separation, and assignment of a speech signal to its corresponding face, which, thus far, have been tackled separately [Hoover et al. 2017; Hu et al. 2015; Monaci 2011].

Vision and speech. There is increased interest in using neural networks for multi-modal fusion of auditory and visual signals to solve various speech-related problems. These include audio-visual speech recognition [Feng et al. 2017; Mroueh et al. 2015; Ngiam et al. 2011], predicting speech or text from silent video (lipreading) [Chung et al. 2016; Ephrat et al. 2017], and unsupervised learning of language from visual and speech signals [Harwath et al. 2016]. These methods leverage natural synchrony between simultaneously recorded visual and auditory signals.

Audio-visual (AV) methods have also been used for speech separation and enhancement [Hershey et al. 2004; Hershey and Casey 2002; Khan 2016; Rivet et al. 2014]. Casanovas et al. [2010] perform AV source separation using sparse representations, which is limited due to dependence on active-alone regions to learn source characteristics, and the assumption that all the audio sources are seen on-screen. Recent methods have used neural networks to perform the task. Hou et al. [2018] propose a multi-task CNN-based model which outputs a denoised speech spectrogram as well as a reconstruction of the input mouth region. Gabbay et al. [2017] train a speech enhancement model on videos where other speech samples of the target speaker are used as background noise, in a scheme they call “*noise-invariant training*”. In concurrent work, Gabbay et al. [2018] use a video-to-sound synthesis method to filter noisy audio.

The main limitation of these AV speech separation approaches is that they are *speaker-dependent*, meaning a dedicated model must be trained for each speaker separately. While these works make specific design choices that limit their applicability only to the speaker-dependent case, we speculate that the main reason a speaker-independent AV model hasn’t been pursued widely so far is the lack of a sufficiently large and diverse dataset for training such models — a dataset like the one we construct and provide in this work. To the best of our knowledge, our paper is the first to address the problem of *speaker-independent* AV speech separation. Our model is capable of separating and enhancing speakers it has never seen before, speaking in languages that were not part of the

Fig. 2. **AVSPEECH dataset:** We first gathered a large collection of 100,000 high-quality, online public videos of talks and lectures (a). From these videos we extracted segments with clean speech (e.g. no mixed music, audience sounds or other speakers), and with the speaker visible in the frame (see Section 3 and Figure 3, for details of the processing). This resulted in 1500 hours of video clips, each of a single person talking with no background interference (b). This data spans a wide variety of people, languages, and face poses, with distributions shown in (c) (age and head angles estimated with automatic classifiers; language based on YouTube metadata). For a detailed list of video sources in our dataset please refer to the project web page.

training set. In addition, our work is unique in that we show high quality speech separation on real world examples, in settings that previous audio-only and audio-visual speech separation work did not address.

A number of independent and concurrent works have recently emerged which address the problem of audio-visual sound source separation using deep neural networks. [Owens and Efros 2018] train a network to predict whether audio and visual streams are temporally aligned. Learned features extracted from this self-supervised model are then used to condition an on/off screen speakers source separation model. Afouras et al. [2018] perform speech enhancement by using a network to predict both magnitude and phase of denoised speech spectrograms. Zhao et al. [2018] and Gao et al. [2018] addressed the closely related problem of separating the sound of multiple on-screen objects (e.g. musical instruments).

Audio-visual datasets. Most existing AV datasets comprise videos with only a small number of subjects, speaking words from a limited vocabulary. For example, the CUAVE dataset [Patterson et al. 2002] contains 36 subjects saying each digit from 0 to 9 five times each, with a total of 180 examples per digit. Another example is the Mandarin sentences dataset, introduced by Hou et al. [2018], which contains video recordings of 320 utterances of Mandarin sentences spoken by a native speaker. Each sentence contains 10 Chinese characters with equally distributed phonemes. The TCD-TIMIT dataset [Harte and Gillen 2015] consists of 60 volunteer speakers with around 200 videos each. The speakers recite various sentences

from the TIMIT dataset [S Garofolo et al. 1992], and are recorded using both front-facing and 30-degree cameras. We evaluate our results on these three datasets in order to compare to previous work.

Recently, the large-scale Lip Reading Sentences (LRS) dataset was introduced by Chung et al. [2016], which includes both a wide variety of speakers and words from a larger vocabulary. However, not only is that dataset not publicly available, but the speech in LRS videos is not guaranteed to be clean, which is crucial for training a model for speech separation and enhancement.

3 AVSPEECH DATASET

We introduce a new, large-scale audio-visual dataset comprising speech clips with no interfering background signals. The segments are of varying length, between 3 and 10 seconds long, and in each clip the only visible face in the video and audible sound in the soundtrack belong to a single speaking person. In total, the dataset contains roughly 1500 hours of video segments, spanning a wide variety of people, languages and face poses. Representative frames, audio waveforms and some dataset statistics are shown in Figure 2.

We collected the dataset automatically, since for assembling a corpus of this magnitude it was important not to rely on substantial human feedback. Our dataset creation pipeline collected clips from roughly 100,000 YouTube videos of lectures (e.g. TED talks) and how-to videos. For such channels, most of the videos comprise a single speaker, and both the video and audio are generally of high quality.

Fig. 3. **Video and audio processing for dataset creation:** (a) We use face detection and tracking to extract speech segment candidates from videos and reject frames in which faces are blurred or not sufficiently frontal-facing. (b) We discard segments with noisy speech by estimating speech SNR (see Section 3). The plot is intended to show the accuracy of our speech SNR estimator (and thus the quality of the dataset). We compare true speech SNR with our predicted SNR for synthetic mixtures of clean speech and non-speech noise at known SNR levels. Predicted SNR values (in dB) are averaged over 60 generated mixtures per SNR bin, with error bars representing 1 std. We discard segments for which the predicted speech SNR is below 17 dB (marked by the gray dotted line in the plot).

