



Can you Remove the Downstream Model for Speaker Recognition with Self-Supervised Speech Features?

Zakaria Aldeneh¹, Takuya Higuchi¹, Jee-weon Jung², Skyler Seto¹, Tatiana Likhomanenko¹, Stephen Shum¹, Ahmed Hussien Abdelaziz¹, Shinji Watanabe², Barry-John Theobald¹

¹Apple, USA ²Carnegie Mellon University, USA

{zaldeneh, takuya.higuchi}@apple.com

Abstract

Self-supervised features are typically used in place of filter-bank features in speaker verification models. However, these models were originally designed to ingest filter-banks as inputs, and thus, training them on self-supervised features assumes that both feature types require the same amount of learning for the task. In this work, we observe that pre-trained self-supervised speech features inherently include information required for a downstream speaker verification task, and therefore, we can simplify the downstream model without sacrificing performance. To this end, we revisit the design of the downstream model for speaker verification using self-supervised features. We show that we can simplify the model to use 97.51% fewer parameters while achieving a 29.93% average improvement in performance on SUPERB. Consequently, we show that the simplified downstream model is more data efficient compared to the baseline—it achieves better performance with only 60% of the training data.

Index Terms: Self-supervised learning, representation learning, speaker recognition, speaker verification

1. Introduction

Self-Supervised Learning (SSL) for speech (e.g., wav2vec 2.0 [1], HuBERT [2], w2v-BERT [3], BEST-RQ [4]) enables learning powerful representations using a large amount of unlabeled data. Once trained on unlabeled data, SSL models can be fine-tuned on labeled data to achieve remarkable performance on target downstream tasks (e.g., automatic speech recognition, speaker recognition, language identification, emotion recognition) [5–10]. Fine-tuning a pre-trained SSL model for each downstream task, however, can be costly due to computation and memory constraints. A more appealing setup is using an SSL model as a *general-purpose feature extractor*, where the pre-trained model is frozen, and features extracted from this frozen model are used with smaller, task-dependent downstream models [11]. In this work, we study the role the downstream model plays when using general-purpose SSL features for speaker verification.

Downstream speaker verification models (such as x-vector [12] and ECAPA-TDNN [13]) used in prior works were originally designed to ingest filter-bank features as inputs, whereas state-of-the-art SSL models operate on raw waveforms and the features that the SSL models extract are learned in an end-to-end manner using a Transformer architecture [14–16]. An important difference here is that unlike features derived from filter-banks, features from SSL models capture long-form contextual information in their representation, and it has been shown that this information is useful for predictive speech processing tasks [17–19]. In addition, SSL representations capture

speaker information as it was shown that explicitly disentangling speaker information during SSL pre-training results in improved performance on content related tasks [20]. These findings suggest that SSL pre-training may have already done some of the learning required for downstream speaker-related tasks, whereas models trained on top of filter-bank features must still extract all of this information from the network inputs.

Given the contrast between filter-banks and SSL features, we explore the role the downstream model plays in speaker verification with SSL features. We first seek to understand the capability of several SSL models when conducting speaker verification without a downstream model (i.e., zero-shot capability) and use the findings from our analyses to revisit the design of the downstream speaker verification model. Specifically, we show that we can reduce the capacity of the downstream speaker verification model by 97.51% and still obtain a 29.93% average improvement in performance on the SUPERB [11] benchmark. Additionally, we show that the simplified downstream model is especially effective in limited training data scenarios, outperforming its baseline counterpart with only 60% of the training data.

2. Related Work

In this section, we discuss relevant prior works that looked at the intersection of SSL and speaker recognition. Specifically, we focus on SSL approaches that learn *generic* representations rather than approaches designed to extract specialized representations for speaker verification (e.g., [22–24]). We refer the reader to [25] for a thorough review on SSL representations.

