



Target Active Speaker Detection with Audio-visual Cues

Yidi Jiang¹, Ruijie Tao¹, Zexu Pan¹, Haizhou Li^{1,2}

¹National University of Singapore, Singapore

²Shenzhen Research Institute of Big Data, School of Data Science,
The Chinese University of Hong Kong, Shenzhen, China

{yidi_jiang, ruijie.tao, pan_zexu}@u.nus.edu, haizhouli@cuhk.edu.cn

Abstract

In active speaker detection (ASD), we would like to detect whether an on-screen person is speaking based on audio-visual cues. Previous studies have primarily focused on modeling audio-visual synchronization cue, which depends on the video quality of the lip region of a speaker. In real-world applications, it is possible that we can also have the reference speech of the on-screen speaker. To benefit from both facial cue and reference speech, we propose the Target Speaker TalkNet (TS-TalkNet), which leverages a pre-enrolled speaker embedding to complement the audio-visual synchronization cue in detecting whether the target speaker is speaking. Our framework outperforms the popular model, TalkNet on two datasets, achieving absolute improvements of 1.6% in mAP on the AVA-ActiveSpeaker validation set, and 0.8%, 0.4%, and 0.8% in terms of AP, AUC and EER on the ASW test set, respectively. Code is available at <https://github.com/Jiang-Yidi/TS-TalkNet/>.

Index Terms: Active speaker detection, target speaker, audio-visual, speaker recognition

1. Introduction

The study of active speaker detection (ASD) is to determine whether an on-screen person is speaking in each video frame of an audio-visual scene [1]. It plays a critical frontend role for various speech processing tasks, such as audio-visual speaker localization [2], speaker verification [3, 4, 5], speaker extraction [6, 7], speech recognition [8, 9], among others [10, 11].

Typically, ASD methods rely on modeling the synchronization between the audio and visual modalities to detect whether a person is talking [12, 13, 14]. Early approaches focused on short temporal segments of speech activity [15, 16], lip movement [17] and audio-visual synchronization [18, 19, 20] to identify the active speaker. These methods have an efficient structure, but lack the ability to capture the long-term temporal association between the audio and visual signals. The recent work, TalkNet [21], incorporated long-term temporal information and demonstrated strong performance on the ASD task. Following this pipeline, other works attempted to model audio-visual synchronization in a more efficient way, such as Graph Neural Networks (GNN) [22, 23]. However, obtaining high-resolution video of the lip regions is not always feasible, which can hinder the accurate determination of audio-visual synchronization.

Human has an inherent ability to focus on a known speaker in complex social environments, known as selective auditory attention in a “cocktail party” [24, 25, 26]. In other words, when we are already familiar with a person’s voice, i.e., the target speaker, we can perform selective listening by comparing the heard voice with the reference voice.

In target speaker extraction task, the voice of a specific speaker is pre-enrolled in form of a speaker embedding. This

speaker embedding is then used as a reference for target speaker extraction [27, 28]. Similarly, target speaker voice activity detection maps unknown speech to extracted target speaker embeddings to perform speaker diarization task [15, 29]. In [30], speaker diarization task was reformulated to a supervised classification problem by leveraging audio-visual correlation and speaker models for each speaker. These studies highlight the benefits of understanding a target speaker’s voice characteristic in speech processing. Motivated by these studies, we believe that the reference speech could be informative in ASD task as well, in particular, in challenging acoustic environments.

In this work, we propose the Target Speaker TalkNet (TS-TalkNet) framework as an extension of TalkNet [21], introducing the concept of the “target speaker” to the ASD task. TS-TalkNet first enrolls the on-screen speaker as the target speaker and then performs auditory attention with the corresponding pre-enrolled speaker embedding for the ASD task. To achieve this, we construct a face-speaker library where each target speaker’s face is associated with their pre-enrolled reference speech (if exists), providing the target speaker embedding. TS-TalkNet explores the voice characteristic cue to determine whether the audio signals resemble the target speaker embedding, which complements the audio-visual synchronization cue. Noted that if reference speech does not exist, TS-TalkNet only utilizes the synchronization cue. We also explore various multi-modal fusion methods to investigate the efficacy of leveraging speaker embedding in TS-TalkNet. The contributions of this paper can be summarised as follows:

