



Personal VAD: Speaker-Conditioned Voice Activity Detection

Shaojin Ding^{*2} Quan Wang^{*1} Shuo-yiin Chang¹ Li Wan¹ Ignacio Lopez Moreno¹

¹Google Inc., USA ²Texas A&M University, USA

shjd@tamu.edu { quanw, shuoyiin, liwan, elnota } @google.com

Abstract

In this paper, we propose “personal VAD”, a system to detect the voice activity of a target speaker at the frame level. This system is useful for gating the inputs to a streaming on-device speech recognition system, such that it only triggers for the target user, which helps reduce the computational cost and battery consumption, especially in scenarios where a keyword detector is unpreferable. We achieve this by training a VAD-alike neural network that is conditioned on the target speaker embedding or the speaker verification score. For each frame, personal VAD outputs the probabilities for three classes: non-speech, target speaker speech, and non-target speaker speech. Under our optimal setup, we are able to train a model with only 130K parameters that outperforms a baseline system where individually trained standard VAD and speaker recognition networks are combined to perform the same task.

1. Introduction

In modern speech processing systems, voice activity detection (VAD) usually lives in the upstream of other speech components such as speech recognition and speaker recognition. As a gating module, VAD not only improves the performance of downstream components by discarding non-speech signals, but also significantly reduces the overall computational cost due to its relatively small size.

A typical VAD system uses a frame-level classifier on acoustic features to make speech/non-speech decisions for each audio frame (*e.g.* with 25ms width and 10ms step). Poor VAD systems could either mistakenly accept background noise as speech or falsely reject speech. False accepting non-speech as speech largely slows down the downstream automatic speech recognition (ASR) processing. It is also computationally expensive as ASR models are normally much larger than VAD models. On the other hand, false rejecting speech leads to deletion errors in ASR transcriptions (a few milliseconds of missed audio could remove an entire word). A good VAD model needs to work accurately in challenging environments, including noisy conditions, reverberant environments and environments with competing speech. Significant research has been devoted to finding the optimal VAD features and models [1, 2, 3, 4, 5, 6, 7, 8, 9]. In the literature, LSTM-based VAD is a popular architecture for sequential modeling of the VAD task, showing state-of-the-art performance [7, 8, 9].

In many scenarios, especially on-device speech recognition [10], the computational resources such as CPU, memory, and battery are typically limited. In such cases, we wish to run the computationally intensive components such as speech recognition only when the target user is talking to the device.

* Equal contribution. Shaojin performed this work as an intern at Google.

False triggering such components in the background while only speech signals from other talkers or TV noises are present would cause battery drain and bad user experience. Although such concerns usually can be easily addressed by introducing a keyword detection [11] (*a.k.a.* wake word detection [12]) model, in many applications, the users would largely prefer a more seamless and natural interaction with the voice assistant without having to speak a predefined keyword. Thus, having a tiny model that only passes through speech signals from the target user is very necessary, which is our motivation of developing the personal VAD system.

Although standard speaker recognition and speaker diarization techniques [13, 14, 15, 16, 17, 18] can be directly used for the same task, we argue that the personal VAD system is largely preferred here for a couple of reasons:

1. To minimize the latency of the whole system, an accept/reject decision is needed upon the arrival of each frame immediately, which prefers frame-level inference of the model. However, many state-of-the-art speaker recognition and diarization systems usually require window-based or segment-based inference, or even offline full-sequence inference.
2. To minimize battery consumption on the device, the model must be very small, while most speaker recognition and diarization models are pretty big (typically millions of parameters).
3. Unlike speaker recognition or diarization, in personal VAD, it is unnecessary to distinguish between different non-target speakers, as we only trigger downstream components for the target speaker.

In fact, we implemented a baseline system by directly combining a standard speaker verification model and a standard VAD model for the personal VAD task, as described in Section 2.2.1, and found that its performance is worse than a dedicated personal VAD model. To the best of our knowledge, this work is the first lightweight solution that aims at directly detecting the voice activity of a target speaker in real time.

The proposed personal VAD is a VAD-alike neural network, conditioned on the target speaker embedding or the speaker verification score. Instead of determining whether a frame is speech or non-speech in standard VAD, personal VAD extends the determination to three classes: non-speech, target speaker speech, and non-target speaker speech.

