



Why does Self-Supervised Learning for Speech Recognition Benefit Speaker Recognition?

Sanyuan Chen¹, Yu Wu², Chengyi Wang², Shujie Liu², Zhuo Chen², Peidong Wang², Gang Liu², Jinyu Li², Jian Wu², Xiangzhan Yu¹, Furu Wei²

¹Harbin Institute of Technology, China
²Microsoft Corporation

{yuwul, t-chewang, shujliu, zhuc, peidongwang, liu.gang, jinyli, fuwei}@microsoft.com
sychen@ir.hit.edu.cn, jianwu@exchange.microsoft.com, yxz@hit.edu.cn

Abstract

Recently, self-supervised learning (SSL) has demonstrated strong performance in speaker recognition, even if the pre-training objective is designed for speech recognition. In this paper, we study which factor leads to the success of self-supervised learning on speaker-related tasks, e.g. speaker verification (SV), through a series of carefully designed experiments. Our empirical results on the Voxceleb-1 dataset suggest that the benefit of SSL to SV task is from a combination of mask speech prediction loss, data scale, and model size, while the SSL quantizer has a minor impact. We further employ the integrated gradients attribution method and loss landscape visualization to understand the effectiveness of self-supervised learning for speaker recognition performance.

Index Terms: Self-Supervised Learning, Speaker Recognition, Speaker Verification

1. Introduction

Recently, self-supervised learning (SSL) has achieved the state-of-the-art results on a diverse array of downstream speech tasks [1, 2, 3, 4, 5, 6, 7]. Typical SSL methods either discriminate the correlated positive samples from the negative ones (e.g. wav2vec 2.0) [1], or predict discrete pseudo-labels on the masked regions (e.g. HuBERT) [2]. Both methods try to implicitly learn short-time phonetic information from a huge amount of unlabeled speech, and mainly target at self-supervised learning for automatic speech recognition task (SSL4ASR).

Due to the high correlation with phoneme units, it is straightforward to understand that SSL4ASR has the potential to drastically improve the speech recognition task. Interestingly, SSL4ASR also achieves state-of-the-art performance on the speaker-related tasks, e.g. speaker verification (SV). For instance, WavLM [6] and BigSSL [3] show the best performance on different partitions of VoxCeleb1 dataset [8], and the ensemble of WavLM model and Res2Net [9, 10] ranks at the top position on VoxSRC 2021 speaker verification permanent leaderboard¹ with the team name Strasbourg-Spk.

In this work, our goal is to understand *which factor leads to the success of SSL4ASR in speaker recognition*. Specifically, we try to answer the following questions:

1. Can supervised ASR model benefit the SV task?
2. How does SSL benefit the SV task?
3. What is the best SSL setup for the SV task?

¹<https://competitions.codalab.org/competitions/34066#results>.

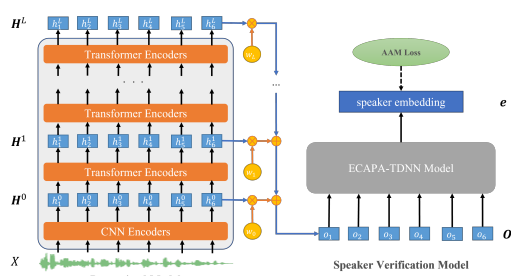


Figure 1: Self-supervised learning for speaker verification task.

To this end, we carefully design and conduct a series of experiments to investigate what is the indispensable part of SSL. We also perform Integrated Gradients attribution analysis and loss landscape visualization to further understand the contribution of SSL to SV performance.

The main finding is three-fold as follows. First, SSL4ASR models have significantly better transferability than supervised ASR models in an apple-to-apple comparison, indicating the SSL objective function is a key ingredient for achieving excellent transferability. Second, the HuBERT style loss, mask speech prediction, is slightly better than other SSL losses, such as contrastive learning and Mean Squared Error (MSE) loss, while how to generate pseudo-labels has minor impacts on the performance of HuBERT style models. Even pre-training with simple clustering methods on raw inputs could provide good performance on the SV task. Data augmentation proposed in WavLM [6] is very helpful, even if the pre-train data is scaled up to 94k hours. In addition, data scale and model scale have a strong correlation to model transferability. Third, our analysis shows that SSL models only learn speaker related knowledge with shallow layers in pre-training stage, while fine-tuning stage could unleash the full capability of the model. We observe that an SSL model could provide a wider optima in fine-tuning, which enables better resistance against small perturbation, stronger generalization capability, and easier SV model optimization.

