

UniCon: Unified Context Network for Robust Active Speaker Detection

Yuanhang Zhang*[†]
Key Laboratory of Intelligent
Information Processing, Institute of
Computing Technology, Chinese
Academy of Sciences (CAS)
Beijing, China
zhangyuanhang15@mails.ucas.ac.cn

Susan Liang*[†]
Key Laboratory of Intelligent
Information Processing, Institute of
Computing Technology, CAS
Beijing, China
liangsusan18@mails.ucas.ac.cn

Shuang Yang[†]
Key Laboratory of Intelligent
Information Processing, Institute of
Computing Technology, CAS
Beijing, China
shuang.yang@ict.ac.cn

Xiao Liu
Tomorrow Advancing Life
Beijing, China
ender.liux@gmail.com

Zhongqin Wu
Tomorrow Advancing Life
Beijing, China
30388514@qq.com

Shiguang Shan[†]
Key Laboratory of Intelligent
Information Processing, Institute of
Computing Technology, CAS
Beijing, China
sgshan@ict.ac.cn

Xilin Chen[†]
Key Laboratory of Intelligent
Information Processing, Institute of
Computing Technology, CAS
Beijing, China
xlchen@ict.ac.cn

ABSTRACT

We propose a new efficient framework, the Unified Context Network (UniCon), for robust active speaker detection (ASD). Traditional methods for ASD usually operate on each candidate's pre-cropped face track separately and do not sufficiently consider the relationships among the candidates. This potentially limits performance, especially in challenging scenarios with low-resolution faces, multiple candidates, etc. Our solution is a novel, unified framework that focuses on jointly modeling multiple types of *contextual* information: spatial context to indicate the position and scale of each candidate's face, relational context to capture the visual relationships among the candidates and contrast audio-visual affinities with each other, and temporal context to aggregate long-term information and smooth out local uncertainties. Based on such information, our model optimizes all candidates in a unified process

for robust and reliable ASD. A thorough ablation study is performed on several challenging ASD benchmarks under different settings. In particular, our method outperforms the state-of-the-art by a large margin of about 15% mean Average Precision (mAP) absolute on two challenging subsets: one with three candidate speakers, and the other with faces smaller than 64 pixels. Together, our UniCon achieves 92.0% mAP on the AVA-ActiveSpeaker validation set, surpassing 90% for the first time on this challenging dataset at the time of submission. Project website: <https://unicon-asd.github.io/>.

CCS CONCEPTS

• **Computing methodologies** → **Activity recognition and understanding**.

KEYWORDS

active speaker detection, audio-visual, speech, computer vision

ACM Reference Format:

Yuanhang Zhang, Susan Liang, Shuang Yang, Xiao Liu, Zhongqin Wu, Shiguang Shan, and Xilin Chen. 2021. UniCon: Unified Context Network for Robust Active Speaker Detection. In *Proceedings of the 29th ACM International Conference on Multimedia (MM '21), October 20–24, 2021, Virtual Event, China*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3474085.3475275>

1 INTRODUCTION

Active Speaker Detection (ASD) refers to the task of identifying when each visible person is speaking in a video, typically through

*Both authors contributed equally to this research.

[†]Also with School of Computer Science and Technology, University of Chinese Academy of Sciences.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '21, October 20–24, 2021, Virtual Event, China

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8651-7/21/10...\$15.00

<https://doi.org/10.1145/3474085.3475275>

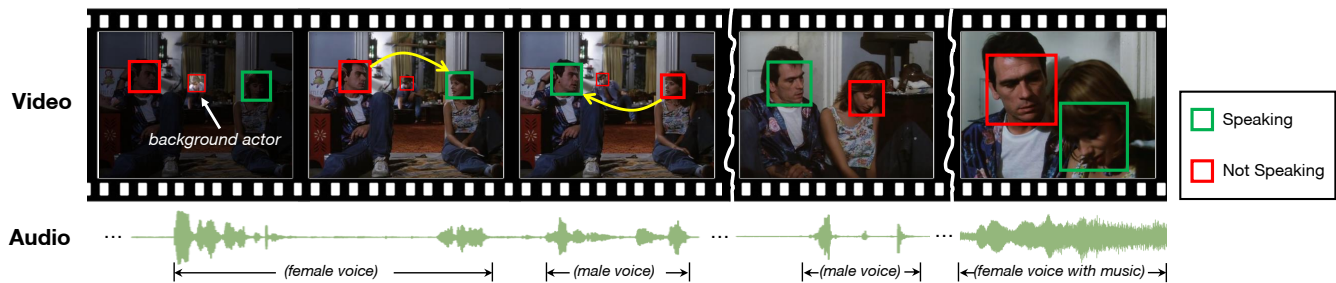


Figure 1: What marks an active speaker out from others? Admittedly, face motion and its synchrony with the audio are the most obvious clues; however, as shown in the figure above, the underlying signals can be highly ambiguous, especially in hard scenarios with poorly lit, low-resolution faces and noisy acoustics, etc. We observe that three additional types of *contextual* clues can help identify the active speaker: (a) higher correspondence between his/her face track and the audio relative to other candidates, (b) a form of contrast with the remaining candidates in terms of visual saliency, the amount of attention he/she receives from others, etc. and (c) temporal continuity. In this paper, we propose a novel Unified Context Network (UniCon) which models such context information to jointly optimize all candidates for robust active speaker detection.

careful joint analysis of face motion and voices. It has a wide range of modern practical applications, such as video re-targeting [14], speech diarization [9], automatic video indexing, human-robot interaction [6], speech enhancement, and the automatic curation of large-scale audio-visual speech datasets [10, 18, 19].

Although research in this direction dates back almost two decades [13, 39], most early work only experiment with simple, short sequences depicting frontal faces, which is a major simplification of practical settings. Recent developments in deep learning for audio-visual tasks have led to renewed interest in the ASD problem [11, 26]. In 2020, Roth *et al.* [35] introduced AVA-ActiveSpeaker, the first large-scale, challenging ASD benchmark extracted from diverse YouTube movies, which significantly boosted subsequent research. Several effective methods have been proposed, building upon 2D and 3D convolutional neural networks (CNNs) and recurrent neural networks (RNNs) [2, 8, 35, 42]. Among them, most closely related to our work is the Active Speaker Context (ASC) model proposed in 2020 [2], which combines cropped face tracks from all candidates with the audio using a self-attention module [40] to model relationships among the candidates. However, relationships learned in this manner are weak, and lack clear interpretation. In addition, like most existing methods, the ASC model also focuses on optimizing each candidate separately. As a consequence, in complex scenarios with multiple candidates, low-resolution faces, occlusion, noisy acoustics etc., current ASD models often lack robustness and cannot yield satisfactory results.