Fan et al. [5] studied the effectiveness of a wav2vec 2.0 model on speaker verification and language identification. The authors visualized the features extracted from the model to show that the features capture speaker and language information. The authors then attached a fully-connected layer to the top of the model and ran experiments (both with and without fine-tuning the full model) to quantitatively demonstrate effectiveness of the pre-trained model on the downstream tasks. In contrast to [5], our work presents a comparative study that quantifies (not visualizes) the speaker information captured by several state-of-the-art SSL models (not just wav2vec 2.0). In addition, our work presents a study into the role the downstream model plays when performing speaker verification using SSL features.

Chen et al. [7] ran experiments to understand the components of pre-training that affect the performance of SSL models when fine-tuned on the speaker verification task. The authors used a weighted average of the hidden states, which is then passed to a downstream model for learning the speaker embeddings. Their results suggested that SSL models provide better features for speaker verification compared to those extracted from a model trained on the same dataset to perform

Table 1: Zero-shot (i.e., no downstream model) speaker verification performance of SSL models. “Params.” denotes the number of parameters in the model; “Data” denotes the data used for training the model; Δ denotes the relative (%) improvement that SSL features provide over using filter-bank (FBank) features. The values for “Params.” and “Data” columns are taken from [11].

Model	Params.	Data	LibriSpeech (in-domain)		VoxCeleb1 (out-of-domain)	
			EER (%) ↓	Δ (%) ↑	EER (%) ↓	Δ (%) ↑
FBank	—	—	7.2	0.0	40.4	0.0
HuBERT (base)	94.68M	LS 960 hr	4.4	38.9	32.0	20.8
HuBERT (large)	316.61M	LL 60k hr	3.1	56.9	31.7	21.5
wav2vec 2.0 (base)	95.04M	LS 960 hr	5.5	23.6	33.2	17.8
wav2vec 2.0 (large)	317.38M	LS 960 hr	2.6	63.9	30.7	24.0
wav2vec 2.0 (large)	317.38M	LL 60k hr	2.8	61.1	27.2	32.7
wav2vec 2.0 (large)	317.38M	VoxPopuli 100k hr	3.1	56.9	32.7	19.1
WavLM (base)	94.70M	LS 960 hr	4.7	34.7	31.3	22.5
WavLM (base+)	94.70M	Mix 94k hr*	4.0	44.4	31.3	22.5
WavLM (large)	316.62M	Mix 94k hr*	2.5	65.3	23.0	43.1
wav2vec	32.54M	LS 960 hr	5.3	26.4	30.7	24.0
vq-wav2vec	34.15M	LS 960 hr	11.4	−58.3	37.8	6.4
Modified CPC	1.84M	LL 60k hr	3.5	51.4	27.9	30.9

* The dataset contains GigaSpeech [21], which includes samples collected from YouTube.

Table 2: An unconstrained speaker verification setup yields a lower equal error rate (EER, %) on VoxCeleb1 compared to SUPERB. We fine-tune both the WavLM and ECAPA-TDNN models for the unconstrained setup; we include the VoxCeleb2 during fine-tuning and apply training-time augmentations.

Model	SUPERB	Unconstrained
WavLM+ECAPA-TDNN	2.03	0.39

automatic speech recognition. Chen et al. [8] examined the impact of different pre-training methods, SSL model sizes, and training datasets on the downstream performance. The results reaffirmed the benefit of SSL features over filter-bank features; and the importance of data augmentation for achieving state-of-the-art performance when training the downstream model. In contrast to [7] and [8], our experiments do not use a fixed downstream architecture; instead, we focus on re-designing the downstream model given the differences between filter-banks and SSL features.

Peng et al. [10] studied parameter-efficient fine-tuning to adapt pre-trained SSL models for speaker verification. They showed that using adaptors is better than fine-tuning the full SSL model in low-resource settings. Stafylakis et al. [9] proposed correlation pooling, an approach for aggregating frame-level SSL features across time to induce fixed-size utterance-level features. The authors showed that replacing statistics pooling with correlation pooling improved the performance of speaker verification when using SSL features. In contrast to [10] and [9], our work does not study adaptors (i.e., the addition of modules between the layers of the pre-trained SSL model)—we use SSL models as generic feature extractors and focus on the design of the full downstream model, not just the pooling mechanism.