2. Background

Self-supervised learning (SSL) has been shown to be an effective means of improving state-of-the-art results on SV task [5, 11, 6]. A common practice is: We first optimize the pre-train model with SSL objective on the large-scale unsupervised data, then fine-tune the pre-trained model along with a downstream SV model on the annotated dataset.

The typical SSL objectives are designed for automatic

speech recognition task (SSL4ASR) by implicitly learning short-time phonetic information from unlabeled speech [1, 2]. Specifically, given a raw audio X , a latent representation $\mathbf{H}^0 = \{\mathbf{h}_t^0\}_{t=1}^T$ is obtained by a CNN feature extractor, where T is the number of frames. Then the representation is fed to an L layer Transformer model, yielding hidden states $\mathbf{H}^l = \{\mathbf{h}_t^l\}_{t=1}^T$, where l denotes the l -th layer in the encoder. During pre-training, we employ the masked-based self-supervised learning methods to optimize the Transformer-based model. Before feeding the latent representation to the Transformer model, SSL methods first mask a proportion of them in some random frames, then minimize a variety of self-supervised objective functions based on the last layer hidden states output \mathbf{H}^L in the masked regions.

During fine-tuning, we weighted average the hidden states of each layer to generate the output representation $\mathbf{O} = \{\mathbf{o}_t = \sum_{l=0}^L w_l \cdot \mathbf{h}_t^l\}_{t=1}^T$, where w_l is a learnable weight for the hidden state of the l -th layer. Then we employ ECAPA-TDNN [12] as the downstream SV model following [11], and feed the output representation into the downstream model to generate the speaker embedding $\mathbf{e} = \text{ECAPA-TDNN}(\mathbf{o}_1, \dots, \mathbf{o}_T)$. We use the additive angular margin (AAM) loss [13] as the supervised objective function, and train the downstream SV model along with the pre-trained model for two stages. In the first stage, we optimize the parameters of the downstream model with the pretrained parameters fixed. In the second stage, we continue to optimize the parameters of the downstream model as well as the pre-trained model. In addition, we can also apply large-margin fine-tuning strategy and score calibration to further improve the speaker verification performance [14].

3. Why does SSL4ASR Benefit the SV task?

3.1. Can supervised ASR model benefit the SV task?

Given the similar modeling unit between SSL4ASR and supervised ASR models, it is a natural question whether the supervised ASR model can also benefit speaker verification task. To verify this hypothesis, we compare the transferability of supervised ASR and SSL4ASR, both of which are trained on LibriSpeech 960h [15] and use the Transformer structure of HuBERT [2].

The ASR model is trained with the Connectionist Temporal Classification (CTC) loss function [16] in a supervised way. We use the character sequence as the target golden labels, and require the ASR model to predict the golden label Y^{golden} given the hidden states of the last encoder layer \mathbf{H}^L : $\mathcal{L}_{\text{CTC}} = -\log p_{\text{CTC}}(Y^{\text{golden}}|\mathbf{H}^L)$. Spec-augmentation is also applied following [17].

HuBERT, based on masked pseudo-label prediction loss, is selected as the SSL4ASR model for the comparison [2]. The pseudo labels are generated by iterative clustering. At the first iteration, we conduct an offline clustering step on the MFCC feature of the input audio, where the clustering center of each frame is indexed as the pseudo label. Then we use the hidden states $\mathbf{H}^L = \{\mathbf{h}_t^L\}_{t=1}^T$ to predict the embeddings $\mathbf{E} = \{\mathbf{e}_{y_t}\}_{t=1}^T$ corresponding to the pseudo labels $Y = \{y_t\}_{t=1}^T$ with the cross-entropy loss function in the masked regions M :

$$\mathcal{L}_{\text{HuBERT}} = \sum_{t \in M} -\log \frac{\exp(\text{sim}(\mathbf{h}_t^L \mathbf{W}, \mathbf{e}_{y_t})/\tau)}{\sum_{\hat{y}=1}^K \exp(\text{sim}(\mathbf{h}_t^L \mathbf{W}, \mathbf{e}_{\hat{y}})/\tau)},$$

where \mathbf{W} is the projection matrix, $\text{sim}(\cdot, \cdot)$ denote the cosine similarity function, τ is a pre-defined hyperparameter, and K is the number of clusters. Starting from the second iteration,