In this paper, we present a fundamentally different approach where we jointly model multiple sources of *contextual* information and simultaneously optimize all candidates in the scene in a unified process for robust ASD. As shown in Fig. 1, we observe that when face motion and audio signals are ambiguous, three other types of contextual clues often help mark the true active speaker out from others: (a) *higher* correspondence between his/her face track and the audio relative to others, e.g. hearing a male’s voice helps eliminate female candidates; (b) a certain form of contrast with the remaining candidates in terms of his/her visual saliency, the amount of attention he/she received from others, etc., for instance, one person speaks while the rest look at him/her; and (c) temporal continuity, *i.e.* by observing each candidate’s behavior in a temporal

neighborhood, we can smooth out erroneous predictions that arise from instantaneous frame-level decisions.

In light of these observations, we propose a novel framework that seeks to resolve local ambiguities and achieve robust ASD by jointly modeling different kinds of contextual evidence. Specifically, we introduce face position and scale information as global *spatial context* to implicitly reflect the visual saliency and the visual focus of attention (VFoA) of each candidate in the scene. Building upon this information, we then construct a powerful and efficient multi-speaker *relational context* where each candidate is contrasted with others from both visual and audio-visual perspectives. Finally, we integrate *temporal context* into the previous relational context to aggregate long-term information and remove noises in frame-level predictions. By incorporating spatial, relational, and temporal context into a unified framework, our method can jointly optimize all candidates in the scene end-to-end. In the experiments, we show that our method significantly improves ASD performance, especially in challenging scenarios.

In summary, the main contributions of our work are:

- We propose a well-motivated, compact framework, the Unified Context Network (UniCon) for robust ASD, which unifies spatial, relational, and temporal context to jointly optimize all candidates in the scene.
- We highlight the use of face position and scale in constructing each candidate’s global spatial context, and explore the modeling of the relationships among the candidates.
- We conduct thorough ablation studies that prove the effectiveness of our method, particularly under challenging multi-speaker and low-resolution settings. Moreover, we outperform the current state-of-the-art by a large margin on challenging ASD benchmarks.

2 RELATED WORK

Active speaker detection (ASD). Current approaches for ASD can be summarized into two paradigms: unsupervised and supervised. *Unsupervised* methods usually operate under one of two assumptions: (a) the active speaker’s face is the one with the highest correspondence between the audio and its visual information, or (b) the active speaker’s face is the one that co-occurs most frequently with the person’s idiosyncratic voice. Under assumption (a), one

typical practice is to learn an audio-visual two-stream network using a contrastive loss between the two modalities [11]. During training, pairs of temporally aligned audio and visual signals are pulled close, while misaligned ones are pushed apart. During inference, the active speaker is naturally the one with lowest distance between audio and visual features. Recently, the method has been improved by better training objective designs, such as multi-way classification [12] and multinomial loss [17]. Methods based on assumption (b) require strong face and voice embeddings for robust clustering and matching [5, 26]. For example, Hoover *et al.* [26] first detect speech segments and face tracks, and cluster them independently using pre-trained speech embeddings and FaceNet [36]. They then apply bipartite matching to assign each speech cluster to the face track cluster with maximum temporal overlap. However, unsupervised methods are not robust enough when the assumptions are not met or clustering results are noisy. *Supervised* methods usually formulate ASD as a binary classification problem, and directly learn a frame-level binary classifier on fused audio-visual features [8, 35]. Some works [27, 42] apply multi-task learning to boost performance, by introducing an additional contrastive loss on audio and visual features. An interesting recent work [2] constructs an Active Speaker Context ensemble to improve local predictions by querying adjacent time intervals and other speakers in the scene. Our work is a new and fundamentally different attempt in this direction, which proposes a methodical process to better incorporate different types of meaningful context information. We leverage such information in a unified framework to optimize all candidates jointly rather than independently for more robust ASD.

Context-aware action recognition and detection. There has been a large body of work that focuses on how to integrate *contextual* information in the field of action recognition and spatio-temporal action detection. For example, [23] uses a modified Transformer [40] architecture to aggregate context from other people to recognize the target person’s actions. [34] proposes a high-order relation reasoning operator that enhances action recognition by modeling Actor-Context-Actor relations. [41] improves spatio-temporal action detection by expanding actor bounding boxes to include more background context. We draw inspiration from these successful works and explore the use of contextual information in ASD.

Gaze and social interaction. Evidence from psychology suggests that gaze direction plays an important role in social interactions [30], such as conversations. Since the VFoA of conversation participants typically converge on the active speaker, some previous work on speaker turn detection [29] and audio-visual diarization of meetings [16, 20, 21] have reported good results by explicitly modeling head pose or gaze information with ad-hoc modules. In this work, we implicitly utilize such information through our spatial and relational contexts, which model the relationships among the candidates in an efficient and unified process.

3 THE UNICON MODEL

In this section, we describe our UniCon model, which integrates spatial, relational, and temporal context in a unified framework. First, for each candidate, the scale and position of all candidates’ faces are introduced as global *spatial* context to complement facial information and help learn the relationships among the speakers. Each candidate is then contrasted with others from both a visual and

audio-visual perspective in the *relational* context modeling component. To further improve the robustness of the model’s predictions, *temporal* context is integrated. Finally, based on the aggregated contextual features, we generate speaker activity predictions for all speakers simultaneously with a shared prediction layer. Fig. 2 provides an overview of our proposed approach.

3.1 Encoders

Given an input video clip, we first crop face tracks for each candidate speaker and transform them into low-dimensional speaker descriptors for further analysis. Likewise, we encode each audio frame into a low-dimensional audio descriptor.

Face Track Encoder: To encode short-term temporal dynamics, at each time step $t = 1, 2, \dots, T$, where T is the total number of time steps (i.e. frames), the i -th candidate’s input ($i = 1, 2, \dots, N$ where N is the number of visible candidates at time step t) is a stack of k consecutive face crops $\mathbf{x}_i^t = (x_i^{t-[k/2]}, \dots, x_i^{t+[k/2]})$ between step $t - [k/2]$ and $t + [k/2]$. Then an encoder ϕ , which is a ResNet-18 [25] is used to produce an average-pooled feature vector $\phi(\mathbf{x}_i^t) \in \mathbb{R}^d$ over the k face crops. To keep the computational cost manageable, we choose $k = 5$ for our experiments and reduce the dimensionality of each $\phi(\mathbf{x}_i^t)$ to $d' = 128$ using a single shared fully-connected layer to obtain each candidate’s final face track features:

$$v_i = (\text{FC}(\phi(\mathbf{x}_i^1)), \text{FC}(\phi(\mathbf{x}_i^2)), \dots, \text{FC}(\phi(\mathbf{x}_i^T))) \in \mathbb{R}^{T \times d'}$$

Audio Encoder: At each time step t ($t = 1, 2, \dots, T$), we obtain audio representations from a 400ms window preceding t . A ResNet-18 encoder ψ takes 13-dimensional MFCCs of the window as input and outputs a 512-dimensional average-pooled feature. These features are also dimension-reduced to 128 using a single fully-connected layer. We denote the final audio features by $a = (a_1, a_2, \dots, a_T)$.

