

# VOICEVECTOR: MULTIMODAL ENROLMENT VECTORS FOR SPEAKER SEPARATION

Akam Rahimi, Triantafyllos Afouras, Andrew Zisserman

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

## ABSTRACT

We present a transformer-based architecture for voice separation of a target speaker from multiple other speakers and ambient noise. We achieve this by using two separate neural networks: (A) An enrolment network designed to craft speaker-specific embeddings, exploiting various combinations of audio and visual modalities; and (B) A separation network that accepts both the noisy signal and enrolment vectors as inputs, outputting the clean signal of the target speaker. The novelities are: (i) the enrolment vector can be produced from: audio only, audio-visual data (using lip movements), or visual data alone (using lip movements from silent video); and (ii) the flexibility in conditioning the separation on multiple positive and negative enrolment vectors. We compare to previous methods and obtain superior performance.

*Index Terms*— Speech separation, speaker embedding

## 1. INTRODUCTION

Speech and audio processing applications, such as hearing aids, voice-activated systems, and video conferencing, hinge on voice separation performance, especially in noisy, multi-speaker environments. The complexity of separating voices is accentuated when confronted with diverse acoustic settings, ambient noise, and overlapping speech.

A notable set of methods represented by [1–3], address this challenge by conditioning the separation process on speaker embeddings that are derived from clean, noise-free audio. These embeddings capture the distinctive vocal characteristics of an individual, and while they have shown significant promise in controlled environments, their dependency on such clean audio can hinder their effectiveness in real-world, noise-rich scenarios where pre-recorded clean audio of the target speaker is not available.

On the other hand, a subset of methodologies, such as [4–6], harness visual cues, specifically synchronized lip movements, to enhance the voice separation task. These audio-visual methods, while effective, are not without limitations. Their optimal performance is often contingent on the uninterrupted availability of visual cues from the very recording

undergoing separation. This becomes a challenge, especially when visual data is occluded or missing.

In this paper, we present a novel method that marries the strengths of both audio and visual conditioning paradigms while efficiently sidestepping their inherent challenges. Our approach is structured around two distinct phases:

**Enrolment Phase:** In this phase, we harness auditory and/or visual information to craft ‘enrolment vectors’. These descriptors enable the separation of a target speaker’s voice in noisy settings. Although our method can derive these vectors from clean audio (Fig. 1a), they can equally be crafted from noisy audio paired with visual cues or visual cues alone (Fig. 1b). Notably, unlike other audio-visual methods, ours doesn’t mandate sourcing visual cues directly from the intended recording being separated. Furthermore, our approach uses not just positive examples from the target speaker, but also negative examples from other speakers, serving as a contrastive measure. This dual conditioning significantly improves the separation performance of the model.

**Separation Phase:** For this stage, we only use auditory data (Fig. 1c). By design, we circumvent the reliance on visual cues during separation, rendering our system both computationally efficient and robust against the variability or absence of visual information, thus addressing the limitation of prior audio-visual techniques.

In summary, our contributions include: (i) Introducing speaker-discriminative embeddings using a specialised enrolment network that is amenable to audio, audio-visual, and uniquely, video-only inputs. (ii) Conditioning voice separation with both positive and negative enrolment vectors.

