

Learning Lip-Based Audio-Visual Speaker Embeddings with AV-HuBERT

Bowen Shi¹, Abdelrahman Mohamed², Wei-Ning Hsu²

¹Toyota Technological Institute at Chicago

²Meta AI

bshi@ttic.edu, {abdo,wnhsu}@fb.com

Abstract

This paper investigates self-supervised pre-training for audio-visual speaker representation learning where a visual stream showing the speaker’s mouth area is used alongside speech as inputs. Our study focuses on the Audio-Visual Hidden Unit BERT (AV-HuBERT) approach, a recently developed general-purpose audio-visual speech pre-training framework. We conducted extensive experiments probing the effectiveness of pre-training and visual modality. Experimental results suggest that AV-HuBERT generalizes decently to speaker related downstream tasks, improving label efficiency by roughly ten fold for both audio-only and audio-visual speaker verification. We also show that incorporating visual information, even just the lip area, greatly improves the performance and noise robustness, reducing EER by 38% in the clean condition and 75% in noisy conditions.¹

Index Terms: audio-visual, speaker verification and recognition, representation learning, self-supervised pre-training

1. Introduction

Personalizing user experiences is essential in spoken language technology systems, e.g., smart speakers and personal banking applications. Robust speaker verification (SV) and recognition models are crucial for enabling authentication and conversational experiences, as well as many other tasks like speaker diarization [1], voice conversion [2] and source separation [3].

Supervised speaker representation methods made significant progress over the past decade [4, 5, 6, 7]; however, they require a non-trivial amount of human annotations of speaker identity, which might not comply with the evolving privacy-preserving standards. Furthermore, it is challenging to provide speaker labels for multi-speaker dialogues or when speakers’ voices alternate between whispering and shouting [8]. Self-supervised speaker representation approaches, which work around these challenges, have recently gained popularity. One family of self-supervised speaker representation methods relies on contrastive learning, which constructs positive samples by either augmenting the same speech segment or assuming a single speaker is recorded per utterance [9, 10, 11]. They present solid downstream performance, where the unsupervised state-of-the-art approach [10] achieves EER of 1.66%, close to some of the SOTA supervised systems (e.g., 0.41% from [12]). However, one downside of these approaches is that they are tailored solely for speaker embedding tasks. In contrast, general self-supervised speech representation learning approaches, e.g., wav2vec 2.0 [13] and HuBERT [14], were found to capture enough speaker information to be competitive on SV while excelling at many other downstream tasks [15].

The Audio-Visual Hidden Unit BERT (AV-HuBERT) was recently introduced as a general audio-visual representation learning approach. It learns joint representations over speech and lip-movement streams by alternating between clustering representations using a small codebook mimicking broad phonetic units and learning latent contextual representations through the masked prediction loss. AV-HuBERT achieves SOTA results on lip-reading and audio-visual speech recognition (AVSR) under adverse noise conditions [16, 17], thanks to the noise-immune visual modality.

This paper goes beyond single modality speaker representations to work with audio and lip-movement information to learn noise-robust speaker embeddings. We extend the representation learned by the AV-HuBERT approach to study their effectiveness for speaker-based downstream tasks in multi-modal settings. Compared to recent specialized unsupervised speaker representation methods [9, 10, 11], one advantage of utilizing a general approach like AV-HuBERT is its ability to simultaneously serve other downstream tasks beyond speaker embedding. Prior work on audio-visual speaker representation learning focused on the consistency between the audio and the visual information, either by learning speaker embedding via audio-visual synchronization and identity matching [18] or by multi-way matching in a joint audio-visual embedding space [19]. AV-HuBERT offers a more stable training procedure than methods utilizing contrastive and consistency objectives since its masked prediction loss is computed over offline learned discrete units.

In our experiments, AV-HuBERT representations are used either in an ELMo-style feature combination protocol [20] or through fine-tuning the whole network for the target downstream task. We report our results on speaker classification and verification tasks under four types of interfering noise and five different signal-to-noise ratios (SNR). Our audio-visual models improve label efficiency by 10 folds from supervised models, and offer 38% and 75% relative equal error rate (EER) reduction for SV under clean and noisy conditions compared to audio-only pre-trained models.