3.2 Spatial Context

The purpose of modeling spatial context is twofold. First, active speakers usually occupy a central position in the scene and depict higher visual saliency, especially in movies (commonly known as "the language of the lens"). Hence knowing the scale, position, and trajectory of a face in the video can help eliminate unlikely candidates. Second, people tend to look at the active speaker as they listen. Therefore, we wish to reflect such gaze-related information to some degree, by providing the model with both facial features and the relative positions of the candidates in the scene.

Specifically, we encode head positions of all candidates in the scene using 64×64 coordinate-normalized maps of 2D Gaussians, which is motivated by [33]. The position and radius of each Gaussian represent the relative position and size of each candidate’s face. Next, as shown in Fig. 3, to further indicate the candidates’ relationships, for each candidate i ($i = 1, 2, \dots, N$) we construct his/her person-specific head map H_i by generating a color-coded version of the initial Gaussian map in the following principle: yellow denotes candidate i , and blue denotes other candidates in the scene. To further facilitate the subsequent modeling of relationships between candidate i and every other candidate j ($j = 1, 2, \dots, N, j \neq i$), we tweak the color coding scheme to construct paired variants H_{ij} as follows: red denotes candidate i , green denotes candidate j and blue denotes the rest. Alternatively, H_{ij} can be thought of an RGB image: the red channel shows the "subject", i.e. candidate i ; the green channel shows the "object", i.e. candidate j ; and the blue channel

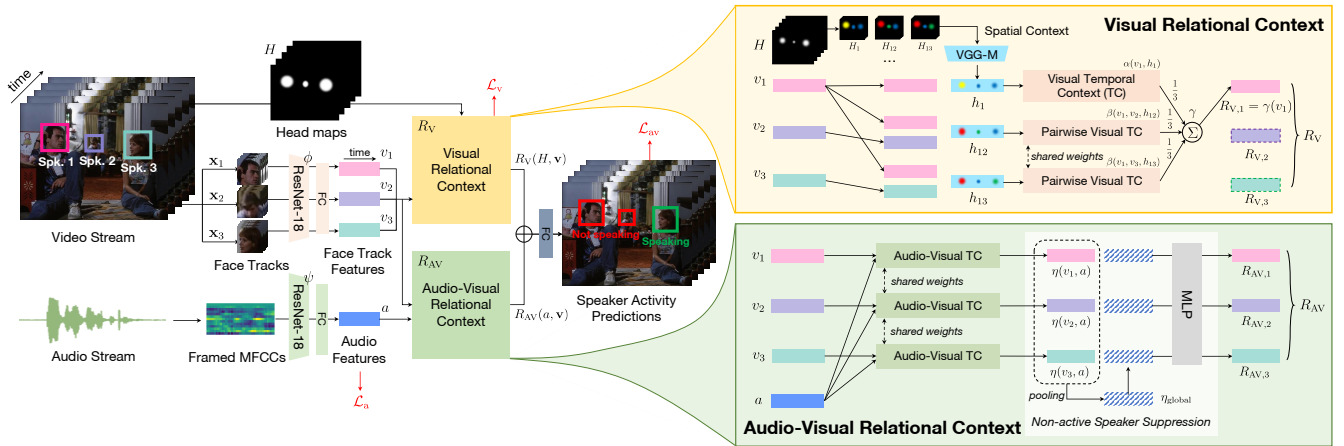


Figure 2: Model overview. Given a video, we first extract face track features and audio features. Face scale and position information are encoded as 2D Gaussians and embedded with CNN layers, which we refer to as *spatial context*. Next, we construct contextual audio-visual representations for each candidate, through intermediate *visual relational context* and *audio-visual relational context* modules. For the visual relational context, we introduce a permutation-equivariant layer to refine each speaker’s visual representation by incorporating pairwise relationships and long-term *temporal context*. For the audio-visual relational context, we model audio-visual affinity over time for each candidate and then suppress non-active speakers by contrasting their affinity features with others. The final refined visual and audio-visual representations are concatenated and passed through a shared prediction layer to estimate a confidence score for each visible candidate.

shows other "context" candidates, i.e. those other than i and j (H_i is identified with H_{ii}). Then a VGG-M-inspired network with four 2D convolutional layers [33] is used to embed these colored head maps H_i and H_{ij} into 64-dimensional vectors per frame, h_i and h_{ij} for each candidate i and candidate pair (i, j) with $i \neq j$. We term the resulting embeddings h_i and h_{ij} candidate i 's *spatial context*.

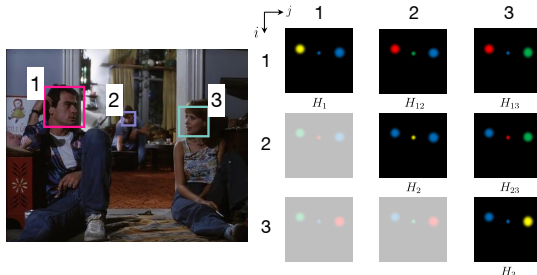


Figure 3: Head map construction. For each candidate i , we construct his/her speaker-centric head map H_i and pairwise head maps $\{H_{ij} \mid i \neq j\}$ for each candidate pair (i, j) . Note that the lower-triangular elements with $i > j$ are not used due to the skew-symmetric implementation (see Sec. 3.3).

3.3 Relational Context

After obtaining face track and audio descriptors as well as the spatial context embeddings, we jointly model all candidates in the scene and refine each candidate’s representations for robust ASD. Our motivation for modeling relational context is the inevitable presence of local ambiguities which make directly matching face motion (e.g. lip movements) and audio non-trivial. Therefore, we integrate different types of contextual clues and consider all candidates in the scene holistically to facilitate the ASD task. For example, a person who is speaking tends to receive more attention from others and is usually portrayed with higher visual saliency (especially in movies).

Specifically, our relational context component is designed to complete two natural sub-tasks for ASD: learning a contextual visual representation for *visual voice activity detection*, and a contextual audio-visual representation for *audio-visual affinity modeling*. The two resulting representations will then be fused for joint analysis to produce the final prediction. As shown in Fig. 2, our relational context contains two parts: **visual** (R_V) and **audio-visual relational context** (R_{AV}). This special decoupled design strengthens the model’s robustness to synchronization errors between audio and video. We now elaborate on the design of R_V and R_{AV} .

Visual Relational Context: Our key idea is to represent each candidate’s visual activity by aggregating his/her locally perceived activity and all pairwise interactions with other candidates in the scene. Specifically, we design a *permutation-equivariant* layer γ to process all candidates simultaneously, while preserving the order in which they are provided to the network. Denote the input visual feature stack by $\mathbf{v} = (v_1, v_2, \dots, v_N) \in \mathbb{R}^{N \times T \times d}$. For each candidate $i = 1, 2, \dots, N$, the output in the i -th position of $R_V(\mathbf{v})$ is

$$R_{V,i} = \gamma(v_i) = \frac{1}{N} \left[\alpha(v_i, h_i) + \sum_{j \neq i} \beta(v_i, v_j, h_{ij}) \right], \quad (1)$$

where $\alpha(\cdot, \cdot)$ and $\beta(\cdot, \cdot, \cdot)$ are two networks that respectively model visual activity and pairwise visual interactions, v_i denotes candidate i 's visual features, h_i is the head map embedding for candidate i , and h_{ij} is the head map embedding for the candidate pair (i, j) .