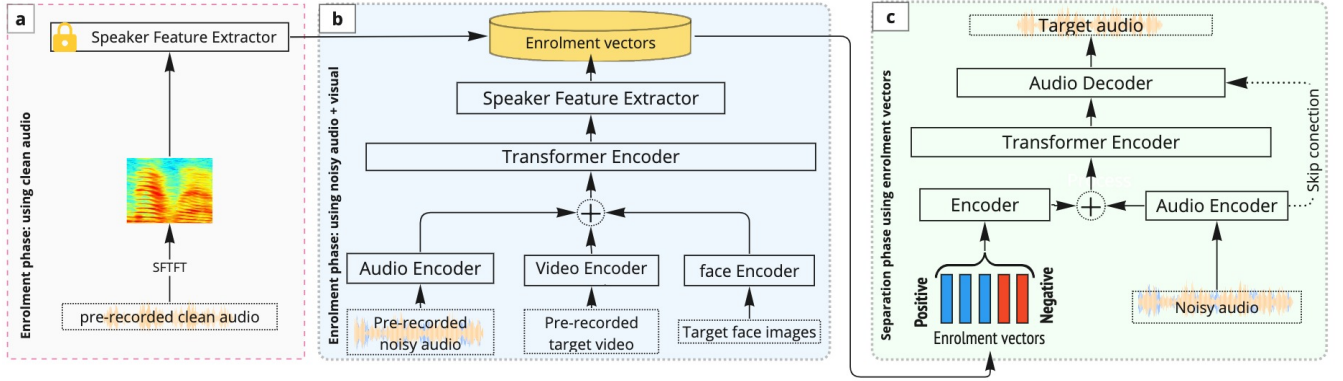
## 2. RELATED WORK

Advances in target speaker separation have predominantly used two paradigms: multi-modal conditioning, and speaker-specific embeddings [7]. Several methods represented by [4–6, 8–11], use visual cues, mainly lip movements, to inform separation. While these methods require these cues throughout the separation, our separation phase operates entirely without them, offering a significant advantage.

Turning to speaker-conditioned models, a common approach is to condition the separation on speaker embeddings generated from clean speech samples of target speakers [1–3, 12]. In a similar manner to our video-only (no audio) enrolment approach, face images of target speakers have also

---

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



**Fig. 1.** The VoiceVector Architecture Overview: There are three networks for two tasks: ‘enrolment vector’ generation networks (a,b); and a speaker separation network (c). In the enrolment phase, network **a** is used for generating enrolment vectors from clean audio, while network **b** accepts other data modalities, including: various combinations of lip-reading, facial images, or noisy audio paired with visual cues. In **b** a combination of available modalities is concatenated in time dimension and passed to a transformer encoder which aggregates the concatenated input. Notably, during training the speaker feature extractor in network **b** is trained to mirror the output of network **a** for separate recordings of the same speaker. In the separation phase **c**, a U-Net structure augmented with a transformer, processes noisy audio conditioned on a set of positive and negative enrolment vectors generated from **a** or **b**. Leveraging these vectors, the model separates the waveform of the target speaker.

been leveraged for speaker separation [13]. Previous research by [14] has used negative vectors for audio separation. In contrast to their approach, our model can generate the enrolment vectors from noisy audio and does not require distinct attention mechanisms for positive and negative vectors.

Attempts at target speaker separation using only audio [15, 16], are challenged by the permutation problem. While permutation invariant training [17, 18] provides partial alleviation, the inference stage reintroduces the problem due to ambiguous channel assignment.

### 3. METHOD

In this section, we describe our method, named ‘VoiceVec-tor’, for separating a target speaker from other simultaneous speakers and ambient noise. When presented with a noisy speech signal, the aim is to separate the component representing the distinct acoustic characteristics of a specified speaker, whilst concurrently filtering components such as competing voices or ambient noise. Refer to Fig. 1 for a graphical representation of the architecture.

#### 3.1. Architecture

The overall architecture consists of two separate but interconnected networks, one to generate enrolment vectors and the other tasked with speaker separation.

##### 3.1.1. Speaker enrolment network.

We have devised three distinctive configurations of the enrolment network, each tailored for specific input modalities:

**Audio-Only Network (Fig 1(a)):** This configuration makes use of an off-the-shelf speaker feature extractor network named ECAPA-TDNN [19]. It processes clean audio  $a_c \in \mathbb{R}^{t_a}$  represented in the spectrogram format,  $S(f,t) = STFT(a_c)$ , and outputs speaker embeddings,  $S_{a_c} \in \mathbb{R}^{192}$ . Here  $STFT$  denotes the Short-Time Fourier Transform.

**Audio-Visual Network (Fig 1(b)):** This network generates speaker embeddings from noisy audio by leveraging asso-

ciated visual elements. This includes video footage capturing the lip motions and/or still images of the target speaker. The noisy audio waveforms  $a \in \mathbb{R}^{t_a}$  are encoded as  $E_a \in \mathbb{R}^{t_a \times 768}$ , and concurrently the video stream (lip motion)  $v \in \mathbb{R}^{3 \times t_v \times H \times W}$  is encoded as video embeddings  $E_v \in \mathbb{R}^{t_v \times 512}$ . Optionally, this can be further enriched by introducing static face images  $f \in \mathbb{R}^{3 \times H \times W}$  of the target speaker encoded into image embeddings  $E_f \in \mathbb{R}^{128}$ .  $t_a$  and  $t_v$  represent the temporal dimensions of audio and video respectively. These encodings are passed through separate convolution layers to have the same channel dimensions and concatenated to form a joint feature representation:

$$F(E_a, E_v, E_f) = (E_a; E_v; E_f) \in \mathbb{R}^{(t_a+t_v+1) \times 768}$$

The aggregated features are passed through a three-layer transformer encoder. The resulting encoded representation is then fed into a modified speaker feature extractor, yielding the final enrolment vectors  $S_{avf} \in \mathbb{R}^{192}$ .