3. Experiments

This work explores the design of the downstream speaker verification architecture given the contrasting nature of filter-banks and SSL features. We begin with an investigation into the speaker information that is captured by state-of-the-art SSL

methods (Section 3.1). We then run an ablation on the downstream architecture to understand the role different components play in speaker verification using SSL features (Section 3.2).

3.1. What Speaker Information is Captured by SSL?

Motivation. Results from SUPERB [11] suggest that SSL models capture speaker information even though these models were trained with frame-level objectives that resemble the objective of automatic speech recognition. The captured speaker information can be undesirable if the goal is to learn models that focus on content related tasks [20]. We identified two limitations in prior analyses that we address in our work. First, prior analyses focused only on reporting the speaker verification performances of SSL models on the VoxCeleb1 dataset [27]. However, VoxCeleb1 is an out-of-domain data for several SSL models that are evaluated in the literature—these SSL models were trained on audiobooks. Thus, it is unclear if the performance degradation comes from out-of-domain data, or from the task itself. Second, prior analyses used SSL features along with either x-vector or ECAPA-TDNN architectures for the verification task. However, there is evidence suggesting that the performance and the ranking is highly sensitive to the choice of the downstream model [28]. To this end, we seek to understand the capability of SSL models to do speaker verification *without a downstream model* (i.e., *zero-shot capability*) on both in-domain and out-of-domain data.

Approach. We extract frame-level features, $\mathbf{H}^l = \{\mathbf{h}_1^l, \mathbf{h}_2^l, \dots, \mathbf{h}_T^l\}$, from layer $l \in \{1, \dots, L\}$, and then induce a fixed-dimensional representation by aggregating all T frame-level features by computing the mean and standard deviation across time; where L is the number of layers in the SSL model and T is the number of frames in the representation. We use the cosine score to measure the similarity for trial pairs.

Setup. We assess the zero-shot speaker verification capability of SSL methods on LibriSpeech (in-domain) and VoxCeleb1 (out-of-domain). We use the Vox1-O evaluation protocol and the `test-clean` and `test-other` sets of LibriSpeech. We create a verification split for LibriSpeech by: (1) sampling all utterances that are $8 < x < 12$ seconds; (2) creating a

Table 3: We can simplify the downstream model for speaker verification when using SSL features. The “base” SSL models were trained on LibriSpeech 960 hr, and the “large” SSL models were trained on Libri-Light 60k hr. The downstream models were trained on the development set of VoxCeleb1. The equal error rates (EERs, %) are reported on the verification set of VoxCeleb1.

Downstream	Params.	FBank	EER (%) ↓						
			HuBERT base	wav2vec 2.0 base	WavLM base	HuBERT large	wav2vec 2.0 large	WavLM large	
A	x-vector [11,26]	5.7M	9.56	5.11	6.02	4.69	5.98	5.65	3.77
B	x-vector (our setup)	5.7M	9.95	4.41	5.12	4.97	5.28	6.04	4.04
C	Stafylakis et al. [9]	—	—	—	—	—	4.8	—	3.8
D.0	ECAPA-TDNN	8M	8.78	3.81	4.63	3.67	3.14	6.35	2.03
	→ w/o frame-level encoder								
D.1	→ w/ channel & context stats. pooling	725K	8.75	2.95	3.41	2.71	2.46	2.60	1.55
D.2	→ w/ attentive stats. pooling	725K	9.26	2.87	3.09	2.51	2.37	2.56	1.70
D.3	→ w/ stats. pooling	199K	10.74	3.11	3.19	2.83	2.34	2.63	1.56
$\Delta(D.0, D.3)$ (%) ↑		97.51	-22.32	18.37	31.10	22.89	25.48	58.58	23.15

list of all possible pairs; and (3) down-sampling the negative class samples such that we retain a 1:5 positive-to-negative ratio in the trial list. We follow above process separately for the `test-clean` and `test-other` sets of LibriSpeech and then merge the two to obtain a list with 73 speakers and 1908 pairs.