2. Method

2.1. Overview of AV-HuBERT

Audio-Visual Hidden Unit BERT (AV-HuBERT) is a self-supervised model that learns from unlabeled audio-visual speech data. Similar to its audio counterpart — HuBERT [14], AV-HuBERT was initially benchmarked on speech recognition tasks and achieved state-of-the-art performances on unimodal (audio-only and video-only) [17] and multimodal (audio-visual) [16] setups. As depicted in Figure 1, AV-HuBERT comprises four modules: a feed-forward network (FFN) audio feature extractor, a modified ResNet [21, 22] video feature extractor, a fusion module, and a Transformer [23] backend. The two feature extractors generate frame-level representation for the corresponding stream, which are frame-wise concatenated

¹The analysis and all work described in this paper was performed by the author at Toyota Technological Institute at Chicago. Abdelrahman Mohamed and Wei-Ning Hsu served as advisors to the project.

Figure 1: AV-HuBERT for learning speaker embedding. Dashed box: added during fine-tuning.

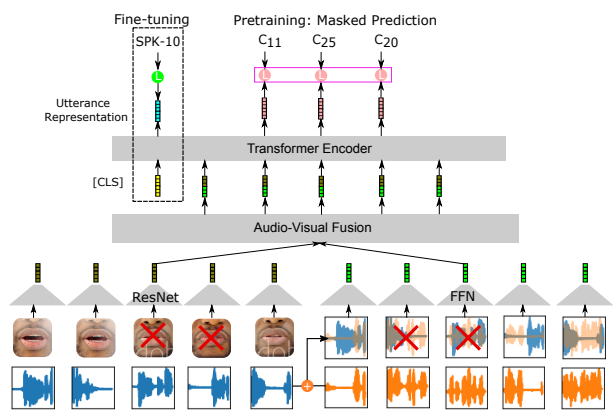


Figure 2: Example lip images, from LRS3 dataset [28]. Images in the same row are from the same speaker.



by the fusion module to form initial audio-visual features. The Transformer backend takes these features and produces the contextualized frame-level audio-visual representations. The entire model is optimized to perform masked prediction, where random segments are masked for each stream independently (denoted by the red crosses in Figure 1), and the model learns to predict the cluster assignment of the masked frames (the middle three frames in the example). The cluster assignment are iteratively refined: it is produced by clustering MFCC features in the first iteration, and by clustering previous iteration’s AV-HuBERT representations in the subsequent iterations.

While the self-supervised learning objective of HuBERT and AV-HuBERT (masked prediction of cluster assignments) resembles the task of automatic speech recognition (ASR) [17], several analyses observed that such models still learn rich speaker information especially in earlier layers [24, 15, 25]. These observations can be linked to many studies in the ASR and SV literature. For example, researchers have found that speaker information can improve ASR performance by encoding the effects to which an ASR system should be invariant [26]. Hence, AV-HuBERT may learn to extract speaker information at early layers to better infer phonetic information in the subsequent layers. In addition, statistics of the acoustic features extracted from ASR models have been widely used for SV [27], which can also explain why HuBERT features benefits SV. Motivated by the strong results of adapting HuBERT for SV, we explore learning audio-visual speaker embeddings from a pre-trained AV-HuBERT model with emphasis on the noise robustness aspect strengthened by the addition of visual stream.

2.2. Learning multimodal speaker embedding

We describe in this section how to learn speaker embeddings with pairs of audio-visual speech \mathbf{X} and speaker label y . Given a pre-trained AV-HuBERT model $f_{\theta_{avh}}$ parameterized by θ_{avh} , the goal is to leverage it to learn a speaker embedder $f_{\theta_{spk}}$ such that $f_{\theta_{spk}}(\mathbf{X}_i)$ are similar to $f_{\theta_{spk}}(\mathbf{X}_j)$ if $y_i = y_j$, and are dissimilar otherwise.

Two learning protocols are considered, which differs in whether $f_{\theta_{avh}}$ is frozen or not. The former is used to evaluate representation quality by treating $f_{\theta_{avh}}$ as a fixed feature extractor. Following [15], we consider ELMo [29] style fine-tuning, where frame-level AV-HuBERT features at each Transformer layer $\mathbf{C}^{(l)} = f_{\theta_{avh}}^{(l)}(\mathbf{X}) \in \mathbb{R}^{D_{avh} \times T}$ are weighted summed using learnable non-negative weights $w^{(l)} \in \mathbb{R}_{\geq 0}$ that sums to one: $\sum_{l=1}^L w^{(l)} = 1$. This representation is then passed to the downstream model f_{θ_d} to produce a fixed-sized speaker embedding: $\mathbf{s} = f_{\theta_d}(\sum_{l=1}^L w^{(l)} \mathbf{C}^{(l)}) \in \mathbb{R}^{D_{spk}}$. Only $w^{(l)}$ and θ_d are updated during fine-tuning. In this protocol, we adopt the same downstream models in [15].