**Video-Only Network (Fig 1(b)):** In this case there is no audio input. The video encoder processes silent video stream and outputs visual embeddings  $E_v \in \mathbb{R}^{t_v \times 512}$ , while the face encoder transforms still images into face embeddings  $E_f \in \mathbb{R}^{128}$ . The network can utilise either or both embeddings. When both are used, they are aligned by channel dimensions and concatenated. This combined data undergoes processing through three transformer encoder layers, subsequently feeding into the modified speaker feature extractor to produce enrolment vectors  $S_{vf} \in \mathbb{R}^{192}$ .

$$S_{vf} = \text{SpeakerExtractor}(\text{Transformer}([E_v; E_f]))$$

##### 3.1.2. Speaker separation network.

The second network shown in Fig 1(c) is a VoiceFormer-inspired model [4], consisting of an audio encoder-decoder and a speaker embedding encoder. This network is fed with noisy audio waveforms and corresponding multiple positive and negative enrolment vectors. It utilises a three-layer

Transformer encoder to combine the audio and speaker encodings. Attention is increased for features that match the target speaker (positive embeddings) while downplaying features from non-target speakers (negative embeddings). Finally, skip connections between the audio encoder and decoder help retain both low- and high-level information, leading to improved performance. The output of this network is the separated target speaker waveform signals.

### 3.2. Training

**Training objectives:** For the *Speaker Enrolment Network* given a dataset  $\mathcal{D}$  of tuples  $(a, v, f)$ , the network uses pre-recorded noisy audio waveforms  $a \in \mathbb{R}^{t_a}$ , conditioned on lip motions alone  $v \in \mathbb{R}^{3 \times t_v \times H \times W}$ , face images alone  $f \in \mathbb{R}^{3 \times H \times W}$  or a combination of both, to predict an enrolment vector  $S_p \in \mathbb{R}^{192}$ . The objective is to minimize the L1 discrepancy between the predicted  $S_p$  and the speaker embedding  $S_{a_c}$  of the target speaker generated by the audio-only network shown (Fig 1(a)). Crucially  $S_{a_c}$  is generated from a distinct recording of the target speaker.  $\mathcal{L}_{\text{enrolment}} = \mathbb{E}_{(a,v,f) \in \mathcal{D}} \|S_{a_c} - S_p\|_1$ .

For the *Speaker Separation Network*, given a dataset  $\mathcal{D}$  of tuples  $(a_m, S_p, a_t)$ , the network conditions the noisy audio mixture  $a_m$  on the predicted enrolment vectors  $S_p$  and aims to align the predicted audio  $\hat{a}_t$  with the clean target waveforms  $a_t$ , formalized using the L1 loss:  $\mathcal{L}_{\text{separate}} = \mathbb{E}_{(a_m, S_p, a_t) \in \mathcal{D}} \|a_t - \hat{a}_t\|_1$ .

**Training stages:** The training is structured in a step-by-step process. Initially, as shown in Fig 1(a), clean audio is used to create target enrolment vectors. Next, as shown in Fig 1(b), the enrolment network is trained using knowledge distillation, such that it aligns its outputs with the speaker feature extractor model in Fig 1(a). This network is further enhanced by training it with a mix of noisy audio with corresponding visual elements or just visual information. The video-only version capitalizes on either the lip motion of the speaker or by integrating static full-face images of the speaker in tandem with the corresponding video clips. The subsequent step involves training the separation network, represented in Fig 1(c). Here, we use noisy audio that includes the voices of the target speaker, other speakers, and background noise. This audio is paired with a random combination of both positive and negative enrolment vectors that serve as conditioning signals, helping to emphasize the desired voice while suppressing others in the mix. Finally, we implemented an end-to-end fine-tuning phase, refining both the enrolment and separation networks.

### 3.3. Implementation details

The model is trained using Pytorch. Video data underwent facial cropping to the mouth region of the speaker with a uniform 25 FPS. Audio was monophonized via channel averaging and resampled to 16kHz.