Here, the parameters of α are shared across all candidates, and the parameters of β are shared across all candidate pairs, so Eq. (1) is directly applicable to any number of co-occurring candidates N . This in turn means that our network can analyze a theoretically unlimited number of candidates at test time. On the surface, Eq. (1) appears to suggest that obtaining $R_V(\mathbf{v})$ requires $N^2 + N$ evaluations in total. In our implementation, we reduce this number by half to

$N(N+1)/2$ by parameterizing $\beta(\cdot, \cdot, \cdot)$ as a *skew-symmetric* function:

$$\beta_{ji} = \beta(v_j, v_i, h_{ji}) = -\beta(v_i, v_j, h_{ij}) = -\beta_{ij}. \quad (2)$$

Our intuition is that β_{ij} should learn to represent the amount of attention that candidate i directs at candidate j , and rank the relative “activeness” and visual saliency of candidate i against candidate j . In this case, it suffices to express the direction of such relationships using a single positive or negative sign for each component.

Audio-Visual Relational Context: We model the affinity between the audio and each speaker’s face track in terms of both synchrony and cross-modal biometrics. We fuse per-frame face track features with the audio features via concatenation, and then pass them through a shared network η to compute local A-V affinity features $\eta(v_i, a)$. Considering that local A-V affinity estimation is vulnerable to signal ambiguity (e.g. low-resolution faces, profile faces, noisy audio etc.), we aggregate evidence from all candidates to make more reliable predictions: if one or some candidates display strong A-V agreement, then the network can accordingly lower its affinity predictions for other candidates to mitigate false positives. We therefore design a learnable module to adaptively suppress only the non-active candidates whose A-V affinities are of significantly lower magnitudes, while leaving the active candidates’ features “as-is”. To this end, we introduce an element-wise max-pooling operation over all A-V affinity features to obtain a global representation. Specifically, the pooled feature is computed as $\eta_{\text{global}} = \max_{1 \leq i \leq N} \eta(v_i, a)$, and then concatenated to each candidate’s initial features $\eta(v_i, a)$. The concatenated features are further processed with two fully connected layers to generate each candidate’s final contextual audio-visual representation $R_{AV,i}$. We name this operation *non-active speaker suppression* after its effect.

3.4 Temporal Context

As shown in Fig. 2, we incorporate temporal context in the networks α , β , and η inside R_V and R_{AV} . This benefits ASD in two ways: first, it improves the consistency of the relational context modeling process and smoothes out local, instantaneous noises; second, it helps to alleviate synchronization errors between the audio and the video stream, an issue that is ubiquitous with in-the-wild videos.

We choose two basic architectures as our temporal modeling back-end: temporal convolutions (1D CNNs), and bi-directional Gated Recurrent Units (Bi-GRUs). The former is favorable since it has a fixed temporal receptive field, and can be easily adapted for online settings. The latter is more reliable since it has access to both past and future information. Hence, for the rest of the paper, by mentioning **temporal context** we refer to the usage of a one-layer Bi-GRU backend with 256 cells. Otherwise, we apply a stack of two 1D convolutional layers of kernel size 3 interleaved with Batch Normalization [28] and ReLU activation by default.

Finally, the refined contextual representations R_V and R_{AV} are concatenated and fed to a fully-connected layer that is shared across candidates. The outputs are passed through sigmoid activation, resulting in real values between 0 and 1 that indicate each candidate’s probability of being the active speaker.

3.5 Losses

The model is trained end-to-end with a multi-task loss formulation. Since the goal is to predict a binary speaking/not speaking label, for each loss term we apply the standard binary cross-entropy (BCE)

loss, averaged over all time steps. The BCE loss is defined as

$$\mathcal{L}_{\text{BCE}}(y, \hat{y}) = -\hat{y} \log y - (1 - \hat{y}) \log(1 - y), \quad (3)$$

where y and \hat{y} are the predictions and ground truth labels, respectively. For a scene with T frames and N candidates, let A_{aux} and V_{aux} denote the audio and visual auxiliary prediction layer, AV_{pred} the final audio-visual prediction layer, and $\hat{y}_{v,i}^{1,2,\dots,T}$ and $\hat{y}_{av,i}^{1,2,\dots,T}$ the corresponding visual and audio-visual ground truths of the i -th candidate. We add two auxiliary losses \mathcal{L}_a and \mathcal{L}_v to learn discriminative features for both visual and audio streams. The *audio* prediction loss is

$$\mathcal{L}_a = \frac{1}{T} \sum_{t=1}^T \mathcal{L}_{\text{BCE}}(\sigma(A_{\text{aux}}(a^t)), \hat{y}_a^t), \quad (4)$$

where $\hat{y}_a^t = \max_{1 \leq i \leq N} \hat{y}_{av,i}^t$ is the audio ground truth indicating whether at least one of the candidates is speaking, and $\sigma(\cdot)$ is the sigmoid function.

When training on a single candidate (*i.e.* without employing relational context), the *visual* prediction loss and main *audio-visual* prediction loss are given by

$$\begin{aligned} \mathcal{L}_v &= \frac{1}{T} \sum_{t=1}^T \mathcal{L}_{\text{BCE}}(\sigma(V_{\text{aux}}(v_i^t)), \hat{y}_{v,i}^t), \\ \mathcal{L}_{av} &= \frac{1}{T} \sum_{t=1}^T \mathcal{L}_{\text{BCE}}(\sigma(AV_{\text{pred}}(a^t \oplus v_i^t)), \hat{y}_{av,i}^t), \end{aligned} \quad (5)$$

where \oplus denotes feature concatenation. When training with relational context on multiple candidates, the losses are slightly different. To compute the visual and audio-visual losses, we replace the inputs to the auxiliary prediction layers with (concatenated) contextual representations, and aggregate losses from all candidates:

$$\begin{aligned} \mathcal{L}_v &= \frac{1}{N} \sum_{i=1}^N \frac{1}{T} \sum_{t=1}^T \mathcal{L}_{\text{BCE}}(\sigma(V_{\text{aux}}(R_{V,i}^t)), \hat{y}_{v,i}^t), \\ \mathcal{L}_{av} &= \frac{1}{N} \sum_{i=1}^N \frac{1}{T} \sum_{t=1}^T \mathcal{L}_{\text{BCE}}(\sigma(AV_{\text{pred}}(R_{V,i}^t \oplus R_{AV,i}^t)), \hat{y}_{av,i}^t). \end{aligned} \quad (6)$$

Finally, the *total* loss is defined as follows:

$$\mathcal{L} = \mathcal{L}_a + \mathcal{L}_v + \mathcal{L}_{av}. \quad (7)$$

Note that in practice, N may vary over time as candidates enter and leave the scene.

