



SVVAD: Personal Voice Activity Detection for Speaker Verification

Zuheng Kang, Jianzong Wang*, Junqing Peng, Jing Xiao

Ping An Technology (Shenzhen) Co., Ltd., Shenzhen, China

{kangzuheng896, wangjianzong347, pengjq, xiaojing661}@pingan.com.cn

Abstract

Voice activity detection (VAD) improves the performance of speaker verification (SV) by preserving speech segments and attenuating the effects of non-speech. However, this scheme is not ideal: (1) it fails in noisy environments or multi-speaker conversations; (2) it is trained based on inaccurate non-SV sensitive labels. To address this, we propose a speaker verification-based voice activity detection (SVVAD) framework that can adapt the speech features according to which are most informative for SV. To achieve this, we introduce a label-free training method with triplet-like losses that completely avoids the performance degradation of SV due to incorrect labeling. Extensive experiments show that SVVAD significantly outperforms the baseline in terms of equal error rate (EER) under conditions where other speakers are mixed at different ratios. Moreover, the decision boundaries reveal the importance of the different parts of speech, which are largely consistent with human judgments.

Index Terms: voice activity detection, personal VAD, speaker verification

1. Introduction

Voice activity detection (VAD) is a task that identifies whether human speech is present or absent and is often used upstream of other speech components such as automatic speech recognition (ASR), speaker verification (SV), and speaker diarization (SD). It aims to reduce the impact of non-speech on downstream speech tasks and indirectly improve their performance. However, their goals are different. ASR and SD models need to efficiently and accurately determine the boundary between speech and non-speech to avoid missing content. SV is much more complicated because there are more factors involved.

A typical VAD framework is considered to be a gating module that makes a speech/non-speech decision for each frame. Early studies focused on signal and statistical analysis and feature engineering [1, 2, 3, 4]. More recently, the use of conventional deep learning methods, such as convolutional neural networks (CNN) and recurrent neural networks (RNN), has shown significant improvements in detection performance at low signal-to-noise ratios (SNR) [5, 6, 7]. Later, with the introduction of the attention mechanism, the model can automatically compare the characteristics of speech and noise within audio to derive more accurate judgments [8, 9]. Some authors use audiovisual information for VAD detection [10, 11]. Although it improves the performance of SV by VAD to some extent, these frameworks are insufficient because non-target speakers are also identified and retained as speech labels. Since the SV model uses only a single speaker for supervised training and

does not consider multiple speakers, this leads to a significant drop in verification performance. Therefore, a target speaker-only VAD framework is required.

The personal VAD (PVAD) framework solves this problem by extending the traditional VAD to recognize only the target speaker part and ignore the non-target part. The feasibility of PVAD has been demonstrated in several studies. The author of [12, 13] proposed and improved the concept of PVAD to solve the problem of “always running” models on devices. [14, 15] extends the PVAD to make it easier to use in ASR applications. Moreover, the concept of PVAD has also been applied to the task of speech enhancement as a secondary task to improve the performance of separation [16, 17, 18]. However, these methods were not applicable to SV. Although the author of [19, 20, 21, 22, 23] tries to employ SD to recognize the speech of different speakers and then find the target speaker by some rules, this type of framework is inefficient for downstream speech tasks due to the complexity of its process.

Traditional PVAD frameworks are inadequate for SV tasks for several reasons: (1) Traditional PVAD models are trained with frame-by-frame supervision based on human-assigned or ASR forced alignment labels. However, not every frame predicted by the VAD model has a positive impact on the SV model. In practice, segments identified as speech by traditional VAD are sometimes associated with low SNR or multiple speakers talking simultaneously, which can severely degrade the performance of SV; (2) These labels are usually set as hard labels. However, soft labels are more suitable for SV scenarios because different speech segments contribute differently; (3) In SV, the VAD model faces a more complex situation and it will be more challenging to take all factors into account.