Table 1: *Transferability of supervised ASR and SSL4ASR*

Model	EER (%)		
	Vox1-O	Vox1-E	Vox1-H
FBank [12]	1.01	1.24	2.32
Random	3.696	3.71	6.034
CTC	1.159	1.256	2.434
HuBERT	0.84	0.879	1.726

Table 2: *SSL with different objective functions*

Model	EER (%)		
	Vox1-O	Vox1-E	Vox1-H
MSE	0.979	1.075	1.98
wav2vec 2.0	0.973	0.933	1.831
HuBERT	0.84	0.879	1.726

we perform the offline clustering step with the hidden states extracted from the last iteration pretrained HuBERT model, and then train a new HuBERT model with the pseudo labels obtained by the new clustering centers.

We also use a random initialized Transformer model as a baseline to get rid of the effect of the additional parameters introduced by the pre-trained model and focus on the performance of different pre-training methods.

Table 1 shows SSL4ASR model can provide a better representation than the handcrafted FBank feature, while the representations from the ASR model with CTC loss and the random initialized Transformer model are inferior to the FBank feature. It indicates that the key to the success of SSL4ASR on SV task is neither the Transformer structure nor the fine-tuning pipeline, but the self-supervised learning procedure.

3.2. What is the best SSL objective for the SV task?

Besides HuBERT, which is based on masked pseudo-label prediction loss, we also evaluate the transferability of wav2vec 2.0 [1] and Mean Squared Error (MSE) loss based pre-training method. It should be noted that all the three methods use the same mask setting proposed in HuBERT.

MSE firstly calculates the FBank feature $\mathbf{F} = \{\mathbf{f}_t\}_{t=1}^T$ of the raw audio, then measures the mean square error between the FBank feature and the linear projection of last layer hidden states output $\mathbf{H}^L = \{\mathbf{h}_t^L\}_{t=1}^T$ in the masked regions M as the objective function: $\mathcal{L}_{\text{MSE}} = \sum_{t \in M} \|\mathbf{f}_t - \mathbf{h}_t^L\|^2$.

Wav2vec 2.0 firstly discretizes the latent representation \mathbf{h}_t^0 of each masked timestep t to the quantized latent representation \mathbf{q}_t , then uses the context representation \mathbf{h}_t^L to identify the true quantized latent representation \mathbf{q}_t out of a set of candidate representations $\hat{\mathbf{q}} \in \mathbf{Q}_t$ with contrastive loss function:

$$\mathcal{L}_{\text{wav2vec 2.0}} = \sum_{t \in M} -\log \frac{\exp(\text{sim}(\mathbf{h}_t^L \mathbf{W}, \mathbf{q}_t)/\tau)}{\sum_{\hat{\mathbf{q}} \in \mathbf{Q}_t} \exp(\text{sim}(\mathbf{h}_t^L \mathbf{W}, \hat{\mathbf{q}})/\tau)}.$$

Table 2 demonstrates that all the three SSL methods can provide better representation than the FBank feature, which is attributed to the contextual speech representation learning from the masked speech. HuBERT achieves the best performance, indicating the better generalization and effectiveness of pseudo-label prediction loss than contrastive loss and MSE loss.

3.3. What is the best SSL quantizer for the SV task?

Since HuBERT style loss is better than others, we explore the performance of different pseudo-label creation methods (quantizers) for HuBERT loss. Besides the MFCC Clustering and Hidden State Clustering introduced by HuBERT, we also ex-

Table 3: HuBERT style loss with different quantizers.

Model	EER (%)		
	Vox1-O	Vox1-E	Vox1-H
MFCC Clustering	0.872	0.917	1.766
Hidden State Clustering	0.840	0.879	1.726
Random Projection (500 codes)	0.899	0.95	1.775
Random Projection (8192 codes)	0.883	0.903	1.675
VQ-VAE	0.824	0.899	1.655
Phoneme	0.867	0.918	1.776

periment with the labels obtained by Random Projection [18], VQ-VAE quantizers [19], and frame-phoneme alignment.