Dataset creation pipeline. Our dataset collection process has two main stages, as illustrated in Figure 3. First, we used the speaker tracking method of Hoover et al. [2017] to detect video segments of a person actively speaking with their face visible. Face frames that were blurred, insufficiently illuminated or had extreme pose were discarded from the segments. If more than 15% of a segment’s face frames were missing, it was discarded altogether. We used Google Cloud Vision API¹ for the classifiers in this stage, and to compute the statistics in Figure 2.

The second step in building the dataset is refining the speech segments to include only clean, non-interfered speech. This is a crucial component because such segments serve as ground truth during training. We perform this refinement step automatically by estimating the speech SNR (the log ratio of the main speech signal to the rest of the audio signal) of each segment as follows.

We used a pre-trained audio-only speech denoising network to predict the SNR of a given segment using the denoised output as an estimation of the clean signal. The architecture of this network is the same as the one implemented for the audio-only speech enhancement baseline in Section 5, and it was trained on speech from the LibriVox collection of public domain audio books. Segments for which the estimated SNR is below a threshold were rejected. The threshold was set empirically using synthetic mixtures of clean speech and non-speech interfering noise at different, known SNR

levels.² These synthetic mixtures were fed into the denoising network and the estimated (denoised) SNR was compared to the ground truth SNR (see Figure. 3(b)).

We found that at low SNRs, on average, the estimated SNR is very accurate, thus can be considered a good predictor of the original noise level. At higher SNRs (i.e. segments with little-to-no interference of the original speech signal), the accuracy of this estimator diminishes because the noise signal is faint. The threshold at which this occurs is at around 17 dB, as can be seen in Figure 3(b). We listened to a random sample of 100 clips which passed this filtering, and found that none of them contained noticeable background noise. We provide sample video clips from our dataset in the supplementary material.

4 AUDIO-VISUAL SPEECH SEPARATION MODEL

At a high-level, our model is comprised of a multi-stream architecture which takes visual streams of detected faces and noisy audio as input, and outputs complex spectrogram masks, one for each detected face in the video (Figure 4). The noisy input spectrograms are then multiplied by the masks to obtain an isolated speech signal for each speaker, while suppressing all other interfering signals.

4.1 Video and Audio Representation

Input features. Our model takes both visual and auditory features as input. Given a video clip containing multiple speakers, we use an off-the-shelf face detector (e.g. Google Cloud Vision API) to find faces in each frame (75 face thumbnails altogether per speaker, assuming 3-second clips at 25 FPS). We use a pretrained face recognition model to extract one face embedding per frame for each of the detected face thumbnails. We use the lowest layer in the network that is not spatially varying, similar to the one used by Cole et al. [2016] for synthesizing faces. The rationale for this is that these embeddings retain information necessary for recognizing millions of faces, while discarding irrelevant variation between images, such as illumination. In fact, recent work also demonstrated that it is possible to recover facial expressions from such embeddings [Rudd et al. 2016]. We also experimented with raw pixels of the face images, which did not lead to improved performance.

As for the audio features, we compute the short-time Fourier transform (STFT) of 3-second audio segments. Each time-frequency (TF) bin contains the real and imaginary parts of a complex number, both of which we use as input. We perform power-law compression to prevent loud audio from overwhelming soft audio. The same processing is applied to both the noisy signal and the clean reference signal.

At inference time, our separation model can be applied to arbitrarily long segments of video. When more than one speaking face is detected in a frame, our model can accept multiple face streams as input, as we will discuss shortly.

Output. The output of our model is a multiplicative spectrogram mask, which describes the time-frequency relationships of clean speech to background interference. In previous work [Wang and Chen 2017; Wang et al. 2014], multiplicative masks have been observed to work better than alternatives such as direct prediction

¹<https://cloud.google.com/vision/>

²Such mixtures simulate well the type of interference in our dataset, which typically involves a single speaker interfered by non-speech sounds like audience clapping or intro music.

Fig. 4. **Our model's multi-stream neural network-based architecture:** The visual streams take as input thumbnails of detected faces in each frame in the video, and the audio stream takes as input the video's soundtrack, containing a mixture of speech and background noise. The visual streams extract face embeddings for each thumbnail using a pretrained face recognition model, then learn a visual feature using a dilated convolutional NN. The audio stream first computes the STFT of the input signal to obtain a spectrogram, and then learns an audio representation using a similar dilated convolutional NN. A joint, audio-visual representation is then created by concatenating the learned visual and audio features, and is subsequently further processed using a bidirectional LSTM and three fully connected layers. The network outputs a complex spectrogram mask for each speaker, which is multiplied by the noisy input, and converted back to waveforms to obtain an isolated speech signal for each speaker.

Table 1. Dilated convolutional layers comprising our model's audio stream.

	conv1	conv2	conv3	conv4	conv5	conv6	conv7	conv8	conv9	conv10	conv11	conv12	conv13	conv14	conv15
Num Filters	96	96	96	96	96	96	96	96	96	96	96	96	96	96	8
Filter Size	1×7	7×1	5×5	5×5	5×5	5×5	5×5	5×5	5×5	5×5	5×5	5×5	5×5	5×5	1×1
Dilation	1×1	1×1	1×1	2×1	4×1	8×1	16×1	32×1	1×1	2×2	4×4	8×8	16×16	32×32	1×1
Context	1×7	7×7	9×9	13×11	21×13	37×15	69×17	133×19	135×21	139×25	147×33	163×49	195×81	259×145	259×145

Table 2. Dilated convolutional layers comprising our model's visual streams.

	conv1	conv2	conv3	conv4	conv5	conv6
Num Filters	256	256	256	256	256	256
Filter Size	7×1	5×1	5×1	5×1	5×1	5×1
Dilation	1×1	1×1	2×1	4×1	8×1	16×1
Context	7×1	9×1	13×1	21×1	37×1	69×1

of spectrogram magnitudes or direct prediction of time-domain waveforms. Many types of masking-based training targets exist in the source separation literature [Wang and Chen 2017], of which we experiment with two: ratio mask (RM) and complex ratio mask (cRM).