The enrolment network employs a visual backbone from [20], integrating a VTP network with transformer units over a hybrid 3D/2D residual CNN. Face embeddings are produced via an off-the-shelf library based on the VGG-Face model [21].

The audio encoder uses a sequence of 1D CNN layers, processing audio waveforms directly. The Transformer comprises 3 layers and 8 attention heads, with a model dimension of 532. We use 768-dimensional embeddings across modalities throughout, ensuring alignment with post-audio encoding channel dimensions. The transformer’s intermediate outputs within this network are channelled into an adapted ECAPA-TDNN model [19]. Note that, unlike the traditional ECAPA-TDNN that consumes spectrograms, our modified version leverages features from the transformer’s output.

For the speaker separation network, the architecture drew insights from the design in [4], but adapted to accommodate enrolment vectors as opposed to video or text. This module integrates an audio encoder-decoder with a dedicated encoder for enrolment vectors. The model ingests a varying number of positive and negative enrolment vectors, always ensuring at least one positive enrolment vector during training. A three-layer Transformer bottleneck fuse encoded noisy audio and enrolment vectors, attending to relevant audio features based on the positive and negative embeddings.

Each enrolment vector is derived from a 4-second audio sample, processed using the enrolment network depicted in Fig 1(b). During training, the separation network was exposed to a varied number of positive and negative vectors.

## 4. EXPERIMENTS

### 4.1. Synthetic sequences

Consistent with previous research [5, 8], we train and test our models using synthetic noisy samples, formed by combining waveforms from two separate speakers. For a diverse training set, we mix speech from multiple speakers and background noise from the DNS dataset, including traffic, wind, and other general noises. These 4-second audio mixtures are randomly cropped from extended tracks and undergo augmentations like speed, pitch, and decibel adjustments to mimic varied acoustic scenarios. It is worth highlighting that despite our model’s training and evaluation on synthesized audio mixtures, it remains adept at managing authentic noisy sequences, given the small domain gap between synthetic and real-world audio samples.

### 4.2. Datasets, training and evaluation protocol

**Data.** Our evaluation uses two well-known datasets: LRS3 [22] and Librispeech [23]. LRS3 derived from public TEDx videos, boasts a vast array of audio-visual clips covering diverse speaking styles and subjects. Traditionally favoured for lip reading and speech recognition, its comprehensive nature aligns perfectly with our objectives. Librispeech, a vast audio-only dataset sourced from public-domain audiobooks, is meticulously segmented, proving invaluable for tasks like automatic speech recognition and speaker identification.

Both the LRS3 and Librispeech datasets cover a diverse range of speakers of varying genders, accents, and speech patterns, offering a rich and varied testing environment. Cru-

cially, the speakers in the test sets are not encountered during training. Synthetic test samples were crafted by randomly mixing two distinct speakers from the evaluation subsets of LRS3 and Librispeech, with the addition of ambient noise.

**Evaluation metrics.** We use three standard metrics: (i) Signal-to-Distortion Ratio (SDR) to quantitatively measure the quality of the separated output; (ii) Short-Time Objective Intelligibility (STOI) quantifies signal intelligibility; and (iii) the Perceptual Evaluation of Speech Quality (PESQ) offers a perceptual score reflecting the listener’s perception. Together, these metrics offer a comprehensive evaluation of our method and allow comparison with existing benchmarks.

### 4.3. Results

In this section, we evaluate our proposed method. We first assess the separation performance under various combinations of positive and negative embeddings. Next, we evaluate the model based on enrolment vectors from diverse modalities. We conclude by benchmarking against the state-of-the-art in speaker separation and speech enhancement tasks.

**Positive and negative embeddings.** The positive vectors are crafted from distinct recordings of the target speaker, ensuring they are different from the immediate audio being separated. On the other hand, the negative vectors are sourced from different recordings of the non-target speaker present in the mixed audio. Table 1 illustrates the effectiveness of introducing both positive and negative enrolment vectors. The results confirm that increasing the number of enrolment vectors enhances the model’s separation accuracy.

**Table 1. Speaker Separation performance with a varied number of enrolment vectors.** Results on the LRS3 test set, as the numbers of **P** positive and **N** negative enrolment vectors are varied. It is evident that increasing the number of positive and negative enrolment vectors enhances performance. Higher score is better.

