

Looking into Your Speech: Learning Cross-modal Affinity for Audio-visual Speech Separation

Jiyoung Lee^{1*} Soo-Whan Chung^{1,2*} Sunok Kim³
 easy00@yonsei.ac.kr soowhan.chung@navercorp.com sunok.kim@kau.ac.kr

Hong-Goo Kang^{1†} Kwanghoon Sohn^{1†}
 hgkang@yonsei.ac.kr khsohn@yonsei.ac.kr

¹Department of Electrical & Electronic Engineering, Yonsei University, Korea

²Naver Corporation, Korea ³Korea Aerospace University, Korea

<https://caffnet.github.io/>

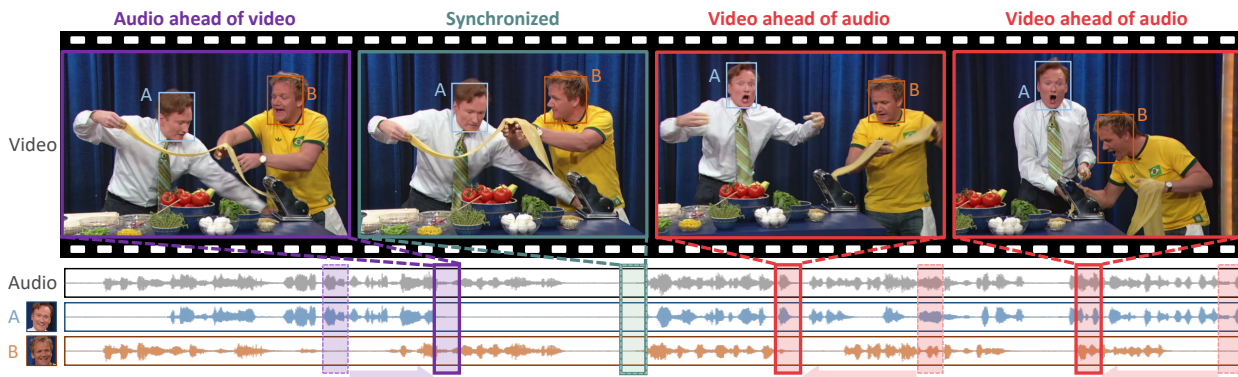


Figure 1: We wish to hear individual speech of a desired speaker only even if there is frame discontinuity in the audio-visual data. When audio and video segments are taken from different points in time (**solid box**), it is intuitively difficult to separate speech of each speaker compared to the aligned cases (**dashed box**). Best viewed in color.

Abstract

In this paper, we address the problem of separating individual speech signals from videos using audio-visual neural processing. Most conventional approaches utilize frame-wise matching criteria to extract shared information between co-occurring audio and video. Thus, their performance heavily depends on the accuracy of audio-visual synchronization and the effectiveness of their representations. To overcome the frame discontinuity problem between two modalities due to transmission delay mismatch or jitter, we propose a cross-modal affinity network (CaffNet) that learns global correspondence as well as locally-varying affinities between audio and visual streams. Given that the global term provides stability over a temporal sequence at the

utterance-level, this resolves the label permutation problem characterized by inconsistent assignments. By extending the proposed cross-modal affinity on the complex network, we further improve the separation performance in the complex spectral domain. Experimental results verify that the proposed methods outperform conventional ones on various datasets, demonstrating their advantages in real-world scenarios.

1. Introduction

Humans have a remarkable auditory system that can perceive sound sources separately in their conversations even in the presence of many surrounding sounds, including background noise, crowded babbling, thumping music, and sometimes other loud voices [1, 2]. However, reliably separating a target speech signal for human-computer interaction (HCI) systems such as speech recognition [3, 4, 5], speaker recognition [6, 7, 8], and emotion recognition [9, 10] is still a challenging task because it is an ill-posed problem.

* Both authors contributed equally to this work

† Corresponding authors

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT). (NRF-2021R1A2C2006703).

With the impressive advent of deep learning technologies that utilize high-dimensional embeddings [11, 12, 13], it is possible nowadays to simultaneously analyze the unique acoustic characteristics of different speakers even from mixed signals. Although these deep learning-based methods are effective compared to conventional statistical signal processing-based ones, they are prone to a label permutation (or ambiguity) error due to their frame-by-frame or short segment-based processing paradigm [11, 14]. In order to address this problem, permutation invariant training [15, 16] that utilizes a permutation loss criterion was presented, but the label ambiguity problem still occurs at the inference stage, especially for unseen speakers.

Leveraging the visual streams of target speech signals can be one of the best alternatives. In psychology, several experiments have proved that looking at speakers' faces is helpful for auditory perception under background noise environments [17, 18]. For example, lip reading, which matches lip movements onto utterances, is widely used to recognize others' words better [19]. In audio-visual speech separation (AVSS) systems, audio and visual features are used together or complement each other to derive unique characteristics [20, 21, 22, 23, 24, 25, 26]. Mostly, AVSS first extracts the common correspondence features between speaker/linguistic information of speech signals and face/articulatory lip movements of video signals, after which the extracted features are exploited for the following source separation task. Consequently, the AVSS problem can be viewed as a local matching (*i.e.* frame-wise matching) task, where segmented visual features are matched with frames of specific sounds. Thus, the separation performance highly depends on the alignment accuracy between audio and video streams.

In real-world scenarios, however, audio and video are recorded from different devices with their own specifications, and they are transmitted through independent communication channels and saved with different codec protocols. These practical issues frequently cause mutually unaligned states in talking videos. Fig. 1 shows an example of a video with a speech that has physical errors in its video contents, where sometimes audio plays ahead of video and vice versa. When there are even subtle data transformations caused by jitters, omissions, and out-of-synchronization in video streams, conventional local matching strategies [20, 23, 25] are vulnerable. This issue can be detrimental to the performance of AVSS systems in video-telephony, broadcasting, video conferencing, or filming.

