



Towards Supervised Performance on Speaker Verification with Self-Supervised Learning by Leveraging Large-Scale ASR Models

Victor Miara, Theo Lepage, Reda Dehak

EPITA Research Laboratory (LRE), France

victor.miara@epita.fr, theo.lepage@epita.fr, reda.dehak@epita.fr

Abstract

Recent advancements in Self-Supervised Learning (SSL) have shown promising results in Speaker Verification (SV). However, narrowing the performance gap with supervised systems remains an ongoing challenge. Several studies have observed that speech representations from large-scale ASR models contain valuable speaker information. This work explores the limitations of fine-tuning these models for SV using an SSL contrastive objective in an end-to-end approach. Then, we propose a framework to learn speaker representations in an SSL context by fine-tuning a pre-trained WavLM with a supervised loss using pseudo-labels. Initial pseudo-labels are derived from an SSL DINO-based model and are iteratively refined by clustering the model embeddings. Our method achieves 0.99% EER on VoxCeleb1-O, establishing the new state-of-the-art on self-supervised SV. As this performance is close to our supervised baseline of 0.94% EER, this contribution is a step towards supervised performance on SV with SSL.

Index Terms: Speaker Recognition, Self-Supervised Learning, Speech Representations, ASR

1. Introduction

Speaker Recognition (SR) is the process of recognizing the identity of the person speaking in an audio speech utterance. A typical application is Speaker Verification (SV), which aims to determine whether two audio utterances originate from the same speaker. The objective of SR systems is to learn speaker representations by maximizing inter-speaker distances while minimizing intra-speaker variances and being robust to extrinsic variabilities. With the emergence of deep learning, conventional machine learning methods such as i-vectors [1] have been outperformed by DNN models [2, 3] such as the d-vectors [4], the x-vectors [5], and more recently ECAPA-TDNN [6]. These methods are trained to match an input utterance to the corresponding speaker identity in a supervised manner. However, this training scheme represents a notable constraint due to the scarcity and expense of getting annotated training samples determining the performance of DNN models.

Self-Supervised Learning (SSL) methods have emerged as a solution to this problem by learning meaningful representation directly from the input data. SSL has been shown to enhance the scalability potential of models by leveraging the abundance of unlabeled data available. When applied to SV, most frameworks rely on contrastive learning, such as AP+AAAT [7], SimCLR [8] and MoCo [9]. They assume that two utterances, from distinct

audio files sampled randomly in the training set do not belong to the same speaker. Based on this assumption, they aim to maximize the similarity in the embedding space between positive pairs and minimize the similarity between negative pairs (constructed by sampling utterances from different audio files). Other methods have been explored and successfully applied to SV [10, 11, 12], notably techniques relying on knowledge distillation such as DINO which represents the state-of-the-art [13, 14, 15]. DINO is based on a self-distillation training process that iteratively maximizes the cross-entropy between the output of a student and a teacher branch on multiple different views of the same utterance. All these methods rely on extensive data augmentation to avoid encoding channel characteristics as positive segments are extracted from the same audio files.

The field of Automatic Speech Recognition (ASR) has witnessed significant advancements by pre-training large Transformer-based models in a self-supervised way on speech prediction with masked units, inspired by NLP methods [16]. wav2vec [17, 18], HuBERT [19], and WavLM [20] have been major milestones in ASR. Training these large models from scratch requires immense computational capacity. Nonetheless, due to their massive number of trainable parameters and extensive pre-training, they have been shown to generalize well to SR when fine-tuned [21, 22, 23]. WavLM stands out as its training objective compels the model to encode information relevant to other speech-processing tasks, enabling it to achieve state-of-the-art in SV when fine-tuned with an ECAPA-TDNN. WavLM MHFA [24] replaces the ECAPA-TDNN with a lightweight attention-based back-end while reducing the number of parameters. Therefore, we observe a paradigm shift from relying on task-specific models trained on small datasets to harnessing speech representations for SR.