	1P — 0N	1P — 1N	3P — 2N	3P — 3N
SDR↑	13.8	14.0	14.4	<b>14.5</b>

**Modalities comparison.** Table 2 illustrates our model’s effectiveness in the interplay between audio and visual modalities. When operated solely on clean audio, the model stands robust, but the introduction of visual cues, specifically lip motions, to the clean audio offers a marginal advantage. The model stands resilient in cases where we do not have access to any clean audio from the target speaker and the positive vectors are generated from a combination of noisy audio and visual cues. Interestingly, the reliance solely on lip motions, yields impressive results, highlighting the importance of visual cues. Furthermore, incorporating the facial data amplifies performance. This harmony between auditory and visual data shows the strength of our multi-modal approach in scenarios where one of the modalities might not be available.

**Table 2. Speaker Separation: Enrolment Vectors Across Modalities.** Symbols ✓ and ✗ indicate the modality’s presence or absence in the enrolment network. A stands for audio and V for visual cues. We assess speaker separation using enrolment vectors generated from various modality combinations in the synthetic LRS3 test set. Using 3 positive and 2 negative vectors. Higher is better.

Modality	A	V	SDR↑	STOI↑	PESQ↑
Clean audio-only	✓	✗	14.4	91	2.52
Clean audio + lip motions	✓	✓	14.5	91	2.55
Noisy audio only	✓	✗	6.3	58	1.82
Noisy audio + lip motions	✓	✓	13.7	88	2.45
Lip motions only	✗	✓	11.1	77	2.25
Lip motions + face images	✗	✓	12.0	80	2.35
Face images only	✗	✓	10.8	79	2.20

**State-of-the-Art Comparison.** Table 3 compares the performance of our approach in speaker separation with preceding methods. For benchmarking, we consider the speech separation conditioned on speaker embeddings from [1] and a selection of recent audio-visual approaches [4, 8, 24]. VoiceVector consistently surpasses or performs on par with earlier methods. It is worth highlighting that, unlike the other audio-visual methods, our model does not access visual cues during separation, nor has it encountered synchronised visual cues from the mixture undergoing separation.

**Table 3. Comparison to the state-of-the-art on speaker separation.** Evaluation on the synthetic LRS3 and Librispeech test sets. Our best model outperforms the previous work. A stands for audio and V for visual cues. ✓ and ✗ indicate the presence or absence of the modality respectively. †Our model has not been trained on Librispeech, demonstrating generalisation across datasets. Higher is better.

Model	A	V	Testset	SDR↑	STOI↑	PESQ↑
Conversation [8]	✓	✓	LRS3	11.5	87	2.18
VisualVoice [24]	✓	✓	LRS3	11.9	90	2.46
VoiceFormer [4]	✓	✓	LRS3	14.4	<b>92</b>	2.42
Ours (VoiceVector)	✓	✓	LRS3	<b>14.5</b>	91	<b>2.52</b>
VoiceFilter [1]	✓	✗	Librispeech	12.6	—	—
Ours (VoiceVector)†	✓	✗	Librispeech	<b>13.1</b>	<b>89</b>	<b>2.12</b>

## 5. DISCUSSION

In summary, we introduce a two-phase approach to voice separation that leverages combinations of auditory and visual data to generate enrolment vectors, followed by a separation phase that exclusively relies on auditory data. This strategy ensures robustness against visual data variability in comparison to previous audio-visual separation approaches. The use of both positive and negative vectors exemplifies a nuanced angle to the separation task while improving performance.