The rest of the paper is organized as following. In Section 2.1, we first briefly describe our speaker verification system, which will be used during the training of personal VAD. Then in Section 2.2, we propose four different architectures to achieve personal VAD. In the training of personal VAD, we first treat it as a three-class classification problem and use cross entropy loss to optimize the model. In addition, we noticed that the discriminativity between non-speech and non-target speaker

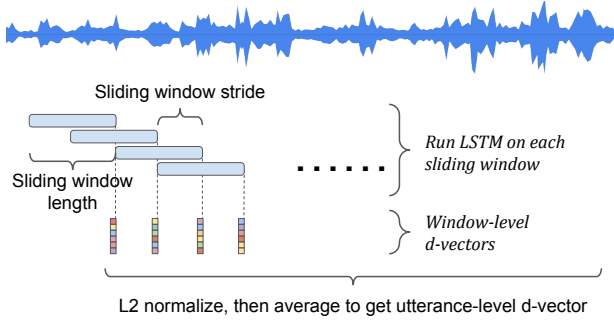


Figure 1: The speaker verification system [13] produces an utterance-level d-vector by aggregating window-level embeddings.

speech is relatively less important than between target speaker speech and the other two classes in personal VAD. Therefore, we further propose a *weighted pairwise loss* to enforce the model to learn these differences, as introduced in Section 2.3. We evaluate the model on an augmented version of the LibriSpeech dataset [19], with experimental setup described in Section 3.2, model configuration described in Section 3.3, metrics explained in Section 3.4, and results presented in Section 3.5. Conclusions are drawn in Section 4.

2. Approach

2.1. Recap of speaker verification system

Personal VAD relies on a pre-trained text-independent speaker recognition/verification model to encode the speaker identity into embedding vectors. In this work, we use the “d-vector” model introduced in [13], which has been successfully applied to various applications including speaker diarization [16, 17], speech synthesis [20, 21], source separation [22], speech translation [23], and audio voice preservation tests [24, 25]. We retrained the 3-layer LSTM speaker verification model using data from 8 languages for language robustness and better performance. During inference, the model produces embeddings on sliding windows, and a final aggregated embedding named “d-vector” is used to represent the voice characteristics of this utterance, as illustrated in Fig. 1. The cosine similarity between two d-vector embeddings can be used to measure the similarity of two voices.

In a real application, users are required to follow an *enrollment* process before enabling speaker verification or personal VAD. During enrollment, d-vector embeddings are computed from the target user’s recordings, and stored on the device. Since the enrollment is a one-off experience and can happen on server-side, we can assume that the embeddings of the target speakers are available at runtime with no cost.

2.2. System architecture

A personal VAD system should produce frame-level class labels for three categories: non-speech (ns), target speaker speech (tss), and non-target speaker speech (ntss). We implemented four different architectures to achieve personal VAD, as illustrated by Fig. 2. All four architectures rely on the embedding of the target speaker, which is acquired via the enrollment process.

2.2.1. Score combination (SC)

Our first approach to implement personal VAD is to simply combine a standard pre-trained speaker verification system and a standard VAD system, as shown in Fig. 2(a). We use this implementation as a **baseline** for other approaches, since it does not require training any new model.

We denote the frame of the input acoustic features at time t as $\mathbf{x}_t \in \mathbb{R}^D$, where D is the dimensionality of the acoustic features. For example, we use 40-dimensional log Mel-filterbank energies as the features. We use subscript $[t]$ to denote the subsequence ending at time t , i.e. $\mathbf{x}_{[t]} = (\mathbf{x}_1, \dots, \mathbf{x}_t)$. A standard VAD model $f_{\text{VAD}}(\cdot)$ and a speaker verification model $f_{\text{SV}}(\cdot)$ run independently on the acoustic features. The standard VAD produces unnormalized probabilities of speech (s) and non-speech (ns) for each frame:

$$\mathbf{z}_t = f_{\text{VAD}}(\mathbf{x}_{[t]}), \quad (1)$$

where $\mathbf{z}_t = [z_t^s, z_t^{\text{ns}}]$. The speaker verification model produces an embedding \mathbf{e}_t at each frame:

$$\mathbf{e}_t = f_{\text{SV}}(\mathbf{x}_{[t]}), \quad (2)$$

then the embedding is verified against the target speaker embedding $\mathbf{e}^{\text{target}}$, which was acquired during enrollment process:

$$s_t = \cos(\mathbf{e}_t, \mathbf{e}^{\text{target}}). \quad (3)$$

To transform the standard VAD probability z_t^s to personal VAD probabilities z_t^{tss} and z_t^{ntss} , we combined it with the resulting speaker verification cosine similarity score s_t , such that:

$$z_t^k = \begin{cases} s_t \cdot z_t^s & \text{if } k = \text{tss}; \\ (1 - s_t) \cdot z_t^s & \text{if } k = \text{ntss}; \\ z_t^{\text{ns}} & \text{if } k = \text{ns}. \end{cases} \quad (4)$$

There are two major disadvantages of this architecture. First, it is running a window-based speaker verification model at a frame level without any adaptation, and such inconsistency could cause significant performance degradation. However, training frame-level speaker verification models is often unscalable due to the difficulties to batch utterances of different length. Second, this architecture requires running a speaker verification system at runtime, which can be expensive since speaker verification models are usually much bigger than VAD models.