We propose a strategy to learn speaker representations by fine-tuning a pre-trained WavLM with the MHFA back-end using a self-supervised strategy. Because of the impracticability of using an SSL contrastive objective in an end-to-end fashion, we generate pseudo-labels from DINO to fine-tune the ASR model with a supervised loss. Compared to [20, 24], our method is fully self-supervised as our fine-tuning does not rely on speaker identity labels from the train set. While other equivalent SSL techniques [25, 26] use standard SR models, we use an ASR model as the backbone. Thus, this contribution is the first to apply SSL to a large-scale pre-trained ASR model for SV. Our approach achieves 0.99% EER on VoxCeleb1-O, representing the new state-of-the-art on this task with SSL, and narrows the gap with supervised performance.

Our complete framework is described in Section 2.4. Section 3 describes our experimental setup. The results of our system on SV are presented and compared to other methods from the literature in Section 4. Section 5 concludes the article.

The code associated with this article is publicly available at https://github.com/theolepage/wavlm_ssl_sv.

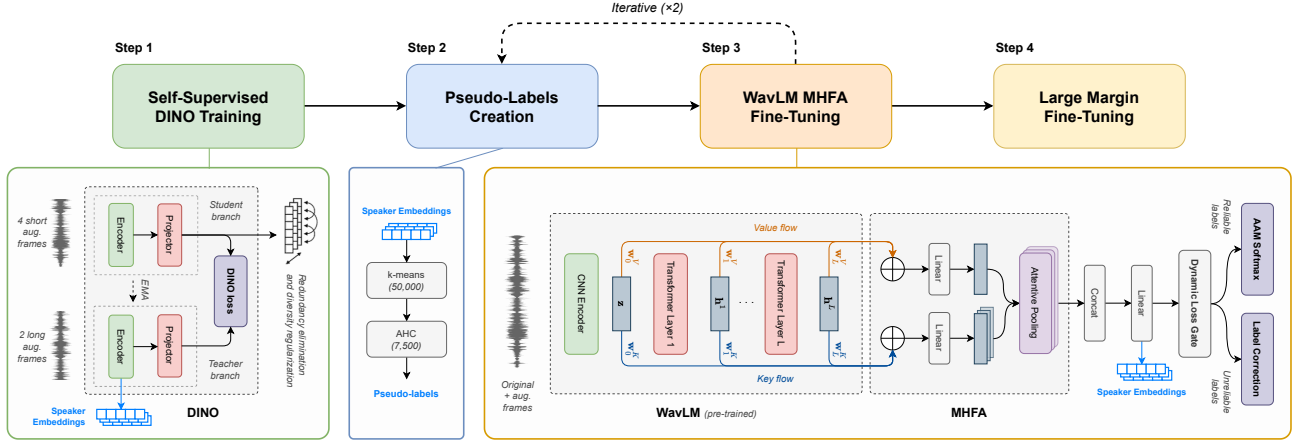


Figure 1: Our complete training framework to leverage speech representations for SV using a self-supervised strategy.

2. Methods

2.1. Self-supervised learning for speaker verification

2.1.1. Contrastive learning (SimCLR)

SSL contrastive learning assumes that two randomly sampled utterances belong to different speakers. The learning task aims at maximizing the similarity within positive pairs, derived from the same utterances, while maximizing the distance between negative pairs, sampled from the mini-batch. Following, SimCLR [27], we rely on a siamese architecture to produce a pair of embeddings for a given unlabeled utterance and apply extensive data-augmentations which is fundamental to produce robust representations. Let $i \in I \equiv \{1 \dots 2N\}$ be the index of an utterance from the mini-batch of size N , $j(i)$ be the index of the other augmented frame from the same audio and $A(i) \equiv I \setminus \{i\}$. *NT-Xent* loss is defined as

$$\mathcal{L}_{\text{NT-Xent}} = -\frac{1}{2N} \sum_{i \in I} \log \frac{\ell(\mathbf{z}_i, \mathbf{z}_{j(i)})}{\sum_{a \in A(i)} \ell(\mathbf{z}_i, \mathbf{z}_a)}, \quad (1)$$

where \mathbf{z}_i is an embedding and $\ell(\mathbf{u}, \mathbf{v}) = e^{\cos(\theta_{\mathbf{u}, \mathbf{v}})/\tau}$.

