VoViT: Low Latency Graph-based Audio-Visual Voice Separation Transformer

Juan F. Montesinos[®], Venkatesh S. Kadandale[®], and Gloria Haro[®]

Universitat Pompeu Fabra, Carrer Roc Boronat, 138 08018 Barcelona, Spain {juanfelipe.montesinos,venkatesh.kadandale,gloria.haro}@upf.edu

Abstract. This paper presents an audio-visual approach for voice separation which produces state-of-the-art results at a low latency in two scenarios: speech and singing voice. The model is based on a two-stage network. Motion cues are obtained with a lightweight graph convolutional network that processes face landmarks. Then, both audio and motion features are fed to an audio-visual transformer which produces a fairly good estimation of the isolated target source. In a second stage, the predominant voice is enhanced with an audio-only network. We present different ablation studies and comparison to state-of-the-art methods. Finally, we explore the transferability of models trained for speech separation in the task of singing voice separation. The demos, code, and weights are available in https://ipcv.github.io/VoViT/.

Keywords: Audio-visual, source separation, speech, singing voice.

1 Introduction

Human voice is usually found together with other sounds. Think of people speaking in a cafeteria or in a social gathering, a journalist reporting on the scene, or an artist singing on a stage. In these situations we can find: multiple concurrent speeches, speech with background noise or a single or multiple singing voices with music accompaniment among others. Our brain is capable of understanding and concentrating on the voice of interest [3]. This cognitive process does not only rely on the hearing. Some works have shown the sight helps to focus on the voice of interest [12] or to resolve ambiguities in a noisy environment [20]. In this paper we address the voice separation and enhancement problems from a multimodal perspective, leveraging the motion information extracted from the visual stream to guide the resolution of the problem.

We propose an audio-visual (AV) voice separation model that produces stateof-the-art results. It is based on a two-stage approach. The first stage estimates a fairly good separation by combining audio and motion features with a transformer. Motion cues are crucial when the sound mixture contains different predominant voices. We extract those cues with a graph convolutional network (CNN) that processes a sequence of face landmarks. The audio-visual features are aligned in the feature dimension and preserve the time resolution. They are

processed by a multimodal spectro-temporal transformer that estimates the isolated voice corresponding to the target face landmarks. In a second stage, the predominant voice is enhanced by a small audio-only U-Net that takes as input just the pre-estimated audio. The voice of interest is predominant in the first estimation and thus an audio-only network is capable of modelling it and cancelling the sparse and mild interferences present in the pre-estimation. The paper includes an ablation study of different configurations of the multimodal transformer, its number of blocks and design of the lead voice enhancer network. The proposed method is compared to state-of-the-art methods in two different scenarios: speech and singing voice separation, showing successful results in both cases.

The contributions of this work are several: i) We propose an audio-visual network based on a transformer which performs better than current state-of-the-art models in speech and singing voice separation. ii) We show that a landmarkbased approach for extracting motion information can be a lightweight competitive alternative to processing raw video frames. iii) We show how an enhancement stage based on a light network can boost the performance of AV models over larger complex models, reducing the computational cost and the required time for training. iv) We reveal that AV models trained in speech separation do not generalise good enough for the separation of singing voice because of the different voice characteristics in each case and that a dedicated training with singing voice examples clearly boosts the results. Finally, v) our method is an end-toend gpu-powered system which is capable of isolating a target voice in real time (including the pre-processing steps).

2 Related work

In the last years there has been a fast evolution of deep-learning-based audiovisual works for speech separation and enhancement (we refer the reader to a recent review in [23]).

Back in 2016, we can find one of the first works in exploiting visual features for speech enhancement [37]. In this work, the authors proposed a CNN to process the visual signal and a fully connected layer to process the raw waveforms. Both modalities were fused by a BiLSTM network. This network had approximately 3M parameters (M for millions), far from the 80M of the most recent work [10]. A two-tower stream for processing audio and video features and then fused with a BiLSTM module that predicted complex masks was proposed in [7].

A two-step enhancement process was proposed in [1]. In the first step, a two-tower stream processed the audio-visual information to extract a binary mask that performed separation on the magnitude spectrogram. Afterwards, the phase of the spectrogram was predicted by passing the estimated magnitude spectrogram together with the noisy phase spectrogram through a 1D-CNN. A similar idea was developed in [8], where a two-tower stream encoder generated an embedding of audio-visual features from which the enhanced speech spectrogram was recovered. On the other hand, in [17] not only the enhanced spectrogram was reconstructed but the input frames as well.

New approaches and explorations different from the two-tower CNNs appeared recently. Variational auto-encoders [28] for speech enhancement joined the scene. Concurrently, [36] developed a time-domain model for speech separation, in contrast to most of the works which usually posed the problem in the time-frequency domain. Multi-channel audio-visual speech separation was addressed in [14] in a four-tower stream fashion. The mixture spectrogram was constrained with directional features from the visual stream of the speaker. A temporal CNN extracted visual features from the lips motion. The audio and visual embeddings were concatenated together with a speaker embedding extracted from the clean audio(s). A different mechanism was used in [19,29,31], where the audio-visual fusion was done with an attention module: or in [38], where the system was trained in a GAN manner so that the discriminator modeled the distribution of the clean speech signals. Transformers have been used in audio-only source separation [40]. Very recently, audio-visual transformers were investigated in [32] for main speaker localization and separation of its corresponding audio. In [33] an audio-visual transformer was used for classification in order to guide an unsupervised source separation model. Finally, in [2] a transformer was used for audio-visual synchronisation.