2.2.2. Score conditioned training (ST)

As shown in Fig. 2(b), our second approach uses the speaker verification model to produce a cosine similarity score s_t for each frame, as explained in Eq. (3), then concatenates this cosine similarity score to the acoustic features:

$$\hat{\mathbf{x}}_t = [\mathbf{x}_t, s_t], \quad (5)$$

The concatenated feature vector $\hat{\mathbf{x}}_t$ is 41-dimensional, as \mathbf{x}_t represents the 40-dimensional log Mel-filterbank energies. We train a new personal VAD network that takes the concatenated features as input, and outputs the probabilities of the three class labels for each frame:

$$\mathbf{z}_t = f_{\text{PVAD}}(\hat{\mathbf{x}}_{[t]}), \quad (6)$$

where $\mathbf{z}_t = [z_t^{\text{tss}}, z_t^{\text{ntss}}, z_t^{\text{ns}}]$.

This approach still requires running the speaker verification model at runtime. However, since it retrains the personal VAD model based on the speaker verification scores, it is expected to perform better than simply combining the scores of two individually trained systems.

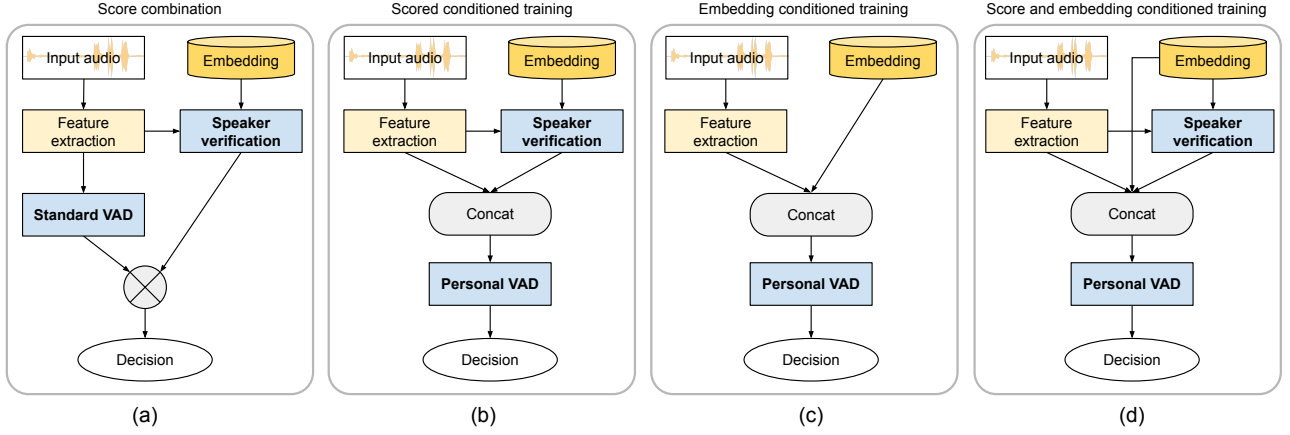


Figure 2: Four different architectures to implement personal VAD: (a) **SC**: Run standard VAD and frame-level speaker verification independently, and combine their results. This is used as a baseline for other approaches. (b) **ST**: Concatenate frame-level speaker verification score with acoustic features to train a personal VAD model. (c) **ET**: Concatenate speaker embedding with acoustic features to train a personal VAD model. (d) **SET**: Concatenate both speaker verification score and speaker embedding with acoustic features to train a personal VAD model.

2.2.3. Embedding conditioned training (ET)

As shown in Fig. 2(c), the third approach directly concatenates the target speaker embedding (acquired in the enrollment process) with the acoustic features:

$$\hat{\mathbf{x}}_t = [\mathbf{x}_t, \mathbf{e}^{\text{target}}]. \quad (7)$$

Since our embedding is 256-dimensional, the concatenated feature vector here is 296-dimensional. Then we train a new personal VAD network, which outputs the probabilities of three class at the frame level similar to Eq. (6).

This approach is similar to a knowledge distillation [26] process. The large speaker verification model was pre-trained on a large-scale dataset individually. Following this, when we train the personal VAD model, we use the speaker embeddings of the target speaker to “distill the knowledge” from the large speaker verification model to the small personal VAD model. As a result, it does not require running the large speaker verification model at runtime, which becomes the most **lightweight** solution among all architectures.