Results. The zero-shot speaker verification capability for several SSL models is shown in Table 1. We use zero-shot performance on filter-bank features as our baseline and discuss our findings below.

Do SSL features highlight speaker characteristics beyond what filter-bank features highlight? Features from all pre-trained SSL models (except vq-wav2vec on in-domain test set) provide improvements over filter-bank features. WavLM (large) features improve the baseline performance by 65.3% and 43.1% on LibriSpeech and VoxCeleb1, respectively. Even though vq-wav2vec features degrade the baseline performance by 58.3% on LibriSpeech, it outperforms the baseline by 6.4% on VoxCeleb1. This result shows that SSL features can provide an improvement over filter-bank features in the zero-shot setting. Furthermore, the results reaffirm that the augmentation strategy used in WavLM training effectively incorporates more speaker information in the representation.

Does the domain mis-match impact the relative improvements from SSL features compared to filter-banks? SUPERB evaluates the quality of speaker verification models on the VoxCeleb1 dataset. However, several of the SSL models are trained on audiobook data. Our results show that we obtain higher relative improvements on LibriSpeech compared to VoxCeleb1. We compute Spearman’s ρ between the relative improvements on the two datasets and obtain 0.66 ($p = 0.019$). This result suggests that, while there is a strong correlation between the two domains, the rankings of the SSL models are not identical and they depend on the domain of the downstream data (even for the same task).

Are bigger SSL models better at capturing speaker information compared to smaller models? Increasing the capacity of wav2vec 2.0 from 95.04M to 317.38M while using the same 960 hours of training data increases the relative improvements from 23.6% to 63.9% and from 17.8% to 24.0% on LibriSpeech and VoxCeleb1, respectively. This finding is also true for WavLM, where increasing the capacity from 94.70M to 316.62M while using the same 94k hours of data increases the relative improvements from 44.4% to 65.3% and from 22.5% to 43.1% on LibriSpeech and VoxCeleb1, respectively. Our results suggest that increasing the model size provides the model with capacity to capture more information, including speaker information.

Is the prior from SSL model architecture adequate for capturing speaker characteristics? In the vision domain, Ulyanov et al. [29] showed that the structure of convolutional networks provides a strong prior for learning. We ask whether or not SSL speech model architecture alone (i.e., no learning) provides appropriate priors for capturing speaker information from raw waveforms. We find that random SSL models, on average, drop the performance by 131.6% for the LibriSpeech setup compared to baseline and drop the performance by 8.4% for the VoxCeleb1 setup, suggesting that the architecture alone is insufficient for capturing speaker characteristics.

3.2. Can we Simplify the Downstream Model?

Motivation. The results from Section 3.1 reaffirm that SSL models capture speaker information beyond what is captured with filter-bank features. This finding suggests that we may need a different model for downstream task because the original downstream models were designed for filter-banks. We re-visit the ECAPA-TDNN architecture [13], noting the impact of its components on the downstream speaker verification task when used with SSL features, and propose a simple yet effective speaker verification architecture suitable for SSL features.

Downstream Model. The ECAPA-TDNN model introduced enhancements to the Time Delay Neural Network (TDNN) architecture and the attentive statistics pooling layer to address the local spatial modeling limitations of convolutional networks. We study the utility of the frame-level encoder and the pooling mechanism when used with SSL features.

Frame-level Encoder. The frame-level encoder includes Res2Net blocks, Squeeze-Excitation (SE) blocks, and Multi-layer Feature Aggregation (MFA). The SE blocks were introduced to the the ECAPA-TDNN architecture to “rescale the frame-level features given global properties of the recording” [13]. MFA was introduced so the model can exploit information from multiple layers before pooling.