To address these issues, we have made the following contributions: (1) We propose a speaker verification-based voice activity detection (SVVAD) framework, which manipulates speech features using FiLM [24] according to their relevance to the SV model; (2) We propose a novel label-free training method that uses triplet-like losses to avoid the performance degradation caused by inaccurate human labeling; (3) Extensive experiments demonstrate that SVVAD achieves significant improvements over the baseline model in terms of equal error rate (EER) under various conditions where other speakers are mixed at different ratios, and that the model-generated VAD decision boundary is highly consistent with human judgment.

2. Methodology

2.1. Recap of Personal VAD Model and Motivation

To show the advancement of our proposed voice activity detection (VAD) framework, we need to discuss previous approaches. In the conventional personal VAD (PVAD) frame-

* Corresponding author: Jianzong Wang, jzwang@188.com




Figure 1: The overview of PVAD framework.

work [12, 13, 19] (shown in Figure 1), first, a pre-trained speaker verification (SV) model computes the speaker embedding of the enrolled target user $v^{(\text{enr})}$ from his/her recordings. Then, the PVAD model takes this SV embedding as a priori and the speech features of the audio to be tested $x_t^{(\text{tst})}$ as inputs to make predictions p_t at frame t . It can be trained using binary cross entropy (BCE) loss. Eventually, this model is capable of determining whether the frame contains the target speaker.

To further discuss this approach, there is an *Assumption of SV*: (1) the enrolled speech is restricted to contain only one speaker’s voice, and usually stored in the form of speaker embeddings; (2) the tested audio includes not only the target speaker’s speech, but also the non-target speaker’s speech and non-speech. Consequently, the PVAD model is required to identify whether the tested speech contains the target speaker’s speech (tss) or its opposite (ntss) with its speaker embedding $v^{(\text{enr})}$, expressed as Equation 1. Where $p_t = [p_t^{\text{tss}}, p_t^{\text{ntss}}]$ can be represented as a posteriori for two categories. $x^{(\text{tst})} \in \mathbb{R}^{T \times F}$, $v^{(\text{enr})} \in \mathbb{R}^E$. E, F, T denote the embedding dimension, feature dimension and sequence length respectively.

$$p_t = \text{PVAD}(x_t^{(\text{tst})}, v^{(\text{enr})}) \quad (1)$$

Although the vanilla PVAD model can identify the target speaker’s speech, it is originally designed for non-SV scenarios. This framework is not optimal for SV because it is based on the hypothesis that retaining only the target speaker’s speech and ignoring other sounds can improve the performance of SV. It is sub-optimal because not all speech of the target speaker contributes positively to SV performance. Besides, this framework relies heavily on accurate labels. Therefore, a more advanced VAD model for SV is needed that can account for more complex situations.

2.2. Improving Personal VAD with Speaker Verification

In order to solve the issue described in § 2.1, we propose a speaker verification-based voice activity detection (SJVAD) framework with two novel approaches: (1) The FiLM layer [24] is used to construct soft VAD decisions to modify the speech features to automatically identify the most informative parts of the SV model; (2) A label-free training method is introduced to avoid considering overly complex situations.

Network Architecture for SVVAD Backbone: Figure 2 outlines the network architecture of the SVVAD backbone. It consists of a PVAD backbone and a FiLM [24] backend. Seen from the input side, this architecture has two branches that receive the speech features to be tested $x^{(\text{tst})}$ and the enrolled target speaker embedding $v^{(\text{enr})}$ inputs respectively. For the speech feature branch, inspired by [13, 22, 25], we use Conformer [26] (with N_{conf} encoder layers) as the speech feature extractor because of its ability to distinguish the features of different speakers more efficiently. To ensure consistency in the time dimension, we removed the subsampling operation from the original Conformer header. In this way, both the input $x_t^{(\text{tst})}$ and out-

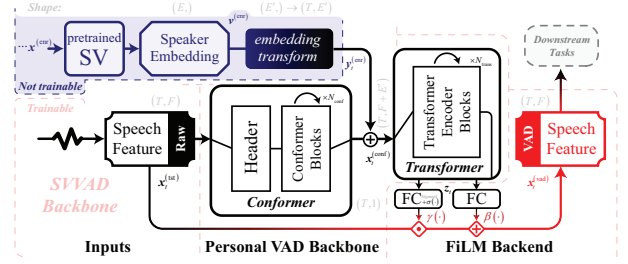


Figure 2: The network architecture of SVVAD backbone.