The ideal ratio mask is defined as the ratio between the magnitudes of the clean and noisy spectrograms, and is assumed to lie between 0 and 1. The complex ideal ratio mask is defined as the ratio of the complex clean and noisy spectrograms. The cRM has a real component and an imaginary component, which are separately estimated in the real domain. Real and imaginary parts of the complex mask will typically lie between -1 and 1, however, we

use sigmoidal compression to bound these complex mask values between 0 and 1 [Wang et al. 2016].

When masking with cRM, denoised waveforms are obtained by performing inverse STFT (ISTFT) on the complex multiplication of the predicted cRM and noisy spectrogram. When using RM, we perform ISTFT on the point-wise multiplication of the predicted RM and noisy spectrogram magnitude, combined with the noisy original phase [Wang and Chen 2017].

Given multiple detected speakers' face streams as input, the network outputs a separate mask for each speaker, and one for background interference. We perform most of our experiments using cRM, as we found that output speech quality using it was significantly better than RM. See Table 6 for a quantitative comparison of the two methods.

4.2 Network architecture

Fig. 4 provides a high-level overview of the various modules in our network, which we will now describe in detail.

Audio and visual streams. The audio stream part of our model consists of dilated convolutional layers, the parameters of which are specified in Table 1.

The visual stream of our model is used to process the input face embeddings (see Section 4.1), and consists of dilated convolutions as detailed in Table 2. Note that “spatial” convolutions and dilations in the visual stream are performed over the temporal axis (not over the 1024-D face embedding channel).

To compensate for the sampling rate discrepancy between the audio and video signals, we upsample the output of the visual stream to match the spectrogram sampling rate (100 Hz). This is done using simple nearest neighbor interpolation in the temporal dimension of each visual feature.

AV fusion. The audio and visual streams are combined by concatenating the feature maps of each stream, which are subsequently fed into a BLSTM followed by three FC layers. The final output consists of a complex mask (two-channels, real and imaginary) for each of the input speakers. The corresponding spectrograms are computed by complex multiplication of the noisy input spectrogram and the output masks. The squared error (L2) between the power-law compressed clean spectrogram and the enhanced spectrogram is used as a loss function to train the network. The final output waveforms are obtained using ISTFT, as described in Section 4.1.

Multiple speakers. Our model supports isolation of multiple visible speakers in a video, each represented by a visual stream, as illustrated in Fig. 4. A separate, dedicated model is trained for each number of visible speakers, e.g. a model with one visual stream for one visible speaker, double visual stream model for two, etc. All the visual streams share the same weights across convolutional layers. In this case, the learned features from each visual stream are concatenated with the learned audio features before continuing on to the BLSTM. It should be noted that in practice, a model which takes a single visual stream as input can be used in the general case in which either the number of speakers is unknown, or a dedicated multi-speaker model is unavailable.

4.3 Implementation details

Our network is implemented in TensorFlow, and its included operations are used for performing waveform and STFT transformations. ReLU activations follow all network layers except for last (mask), where a sigmoid is applied. Batch normalization [Ioffe and Szegedy 2015] is performed after all convolutional layers. Dropout is not used, as we train on a large amount of data and do not suffer from overfitting. We use a batch size of 6 samples and train for 5 million steps (batches) with a learning rate of $3 \cdot 10^{-5}$ which is reduced by half every 1.8 million steps.

All audio is resampled to 16kHz, and stereo audio is converted to mono by taking only the left channel. STFT is computed using a Hann window of length 25ms, hop length of 10ms, and FFT size of 512, resulting in an input audio feature of $257 \times 298 \times 2$ scalars. Power-law compression is performed with $p = 0.3$ ($A^{0.3}$, where A is the input/output audio spectrogram).

We resample the face embeddings from all videos to 25 frames-per-second (FPS) before training and inference by either removing or replicating embeddings. This results in an input visual stream of 75 face embeddings. Face detection, alignment and quality assessment

is performed using the tools described by Cole et al. [2016]. When missing frames are encountered in a particular sample, we use a vector of zeros in lieu of a face embedding.

5 EXPERIMENTS AND RESULTS

We tested our method in a variety of conditions and also compared our results to state-of-the-art audio-only (AO) and audio-visual (AV) speech separation and enhancement, both quantitatively and qualitatively.

Comparison with Audio-Only. There are no publicly available state-of-the-art audio-only speech enhancement/separation systems, and relatively few publicly available datasets for training and evaluating audio-only speech enhancement. And although there is extensive literature on “blind source separation” for audio-only speech enhancement and separation [Comon and Jutten 2010], most of these techniques require multiple audio channels (multiple microphones), and are therefore not applicable to our task. For these reasons, we implemented an AO baseline for speech enhancement which has a similar architecture to the audio stream in our audio-visual model (Fig. 4, when stripping out the visual streams). When trained and evaluated on the CHiME-2 dataset [Vincent et al. 2013], which is widely used for speech enhancement work, our AO baseline achieved a signal-to-distortion ratio of 14.6 dB, nearly as good as the state-of-the-art single channel result of 14.75 dB reported by Erdogan et al. [2015]. Our AO enhancement model is therefore deemed a near state-of-the-art baseline.

In order to compare our separation results to those of a state-of-the-art AO model, we implemented the *permutation-invariant training* introduced by Yu et al. [2017]. Note that speech separation using this method requires *a priori* knowledge of the number of sources present in the recording, and also requires *manual* assignment of each output channel to the face of its corresponding speaker in the video (which our AV method does automatically).

We use these AO methods in all our synthetic experiments in Section 5.1, and also show qualitative comparisons to it on real videos in Section 5.2.

Comparison with Recent Audio-Visual Methods. Since existing AV speech separation and enhancement methods are speaker dependent, we could not easily compare to them in our experiments on synthetic mixtures (Section 5.1), or run them on our natural videos (Section 5.2). However, we show quantitative comparisons with those methods on existing datasets by running our model on videos from those papers. We discuss this comparison in more detail in Section 5.3. In addition, we show qualitative comparisons in our supplementary material.