In this paper, we highlight those limitations and tackle the alignment problems in AVSS processing. We propose a novel cross-modal affinity network for robust speech separation, referred to as *CaffNet*, by utilizing visual cues in consideration of relative timing information. Affinity, *i.e.* mutual correlation, learned in *CaffNet* compensates for abrupt

discontinuities in audio-visual data without external information or additional supervision. Furthermore, we propose an affinity regularization module that tiles the diagonal term of the affinity matrix to match audio-visual sequences at the utterance level. Since the affinity regularization provides a global positional constraint, it avoids the label permutation problem that occurred by inconsistent assignment over time of the speech signals to the visual target. In addition, considering the estimation of the magnitude mask in tandem with the phase mask is one of the keys to reasonable speech reconstruction because such factors are correlated with each other [27, 28]. To accomplish this, we extend *CaffNet* to have a complex-valued convolution network architecture [29, 30, 31] such that speech quality is indeed increased by restoring the mask of the magnitude and phase spectrum together. We demonstrate the effectiveness of the proposed networks with extensive experiments, achieving large improvements in unconditioned scenarios on several benchmark datasets [32, 33, 34].

2. Related Work

Audio-visual Speech Separation. In terms of multi-sensory integration, it has been proved that looking at talking faces during conversation is helpful for speech perception [35]. Inspired by this psychological mechanism, numerous works have tried to effectively utilize visual context on speech separation tasks [36, 37, 25, 20, 38, 24, 21, 22, 39, 23]. With the emergence of deep neural networks and the availability of new large-scale datasets [19, 33, 34, 24], a series of works [38, 40] have been published in the past few years on audio-visual speech separation (AVSS). Although they have shown promising results in various speech-oriented applications, these methods have concentrated on isolating the magnitude of speech only, which restricts their applicability. To alleviate this limitation, several works have been proposed to generate both magnitude and phase masks [20, 24, 23, 22]. However, these methods generate phase masks with estimated magnitude masks and noisy phase without learning complex-valued representations, which requires them to consider the correlation between each complex component. This limitation has manifested itself in significantly degraded performance under extremely noisy conditions [31]. Furthermore, the aforementioned methods have all assumed one-to-one correspondence between the audio-visual segments.

Most recently, this problem has been tackled in [21, 39] for situations in which visual face sequences were not fully reserved. In [21], although they considered the case that visual cues abruptly disappear due to occlusion, it still required video sequences aligned to audio. Even though the speaker identity extracted from the still face image seemed to provide a promising visual cue for the separation [39], it showed far less separation performance than ones using

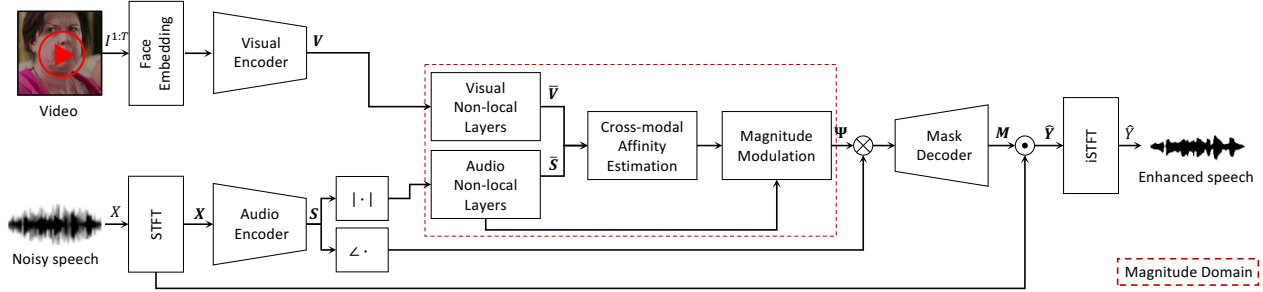


Figure 2: Overall network configuration: (1) encoding audio and visual features; (2) learning cross-modal affinity; (3) predicting spectral soft mask M to reconstruct target speech \hat{Y} . A red-dotted box means the magnitude operation processing.

visual sequences since they only regulated global information rather than local contexts. In this paper, we leverage sequential audio-visual frames as the inputs to our networks under the assumption that locally misaligned visual frames with audio frames can still provide local context and speaker identity for robust speech separation.

Cross-modal Alignment. As audio and video sequences are recorded using different devices, synchronization problem often appears in recordings. Most recent audio-visual synchronization methods rely on cross-modal representation techniques that measure the linguistic similarity between audio-visual embedding pairs [41, 42, 43]. However, there has been little work on exploring the problem of synchronization along with mismatched audio and video pairs because prior works generally assumed that a paired audio and video set has only one speech. More related to our work are affinity-based multi-modal approaches in various other challenging tasks, such as music sound separation [44], emotion recognition [45], language understanding [46], and self-supervised learning [47]. We further extend the cross-modal affinity learning to generate time-independent audio-visual representations using an affinity regularization with an utterance-wise matching criterion.

3. Approach

3.1. Motivation and Overview

In general, humans experience severe confusion when there is a linguistic discrepancy between what we see and what we hear, *i.e.* a difference between the perceived words from the mouth and actual speech. This is called cognitive dissonance, which is known as the McGurk effect [48]. This effect could be observed in previous frame-wise matching based methods [32, 20, 23, 38, 25] inducing poor performance when the cross-modal data is conflicted. To deal with such inconsistency problem, we introduce a CaffNet to estimate time-frequency soft masks to isolate a single speech signal from a mixture of sounds (such as other speakers and background noise), taking into account time-agnostic mutual correlation.

Concretely, our model is split into three parts, including

an audio-visual encoder, learning cross-modal affinity, and soft-mask estimation, as shown in Fig. 2. The key idea of CaffNet is to learn cross-modal affinity between the audio and video streams even if they have different sampling rates in the wild environments. By this, we mean that information from the video stream stretches or compresses to match the audio signal for the reconstruction of the target speaker’s speech regardless of the frame discontinuity problem. Due to matching ambiguity generated in parts that are muted or that have similar pronunciations from simultaneous speakers, the initially computed affinity includes erroneous values and causes the label permutation problem while degrading the separation performance. We resolve this problem by suggesting an affinity regularization to induce global consistency of cross-modal affinity. Furthermore, we extend this approach to complex-valued neural networks, estimating the magnitude and phase components jointly.