put $x_t^{(\text{conf})}$ have the same dimension $\mathbb{R}^{T \times F}$. For the speaker embedding branch, first, a pre-trained SV model is needed to convert the enrolled speech feature $x^{(\text{enr})}$ into speaker embedding $v^{(\text{enr})}$ with dimension \mathbb{R}^E . Since the embedding size is too large, it can then be reduced to $\mathbb{R}^{E'}$ by a fully connected layer. To merge with another branch, this embedding will be replicated T times to generate $y_t^{(\text{enr})}$ with dimension $\mathbb{R}^{T \times E'}$.

The information from these two branches is then combined by concatenation \oplus , with dimension $\mathbb{R}^{T \times (E' + F)}$. Since the target speaker’s information needs to affect every frame in the tested speech, this problem can be solved by using a self-attentive mechanism, i.e., through several layers of stacked Transformer encoder blocks (with N_{trans} layers). The output z_t is shown in Equation 2.

$$z_t = \text{Transformer}(y_t^{(\text{enr})} \oplus x_t^{(\text{conf})}) \quad (2)$$

The fused information z_t is then sent to two separate fully connected layers (FC) to create VAD decisions with $\gamma(z_t)$ and $\beta(z_t)$ with dimension $\mathbb{R}^{T \times 1}$. These decisions are then applied to the original speech features $x_t^{(\text{tst})}$ by the FiLM layer [24] to create $x_t^{(\text{vad})}$ with dimension $\mathbb{R}^{T \times F}$, in Equation 3. Since they have different dimensions, it is necessary to use one-to-many operations.

$$x_t^{(\text{vad})} = \text{FiLM}(z_t) = \gamma(z_t) \cdot x_t^{(\text{tst})} + \beta(z_t) \quad (3)$$

Label-free Optimization with Triplet-like Loss: To completely avoid constructing inaccurate strong VAD labels, it is effective to use only the speaker’s label for VAD-label-free training. Figure 3 summarizes the main structure of the SVVAD framework: the modified speech features $x^{(\text{vad})}$ are directly sent into the SV model to obtain the embedding $v^{(\text{vad})}$. To achieve label-free optimization during training, based on the assumptions of SV, we only need to set different optimization objectives for different cases of the test speech. Specifically, if the speech to be tested contains the voice of the target speaker, it is considered a positive sample, otherwise it is a negative sample. This idea is consistent with the triplet-loss (TL) mechanism [27], shown in Equation 4. Where α is the margin, B is the batch size. x_i^a, x_i^p and x_i^n is the anchor, positive and negative sample respectively at sample i , and can then be converted to speaker embeddings v_i^a, v_i^p and v_i^n by SV model. $\|\cdot\|_2$ is 2-norm. $[\cdot]_+$ is equivalent to $\max(\cdot, 0)$ or $\text{ReLU}(\cdot)$.




Figure 3: The overview of SVVAD framework.

$$\mathcal{L}_{\text{TL}} = \sum_i^B [\|v_i^a - v_i^p\|_2^2 - \|v_i^a - v_i^n\|_2^2 + \alpha]_+ \quad (4)$$

Since in SV, scoring uses the cosine distance instead of mean square error (MSE) in Equation 4, the triplet loss can be rewritten in terms of the cosine distance [28] (denoted as TLcos), in Equation 5. The cosine distance $\cos(\cdot)$ is expressed as 1 minus its cosine similarity.

$$\mathcal{L}_{\text{TLcos}} = \sum_i^B [\cos(v_i^a, v_i^p) - \cos(v_i^a, v_i^n) + \alpha]_+ \quad (5)$$

Since these losses only consider the distance between individual samples and not their overall statistical properties, they may lead to suboptimal performance. To this end, Lin’s concordance correlation coefficient (CCC) [29] loss is more reliable and can be used to stabilize the training process, so as to improve the VAD performance. Then, the Equation 4 can be changed to Equation 6 (denoted as TLccc), and Equation 5 can be modified to Equation 7, 8 (denoted as TLccos).