2.2.4. Score and embedding conditioned training (SET)

As shown in Fig. 2(d), this approach concatenates both the frame-level speaker verification score and the target speaker embedding to the acoustic features to train a new personal VAD model:

$$\hat{\mathbf{x}}_t = [\mathbf{x}_t, \mathbf{e}^{\text{target}}, s_t]. \quad (8)$$

The concatenated feature vector in this approach is 297-dimensional. This approach makes use of the most information from the speaker verification system. However, it still requires running the speaker verification model at runtime, so it’s not a lightweight solution.

2.3. Weighted pairwise loss

With an input frame \mathbf{x} and the corresponding ground truth label $y \in \{\text{ns}, \text{tss}, \text{ntss}\}$, personal VAD can be thought of as

a ternary classification problem.¹ The network outputs the unnormalized distribution of \mathbf{x} over the three classes, denoted as $\mathbf{z} = f_{\text{PVAD}}(\mathbf{x})$. We use z^k to denote the unnormalized probability of the k -th class. To train the model, we minimize the cross entropy loss as:

$$L_{\text{CE}}(y, \mathbf{z}) = -\log \frac{\exp(z^y)}{\sum_k \exp(z^k)}, \quad (9)$$

where $k \in \{\text{ns}, \text{tss}, \text{ntss}\}$.

However, in personal VAD, our goal is to detect the voice activity from only the target speaker. Audio frames that are classified into class ns and ntss will be discarded similarly by downstream components. As a result, confusion errors between $\langle \text{ns}, \text{ntss} \rangle$ have less impact to the system performance than errors between $\langle \text{tss}, \text{ntss} \rangle$ and $\langle \text{tss}, \text{ns} \rangle$. Inspired by Tuplemax loss [27], here we propose a *weighted pairwise loss* to model the different tolerance to each class pair. Given \mathbf{z} and y , we define weighted pairwise loss as:

$$L_{\text{WPL}}(y, \mathbf{z}) = -\mathbb{E}_{k \neq y} \left[w_{\langle k, y \rangle} \cdot \log \frac{\exp(z^y)}{\exp(z^y) + \exp(z^k)} \right], \quad (10)$$

where $w_{\langle k, y \rangle}$ is the weight between class k and class y . By setting lower weight to $\langle \text{ns}, \text{ntss} \rangle$ errors than $\langle \text{tss}, \text{ntss} \rangle$ and $\langle \text{tss}, \text{ns} \rangle$ errors, we can enforce the model to be more tolerant to the confusion between $\langle \text{ns}, \text{ntss} \rangle$ and to focus on distinguishing tss from ns and ntss .

3. Experiments

3.1. Datasets

An ideal dataset to train and evaluate personal VAD would be a dataset such that: (1) each utterance in it contains natural speaker turns; and (2) it contains enrollment utterances for each individual speaker. Unfortunately, to the best of our knowledge,

¹Without loss of generality, we ignore the subscript for the time dimension, and use \mathbf{x} to represent both original and concatenated input features in our notations here.

no public dataset in the community really satisfies both requirements. Although some datasets for speaker diarization [16] have natural speaker turns, they do not provide enrollment utterances for individual speakers. Alternatively, datasets containing enrollment utterances for individual speakers usually do not have natural speaker turns.

To address this limitation, we conducted experiments on an augmented version of the LibriSpeech dataset [19]. To simulate speaker turns, we concatenate single-speaker utterances from different speakers into multi-speaker utterances (see Section 3.2.1). We also noisify the concatenated utterances with reverberant room simulators to mitigate the concatenation artifacts (see Section 3.2.2).

In the LibriSpeech dataset, the training set contains 960 hours of speech, where 460 hours of them are “clean” speech and the other 500 hours are “noisy” speech. The testing set also consists of both “clean” and “noisy” speech. In all the experiments, we use the concatenated LibriSpeech training set to train the models. We use both the original LibriSpeech testing set and the concatenated LibriSpeech testing set for evaluation, as described in the following sections. For all the datasets, to produce the frame-level ground truth personal VAD labels used in training and evaluation, we run forced alignment with a pre-trained speech recognition model.

3.2. Experimental settings

3.2.1. Utterance concatenation

In the training corpora of standard VAD, each utterance usually only contains the speech from one single speaker. However, personal VAD aims to find the voice activity of a target speaker in a conversation where multiple speakers could be engaged. Therefore, we cannot directly use the standard VAD training corpora to train personal VAD. To simulate the conversational speech, we concatenate utterances from multiple speakers into a longer utterance, and then we randomly select one of the speakers as the target speaker in the concatenated utterance.