Another interesting proposal is [4], where the authors were concerned about the extra computational cost of processing the visual features and the possible privacy problems arised from it. On the other hand, to our knowledge, there are only two works using face landmarks, instead of video frames, for source separation. In [25] they process face landmarks with fully connected layers and then use BiLSTMs to predict the masks for the target source. In [24] a U-Net conditioned by a graph convolutional network that processed face landmarks was used for audio-visual singing voice separation. The work in [22] compared different training targets and loss functions for audio-visual speech enhancement.

Most recent algorithms made use of lips motion as well as appearance information, usually implementing cross-modal losses to pull together corresponding audio-visual features [10,21].

3 Approach

In audio-visual voice separation, given an audio-visual recording with several speaking/singing faces, and other sound sources, the goal is to recover their isolated voices by guiding the voice separation with the visual information present in the video frames. More formally, given the audio signal of each speaker, $s_i(t)$ (where t denotes time), the mixture of sounds can be defined as $x(t) = \sum_i s_i(t) + n(t)$ where n(t) denotes any other sound present in the mixture, i.e. background sounds. Therefore, the task of interest can be defined as the estimation of each individual voice $\hat{s}_i(t)$. In our approach $\hat{s}_i(t) = F(x(t), v_i(t))$, where F is a function represented by a neural network. The network receives the visual information of the speaker of interest, $v_i(t)$, and estimates its isolated voice $\hat{s}_i(t)$.

3.1 The AV Voice Separation Network

Our solution comprises of a two-stage neural network that operates in the timefrequency domain. The first stage consists of an AV voice separation network which can isolate the target voice at a good quality. However, this network is the most demanding one in terms of computational cost. To alleviate this, we propose to use downsampled spectrograms in this stage. The second stage consists of a recursive lead voice enhancer network that works with full resolution spectrograms. In Section 5.2, we experimentally show that this two-stage design leads to a higher performance than using larger AV models. To achieve this modularity, the networks at both stages are trained independently. The whole model is presented in Fig.1.



Fig. 1. Audio-visual voice separation network. Audio and video features are concatenated in the channel dimension before being fed to the transformer.

Stage 1: Audio-Visual Voice Separation. For simplicity, we seek to isolate the voice (denoted by s(t) and its corresponding spectrogram S(f,t)) corresponding to a single face at a time. The audio waveform of the mixture, x(t), is transformed into a complex spectrogram X(f,t) applying the Short-Time Fourier Transform (STFT). Once the waveform is mapped to the time-frequency domain, we can define a complex mask M(f,t) that allows to recover the spectrogram of the estimated source with a complex product, denoted as *, that is: S(f,t) = X(f,t) * M(f,t) Then, the goal of the network in the first stage is to estimate the complex mask $\hat{M}(f,t)$. The optimal set of parameters of the network is found by minimising the following loss:

$$\mathcal{L}_1 = \|G \odot (M_b - \hat{M}_b)\|^2$$

where M_b and \hat{M}_b are, respectively, the ground truth and estimated bounded complex masks, \odot denotes the element-wise product, $\|\cdot\|$ is the L2-norm and Gis a gradient penalty term which weights the time-frequency points of the mask according to the energy of the analogous point in the mixture spectrogram X:

$$G(f,t;X) = \max(\min(\log(1 + ||X(f,t)||), 10), 10^{-3}).$$
(1)

Note that, by definition, the ground truth mask M is not bounded. In order to stabilise the training, we bound the complex masks by applying a hyperbolic tangent [35]: $M_b = \tanh M^r + i \tanh M^i$, where M^r and M^i , denote the real and imaginary parts, respectively. The audio waveform of the estimated source can be computed through the inverse STFT of the estimated spectrogram $\hat{S}(f,t) = X(f,t) * \hat{M}(f,t)$.

To solve the AV voice separation problem, we propose to leverage the face motion information present in the video frames of the target person whose voice we want to isolate. For that, we use a spatio-temporal graph neural network that processes the face landmarks to generate motion features. On the other hand, the audio features are generated by a CNN encoder, denoted as *Spec2vec*. Both audio and motion features preserve the temporal resolution and are concatenated in the channel dimension, then they are fed into a transformer. All the submodules have been carefully designed to achieve a high-performance low-latency neural network.

Spatio-temporal graph CNN: Many AV speech separation or enhancement methods rely on lips motion extracted from raw video frames to guide the task. To reduce the computational cost of the visual stream, we propose to use face landmarks together with a spatio-temporal graph CNN [39]. This network, similar to that in [24], was redesigned to preserve the temporal resolution. It consists of a set of blocks which apply a graph convolution over the spatial dimension followed by a temporal convolution. This way we can considerably reduce the amount of data to process and to store, from $96 \times 96 \times 3 \approx 3 \cdot 10^4$ values per frame to $68 \times 2 \approx 10^2$. This supposes a substantial reduction in the storage necessities when working with large audio-visual datasets. For example, *Voxceleb2*'s grayscale ROIs occupy 1Tb, the raw uncompressed dataset occupies several Tb while storing face landmarks only requires 70 Gb.