$$\mathcal{L}_{\text{TLccc}} = [\text{CCC}(v^a, v^p) - \text{CCC}(v^a, v^n) + \alpha]_+ \quad (6)$$

$$\mathcal{L}_{\text{TLccos}} = \begin{cases} \text{CCC} \left(\begin{array}{c} \cos(v^a, v^p) + \alpha, \\ \cos(v^a, v^n) \end{array} \right) & \text{cond} > 0 \\ 0 & \text{cond} \leq 0 \end{cases} \quad (7)$$

$$\text{cond} = \cos(v^a, v^p) - \cos(v^a, v^n) + \alpha \quad (8)$$

2.3. Training and Inference

Let’s start by discussing the inference phase, the model first accepts the target speaker’s embedding $v^{(\text{enr})}$ and the speech features to be tested $x^{(\text{tst})}$ to construct a VAD decision $\gamma(z_t)$ and $\beta(z_t)$. These decisions modify the original speech feature $x^{(\text{tst})}$ by FiLM rules to form $x^{(\text{vad})}$. It can then be converted into a speaker embedding $v^{(\text{vad})}$ by a pre-trained SV model. Finally, $v^{(\text{enr})}$ and $v^{(\text{vad})}$ will be scored by cosine similarity.

In the training phase, except for the same operations as in the inference phase, the data-loader will simultaneously construct the anchor, positive and negative samples x^a , x^p and x^n in real-time. These samples are fed into the trainable VAD and non-trainable SV model to obtain the speaker embeddings v^a , v^p and v^n , which are then trained by the triplet-like losses through Equation 4 to 7. However, since there is no existing dataset of multi-speaker speech, we will adopt a special method to generate these samples.

2.4. Training Data Generation

Multi-speaker samples are generated by concatenating samples from the single-speaker dataset, shown in pseudo-code of Algorithm 1. Where the sign $*$ is either positive p or negative n . D_* is speaker IDs. S_* is multi-speaker speech samples. p_{spk} and p_{overlap} are the probability of speaking and the speech overlap. ϕ represents the empty set or the silence. $\text{length}(\cdot)$ is to obtain the length of the sample. $\text{random}()$ is a value sampled uniformly between 0 and 1. $\text{get}(\cdot)$ is to get audio samples of random duration by the given id or silence ϕ . $\text{select}(\cdot)$ is to randomly select speaker ids. $\text{concat}(\cdot, \cdot)$ is to concatenate the former and the latter to form a longer sequence. $\text{overlapConcat}(\cdot)$ is to operate concatenating but overlapped by a certain ratio. Audio augmentation is performed as follows: all samples are augmented by probabilistically adding noise with SNR of 10 dB to 30 dB, and RIR reverberation. With this approach, the anchor, positive and negative samples (S_a , S_p

Algorithm 1 Training Data Generation Policy

Input: Dataset with single-speaker samples
Output: Multi-speaker samples with anchor, positive, negative

- 1: **for** each iteration in training **do**
- 2: Randomly select the anchor speaker ID as d_a
- 3: Create anchor samples $S_a \leftarrow \text{get}(d_a)$
- 4: Randomly select the number of speakers (≤ 3) for positive or negative, and their corresponding speaker IDs as D_* , where $d_a \in D_p$ and $d_a \notin D_n$
- 5: Create empty samples for positive and negative $S_* \leftarrow \phi$
- 6: **while** $\text{length}(S_*) <$ a fixed duration **do**
- 7: */* the previous and current states are different */*
- 8: **if** $\text{random}() < p_{\text{spk}}$ **then**
- 9: $id \leftarrow \text{select}(D_*)$
- 10: $s_{id} \leftarrow \text{get}(id)$
- 11: **if** $\text{random}() < p_{\text{overlap}}$ **then**
- 12: $S_* \leftarrow \text{overlapConcat}(S_*, s_{id})$
- 13: **else**
- 14: $S_* \leftarrow \text{concat}(S_*, s_{id})$
- 15: **end if**
- 16: **else**
- 17: $s_\phi \leftarrow \text{get}(\phi)$
- 18: $S_* \leftarrow \text{concat}(S_*, s_\phi)$
- 19: **end if**
- 20: **end while**
- 21: Audio augmentation for sample S_a and S_*
- 22: **end for**

return S_a, S_p, S_n

and S_n) are created. Since the SV model is well trained and SV-VAD requires an explicit learning target, the sample x^a is used directly to construct x^p .