Given a noisy time-domain speech X , CaffNet is trained to isolate a clean speech Y from X with corresponding a user-chosen speaker’s face video $I^{1:T}$, where T is a length of the video stream. The noisy sound $X = Y + H$ is assumed to be a sum of clean speech Y and natural environmental factors H such as background noise, distortions in speech, and sound from other speakers. As it has been a common practice to transform a time-domain speech to a time-frequency representation (*i.e.* spectrogram) via short-time Fourier transform (STFT), each of the corresponding time-frequency representations for X , H , and Y is computed by 512-point STFT and denoted by $\mathbf{X} \in \mathbb{C}$, $\mathbf{H} \in \mathbb{C}$ and $\mathbf{Y} \in \mathbb{C}$, respectively.

3.2. Cross-modal Affinity Network

Audio-visual Encoder. As in [20, 21, 44], the audio-visual encoder has a two-stream architecture consisting of an audio encoder stream and a video encoder stream, which take noisy audio and video frames containing the target face, respectively. At first, the audio and video encoders generate their own embedding features independently. In specific, the audio encoder \mathcal{F}_s takes the magnitude spectrum of consecutive input frames, $|\mathbf{X}|$. The audio embedding features are extracted by stacked 1D convolutional lay-




Figure 3: Illustration of cross-modal affinity estimation module. It takes speech feature $\bar{\mathbf{S}}$ and video feature $\bar{\mathbf{V}}$ to calculate the affinity matrix \mathbf{A} . The cross-modal identity matrix $\mathbf{\Gamma}$ regularizes the affinity matrix \mathbf{A} to maintain global correspondence.

ers $\mathbf{S} = \mathcal{F}_s(|\mathbf{X}|)$, where $\mathbf{S} \in \mathbb{R}^{N \times C}$ is a speech embedding feature, C and N indicate the dimension of a channel and the temporal length of the spectrogram, respectively. Besides, visual features are extracted from the temporal stack of five consecutive video frames via the state-of-the-art audio-visual synchronization model $\mathcal{E}(\cdot)$ [41] by a feed-forward process. Finally, a visual embedding feature $\mathbf{V} \in \mathbb{R}^{M \times C}$ is obtained through the visual encoder \mathcal{F}_v :

$$\mathbf{V} = \mathcal{F}_v(\Pi(\mathcal{E}_f(I^{1:5}), \mathcal{E}_f(I^{2:6}), \dots, \mathcal{E}_f(I^{T-4:T}))), \quad (1)$$

where I denotes video frames, T is the number of video frames, $\Pi(\cdot)$ indicates a concatenation operator, and $M = T - 4$ is an entire length of clips tied 5 frames. As the audio and video have different sampling rates, it requires either up-sampling or down-sampling process to equalize the temporal resolution of audio and video embedding matrices [20, 23]. However, our network is fully-convolutional network that effectively learns the audio-visual affinity regardless of the size of each embedding matrix.

Learning Cross-modal Affinity. We assume that audio and video embeddings are naturally out of joint in the unconstrained environments due to temporal mismatch between two media. Considering the fact that learning affinity can draw linguistic dependencies between audio and visual features, it is possible to model relative timing dependencies without considering their distances.

More specifically, we first extract audio feature $\bar{\mathbf{S}}$ and visual feature $\bar{\mathbf{V}}$ by feeding outputs of the audio-visual encoder to modality-separated two non-local layers [49] to measure the affinity on the nearest embedding space as possible. Then, an affinity matrix $\mathbf{A}_{i,j}$ between i -th audio feature and j -th visual feature is computed using cosine similarity with L_2 normalization on an embedding space:

$$\mathbf{A}_{i,j} = \text{softmax}\left(\frac{\langle \bar{\mathbf{S}}_i \mathbf{w}_s, \bar{\mathbf{V}}_j \mathbf{w}_v \rangle}{\|\bar{\mathbf{S}}_i \mathbf{w}_s\|_2 \|\bar{\mathbf{V}}_j \mathbf{w}_v\|_2}\right), \quad (2)$$




Figure 4: Visualization of affinity matrices. (a) When audio contains only one voice, it shows monotonous pattern with audio-visual correspondence. (b) When there is mixture input, the pattern is tangled in absence of sequential consistency. (c) Final affinity is arranged with the help of affinity regularization.

where \mathbf{w}_s and \mathbf{w}_v are embedding weights of audio and visual features, respectively, as illustrated in Fig. 3. In (2), softmax activation function is applied row-wise for normalization to obtain an affinity matrix $\mathbf{A} \in \mathbb{R}^{N \times M}$. Ideally, the linear pattern can be found along the diagonal of the affinity matrix as shown in Fig. 4(a). However, we observe that the affinity weights in different frames are often similar in the silent speech or regions having similar pronunciation between different speakers as depicted in Fig. 4(b), which can cause the label permutation problem. To resolve this problem, we propose an affinity regularization to penalize the probabilistic affinity matrix corresponding to global alignment context to reliably infer the spectrogram mask of an interest such that

$$\mathbf{\Gamma}_{i,j} = \frac{\exp(\sum_{i,j} \mathcal{D}_{i,j}^k \mathbf{A}_{i,j} / \tau)}{\sum_k \exp(\sum_{i,j} \mathcal{D}_{i,j}^k \mathbf{A}_{i,j} / \tau)}, \quad (3)$$

where $\mathbf{\Gamma}$ is a cross-modal identity matrix, $k \in \mathcal{N}_k$ is the search window of offset range across the diagonal term, τ is a softening parameter [50] set as 0.1, and \mathcal{D} is a diagonal mask which satisfies:

$$\mathcal{D}_{i,j}^k = \begin{cases} 1, & \text{if } \frac{f_a}{f_v}(j-k) + 1 \leq i \leq \frac{f_a}{f_v}(j-k+1) \\ 0, & \text{otherwise,} \end{cases} \quad (4)$$

where f_* is a sampling rate of each segment (e.g., $f_a/f_v = 4$ in our experimental setting). In our experiments, we search for the offsets over $[-9, +9]$ frame range, where negative offset means that audio is ahead of video and vice versa. We set $\mathcal{N}_k \in [0, \dots, 19]$ and if the input pair is matched in a timely manner, diagonal term appears from the 9-th video frame index as shown in Fig. 4(c). Fig. 4 clearly shows that the regularization encourages the model to maintain temporal consistency in the matching process. Then the final visual features $\hat{\mathbf{V}} \in \mathbb{R}^{N \times C'}$ are then represented as follows:

$$\hat{\mathbf{V}}_i = \sum_j (\mathbf{A}_{i,j} + \gamma \mathbf{\Gamma}_{i,j}) \cdot (\mathbf{w}_o \bar{\mathbf{V}}_j)^\top, \quad (5)$$

where the identity matrix $\mathbf{\Gamma}$ is added to the initial affinity matrix \mathbf{A} without breaking its global behavior [49, 51]. w_o is a projection parameter which has C' output channel. Balance parameter γ is set to 1.0 in our experiments.