1. We explore the use of reference speech to assist active speaker detection.
2. We propose the TS-TalkNet framework that fully makes use of both audio and visual cues, i.e., voice characteristic signal and audio-visual synchronization information.
3. We conduct experiments on the AVA-ActiveSpeaker (AVA) and the Active Speakers in the Wild (ASW) datasets with improvements: 1.6% in mAP on the AVA validation set; 0.8%, 0.4%, and 0.8% in AP, AUC and EER, respectively, on the ASW test set, confirming the superiority of TS-TalkNet.

2. TS-TalkNet: Target Active Speaker Detection

As depicted in Figure 1, our proposed TS-TalkNet framework takes the cropped face track video, corresponding audio signal and the target speaker’s enrolled speech signal as inputs, and outputs the binary decision of the speech activity for each video frame. TS-TalkNet consists of a feature representation frontend and a speaker detection backend.

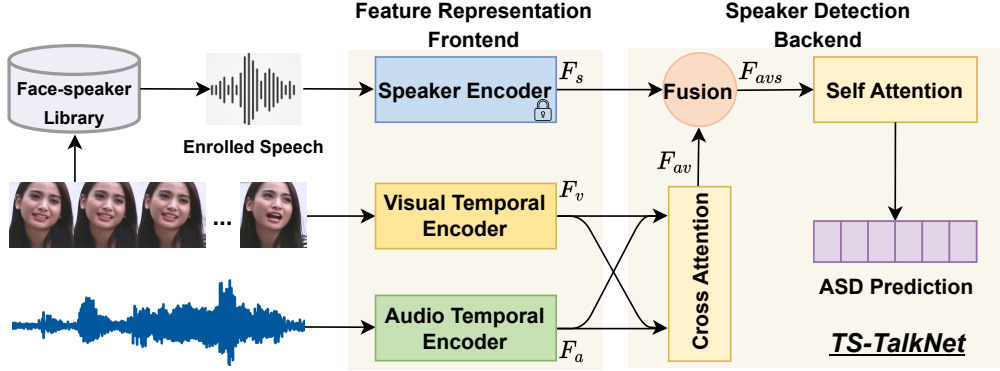


Figure 1: The overview framework of our TS-TalkNet. It consists of a feature representation frontend and a speaker detection backend classifier. The feature representation frontend includes audio and visual temporal encoders, and speaker encoder. The speaker detection backend comprises a cross-attention and a fusion module to combine the audio, visual and speaker embeddings, and a self-attention module to predict the ASD scores. The lock represents the speaker encoder is frozen in our framework.

2.1. Feature representation frontend

The feature representation frontend includes an audio temporal encoder, a visual temporal encoder and a speaker encoder. Following TalkNet, the audio and visual temporal encoders are used to learn audio and visual embeddings from the audio and face track inputs. The speaker encoder is proposed to supplement the target speaker information.

Visual temporal encoder. The visual temporal encoder is designed to learn the long-term representations of facial expression dynamics by encoding the visual stream into a sequence of visual embeddings F_v with consistent time resolution. To achieve this, the visual frontend captures spatial information within each video frame and encodes the video frame stream into a sequence of frame-based embeddings. The visual temporal network then employs a video temporal convolutional block, comprising of five residual connected ReLU, batch normalization, and depth-wise separable convolutional layers, followed by a Conv1D layer that reduces the embedding dimension. The purpose of this encoder is to capture the temporal content within a long-term visual spatio-temporal structure.

Audio temporal encoder. The audio temporal encoder is a ResNet34 network [31] with a squeeze-and-excitation (SE) module [32]. This network extracts audio content representation from the temporal dynamics of a sequence of audio frames, which are initially represented by a vector of Mel-frequency cepstral coefficients (MFCCs). The output of this network is a sequence of audio embeddings, denoted as F_a . F_a matches the time resolution of the visual embeddings, F_v .