The other protocol evaluates pre-training [30, 13, 14, 31], where θ_{avh} is updated along with new parameters during fine-tuning. We prepend a learnable [cls] embedding $\mathbf{e}_{cls} \in \mathbb{R}^{D_{avh}}$ to the input sequence of the Transformer encoder $\mathbf{Z} \in \mathbb{R}^{D_{avh} \times T}$. The augmented input $\tilde{\mathbf{Z}} \in \mathbb{R}^{D_{avh} \times (T+1)}$ is then passed to the Transformer encoder to produce $\tilde{\mathbf{C}}^{(L)} \in \mathbb{R}^{D_{avh} \times (T+1)}$ at the final layer L . The contextualized embedding corresponding to the [CLS] position, $\tilde{\mathbf{C}}_{cls}^{(L)} \in \mathbb{R}^{D_{avh}}$ serves as the speaker embedding \mathbf{s} .

For both protocols, we append a softmax layer that takes a speaker embedding as input and outputs the posterior $\text{softmax}(\mathbf{A}\mathbf{s}) \in [0, 1]^C$ over C training set speakers. Learnable parameters along with the softmax layer are fine-tuned to minimize a cross entropy loss with respect to speaker labels.

2.3. Trade-off between lip and face-based visual input

Visual input (lip or face video) can improve speaker embeddings on top of audio input from two aspects: first, it provides appearance features that is complementary to voice features, resulting in gain in clean and noisy conditions; second, it helps anchor the target speech from a mixture of sounds, making embeddings more resilient to noise. Most prior work use whole-face videos as visual input [32, 33]. While it can in principle lead to better performance, such a system also harvests more potentially sensitive data. In this work, we present a lip-based system (see Figure 2) that enjoys superior noise robustness (Table 2) while demanding less information. The accuracy trade-off is studied in Section 3.4.

3. Experiments

3.1. Setup

Pre-training We pre-training a 12-layer BASE and 24-layer LARGE AV-HuBERT following [17] and [16] with only one change. The dimension of features from BASE and LARGE model are 768 and 1,024 respectively. In [17], models are pre-trained with 433 hours of LRS3 [28] and the English portion of VoxCeleb2 (VC2) [34] from both the dev and test split. In this paper, since we evaluate the speaker embeddings on the VC2 test split and it contains both English and non-English speakers, we combine LRS3 with all VC2 data, English or non-English, but exclude the VC2 test split for pre-training, which sums up to roughly 2,800 hours. Both LRS3 and VC2 are sampled at a frame rate of 25Hz. The AV-HuBERT model also produces representations at this frame rate. To improve noise robustness, noise randomly sampled from MUSAN [35] is added to the audio stream following [16]. The audio and video preprocessing steps remain the same as [17].

Fine-tuning For the frozen protocol, to compare with results reported in [15], we adopt the same prediction heads

as [15], which is an average pooling layer for SC, and an x-vector model [36] for SV. Two widely used audio-visual speech recognition datasets, VoxCeleb1 (VC1, 352 hours/1,251 speakers) [37]² and VC2 (2,442 hours/6,112 speakers), are adopted for supervised fine-tuning. Noise-augmented fine-tuning with MUSAN following [16] is explored in §3.3. All models are optimized with Adam [39] with learning rate warmed up to 0.001 for one third of training steps and then linearly decayed. The pre-trained parameters are frozen for certain steps before being updated. For {5h, 50h, 500h, VC2} setups in Table 1, we train the model for {20, 30, 90, 75} K-steps with a batch size of {100, 100, 100, 400} and {60, 60, 120, 240} for audio-only and audio-visual setting. Whenever the model is used in audio-only setting, the visual feature (ResNet output) is replaced by an all-zero vector.