Spec2vec: It is well known that transformers need proper embeddings to achieve high performance. We use the audio encoder of [7] to generate embeddings without losing temporal resolution.

AV spectro-temporal transformer: The traditional AV source separation methods comprise of a two-tower stream architecture. We can find two major variants: either encoder-decoder CNNs (usually with a U-Net as backbone) (e.g. [10,9,24,30,42,41]) or recurrent neural networks (RNNs), both conditioned on visual features (e.g. [7,25,37]). The major drawback of the latter is that RNNs are sequential, introducing bottlenecks in the processing pipeline. Transformers appeared as an efficient solution, reaching the same performance than RNNs and CNNs in large datasets. They are trained with a masking system allowing to process all the timesteps of a sequence in parallel. However, these architectures operate sequentially at the time of inference, like the RNNs. To overcome this issue we use an encoder-decoder transformer, which can solve the source separation problem in a single forward pass.

Transformers were originally designed to work with two unimodal signals. We study three different possible configurations for the transformer. The first proposal is to use the transformer as an auto-encoder, being fed with an audio-

visual signal directly. This way we ease the task for the transformer as audio and visual features are temporally aligned by construction. Then, it just has to find relationships through the multi-head self-attention. The second proposal is to pass visual features to the encoder and audio features (from the mixture) to the decoder so that the network can find audio-visual interdependencies via multi-head attention. Nevertheless, we hypothesise the dependencies between video and audio are local as audio events mostly occur at the same time than visual events. Lastly, we feed the encoder with an audio-visual signal and the decoder with the ground-truth separated audio. Note that this model is slower than previous ones as the model runs recurrently at inference time, going from a time complexity of $\mathcal{O}(n)$ to $\mathcal{O}(n^2)$ where *n* is the length of the sequence. From the ablation study in Section 5.1 and Table 1, we conclude that the best model is the first one, i.e. the one that uses an audio-visual signal as input, we denote it as AV ST-transformer.



Fig. 2. Three proposed ways to feed a transformer with an audio-visual signal. Left: audio-visual signal, middle: video to the encoder and audio mixture to the decoder, right: audio-visual signal to the encoder and clean audio to the decoder.

We design our AV ST-transformer encoder upon the findings of [40]. The AV ST-transformer has 512 model features across 8 heads. We tried 256 features but it works worse. The compression layer is nothing but a fully connected layer followed by GELU [16] activation which maps the C incoming channels to the 512 channels required by the architecture. It is composed by M encoders and M decoders. The encoder is a set of two traditional encoders in parallel, which processes the signal from a temporal and a spectral point of view [40].

Stage 2: Lead voice enhancer. Although lips motion is correlated with the voice signal and may help in source separation, it is not always accessible or reliable. For example, the scenarios involving a side view of the speaker or a partial occlusion of the face or an out-of-sync audio-visual pair make it challeng-

ing to incorporate the lips motion information in a useful way; all such scenarios may appear in unconstrained video recordings. In [24], the authors show that audio-only models tend to predict the predominant voice in a mixture when there is no prior information about the target speaker. Based on this idea, we hypothesise that, if the first stage of the AV voice separation network outputs a reasonable estimation of the target voice, this voice will be predominant in the estimation. Upon this idea, we use an audio-only network which identifies the predominant voice and enhances the estimation without relying on the motion, just on the pre-estimated audio. To do so, we simply use a small U-Net which takes as input the estimated magnitude spectrogram (at its original resolution) and returns a binary mask. The ground truth binary mask can be obtained from the ground truth spectrogram S and the spectogram to be refined, \hat{S} , which is the one estimated in the stage 1:

$$M(f,t) = \begin{cases} 1, & \text{if } ||S(f,t)|| \ge ||\hat{S}(f,t) - S(f,t)||, \\ 0, & \text{otherwise.} \end{cases}$$
(2)

Notice that the difference $\hat{S}(f,t) - S(f,t)$ are the remaining sources that need to be removed in the refinement stage.

There are different reasons to use binary masks. On the one hand, we found qualitatively, by inspecting the results, that the secondary speaker is often attenuated but not completely removed. In [13], the authors show that binary masks are particularly good at reducing interferences. On the other hand, complex masks appeared as an evolution of binary masks and ratio masks, as a way of estimating, not only the magnitude spectrogram, but the phase too. Note that these masking systems usually reconstruct the estimated waveform with the phase of the mixture as they estimate the magnitude only. In our case, the phase has already been estimated by using complex masks in the previous stage. Lastly, by using binary masks, we are changing the optimisation problem and easing the task since it is simpler to take a binary decision than orienting and modulating a vector.

Note that this refinement network can run recursively, although we empirically found (see Table 2) that applying the refinement network once leads to the best results in terms of SDR and a considerable boost in SIR. Further iterations reduce the interferences (at a lesser extent) but at the cost of introducing more distortion.

Let us denote by \hat{M} the binary mask estimated by the lead voice enhancer network. We trained this network to optimise a weighted binary cross entropy loss:

$$\mathcal{L}_{2} = \sum_{f=1}^{F} \sum_{t=1}^{T} \frac{G(f,t;\hat{S})}{FT} \left(M(f,t) \log \|\hat{M}(f,t)\| + (1 - M(f,t))(1 - \log \|\hat{M}(f,t)\|) \right)$$

where the weights G are defined in (1).