Speaker encoder. To begin, we construct the face-speaker library, which contains multiple reference speeches for each target speaker. Further details regarding the construction process can be found in Subsection 3.2.

We designate the individual of each face track input as the ‘target speaker’. Then for each face track, we randomly select one pre-enrolled speech from all the associated speeches in the face-speaker library. The pre-enrolled speech is then used as the input of the speaker encoder. To incorporate target speaker characteristics into our TS-TalkNet, we leverage a pre-trained speaker recognition model to obtain the target speaker embedding. To guarantee robust performance, the ECAPA-TDNN model [33] is used as the speaker encoder. It has demonstrated reliable performance for speaker recognition task.

The ECAPA-TDNN model employs emphasized channel attention to selectively focus on critical parts of the speech sig-

nal, propagating that information through the network and aggregating it to make a final decision. From the variable lengths of input utterances, the output speaker embedding has the fixed dimension. This capability enables the model to handle speech signals in various scenarios and generate robust speaker representations more effectively.

From the enrolled speech input, we obtain the target speaker embedding F_s and convey the target speaker information to the speaker detection backend. We freeze the parameters of this module in our framework, since our purpose is to obtain robust embedding for the target speaker.

2.2. Speaker detection backend

The speaker detection backend comprises a cross-attention module to achieve audio-visual synchronization, a fusion module to combine the speaker embedding with audio-visual embeddings, in addition to a self-attention module to capture speaking activities from the temporal context at the utterance level. We aim to combine three embeddings generated from the feature representation frontend: F_a for audio modality, F_v for visual modality, and F_s for speaker characteristic, to provide a comprehensive representation of the speaking activities for ASD prediction.

Cross-attention module. Firstly, we use a cross-attention structure along the temporal dimension to achieve audio-visual synchronization and interaction. The inputs are the vectors of query (Q_a, Q_v), key (K_a, K_v), and value (V_a, V_v) from audio and visual embeddings, F_a and F_v respectively, projected by a linear layer. As formulated in Equation 1 and Equation 2, the outputs are the audio attention embedding $F_{a \rightarrow v}$ and visual attention embedding $F_{v \rightarrow a}$, where d denotes the dimension of Q, K and V . The outputs are concatenated together along the temporal direction as audio-visual embedding F_{av} .

$$F_{a \rightarrow v} = \text{softmax}\left(\frac{Q_v K_a^T}{\sqrt{d}}\right) V_a \quad (1)$$

$$F_{v \rightarrow a} = \text{softmax}\left(\frac{Q_a K_v^T}{\sqrt{d}}\right) V_v \quad (2)$$

Fusion module. The fusion module is used to combine the speaker embedding with the audio-visual embedding to obtain the overall audio-visual-speaker representation F_{avs} .

Noted that F_s is not a temporal embedding, while F_{av} , $F_{v \rightarrow a}$ and $F_{a \rightarrow v}$ are all temporal embedding sequences. For alignment, we replicate F_s along the time dimension to \hat{F}_s to implement the fusion module.

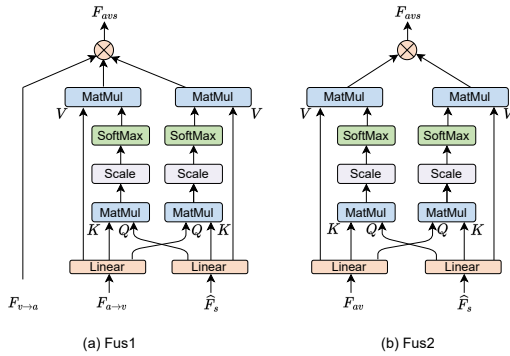


Figure 2: The different fusion approaches for audio, visual and speaker embeddings. We denote the structure (a) and (b) as ‘Fus1’ and ‘Fus2’, respectively. \otimes represents embedding concatenation along the time dimension.