With random projection quantizer, we first extract the FBank features $\mathbf{F} = \{\mathbf{f}_t\}_{t=1}^T$ of the input audio, project \mathbf{f}_t to the vector $\mathbf{A}\mathbf{f}_t$ with a random initialized matrix \mathbf{A} , and then find the closest vector from a set of random initialized vectors $\mathbf{V} = \{\mathbf{v}_i\}_{i=1}^K$, where K is the vector (code) numbers. The pseudo label of t -th frame is defined as the index of the closest vector: $y_t = \arg \min_i \|\mathbf{v}_i - \mathbf{A}\mathbf{f}_t\|$.

With VQ-VAE quantizer, we first extract the FBank features $\mathbf{F} = \{\mathbf{f}_t\}_{t=1}^T$ of the input audio, and train a VQ-VAE model [19] to reconstruct the FBank feature on LibriSpeech 960h [15]. Given the latent variable $\mathbf{Z} = \{\mathbf{z}_t\}_{t=1}^T$ obtained by a 6-layer Transformer-based encoder, we discretize it with the closest vector $\{\hat{\mathbf{z}}_t = \arg \min_{\mathbf{v}_i} \|\mathbf{v}_i - \mathbf{z}_t\|\}_{t=1}^T$ in a latent embedding space $\mathbf{V} = \{\mathbf{v}_i\}_{i=1}^K$, where K is the embedding numbers, and then reconstruct the features $z_q(\mathbf{F})$ with a 6-layer Transformer-based decoder. The training loss of VQ-VAE is to minimize the mean squared error between the reconstructed features and the input features, along with the difference between the encoded variable and the discrete variable:

$$\mathcal{L}_{\text{VQ-VAE}} = \|z_q(\mathbf{F}) - \mathbf{F}\|^2 + \|\text{sg}[\mathbf{Z}] - \hat{\mathbf{Z}}\|^2 + \beta \|\mathbf{Z} - \text{sg}[\hat{\mathbf{Z}}]\|^2,$$

where $\text{sg}[\cdot]$ is the stopgradient operator and β is a pre-defined hyperparameter. During inference, the pseudo label of t -th frame is defined as the index of the discrete latent variables in the latent embedding space: $y_t = \arg \min_i \|\mathbf{v}_i - \mathbf{z}_t\|$.

In addition, we also consider using the phoneme sequence of the input audio as the pseudo label to see if ASR-related pseudo label can benefit the SV performance. Here, we use force-alignment tool [20] to get the frame-phoneme pairs on LibriSpeech 960h data.

Table 3 shows that all the quantizers have similar performance on the speaker verification task. Even when we use the phone sequence as the pseudo label, which is irrelevant to the speaker information, we can still obtain a well-performed speaker verification model with the masked pseudo-label prediction SSL method.

3.4. Large-Scale SSL on SV task

Moreover, we also leverage the data augmentation and scale-up strategy to further strengthen the self-supervised learning for speaker verification task. Following WavLM [6], we employ the masked speech denoising and prediction framework as the data augmented self-supervised learning method to improve pre-trained model robustness for complex acoustic environments and the preservation of speaker identity. We also scale up unlabeled pre-training data to 94k hours of public audios [6], including 60k hours of Libri-Light [21], 10k hours of GigaSpeech [22], and 24k hours of VoxPopuli [23], and enlarge the model to 24 layer Transformers with 316M parameters.

Table 4 shows that the data augmentation strategy used in WavLM can successfully benefit the self-supervised learning

Table 4: Data and Model Scale Up. * means using large margin finetune and calibration

Model	EER (%)		
	Vox1-O	Vox1-E	Vox1-H
HuBERT 960h	0.84	0.879	1.726
WavLM 960h	0.777	0.829	1.629
HuBERT 94kh	0.734	0.847	1.725
WavLM 94kh	0.739	0.742	1.483
WavLM 94kh Large	0.505	0.579	1.176
WavLM 94kh Large*	0.308	0.462	0.906

for SV task. The performance improvement would be more significant if we scale up the pre-training data to 94kh. Thanks to the larger parameter capacity, the WavLM Large model can bring more than 20% EER reduction compared to the WavLM Base model. With the large-margin fine-tuning strategy and score calibration methods, the WavLM Large model can achieve 33.2%, 27.1%, and 8.8% relatively EER reduction compared to the state-of-the-art supervised model (Vox1-O: 0.461, Vox1-E: 0.634, Vox1-H: 0.993) [24] on all the three VoxCeleb1 trial lists.