3.2 Low-latency data pre-processing

Many audio-visual works rely on expensive pipelines to pre-process data, which makes the proposed systems unusable in a real-world scenario unless a great amount of time is invested in optimisation. Pursuing the real applicability of our model, we curated an end-to-end gpu-powered system which can pre-process (from raw audio and video) and isolate the target voice of 10s of recordings in less than 100ms using floating-point 32 precision, and in less than 50 ms using floating-point 16 precision.

Face landmarks: The most common approach in speech separation is to align the faces in the different frames via 2D face landmark estimation together with image warping (e.g. [10,24]). This step removes eventual head motions. In order to achieve real-time audio-visual source separation, we estimate the 3D face landmarks using an optimised version of [15] and an aligned frontal view by applying a rigid transformation, skipping the image warping step. This optimised preprocessing takes around 10 ms to process 10s of video. Thanks to the 3D information, we can recover lips motion from side views by estimating 3D landmarks, as shown in Fig. 3. To do the registration, we use the Kabsch algorithm [18]. Finally, we drop the depth coordinate and consider just the first two spatial coordinates in the nodes of the graph.

Audio: Waveforms are re-sampled to 16384 Hz. Then, we compute the STFT with a window size of 1022 and a hop length of 256. This leads to a $512 \times 64 n$ complex spectrogram where n is the duration of the waveform in seconds. To reduce the computational cost of both training and inference we downsample the spectrogram in the frequency dimension by 2 in Stage 1.



Fig. 3. Frame example from *Voxceleb2* [5] with partial occlusions. Thanks to the landmark estimation together with the registration we can estimate the unoccluded lips.

4 Datasets

Experiments are carried out in two different datasets: Voxceleb2 [5], a dataset of celebrities speaking in a broad range of scenarios; and Acappella [24], a dataset of solo-singing videos. Both datasets are a collection of YouTube recordings which are publicly available. We also consider Audioset [11] and MUSDB18 [27] for

sampling extra audio sources that can be added to the singing voice signal as accompaniment.

Voxceleb2 contains 1 million utterances, most of them of a duration between 4 and 6 seconds, consisting of celebrities covering a wide range of ethnicities, professions and ages. The dataset is formed by in-the-wild videos that include several challenging scenarios, such as: different lightning, side-face views, motion blur and poor image quality. They also span across different scenarios like red carpets, stadiums, public speeches, etc. The dataset provides a test set which contains both, seen-heard and unseen-unheard speakers together. From this test set we selected the unseen-unheard samples and curated two different subsets. The first one, denoted as unheard-unseen wild test set consists of 1,000 samples randomly selected, reflecting the aforementioned challenges. The second one, denoted as unheard-unseen clean test set, is a subset of 1,000 samples, from which 500 of them have a high-quality content with the following characteristics: frontal or almost frontal point of view, low background noise and perceptual image quality above the average of the dataset. The samples were selected manually from the whole unseen-unheard test set, trying to include as many different speakers as possible. The target voice is sampled from the subset of 500 highquality videos in the *clean set*, while the second voice is sampled from the rest of 500 videos. This way we ensure that the video content is good enough to estimate motion features from it and that the ground truth separated audio is reliable, in the sense that it does not contain background sounds that may produce unfounded performance metrics.

Acappella is a 46-hours dataset of a cappella solo singing videos. The videos are divided in four language categories: English, Spanish, Hindi and others. These videos are recorded in a frontal view with no occlussions. It also provides two test sets: the seen-heard test set and the unseen-unheard test set. The former contains videos sampled from the same singers and in the same languages than the training set, whereas the latter contains recordings sampled from new singers in the four language categories plus some new languages. In the test set all the categories are equally represented across languages and gender. This way the algorithms can be tested in challenging real-world scenarios.

Audioset [11] is an in-the-wild large-scale dataset of audio events across more than 600 categories. We gathered the categories related to the human voice and some typical accompaniments. These categories are: acappella, background music, beatboxing, choir, drum, lullaby, rapping, theremin, whistling and yodelling.

Finally, MUSDB18[27] is an audio-only dataset of 150 full-track songs of different styles that includes original sound sources.

5 Experiments

The experiments were carried out in a single RTX 3090 GPU. Each experiment takes around 20 days of training. We used SGD with 0.8 momentum, 10^{-5} weight decay and a learning rate of 0.01. The metrics used for comparing results are Source-to-Distortion Ratio (SDR) and Source-to-Interferences Ratio (SIR) [34].

5.1 Audio-visual transformer

In this experiment, we compare three different versions of the transformer (shown in Fig. 2) in the *Acappella* dataset. The goal is two-fold: i) Compare the proposed architecture against the state-of-the-art model in singing voice separation [24]; and ii) compare the performance of different transformers for the task of singing voice separation.

For the sake of comparison, we train our models the same way as in [24]. In short, we create artificial mixtures of 4s of duration by mixing a voice sample from *Acappella* together with an accompaniment sample sourced either from *Audioset* or *MUSDB18*. Additionally, a second voice sample from *Acappella* is added 50% of the times. This results in mixtures that contain one or more voices plus musical accompaniment. For this dataset we take 4s audio excerpts and the corresponding 100 video frames from which we extract the face landmarks.