2.1.2. Knowledge distillation (DINO)

DINO [28] is a self-distillation framework where a student network is trained to match the output of a teacher network. From an utterance x , a set of different views V are generated with data-augmentation resulting in four *local* segments and two *global* segments (x_1^g and x_2^g). All views are fed through the student branch, while only global views are fed through the teacher network to encourage "local-to-global" correspondences. Both networks share the same architecture, composed of an encoder, where speaker embeddings are extracted, and a projector. The student and teacher output a K dimensional feature (P_s and P_t , respectively) that is normalized with a temperature softmax over the feature dimension. The objective function minimizes the cross-entropy between the two probability distributions and is defined as

$$\mathcal{L}_{\text{DINO}} = \sum_{x \in \{x_1^g, x_2^g\}} \sum_{\substack{x' \in V \\ x' \neq x}} H(P_t(x), P_s(x')), \quad (2)$$

where $H(a, b) = -a \log b$. During training, the teacher network is updated with an Exponential Moving Average (EMA)

of the student weights. To avoid collapse (i.e. trivial solutions), *sharpening* and *centering* are applied to the teacher outputs. Following RDINO [13], we implement redundancy elimination and diversity regularizations.

2.2. WavLM-based speaker recognition

WavLM [20] is a Transformer-based ASR model pre-trained in a self-supervised way that also encodes non-ASR information such as speaker identity. During the pre-training phase, a multi-layer convolutional feature encoder takes as input the sequence $\{\mathbf{x}_t\}_{t=1}^T$ of T time windows from a raw audio sample to produce $\{\mathbf{z}_t\}_{t=1}^T$. These representations are simulated noisy/overlapped with masks and fed into the Transformer encoder which outputs a series of hidden states $\{\mathbf{h}^l\}_{l=1}^L$, with L the number of Transformer layers. Furthermore, the model uses gated relative position bias, enhancing its ability to focus on relevant speech characteristics. WavLM is trained on a masked speech denoising and prediction task which implicitly models speaker-related information as the objective is to predict the pseudo-labels of the original speech on masked regions. The pseudo-labels are generated by first clustering the Mel-Frequency Cepstral Coefficients (MFCCs) of the training data, and then the latent representations.

2.2.1. WavLM with Multi-Head Factorized Attention

Multi-Head Factorized Attention (MHFA) [24] back-end consists of aggregating layer-wise outputs from WavLM's transformer layers into an attentive pooling mechanism. This method relies on two components: (1) L2 regularization between the current weights and the initial pre-trained weights, preventing overfitting due to a large number of parameters; (2) layer-wise learning rate decay, following [29]. As there is a progressive abstraction of information throughout the Transformer layers [20], this latter technique allows more flexible weight modifications in higher layers to refine ASR capabilities, while ensuring lower layers preserve speaker-related information.

2.3. Unfeasibility of end-to-end self-supervised fine-tuning

Our experiments showed that fine-tuning WavLM directly with a self-supervised contrastive loss leads to a sub-optimal solution. Using the DINO framework was not considered due to memory limitations. Figure 2 reveals weight changes in Trans-

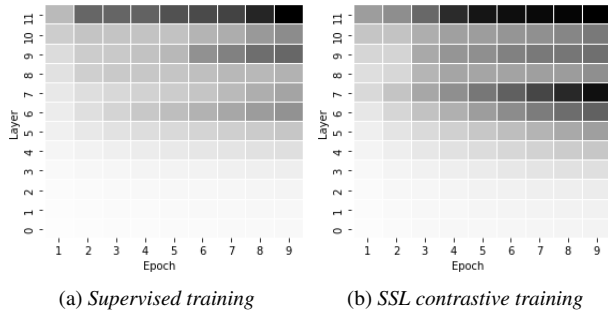


Figure 2: l_2 distance between the pre-trained and fine-tuned weights of the WavLM at different layers and epochs.

former layers during supervised (AAM-Softmax loss) and self-supervised (NT-Xent loss) training per epoch. We notice that optimizing an SSL contrastive loss leads to greater changes, especially in early layers which contain more speaker information. Considering that SSL contrastive methods are prone to learn channel-related information, as positive pairs are sampled from the same utterances, we assume that these weight changes reflect a focus on features extrinsic to speaker identity. This misalignment between the SSL contrastive loss and the downstream task, inherent to the self-supervision and amplified by WavLM learning capabilities, leads to poorer performance.