Channel- and Context-dependent Statistics Pooling. The ECAPA-TDNN architecture extends the temporal attention statistics from [30] to also depend on the channel dimension. This change allows the model to attend to different time frames for different features. The attention module calculates a scalar score for each frame given the channel: $z_{t,c} = f_c(\mathbf{h}_t)$, where \mathbf{h}_t are the hidden states from the previous layer at time t ; and f_c is the channel-dependent non-linear transformation. The scalar scores are normalized across time per channel. The temporal context of attentive pooling is also extended by concatenating the global context features.

The architectural enhancements that were introduced in

downstream models can be unnecessary when these models are used with SSL features. State-of-the-art SSL models use the transformer model in their architecture. The output at each frame from a transformer already captures the full context of the utterance and the multi-head attention mechanism enables each head to focus on different aspects of the utterance. To this end, we ablate the ECAPA-TDNN architecture’s structure to study the enhancements’ impact on the overall performance when using SSL features.

Setup. We follow the SUPERB setup for speaker verification using SSL features. We pass a waveform through a frozen SSL model and take the weighted sum of the hidden states from each layer to produce the output sequence: $\mathbf{o}_t = \sum_{l=1}^L w^l \cdot \mathbf{h}_t^l$, where \mathbf{h}_t^l are the hidden states from layer l at time t ; and w^l is the normalized scalar weight for layer l . The output sequence is then fed into a downstream ECAPA-TDNN model to produce the embeddings: $\mathbf{e} = \text{ECAPA-TDNN}(\mathbf{o}_1, \mathbf{o}_2, \dots, \mathbf{o}_T)$. The summation weights and the downstream model parameters are trained to classify speakers with an additive margin softmax [31] loss, where a scale of 30 and a margin of 0.4 are used. We use 512 channels for the convolutional frame layers in ECAPA-TDNN, and use the following hyper-parameters: optim = AdamW [32]; lr = $5.0e^{-5}$; batch size = 40. We train the models for 100 000 steps and create a checkpoint every 5000 steps. We report the best performance from all checkpoints in accordance with SUPERB.

Note. Our work aims to study how well SSL models capture speaker information and how to extract this information effectively from frozen SSL models. Our goal is *not* to achieve the best possible speaker verification performance—a goal achievable by fine-tuning the SSL model on the downstream task. We run an experiment to highlight the difference in performance on VoxCeleb1 when fixing the SSL model (WavLM) according to SUPERB and when jointly fine-tuning both WavLM and downstream models. Table 2 shows that we can achieve an EER of 0.39% when the SSL and downstream models are fine-tuned simultaneously, highlighting the performance differences due to the SUPERB setup. Despite this difference in performance, we use the SUPERB setup because it is a widely used setup for benchmarking the quality of SSL features, and our focus is on limited training data scenarios for the downstream task.

Results. The results of our downstream ablation are reported in Table 3 and we discuss our findings below.

First, we replicate the x-vector setups reported in prior works to ensure a fair comparison. Our x-vector setups provide, on average, a 2.32% relative improvement in performance compared to x-vector performance reported in [11, 26]. This result establishes that our setup is competitive and reflects state-of-the-art performance on SUPERB. Replacing the x-vector with the ECAPA-TDNN model improves the performance by 11.76% given the filter-bank features and 22.41% on average for the SSL setups, reaffirming the utility of the structural enhancements employed in ECAPA-TDNN.

We remove the frame-level encoder from the architecture and evaluate three different pooling mechanisms: channel- and context-dependent statistics pooling proposed with ECAPA-TDNN model, attentive statistics pooling from [30], and statistics pooling from [12]. Removing the frame-level encoder reduces the model’s capability to use features from multiple layers and some of its capability to exploit contextual information. ECAPA-TDNN’s channel-attentive pooling mechanism without the frame-level encoder improves filter-bank performance by 0.34% but improves the average SSL performance by 29.91%.




Figure 1: The simplified downstream model (D.3 from Table 3) performs better with less data compared to full downstream model (D.0 from Table 3). The equal error rate (EER, %) on VoxCeleb1 is reported under three data conditions: 20%, 60%, and 100% of the training speakers.