We have investigated three different fusion structures to demonstrate that our idea of leveraging voice characteristics isn’t restricted to a specific network architecture. As shown in Fig 2 (a) and (b), ‘Fus1’ structure uses $F_{a \to v}$ and \hat{F}_s as inputs for cross-attention. ‘Fus2’ structure applies F_{av} to achieve the cross-attention process with \hat{F}_s . The third structure is to concatenate the \hat{F}_s and F_{av} along the temporal dimension, which is denoted as ‘Concat’.

Self-attention module. After obtaining the audio-visual-speaker representation F_{avs} , we follow the same strategy as TalkNet [21] and apply a self-attention structure to model F_{avs} temporal information to distinguish between speaking and non-speaking frames. This structure is similar with the cross-attention structure, except that the query, key and value in the attention layer all come from the joint embedding F_{avs} .

2.3. Loss function

Treating ASD as a frame-level classification task, we project the output of the self-attention structure to an ASD label sequence using a fully connected layer and the softmax operation. We then compare the predicted label sequence to the ground truth label sequence using cross-entropy loss, which is given in Equation 3. Here, \hat{y}_i and y_i represent the predicted and ground truth ASD labels for the i^{th} video frame, respectively, where $i \in [1, T]$. The value of T represents the number of video frames.

$$\mathcal{L} = -\frac{1}{T} \sum_{i=1}^T (y_i \cdot \log \hat{y}_i) + (1 - y_i) \cdot \log(1 - \hat{y}_i) \quad (3)$$

3. Experiments

In this section, we describe the datasets and experimental details to evaluate the proposed TS-TalkNet framework.

3.1. Dataset

AVA-ActiveSpeaker (AVA) dataset [1] is a large-scale audio-visual active speaker detection dataset. It consists of 262 videos extracted from Hollywood movies, with 120 for training, 33 for validation, and 109 for testing. This dataset contains about 3.65 million human-labeled video frames or 38.5 hours of face tracks, and the corresponding audio.

The Active Speakers in the Wild (ASW) dataset [34] consists of 212 videos randomly selected from the VoxConverse dataset [35]. It contains 13.4 hours videos in which 56.7% are active in the training set, 9.6 hours videos where 60.4% are ac-

Table 1: Comparison with the state of the arts on the AVA val set in terms of mAP. The bold rows represent the results of our proposed methods.

| Method | mAP (%) \uparrow |
|----------------------------|--------------------|
| Roth et al. [13, 1] | 79.2 |
| Zhang et al. [14] | 84.0 |
| MAAS-LAN [13] | 85.1 |
| Alcazar et al. [37] | 87.1 |
| Chung et al [19] | 87.8 |
| MAAS-TAN [13] | 88.8 |
| UniCon [38] | 92.0 |
| TalkNet [21] | 92.3 |
| ASDNet [39] | 93.5 |
| Light-ASD [40] | 94.1 |
| UniCon+ [41] | 94.5 |
| SPELL+ [23] | 94.9 |
| TS-TalkNet (Fus1) | 93.3 |
| TS-TalkNet (Fus2) | 93.5 |
| TS-TalkNet (Concat) | 93.9 |

tive in the validation set, and 7.9 hours videos where 57.0% are active in the test set.

3.2. Face-speaker library

We pre-process the dataset to build the face-speaker library by searching the face tracks from the same person using a face recognition module. We utilize the pre-trained ResNet50 model on the Glint360K dataset [36] as the face recognition module and set the face similarity threshold at 0.7. The enrolled speeches come from the corresponding audio of the found tracks, which are located by extracting the active segments using ground-truth labels. Then for each input, we can retrieve the pre-enrolled speeches of the target speaker from the face-speaker library according to the input faces. If the reference speeches exist, we randomly choose one as the target speaker enrolled speech and feed it into TS-TalkNet. Otherwise, we set the enrolled speech to zero vectors to process detection. In this way, our TS-TalkNet can handle different situations, whether there is an enrolled speech record or not.