3. Experiments

3.1. Experimental Setup

To validate our proposed SVVAD framework, we employ a pre-trained SV model¹ that shares the same experimental setup with our framework. That is, the SV model uses ECAPA-TDNN [30] from SpeechBrain [31] trained on the Voxceleb dataset [32]. The SVVAD has trained with Voxceleb 1+2 dataset and tested on the test set constructed from the Voxceleb1 (cleaned) verification set by concatenating audio segments from different speakers with varying proportions \mathcal{P} (percentage of the duration of other speakers’ voices). Considering the variability inherent in the generation of synthetic data, each set of experimental data will be constructed 3 times. The test metrics are similar to the SV frameworks and are reported as the equal error rate (EER) in percentage (%) and minimum decision cost function (minDCF) at $P_{\text{target}} = 0.01$ with $C_{\text{FA}} = C_{\text{Miss}} = 1$ [30].

Two baseline methods were added to the experiment for comparison to assess the impact of the different methods on SV performance: (1) WebRTC-VAD [33] with default settings; (2) The traditional PVAD model (in § 2.1). For the SVVAD framework, the SV performance of the four variants of the triplet-like loss (in Equation 4 to 8) will be measured.

3.2. Implementation Details

For speech features, the entire framework uses the same Log-Mel-Filterbanks extractor as the pre-trained SV model [31]. In SVVAD, the size of the Conformer encoder is $F = 256$. The number of Conformer and Transformer layers in SVVAD is $N_{\text{conf}} = 4$ and $N_{\text{trans}} = 3$ respectively, and both have a feed-forward size of 256. The size of the speaker embedding and the shrunken embedding are $E = 192$ and $E' = 64$ respectively. The parameter α in triplet-like loss is 0.9, 0.5, 0.55, 0.55 in TL, TLcos, TLccc and TLccos respectively. The probability $p_{\text{spk}} = 0.9$ and $p_{\text{overlap}} = 0.3$.

¹<https://huggingface.co/speechbrain/spkrec-ecapa-voxceleb>

Table 1: Overall comparison of the proposed SVVAD with 4 triplet-like losses and the baseline models on EER (%) and minDCF (denoted as $C_{0.01}$). + denotes adding methods.

\mathcal{P}	0%		30%		50%		70%	
	EER	$C_{0.01}$	EER	$C_{0.01}$	EER	$C_{0.01}$	EER	$C_{0.01}$
Baselines								
w/o VAD	0.90	0.1104	2.27	0.2284	14.84	0.6950	34.70	0.9942
WebRTC	1.14	0.1215	2.31	0.2332	14.50	0.6733	33.94	0.9873
PVAD	6.22	0.5084	7.95	0.5621	9.89	0.6080	16.29	0.7893
SVVAD + Triplet-like Loss (ours)								
+TL	4.63	0.4412	6.39	0.4794	8.62	0.5893	18.32	0.8621
+TLcos	5.48	0.5253	7.15	0.5336	9.73	0.6012	21.37	0.9157
+TLccc	1.41	0.1201	1.89	0.1996	6.91	0.4912	13.42	0.7396
+TLccos	2.06	0.1926	4.30	0.3729	7.26	0.5433	16.30	0.7605

Table 2: Ablation study of the effect of the combination of two types of losses on SV performance on EER (%).

$\mathcal{L} / \mathcal{P}$	0%	30%	50%	70%
TL	4.63	6.39	8.62	18.32
+TLcos	4.17	5.76	7.75	17.29
TLccc	1.41	1.89	6.91	13.42
+TLccos	1.18	1.70	5.81	12.20

In the first stage, the model is optimized by SGD optimizer for fast convergence, with the learning rate of $1e-2$, the momentum of 0.9, and the weight decay of $4e-4$. The speech durations of training of anchor, positive and negative are 6, 8, 8 seconds. The batch size is $B = 8$. In the fine-tuning stage, the optimizer is switched to AdamW [34], and the learning rate is changed to $1e-4$, and the weight decay is $2e-5$. The speech durations of these three are changed to 8, 12, 12 seconds. The batch size is $B = 64$ with gradient accumulation.