5.1 Quantitative Analysis on Synthetic Mixtures

We generated data for several different single-channel speech separation tasks. Each task requires its own unique configuration of mixtures of speech and non-speech background noise. We describe below the generation procedure for each variant of training data, as well as the relevant models for each task, which were trained from scratch.

In all cases, clean speech clips and corresponding faces are taken from our AVSPEECH (AVS) dataset. Non-speech background noise is obtained from AudioSet [Gemmeke et al. 2017], a large-scale dataset

Table 3. **Quantitative analysis and comparison with audio-only speech separation and enhancement:** Quality improvement (in SDR, see Section A in the Appendix) as function of the number of input visual streams using different network configurations. First row (audio-only) is our implementation of a state-of-the-art speech separation model, and shown as a baseline.

	1S+Noise	2S clean	2S+Noise	3S clean
AO [Yu et al. 2017]	16.2	8.8	10.0	8.5
AV - 1 face	16.3	10.0	10.1	9.1
AV - 2 faces	-	10.8	10.6	9.1
AV - 3 faces	-	-	-	9.9

of manually-annotated segments from YouTube videos. Separated speech quality is evaluated using signal-to-distortion ratio (SDR) improvement from the BSS Eval toolbox [Vincent et al. 2006], a commonly used metric for evaluating speech separation quality (see Section A in the Appendix).

We extracted 3-second non-overlapping segments from the varying-length segments in our dataset (e.g. a 10-sec segment would contribute 3 3-second segments). We generated 1.5 million synthetic mixtures for all the models and experiments. For each experiment, 90% of the generated data was taken to be the training set, and the remaining 10% was used as the test set. We did not use any validation set as no parameter tuning or early stopping were performed.

One speaker + noise (1S+Noise). This is a classic speech enhancement task, for which the training data was generated by a linear combination of unnormalized clean speech and AudioSet noise: $Mix_i = AVS_j + 0.3 * AudioSet_k$ where AVS_j is one utterance from AVS, $AudioSet_k$ is one segment from AudioSet with its amplitude multiplied by 0.3, and Mix_i is a sample in the generated dataset of synthetic mixtures. Our audio only model performs quite well in this case, because the characteristic frequencies of noise are typically well separated from the characteristic frequencies of speech. Our audio-visual (AV) model performs as well as the audio-only (AO) baseline with SDR of 16.3 dB (first column of Table 3).

Two clean speakers (2S clean). The dataset for this two-speaker separation scenario was generated by mixing clean speech of two different speakers from our AVS dataset: $Mix_i = AVS_j + AVS_k$, where AVS_j and AVS_k are clean speech samples from different source videos in our dataset, and Mix_i is a sample in the generated dataset of synthetic mixtures. We trained two different AV models on this task, in addition to our AO baseline:

(i) A model which takes only one visual stream as input, and outputs only its corresponding denoised signal. In this case, at inference, the denoised signal of each speaker is obtained by two forward passes in the network (one for each speaker). Averaging the SDR results of this model gives an improvement of 1.2 dB over our AO baseline (second column of Table 3).

(ii) A model which takes visual information from both speakers as input, in two separate streams (as explained in Section 4). In this case, the output consists of two masks, one for each speaker, and inference is done with a single forward pass. An additional boost of 0.8 dB is obtained using this model, resulting in a 10.8 dB total SDR improvement. Intuitively, jointly processing two visual streams

Fig. 5. **Input SDR vs. output SDR improvement:** A scatter plot showing separation performance (SDR improvement) as a function of original (noisy) SDR for the task of separating two clean speakers (2S clean). Each point corresponds to a single, 3-second audio-visual sample from the test set.

Fig. 6. **Example of input and output audio:** The top row shows the audio spectrogram for one segment in our training data, involving two speakers and background noise (a), together with the ground truth, separate spectrograms of each speaker (b, c). In the bottom row we show our results: the masks our method estimates for that segment, superimposed on one spectrogram with a different color for each speaker (d), and the corresponding output spectrograms for each speaker (e, f).

provides the network with more information and imposes more constraints on the separation task, hence improving the results.

Fig. 5 shows the SDR improvement as a function of input SDR for this task, for both the audio-only baseline and our two-speaker audio-visual model.

Two speakers + noise (2S+Noise). Here, we consider the task of isolating one speaker's voice from a mixture of two speakers and non-speech background noise. To the best of our knowledge, this audio-visual task has not been addressed before. The training data was generated by mixing clean speech of two different speakers (as generated for the 2S clean task) with background noise from AudioSet: $Mix_i = AVS_j + AVS_k + 0.3 * AudioSet_l$.

In this case we trained the AO network with three outputs, one for each speaker and one for background noise. In addition, we trained two different configurations of our model, with one and two visual streams received as input. The configuration of the one-stream AV model is the same as in model (i) in the previous experiment. The two-stream AV outputs three signals, one for each speaker and one

Table 4. **Same-gender separation.** The results in this table, from the 2S *clean* experiment, show that our method is robust to separation of speech from same-gender mixtures.

	SDR
Male-Male	10.2
Female-Female	11.1
Male-Female	11.0

for background noise. As can be seen in Table 3 (third column), the SDR gain of our one-stream AV model over the audio only baseline is 0.1 dB, and 0.6 dB for two streams, bringing the total SDR improvement to 10.6 dB. Fig. 6 shows the inferred masks and output spectrograms for a sample segment from this task, along with its noisy input and ground truth spectrograms.

Three clean speakers (3S clean). The dataset for this task is created by mixing clean speech from three different speakers: $Mix_i = AVS_j + AVS_k + AVS_l$. In a similar manner to the previous tasks, we trained our AV model with one, two and three visual streams as input, which output one, two and three signals, respectively.

We found that even when using a single visual stream, the AV model performs better than the AO model, with a 0.6 dB improvement over it. The two visual stream configuration gives the same improvement over the AO model, while using three visual streams leads to a gain of 1.4 dB, attaining a total 9.9 dB SDR improvement (fourth column of Table 3).