4. Discussion and Analysis

4.1. Contribution Attribution

We employ the Integrated Gradients (IG) attribution method [25] to demonstrate how each layer of the pre-trained model contributes to the final SV performance. Compared with method in [5, 6], IG better models contribution estimation as it consider not only the layer weight, but also the magnitude of each layer’s hidden states. Specifically, given a well-trained downstream model $F(\cdot)$, the hidden states $\{\mathbf{H}^i\}_{i=1}^L$ extracted from all layers, and the corresponding learned weights $\{w_i\}_{i=1}^L$, the attribution score of l -th layer hidden states is assigned as:


$$IG(\mathbf{H}^l) = \sum_{t,f} \left(\mathbf{H}^l \times \int_{\alpha=0}^1 \frac{\partial F(\alpha(\sum_{i=0}^L w_i \cdot \mathbf{H}^i))}{\partial \mathbf{H}^l} d\alpha \right),$$

where \times denotes Hadamard product, and $\sum_{t,f}(\cdot)$ denotes the summation over the time and feature dimensions. The larger attribution score indicates the more importance of the corresponding hidden states. The summation of the attribution scores of all the hidden states indicates the final prediction of the SV model, i.e., $\sum_{l=0}^L IG(\mathbf{H}^l) = F(\sum_{l=0}^L w_l \cdot \mathbf{H}^l) - F(0)$. Due to the intractability, we approximate $IG(\cdot)$ with the gradients summation as:


$$IG^{\text{approx}}(\mathbf{H}^l) = \sum_{t,f} \left(\mathbf{H}^l \times \frac{1}{K} \sum_{k=0}^K \frac{\partial F(\frac{k}{K}(\sum_{i=0}^L w_i \cdot \mathbf{H}^i))}{\partial \mathbf{H}^l} \right),$$

where K is the number of approximation steps for computing integrated gradients. We set K to 50 in our experiment.

Figure 2 shows the contribution attribution from each layer of different pre-trained models. As for the first stage of fine-tuning, where we train the downstream model with the pre-trained parameters fixed, the contribution mostly comes from the output of the CNN feature extractor and the first encoder layer for all the pre-trained models. It indicates that only the shallow layers of pre-trained models learn the speaker-related information during the self-supervised learning procedure. If the hidden states are extracted from the ASR model, which is supervised trained with CTC loss, only the latent feature extracted by CNN extractor contributes to the final prediction. And if the hidden states are extracted from the SSL4ASR model, such as wav2vec 2.0 and HuBERT, the contribution is also dominated by the CNN extracted feature. In contrast, if



(a) Stage 1: we fix the pre-trained model and only train the downstream model.



(b) Stage 2: we train both the pre-trained model and the downstream model.

Figure 2: Contribution attributed to each layer of each pre-trained model

we pre-train HuBERT with data augmentation or the phoneme-independent quantizer, such as MFCC clustering or random projection, there are more contributions from the hidden states encoded by Transformer layers.

As for the second stage of fine-tuning, we update the parameters of the downstream model as well as the pre-trained parameters. Since we unleash the full capability of the pre-trained model, the higher Transformer-based encoder layers can also learn to model the speaker information with the SV training objective, and make more contribution to the final prediction than in the first stage, leading to better speaker verification performance .

4.2. Loss Landscape Visualization

To better understand how self-supervised learning benefits the SV task, we visualize and compare the two-dimensional loss landscapes along with the optimization trajectories of different SV models. For better comparison of different input features, we plot the parameters of the downstream models, and the optimization trajectories in the first fine-tuning stage where the pre-trained parameters are kept frozen.

Following [26, 27], we first define the origin and two axes of the loss surface as the random initialized downstream model’s parameters and two directions in the parameter space, respectively. Then, we uniformly sample multiple points around the initialized parameters, and plot the training loss of the downstream model with the parameters of each sampled point and the input feature from the pre-trained model.