Soft Mask Estimation. The mask decoder \mathcal{F}_m takes both the transformed visual features $\hat{\mathbf{V}}$ and corresponding audio features $\hat{\mathbf{S}}$ to generate a soft mask [52], which filters the mixture spectrogram to produce the enhanced spectrogram. Audio-visual features are concatenated over the channel dimension to compute an integrated feature $\Psi = \Pi(\hat{\mathbf{S}}, \hat{\mathbf{V}})$. In this way, each audio feature is associated with corresponding visual features, which will be used to recover clean speech. We employ the similar mask decoder architecture used in [20] as the residual building block of our decoder. The sequentially up-scaled output to the original size of the input spectrogram is then passed through sigmoid activation to regularize output values between 0 and 1. Finally, the estimated speech spectrogram $\hat{\mathbf{Y}} = \mathbf{M} \odot |\mathbf{X}|$ is computed by element-wise multiplying the estimated mask $\mathbf{M} = \mathcal{F}_m(\Psi)$ on the input spectrogram $|\mathbf{X}|$. Then, the estimated speech \hat{Y} is computed by inverse STFT. We note that the architectural detail of CaffNet is explained in supplementary materials.

Training. The terminal objective of our model is to estimate a target speech Y of an interest person associated visual inputs. During training, while previous local matching methods assume that audio is correctly aligned to video [20, 23, 25], we consider that sometimes the audio stream leads the video or sometimes the video stream is going ahead of the audio. To accomplish these cases in the training scheme, CaffNet leverages a video clip that is recorded a little longer than the randomly sampled audio in the datasets. However, there is no additional label on what time the audio will be matched to in the video.

To train CaffNet, we minimize the magnitude loss \mathcal{L}_{MAG} that makes the magnitude of enhanced spectrogram be similar to that of clean spectrogram on a logarithmic scale [39]:

$$\mathcal{L}_{\text{MAG}}(\mathbf{Y}, \hat{\mathbf{Y}}) = \|\log(|\mathbf{Y}|/|\hat{\mathbf{Y}}|)\|_2. \quad (6)$$

3.3. Complex Cross-modal Affinity Network

In this section, we explain how to further improve the separation ability of CaffNet generating complex ratio mask that considers magnitude and phase simultaneously with simple modifications. The complex model, CaffNet-C, has a similar architecture configuration as that of the CaffNet. Although using only CaffNet provides satisfactory performance, we upbuild inflated complex CaffNet (CaffNet-C) based on complex-valued building blocks [30] to handle complex matrices represented in the spectrograms. In tasks related to audio signal reconstruction, such as speech enhancement [31] and separation [29], it is ideal to perform correct estimation of both components. Details on batch

normalization and weight initialization for complex networks can be found in [30, 31].

Different from CaffNet, which solely takes the magnitude of spectrogram as input, the audio encoder stream of CaffNet-C leverages the whole amount of complex-valued spectrogram to extract the audio embedding feature with stacked complex-valued convolutional layers such that $\mathbf{S} = \mathcal{F}_s^c(\mathbf{X})$, where $\mathbf{S} \in \mathbb{R}^{N \times D \times 2}$ contains the real and imaginary parts of a complex number, and \mathcal{F}_s^c denotes complex audio encoder.

Inspired by [31], we decode the corresponding phase along with both the noisy phase and magnitude from the feature representation step. This solution makes the complex-valued mask \mathbf{M} estimated using the magnitude feature and noise phase feature at the same time. The noise phase is refined to clean phase with the complex mask decoder \mathcal{F}_m^c :

$$\mathbf{M} = \mathcal{F}_m^c(\Pi(|\hat{\mathbf{S}}|, \hat{\mathbf{V}}) \cdot e^{i\theta_{\hat{\mathbf{S}}}}). \quad (7)$$

The estimated speech spectrogram $\hat{\mathbf{Y}}$ is computed by multiplying the estimated mask \mathbf{M} on the input spectrogram \mathbf{X} :

$$\hat{\mathbf{Y}} = \mathbf{M} \odot \mathbf{X} = |\mathbf{M}| \cdot |\mathbf{X}| \cdot e^{i(\theta_{\mathbf{M}} + \theta_{\mathbf{X}})}. \quad (8)$$

Finally, we compute the estimated speech \hat{Y} with inverse STFT. By inducing complex convolutions in CaffNet-C, we use the scale-invariant source-to-distortion ratio (SI-SDR) to the objective function such as

$$\mathcal{L}_{\text{SI-SDR}}(Y, \hat{Y}) = -\langle Y, \hat{Y} \rangle / \|Y\| \|\hat{Y}\|, \quad (9)$$

where it makes more phase sensitive, as inverted phase gets penalized as well. Combining (6) and (9), the final overall objective function in the CaffNet-C is given by $\mathcal{L}_{\text{ALL}} = \mathcal{L}_{\text{MAG}} + \alpha \mathcal{L}_{\text{SI-SDR}}$, where α is a hyper parameter to balance two objective functions and we set it to 1.0.