Results are shown in Table 2. From the ablation on the three versions of the transformer, we can conclude that the AV ST-transformer is the best model in terms of both performance and time complexity. Moreover, it can be observed that the three versions of the transformer greatly outperform the results of [24] in terms of SDR, while the AV ST-transformer also outperforms in SIR.

5.2 Speech separation

In Section 5.1 we found the AV ST-transformer was the best model in terms of time complexity and performance. All the remaining experiments will be carried out with this model. Now we consider the task of AV speech separation and work with *Voxceleb2* dataset. We use 2s audio excerpts which correspond to 50 video frames from which we extracted their face landmarks. In this case, we mix two voice samples from *Voxceleb2* which are normalised with respect to their absolute maximum, so that a mixture is $x(t) = (s_1(t) + s_2(t))/2$. This normalisation aims to have two voices which are codominant in the mixture and that the waveforms of the mixtures are bounded between -1 and 1. Note that the former characteristic is not always true as *Voxceleb2* samples are sometimes accompanied by other voices or sorts of interference (clapping, music, etc.). As *Voxceleb2* is a large-scale dataset, and for the sake of comparison, we extended the size of the AV ST-transformer up to 10 encoder blocks and 10 decoder blocks so that the number of parameters of the audio subnetwork is comparable to that

Model |Y-Net [24] AV ST-transformer |V_A transformer |AV_A transformer

$SDR \uparrow$	6.41	10.63 ± 5.86	8.64 ± 5.89	9.98 ± 5.70
$SIR \uparrow$	17.38	$17.67~\pm~7.73$	14.70 ± 7.88	16.11 ± 7.42

Table 1. Ablation study: performance of different ways of feeding a transformer with an audio-visual signal and comparison to Y-Net model [24]. Evaluated in *Acappella's* unseen-unheard test set. Y-Net metrics taken from *Acappella*. In this table N = 4 (the number of blocks in the transformers) in order to adapt the number of parameters to the size of *Acappella* dataset.

of Visual Voice [10]. We tested the performance of each model in the *unheard-unseen wild* test set and in the *unheard-unseen clean* test set (both described in Section 4). For each test set we randomly made 500 pairs out of the 1,000 samples, ensuring no sample is used more than once.

Lead Voice Enhancer. The first experiment is an ablation designed to address three main questions. i) Compare two different versions of the lead voice enhancer: the audio backbone of Y-Net [24], which is a 7M-parameter U-Net; and the audio backbone of Visual Voice [10], yet another U-Net but with 50M parameters because of a different design. ii) Evaluate the effect of recurrent iterations of the lead voice enhancer. And iii) comparing the results of the 10block 2-stage AV ST-Transformer against a 18-block 1-stage AV ST-Transformer transformer. The details of this subnetwork are explained in Section 3.1. We denote our Voice-Visual Transformer as VoViT (the whole network with two stages) and VoViT-s1 the network without the second stage.

The results are shown in Table 2. As we can see, the refinement network improves the results substantially for the 10-block AV ST-Transformer. Successive iterations of the refinement module further reduce the interferences, but the best SDR is achieved with just one iteration. For the lead voice enhancer, we tried two possible audio-only U-Nets: the U-Net from the Y-Net model [24] and the larger U-Net from Visual Voice [10]. A much larger U-Net does not outperform the smaller one by a large margin. Interestingly, we can observe that adding this module performs better than using the 18-block AV ST-transformer (with around 2 times more parameters). Moreover, this subnetwork can be trained within a day, whereas the 18-block transformer required around a month to train. The reasons behind the lack of improvement of the 18-block transformer are unknown. We observed a phenomena similar to the so called "double descent" [26] while training the 10-block transformer, which may be indicative of a complex optimisation process which is worsened in the 18-block case exceeding our computational resources. In the same line, we trained a larger graph convolutional network, comparable in number of parameters to the motion subnetwork of Visual Voice, however the performance dropped. From this ablation, we can conclude that a 10-block AV ST-transformer with a small U-Net as lead voice enhancer is the best option in terms of performance-latency trade-off.

Comparison to state-of-the-art methods. Next we are going to compare the 10-block AV ST-Transformer to a state-of-the-art AV speech separation model and audio baselines in the *Voxceleb2* dataset. The Visual Voice network [10] is the current state of the art in speech separation. This network uses 2.55s excerpts, the corresponding 64 video frames cropped around the lips and an image of the whole face of the target speaker. Apart from using lips motion features, it extracts cross-modal face-voice embeddings that complement the motion features and are especially useful when the motion is not reliable or when the appearance of the speakers is different. We also compare the results against Y-Net [24] as it is one of the few papers proposing face landmarks. The original work uses 4s excerpts. As around 160k samples for *Voxceleb2* are shorter, we just adapted the model for working with 2s samples.