To generate a concatenated utterance, we draw a random number n indicating the number of utterances used for concatenation from a uniform distribution:

$$n \sim \text{Uniform}(a, b), \quad (11)$$

where a and b are the minimal and maximal numbers of utterances used for concatenation. The waveforms from the n randomly selected utterances are concatenated, and one of the speakers is assumed as the target speaker of the concatenated utterance. At the same time, we modify the VAD ground truth label of each frame according to the target speaker: “non-speech” frames remain the same, while “speech” frames are modified to either “target speaker speech” or “non-target speaker speech” according to whether the source utterance is from the target speaker.

In our experiments, we generated 300,000 concatenated utterances for training set and 5,000 concatenated utterances for testing sets. We use $a = 1$ and $b = 3$ for both sets, to cover both single-speaker and multi-speaker scenarios.

3.2.2. Multistyle training

For both training and evaluations, we apply a data augmentation technique named “multistyle training” (MTR) [28, 29, 30] on our datasets to avoid domain overfitting and mitigate concatenation artifacts. During MTR, the original (concatenated) source utterance is noisified with multiple randomly selected




Figure 3: Histogram of SNR (dB) of our multistyle training.

noise sources, using a randomly selected room configuration. Our noise sources include:

- 827 audios of ambient noises recorded in cafes;
- 786 audios recorded in silent environments;
- 6433 YouTube segments containing background music or noise.

We generated 3 million room configurations using a room simulator to cover different reverberation conditions. The distribution of the signal-to-noise ratio (SNR) of our MTR is shown in Fig. 3.

3.3. Model configuration

The acoustic features are 40-dimensional log Mel-filterbank energies, extracted on frames with 25ms width and 10ms step. For both standard VAD model and personal VAD model, we used a 2-layer LSTM network with 64 cells, followed by a fully-connected layer with 64 neurons. We also tried larger networks but did not see performance improvements, possibly due to the limited variety in training data. We used TensorFlow [31] for training and inference. During training, we used Adam optimizer [32] with a learning rate of 5×10^{-5} . For the models with weighted pairwise loss, we set $w_{\langle \text{tss}, \text{ns} \rangle} = w_{\langle \text{tss}, \text{ntss} \rangle} = 1$ and explored different values for $w_{\langle \text{ns}, \text{ntss} \rangle} \in \{0.01, 0.05, 0.1, 0.5, 1.0\}$.

To reduce the model size and accelerate the runtime inference, we quantized the parameters of the model to 8-bit integer values following [33]. With this quantization, our model using the ET architecture, which has only around 130 thousand parameters and is the smallest among all architectures (see Table 1), will be only 130 KB in size.

3.4. Metrics

To evaluate the performance of the proposed method, we computed the Average Precision (AP) [34] for each class and the mean Average Precision (mAP) over all the classes. AP and mAP are most common metrics for multi-class classification problems. AP summarizes a precision-recall curve as the weighted mean of precisions achieved at each threshold, with the increase in recall from the previous threshold used as the weight. AP can be computed as:

$$AP = \sum_n (R_n - R_{n-1}) \cdot P_n, \quad (12)$$

where R_n and P_n are the recall and precision at the n -th threshold, respectively. We adopted the micro-mean² over all the

²https://scikit-learn.org/stable/modules/generated/sklearn.metrics.average_precision_score.html

Table 1: Architecture and loss function comparison results. **SC**: Score combination, the baseline system. **ST**: Score conditioned training. **ET**: Embedding conditioned training. **SET**: Score and embedding conditioned training. **CE**: Cross entropy loss. **WPL**: Weighted pairwise loss ($w_{\langle ns, ntss \rangle} = 0.1$). We report the Average Precision (AP) for each class, and the mean Average Precision (mAP) over all the classes. Network parameters include 4.88 million parameters from the speaker verification (SV) model, if it is used during inference.

Method	Loss	Without MTR				With MTR				Network parameters (million)
		t _{ss}	ns	nt _{ss}	mean	t _{ss}	ns	nt _{ss}	mean	
SC (baseline)	N/A*	0.886	0.970	0.872	0.900	0.777	0.908	0.768	0.801	4.88 (SV) + 0.06 (VAD)
ST	CE	0.956	0.968	0.956	0.957	0.905	0.885	0.905	0.901	4.88 (SV) + 0.06 (PVAD)
ET		0.932	0.962	0.946	0.946	0.878	0.873	0.890	0.883	0.13 (PVAD)
SET		0.970	0.969	0.972	0.969	0.938	0.888	0.938	0.928	4.88 (SV) + 0.13 (PVAD)
ET	WPL	0.955	0.965	0.961	0.959	0.916	0.883	0.920	0.912	0.13 (PVAD)

* The baseline system does not require training any new model.

classes when computing mAP to take class imbalance into account, which averages APs over all the samples.