4. Experiments

4.1. Setup

Datasets. Our networks are evaluated on three commonly used AVSS benchmarks: Lip Reading Sentences 2 (LRS2) [53, 19, 54], Lip Reading Sentences 3 (LRS3) [33], and VoxCeleb2 [34] datasets. LRS2 and LRS3 include 224 and 475 hours of videos respectively, along with cropped face tracks of the speakers. While LRS2 is sourced from British television broadcasts, LRS3 contains TED and TEDx videos. Following [21], we remove the few speakers from the LRS3 training set that also appear in the test set, so that there is no overlap of identities between the two subsets. Hence, the test set includes only unseen and unheard speakers during training and is suitable for a speaker-agnostic evaluation of our methods. VoxCeleb2 contains over 1 million utterances spoken by 5,994 speakers. They

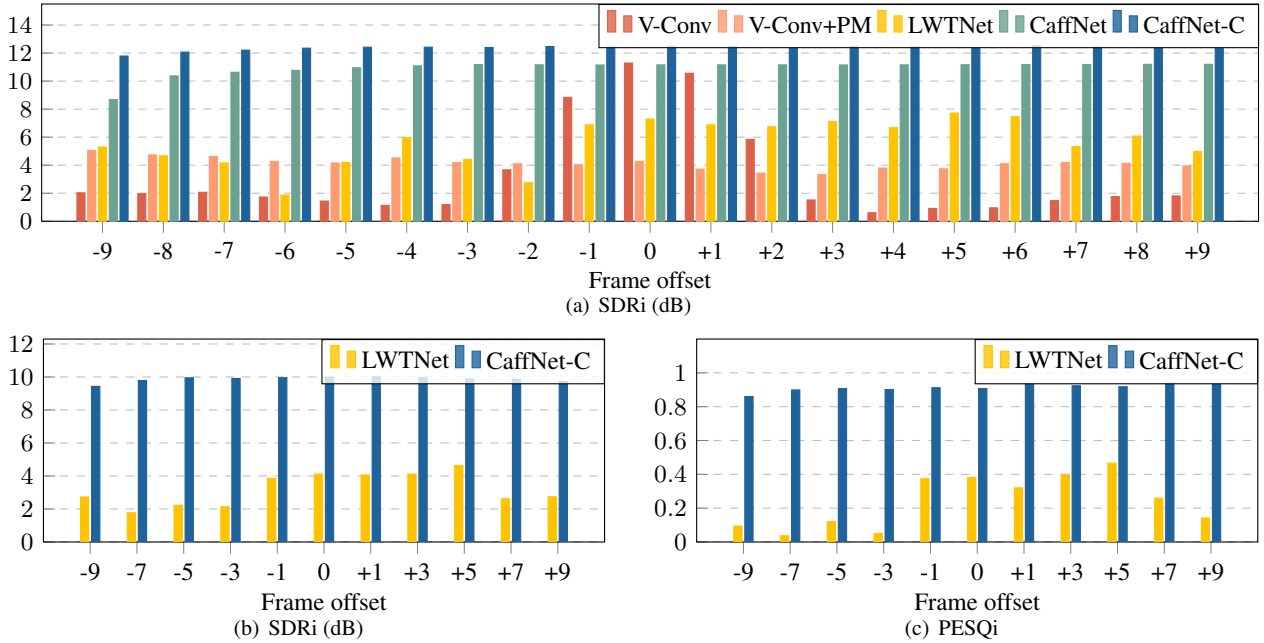


Figure 5: Evaluation of AVSS performance with respect to each delay offset between audio and visual streams on LRS2 dataset. The frame offset unit is 40ms which is the duration length between consecutive video frames. (a) reports the SDRi evaluation using ground-truth phase with estimated magnitude spectrum. (b) and (c) report the SDRi and PESQi evaluation results on the predicted phase as well as estimated magnitude spectrum, respectively.

provide the pre-train set and test set, and we follow this setting in our experiments. It is assumed that all datasets are well-synchronized [20, 21], so we adapt them to our purposes by augmenting data. Furthermore, VoxCeleb2 is divided into training and test sets according to the speaker’s identity to explicitly assess whether our model can generalize to unseen speakers during training. Thus training and testing sets are disjoint in terms of speaker’s identities.

Data Sampling Protocol. As mentioned in Sec. 3, we premise that audio-visual data is obtained in asynchronous circumstances. We assume that any given frame possesses the same time shift, so the visual stream is randomly shifted by -9 to 9 frames. Although we randomly shift video frames to assume the discontinuities of training data, the corresponding video clips contain all the audio-response information. For example, if the audio is sampled in the time duration $[T, T + \delta]$ and the video is shifted by -9 frames, *i.e.* 0.36 s, we extract the video frames during the time duration $[T - 0.36, T + \delta]$. If the selected time offset is 9 , the video frames are extracted within the time duration $[T, T + \delta + 0.36]$. For consistency and fair evaluation, we follow the similar evaluation settings in the previous work [20]. To generate synthetic training examples, we first select a source pair consisting of visual and audio features by sampling 2 seconds randomly. Source speech is mixed with randomly selected other speaker’s speech signal in the time domain, to simulate multi-talker backgrounds signals.

Features. We use a recent audio-visual synchronization model [41] for extracting visual features to serve as its visual input. The input to the visual stream is a video of cropped facial frames, with a frame rate of 25 fps. For every video frame, it outputs a compact 512-dimensional feature vector. For audio features, we use a time-frequency representation via STFT with a 25ms window length and a 10ms hop length as a sampling rate of 16kHz. Note that the extraction of face embeddings follows prior work [22].

Evaluation Metrics. We use three metrics to compare the results of our method to previous methods [20, 22]. First, the signal-to-distortion ratio (SDR) [55] is commonly used metric in recent works [20, 21, 38] to investigate the quality of enhanced speech. Following the previous work [20], we also report results on the perceptual evaluation of speech quality (PESQ) [56] varying from -0.5 to 4.5 and the short-time objective intelligibility (STOI) [57], which is correlated with the intelligibility of degraded speech signals. In the following experiments, we report SDR improvement (SDRi) and PESQ improvement (PESQi) for a fair comparison with other methods since the testing samples are randomly generated by combinations of the test set.