2.4. Self-supervised fine-tuning by leveraging pseudo-labels

Based on this impracticability, we opt to fine-tune WavLM MHFA with pseudo-labels, initially extracted from a self-supervised model, instead of training the whole framework in an end-to-end self-supervised way. Our training framework is depicted in Figure 1 and detailed in this section.

2.4.1. Fine-tuning WavLM MHFA with refined pseudo-labels

After the DINO-based model training (**Step 1**), initial pseudo-labels are generated by applying a clustering algorithm on training samples’ embeddings (**Step 2**). We use k-means with 50,000 clusters followed by Agglomerative Hierarchical Clustering (AHC) with 7,500 clusters, similar to [14]. Each cluster is treated as a distinct speaker class, resulting in a list of sample-label pairs used to fine-tune WavLM MHFA (**Step 3**). Then, we iteratively generate refined pseudo-labels from the WavLM MHFA of the last iteration (**Step 2** and **Step 3**). This method allows training our framework in a self-supervised way, without ever needing the original training labels. After two iterations, we increase the input audio length and the AAM-Softmax margin to perform the Large Margin Fine-Tuning (LMFT) (**Step 4**).

2.4.2. Handling unreliable pseudo-labels

Similarly to [30], samples with incorrect pseudo-labels typically have higher loss values after a few epochs during the training. This can be attributed to the model’s fast adaptation and generalization to samples with reliable pseudo-labels, representing the majority of the training data. Therefore, we apply Dynamic Loss-Gate (DLG), which ignores unreliable samples having a loss higher than a defined threshold during the training. Following [25], we implement Label Correction (LC) to avoid discarding all unreliable samples. This involves maximizing the similarity between the output distributions of clean and augmented versions of utterances, with a cross-entropy loss.

3. Experimental setup

3.1. Datasets

The fine-tuning of the WavLM MHFA is performed on the VoxCeleb2 [31] *dev* set, which contains 1,092,009 utterances from 5,994 speakers. We do not rely on the original speaker identity labels at any stage of the training. We apply online data augmentation with the MUSAN [32] noise and the Room Impulse Response and Noise Database [33] reverberation corpora, similar to [11]. All systems are evaluated on VoxCeleb1 [34] test sets comprising 4,874 utterances from 40 speakers.

3.2. Implementation details

3.2.1. DINO

DINO is trained from scratch in a self-supervised way, using the setup described in [13]. The encoder is an ECAPA-TDNN [6] which outputs a 512-dimensional embedding.

3.2.2. WavLM MHFA

The pre-trained WavLM base+¹ [20] and the MHFA backend [24] have respectively 94M and 2M learnable parameters. WavLM base+ comprises a CNN encoder, acting as a feature extractor, and 12 Transformer layers. The dimension of the Transformer’s output is 768, and the dimension of the MHFA backend’s output, used for speaker verification, is 256. WavLM MHFA is trained on $2 \times$ NVIDIA A100 80 GB GPUs for 15 epochs. The mini-batch size is set to 120, the optimizer is Adam, and the loss is AAM-Softmax with a margin of 0.2 and a scale of 30. The learning rate is decreased by 5% at each epoch. The input raw waveform duration is 3 s. Dynamic Loss-Gate (DLG) is not applied during the first 5 epochs. Onwards, the losses for each sample are recorded to dynamically compute the loss gate value τ_1 . As described in [25], a Gaussian Mixture Model (GMM) with two components is used to fit the loss distribution, and the threshold τ_1 is set at the intersection of the two distributions. Label Correction (LC) is applied solely to unreliable samples, which have a loss value higher than τ_1 and at least one output class probability greater than τ_2 . LC is activated 3 epochs after enabling DLG, allowing the loss distribution to have two distinct components, reliable and unreliable samples. Similar to [25], the threshold τ_2 is set to 0.5 and the sharpen parameter ϵ_c to 0.1. For the Large Margin Fine-Tuning (LMFT), we continue the training for 2 epochs, and we set the input audio duration to 5 s and the margin to 0.5.