4 EXPERIMENTS

To validate the effectiveness of our model, we thoroughly evaluate and analyze the performance of UniCon from two aspects. First, we perform comprehensive ablation studies and comparison on a popular and challenging ASD benchmark derived from diverse, in-the-wild movies, AVA-ActiveSpeaker. Second, we perform cross-dataset testing on three additional datasets, Columbia, RealVAD, and AVDIAR. Example frames from each dataset are shown in Fig. 4.

4.1 Experimental Setup

Datasets: The **AVA-ActiveSpeaker** dataset [35] consists of 262 YouTube movies from film industries around the world. Each video is annotated from minutes 15 to 30. The dataset provides face bounding boxes that are linked into face tracks and labeled for both speech

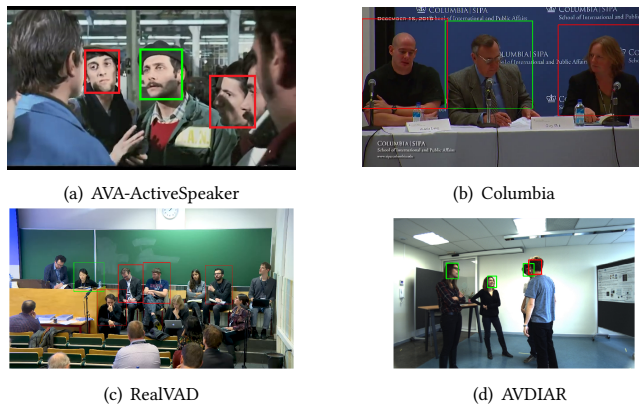


Figure 4: Example frames from the datasets used in this paper. Green boxes denote the active speakers.

activity and whether the speech is audible. As shown in Fig. 4a, the dataset is highly challenging as it contains occlusions, low-resolution faces, low-quality audio, and varied lighting conditions. It has become a mainstream benchmark for the ASD task.

The **Columbia** dataset [7] annotates 35 minutes of a one-hour panel discussion video featuring 6 unique panelists (Bell, Bollinger, Lieberman, Long, Sick, Abbas). The panelists directly face the camera, but occasionally look elsewhere, or perform spontaneous activities (e.g. drinking water). Upper body bounding boxes and voice activity ground truth labels are provided, as shown in Fig. 4b.

The **RealVAD** dataset [4] (Fig. 4c) is also constructed from a panel discussion lasting approximately 83 minutes. The video is recorded using a static, mounted camera, capturing the 9 panelists in a full shot. The panelists sit in two rows, and at any given time can be looking anywhere. In addition, the panelists perform natural, spontaneous actions that may result in partial occlusion of the face and mouth (e.g. touching faces, cupping chins). Similar to Columbia, upper body bounding boxes and voice activity labels are provided. These two datasets are often adopted by works on visual voice activity detection (V-VAD) and ASD [1, 11, 37, 38].

The **AVDIAR** dataset [22] (Fig. 4d) provides 27 recordings of informal conversations involving one to four participants. The dataset is extremely challenging due to varying levels of speaker movement and speech overlap, occlusions, profile faces, and even back-facing speakers. The dataset was originally used to evaluate speaker diarization and tracking systems. We re-purpose the dataset to evaluate our ASD model by aggregating per-frame ASD results into diarization results.

Data Preprocessing: We preprocess each dataset by cropping the face tracks of each visible candidate. For the **AVA-ActiveSpeaker** dataset, to group face tracks into “scenes” for training and testing, each video is first segmented into shots with ffmpeg [31]. Within each shot, disconnected tracks belonging to the same candidate are then merged based on bounding box overlap with an IoU threshold of 0.8. All videos are resampled to 25fps, and ground truth labels are computed according to frame timestamps via nearest-neighbor interpolation. The **Columbia** and **RealVAD** datasets do not provide face bounding boxes, so face detection is performed with an off-the-shelf RetinaFace detector [15]. The resulting bounding boxes are expanded by 1.3x to mimic AVA-style detections. For the **AVDIAR** dataset, the provided bounding boxes are used directly.

Train/test Split and Evaluation Metric: For **AVA-ActiveSpeaker**, the official train/validation/test split is adopted. The test set is held out for the ActivityNet challenge and unavailable, so we perform our analysis on the validation set instead, as several previous methods do [2, 8, 42]. The **Columbia** and **AVDIAR** datasets do not provide training sets, so we only perform zero-shot testing. For Columbia, following previous practice [7, 11], we report results on all but one speaker (Abbas) which is usually held out for validation. For **RealVAD**, in addition to zero-shot results, we also report leave-one-out cross-validation results for each of the 9 panelists to compare with previous work. We only train on scenes in which the test speaker does not co-occur with the training speakers.

Finally, we report the metric that was commonly adopted by previous work for each dataset: mean Average Precision (mAP) and Area Under Receiver Operating Characteristic Curve (AUROC) for AVA-ActiveSpeaker, frame-wise F_1 score for Columbia and RealVAD, and Diarization Error Rate (DER) for AVDIAR, which is defined as the sum of false alarm rate, missed speech rate, and speaker confusion rate.

Implementation Details: We implement our model with PyTorch and the pytorch-lightning package. All models are trained using the AdamW optimizer [32]. The network parameters are initialized using He initialization [24]. During training, we augment the data via random horizontal flipping and uniform corner cropping along the input face tracks, followed by random adjustments to brightness, contrast, and saturation. All cropped face tracks are resized to 144×144 , and randomly cropped to 128×128 for training. We use a central 128×128 patch for testing. We only run one trial with a fixed random seed for all experiments to ensure the results are comparable, but find the results to be stable under different network initializations and data augmentations. During inference, occasional long segments that do not entirely fit into GPU memory are split into shorter, fixed-size chunks and the predictions are re-combined later.

To facilitate training, we apply curriculum learning [3] by first training a single-candidate model without relational context according to Eq. (7), and then continuing to train on up to 3 candidates, i.e. $N = 3$ in Eq. (1), but without the \mathcal{L}_a term in Eq. (7). This is theoretically and empirically enough to learn the relevant parameters since R_V captures *pairwise* interactions, and we observe diminishing returns by sampling 4 or more candidates (note that we can still test on $N > 3$ candidates due to parameter sharing). In addition, we follow the sampling strategy in [2] during training, and randomly sample a 1.12s segment (28 frames at 25fps)¹ from every training example in the dataset for each epoch. Therefore, our epoch size correlates with the number of identities or scenes rather than face detections, which prevents over-fitting.

4.2 Ablation Studies

In this subsection, we provide a thorough ablation study on the modeling components on the challenging AVA-ActiveSpeaker dataset.