Let θ_0, θ_1 denote the random initialized parameters and well-trained parameters of the SV downstream model respectively, we can define one of the axes as the optimization direction $\delta_1 = \theta_1 - \theta_0$. The other axis is set as a random direction $\delta_2 = \theta_2 - \theta_0$, where θ_2 is the randomly generated parameters. Due to the high-dimensional parameter space, experimental results confirm that the two axes δ_1 and δ_2 are divergent and orthogonal to each other. Then, the 2-D loss surface can be plotted with the function: $f(\alpha, \beta) = \mathcal{L}(\theta_1 + \alpha\delta_1 + \beta\delta_2)$, where α, β are scalar values and \mathcal{L} is the loss function of the SV model training. For better visualization, we scale the second direction

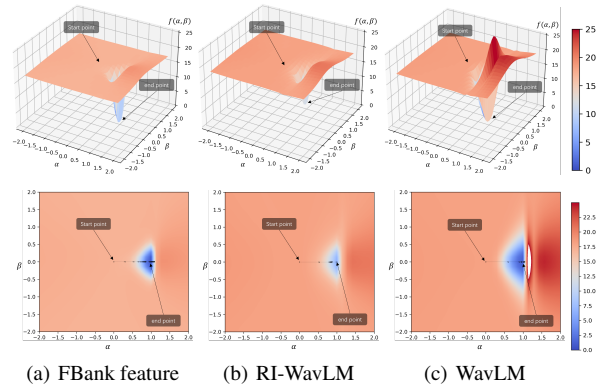


Figure 3: Visualization of the loss landscape and optimization trajectories of SV model with different input features (FBank feature, Random-Initialized WavLM feature, and WavLM feature). The figure below is a top view of the figure above.

vector δ_2 to the same norm as the first one δ_1 by $\delta_2 = \frac{\|\delta_1\|}{\|\delta_2\|}\delta_2$, where $\|\cdot\|$ is the Euclidean norm. We set the range of α and β to $[-2, 2]$, and uniformly sample 29 points for each axis. In addition, we also project the optimization trajectory of the SV downstream model onto the two-dimensional loss surface. Specifically, for the parameters of the downstream model θ^i at i -th training epoch, $\delta^i = \theta^i - \theta_0$ denotes the optimization direction at the i -th epoch, we can calculate the cosine similarity between the optimization direction δ^i and each of the projected directions δ as $\cos(\delta^i, \delta) = \frac{\delta^i \cdot \delta}{\|\delta^i\| \|\delta\|}$. Then, the corresponding projected point (α^i, β^i) in the 2-D loss surface of θ^i can be calculated as: $\alpha^i = \frac{\|\delta^i\|}{\|\delta_1\|} \cos(\delta^i, \delta_1) = \frac{\delta^i \cdot \delta_1}{\|\delta_1\|^2}, \beta^i = \frac{\delta^i \cdot \delta_2}{\|\delta_2\|^2}$.

Figure 3 shows the visualization of speaker verification downstream model with different input features. Compared with the FBank feature, we can find that the representation with the Random-Initialized WavLM model can provide a wider optima, which enables better resistance against some small perturbation, and leads to easier SV model optimization. However, without the self-supervised pretraining, the speaker verification model would stuck into a poor local minima with worse speaker verification performance. With large-scale self-supervised learning, the pretrained WavLM representation can provide a better initial point with a much broader and deeper optimum area. Even with some small disturbance, the WavLM input feature enables the downstream model to converge to the expected optimal region, and prevent it from skipping the optimal region with a steep loss hill.

5. Conclusion

Our experimental results demonstrate that the self-supervised learning procedure is the key to the success on SV task. Among a variety of SSL methods, the masked pseudo-label prediction loss can provide the representation with best generalization capability on SV task, regardless of the pseudo-label creation methods. We also show that data augmentation and model scale-up can further strengthen SSL for SV task. Moreover, our analyses show that two-stage fine-tuning can make use of the full capacity of SSL models, and that SSL models can facilitate the SV model optimization with a better initial point with a broader and deeper optimum area.