Same-gender separation. Many previous speech separation methods show a drop in performance when attempting to separate speech mixtures containing same-gender speech [Delfarah and Wang 2017; Hershey et al. 2016]. Table 4 shows a breakdown of our separation quality by the different gender combinations. Interestingly, our model performs best (by a small margin) on female-female mixtures, but performs well on the other combinations as well, demonstrating its gender robustness.

5.2 Real-World Speech Separation

In order to demonstrate our model’s speech separation capabilities in real-world scenarios, we tested it on an assortment of videos containing heated debates and interviews, noisy bars and screaming children (Fig. 7). In each scenario we use a trained model whose number of visual input streams matches the number of visible speakers in the video. For example, for a video with two visible speakers, a two-speaker model was used. We performed separation using a single forward pass per video, which our model supports, since our network architecture never enforces a specific temporal duration. This allows us to avoid the need to post-process and consolidate results on shorter chunks of the video. Because there is no clean reference audio for these examples, these results and their comparison to other methods are evaluated qualitatively; they are presented in our supplementary material. It should be noted that our method does not work in real-time, and, in its current form, our speech enhancement is better suited for the post-processing stage of video editing.

The synthetic “Double Brady” video in our supplementary material highlights the utilization of visual information by our model, as it is very difficult to perform speech separation in this scenario using only characteristic speech frequencies contained in the audio.

The “Noisy Bar” scene shows a limitation of our approach in separating speech from mixtures with low SNR. In this case, the background noise is almost entirely suppressed, however output

Table 5. **Comparison with existing audio-visual speech separation work.** We compare our speech separation and enhancement results on several datasets to those of previous work, using the evaluation protocols and objective scores reported in the original papers. Note that previous approaches are *speaker-dependent*, whereas our results are obtained by using a general, *speaker-independent* model.

Mandarin (Enhancement)			
	Gabbay et al. [2017]	Hou et al. [2018]	Ours
PESQ	2.25	2.42	2.5
STOI	-	0.66	0.71
SDR	-	2.8	6.1
TCD-TIMIT (Separation)			
	Gabbay et al. [2017]	Ours	
SDR	0.4	4.1	
PESQ	2.03	2.42	
CUAVE (Separation)			
	Casanovas et al. [2010]	Pu et al. [2017]	Ours
SDR	7	6.2	12.6

speech quality is noticeably degraded. Sun et al. [2017] observed that this limitation stems from the use of a masking-based approach for separation, and that in this scenario, directly predicting the denoised spectrogram could help overcome this problem. In cases of classic speech enhancement, i.e. one speaker with non-speech background noise, our AV model obtains similar results to those of our strong AO baseline. We suspect this is because the characteristic frequencies of noise are typically well separated from the characteristic frequencies of speech, and therefore incorporating visual information does not provide additional discrimination capabilities.

5.3 Comparison with Previous Work in Audio-Visual Speech Separation and Enhancement

Our evaluation would not be complete without comparing our results to those of previous work in AV speech separation and enhancement. Table 5 contains these comparisons on three different AV datasets, Mandarin, TCD-TIMIT and CUAVE, mentioned in Section 2, using the evaluation protocols and metrics described in the respective papers. The reported objective quality scores are PESQ [Rix et al. 2001], STOI [Taal et al. 2010] and SDR from the BSS eval toolbox [Vincent et al. 2006]. Qualitative results of these comparisons are available on our project page.

It is important to note that these prior methods require training a dedicated model for each speaker in their dataset (*speaker dependent*), whereas our evaluation on their data is done using a model trained on our general AVS dataset (*speaker independent*). Despite having never encountered these particular speakers before, our results are significantly better than those reported in the original papers, indicating the strong generalization capability of our model.

5.4 Application to Video Transcription

While our focus in this paper is speech separation and enhancement, our method can also be useful for automatic speech recognition (ASR) and video transcription. As a proof of concept, we

Fig. 7. **Speech separation in the wild:** Representative frames from natural videos demonstrating our method in various real-world scenarios. All videos and results can be found in the supplementary material. The “Undisputed Interview” video is courtesy of Fox Sports.

performed the following qualitative experiment. We uploaded our speech-separated results for the “Stand-Up” video to YouTube, and compared the resulting captions produced by YouTube’s automatic captioning³ with those it produced for the corresponding source videos with mixed speech. For parts of the original “Stand-Up” video, the ASR system was unable to generate any captions in mixed speech segments of the video. The results included speech from both speakers, resulting in hard-to-read sentences. However, captions produced on our separated speech results were noticeably more accurate. We show the full captioned videos in our supplementary material.

5.5 Additional Analysis

We also conducted extensive experiments to better understand the model’s behavior and how its different components affect the results.

Ablation study. In order to better understand the contribution of different parts of our model, we performed an ablation study on the task of speech separation from a mixture of two clean speakers (2S Clean). In addition to ablating several combinations of network modules (visual and audio streams, BLSTM and FC layers), we also investigated higher-level changes such as a different output mask (magnitude), the effect of reducing the learned visual features to one scalar per timestep, and a different fusion method (early fusion).

In the early fusion model, we do not have separate visual and audio streams, but rather combine the two modalities at the input. This is done by first using two fully connected layers to reduce the dimensionality of each visual embedding to match the spectrogram dimension at each timestep, then stacking the visual features as a

³<https://support.google.com/youtube/answer/6373554?hl=en>

Table 6. **Ablation study:** We investigate the contribution of different parts of our model on the scenario of separating mixtures of two clean speakers. SDR correlates well with noise suppression, and ViSQOL indicates level of speech quality (see Section A in the Appendix)

	SDR	ViSQOL
Our full model (cRM)	10.8	3.17
- No FC	10.1	2.88
- No BLSTM	9.1	2.89
- Audio-only (input)	8.8	2.76
- No BLSTM or FC	1.0	1.87
- Visual-only (input)	1.6	2.05
Magnitude mask (RM)	10.0	2.81
Bottleneck (cRM)	10.0	3.08
Early fusion (cRM)	8.2	2.44
Oracle RM + noisy ph.	11.1	3.52
Oracle cRM	15.1	3.77
Oracle RM + oracle ph.	18.9	4.86

third spectrogram “channel” and processing them jointly throughout the model.