Evaluation Two tasks from [15], speaker verification (SV) and speaker classification (SC), are used. SC evaluates prediction accuracy of a closed set speaker seen during training. We report speaker accuracy (Acc) on VC1 using the official speaker identification split. SV considers a speaker-independent setup, where a set of testing trials is provided, each of which contains two utterances and a label indicating if the two are from the same speaker. We report EER on VC1 using its official test trials (37,611 trials) and on VC2 by sampling one positive trial and one negative trial for each test set utterance (72,474 trials). In terms of trial scoring, we adopt the protocol used in [5]: for any pair of audios \mathbf{X}_i and \mathbf{X}_j , ten 4-second segments ($\mathbf{X}_{i,n}$ and $\mathbf{X}_{j,m}$ for $n, m \in 1, \dots, 10$) are evenly sampled from each audio and we use average cosine similarity of the segment speaker embeddings ($\mathbf{s}_{i,n}$ and $\mathbf{s}_{j,m}$) between all combinations, $avg(\sum_{n=1}^{10} \sum_{m=1}^{10} \cos(\mathbf{s}_{i,n}, \mathbf{s}_{j,m}))$, as their matching score.

To probe noise robustness, we follow [16] to create 20 noisy test sets for VC1 and VC2, where each clean set is mixed with {Babble, Speech, Music, Other} noise at a SNR in {-10, -5, 0, 5, 10} dB. By default, we fine-tune all models on VC2 without noise augmentation and with the protocol where AV-HuBERT parameters are updated, and report EER on the clean set. We use AV-HuBERT Base for analysis purpose (sec 3.2-3.4).

3.2. Effectiveness of AV-HuBERT pre-training

We first study label efficiency of AV-HuBERT pre-training for audio-only and audio-visual speaker verification, and evaluate their performance on both the clean and noisy test trials (Table 1). In the “noisy” columns, we report the average EER over the 20 testing configurations. Three sizes of subsets are considered: 20%/2%/0.2% of the labeled data. For each size, we generate subsets in two ways: one is by sampling $x\%$ of the utterances, and the other is by sampling $x\%$ of the speakers and select all their utterances. For a given size, the former would contain more speakers but less utterances per speaker.

For each labeled subset, AV-HuBERT outperforms models trained from scratch (PT=None) when the same modalities are used as input. Comparing across subsets, we find AV-HuBERT often matches or outperforms a randomly initialized model that use 10 times more data (e.g., AV-HuBERT is 16.7% on VC2 clean set with VC2-5h data and AV input, while the baseline is 21.5% with VC2-50h data and AV input).

²The video files from [38] was down-sampled by a factor of six (~ 4.17 Hz). To tackle this, we let the ResNet video feature extractor to process the downsampled video as is and upsample its output by a factor of six before passing it to the Transformer. Empirically, this leads to similar performance as upsampling the video at input to the ResNet while reducing the memory and the compute.

Table 1: SV performance on clean and noisy test sets when fine-tuned with various VC2 subsets. The EER averaged over 20 setups (5 SNRs \times 4 types) is reported for the noisy test sets.

PT	FT	Mod.	VC2 EER (%)		VC1 EER (%)	
			clean	noisy	clean	noisy
None	VC2-15spk (5h)	A	26.8	39.2	25.1	39.2
None		AV	29.8	35.9	24.6	28.7
VC2+LRS3		A	23.3	33.9	20.0	33.0
VC2+LRS3		AV	22.6	28.0	19.4	21.9
None	VC2-156spk (50h)	A	18.5	34.5	16.1	34.6
None		AV	16.4	24.7	13.1	17.7
VC2+LRS3		A	11.8	28.9	9.4	29.1
VC2+LRS3		AV	9.3	18.8	7.8	12.5
None	VC2-1200spk (485h)	A	11.1	31.6	8.6	30.5
None		AV	9.3	17.6	7.0	9.9
VC2+LRS3		A	7.2	26.1	4.9	25.2
VC2+LRS3		AV	5.7	12.6	3.8	6.1
None	VC2-5h (1740spk)	A	24.4	39.4	21.7	39.5
None		AV	32.8	41.0	30.2	40.3
VC2+LRS3		A	20.1	34.7	17.7	34.5
VC2+LRS3		AV	16.7	28.6	13.9	23.0
None	VC2-50h (5113spk)	A	20.2	35.5	16.1	34.7
None		AV	21.5	26.3	15.7	16.4
VC2+LRS3		A	10.7	29.7	8.0	28.7
VC2+LRS3		AV	7.4	19.8	4.8	11.4
None	VC2-500h (5992spk)	A	10.6	33.1	8.0	31.4
None		AV	6.5	14.5	5.3	7.8
VC2+LRS3		A	4.9	23.7	3.0	22.8
VC2+LRS3		AV	3.7	9.2	1.7	3.9
None	VC2 (5994spk)	A	7.3	29.2	5.1	27.8
None		AV	5.1	11.3	2.9	4.7
VC2+LRS3		A	3.4	20.9	1.9	20.0
VC2+LRS3		AV	2.4	7.8	1.0	2.5