3.3. Implementation details

The proposed TS-TalkNet is implemented by PyTorch using the Adam optimizer. We set the initial learning rate to 10^{-4} and decrease it by 5% for each epoch. The MFCC dimension is set to 13, and all faces are resized into 112×112 pixels. Both audio and visual embeddings are set to 128 dimensions. For both cross-attention and self-attention structures, a single transformer layer with 8 attention heads is used. We apply visual augmentation such as random flipping, rotating, and cropping the original images to improve framework performance.

3.4. Evaluation metric

For the AVA dataset, we use the official ActivityNet evaluation tool to compute mean average precision (mAP) and evaluate on the validation (val) set. For the ASW dataset, following [34, 21], we compute three metrics: mean precision (AP), area under the receiver operating characteristic (AUC), and equal error rate (EER) using sklearn package and evaluate on the ASW val and test set.

4. Results and Analysis

In this section, we present the performance comparison of our proposed TS-TalkNet with state-of-the-art methods on the AVA

Table 2: Comparison with the state of the arts on the ASW val and test set in terms of AP, AUC and EER. The bold rows represent the results of our proposed methods.

| Set | Method | AP (%)↑ | AUC (%)↑ | EER (%)↓ |
|------|----------------------------|-------------|-------------|------------|
| Val | TalkNet [21] | 96.4 | 98.2 | 6.0 |
| | TS-TalkNet (Fus1) | 97.3 | 98.6 | 5.4 |
| | TS-TalkNet (Fus2) | 97.5 | 98.7 | 5.1 |
| | TS-TalkNet (Concat) | 97.7 | 98.7 | 5.1 |
| Test | RothNet [1] | 89.7 | - | - |
| | SyncNet [20] | 92.4 | - | - |
| | ASW-SSL [34] | 90.5 | - | - |
| | LoCoNet [42] | 93.4 | 95.1 | 9.8 |
| | ASW-BGRUs [34] | 96.6 | 97.2 | 6.2 |
| | TalkNet [21] | 97.7 | 98.6 | 5.1 |
| | TS-TalkNet (Fus1) | 98.3 | 98.9 | 4.8 |
| | TS-TalkNet (Fus2) | 98.3 | 98.9 | 4.4 |
| | TS-TalkNet (Concat) | 98.5 | 99.0 | 4.3 |

Table 3: Results comparison of different percentages of the frames with active labels on the ASW val set. For example, ‘0-20’ represents the tracks with less than 20% active labels.

| Active Frame (%) | Method | AP (%)↑ | AUC (%)↑ | EER (%)↓ |
|------------------|-------------------|--------------|--------------|--------------|
| ‘0 - 20’ | TalkNet | 8.29 | 95.92 | 9.45 |
| | TS-TalkNet | 11.54 | 96.65 | 7.63 |
| ‘20 - 40’ | TalkNet | 83.47 | 94.85 | 11.22 |
| | TS-TalkNet | 89.37 | 96.06 | 9.41 |
| ‘40 - 60’ | TalkNet | 92.54 | 93.33 | 13.95 |
| | TS-TalkNet | 93.39 | 93.94 | 13.35 |
| ‘60 - 80’ | TalkNet | 96.95 | 93.16 | 14.37 |
| | TS-TalkNet | 97.30 | 93.99 | 12.97 |
| ‘80 - 100’ | TalkNet | 99.71 | 93.44 | 13.89 |
| | TS-TalkNet | 99.75 | 94.01 | 13.28 |

and ASW datasets to show its efficiency. We also explore the architectures of the fusion module and analyze the results to find out the improvement under the different test conditions.

4.1. AVA results

We firstly report the performance of TS-TalkNet on the AVA val dataset in Table 1. With the inclusion of enrolled speech, we observe that TS-TalkNet (Concat) achieves 93.9% mAP and 1.6% mAP improvements over TalkNet on the AVA val set. Moreover, for TS-TalkNet (Fus1) and TS-TalkNet (Fus2) frameworks, we obtain the improvements by 1.0% and 1.2% mAP, respectively.