3.5. Results

We conducted three groups of experiments to evaluate the proposed method. First, we compared the four architectures for personal VAD. Following this, we examined the effectiveness of weighted pairwise loss and compared it against conventional cross entropy loss. Finally, we evaluated personal VAD on a standard VAD task, to see if personal VAD can replace standard VAD without performance degradation.

3.5.1. Architecture comparisons

In the first group of experiments, we compared the performance of four personal VAD architectures described in Fig. 2. We evaluated these systems on the concatenated LibriSpeech testing set. Additionally, to explore the performance of personal VAD on noisy speech, we also applied data augmentation technique (MTR) on the testing set. In personal VAD tasks, the most important metric is the AP for class t_{ss}, as downstream processes will only be applied to the speech produced by the target speaker.

We reported the evaluation results on the testing set with and without MTR, as shown in Table 1. Results show that ST, ET, and SET significantly outperform the baseline SC system in all cases. When applying MTR to the testing set, we observed an even larger performance gain between the proposed methods and the baseline. Among the proposed systems, SET achieved the highest AP for t_{ss}, and ST slightly outperforms ET. However, both ST and SET require to run speaker verification model (4.88 million parameters) to compute the cosine similarity score during inference time, which would largely increase both the number of parameters in the system and inference computational cost. By contrast, ET obtained 0.932 (without MTR) / 0.878 (with MTR) AP for class t_{ss} on the testing set with a model of only 0.13 million parameters (~ 40 times smaller), which is more appropriate for on-device applications.

3.5.2. Loss function comparisons

In the second group of experiments, we compared the proposed weighted pairwise loss against the conventional cross entropy loss. Here we only consider the ET architecture, as it is much more lightweight while achieving reasonably good performance. Similarly, we evaluated the systems on the concatenated LibriSpeech testing set with and without MTR.

In Fig. 4, we plot the AP for t_{ss} against different val-

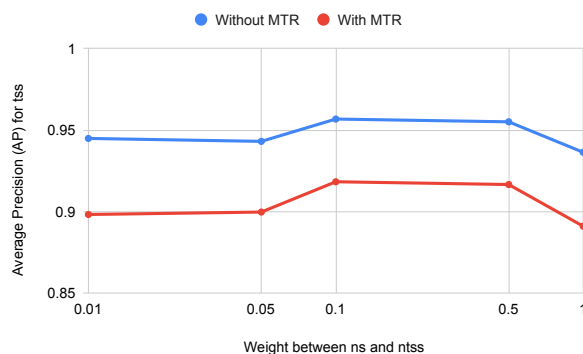


Figure 4: Mean Average Precision (mAP) of personal VAD (ET) with different values of $w_{\langle ns, ntss \rangle}$ in weighted pairwise loss. The weight between ns and nt_{ss} is displayed in log scale.

ues of $w_{\langle ns, ntss \rangle}$ in weighted pairwise loss. From the results, we observed that using a smaller value of $w_{\langle ns, ntss \rangle}$ than $w_{\langle tss, ns \rangle}$ and $w_{\langle tss, ntss \rangle}$ will improve the performance, which demonstrates that confusion errors between $\langle ns, ntss \rangle$ have less impact to the system performance than errors between $\langle tss, ntss \rangle$ and $\langle tss, ns \rangle$.

However, when $w_{\langle ns, ntss \rangle}$ becomes too small (e.g. 0.05 or 0.01), we found performance degradations from the curve. This result shows that completely ignoring the difference between nt_{ss} and ns is harmful to the system performance as well. In another word, it is insufficient to simply treat personal VAD task as a binary classification problem (target speaker speech v.s. other). The best performance is reached when setting $w_{\langle ns, ntss \rangle} = 0.1$, with detailed results listed in Table 1.

3.5.3. Personal VAD on standard VAD tasks

If we want to replace a standard VAD component with personal VAD, we also need to guarantee that the performance degradation on a standard speech/non-speech task is minimal. Finally, we conducted an experiment for personal VAD on standard VAD tasks. We evaluated two personal VAD models (ET architecture with cross entropy loss, and ET architecture with weighted pairwise loss) on the non-concatenated LibriSpeech testing data (so each utterance only has the target speaker). For comparison purpose, we also implemented a standard VAD model with the same network structure (2-layer LSTM network

with 64 cells, followed by a fully-connected layer with 64 neurons).

The results are shown in Table 2. We can see that the AP for class speech (s) is very close between personal VAD models and the standard VAD model, which justifies replacing standard VAD by personal VAD. Additionally, the architectures of personal VAD models and the standard VAD model are the same in this experiment, so replacing standard VAD by personal VAD will not increase the model size or computational cost at inference time.