This result suggests that SSL models do not require the same frame-level processing that filter-banks require for extracting speaker information.

Replacing channel- and context-dependent statistics pooling with attentive statistics pooling from [30] drops the performance by 5.83% for the filter-bank model but improves the average performance by 2.5% for the SSL models. This result suggests that channel-attention is important when using filter-bank features but not when using SSL features. Finally, replacing the attentive statistics pooling with non-weighted statistics pooling reduces filter-bank performance by 15.98% and reduces the average SSL performance by 2.93%, suggesting that the attention mechanism is less important for SSL models compared to filter-banks for speaker verification.

Data Efficiency. We further create a random subset of the VoxCeleb1 containing only 20% of the training speakers to study the data efficiency of the simplified downstream model similar to [33]. Then, we gradually increase the number of training speakers by adding randomly selected speakers until we cover 100% of the data. We evaluate the downstream models using three features: filter-banks, HuBERT (base), and HuBERT (large). We focus on HuBERT in our analysis because it is widely used and it is used for pre-training WavLM. The results in Figure 1 show that the simplified downstream model (D.3 from Table 3) is more data efficient; the model achieves better performance when using only 60% of the data for both HuBERT setups.

4. Conclusion

We observed that state-of-the-art SSL features used in downstream speaker verification models were designed originally for filter-bank features. We hypothesized that downstream models can be simplified because SSL models have potentially done some of the learning required for the task. Our results suggest that, although we can’t completely remove the downstream model when using SSL features, we can simplify the model to use 97.51% fewer parameters and obtain a 29.93% average improvement in performance compared to the original model on SUPERB. We also showed that the simplified downstream model requires less training data—the model uses 60% of the original data to achieve the same or better performance compared to the full model.