4.2. ASW results

The results of TS-TalkNet on the ASW val and test dataset are reported in Table 2. Specially, the proposed TS-TalkNet (Concat) framework outperforms TalkNet by 1.3%, 0.6% and 0.9% in terms of AP, AUC and EER, respectively, on the ASW val set; and 0.8%, 0.4% and 0.8% in terms of AP, AUC and EER, respectively, on the ASW test set. Similarly, both TS-TalkNet (Fus1) and TS-TalkNet (Fus2) frameworks outperform TalkNet by a large margin.

We visualize the ASD results in Figure 3 for two videos from the ASW dataset. Our results verify the benefits of speaker embedding on ASD task and demonstrate the superiority of TS-TalkNet.

4.3. Qualitative studies

Fusion structure. We implement three fusion approaches to combine the audio-visual embeddings and speaker embedding,

as mentioned in Subsection 2.2. As shown in Table 2, for TS-TalkNet (Fus1), we obtain the improvements by 0.9%, 0.4% and 0.6% in terms of AP, AUC and EER, respectively, over the TalkNet on the ASW val set; and 0.6%, 0.3% and 0.3% in terms of AP, AUC and EER, respectively, on the ASW test set. For TS-TalkNet (Fus2), we achieve the improvements by 1.1%, 0.5% and 0.9% in terms of AP, AUC and EER, respectively, over TalkNet on the ASW val set; and 0.6%, 0.3% and 0.7% in terms of AP, AUC and EER respectively on the ASW test set. Similar observations can be found in Table 1.

These results indicate that incorporating the speaker embedding can improve the performance of TS-TalkNet, regardless of the fusion approaches. Therefore, our idea of using speaker characteristics to assist ASD is not restricted to a specific network architecture.

Analysis study. We aim to investigate the effect of speaker embedding incorporation for face tracks with varying proportions of active frames. So we analyze the performance improvement of TS-TalkNet on the ASW val set by dividing the instances into five sections based on the percentage of the frames with active labels of each track input. As shown in Table 3, TS-TalkNet can improve the detection accuracy for each section, indicating that the speaker-specific characteristics are valuable for detecting the target speaker’s active frames across different scenarios.




Figure 3: The results of TS-TalkNet for real-world videos with one person (a) and multiple persons (b) on the screen. The green box denotes the active speaker. The red box denotes the inactive speaker. As demonstrated in (a), despite the occlusion of the lip in certain video frames, which compromises the audio-visual synchronization cue, TS-TalkNet can still accurately detect the speaker with the aid of speaker embedding complementation.

5. Conclusion

In this paper, we propose the Target Speaker TalkNet (TS-TalkNet) for active speaker detection by incorporating target speaker embeddings. Our experiments and studies demonstrate that speaker characteristic plays a crucial role in solving the ASD problem. Our work provides a new perspective and potential solution for improving the performance of ASD systems in complicated acoustic environments. In future work, we can investigate the potential applications of TS-TalkNet in other speech-related tasks.

6. Acknowledgements

This work is funded by 1) Huawei Noah’s Ark Lab; 2) National Natural Science Foundation of China (Grant No. 62271432); 3) Agency for Science, Technology and Research (A*STAR) under its AME Programmatic Funding Scheme (Project No. A18A2b0046)