3.3. Evaluation Results

Overall Evaluation: Table 1 compares the EER and minDCF of various methods at different \mathcal{P} . As for the baseline, the WebRTC-VAD does not enhance the SV performance, sometimes even worse. The PVAD model can improve the SV performance to some extent, but the improvement is limited. For SVVAD, among 4 triplet-like losses, the addition of CCC leads to a significant improvement in the SV performance. But surprisingly, the cosine distance-based loss, which is theoretically more suitable because it is used for scoring, perform worse than the other losses. In our analysis, this may be due to the fact that the cosine distance-based loss considers only one value as the learning target instead of using the E -dimensional embedding for optimization, which may lead to a large amount of loss of embedding information when training SVVAD, thus making it difficult to converge. Furthermore, there is a noticeable drop in SV performance when no one else is speaking ($\mathcal{P} = 0\%$), which is unsatisfactory.

Ablation Study: To solve the above problem, we combine two types of losses, which can not only effectively reduce the ambiguity of the learning objectives, but also optimize the cosine distance for scoring. The experimental results are reported in Table 2. For faster training, the model is trained on the basis of TL and TLccc. The total loss \mathcal{L} consists of two sub-losses, they are $(\mathcal{L}_{TLcos} + \xi \cdot \mathcal{L}_{TL})$ or $(\mathcal{L}_{TLccos} + \xi \cdot \mathcal{L}_{TLccc})$, where ξ is the hyperparameter determined by the experiment (set as 0.1). The results demonstrate that by integrating these two sub-losses, the SV performance is further improved when \mathcal{P} is larger. In addition, there is no major performance degradation at $\mathcal{P} = 0\%$.

Figure 4: The speech features of $\mathbf{x}^{(tst)}$ and $\mathbf{x}^{(vad)}$ with their VAD decision boundaries (click on the figure to hear the sound).


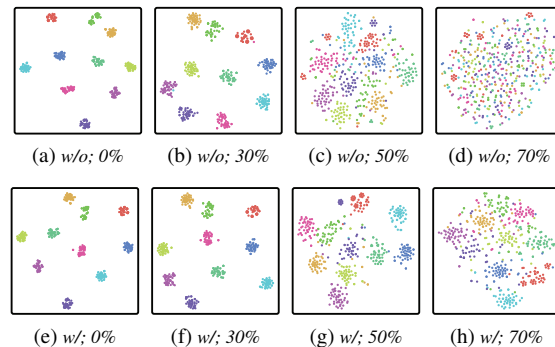


Figure 5: The t-SNE plots of $\mathbf{v}^{(vad)}$ for 10 example speakers in conditions of (1) with or without SVVAD; (2) the different \mathcal{P} .



The best-performing model achieves relative EER reductions over PVAD of 78.6%, 41.3% and 25.1% when \mathcal{P} is 30%, 50% and 70% respectively.

Generated Samples: According to Figure 5, SVVAD produces more aggregated speaker embeddings when \mathcal{P} is large, indicating that the model can achieve a greater improvement in SV performance when the input is mixed with more speech from non-target speakers. Conversely, when \mathcal{P} is small, the compactness of the cluster looks similar, suggesting that the SV performance is nearly the same. Figure 4 illustrates the decision boundaries generated by SVVAD, indicating that, similar to human perception, the model attempts to mask out the speech of non-target speakers and reduce strong noise, thereby suppressing their impact on SV performance degradation. However, the ambient noise may be negligible probably because the SV model is inherently robust to it.