Table 6 shows the results of our ablation study. The table includes evaluation using SDR and ViSQOL [Hines et al. 2015], an objective measure intended to approximate human listener mean opinion scores (MOS) of speech quality. The ViSQOL scores were calculated on a random 2000 sample subset of our testing data. We found that SDR correlates well with the amount of noise left in the separated audio, and ViSQOL is a better indicator of output speech quality. See Section A in the Appendix for more details on these scores. “Oracle” RMs and cRMs are masks obtained as described in Section 4.1, by using the ground truth real-valued and complex-valued spectrograms, respectively.

The most interesting findings of this study are the drop in MOS when using a real-valued magnitude mask rather than a complex one, and the surprising effectiveness of squeezing the visual information into one scalar per timestep, described below.

Bottleneck features. In our ablation analysis we found that a network which squeezes the visual information into a bottleneck of one scalar per timestep (“Bottleneck (cRM)”) performs almost as well (only 0.8 dB less) as our full model (“Full model (cRM)”) that uses 64 scalars per timestep.

How does the model utilize the visual signal? Our model uses face embeddings as the input visual representation (Section 4.1). We want to gain insights on the information captured in these high-level features and to identify which regions of the input frames are used by the model for separating the speech. To this end, we follow a similar protocol as in [Zeiler and Fergus 2014; Zhou et al. 2014] for visualizing receptive fields of deep networks. We extend that protocol from 2D images to 3D (space-time) video. More specifically, we use a space-time patch occluder ($11px \times 11px \times 200ms$ patch⁴) in a sliding window fashion. For each space-time occluder, we feed-forward the occluded video into our model and compare the speech separation result, S_{occ} , with the one obtained on the original (non-occluded) video, S_{orig} . To quantify the difference between the network outputs, we use SNR, treating the result without

⁴We use 200ms length to cover the typical range of phoneme duration: 30-200 ms.

Fig. 8. **How does the model utilize the visual signal?** We show heat maps overlaid on representative input frames from several videos, visualizing the contribution of different regions of the frames to our speech separation result (in dB, see text), from blue (low contribution) to red (high contribution).

the occluder as the “signal”⁵. That is, for each space-time patch, we compute:

$$E = 10 \cdot \log \left(\frac{S_{orig}^2}{(S_{occ} - S_{orig})^2} \right). \quad (1)$$

Repeating this process for all space-time patches in a video results in a heat map for each frame. For visualization purposes we normalize the heat maps by the maximum SNR for the video: $\tilde{E} = E_{max} - E$. In $\tilde{E}$, high values correspond to patches with high impact on the speech separation result.

In Fig. 8 we show the resulting heat maps for representative frames from several videos (the full heat map videos are available on our project page). As expected, the facial regions that contribute the most are located around the mouth, yet the visualization reveals that other areas such as the eyes and cheeks contribute as well.

Effect of missing visual information. We further tested the contribution of visual information to the model by gradual elimination of visual embeddings. Specifically, we start by running the model and evaluating the speech separation quality using visual information for the full 3 second video. We then gradually discard embeddings from both ends of the segment, and re-evaluate the separation quality with visual durations of 2, 1, 0.5 and 0.2 seconds.

The results are shown in Fig. 9. Interestingly, the speech separation quality is reduced by only 0.5 dB on average when dropping as much as 2/3 of the visual embeddings in the segments. This shows the robustness of the model to missing visual information, which may occur in real world scenarios due to head motion or occlusions.

6 CONCLUSION

We proposed an audio-visual neural network-based model for single-channel, speaker-independent speech separation. Our model works

⁵We refer the reader to the supplementary material to validate that our separated speech on the non-occluded video, which we treat as “correct” in this example, is indeed accurate.

Fig. 9. **The effect of missing visual information:** This graph shows the impact of the duration of the visual information on output SDR improvement in the 2S clean scenario. We test this by gradually zeroing-out input face embeddings from both ends of the sample. The results show that even a small number of visual frames are sufficient for high-quality separation.

well in challenging scenarios, including multi-speaker mixtures with background noise. To train the model, we created a new audio-visual dataset with thousands of hours of video segments containing visible speakers and clean speech we collected from the Web. We showed state-of-the-art results on speech separation as well as a potential application to video captioning and speech recognition. We also conducted extensive experiments to analyze the behavior of our model and its components.

ACKNOWLEDGEMENTS

We would like to thank Yossi Matias and Google Research Israel for their support for the project, and John Hershey for his valuable feedback. We also thank Arkady Ziefman for his help with figure design and video editing, and Rachel Soh for helping us procure permissions for video content in our results.