Next, we observe that with pre-training, audio-visual models always outperform audio-only ones. The consistent gain on the clean set shows that lip videos bring complementary information for recognizing speakers, which is different from speech recognition where visual information helps very little in clean conditions [16, 40]. On the other hand, audio-visual based models bring substantial gains in noisy conditions compared to audio-based models (e.g., reducing EER from 20.0% to 2.5% on noisy VC1 test set when pre-trained and fine-tuned on the entire VC2). This demonstrates the enhancement in noise robustness by incorporating lip information.

Comparing the different subset sampling strategies for 0.2% VC2 (VC2-5h versus VC2-15spk), it shows that pre-trained models benefit from having more speakers (w/ PT, 16.7% for VC2-5h vs. 22.6% for VC2-15spk on VC2 clean with AV input),³ while the supervised audio-visual baselines struggle with too few examples per speaker (w/o PT 32.8% vs. 29.8% on VC2 clean with AV input). This implies that a visual encoder trained from scratch hardly benefits from increasing number of speakers given only few utterances per speaker. In the 5h-setting (few utterances per speaker), the extra visual modality is only effective in case of pre-training whereas an audio-visual model trained from scratch is consistently worse than its audio-only counterpart regardless of auditory conditions (PT=None, A vs. AV). To conclude, we discover pre-training can lead to better generalization with few-shot learning.

³The model achieves similar performance on noisy test sets for both subset sampling strategies, which is caused by the mismatch in training/testing audio conditions (clean vs. noisy). With noise-augmented fine-tuning, the same trend emerges again (w/ PT, 5h vs. 15spk: 26.1%/13.8% vs. 30.7%/21.6% with A/AV input on VC1 noisy).

3.3. Noise-augmented fine-tuning

Observing the gap between audio-based and audio-visual models on SV in the previous section, we study if applying noise augmentation during fine-tuning can bridge the gap. We fine-tune AV-HuBERT on VC2-500h audio and audio-visual data with and without noise augmentation described in §3.1

The performance of the four models (A or AV, with or without noise augmentation) on the 20 noisy test sets are presented in Table 2. Unsurprisingly, all models perform worse on lower SNR conditions. Nevertheless, we observe much bigger degradation for audio-based model, especially when corrupted with speech (S) or babble (B) noise, because audio models can not determine who the target speaker is from a mixture of speech. In contrast, audio-visual based models suffer very minor performance degradation in noisier conditions, because lip videos can help identify the target to infer speaker embedding for.

We also see noise augmentation reduces the average EER at -10dB from 43.7% to 30.8% for audio-only models, and from 6.3% to 3.3% for audio-visual models. The results suggest that while noise-augmentation is beneficial, adopting it alone can not bridge the gap between audio and audio-visual models pre-trained and fine-tuned on the same amount of data.

Table 2: AV-HuBERT fine-tuned on VC2-500h with audio (A) or audio-visual (AV) input, with or without noise augmentation. The abbreviations used are B: Babble, S: Speech, M: Music, and O: Other.

Noise Aug?	Noise Type	A, VC1 EER (%), SNR (dB)=					AV, VC1 EER (%), SNR (dB)=				
		-10	-5	0	5	10	-10	-5	0	5	10
N	B	48.2	36.4	18.5	9.6	6.0	4.4	3.9	3.4	2.6	2.2
	S	48.8	46.5	36.5	18.3	8.5	8.6	6.8	4.8	3.4	2.5
	M	39.3	26.9	14.5	8.3	5.5	6.2	4.3	3.1	2.4	2.0
	O	34.0	23.3	13.7	8.8	5.9	6.0	4.3	3.2	2.6	2.3
Y	B	48.1	27.2	12.7	7.3	5.2	3.4	3.2	2.5	2.2	2.0
	S	24.4	14.9	11.8	12.3	9.6	3.2	2.8	2.6	2.3	2.0
	M	27.3	14.3	8.2	5.6	4.4	3.5	2.8	2.4	2.0	1.8
	O	23.6	13.0	8.0	5.8	4.7	3.1	2.6	2.3	2.1	2.0