4. Conclusions

In this paper, we present SVVAD, a speaker verification-based voice activity detection framework according to which speech features are most beneficial for SV. We also introduce a label-free training method that optimizes with the triplet-like loss approach, without relying on human labeling. Extensive experiments on various scenarios show that SVVAD outperforms the baselines in terms of EER in noisy and multi-speaker conditions. Furthermore, the decision boundaries reveal the importance of different parts of speech for SV, which is largely consistent with human perception. This work opens up new possibilities for developing more robust and accurate VADs for SV systems in real-world applications.

5. Acknowledgements

Supported by the Key Research and Development Program of Guangdong Province (grant No. 2021B0101400003) and Corresponding author is Jianzong Wang (jzwang@188.com).

6. References

- [1] J. Sohn, N. S. Kim, and W. Sung, "A statistical model-based voice activity detection," *IEEE signal processing letters*, vol. 6, no. 1, pp. 1–3, 1999.
- [2] Z.-H. Tan, N. Dehak *et al.*, "rvad: An unsupervised segment-based robust voice activity detection method," *Computer speech & language*, vol. 59, pp. 1–21, 2020.
- [3] T. Kinnunen and P. Rajan, "A practical, self-adaptive voice activity detector for speaker verification with noisy telephone and microphone data," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2013, pp. 7229–7233.
- [4] E. Sarkar, R. Prasad, and M. Magimai-Doss, "Unsupervised voice activity detection by modeling source and system information using zero frequency filtering," in *International Speech Communication Association (INTERSPEECH)*, 2022.
- [5] A. Vafeiadis, E. Fanioudakis, I. Potamitis, K. Votis, D. Giakoumis, D. Tzovaras, L. Chen, and R. Hamzaoui, "Two-dimensional convolutional recurrent neural networks for speech activity detection," in *International Speech Communication Association (INTERSPEECH)*. IEEE, 2019.
- [6] G.-B. Wang and W.-Q. Zhang, "An rnn and crnn based approach to robust voice activity detection," in *2019 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2019, pp. 1347–1350.
- [7] H. Dinkel, Y. Chen, M. Wu, and K. Yu, "Voice Activity Detection in the Wild via Weakly Supervised Sound Event Detection," in *International Speech Communication Association (INTERSPEECH)*, 2020, pp. 3665–3669.
- [8] Z. Zheng, J. Wang, N. Cheng, J. Luo, and J. Xiao, "Mlnet: An adaptive multiple receptive-field attention neural network for voice activity detection," in *International Speech Communication Association (INTERSPEECH)*, 2020.
- [9] J. Kim and M. Hahn, "Voice activity detection using an adaptive context attention model," *IEEE Signal Processing Letters*, vol. 25, no. 8, pp. 1181–1185, 2018.
- [10] I. Ariav and I. Cohen, "An end-to-end multimodal voice activity detection using wavenet encoder and residual networks," *Journal of Selected Topics in Signal Processing*, vol. 13, no. 2, pp. 265–274, 2019.
- [11] F. Tao and C. Busso, "End-to-end audiovisual speech activity detection with bimodal recurrent neural models," *Speech Communication*, vol. 113, pp. 25–35, 2019.
- [12] S. Ding, Q. Wang, S.-Y. Chang, L. Wan, and I. Lopez Moreno, "Personal VAD: Speaker-Conditioned Voice Activity Detection," in *Proc. Odyssey 2020 The Speaker and Language Recognition Workshop*, 2020, pp. 433–439. [Online]. Available: <http://dx.doi.org/10.21437/Odyssey.2020-62>
- [13] S. Ding, R. Rikhye, Q. Liang, Y. He, Q. Wang, A. Narayanan, T. O'Malley, and I. McGraw, "Personal vad 2.0: Optimizing personal voice activity detection for on-device speech recognition," in *International Speech Communication Association (INTERSPEECH)*, 2022.
- [14] A. Jayasimha and P. Paramasivam, "Personalizing speech start point and end point detection in asr systems from speaker embeddings," in *Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 771–777.