3.3. Evaluation protocol

To perform SV, we extract embeddings from 15 evenly spaced frames of 3 s for each utterance. The trial score is the average of the cosine similarity for all combinations of l_2 -normalized embeddings pairs. We report the performance in terms of Equal Error Rate (EER) and minimum Detection Cost Function (minDCF) with $P_{target} = 0.01$, $C_{miss} = 1$ and $C_{fa} = 1$.

4. Results and discussions

4.1. Fine-tuning WavLM with an SSL contrastive loss

In Table 1, we present the results obtained on SV when fine-tuning WavLM with an SSL contrastive loss via the SimCLR framework (NT-Xent loss). This experiment used an MLP on

¹<https://github.com/microsoft/unilm>

Table 3: Evaluation of different self-supervised SV methods on VoxCeleb1 trials (Original, Extended, Hard). The results for the top rows are drawn from the literature. "Ours" represents the system described in this article (Section 2.4). "Supervised" is our baseline which corresponds to our reproduction of WavLM MHFA [24]. The number of iterations includes the first pseudo-labeling from DINO.

Method	# of iterations	VoxCeleb1-O		VoxCeleb1-E		VoxCeleb1-H	
		EER (%)	minDCF _{0.01}	EER (%)	minDCF _{0.01}	EER (%)	minDCF _{0.01}
JHU [26]	4	1.89	-	-	-	-	-
DKU [35]	4	1.81	-	-	-	-	-
SNU [36]	4	1.66	-	-	-	-	-
LGL [30]	5	1.66	-	2.18	-	3.76	-
DLG-LC [25]	5	1.47	-	1.78	-	3.19	-
Ours	3	0.99	0.0905	1.21	0.1263	2.35	0.2214
Supervised	-	0.94	0.1179	0.93	0.1066	1.94	0.1919

Table 1: Fine-tuning WavLM with an SSL contrastive loss. Positive pairs are either sampled from the same audio files (SimCLR) or distinct audio files (supervised). EER and minDCF are evaluated on SV (VoxCeleb1-O).

Positive pairs sampling	EER (%)	minDCF _{0.01}
Same audio files (SimCLR)	15.13	0.9586
Different audio files	10.81	0.9377

top of the WavLM and trained the system for 50 epochs on VoxCeleb1. Sampling positive pairs from the same utterances, following contrastive SSL frameworks, results in an EER that is not on par with equivalent methods. However, using labels to sample positive pairs from different utterances improves downstream performance significantly, which reinforces our assumption that the strong reliance of contrastive SSL on channel characteristics prevents the WavLM from focusing on speaker information during fine-tuning, as discussed in Section 2.3.

4.2. Effect of the different components of our framework

In Table 2, we report the incremental improvement of the different components of our framework on speaker verification (VoxCeleb1-O). Training the WavLM MHFA on DINO's pseudo-labels results in a 52.5% relative reduction of the EER. This shows the benefit of leveraging speech representations by fine-tuning the ASR model with a supervised objective on the pseudo-labels obtained in a self-supervised way. The use of MHFA is motivated by the fact that speaker information is not uniformly distributed among the Transformer layers. In addition, MHFA lowers training time by limiting the number of trainable parameters compared to other back-ends. Dynamic Loss-Gate (DLG) and Label Correction (LC) improve the downstream performance to 1.22% EER, showing the importance of handling unreliable pseudo-labels. The Iterative Clustering (IC), generating refined pseudo-labels, allows reaching 1.01% EER after the second iteration. We opt to discard the third iteration as it results in equivalent performance. Finally, we perform the Large Margin Fine-Tuning (LMFT) step and achieve 0.99% EER. To assess the quality of our pseudo-labels, we rely on the Adjusted Rand Index (ARI) and the Normalized Mutual Information (NMI). Initial pseudo-labels (DINO) results in 0.81 ARI and 0.95 NMI while pseudo-labels obtained from our final system results in 0.90 ARI and 0.98 NMI. This improvement proves the efficiency of our system.