REFERENCES

- T. Afouras, J. S. Chung, and A. Zisserman. 2018. The Conversation: Deep Audio-Visual Speech Enhancement. In *arXiv:1804.04121*.
- Anna Llagostera Casanovas, Gianluca Monaci, Pierre Vanderghenynst, and Rémi Grignonval. 2010. Blind audiovisual source separation based on sparse redundant representations. *IEEE Transactions on Multimedia* 12, 5 (2010), 358–371.
- E Colin Cherry. 1953. Some experiments on the recognition of speech, with one and with two ears. *The Journal of the acoustical society of America* 25, 5 (1953), 975–979.
- Joon Son Chung, Andrew W. Senior, Oriol Vinyals, and Andrew Zisserman. 2016. Lip Reading Sentences in the Wild. *CoRR* abs/1611.05358 (2016).
- Forrester Cole, David Belanger, Dilip Krishnan, Aaron Sarna, Inbar Mosseri, and William T Freeman. 2016. Synthesizing normalized faces from facial identity features. In *CVPR’17*.
- Pierre Comon and Christian Jutten. 2010. *Handbook of Blind Source Separation: Independent component analysis and applications*. Academic press.
- Masood Delfarah and DeLiang Wang. 2017. Features for Masking-Based Monaural Speech Separation in Reverberant Conditions. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 25 (2017), 1085–1094.
- Ariel Ephrat, Tavi Halperin, and Shmuel Peleg. 2017. Improved Speech Reconstruction from Silent Video. In *ICCV 2017 Workshop on Computer Vision for Audio-Visual Media*.
- Hakan Erdogan, John R. Hershey, Shinji Watanabe, and Jonathan Le Roux. 2015. Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2015).
- Wei Jiang Feng, Naiyang Guan, Yuan Li, Xiang Zhang, and Zhigang Luo. 2017. Audio visual speech recognition with multimodal recurrent neural networks. In *Neural Networks (IJCNN), 2017 International Joint Conference on*. IEEE, 681–688.
- Aviv Gabbay, Ariel Ephrat, Tavi Halperin, and Shmuel Peleg. 2018. Seeing Through Noise: Speaker Separation and Enhancement using Visually-derived Speech. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2018).
- Aviv Gabbay, Asaph Shamir, and Shmuel Peleg. 2017. Visual Speech Enhancement using Noise-Invariant Training. *arXiv preprint arXiv:1711.08789* (2017).
- R. Gao, R. Feris, and K. Grauman. 2018. Learning to Separate Object Sounds by Watching Unlabeled Video. *arXiv preprint arXiv:1804.01665* (2018).
- Jort F. Gemmeke, Daniel P. W. Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R. Channing Moore, Manoj Plakal, and Marvin Ritter. 2017. Audio Set: An ontology and human-labeled dataset for audio events. In *Proc. IEEE ICASSP 2017*.

- Elana Zion Golumbic, Gregory B Cogan, Charles E. Schroeder, and David Poeppel. 2013. Visual input enhances selective speech envelope tracking in auditory cortex at a "cocktail party". *The Journal of neuroscience : the official journal of the Society for Neuroscience* 33 4 (2013), 1417–26.
- Naomi Harte and Eoin Gillen. 2015. TCD-TIMIT: An audio-visual corpus of continuous speech. *IEEE Transactions on Multimedia* 17, 5 (2015), 603–615.
- David F. Harwath, Antonio Torralba, and James R. Glass. 2016. Unsupervised Learning of Spoken Language with Visual Context. In *NIPS*.
- John Hershey, Hagai Attias, Nebojsa Jojic, and Trausti Kristjansson. 2004. Audio-visual graphical models for speech processing. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- John R Hershey and Michael Casey. 2002. Audio-visual sound separation via hidden Markov models. In *Advances in Neural Information Processing Systems*. 1173–1180.
- John R. Hershey, Zhuo Chen, Jonathan Le Roux, and Shinji Watanabe. 2016. Deep clustering: Discriminative embeddings for segmentation and separation. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2016), 31–35.
- Andrew Hines, Eoin Gillen, Damien Kelly, Jan Skoglund, Anil C. Kokaram, and Naomi Harte. 2015. ViSQOLAudio: An objective audio quality metric for low bitrate codecs. *The Journal of the Acoustical Society of America* 137 6 (2015), EL449–55.
- Andrew Hines and Naomi Harte. 2012. Speech Intelligibility Prediction Using a Neurogram Similarity Index Measure. *Speech Commun.* 54, 2 (Feb. 2012), 306–320. DOI: <http://dx.doi.org/10.1016/j.specom.2011.09.004>
- Ken Hoover, Sourish Chaudhuri, Caroline Pantofaru, Malcolm Slaney, and Ian Sturdy. 2017. Putting a Face to the Voice: Fusing Audio and Visual Signals Across a Video to Determine Speakers. *CoRR* abs/1706.00079 (2017).
- Jen-Cheng Hou, Syu-Siang Wang, Ying-Hui Lai, Jen-Chun Lin, Yu Tsao, Hsiu-Wen Chang, and Hsin-Min Wang. 2018. Audio-Visual Speech Enhancement Using Multimodal Deep Convolutional Neural Networks. *IEEE Transactions on Emerging Topics in Computational Intelligence* 2, 2 (2018), 117–128.
- Yongtao Hu, Jimmy SJ Ren, Jingwen Dai, Chang Yuan, Li Xu, and Wenping Wang. 2015. Deep multimodal speaker naming. In *Proceedings of the 23rd ACM international conference on Multimedia*. ACM, 1107–1110.
- Sergey Ioffe and Christian Szegedy. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *ICML*.
- Yusuf Isik, Jonathan Le Roux, Zhou Chen, Shinji Watanabe, and John R Hershey. 2016. Single-Channel Multi-Speaker Separation Using Deep Clustering. *Interspeech* (2016), 545–549.
- Faheem Khan. 2016. *Audio-visual speaker separation*. Ph.D. Dissertation. University of East Anglia.
- Wei Ji Ma, Xiang Zhou, Lars A. Ross, John J. Foxe, and Lucas C. Parra. 2009. Lip-Reading Aids Word Recognition Most in Moderate Noise: A Bayesian Explanation Using High-Dimensional Feature Space. *PLoS ONE* 4 (2009), 233 – 252.
- Josh H McDermott. 2009. The cocktail party problem. *Current Biology* 19, 22 (2009), R1024–R1027.
- Gianluca Monaci. 2011. Towards real-time audiovisual speaker localization. In *Signal Processing Conference, 2011 19th European*. IEEE, 1055–1059.
- Youssef Mroueh, Etienne Marcheret, and Vaibhava Goel. 2015. Deep multimodal learning for audio-visual speech recognition. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*. IEEE, 2130–2134.
- Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y. Ng. 2011. Multimodal Deep Learning. In *ICML*.
- Andrew Owens and Alexei A Efros. 2018. Audio-Visual Scene Analysis with Self-Supervised Multisensory Features. (2018).
- Eric K. Patterson, Sabri Gurbuz, Zekeriya Tufekci, and John N. Gowdy. 2002. Moving-Talker, Speaker-Independent Feature Study, and Baseline Results Using the CUAVE Multimodal Speech Corpus. *EURASIP J. Adv. Sig. Proc.* 2002 (2002), 1189–1201.
- Jie Pu, Yannis Panagakis, Stavros Petridis, and Maja Pantic. 2017. Audio-visual object localization and separation using low-rank and sparsity. In *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2901–2905.
- Bertrand Rivet, Wenwu Wang, Syed M. Naqvi, and Jonathan A. Chambers. 2014. Audio-visual Speech Source Separation: An overview of key methodologies. *IEEE Signal Processing Magazine* 31 (2014), 125–134.
- Antony W Rix, John G Beerends, Michael P Hollier, and Andries P Hekstra. 2001. Perceptual evaluation of speech quality (PESQ)-a new method for speech quality assessment of telephone networks and codecs. In *Acoustics, Speech, and Signal Processing, 2001. Proceedings.(ICASSP'01). 2001 IEEE International Conference on*, Vol. 2. IEEE, 749–752.
- Ethan M Rudd, Manuel Günther, and Terrance E Boulton. 2016. Moon: A mixed objective optimization network for the recognition of facial attributes. In *European Conference on Computer Vision*. Springer, 19–35.
- J S Garofolo, Lori Lamel, W M Fisher, Jonathan Fiscus, D S Pallett, N L Dahlgren, and V Zue. 1992. TIMIT Acoustic-phonetic Continuous Speech Corpus. (11 1992).
- Lei Sun, Jun Du, Li-Rong Dai, and Chin-Hui Lee. 2017. Multiple-target deep learning for LSTM-RNN based speech enhancement. In *HSCMA*.
- Cees H Taal, Richard C Hendriks, Richard Heusdens, and Jesper Jensen. 2010. A short-time objective intelligibility measure for time-frequency weighted noisy speech. In *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*. IEEE, 4214–4217.
- Emmanuel Vincent, Jon Barker, Shinji Watanabe, Jonathan Le Roux, Francesco Nesta, and Marco Matassoni. 2013. The second 'chime' speech separation and recognition challenge: Datasets, tasks and baselines. *2013 IEEE International Conference on Acoustics, Speech and Signal Processing* (2013), 126–130.
- E. Vincent, R. Gribonval, and C. Fevotte. 2006. Performance Measurement in Blind Audio Source Separation. *Trans. Audio, Speech and Lang. Proc.* 14, 4 (2006), 1462–1469.
- DeLiang Wang and Jitong Chen. 2017. Supervised Speech Separation Based on Deep Learning: An Overview. *CoRR* abs/1708.07524 (2017).
- Yuxuan Wang, Arun Narayanan, and DeLiang Wang. 2014. On training targets for supervised speech separation. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)* 22, 12 (2014), 1849–1858.
- Ziteng Wang, Xiaofei Wang, Xu Li, Qiang Fu, and Yonghong Yan. 2016. Oracle performance investigation of the ideal masks. In *IWAENC*.
- Felix Weninger, Hakan Erdogan, Shinji Watanabe, Emmanuel Vincent, Jonathan Le Roux, John R. Hershey, and Björn W. Schuller. 2015. Speech Enhancement with LSTM Recurrent Neural Networks and its Application to Noise-Robust ASR. In *LVA/ICA*.
- Dong Yu, Morten Kolbæk, Zheng-Hua Tan, and Jesper Jensen. 2017. Permutation invariant training of deep models for speaker-independent multi-talker speech separation. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2017), 241–245.
- Matthew D Zeiler and Rob Fergus. 2014. Visualizing and understanding convolutional networks. In *European conference on computer vision*. Springer, 818–833.
- Hang Zhao, Chuang Gan, Andrew Rouditchenko, Carl Vondrick, Josh McDermott, and Antonio Torralba. 2018. The Sound of Pixels. (2018).
- Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. 2014. Object detectors emerge in deep scene cnns. *arXiv preprint arXiv:1412.6856* (2014).