Table 2: Evaluation on a standard VAD task. We report the Average Precision (AP) for speech (s) and non-speech (n.s).

Method	Loss	Without MTR		With MTR	
		s	n.s	s	n.s
Standard VAD	CE	0.992	0.975	0.975	0.918
Personal VAD (ET)	CE	0.991	0.965	0.979	0.893
Personal VAD (ET)	WPL	0.991	0.967	0.979	0.901

4. Conclusions

In this paper, we proposed four different architectures to implement personal VAD, a system that detects the voice activity of a target user in real time. Among the different architectures, using a single small network that takes acoustic features and enrolled target speaker embedding as inputs achieves near-optimal performance with smallest runtime computational cost. To model the tolerance to different types of errors, we proposed a new loss function, the weighted pairwise loss, which proves to have better performance than a conventional cross entropy loss. Our experiments also show that personal VAD and standard VAD perform equally well on a standard VAD task. In summary, our findings suggest that, by focusing only on the desired target speaker, a personal VAD can reduce the overall computational cost of speech recognition systems operating in noisy environments.

5. References

- [1] Samuel Thomas, George Saon, Maarten Van Segbroeck, and Shrikanth S Narayanan, “Improvements to the IBM speech activity detection system for the DARPA RATS program,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015.
- [2] Martin Graciarena, Abeer Alwan, Dan Ellis, Horacio Franco, Luciana Ferrer, John HL Hansen, Adam Janin, Byung Suk Lee, Yun Lei, Vikramjit Mitra, et al., “All for one: Feature combination for highly channel-degraded speech activity detection,” in *Interspeech*. ISCA, 2013, pp. 709–713.
- [3] Shuo-Yiin Chang, Bo Li, Gabor Simko, Tara N Sainath, Anshuman Tripathi, Aäron van den Oord, and Oriol Vinyals, “Temporal modeling using dilated convolution and gating for voice-activity-detection,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5549–5553.
- [4] Lee Ngee Tan, Bengt J Borgstrom, and Abeer Alwan, “Voice activity detection using harmonic frequency components in likelihood ratio test,” in *2010 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2010, pp. 4466–4469.
- [5] Seyed Omid Sadjadi and John HL Hansen, “Unsupervised speech activity detection using voicing measures and perceptual spectral flux,” *IEEE Signal Processing Letters*, vol. 20, no. 3, pp. 197–200, 2013.
- [6] Thomas Drugman, Yannis Stylianou, Yusuke Kida, and Masami Akamine, “Voice activity detection: Merging source and filter-based information,” *IEEE Signal Processing Letters*, vol. 23, no. 2, pp. 252–256, 2015.
- [7] Florian Eyben, Felix Weninger, Stefano Squartini, and Björn Schuller, “Real-life voice activity detection with lstm recurrent neural networks and an application to hollywood movies,” in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2013, pp. 483–487.
- [8] Shuo-Yiin Chang, Bo Li, Tara N Sainath, Gabor Simko, and Carolina Parada, “Endpoint detection using grid long short-term memory networks for streaming speech recognition,” in *Interspeech*. ISCA, 2017, pp. 3812–3816.
- [9] Matt Shannon, Gabor Simko, Shuo-Yiin Chang, and Carolina Parada, “Improved end-of-query detection for streaming speech recognition,” in *Interspeech*. ISCA, 2017, pp. 1909–1913.
- [10] Yanzhang He, Tara N Sainath, Rohit Prabhavalkar, Ian McGraw, Raziq Alvarez, Ding Zhao, David Rybach, Anjuli Kannan, Yonghui Wu, Ruoming Pang, et al., “Streaming end-to-end speech recognition for mobile devices,” in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6381–6385.
- [11] Guoguo Chen, Carolina Parada, and Georg Heigold, “Small-footprint keyword spotting using deep neural networks,” in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 4087–4091.
- [12] Kenichi Kumatani, Sankaran Panchapagesan, Minhua Wu, Minjae Kim, Nikko Strom, Gautam Tiwari, and