## 6. REFERENCES

- [1] Quan Wang, Hannah Muckenhirn, Kevin Wilson, Prashant Sridhar, Zelin Wu, John Hershey, Rif A Saurous, Ron J Weiss, Ye Jia, and Ignacio Lopez Moreno, “Voicefilter: Targeted voice separation by speaker-conditioned spectrogram masking,” in *Interspeech*, 2018.
- [2] Kateřina Žmolíková, Marc Delcroix, Keisuke Kinoshita, Tsubasa Ochiai, Tomohiro Nakatani, Lukáš Burget, and Jan Černocký, “Speakerbeam: Speaker aware neural network for target speaker extraction in speech mixtures,” *JSTSP*, vol. 13, no. 4, pp. 800–814, 2019.
- [3] Tsubasa Ochiai, Marc Delcroix, Keisuke Kinoshita, Atsunori Ogawa, and Tomohiro Nakatani, “Multimodal SpeakerBeam: Single Channel Target Speech Extraction with Audio-Visual Speaker Clues,” in *Proc. Interspeech 2019*, 2019, pp. 2718–2722.
- [4] Akam Rahimi, Triantafyllos Afouras, and Andrew Zisserman, “Reading to listen at the cocktail party: Multimodal speech separation,” in *CVPR*, 2022.
- [5] Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinatan Hassidim, William T. Freeman, and Michael Rubinstein, “Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation,” *ACM Trans. Graph.*, 2018.
- [6] Triantafyllos Afouras, Andrew Owens, Joon Son Chung, and Andrew Zisserman, “Self-supervised learning of audio-visual objects from video,” in *ECCV*, 2020.
- [7] Katerina Zmolikova, Marc Delcroix, Tsubasa Ochiai, Keisuke Kinoshita, Jan Černocký, and Dong Yu, “Neural target speech extraction: An overview,” *IEEE Signal Processing Magazine*, vol. 40, no. 3, pp. 8–29, 2023.
- [8] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman, “The conversation: Deep audio-visual speech enhancement,” in *Interspeech*, 2018.
- [9] Héctor Martel, Julius Richter, Kai Li, Xiaolin Hu, and Timo Gerkmann, “Audio-Visual Speech Separation in Noisy Environments with a Lightweight Iterative Model,” in *Interspeech*, 2023.
- [10] Aviv Gabbay, Ariel Ephrat, Tavi Halperin, and Shmuel Peleg, “Seeing through noise: Visually driven speaker separation and enhancement,” *IEEE*, 2018.
- [11] T. Afouras, J. S. Chung, and A. Zisserman, “My lips are concealed: Audio-visual speech enhancement through obstructions,” in *Interspeech*, 2019.
- [12] David Snyder, Daniel Garcia-Romero, Gregory Sell, Daniel Povey, and Sanjeev Khudanpur, “X-vectors: Robust dnn embeddings for speaker recognition,” in *ICASSP*, 2018, pp. 5329–5333.
- [13] Soo-Whan Chung, Soyeon Choe, Joon Son Chung, and Hong-Goo Kang, “Facefilter: Audio-visual speech separation using still images,” in *Interspeech*, 2020.
- [14] Xiong Xiao, Zhuo Chen, Takuya Yoshioka, Hakan Erdogan, Changliang Liu, Dimitrios Dimitriadis, Jasha Droppo, and Yifan Gong, “Single-channel speech extraction using speaker inventory and attention network,” in *ICASSP*, 2019.
- [15] Yi Luo and Nima Mesgarani, “Tasnet: Time-domain audio separation network for real-time, single-channel speech separation,” *ICASSP*, pp. 696–700, 2017.
- [16] John Hershey, Zhuo Chen, Jonathan Le Roux, and Shinji Watanabe, “Deep clustering: Discriminative embeddings for segmentation and separation,” in *ICASSP*, 2016.
- [17] Dong Yu, Morten Kolbæk, Zheng-Hua Tan, and Jesper Jensen, “Permutation invariant training of deep models for speaker-independent multi-talker speech separation,” in *ICASSP*, 2017, pp. 241–245.
- [18] Morten Kolbaek, Dong Yu, Zheng-Hua Tan, Jesper Jensen, Morten Kolbaek, Dong Yu, Zheng-Hua Tan, and Jesper Jensen, “Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks,” *IEEE/ACM TASLP*, 2017.
- [19] Brecht Desplanques, Jenthe Thienpondt, and Kris Demuynck, “Ecapa-tdnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification,” in *Interspeech*, 2020.
- [20] Prajwal K Renukanand, Liliane Momeni, Triantafyllos Afouras, and Andrew Zisserman, “Visual keyword spotting with attention,” in *BMVC*, 2021.
- [21] Omkar M. Parkhi, Andrea Vedaldi, and Andrew Zisserman, “Deep face recognition,” in *BMVC*, 2015.
- [22] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman, “Lrs3-ted: a large-scale dataset for visual speech recognition,” 2018.
- [23] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: An asr corpus based on public domain audio books,” in *ICASSP*, 2015.
- [24] Ruohan Gao and Kristen Grauman, “VisualVoice: Audio-Visual Speech Separation with Cross-Modal Consistency,” in *CVPR*, 2021.