12 Juan F. Montesinos et al.

		Wild t	est set
		$\text{SDR}\uparrow$	SIR \uparrow
	VoViT-s1	9.68	15.75
çk	VoViT (VV in stage 2, $r = 1$)	10.05	18.30
blc	VoViT (VV in stage 2, $r = 2$)	9.77	19.38
10-	VoViT (YN in stage 2, $r = 1$)	10.03	18.18
	VoViT (YN in stage 2, $r = 2$)	9.78	19.09
	18-block VoViT-s1	9.27	15.53

Table 2. Ablation of different variants of the refinement stage and number of blocks in the transformer of the first stage. VoViT-s1 stands for the model with just the first stage, r stands for the number of recurrent passes in stage 2. For the stage 2 we considerered both, the Visual Voice's UNet (VV) [10] and the Y-Net's UNet (YN) [24].

	$\#\ {\rm parameters}$		Wild Test set		Clean Test set	
	Visual	Whole	SDR ↑	SIR ↑	SDR ↑	SIR ↑
	Net.	Net.	SDIC	SIIC	SDR	SIIt
Visual Voice Audio-only	-	46.14	7.7	13.6	-	-
Face Filter [6]	_	_	2.53	-	_	_
The conversation [1]	_	_	8.89	14.8	-	-
Visual Voice Motion-only	9.14	55.28	9.94	17	-	-
Y-Net [24]	1.42	9.7	5.29 ± 5.06	8.45 ± 6.8	5.86 ± 4.78	9.25 ± 6.44
Visual Voice [10]	20.38	77.75	9.92 ± 3.56	16.11 ± 4.8	10.18 ± 3.36	16.49 ± 4.5
VoViT	1.42	58.2	10.03 ± 3.35	18.18 ± 4.72	10.25 ± 2.61	18.65 ± 3.8

Table 3. Evaluation on *Voxceleb2* unheard-unseen test sets (mean \pm standard deviation). VoViT stands for our model with the 10-block AV ST-Transformer with the Y-Net's UNet backbone as the lead voice enhancer. Number of parameters in millions. Results in the first block are taken from the original papers.

Numerical results are shown in Table 3. The 10-block VoViT outperforms all the previous AV speech separation models. Compared to Visual Voice, it achieves a much better SIR and slightly better SDR, both for the wild and clean test sets. In particular, for the clean test set, when the motion cues are more reliable, our model has a much lower standard deviation. Some aspects need to be taken into account:

- The face landmark extractor has been trained with higher quality videos than the ones in *Voxceleb2*. On the contrary, the Visual Voice video network has been trained specifically for *Voxceleb2*.

- Our visual subnetwork, the graph CNN, has 10 times less parameters than its counterpart in Visual Voice.

- Apart from motion cues, Visual Voice takes also into account speaker appearance features which are correlated with voice features, and which can be crucial in poor quality videos where lip motion is unreliable.

Fig. 4 shows SDR and SIR differences between VoViT and Visual Voice in two different test sets: the *wild* and the *clean* set. Each plot is a scatter plot where each point corresponds to a 2s long mixture. As it can be observed, our method especially outperforms Visual Voice in SIR while in SDR both methods



Fig. 4. Scatter plot showing the difference in SDR and SIR, ΔSDR and ΔSIR , as functions of the SDR and SIR of the input mixture in the unseen-unheard wild and clean test sets. The difference is: $\Delta SDR = SDR(VoViT) - SDR(Visual Voice)$ so a positive value means VoViT outperforms Visual Voice.

have a comparable performance. In order to assess the significance of the results of Table 3, we calculated the *p*-values with respect to the Visual Voice results. Only the improvement on SIR is significant (p < 0.05). While the improvement from stage 1 to 2 (Table 2) is significant both in SDR and SIR. In the *wild test set* there are a few samples where our model performs worse than Visual Voice. Those correspond to samples where the audio and video are extremely unsynchronised or samples where the lip motion is mispredicted and the network separates the other speaker. In those cases, the Visual Voice model might be able to alleviate the situation either by relying on the appearance features to guide the separation or by using the motion information present in the raw video despite its poor quality (e.g. blur, compression artefacts, lack of sharpness). There are no such cases in the *clean set*, as those type of samples were filtered out. Audiovisual files with the top K worst performing examples and demos for both, Visual Voice and VoViT models, are provided in the supplementary material.

5.3 Singing voice separation

In this last experiment we consider the task of singing voice. We are interested in exploring how transferable models trained for speech separation are to the case of singing voice. Since speech models were trained with two voices and no extra sounds and in *Voxceleb2*, which contains mainly English, we restricted to similar types of mixtures in singing voice. In particular, we create mixtures of

two singers in English from the unseen-unheard test set of *Acappella*, with no accompaniment. Table 4 compares the results of models trained directly with samples of singing voice (top block of results in Table 4) versus models trained with speech samples (bottom block). In the case of singing voice we used our model with just the first stage and a 4-block AV ST-transformer. We observe that dedicated models for singing voice perform largely better than models trained for speech. This may be explained to particular differences between a speaking and a singing voice. For example, vowels are much more sustained in singing voice, there is much less coarticulation of consonants with surrounding vowels and vibrato is not present in speech. Moreover, singing voice contains varying pitches covering a wider frequency range.

