



# Application of Knowledge Distillation to Multi-Task Speech Representation Learning

\**Mine Kerpicci*<sup>1</sup>    *Van Nguyen*<sup>2</sup>    *Shuhua Zhang*<sup>2</sup>    *Erik Visser*<sup>2</sup>

<sup>1</sup> Georgia Institute of Technology, Atlanta, Georgia

<sup>2</sup> Qualcomm Technologies, Inc., San Diego, California

mkerpicci3@gatech.edu, vanguyen@qti.qualcomm.com, shuhuaz@qti.qualcomm.com,  
evisser@qti.qualcomm.com

## Abstract

Model architectures such as wav2vec 2.0 and HuBERT have been proposed to learn speech representations from audio waveforms in a self-supervised manner. When they are combined with downstream tasks such as keyword spotting and speaker verification, they provide state-of-the-art performance. However, these models use a large number of parameters, the smallest version of which has 95 million parameters. This constitutes a challenge for edge AI device deployments. In this paper, we investigate the application of knowledge distillation to speech representation learning (SRL) models followed by joint fine-tuning with multiple downstream voice-activated tasks. In our experiments on two such tasks, our approach results in nearly 75% reduction in model size while suffering only 0.1% accuracy and 0.9% equal error rate degradation compared to the full-size model. In addition, we show that fine-tuning the SRL models results in a significant performance boost compared to using frozen SRL models.

**Index Terms:** Speech representation learning, multi-task training, wav2vec, HuBERT, knowledge distillation

## 1. Introduction

Speech representation learning (SRL) has been extensively studied in the literature and shown potential for various speech recognition tasks [1, 2, 3]. Deep neural network architectures [4, 5] are the method of choice for SRL. Since sufficient amounts of labeled data are not typically available for supervised learning in real scenarios, this problem is circumvented by using self-supervised learning on unlabeled data [6, 7, 8, 9, 10]. In [6, 7], convolutional networks are combined with transformer encoder-based models to learn speech representations in a self-supervised manner. [11, 12, 13] propose to combine such a model with a downstream network to address classification problems such as emotion recognition or speaker recognition. The resulting models provide competitive performance on these tasks but require training multiple large networks each for a different single task. To enable models for multi-tasking, [14] proposes to learn speech representations from raw audio with wav2vec 2.0 [6] and then fine-tune the model for multiple tasks simultaneously by sampling training instances from the combined task-specific datasets. This method provides even better performance than training on single tasks separately.

However, the memory footprint of the resulting models is large with the number of parameters ranging from 95 million to 1 billion, which makes them difficult to deploy on edge AI devices. This problem is studied with parameter sharing in [15]. Knowledge distillation is another method of choice to address

model complexity reduction in various fields ranging from natural language processing to speech recognition [16, 17, 18]. A common approach in the literature is to reduce model size by maintaining its performance via teacher-student network distillation architectures. [19, 20] apply knowledge distillation to large SRL methods to construct efficient speech recognition networks. The conventional approach is transferring knowledge from the final layer of the teacher network to the final prediction of the student network. However, the middle layers of these SRL models also contain valuable speech representation information [21]. In [22], layer-wise distillation to the HuBERT [7] model is applied where knowledge is transferred from middle layers of the large network to construct a smaller model.

In our work, we discuss generalization of knowledge distillation to use a large network in training of another network with reduced model size, and we use a special knowledge distillation approach as presented in [22] in our experiments. As opposed to this approach, we combine the student model with linear downstream heads and fine-tune all network parameters via a multi-task training scheme as in [14] instead of freezing the SRL model as done in [22]. In our experiments, we address keyword spotting and speaker verification problems in single task and multi-task frameworks. Through our results, we show that the constructed student networks perform as well as the teacher models even though the former model size is 28% of the latter. Moreover, we perform this in two settings where we construct distilled SRL modules with wav2vec 2.0 and HuBERT separately for a comprehensive performance comparison.

To the best of our knowledge, we propose for the first time the combination of knowledge distillation application to SRL and joint fine-tuning of the complete model (SRL module and downstream heads) for multiple downstream voice-activated tasks. This is in contrast to common speech recognition multi-task training approaches where an SRL module is frozen, functioning as a feature extractor, and only the downstream speech recognition tasks are trained [22, 23]. On the contrary, we show that fine-tuning both distilled SRL module and downstream heads for multiple voice-activated tasks simultaneously achieves significantly higher performance on all such tasks compared to fine-tuning only the downstream heads with frozen SRL module.

The main contributions of this work are as follows:

- We successfully combine knowledge distillation and multi-task training to construct a single network (with about 75% reduced size) that can be efficiently used in devices and embedded systems for multiple tasks.
- We show that training the complete network end-to-end (both SRL module and downstream heads) on multiple tasks achieves significantly higher performance than training on the same tasks with the frozen SRL module.