Table 2: Incremental study of the components of our framework. Grayed rows are discarded. LMFT represents our final model. EER and minDCF are evaluated on SV (VoxCeleb1-O).

Method	EER (%)	minDCF _{0.01}
DINO	3.16	0.2233
+ WavLM MHFA	1.50	0.1378
+ DLG	1.27	0.1401
+ LC	1.22	0.1531
+ IC (iter 1)	1.17	0.1351
+ IC (iter 2)	1.01	0.1399
+ IC (iter 3)	1.08	0.1340
+ LMFT	0.99	0.0905

4.3. Comparison to other speaker verification methods

An evaluation of different self-supervised SV methods, including our system, is presented in Table 3. The final performance of our self-supervised framework is 0.99% EER and 0.0640 minDCF on VoxCeleb1-O. Additionally, we achieve 1.21% EER and 2.35% EER on VoxCeleb1 *Extended* and *Hard* test sets, respectively. This result almost reaches the supervised baseline set at 0.94% EER and 0.0655 minDCF on VoxCeleb1-O. Thus, our system significantly narrows the gap between supervised and self-supervised performance. Furthermore, we achieve state-of-the-art performance on self-supervised SV as we outperform other equivalent methods from the literature [26, 35, 36, 30, 25] on all VoxCeleb1 trials. This shows the robustness of our system as we only rely on 3 iterations of iterative clustering and do not apply model fusion or calibration.

5. Conclusions

In this article, we propose a fully self-supervised framework to leverage speech representations for SV. We show that fine-tuning WavLM with an SSL contrastive objective function does not capture valuable speaker information. Thus, we rely on pseudo-labels generated through self-supervision to fine-tune the pre-trained ASR model with the MHFA back-end. Our experiments on the VoxCeleb1 test set show that our method achieves state-of-the-art performance on self-supervised SV with a result close to the supervised baseline. This contribution is a step towards supervised performance with self-supervised learning on this task.

6. Acknowledgements

This work was performed using HPC resources from GENCI-IDRIS (Grant 2023-AD011014623) and has been partially funded by the French National Research Agency (project AP-ATE - ANR-22-CE39-0016-05).