7. References

- [1] J. Roth, S. Chaudhuri, O. Klejch, R. Marvin, A. Gallagher, L. Kaver, S. Ramaswamy, A. Stopczynski, C. Schmid, Z. Xi *et al.*, “AVA active speaker: An audio-visual dataset for active speaker detection,” in *ICASSP*, 2020, pp. 4492–4496.
- [2] X. Qian, M. Madhavi, Z. Pan, J. Wang, and H. Li, “Multi-target DoA estimation with an audio-visual fusion mechanism,” in *ICASSP*, 2021, pp. 4280–4284.
- [3] R. Tao, K. A. Lee, R. K. Das, V. Hautamäki, and H. Li, “Self-supervised training of speaker encoder with multi-modal diverse positive pairs,” *arXiv preprint arXiv:2210.15385*, 2022.
- [4] T. Liu, R. K. Das, K. A. Lee, and H. Li, “MFA: TDNN with multi-scale frequency-channel attention for text-independent speaker verification with short utterances,” in *ICASSP*, 2022, pp. 7517–7521.
- [5] —, “Neural acoustic-phonetic approach for speaker verification with phonetic attention mask,” *IEEE Signal Processing Letters*, pp. 782–786, 2022.
- [6] T. Ochiai, M. Delcroix, K. Kinoshita, A. Ogawa, and T. Nakatani, “Multimodal speakerbeam: Single channel target speech extraction with audio-visual speaker clues,” in *INTERSPEECH*, 2019, pp. 2718–2722.
- [7] Z. Pan, X. Qian, and H. Li, “Speaker extraction with co-speech gestures cue,” *IEEE Signal Processing Letter*, pp. 1467–1471, 2022.
- [8] T. Afouras, J. S. Chung, A. Senior, O. Vinyals, and A. Zisserman, “Deep audio-visual speech recognition,” *IEEE transactions on pattern analysis and machine intelligence*, pp. 8717–8727, 2018.
- [9] J. Wang, X. Qian, and H. Li, “Predict-and-update network: Audio-visual speech recognition inspired by human speech perception,” *arXiv preprint arXiv:2209.01768*, 2022.
- [10] X. Yang, J. Ye, and X. Wang, “Factorizing knowledge in neural networks,” in *European Conference on Computer Vision*, 2022.
- [11] X. Yang, D. Zhou, S. Liu, J. Ye, and X. Wang, “Deep model reassembly,” in *Advances in Neural Information Processing Systems*, 2022.
- [12] B. Pouthier, L. Pilati, L. Gudupudi, C. Bouveyron, and F. Precioso, “Active speaker detection as a multi-objective optimization with uncertainty-based multimodal fusion,” in *INTERSPEECH*, 2021, pp. 2381–2385.
- [13] J. L. Alcázar, F. Caba, A. K. Thabet, and B. Ghanem, “MAAS: Multi-modal assignation for active speaker detection,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 265–274.
- [14] Y.-H. Zhang, J. Xiao, S. Yang, and S. Shan, “Multi-task learning for audio-visual active speaker detection,” *The ActivityNet Large-Scale Activity Recognition Challenge*, 2019.
- [15] S. Ding, Q. Wang, S.-Y. Chang, L. Wan, and I.-L. Moreno, “Personal VAD: Speaker-conditioned voice activity detection,” in *Proc. Odyssey The Speaker and Language Recognition Workshop*, 2020, pp. 433–439.
- [16] I. Medennikov, M. Korenevsky, T. Prisyach, Y. Khokhlov, M. Korenevskaya, I. Sorokin, T. Timofeeva, A. Mitrofanov, A. Andrusenko, I. Podluzhny *et al.*, “Target-speaker voice activity detection: A novel approach for multi-speaker diarization in a dinner party scenario,” in *INTERSPEECH*, 2020, pp. 274–278.
- [17] F. Patrona, A. Iosifidis, A. Tefas, N. Nikolaidis, and I. Pitas, “Visual voice activity detection in the wild,” *IEEE Transactions on Multimedia*, pp. 967–977, 2016.
- [18] T. Afouras, A. Owens, J. S. Chung, and A. Zisserman, “Self-supervised learning of audio-visual objects from video,” in *European Conference on Computer Vision*, 2020, pp. 208–224.
- [19] J. S. Chung, “Naver at ActivityNet Challenge 2019–Task B Active Speaker Detection (AVA),” *arXiv preprint arXiv:1906.10555*, 2019.
- [20] J. S. Chung and A. Zisserman, “Out of time: automated lip sync in the wild,” in *ACCV 2016 Workshops*, 2017, pp. 251–263.