Baseline Models. For the fair comparison with CaffNet, we reproduce the magnitude network of ‘V-Conv’ [20], which is trained with the magnitude loss only. Also, since V-Conv model only assumes the synchronization circumstances between audio and visual streams, we combine V-

Dataset	Method	SDRi \uparrow				PESQi \uparrow				STOI \uparrow			
		GT	GL	PR	MX	GT	GL	PR	MX	GT	GL	PR	MX
LRS2	V-Conv [20]	11.28	-4.36	-	6.73	1.35	0.63	-	0.75	0.89	0.85	-	0.86
	LWTNet [22]	6.88	-4.61	4.06	3.77	0.65	0.16	0.32	0.29	0.77	0.72	0.73	0.74
	CaffNet (ours)	11.16	-3.49	-	6.79	1.29	0.63	-	0.73	0.89	0.85	-	0.86
	CaffNet-C (ours)	12.46	-2.54	10.01	7.94	1.15	0.65	0.94	0.73	0.89	0.84	0.88	0.86
LRS3	V-Conv [20]	11.23	-1.37	-	7.00	1.08	0.55	-	0.61	0.86	0.82	-	0.83
	LWTNet [22]	7.71	-3.93	4.83	4.44	0.82	0.35	0.49	0.45	0.84	0.80	0.82	0.82
	CaffNet (ours)	10.22	-2.64	-	6.46	1.06	0.49	-	0.60	0.86	0.82	-	0.84
	CaffNet-C (ours)	12.31	-1.38	9.78	7.92	0.91	0.49	0.71	0.55	0.86	0.82	0.85	0.83

Table 1: Evaluation of AVSS performance on the LRS2 and LRS3 datasets when audio and visual inputs are in synchronous condition. Contrary to other methods [22, 20], CaffNet and CaffNet-C are trained in unconditioned circumstance, *i.e.* training with randomly given frame offsets. GT: ground-truth phase; GL: Griffin-Lim; PR: predicted phase; MX: mixture phase.

Conv with a cutting-edge synchronization method proposed in [41, 58] to deal with asynchronous samples, referred to as ‘V-Conv+PM’. Furthermore, we examine the performance of LWTNet [22] using delayed test samples, where this method includes an independent synchronization module.

4.2. Results

LRS2 and LRS3. In Fig. 5, the proposed models, CaffNet and CaffNet-C, show robust performance in asynchronized environment while the previous methods have significant degradation in performance when the two frames were delayed. Although V-Conv+PM system has the synchronization step before the separation, it is clear that our methods are more robust to these temporal shifts. Because V-Conv+PM is a cascaded-step system, errors in the first step have a negative impact on the second. For that reason, V-Conv+PM even shows less effective results, even compared to V-Conv when audio and visual streams are synchronized. Furthermore, despite the alignment step in LWTNet [22], its accuracy is poorer than our methods on delayed samples.

When the two streams are well-aligned, our method provides competitive performance compared to baseline methods. In Tab. 1, we summarize the comparison results on the LRS2 and LRS3 datasets without a random shift on video streams (*i.e.*, frame offset is 0). Even though our network is trained in an unconditioned environment where audio and visual streams are not synchronized, our method outperforms existing methods [20, 22] which are trained with the synchronized audio-visual streams.

To validate the generality of our models, we investigate the performance on the LRS3 dataset, summarized in Tab. 1. Each model is only trained on the LRS2 dataset and evaluated on the LRS3 dataset without an additional adaptation process. Overall performances are similar to those on the LRS2 dataset. CaffNet-C achieves the best SDRi of 12.31 and still outperforms all the others, where the reconstructions are obtained using the magnitudes predicted by our network and either the ground truth phase. This demonstrates that our methods can be generalized to other datasets. Furthermore, the results show that phase estimation helps

Speaker	Mixture (Source-Reference)	Phase estimation			
		GT	GL	PR	MX
Seen	Male-Female	7.35	-5.46	3.54	3.33
	Female-Male	7.51	-3.63	3.99	3.71
	Male-Male	8.21	-4.61	4.24	3.91
	Female-Female	7.24	-3.41	3.91	3.58
Unseen	Male-Female	8.22	-5.12	4.80	4.42
	Female-Male	7.34	-3.42	3.99	3.71
	Male-Male	7.35	-5.46	3.55	3.33
	Female-Female	6.37	-4.23	3.15	2.92

Table 2: Evaluation of SDRi on CaffNet-C with regarding to gender combinations on the VoxCeleb2 dataset

reconstruct the magnitude. Compared to CaffNet, CaffNet-C improves the SDRi by 2.09dB using ground-truth phase on the LRS3 dataset, respectively.

VoxCeleb2. In order to explicitly investigate whether our model can be generalized to speakers unseen during training, we fine-tune and test on the VoxCeleb2 in Tab. 2. The training and test sets are disjoint in terms of speaker identities. We evaluate performance based on gender combinations, as many previous speech separation methods have shown significant performance degradation when mixtures involve same-gender speech [11, 39]. Although our method shows slight drops in performance on the female-female set, the result on the male-male set is similar to those of subsets including different genders.

4.3. Analysis

We conduct extensive experiments to describe the contribution of our method. Note that all experiments in this section are performed on the LRS2 dataset with CaffNet-C.

Channel Latency. Our premise is that temporal alignment of cross-modal can be achieved in our networks without additional supervision for ground-truth mapping between each pair of input streams. As shown in Fig. 6, CaffNet-C surprisingly finds appropriate frame offsets without any extra supervision. It demonstrates that our method reliably infers clean speech with well-aligned visual streams. Although we obtain the highest SDR in the




Figure 6: Qualitative results of affinity matrices according to the frame delay offsets. If the frame offset is negative, a linear pattern along the diagonal of the affinity matrix is displaced higher than zero offset case, and if it is positive, the pattern is shown lower than (b). (·) is SDR.

AR	Magnitude	Phase	Delay offset		
			-5	0	5
✗	Prediction	GT	8.01	7.77	7.71
✗	Prediction	GL	-3.93	-4.24	-3.96
✗	Prediction	MX	4.37	3.99	3.86
✗	Prediction	PR	5.13	4.75	4.62
✓	Prediction	GT	12.42	12.48	12.41
✓	Prediction	GL	-3.05	-2.67	-2.54
✓	Prediction	MX	7.96	7.97	7.85
✓	Prediction	PR	9.95	9.92	9.88

Table 3: Evaluation of SDRi to demonstrate the effectiveness of affinity regularization. ‘AR’ refers to the presence or absence of regularization.

synchronous setting, there are only slight differences between the performance despite delays in audio and video.

Affinity Regularization. One might also ask whether affinity regularization of our method is helpful. To verify its efficacy, we conducted an ablation study with CaffNet-C in Tab. 3. Since the affinity regularization has induced the global correspondence term for reliable speech separation, SDRi increases about 4 dB. It provides reasonable evidence for the effects of the affinity regularization.