Model	$\mathrm{SDR}\uparrow$	$\mathrm{SIR}\uparrow$
Y-Net [24]	11.08 ± 7.51	17.18 ± 9.68
VoViT-s1 (4 blocks)	$14.85~\pm~7.87$	$\textbf{21.06} \pm \textbf{9.69}$
VoViT-s1	3.89 ± 9.28	5.89 ± 11.15
VoViT	4.04 ± 10.30	$\textbf{7.21} \pm \textbf{13.26}$
Visual Voice [10]	$\textbf{4.52} \pm \textbf{8.64}$	7.03 ± 7.11

Table 4. Singing voice separation. Mixtures of two singers with no additional accompaniment from the test set unseen-unheard (only samples in English) of *Acappella*. Results in top block: models trained directly with samples of singing voice; bottom block: models trained with speech samples.

6 Conclusions & Future work

In this work we present a lightweight audio-visual source separation method which can process 10s of recordings in less than 0.1s in an end-to-end GPU powered manner. Besides, the method shows competitive results to the state-of-theart in reducing distortions while clearly outperforming in reducing interferences. We show that face landmarks are computationally cheaper alternatives to raw video and help to deal with large-scale datasets. For the first time, we evaluate AV speech separation systems in singing voice, showing empirically that the characteristics of the singing voice differ substantially from the ones of speech.

As future work we would like to explore lighter and faster embedding generators for the transformer and different optimisations in its architecture which leads to a fast and powerful system.

Acknowledgments. We acknowledge support by MICINN/FEDER UE project PID2021-127643NB-I00; H2020-MSCA-RISE-2017 project 777826 NoMADS. J.F.M. acknowledges support by FPI scholarship PRE2018-083920.We acknowledge NVIDIA Corporation for the donation of GPUs used for the experiments.

References

- 1. Afouras, T., Chung, J.S., Zisserman, A.: The conversation: Deep audio-visual speech enhancement. In: Interspeech (2018)
- Chen, H., Xie, W., Afouras, T., Nagrani, A., Vedaldi, A., Zisserman, A.: Audiovisual synchronisation in the wild. In: 32nd British Machine Vision Conference, BMVC (2021)
- 3. Cherry, E.C.: Some experiments on the recognition of speech, with one and with two ears. The Journal of the acoustical society of America **25**(5), 975–979 (1953)
- Chuang, S.Y., Tsao, Y., Lo, C.C., Wang, H.M.: Lite audio-visual speech enhancement. In: Proc. Interspeech 2020
- Chung, J.S., Nagrani, A., Zisserman, A.: Voxceleb2: Deep speaker recognition. In: Interspeech (2018)
- Chung, S.W., Choe, S., Chung, J.S., Kang, H.G.: Facefilter: Audio-visual speech separation using still images. arXiv preprint arXiv:2005.07074 (2020)
- Ephrat, A., Mosseri, I., Lang, O., Dekel, T., Wilson, K., Hassidim, A., Freeman, W.T., Rubinstein, M.: Looking to listen at the cocktail party: A speakerindependent audio-visual model for speech separation. In: SIGGRAPH (2018)
- Gabbay, A., Shamir, A., Peleg, S.: Visual speech enhancement. In: Interspeech. pp. 1170–1174. ISCA (2018)
- Gao, R., Grauman, K.: Co-separating sounds of visual objects. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3879–3888 (2019)
- Gao, R., Grauman, K.: Visualvoice: Audio-visual speech separation with crossmodal consistency. In: CVPR (2021)
- Gemmeke, J.F., Ellis, D.P., Freedman, D., Jansen, A., Lawrence, W., Moore, R.C., Plakal, M., Ritter, M.: Audio set: An ontology and human-labeled dataset for audio events. In: IEEE Int. Conf. on Acoustics, Speech and Signal Processing (2017)
- Golumbic, E.Z., Cogan, G.B., Schroeder, C.E., Poeppel, D.: Visual input enhances selective speech envelope tracking in auditory cortex at a "cocktail party". The Journal of Neuroscience 33, 1417 – 1426 (2013)
- Grais, E.M., Roma, G., Simpson, A.J., Plumbley, M.: Combining mask estimates for single channel audio source separation using deep neural networks. In: Interspeech (2016)
- Gu, R., Zhang, S.X., Xu, Y., Chen, L., Zou, Y., Yu, D.: Multi-modal multi-channel target speech separation. IEEE Journal of Selected Topics in Signal Processing 14(3), 530–541 (2020)
- Guo, J., Zhu, X., Yang, Y., Yang, F., Lei, Z., Li, S.Z.: Towards fast, accurate and stable 3d dense face alignment. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK. pp. 152–168 (2020)
- Hendrycks, D., Gimpel, K.: Gaussian error linear units (gelus). arXiv preprint arXiv:1606.08415 (2016)
- Hou, J.C., Wang, S.S., Lai, Y.H., Tsao, Y., Chang, H.W., Wang, H.M.: Audiovisual speech enhancement using multimodal deep convolutional neural networks. IEEE Transactions on Emerging Topics in Computational Intelligence 2(2), 117– 128 (2018)
- Kabsch, W.: A discussion of the solution for the best rotation to relate two sets of vectors. Acta Crystallographica Section A 34(5), 827–828 (1978)
- 19. Li, C., Qian, Y.: Deep audio-visual speech separation with attention mechanism. In: ICASSP 2020 2020 IEEE International Conference on

Acoustics, Speech and Signal Processing (ICASSP). pp. 7314–7318 (2020). https://doi.org/10.1109/ICASSP40776.2020.9054180