3.4. Choice of visual input

Using lip instead of face videos is a key feature of the proposed model as it better addresses data input concerns while exhibits excellent noise robustness. We quantify how much our model degrades by comparing it with a variant of AV-HuBERT pre-trained and fine-tuned on face videos. As a baseline, we also evaluate a widely used face recognition model RetinaFace [41] taking a still image as input, which only achieves an EER of 14.6% on VC2. Table 3 shows that the face-based AV-HuBERT model indeed performs slightly better (0.9% absolute EER reduction), which is a trade-off to be considered.

Table 3: Comparing different input. AV-HuBERT is fine-tuned on VC2-500h.

Model	Input	VC2 EER (%)
AV-HuBERT	audio + face video	2.8
AV-HuBERT	audio + lip video	3.7

3.5. Comparison with prior work

We first compare AV-HuBERT with other self-supervised models using the SUPERB [15] protocol that evaluates the quality of frozen representations (upper half of Table 4), where models are fine-tuned and evaluated on VC1 for SC and SV. With audio-only input, AV-HuBERT-L outperforms wav2vec2-L and HuBERT-L and is inferior to WavLM-L. Outperforming all audio-only pre-trained models, AV-HuBERT achieves significantly better results on the SC and SV tasks with audio-visual input even with an order of magnitude less pre-training data.

Table 4: (Top) Comparison with the prior work following the SUPERB fine-tuning protocol. Models are fine-tuned on VC1. (Bottom) Comparison with prior work that does not follow the SUPERB evaluation protocol.

Method	PT Data	Mod.	VC1	
			SC-Acc	SV-EER
FBANK [15]	-	A	8.5E-4	9.56
wav2vec2-B [15]	LS (960 hr)	A	75.18	6.02
HuBERT-B [15]	LS (960 hr)	A	81.42	5.11
WavLM-B [25]	Mix (94k hr)	A	89.42	4.07
wav2vec2-L [15]	LL (60k hr)	A	86.14	5.65
HuBERT-L [15]	LL (60k hr)	A	90.33	5.98
WavLM-L [25]	Mix (94k hr)	A	95.49	3.77
AV-HuBERT-B		A	80.99	5.85
AV-HuBERT-B	VC2+LRS3	AV	93.90	4.85
AV-HuBERT-L	(2.8k hr)	A	91.56	4.42
AV-HuBERT-L		AV	98.06	2.95

Method	PT Data	FT Data	Mod.	VC1	VC1	VC2
				SC-Acc	SV-EER	SV-EER
Nagrani et al. [18]	20% VC2	VC1	AV	-	9.43	-
WavLM-L [25]	Mix (94k hr)	VC2	A	-	0.38	-
Shon et al. [32]	-	VC2	AV	-	-	5.29
Unimodal [33]	-	VC2	A	-	2.2	3.5
Multi-view [33]	-	VC2	AV	-	1.8	2.4
Feature Fusion [33]	-	VC2	AV	-	1.4	2.0
Ensemble [33]	-	VC2	AV	-	0.7	1.6
AV-HuBERT-B		VC2	A	-	1.92	3.43
AV-HuBERT-B	VC2+LRS3	VC2	AV	-	1.00	2.41
AV-HuBERT-L	(2.8k hr)	VC2	A	-	1.71	3.11
AV-HuBERT-L		VC2	AV	-	0.84	2.29