- Arindam Mandai, “Direct modeling of raw audio with dnns for wake word detection,” in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 252–257.
- [13] Li Wan, Quan Wang, Alan Papir, and Ignacio Lopez Moreno, “Generalized end-to-end loss for speaker verification,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4879–4883.
- [14] Chao Li, Xiaokong Ma, Bing Jiang, Xiangang Li, Xuewei Zhang, Xiao Liu, Ying Cao, Ajay Kannan, and Zhenyao Zhu, “Deep speaker: an end-to-end neural speaker embedding system,” *arXiv preprint arXiv:1705.02304*, 2017.
- [15] David Snyder, Daniel Garcia-Romero, Gregory Sell, Daniel Povey, and Sanjeev Khudanpur, “X-vectors: Robust dnn embeddings for speaker recognition,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5329–5333.
- [16] Quan Wang, Carlton Downey, Li Wan, Philip Andrew Mansfield, and Ignacio Lopez Moreno, “Speaker diarization with LSTM,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5239–5243.
- [17] Aonan Zhang, Quan Wang, Zhenyao Zhu, John Paisley, and Chong Wang, “Fully supervised speaker diarization,” in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6301–6305.
- [18] Arsha Nagrani, Joon Son Chung, and Andrew Senior, “Voxceleb: a large-scale speaker identification dataset,” *arXiv preprint arXiv:1706.08612*, 2017.
- [19] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: An ASR corpus based on public domain audio books,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [20] Ye Jia, Yu Zhang, Ron Weiss, Quan Wang, Jonathan Shen, Fei Ren, Patrick Nguyen, Ruoming Pang, Ignacio Lopez Moreno, Yonghui Wu, et al., “Transfer learning from speaker verification to multispeaker text-to-speech synthesis,” in *Advances in neural information processing systems*, 2018, pp. 4480–4490.
- [21] Xin Wang, Junichi Yamagishi, Massimiliano Todisco, Hector Delgado, Andreas Nautsch, Nicholas Evans, Md Sahidullah, Ville Vestman, Tomi Kinnunen, Kong Aik Lee, et al., “Asvspoof 2019: a large-scale public database of synthetic, converted and replayed speech,” *arXiv*, pp. arXiv–1911, 2019.
- [22] Quan Wang, Hannah Muckenhirn, Kevin Wilson, Prashant Sridhar, Zelin Wu, John R. Hershey, Rif A. Saurous, Ron J. Weiss, Ye Jia, and Ignacio Lopez Moreno, “VoiceFilter: Targeted Voice Separation by Speaker-Conditioned Spectrogram Masking,” in *Interspeech 2019*, 2019, pp. 2728–2732.
- [23] Ye Jia, Ron J. Weiss, Fadi Biadsy, Wolfgang Macherey, Melvin Johnson, Zhifeng Chen, and Yonghui Wu, “Direct Speech-to-Speech Translation with a Sequence-to-Sequence Model,” in *Interspeech 2019*, 2019, pp. 1123–1127.
- [24] W Bastiaan Kleijn, Felicia SC Lim, Alejandro Luebs, Jan Skoglund, Florian Stimberg, Quan Wang, and Thomas C Walters, “Wavenet based low rate speech coding,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 676–680.
- [25] Yutian Chen, Yannis Assael, Brendan Shillingford, David Budden, Scott Reed, Heiga Zen, Quan Wang, Luis C. Cobo, Andrew Trask, Ben Laurie, Caglar Gulcehre, Aaron van den Oord, Oriol Vinyals, and Nando de Freitas, “Sample efficient adaptive text-to-speech,” in *International Conference on Learning Representations (ICLR)*, 2019.
- [26] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean, “Distilling the knowledge in a neural network,” *arXiv preprint arXiv:1503.02531*, 2015.
- [27] Li Wan, Prashant Sridhar, Yang Yu, Quan Wang, and Ignacio Lopez Moreno, “Tuplex loss for language identification,” in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5976–5980.
- [28] Richard Lippmann, Edward Martin, and D Paul, “Multi-style training for robust isolated-word speech recognition,” in *1987 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. IEEE, 1987, vol. 12, pp. 705–708.
- [29] Tom Ko, Vijayaditya Peddinti, Daniel Povey, Michael L Seltzer, and Sanjeev Khudanpur, “A study on data augmentation of reverberant speech for robust speech recognition,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 5220–5224.
- [30] Chanwoo Kim, Ananya Misra, Kean Chin, Thad Hughes, Arun Narayanan, Tara Sainath, and Michiel Bacchiani, “Generation of large-scale simulated utterances in virtual rooms to train deep-neural networks for far-field speech recognition in Google Home,” in *Interspeech*, 2017, pp. 379–383.
- [31] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al., “TensorFlow: Large-scale machine learning on heterogeneous systems,” 2015, Software available from tensorflow.org.
- [32] Diederik P Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [33] Raziel Alvarez, Rohit Prabhavalkar, and Anton Bakhtin, “On the efficient representation and execution of deep acoustic models,” *arXiv preprint arXiv:1607.04683*, 2016.
- [34] Mu Zhu, “Recall, precision and average precision,” *Department of Statistics and Actuarial Science, University of Waterloo, Waterloo*, vol. 2, pp. 30, 2004.