Spatial Context: We compare two different ways to incorporate spatial context. The first approach applies *early fusion* and concatenates speaker-centric spatial embeddings h_i with the i -th candidate’s face track features before dimensionality reduction. In contrast, the second approach applies *late fusion* and concatenates

¹The mean speech segment duration is 1.11s on the AVA-ActiveSpeaker dataset [35].

h_i with the dimension-reduced face track features v_i . As shown in Table 1, both improve upon the simplest baseline that only uses convolutional temporal context (see Sec. 3.4) on audio and face track features, but early fusion performs 0.8% worse than late fusion, likely because information is lost too early as features pass the 128-dim bottleneck. Thus, we adopt late fusion hereafter. As shown in Table 2, simply concatenating the spatial context embeddings to the face track features (denoted by +S) already yields stronger visual representations that improve the baseline (+1.7% mAP).

Table 1: Comparison w.r.t. the stage at which spatial context is introduced. We observe that applying late fusion yields better results since it helps to retain more spatial information for subsequent modeling.

Method	mAP (%)	AUROC
Baseline	84.0	0.93
Early fusion	85.0	0.93
Late fusion	85.7	0.94

Relational Context: We assess the efficacy of our relational context component. As shown in Table 2, incorporating the relational context module yields noticeable improvements of 1.7% absolute over the 1D CNN baseline when used alone (+R), and works in synergy with spatial context (+2.6% mAP for +S+R). We observe that this improvement is especially pronounced in multi-speaker and small face scenarios. This is because UniCon can model the relationships among the candidates in the scene, which leads to a better holistic understanding under these challenging settings. Moreover, we introduce weight sharing strategies within a permutation-equivariant formulation, which enables our model to process a theoretically unlimited number of speakers at test time.

Temporal Context: Finally, incorporating temporal context with Bi-GRUs yields the most substantial improvement, which is reasonable because Bi-GRUs have access to each candidate’s complete track history, and can thus make more reliable predictions. By leveraging long-term temporal context, our model ultimately achieves 92.0% mAP, improving upon our initial baseline by as much as 8.0%.

Non-active Speaker Suppression: In this part, we discuss the role of non-active speaker suppression in ASD. As shown in Table 3, we compare two ways of integrating global information against a baseline that does not apply non-active speaker suppression. One uses mean-pooling to obtain η_{global} , and the other uses max-pooling. Both improve over the model that only introduces visual relational context, and max-pooling performs slightly better, yielding fewer false negatives. One possibility is that max-pooling back-propagates more gradients to the most salient speaker, while mean-pooling results in weaker gradients for the true active speaker, making the model less confident.

Table 2: Ablation results on the AVA-ActiveSpeaker validation set. S: Spatial Context; R: Relational Context; T: Temporal Context. Predictions from models without temporal context are smoothed over 11-frame windows via Wiener filtering.

Method	#Params	mAP (%)	AUROC
Baseline	23.2M	84.0	0.93
+S	23.4M	85.7	0.94
+R	23.8M	85.7	0.94
+S+R	24.0M	86.6	0.94
+T (Bi-GRU)	23.0M	89.6	0.95
+S+T	23.2M	90.3	0.96
+R+T	23.5M	90.4	0.96
+S+R+T (UniCon)	23.8M	92.0	0.97

Table 3: Effect of implementing non-active speaker suppression with different pooling methods.

Pooling	mAP (%)	AUROC
None (no suppression)	90.8	0.96
Mean-pooling	91.6	0.96
Max-pooling	92.0	0.97

As shown in Table 4, our full UniCon outperforms previous audio-visual ASD methods including the state-of-the-art by a large margin. Remarkably, we achieve 92.0% mAP without any pre-training, surpassing 90% for the first time on the AVA-ActiveSpeaker validation set at the time of submission. This overall performance strongly supports the effectiveness and superiority of our UniCon. It is also worth noting that the +S+T model in Table 2 provides a very strong baseline (90.3% mAP) that already exceeds all methods available for comparison here, which further confirms the benefits of our model. As a side note, applying ImageNet pre-training to the encoders further boosts performance by 0.2% to 92.2%.

4.3 Comparison with the State-of-the-Art

Table 4: Comparison with previous work on the AVA-ActiveSpeaker validation set. For each method, we copy the results from its original paper. mAP is calculated using the official evaluation tool, after interpolating our predictions to the timestamps in the original annotations.

Method	Pre-training?	mAP (%)	AUROC
AV-GRU [35]	×	-	0.92
Optical Flow [27]	×	-	0.93
Multi-Task [42]	✓	84.0	-
VGG-LSTM [8]	✓	85.1	-
VGG-TempConv [8]	✓	85.5	-
ASC [2]	✓	87.1	-
Ours (UniCon)	×	92.0	0.97
Ours (UniCon)	✓	92.2	0.97

4.4 Cross-dataset Evaluation

Apart from evaluating on the large-scale AVA-ActiveSpeaker dataset, we also conduct cross-dataset evaluation on three other datasets: Columbia, RealVAD and AVDIAR. These datasets contain challenges that are previously rare, if not unseen in AVA-ActiveSpeaker, such as overlapped speech, reverberation, extreme poses, heavy face occlusion, and different spatial distributions of candidates. Our cross-dataset evaluation protocols include zero-shot testing on all three datasets using the model trained on AVA-ActiveSpeaker, and leave-one-out testing after fine-tuning on RealVAD. We report both visual-only and audio-visual performance. The former is obtained using only the auxiliary visual prediction layer V_{aux} in Sec. 3.5.

Zero-shot Testing: Overall, our UniCon model achieves good zero-shot generalization performance on all three datasets. In particular, it achieves state-of-the-art performance on both Columbia (90.9% average $F-1$ score) and RealVAD (80.30% average $F-1$ score). On AVDIAR, we are the first to report diarization performance using single-channel audio and a single camera viewpoint.

On the **Columbia** dataset, we outperform all previous state-of-the-art under both visual-only and audio-visual settings. It is worth noting that the listed visual-only methods [7, 37, 38] use the entire upper body, while we only use cropped face tracks. We also outperform previous audio-visual models [1, 11] that are pre-trained on datasets a magnitude larger (224 and 606 hours, respectively). Similarly, on **RealVAD**, our zero-shot visual-only and audio-visual results outperform the only existing method available for equal comparison by a large margin (87.22% V, 80.30% AV vs. 53.04%). An interesting observation is that on RealVAD and Columbia, our visual-only model outperforms the audio-visual model in most cases. This is because labels provided with the two datasets are not frame-accurate; our audio-visual model successfully detects short inter-sentence pauses that are not annotated as "not speaking" by the dataset curators. Due to the page limit, please refer to the supplementary video for details.

On **AVDIAR**, our model provides the first baseline that only uses mono microphone input, and a single camera viewpoint – a highly challenging setup. Nevertheless, our audio-visual model still achieves 55.16% DER on average, which is a very promising zero-shot result. Moreover, our full UniCon model reduces DER by 4 to 5% over the +S+T model that does not model relationships among the candidates, which once again supports the efficacy of our relational context in multi-speaker scenarios.