The lower half of Table 4 compares AV-HuBERT with other supervised and self-supervised models reporting results on VC1 or VC2 without following the SUPERB protocol. Specifically, we fine-tune our pre-trained model using the whole VC2 labeled data. As is expected, fine-tuning with a large amount of labeled data improves performance. In audio-visual setting, our best model (0.84%) outperforms [33] (1.8% and 1.4%) with a single model and slightly falls behind its ensembled model (0.7%). Note in contrast to the prior works [18, 33, 32] which uses the whole face, our model only relies on the lip area of the speaker as visual input and achieves a better trade-off between required information and performance. In addition, we acknowledge the gap between our best model and the current SOTA on VC1 ([25]: 0.38%). Such gap is attributed to the relative smaller amount of pre-training data we use as well as the simplicity of our training and evaluation pipeline (e.g., [25] uses Inter-TopK penalty [12], ECAPA-TDNN [6] as an upstream network, second stage large-margin fine-tuning, adaptive s-norm [42, 43] and calibration [44] in evaluation, etc). As the goal of this paper is to study the speaker embedding learned by AV-HuBERT, we stick to a simple and standard pipeline for speaker verification. How to incorporate extra techniques in AV-HuBERT framework is left as our future work.

4. Conclusion

In this paper, we study learning lip-based audio-visual speaker embeddings from AV-HuBERT. We show that the general-purpose audio-visual speech representation learning framework, AV-HuBERT, is able to learn high-quality speaker embeddings that improves the performance on speaker related tasks including classification and verification using less labeled data. The proposed model also greatly improves robustness to a variety of noise while requiring only lip videos instead of whole face videos.