5. References

- [1] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in neural information processing systems*, 2020.
- [2] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2021.
- [3] Y.-A. Chung, Y. Zhang, W. Han, C.-C. Chiu, J. Qin, R. Pang, and Y. Wu, "W2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training," in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2021.
- [4] C.-C. Chiu, J. Qin, Y. Zhang, J. Yu, and Y. Wu, "Self-supervised learning with random-projection quantizer for speech recognition," in *International Conference on Machine Learning (ICML)*, 2022.
- [5] Z. Fan, M. Li, S. Zhou, and B. Xu, "Exploring wav2vec 2.0 on Speaker Verification and Language Identification," in *Proc. Interspeech*, 2021.
- [6] Y. Wang, A. Boumadane, and A. Heba, "A fine-tuned wav2vec 2.0/hubert benchmark for speech emotion recognition, speaker verification and spoken language understanding," *arXiv preprint arXiv:2111.02735*, 2021.
- [7] S. Chen, Y. Wu, C. Wang, S. Liu, Z. Chen, P. Wang, G. Liu, J. Li, J. Wu, X. Yu, and F. Wei, "Why does Self-Supervised Learning for Speech Recognition Benefit Speaker Recognition?" in *Proc. Interspeech*, 2022.
- [8] Z. Chen, S. Chen, Y. Wu, Y. Qian, C. Wang, S. Liu, Y. Qian, and M. Zeng, "Large-scale self-supervised speech representation learning for automatic speaker verification," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.
- [9] T. Stafylakis, L. Mošner, S. Kakouros, O. Plchot, L. Burget, and J. Černocký, "Extracting speaker and emotion information from self-supervised speech models via channel-wise correlations," in *IEEE Spoken Language Technology Workshop (SLT)*, 2023.
- [10] J. Peng, T. Stafylakis, R. Gu, O. Plchot, L. Mošner, L. Burget, and J. Černocký, "Parameter-efficient transfer learning of pre-trained transformer models for speaker verification using adapters," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
- [11] S.-w. Yang, P.-H. Chi, Y.-S. Chuang, C.-I. J. Lai, K. Lakhotia, Y. Y. Lin, A. T. Liu, J. Shi, X. Chang, G.-T. Lin *et al.*, "Superb: Speech processing universal performance benchmark," in *Proc. Interspeech*, 2021.
- [12] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust dnn embeddings for speaker recognition," in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 2018.
- [13] B. Desplanques, J. Thienpondt, and K. Demuynck, "ECAPA-TDNN: Emphasized Channel Attention, Propagation and Aggregation in TDNN Based Speaker Verification," in *Proc. Interspeech*, 2020.
- [14] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, 2017.
- [15] J. D. M.-W. C. Kenton and L. K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of NAACL-HLT*, 2019.
- [16] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.
- [17] R. Masumura, T. Tanaka, T. Moriya, Y. Shinohara, T. Oba, and Y. Aono, "Large context end-to-end automatic speech recognition via extension of hierarchical recurrent encoder-decoder models," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- [18] S. Shon, F. Wu, K. Kim, P. Sridhar, K. Livescu, and S. Watanabe, "Context-aware fine-tuning of self-supervised speech models," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
- [19] Y. Zhao, T. Zhou, Z. Chen, and J. Wu, "Improving deep cnn networks with long temporal context for text-independent speaker verification," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- [20] K. Qian, Y. Zhang, H. Gao, J. Ni, C.-I. Lai, D. Cox, M. Hasegawa-Johnson, and S. Chang, "Contentvec: An improved self-supervised speech representation by disentangling speakers," in *International Conference on Machine Learning*, 2022.
- [21] G. Chen, S. Chai, G. Wang, J. Du, W.-Q. Zhang, C. Weng, D. Su, D. Povey, J. Trmal, J. Zhang *et al.*, "Gigaspeech: An evolving, multi-domain asr corpus with 10,000 hours of transcribed audio," *arXiv preprint arXiv:2106.06909*, 2021.
- [22] J. Cho, R. Pappagari, P. Żelasko, L. Moro-Velazquez, J. Villalba, and N. Dehak, "Non-contrastive self-supervised learning of utterance-level speech representations," *arXiv preprint arXiv:2208.05413*, 2022.
- [23] J. Cho, P. Żelasko, J. Villalba, S. Watanabe, and N. Dehak, "Learning Speaker Embedding from Text-to-Speech," in *Proc. Interspeech*, 2020.
- [24] H. Zhang, Y. Zou, and H. Wang, "Contrastive self-supervised learning for text-independent speaker verification," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021.
- [25] A. Mohamed, H.-y. Lee, L. Borgholt, J. D. Havtorn, J. Edin, C. Igel, K. Kirchhoff, S.-W. Li, K. Livescu, L. Maaløe *et al.*, "Self-supervised speech representation learning: A review," *IEEE Journal of Selected Topics in Signal Processing*, 2022.
- [26] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, 2022.
- [27] A. Nagrani, J. S. Chung, W. Xie, and A. Zisserman, "Voxceleb: Large-scale speaker verification in the wild," *Computer Speech & Language*, 2020.
- [28] S. Zaiem, Y. Kemiche, T. Parcollet, S. ESSID, and M. Ravanelli, "Speech Self-Supervised Representation Benchmarking: Are We Doing it Right?" in *Proc. Interspeech*, 2023.
- [29] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep image prior," in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2018.
- [30] K. Okabe, T. Koshinaka, and K. Shinoda, "Attentive Statistics Pooling for Deep Speaker Embedding," in *Proc. Interspeech*, 2018.
- [31] F. Wang, J. Cheng, W. Liu, and H. Liu, "Additive margin softmax for face verification," *IEEE Signal Processing Letters*, 2018.
- [32] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *International Conference on Learning Representations*, 2018.
- [33] H.-S. Heo, J.-w. Jung, J. Kang, Y. Kwon, Y. J. Kim, B.-J. Lee, and J. S. Chung, "Curriculum learning for self-supervised speaker verification," in *Proc. Interspeech*, 2023.