7. References

- [1] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-End Factor Analysis for Speaker Verification," *IEEE TASLP*, 2011.
- [2] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep Speaker Recognition," in *Interspeech*, 2018.
- [3] 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.
- [4] E. Variani, X. Lei, E. McDermott, I. L. Moreno, and J. Gonzalez-Dominguez, "Deep Neural Networks for Small Footprint Text-Dependent Speaker Verification," in *ICASSP*, 2014.
- [5] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-Vectors: Robust DNN Embeddings for Speaker Recognition," in *ICASSP*, 2018.
- [6] B. Desplanques, J. Thienpondt, and K. Demuyne, "ECAPA-TDNN: Emphasized Channel Attention, Propagation and Aggregation in TDNN Based Speaker Verification," in *Interspeech*, 2020.
- [7] J. Huh, H. S. Heo, J. Kang, S. Watanabe, and J. S. Chung, "Augmentation adversarial training for unsupervised speaker recognition," in *NeurIPS Workshop*, 2020.
- [8] H. Zhang, Y. Zou, and H. Wang, "Contrastive Self-Supervised Learning for Text-Independent Speaker Verification," in *ICASSP*, 2021.
- [9] 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.
- [10] M. Ravanelli and Y. Bengio, "Learning Speaker Representations with Mutual Information," in *Interspeech*, 2019.
- [11] T. Lepage and R. Dehak, "Label-Efficient Self-Supervised Speaker Verification With Information Maximization and Contrastive Learning," in *Interspeech*, 2022.
- [12] M. Mohammadami, D. Matrouf, J.-F. A. Bonastre, S. Dowerah, R. Serizel, and D. Jouviet, "Barlow Twins self-supervised learning for robust speaker recognition," in *Interspeech*, 2022.
- [13] Y. Chen, S. Zheng, H. Wang, L. Cheng, and Q. Chen, "Pushing the limits of self-supervised speaker verification using regularized distillation framework," in *ICASSP*, 2023.
- [14] J. Cho, J. Villalba, L. Moro-Velazquez, and N. Dehak, "Non-Contrastive Self-Supervised Learning for Utterance-Level Information Extraction From Speech," *IEEE JSTSP*, 2022.
- [15] C. Zhang and D. Yu, "C3-DINO: Joint Contrastive and Non-Contrastive Self-Supervised Learning for Speaker Verification," *IEEE JSTSP*, 2022.
- [16] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in *NAACL*, 2018.
- [17] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Unsupervised Pre-Training for Speech Recognition," in *Interspeech*, 2019.
- [18] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations," in *NeurIPS*, 2020.
- [19] 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 TASLP*, 2021.
- [20] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, J. Wu, M. Zeng, X. Yu, and F. Wei, "WavLM: Large-Scale Self-Supervised Pre-Training for Full Stack Speech Processing," *IEEE JSTSP*, 2022.
- [21] 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?" *arXiv preprint arXiv:2204.12765*, 2022.
- [22] S. A. Chowdhury, N. Durrani, and A. Ali, "What do end-to-end speech models learn about speaker, language and channel information? A layer-wise and neuron-level analysis," *Computer Speech & Language*, 2024.
- [23] Z. Fan, M. Li, S. Zhou, and B. Xu, "Exploring wav2vec 2.0 on Speaker Verification and Language Identification," in *Interspeech*, 2021.
- [24] J. Peng, O. Plhot, T. Stafylakis, L. Mošner, L. Burget, and J. Černocký, "An Attention-Based Backend Allowing Efficient Fine-Tuning of Transformer Models for Speaker Verification," in *IEEE SLT*, 2022.
- [25] H. Bing, C. Zhengyang, and Q. Yanmin, "Self-Supervised Speaker Verification Using Dynamic Loss-Gate and Label Correction," in *Interspeech*, 2022.
- [26] J. Cho, J. Villalba, and N. Dehak, "The JHU submission to VoxSRC-21: Track 3," in *VoxSRC Workshop*, 2021.
- [27] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A Simple Framework for Contrastive Learning of Visual Representations," in *ICML*, 2020.
- [28] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, "Emerging Properties in Self-Supervised Vision Transformers," in *ICCV*, 2021.
- [29] S. Chi, Q. Xipeng, X. Yige, and H. Xuanjing, "How to Fine-Tune BERT for Text Classification?" in *CCL*, 2019.
- [30] T. Ruijie, L. Kong Aik, D. Rohan Kumar, H. Ville, and L. Haizhou, "Self-supervised Speaker Recognition with Loss-gated Learning," in *ICASSP*, 2022.
- [31] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep Speaker Recognition," in *Interspeech*, 2018.
- [32] S. David, C. Guoguo, and P. Daniel, "MUSAN: A Music, Speech, and Noise Corpus," *arXiv preprint arXiv:1510.08484*, 2015.
- [33] T. Ko, V. Peddinti, D. Povey, M. L. Seltzer, and S. Khudanpur, "A study on data augmentation of reverberant speech for robust speech recognition," in *ICASSP*, 2017.
- [34] A. Nagrani, J. S. Chung, and A. Zisserman, "VoxCeleb: A Large-Scale Speaker Identification Dataset," in *Interspeech*, 2017.
- [35] C. Danwei and L. Ming, "The DKU-DukeECE System for the Self-Supervised Speaker Verification Task of the 2021 VoxCeleb Speaker Recognition Challenge," in *VoxSRC Workshop*, 2021.
- [36] S. H. Mun, M. H. Han, and N. S. Kim, "SNU-HIL system for the VoxCeleb speaker recognition challenge 2021," in *VoxSRC Workshop*, 2021.