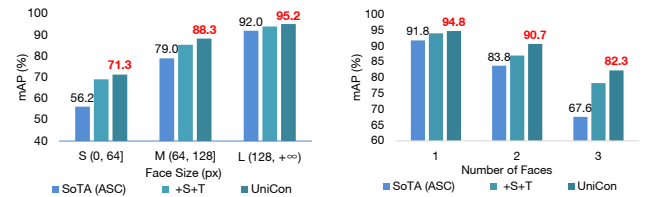
Fine-tuned Results: We also evaluate fine-tuned leave-one-out performance on RealVAD (Columbia and AVDIAR do not provide training splits). After fine-tuning on homogeneous data, our model quickly adapts to the unseen scenario (panel discussion) and noises (e.g. spontaneous actions near the face). The average audio-visual F_1 score is improved by 5.34% from 80.30% to 85.64%, and the average visual F_1 score is improved by 1.47% from 87.22% to 88.69%. More detailed results can be found in the supplemental material.

4.5 Performance Breakdown

In this part, we provide a more in-depth analysis of three known challenging scenarios on AVA-ActiveSpeaker. We provide a comparison with the existing state-of-the-art and highlight some advantages and appealing properties of the UniCon model.

Low-resolution Faces. We summarize performance for different face sizes in Fig. 5a. Following previous evaluation procedures [2, 35], we partition the dataset into three bins by the widths of the detected faces: small (width ≤ 64 pixels), medium (width between 64 and 128 pixels), and large (width > 128 pixels). While the +S+T model provides a strong baseline with 90.3% mAP that already outperforms the previous state-of-the-art [2], our full model which exploits relational context yields further performance gains of around 2% mAP absolute on all subsets. In particular, on the "small" subset, we outperform the previous state-of-the-art (56.2% mAP) by a large margin of 15.1%.

Multiple Candidates. Fig. 5b shows the model performance according to the number of detected faces in a frame. We consistently improve over the previous state-of-the-art for different numbers of detected faces, with the most significant gain being in the subset with 3 faces (+14.7% mAP). Moreover, we surpass 90% on the two-faces subset for the first time. The improvements prove the effectiveness of spatial and relational context modeling, which equip the model with the ability to discern background actors from main



(a) mAP breakdown by face resolution. (b) mAP breakdown by number of on-screen faces.

Figure 5: Performance breakdown. We evaluate the performance of our model on faces of different sizes and on frames with one, two, and three detected faces (see main text). The performance of the previous state-of-the-art [2] (ASC) and our best model are annotated for better comparison.

actors, as well as an enhanced understanding of the relationships between the visible speakers. We refer interested readers to the supplemental video for more interesting qualitative results.

Out-of-sync Audio and Video. Many real-world videos suffer from poor synchronization between audio and video, caused by transmission or re-encoding. To evaluate our model's performance on out-of-sync data, we assume perfect synchronization in the source videos and artificially shift the audio stream by up to 10 frames to mimic out-of-sync videos. We then assess the performance of our model on manually de-synchronized videos. As shown in Fig. 6, our model is fairly resilient to A-V sync errors. Remarkably, even on videos that are shifted by 10 frames (0.4sec), our full UniCon model only loses 0.67% mAP, and still outperforms a simple Bi-GRU baseline by 1.64% absolute, which does not model spatial and relational context (+T in Table 2), and receives *synchronized* inputs. This shows that our model is robust to synchronization error.

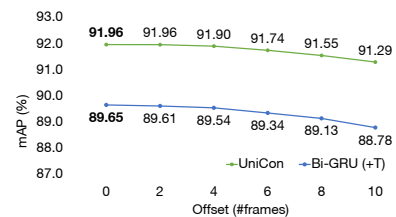


Figure 6: Results on out-of-sync videos. Our model is fairly resilient to A-V synchronization error.

5 CONCLUSION

We have proposed a novel model named UniCon for active speaker detection. Key to our method is a unified modeling framework which efficiently aggregates different types of contextual evidence to make robust decisions, by considering the relationships between each candidate and others. We demonstrate via experiments on the large-scale AVA-ActiveSpeaker dataset and various other real-world datasets that our model successfully tackles challenging cases with multiple candidates and low-resolution faces, and outperforms state-of-the-art methods by a large margin.

ACKNOWLEDGMENTS

This work was partially supported by the National Key R&D Program of China (No. 2017YFA0700804) and the National Natural Science Foundation of China (No. 61876171).