- [21] R. Tao, Z. Pan, R. K. Das, X. Qian, M. Z. Shou, and H. Li, “Is someone speaking? exploring long-term temporal features for audio-visual active speaker detection,” in *ACM Multimedia*, 2021, pp. 3927–3935.
- [22] J. L. Alcázar, M. Cordes, C. Zhao, and B. Ghanem, “End-to-end active speaker detection,” in *European Conference on Computer Vision*, 2022, pp. 126–143.
- [23] K. Min, S. Roy, S. Tripathi, T. Guha, and S. Majumdar, “Learning long-term spatial-temporal graphs for active speaker detection,” in *European Conference on Computer Vision*, 2022, pp. 371–387.
- [24] A. W. Bronkhorst, “The cocktail party phenomenon: A review of research on speech intelligibility in multiple-talker conditions,” *Acta Acustica united with Acustica*, pp. 117–128, 2000.
- [25] —, “The cocktail-party problem revisited: early processing and selection of multi-talker speech,” *Attention, Perception, & Psychophysics*, pp. 1465–1487, 2015.
- [26] N. Mesgarani and E. F. Chang, “Selective cortical representation of attended speaker in multi-talker speech perception,” *Nature*, pp. 233–236, 2012.
- [27] M. Ge, C. Xu, L. Wang, E. S. Chng, J. Dang, and H. Li, “Spex+: A complete time domain speaker extraction network,” in *INTERSPEECH*, 2020.
- [28] Z. Pan, R. Tao, C. Xu, and H. Li, “Selective listening by synchronizing speech with lips,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 1650–1664, 2022.
- [29] Z. Chen, B. Han, S. Wang, and Y. Qian, “Attention-based encoder-decoder network for end-to-end neural speaker diarization with target speaker attractor,” in *INTERSPEECH*, 2023.
- [30] J. S. Chung, B.-J. Lee, and I. Han, “Who said that?: Audio-visual speaker diarisation of real-world meetings,” in *INTERSPEECH*, 2019, pp. 371–375.
- [31] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *CVPR*, 2016, pp. 770–778.
- [32] J. Hu, L. Shen, and G. Sun, “Squeeze-and-excitation networks,” in *CVPR*, 2018, pp. 7132–7141.
- [33] B. Desplanques, J. Thienpondt, and K. Demuyne, “ECAPA-TDNN: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification,” *INTERSPEECH*, 2020.
- [34] Y. J. Kim, H.-S. Heo, S. Choe, S.-W. Chung, Y. Kwon, B.-J. Lee, Y. Kwon, and J. S. Chung, “Look who’s talking: Active speaker detection in the wild,” *INTERSPEECH*, 2021.
- [35] J. S. Chung, J. Huh, A. Nagrani, T. Afouras, and A. Zisserman, “Spot the conversation: Speaker diarisation in the wild,” in *INTERSPEECH*, 2020, pp. 299–303.
- [36] X. An, J. Deng, J. Guo, Z. Feng, X. Zhu, Y. Jing, and L. Tongliang, “Killing Two Birds with One Stone: Efficient and Robust Training of Face Recognition CNNs by Partial FC,” in *CVPR*, 2022, pp. 4042–4051.
- [37] J. L. Alcázar, F. Caba, L. Mai, F. Perazzi, J.-Y. Lee, P. Arbeláez, and B. Ghanem, “Active speakers in context,” in *CVPR*, 2020, pp. 12465–12474.
- [38] Y. Zhang, S. Liang, S. Yang, X. Liu, Z. Wu, S. Shan, and X. Chen, “UniCon: Unified context network for robust active speaker detection,” in *ACM Multimedia*, 2021, pp. 3964–3972.
- [39] O. Köpüklü, M. Taseska, and G. Rigoll, “How to design a three-stage architecture for audio-visual active speaker detection in the wild,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 1193–1203.
- [40] J. Liao, H. Duan, K. Feng, W. Zhao, Y. Yang, and L. Chen, “A light weight model for active speaker detection,” *arXiv preprint arXiv:2303.04439*, 2023.
- [41] Y. Zhang, S. Liang, S. Yang, and S. Shan, “Unicon+: Ictcas-ucas submission to the ava-activespeaker task at activitynet challenge 2022,” *arXiv preprint arXiv:2206.10861*, 2022.
- [42] X. Wang, F. Cheng, G. Bertasius, and D. Crandall, “LoCoNet: Long-short context network for active speaker detection,” *arXiv preprint arXiv:2301.08237*, 2023.