A OBJECTIVE METRICS USED FOR EVALUATING SEPARATION QUALITY

A.1 SDR

The signal-to-distortion ratio (SDR), introduced by Vincent et al. [2006], is one of a family of measures designed to evaluate Blind Audio Source Separation (BASS) algorithms, where the original source signals are available as ground truth. The measures are based on the decomposition of each estimated source signal into a true source part (s_{target}) plus error terms corresponding to interferences (e_{interf}), additive noise (e_{noise}) and algorithmic artifacts (e_{artif}).

SDR is the most general score, commonly reported for speech separation algorithms. It is measured in dB, and is defined as:

$$SDR := 10 \cdot \log_{10} \left(\frac{\|s_{target}\|^2}{\|e_{interf} + e_{noise} + e_{artif}\|^2} \right). \quad (2)$$

We refer the reader to the original paper for details on signal decomposition into its components. We found this measure to correlate well with the amount of noise left in the separated audio.

A.2 ViSQOL

The Virtual Speech Quality Objective Listener (ViSQOL) is an objective speech quality model, introduced by Hines et al. [2015]. The metric models human speech quality perception using a spectro-temporal measure of similarity between a reference (r) and a degraded (d) speech signal, and is based on the Neurogram Similarity Index Measure (NSIM) [Hines and Harte 2012]. NSIM is defined as

$$NSIM(r, d) = \frac{2\mu_r\mu_d + C_1}{\mu_r^2 + \mu_d^2 + C_1} \cdot \frac{\sigma_{rd} + C_2}{\sigma_r\sigma_d + C_2}, \quad (3)$$

where the μ s and σ s are mean and correlation coefficients, respectively, calculated between reference and degraded spectrograms.

In ViSQOL, NSIM is calculated on spectrogram patches of the reference signal and their corresponding patches from the degraded signal. The algorithm subsequently aggregates and translates the NSIM scores into a mean opinion score (MOS) between 1 and 5.