REFERENCES

- [1] Triantafyllos Afouras, Andrew Owens, Joon Son Chung, and Andrew Zisserman. 2020. Self-supervised Learning of Audio-Visual Objects from Video. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XVIII (Lecture Notes in Computer Science, Vol. 12363)*, Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer, 208–224.
- [2] Juan León Alcázar, Fabian Caba, Long Mai, Federico Perazzi, Joon-Young Lee, Pablo Arbeláez, and Bernard Ghanem. 2020. Active Speakers in Context. In *CVPR. IEEE*, 12462–12471.
- [3] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *ICML (ACM International Conference Proceeding Series, Vol. 382)*. ACM, 41–48.
- [4] C. Beyan, M. Shahid, and V. Murino. 2020. RealVAD: A Real-world Dataset and A Method for Voice Activity Detection by Body Motion Analysis. *IEEE Transactions on Multimedia* (2020), 1–1. <https://doi.org/10.1109/TMM.2020.3007350>
- [5] Hervé Bredin and Gregory Gelly. 2016. Improving Speaker Diarization of TV Series using Talking-Face Detection and Clustering. In *ACM Multimedia*. ACM, 157–161.
- [6] J. Cech, R. Mittal, A. Deleforge, J. Sanchez-Riera, X. Alameda-Pineda, and R. Horaud. 2013. Active-speaker detection and localization with microphones and cameras embedded into a robotic head. In *2013 13th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*. 203–210. <https://doi.org/10.1109/HUMANOIDS.2013.7029977>
- [7] Punarjay Chakravarty and Tinne Tuytelaars. 2016. Cross-Modal Supervision for Learning Active Speaker Detection in Video. In *ECCV (5) (Lecture Notes in Computer Science, Vol. 9909)*. Springer, 285–301.
- [8] Joon Son Chung. 2019. Naver at ActivityNet Challenge 2019 - Task B Active Speaker Detection (AVA). *CoRR* abs/1906.10555 (2019).
- [9] Joon Son Chung, Jaesung Huh, Arsha Nagrani, Triantafyllos Afouras, and Andrew Zisserman. 2020. Spot the Conversation: Speaker Diarisation in the Wild. In *INTERSPEECH. ISCA*, 299–303.
- [10] Joon Son Chung, Arsha Nagrani, and Andrew Zisserman. 2018. VoxCeleb2: Deep Speaker Recognition. In *INTERSPEECH. ISCA*, 1086–1090.
- [11] Joon Son Chung and Andrew Zisserman. 2016. Out of Time: Automated Lip Sync in the Wild. In *ACCV Workshops (2) (Lecture Notes in Computer Science, Vol. 10117)*. Springer, 251–263.
- [12] Soo-Whan Chung, Joon Son Chung, and Hong-Goo Kang. 2019. Perfect Match: Improved Cross-modal Embeddings for Audio-visual Synchronisation. In *ICASSP. IEEE*, 3965–3969.
- [13] Ross Cutler and Larry S. Davis. 2000. Look Who's Talking: Speaker Detection using Video and Audio Correlation. In *IEEE International Conference on Multimedia and Expo (III)*. IEEE Computer Society, 1589–1592.
- [14] Ross Cutler, Ramin Mehran, Sam Johnson, Cha Zhang, Adam Kirk, Oliver Whyte, and Adarsh Kowdle. 2020. Multimodal Active Speaker Detection and Virtual Cinematography for Video Conferencing. In *ICASSP. IEEE*, 4527–4531.
- [15] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. 2020. RetinaFace: Single-Shot Multi-Level Face Localisation in the Wild. In *CVPR. IEEE*, 5202–5211.
- [16] Jorge Francisco Madrigal Diaz, Frédéric Lerasle, Lionel Pibre, and Isabelle Ferrané. 2021. Audio-Video detection of the active speaker in meetings. In *IEEE International Conference on Pattern Recognition (ICPR 2020)*.
- [17] Yifan Ding, Yong Xu, Shi-Xiong Zhang, Yahuan Cong, and Liqiang Wang. 2020. Self-Supervised Learning for Audio-Visual Speaker Diarization. In *ICASSP. IEEE*, 4367–4371.
- [18] Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinandan Hasidim, William T. Freeman, and Michael Rubinstein. 2018. Looking to listen at the cocktail party: a speaker-independent audio-visual model for speech separation. *ACM Trans. Graph.* 37, 4 (2018), 112:1–112:11.
- [19] Y. Fan, J. W. Kang, L. T. Li, K. C. Li, H. L. Chen, S. T. Cheng, P. Y. Zhang, Z. Y. Zhou, Y. Q. Cai, and D. Wang. 2020. CN-Celeb: A Challenging Chinese Speaker Recognition Dataset. In *ICASSP. IEEE*, 7604–7608.
- [20] Giulia Garau, Sileye O. Ba, Hervé Bourlard, and Jean-Marc Odobez. 2009. Investigating the use of visual focus of attention for audio-visual speaker diarisation. In *ACM Multimedia*. ACM, 681–684.
- [21] Giulia Garau, Alfred Dielmann, and Hervé Bourlard. 2010. Audio-visual synchronisation for speaker diarisation. In *INTERSPEECH. ISCA*, 2654–2657.
- [22] Israel D. Gebru, Sileye O. Ba, Xiaofei Li, and Radu Horaud. 2018. Audio-Visual Speaker Diarization Based on Spatiotemporal Bayesian Fusion. *IEEE Trans. Pattern Anal. Mach. Intell.* 40, 5 (2018), 1086–1099.
- [23] Rohit Girdhar, João Carreira, Carl Doersch, and Andrew Zisserman. 2019. Video Action Transformer Network. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*. Computer Vision Foundation / IEEE, 244–253.
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. In *ICCV*. IEEE Computer Society, 1026–1034.
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *CVPR. IEEE Computer Society*, 770–778.
- [26] Ken Hoover, Sourish Chaudhuri, Caroline Pantofaru, Ian Sturdy, and Malcolm Slaney. 2018. Using audio-visual information to understand speaker activity: Tracking active speakers on and off screen. In *ICASSP. IEEE*, 6558–6562.
- [27] Chong Huang and Kazuhito Koishida. 2020. Improved Active Speaker Detection based on Optical Flow. In *CVPR Workshops. IEEE*, 4084–4090.
- [28] Sergey Ioffe and Christian Szegedy. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *ICML (JMLR Workshop and Conference Proceedings, Vol. 37)*. JMLR.org, 448–456.
- [29] Kristiina Jokinen, Kazuaki Harada, Masafumi Nishida, and Seiichi Yamamoto. 2010. Turn-alignment using eye-gaze and speech in conversational interaction. In *INTERSPEECH. ISCA*, 2018–2021.
- [30] Adam Kendon. 1967. Some functions of gaze-direction in social interaction. *Acta psychologica* 26 (1967), 22–63.
- [31] Rainer Lienhart. 2001. Reliable Transition Detection in Videos: A Survey and Practitioner's Guide. *Int. J. Image Graph.* 1, 3 (2001), 469–486.
- [32] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- [33] Manuel J. Marin-Jiménez, Vicky Kalogeiton, Pablo Medina-Suarez, and Andrew Zisserman. 2019. LAEO-Net: Revisiting People Looking at Each Other in Videos. In *CVPR. Computer Vision Foundation / IEEE*, 3477–3485.
- [34] Junting Pan, Siyu Chen, Zheng Shou, Jing Shao, and Hongsheng Li. 2020. Actor-Context-Actor Relation Network for Spatio-Temporal Action Localization. *CoRR* abs/2006.07976 (2020).
- [35] Joseph Roth, Sourish Chaudhuri, Ondrej Klejch, Radhika Marvin, Andrew C. Gallagher, Liat Kaver, Sharadh Ramaswamy, Arkadiusz Stopczynski, Cordelia Schmid, Zhonghua Xi, and Caroline Pantofaru. 2020. AVA-ActiveSpeaker: An Audio-Visual Dataset for Active Speaker Detection. In *ICASSP. IEEE*, 4492–4496.
- [36] Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. FaceNet: A unified embedding for face recognition and clustering. In *CVPR. IEEE Computer Society*, 815–823.
- [37] Muhammad Shahid, Cigdem Beyan, and Vittorio Murino. 2019. Voice Activity Detection by Upper Body Motion Analysis and Unsupervised Domain Adaptation. In *2019 IEEE/CVF International Conference on Computer Vision Workshops, ICCV Workshops 2019, Seoul, Korea (South), October 27-28, 2019*. IEEE, 1260–1269.
- [38] Muhammad Shahid, Cigdem Beyan, and Vittorio Murino. 2021. S-VAD: Visual Voice Activity Detection by Motion Segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. 2332–2341.
- [39] Malcolm Slaney and Michele Covell. 2000. FaceSync: A Linear Operator for Measuring Synchronization of Video Facial Images and Audio Tracks. In *NIPS*. MIT Press, 814–820.
- [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *NIPS*. 5998–6008.
- [41] Jianchao Wu, Zhanghui Kuang, Limin Wang, Wayne Zhang, and Gangshan Wu. 2020. Context-Aware RCNN: A Baseline for Action Detection in Videos. In *ECCV (25) (Lecture Notes in Computer Science, Vol. 12370)*. Springer, 440–456.
- [42] Yuan-Hang Zhang, Jingyun Xiao, Shuang Yang, and Shiguang Shan. 2019. Multi-Task Learning for Audio-Visual Active Speaker Detection. *The ActivityNet Large-Scale Activity Recognition Challenge 2019* (2019), 1–4.