Jitter. We then assessed the impact of our method from jitter, another well-known discontinuity problem frequently observed due to independent packetization of the audio-video streams. As a proof of concept, we assumed that no video frames are transmitted from t to $t+\tau$ frames, where τ is chosen as a random integer number such that $\tau \leq 8$ (≈ 0.3 s). Considering a simple frame repetition method that can be used in the jitter situation, we replaced the missing video frames with the $(t-1)$ -th frame. CaffNet-C shows compliant performance with 5.64 dB even in these challenging situations, in terms of SDRi on the testing set of LRS2 dataset. As exemplified in Fig. 7, it shows that our network clearly distinguishes two global terms in the affinity matrix when the frame jitter arises in the visual stream.

Speech Recognition. To verify the intelligibility of the outputs, we further exploit the estimated speech signals on




Figure 7: Example of affinity matrices with jitter problem. In this case, the jitter is happen for 0.4s on the 25th frame, thus diagonal components are disconnected and reappeared from the 35th frame.

Magnitude	Phase	Delay offset			Avg. ↓
		-5	0	5	
Mixture	MX	-	84.91	-	84.91
Ground-truth	GT	-	17.74	-	17.74
Prediction	GT	31.70	32.96	31.90	32.18
Prediction	GL	40.54	39.38	39.74	39.88
Prediction	MX	38.35	37.01	37.73	37.69
Prediction	PR	35.88	35.08	34.83	35.26

Table 4: Automatic speech recognition on the LRS2 dataset.

another task. Specifically, we conduct an additional experiment for automatic speech recognition with enhanced speech signals. To do this, we utilize the speech recognition API of Google cloud system¹ and compute the word error rate (WER) as an automatic metric to evaluate the accuracy of recognition. Firstly, we obtain the WER of 17.74% on the clean ground-truth set, which is the best result that we can achieve in this setup, while the WER on the mixture set is 84.95%. In Tab. 4, CaffNet-C achieves 35.26% error rate when we use separated speech signals with the network setup of phase prediction (PR). It clearly shows that there is no meaningful difference by varying delay offset.

5. Conclusion and Future Work

We presented a novel framework to separate a target speaker’s speech from audio in the wild. To deal with the local matching problem in AVSS, we effectively established a cross-modal affinity between a pair of audio-visual features by modeling relative timing dependencies without supervision. We further suggested a way of affinity regularization to employ global alignment while addressing the label permutation problem. Moreover, we generalized the proposed networks to the complex domain to reconstruct both magnitude and phase masks, enhancing the separation performance. A direction for future work is to explore applications of this generic framework which can be widely plugged into many audio-visual systems.

¹<https://cloud.google.com/speech-to-text>

This research was supported by the Yonsei University Research Fund of 2021(2021-22-0001).

References

- [1] Simon Haykin and Zhe Chen. The cocktail party problem. *Neural computation*, 17(9):1875–1902, 2005.
- [2] Albert S Bregman. *Auditory scene analysis: The perceptual organization of sound*. MIT press, 1994.
- [3] Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdelrahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine*, 29(6):82–97, 2012.
- [4] Dzmitry Bahdanau, Jan Chorowski, Dmitriy Serdyuk, Philemon Brakel, and Yoshua Bengio. End-to-end attention-based large vocabulary speech recognition. In: *ICASSP*, 2016.
- [5] William Chan, Navdeep Jaitly, Quoc V Le, and Oriol Vinyals. Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In: *ICASSP*, 2016.
- [6] Ehsan Variiani, Xin Lei, Erik McDermott, Ignacio Lopez Moreno, and Javier Gonzalez-Dominguez. Deep neural networks for small footprint text-dependent speaker verification. In: *ICASSP*, 2014.
- [7] David Snyder, Daniel Garcia-Romero, Gregory Sell, Daniel Povey, and Sanjeev Khudanpur. X-vectors: Robust dnn embeddings for speaker recognition. In: *ICASSP*, 2018.
- [8] Yoohan Kwon, Soo-Whan Chung, and Hong-Goo Kang. Intra-class variation reduction of speaker representation in disentanglement framework. In: *INTERSPEECH*, 2020.
- [9] Navdeep Jaitly and Geoffrey Hinton. Learning a better representation of speech soundwaves using restricted boltzmann machines. In: *ICASSP*, 2011.
- [10] George Trigeorgis, Fabien Ringeval, Raymond Brueckner, Erik Marchi, Mihalis A Nicolaou, Björn Schuller, and Stefanos Zafeiriou. Adieu features? end-to-end speech emotion recognition using a deep convolutional recurrent network. In: *ICASSP*, 2016.
- [11] John R Hershey, Zhuo Chen, Jonathan Le Roux, and Shinji Watanabe. Deep clustering: Discriminative embeddings for segmentation and separation. In: *ICASSP*, 2016.
- [12] Zhuo Chen, Yi Luo, and Nima Mesgarani. Deep attractor network for single-microphone speaker separation. In: *ICASSP*, 2017.
- [13] Yi Luo and Nima Mesgarani. Tasnet: time-domain audio separation network for real-time, single-channel speech separation. In: *ICASSP*, 2018.
- [14] Chao Weng, Dong Yu, Michael L Seltzer, and Jasha Droppo. Deep neural networks for single-channel multi-talker speech recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(10):1670–1679, 2015.
- [15] 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.
- [16] Morten Kolbæk, Dong Yu, Zheng-Hua Tan, and Jesper Jensen. Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 25(10):1901–1913, 2017.
- [17] W. H. Sumby and Irwin Pollack. Visual contribution to speech intelligibility in noise. *The Journal of the Acoustical Society of America*, 26(2):212–215, 1954.
- [18] Sarah Partan and Peter Marler. Communication goes multi-modal. *Science*, 283(5406):1272–1273, 1999.
- [19] Joon Son Chung, Andrew Senior, Oriol Vinyals, and Andrew Zisserman. Lip reading sentences in the wild. In: *CVPR*, 2017.
- [20] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. The conversation: Deep audio-visual speech enhancement. In: *INTERSPEECH*, 2018.
- [21] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. My lips are concealed: Audio-visual speech enhancement through obstructions. In: *INTERSPEECH*, 2019.
- [22] Triantafyllos Afouras, Andrew Owens, Joon Son Chung, and Andrew Zisserman. Self-supervised learning of audio-visual objects from video. In: *ECCV*, 2020.
- [23] Chenda Li and Yanmin Qian. Deep audio-visual speech separation with attention mechanism. In: *ICASSP*, 2020.
- [24] 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 Transactions on Graphics*, 2018.
- [25] Rui Lu, Zhiyao Duan, and Changshui Zhang. Listen and look: Audio-visual matching assisted speech source separation. *IEEE Signal Processing Letters*, 25(9):1315–1319, 2018.
- [26] Soo-Whan Chung, Hong-Goo Kang, and Joon Son Chung. Seeing voices and hearing voices: Learning discriminative embeddings using cross-modal self-supervision. In: *INTERSPEECH*, 2020.
- [27] Martin Krawczyk and Timo Gerkmann. Stft phase reconstruction in voiced speech for an improved single-channel speech enhancement. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 22(12):1931–1940, 2014.
- [28] Jinkyu Lee and Hong-Goo Kang. A joint learning algorithm for complex-valued t-f masks in deep learning-based single-channel speech enhancement systems. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(6):1098–1108, 2019.
- [29] Yuan-Shan Lee, Chien-Yao Wang, Shu-Fan Wang, Jia-Ching Wang, and Chung-Hsien Wu. Fully complex deep neural network for phase-incorporating monaural source separation. In: *ICASSP*, 2017.
- [30] Chiheb Trabelsi, Olexa Bilaniuk, Ying Zhang, Dmitriy Serdyuk, Sandeep Subramanian, João Felipe Santos, Soroush Mehri, Negar Rostamzadeh, Yoshua Bengio, and Christopher J Pal. Deep complex networks. In: *ICLR*, 2018.