- [15] Q. Wang, I. L. Moreno, M. Saglam, K. Wilson, A. Chiao, R. Liu, Y. He, W. Li, J. Pelecanos, M. Nika *et al.*, "Voicefilter-lite: Streaming targeted voice separation for on-device speech recognition," *arXiv preprint arXiv:2009.04323*, 2020.
- [16] R. Giri, S. Venkataramani, J.-M. Valin, U. Isik, and A. Krishnaswamy, "Personalized perceptnet: Real-time, low-complexity target voice separation and enhancement," in *International Speech Communication Association (INTERSPEECH)*, 2021.
- [17] Y. Ju, W. Rao, X. Yan, Y. Fu, S. Lv, L. Cheng, Y. Wang, L. Xie, and S. Shang, "Tea-pse: Tencent-ethereal-audio-lab personalized speech enhancement system for icassp 2022 dns challenge," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 9291–9295.
- [18] S. E. Eskimez, T. Yoshioka, H. Wang, X. Wang, Z. Chen, and X. Huang, "Personalized speech enhancement: New models and comprehensive evaluation," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 356–360.
- [19] 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 *International Speech Communication Association (INTERSPEECH)*, 2020.
- [20] X. Xiao, N. Kanda, Z. Chen, T. Zhou, T. Yoshioka, S. Chen, Y. Zhao, G. Liu, Y. Wu, J. Wu *et al.*, "Microsoft speaker diarization system for the voxceleb speaker recognition challenge 2020," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5824–5828.
- [21] K. Wang, X. Mao, H. Wu, C. Ding, C. Shang, R. Xia, and Y. Wang, "The bytedance speaker diarization system for the voxceleb speaker recognition challenge 2021," *arXiv preprint arXiv:2109.02047*, 2021.
- [22] M. Cheng, H. Wang, Z. Wang, Q. Fu, and L. Ming, "The whulibaba audio-visual speaker diarization system for the misp challenge 2022," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023.
- [23] M. He, D. Raj, Z. Huang, J. Du, Z. Chen, and S. Watanabe, "Target-speaker voice activity detection with improved i-vector estimation for unknown number of speaker," in *International Speech Communication Association (INTERSPEECH)*, August 2021.
- [24] E. Perez, F. Strub, H. De Vries, V. Dumoulin, and A. Courville, "Film: Visual reasoning with a general conditioning layer," in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, vol. 32, no. 1, 2018.
- [25] Y. Liu, E. Han, C. Lee, and A. Stolcke, "End-to-end neural diarization: From transformer to conformer," in *International Speech Communication Association (INTERSPEECH)*, 2021.
- [26] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented Transformer for Speech Recognition," in *Proceedings of International Speech Communication Association (INTERSPEECH)*, 2020, pp. 5036–5040.
- [27] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2015, pp. 815–823.
- [28] C. Li, X. Ma, B. Jiang, X. Li, X. Zhang, X. Liu, Y. Cao, A. Kannan, and Z. Zhu, "Deep speaker: an end-to-end neural speaker embedding system," *arXiv preprint arXiv:1705.02304*, 2017.
- [29] I. Lawrence and K. Lin, "A concordance correlation coefficient to evaluate reproducibility," *Biometrics*, pp. 255–268, 1989.
- [30] B. Desplanques, J. Thienpondt, and K. Demuyne, "ECAPA-TDNN: emphasized channel attention, propagation and aggregation in TDNN based speaker verification," in *International Speech Communication Association (INTERSPEECH)*, H. Meng, B. Xu, and T. F. Zheng, Eds. ISCA, 2020, pp. 3830–3834.
- [31] M. Ravanelli, T. Parcollet, P. Plantinga, A. Rouhe, S. Cornell, L. Lugosch, C. Subakan, N. Dawalatabad, A. Heba, J. Zhong, J.-C. Chou, S.-L. Yeh, S.-W. Fu, C.-F. Liao, E. Rastorgueva, F. Grondin, W. Aris, H. Na, Y. Gao, R. D. Mori, and Y. Bengio, "SpeechBrain: A general-purpose speech toolkit," 2021, arXiv:2106.04624.
- [32] A. Nagrani, J. S. Chung, W. Xie, and A. Zisserman, "Voxceleb: Large-scale speaker verification in the wild," *Computer Science and Language*, 2019.
- [33] WebRTC. [Online]. Available: <https://webrtc.org/>
- [34] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *International Conference on Learning Representations*, 2017.