5. References

- [1] Q. Wang, C. Downey, L. Wan, P. A. Mansfield, and I. Lopez-Moreno, "Speaker diarization with lstm," 2018, pp. 5239–5243.
- [2] L. Zhang, C. Yu, H. Lu, C. Weng, C. Zhang, Y. Wu, X. Xie, Z. Li, and D. Yu, "DurIAN-SC: Duration informed attention network based singing voice conversion system," in *Interspeech*, 2020.
- [3] Q. Wang *et al.*, "Voicefilter: Targeted voice separation by speaker-conditioned spectrogram masking," in *Interspeech*, 2019.
- [4] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-Vectors: Robust DNN embeddings for speaker recognition," in *ICASSP*, 2018.
- [5] J. S. Chung, J. Huh, S. Mun, M. Lee, H. S. Heo, S. Choe, C. Ham, S. Jung, B.-J. Lee, and I. Han, "In defence of metric learning for speaker recognition," in *Interspeech*, 2020.
- [6] B. Desplanques, J. Thienpondt, and K. Demuynck, "ECAPA-TDNN: Emphasized Channel Attention, propagation and aggregation in TDNN based speaker verification," in *Interspeech*, 2020, pp. 3830–3834.
- [7] K. A. Lee, O. Sadjadi, H. Li, and D. Reynolds, "Two decades into speaker recognition evaluation - are we there yet?" *Computer Speech and Language*, 2020.
- [8] J. Hansen and T. Hasan, "Speaker recognition by machines and humans: A tutorial review," *Signal Processing Magazine, IEEE*, vol. 32, pp. 74–99, 11 2015.
- [9] N. Inoue and K. Goto, "Semi-supervised contrastive learning with generalized contrastive loss and its application to speaker recognition," in *APSIPA ASC*, 2020.
- [10] R. Tao, K. A. Lee, R. Das, V. Hautamäki, and H. Li, "Self-supervised speaker recognition with loss-gated learning," *arXiv preprint arXiv:2110.03869*, 2021.
- [11] W. Xia, C. Zhang, C. Weng, M. Yu, and D. Yu, "Self-supervised text-independent speaker verification using prototypical momentum contrastive learning," in *ICASSP*, 2021.
- [12] M. Zhao, Y. Ma, M. Liu, and M. Xu, "The SpeakIn system for voxceleb speaker recognition challenge 2021," *ArXiv*, vol. abs/2109.01989, 2021.
- [13] A. Baeovski, H. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *NeurIPS*, 2020.
- [14] W.-N. Hsu *et al.*, "HuBERT: Self-supervised speech representation learning by masked prediction of hidden units," in *ICASSP*, 2021.
- [15] S. Yang *et al.*, "SUPERB: Speech processing Universal Performance Benchmark," in *Interspeech*, 2021.
- [16] B. Shi, W.-N. Hsu, and A. Mohamed, "Robust self-supervised audio-visual speech recognition," *arXiv preprint arXiv:2201.01763*, 2022.
- [17] B. Shi, W.-N. Hsu, K. Lakhotia, and A. Mohamed, "Learning audio-visual speech representation by masked multimodal cluster prediction," 2022.
- [18] A. Nagrani, J. S. Chung, S. Albanie, and A. Zisserman, "Disentangled speech embeddings using cross-modal self-supervision," in *ICASSP*, 2020, pp. 6829–6833.
- [19] S.-W. Chung, J. S. Chung, and H.-G. Kang, "Perfect match: Self-supervised embeddings for cross-modal retrieval," *IEEE Journal of Selected Topics in Signal Processing*, vol. 14, pp. 568–576, 2020.
- [20] J. Sarzynska-Wawer *et al.*, "Detecting formal thought disorder by deep contextualized word representations," *Psychiatry Research*, vol. 304, p. 114135, 2021.
- [21] T. Stafylakis and G. Tzimiropoulos, "Combining residual networks with LSTMs for lipreading," in *Interspeech*, 2017.
- [22] B. Martínez, P. Ma, S. Petridis, and M. Pantic, "Lipreading using temporal convolutional networks," *ICASSP*, 2020.
- [23] A. Vaswani *et al.*, "Attention is all you need," in *NeurIPS*, 2017.
- [24] H.-J. Chang, S.-w. Yang, and H.-y. Lee, "Distilhubert: Speech representation learning by layer-wise distillation of hidden-unit bert," *arXiv preprint arXiv:2110.01900*, 2021.
- [25] S. Chen *et al.*, "WavLM: Large-scale self-supervised pre-training for full stack speech processing," *arXiv preprint arXiv:2110.13900*, 2021.
- [26] A. Senior and I. Lopez-Moreno, "Improving DNN speaker independence with i-vector inputs," in *ICASSP*, 2014.
- [27] Y. Lei, N. Scheffer, L. Ferrer, and M. McLaren, "A novel scheme for speaker recognition using a phonetically-aware deep neural network," in *ICASSP*, 2014.
- [28] T. Afouras, J. S. Chung, and A. Zisserman, "LRS3-TED: a large-scale dataset for visual speech recognition," 2018, arXiv:1809.00496.
- [29] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, "Deep contextualized word representations," in *NAACL*, 2018.
- [30] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *NAACL*, 2019.
- [31] A. Baeovski, W.-N. Hsu, Q. Xu, A. Babu, J. Gu, and M. Auli, "Data2vec: A general framework for self-supervised learning in speech, vision and language," *arXiv preprint arXiv:2202.03555*, 2022.
- [32] S. Shon, T.-H. Oh, and J. Glass, "Noise-tolerant audio-visual online person verification using an attention-based neural network fusion," in *ICASSP*. IEEE, 2019, pp. 3995–3999.
- [33] L. Sari, K. Singh, J. Zhou, L. Torresani, N. Singhal, and Y. Saraf, "A multi-view approach to audio-visual speaker verification," 2021.
- [34] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep Speaker Recognition," in *Interspeech*, 2018.
- [35] D. Snyder, G. Chen, and D. Povey, "MUSAN: A Music, Speech, and Noise Corpus," *ArXiv*, vol. abs/1510.08484, 2015.
- [36] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust dnn embeddings for speaker recognition," in *ICASSP*. IEEE, 2018, pp. 5329–5333.
- [37] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: a large-scale speaker identification dataset," *arXiv preprint arXiv:1706.08612*, 2017.
- [38] A. Nagrani, S. Albanie, and A. Zisserman, "Seeing voices and hearing faces: Cross-modal biometric matching," in *CVPR*, 2018.
- [39] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *ICLR*, 12 2014.
- [40] P. Ma, S. Petridis, and M. Pantic, "End-to-end audio-visual speech recognition with conformers," in *ICASSP*, 2021.
- [41] J. Deng, J. Guo, E. Ververas, I. Kotsia, and S. Zafeiriou, "Retinaface: Single-shot multi-level face localisation in the wild," in *CVPR*, June 2020.
- [42] Z. N. Karam, W. M. Campbell, and N. Dehak, "Towards reduced false-alarms using cohorts," *ICASSP*, pp. 4512–4515, 2011.
- [43] S. Cumani, P. D. Bazu, D. Colibro, C. Vair, P. Laface, and V. Vasilakakis, "Comparison of speaker recognition approaches for real applications," in *Interspeech*, 2011.
- [44] J. Thienpondt, B. Desplanques, and K. Demuynck, "The Idlab Voxsrc-20 submission: Large margin fine-tuning and quality-aware score calibration in DNN based speaker verification," in *ICASSP*, 2021.