- [31] Hyeong-Seok Choi, Jang-Hyun Kim, Jaesung Huh, Adrian Kim, Jung-Woo Ha, and Kyogu Lee. Phase-aware speech enhancement with deep complex u-net. *arXiv preprint arXiv:1903.03107*, 2019.
- [32] Triantafyllos Afouras, Joon Son Chung, Andrew Senior, Oriol Vinyals, and Andrew Zisserman. Deep audio-visual speech recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [33] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. Lrs3-ted: a large-scale dataset for visual speech recognition. *arXiv preprint arXiv:1809.00496*, 2018.
- [34] Joon Son Chung, Arsha Nagrani, and Andrew Zisserman. Voxceleb2: Deep speaker recognition. In: *INTERSPEECH*, 2018.
- [35] Nima Mesgarani and Edward F Chang. Selective cortical representation of attended speaker in multi-talker speech perception. *Nature*, 485(7397):233–236, 2012.
- [36] John W Fisher III, Trevor Darrell, William Freeman, and Paul Viola. Learning joint statistical models for audio-visual fusion and segregation. In: *NeurIPS*, 2000.
- [37] Zohar Barzelay and Yoav Y Schechner. Harmony in motion. In: *CVPR*, 2007.
- [38] Andrew Owens and Alexei A Efros. Audio-visual scene analysis with self-supervised multisensory features. In: *ECCV*, 2018.
- [39] Soo-Whan Chung, Soyeon Choe, Joon Son Chung, and Hong-Goo Kang. Facefilter: Audio-visual speech separation using still images. In: *INTERSPEECH*, 2020.
- [40] Jen-Cheng Hou, Syu-Siang Wang, Ying-Hui Lai, Yu Tsao, Hsiu-Wen Chang, and Hsin-Min Wang. Audio-visual speech enhancement using multimodal deep convolutional neural networks. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2(2):117–128, 2018.
- [41] Soo-Whan Chung, Joon Son Chung, and Hong-Goo Kang. Perfect match: Improved cross-modal embeddings for audio-visual synchronisation. In: *ICASSP*, 2019.
- [42] Joon Son Chung and Andrew Zisserman. Out of time: automated lip sync in the wild. In: *ACCV*, 2016.
- [43] You Jin Kim, Hee Soo Heo, Soo-Whan Chung, and Bong-Jin Lee. End-to-end lip synchronisation based on pattern classification. In: *SLT*, 2021.
- [44] Chuang Gan, Deng Huang, Hang Zhao, Joshua B Tenenbaum, and Antonio Torralba. Music gesture for visual sound separation. In: *CVPR*, 2020.
- [45] Jiyoung Lee, Sunok Kim, Seungryong Kim, and Kwanghoon Sohn. Multi-modal recurrent attention networks for facial expression recognition. *IEEE Transactions on Image Processing*, 29:6977–6991, 2020.
- [46] Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. Multimodal transformer for unaligned multimodal language sequences. In: *ACL*, 2019.
- [47] Ying Cheng, Ruize Wang, Zhihao Pan, Rui Feng, and Yuejie Zhang. Look, listen, and attend: Co-attention network for self-supervised audio-visual representation learning. In: *ACM MM*, 2020.
- [48] Harry McGurk and John MacDonald. Hearing lips and seeing voices. *Nature*, 264(5588):746–748, 1976.
- [49] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In: *CVPR*, 2018.
- [50] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In: *NeurIPS*, 2015.
- [51] Jungin Park, Jiyoung Lee, Ig-Jae Kim, and Kwanghoon Sohn. Sumgraph: Video summarization via recursive graph modeling. In: *ECCV*, 2020.
- [52] 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.
- [53] Joon Son Chung and Andrew Zisserman. Lip reading in the wild. In: *ACCV*, 2016.
- [54] Joon Son Chung and Andrew Zisserman. Lip reading in profile. In: *BMVC*, 2017.
- [55] Emmanuel Vincent, Rémi Gribonval, and Cédric Févotte. Performance measurement in blind audio source separation. *IEEE Transactions on Audio, Speech, and Language Processing*, 14(4):1462–1469, 2006.
- [56] Antony W Rix, John G Beerends, Michael P Hollier, and Andries P Hekstra. Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs. In: *ICASSP*, 2001.
- [57] Cees H Taal, Richard C Hendriks, Richard Heusdens, and Jesper Jensen. An algorithm for intelligibility prediction of time-frequency weighted noisy speech. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2011.
- [58] Soo-Whan Chung, Joon Son Chung, and Hong-Goo Kang. Perfect match: Self-supervised embeddings for cross-modal retrieval. *IEEE Journal of Selected Topics in Signal Processing*, 14(3):568–576, 2020.