6. References

- [1] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 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,” *arXiv preprint arXiv:2106.07447*, 2021.
- [3] Y. Zhang, D. S. Park, W. Han, J. Qin, A. Gulati, J. Shor, A. Jansen, Y. Xu, Y. Huang, S. Wang *et al.*, “Bigssl: Exploring the frontier of large-scale semi-supervised learning for automatic speech recognition,” *arXiv preprint arXiv:2109.13226*, 2021.
- [4] S. Yang, P. Chi, Y. Chuang, C. J. Lai, K. Lakhotia, Y. Y. Lin, A. T. Liu, J. Shi, X. Chang, G. Lin, T. Huang, W. Tseng, K. Lee, D. Liu, Z. Huang, S. Dong, S. Li, S. Watanabe, A. Mohamed, and H. Lee, “SUPERB: Speech Processing Universal PERFORMANCE Benchmark,” in *Interspeech*, 2021, pp. 1194–1198.
- [5] S. Chen, Y. Wu, C. Wang, Z. Chen, Z. Chen, S. Liu, J. Wu, Y. Qian, F. Wei, J. Li *et al.*, “Unispeech-sat: Universal speech representation learning with speaker aware pre-training,” *arXiv preprint arXiv:2110.05752*, 2021.
- [6] 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,” *arXiv preprint arXiv:2110.13900*, 2021.
- [7] C. Wang, Y. Wu, S. Chen, S. Liu, J. Li, Y. Qian, and Z. Yang, “Self-supervised learning for speech recognition with intermediate layer supervision,” *arXiv preprint arXiv:2112.08778*, 2021.
- [8] A. Nagrani, J. S. Chung, W. Xie, and A. Zisserman, “Voxceleb: Large-scale speaker verification in the wild,” *Computer Speech & Language*, vol. 60, p. 101027, 2020.
- [9] S.-H. Gao, M.-M. Cheng, K. Zhao, X.-Y. Zhang, M.-H. Yang, and P. Torr, “Res2net: A new multi-scale backbone architecture,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 43, no. 2, pp. 652–662, 2019.
- [10] T. Zhou, Y. Zhao, and J. Wu, “Resnext and res2net structures for speaker verification,” in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 301–307.
- [11] 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,” *arXiv preprint arXiv:2110.05777*, 2021.
- [12] B. Desplanques, J. Thienpondt, and K. Demuynck, “Ecapadnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification,” *arXiv preprint arXiv:2005.07143*, 2020.
- [13] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, “Arcface: Additive angular margin loss for deep face recognition,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 4690–4699.
- [14] 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 *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5814–5818.
- [15] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An ASR corpus based on public domain audio books,” in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [16] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *International Conference on Machine Learning (ICML)*, 2006, pp. 369–376.
- [17] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” in *Interspeech*, 2019, pp. 2613–2617.
- [18] C.-C. Chiu, J. Qin, Y. Zhang, J. Yu, and Y. Wu, “Self-supervised learning with random-projection quantizer for speech recognition,” *arXiv preprint arXiv:2202.01855*, 2022.
- [19] A. Van Den Oord, O. Vinyals *et al.*, “Neural discrete representation learning,” *Advances in neural information processing systems*, vol. 30, 2017.
- [20] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, “Montreal Forced Aligner: Trainable Text-Speech Alignment Using Kaldi,” in *Proc. Interspeech 2017*, 2017, pp. 498–502.
- [21] J. Kahn, M. Rivière, W. Zheng, E. Kharitonov, Q. Xu, P.-E. Mazaré, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen *et al.*, “Libri-light: A benchmark for asr with limited or no supervision,” in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7669–7673.
- [22] G. Chen, S. Chai, G. Wang, J. Du, W. Zhang, C. Weng, D. Su, D. Povey, J. Trmal, J. Zhang, M. Jin, S. Khudanpur, S. Watanabe, S. Zhao, W. Zou, X. Li, X. Yao, Y. Wang, Y. Wang, Z. You, and Z. Yan, “Gigaspeech: An evolving, multi-domain asr corpus with 10,000 hours of transcribed audio,” in *Interspeech*, 2021.
- [23] C. Wang, M. Rivière, A. Lee, A. Wu, C. Talnikar, D. Haziza, M. Williamson, J. Pino, and E. Dupoux, “Voxpopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation,” *arXiv preprint arXiv:2101.00390*, 2021.
- [24] M. Zhao, Y. Ma, M. Liu, and M. Xu, “The speakin system for voxceleb speaker recognition challenge 2021,” *arXiv preprint arXiv:2109.01989*, 2021.
- [25] M. Sundararajan, A. Taly, and Q. Yan, “Axiomatic attribution for deep networks,” in *International conference on machine learning*. PMLR, 2017, pp. 3319–3328.
- [26] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein, “Visualizing the loss landscape of neural nets,” *Advances in neural information processing systems*, vol. 31, 2018.
- [27] Y. Hao, L. Dong, F. Wei, and K. Xu, “Visualizing and understanding the effectiveness of bert,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019, pp. 4143–4152.