- Ma, W.J., Zhou, X., Ross, L.A., Foxe, J.J., Parra, L.C.: Lip-reading aids word recognition most in moderate noise: a bayesian explanation using high-dimensional feature space. PloS one 4(3), e4638 (2009)
- Makishima, N., Ihori, M., Takashima, A., Tanaka, T., Orihashi, S., Masumura, R.: Audio-visual speech separation using cross-modal correspondence loss. In: ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 6673–6677. IEEE (2021)
- Michelsanti, D., Tan, Z.H., Sigurdsson, S., Jensen, J.: On training targets and objective functions for deep-learning-based audio-visual speech enhancement. In: ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 8077–8081. IEEE (2019)
- Michelsanti, D., Tan, Z.H., Zhang, S.X., Xu, Y., Yu, M., Yu, D., Jensen, J.: An overview of deep-learning-based audio-visual speech enhancement and separation. IEEE/ACM Transactions on Audio, Speech, and Language Processing 29, 1368– 1396. https://doi.org/10.1109/TASLP.2021.3066303
- Montesinos, J.F., Kadandale, V.S., Haro, G.: A cappella: Audio-visual singing voice separation. In: 32nd British Machine Vision Conference, BMVC (2021)
- Morrone, G., Bergamaschi, S., Pasa, L., Fadiga, L., Tikhanoff, V., Badino, L.: Face landmark-based speaker-independent audio-visual speech enhancement in multitalker environments. In: ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 6900–6904. IEEE (2019)
- Nakkiran, P., Kaplun, G., Bansal, Y., Yang, T., Barak, B., Sutskever, I.: Deep double descent: where bigger models and more data hurt. Journal of Statistical Mechanics: Theory and Experiment **2021** (2020)
- 27. Rafii, Z., Liutkus, A., Stöter, F.R., Mimilakis, S.I., Bittner, R.: The MUSDB18 corpus for music separation (Dec 2017). https://doi.org/10.5281/zenodo.1117372, https://doi.org/10.5281/zenodo.1117372
- Sadeghi, M., Alameda-Pineda, X.: Mixture of inference networks for vae-based audio-visual speech enhancement. IEEE Transactions on Signal Processing 69, 1899–1909 (2021)
- Sato, H., Ochiai, T., Kinoshita, K., Delcroix, M., Nakatani, T., Araki, S.: Multimodal attention fusion for target speaker extraction. In: 2021 IEEE Spoken Language Technology Workshop (SLT). pp. 778–784. IEEE (2021)
- Slizovskaia, O., Haro, G., Gómez, E.: Conditioned source separation for musical instrument performances. IEEE/ACM Transactions on Audio, Speech, and Language Processing 29, 2083–2095 (2021)
- Sun, Z., Wang, Y., Cao, L.: An attention based speaker-independent audio-visual deep learning model for speech enhancement. In: International Conference on Multimedia Modeling. pp. 722–728. Springer (2020)
- 32. Truong, T.D., Duong, C.N., Pham, H.A., Raj, B., Le, N., Luu, K., et al.: The right to talk: An audio-visual transformer approach. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1105–1114 (2021)
- 33. Tzinis, E., Wisdom, S., Remez, T., Hershey, J.R.: Improving on-screen sound separation for open-domain videos with audio-visual self-attention. arXiv preprint arXiv:2106.09669 (2021)
- Vincent, E., Gribonval, R., Fevotte, C.: Performance measurement in blind audio source separation. IEEE Transactions on Audio, Speech, and Language Processing 14(4), 1462–1469 (2006). https://doi.org/10.1109/TSA.2005.858005

- Williamson, D.S., Wang, Y., Wang, D.: Complex ratio masking for monaural speech separation. IEEE/ACM transactions on audio, speech, and language processing 24(3) (2015)
- Wu, J., Xu, Y., Zhang, S.X., Chen, L.W., Yu, M., Xie, L., Yu, D.: Time domain audio visual speech separation. In: 2019 IEEE automatic speech recognition and understanding workshop (ASRU). pp. 667–673. IEEE (2019)
- Wu, Z., Sivadas, S., Tan, Y.K., Bin, M., Goh, R.S.M.: Multi-modal hybrid deep neural network for speech enhancement. arXiv preprint arXiv:1606.04750 (2016)
- Xu, X., Wang, Y., Xu, D., Peng, Y., Zhang, C., Jia, J., Chen, B.: Vsegan: Visual speech enhancement generative adversarial network. arXiv preprint arXiv:2102.02599 (2021)
- Yan, S., Xiong, Y., Lin, D.: Spatial temporal graph convolutional networks for skeleton-based action recognition. In: Proceedings of the AAAI conference on artificial intelligence. vol. 32 (2018)
- 40. Zadeh, A., Ma, T., Poria, S., Morency, L.P.: Wildmix dataset and spectrotemporal transformer model for monoaural audio source separation. arXiv preprint arXiv:1911.09783 (2019)
- Zhao, H., Gan, C., Ma, W.C., Torralba, A.: The sound of motions. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1735–1744 (2019)
- 42. Zhao, H., Gan, C., Rouditchenko, A., Vondrick, C., McDermott, J., Torralba, A.: The sound of pixels. In: Proceedings of the European conference on computer vision (ECCV). pp. 570–586 (2018)