\* Work performed as an intern at Qualcomm Technologies

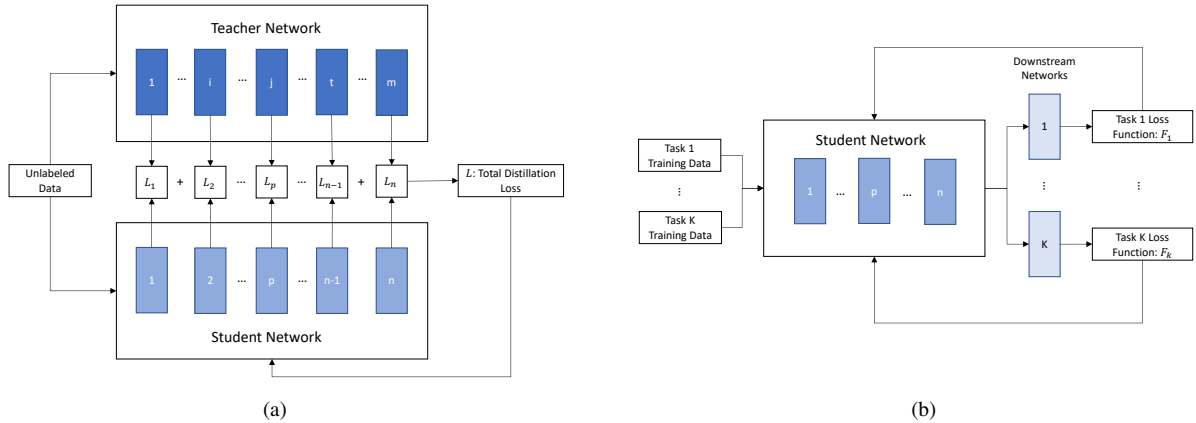


Figure 1: Flowchart of (a) the knowledge distillation process where the student network is trained via hidden layers of the pre-trained teacher network (distillation can be performed among any desired layers with a suitable loss function choice), (b) the fine-tuning step where the distilled network is trained with a multi-task training scheme.

- We discuss generalization of the used methods (knowledge distillation and multi-task training) and show that they can be efficiently applied to different networks where we achieve competitive results with both distilled wav2vec 2.0 and distilled HuBERT models on keyword spotting (KWS) and speaker verification (SV) voice-activated tasks.

The paper is organized as follows. Section 2 describes our approach in detail, Section 3 presents experiments conducted to validate our approach and we conclude with Section 4.

## 2. METHOD

We apply the student-teacher approach [16] to construct our SRL module. We use wav2vec 2.0 and HuBERT models as our teacher network as discussed in Section 2.1 and apply distillation method explained in Section 2.2. Then, we fine-tune the student model and downstream network with the multi-task training approach in Section 2.3. We provide the flowchart of the complete process in Figure 1. Note that the proposed end-to-end training scheme can be applied to various teacher-student models for several tasks. Here, in particular, we investigate its application to SRL models and present our experimental analysis on multiple voice-activated tasks.

### 2.1. Speech Representation Learning

HuBERT [7] and wav2vec 2.0 [6] methods are capable of learning powerful speech representations from audio waveforms in a self-supervised manner. They share the same architecture consisting of CNN and transformer encoder layers. Their BASE model corresponding to the smallest version contains 95 million parameters while LARGE model contains 317 million parameters. HuBERT has an additional X-LARGE model with 1 billion parameters. While sharing a similar structure, differences result from their pre-training schemes. The wav2vec 2.0 model has a single step of training with a quantization process to learn the targets. HuBERT, however, first constructs the pseudo-targets in a separate clustering step and then learns these targets in the second step during training. As a result, each of their hidden layers learns different properties of speech representation and transfers various abilities when they are used for distillation.

In our experiments, we investigate both wav2vec 2.0 and

HuBERT as our teacher network for knowledge distillation to construct separate student networks. Then, we evaluate these in two settings to compare the performance of final student models and to show the applicability of the discussed processes.

### 2.2. Knowledge Distillation

Deep architectures achieve state-of-the-art performance in many tasks due to their high learning capacities. However, their deployment in mobile devices and embedded systems is a challenging problem due to cost and memory limitations. In the literature, this problem is addressed with the idea of transferring knowledge from such models to smaller networks to reduce network size while keeping the performance similar.

Knowledge distillation is a method that transfers knowledge from layers of a large pre-trained network (teacher) to a simpler or smaller model (student) as demonstrated in Figure 1a. As in this generalized representation, the transfer can be performed among any layers of the teacher and the student networks where the teacher guides the student through a defined loss function during pre-training [16]. Several studies have suggested that not only the final layer but also the middle layers of SRL modules contain valuable information [21, 24]. Hence, it is proposed to transfer knowledge from middle layers of teacher as well as its final layer to learn the hidden representations [22]. We use this approach in our pre-training as a special case of the process in Figure 1a. In this method, knowledge is transferred between only some of the student and teacher layers via feature vectors as in the diagram in Figure 2a. The student parameters are initialized with their corresponding layers in the teacher network and trained with the loss function

$$L_p = \sum_{i=1}^T \left( \frac{1}{K} \|\mathbf{f}_{p,i} - \bar{\mathbf{f}}_{p,i}\|_1 - \log \sigma \left( \cos(\mathbf{f}_{p,i}, \bar{\mathbf{f}}_{p,i}) \right) \right).$$

This function is separately calculated for all hidden layers where distillation is performed. Here,  $T$  is the number of time steps and  $K$  is the dimension of feature vectors where  $\mathbf{f}_p$  and  $\bar{\mathbf{f}}_p$  are the feature vectors of the  $p^{\text{th}}$  distillation layer of the teacher and student models, respectively. Moreover,  $\sigma$  is sigmoid function and  $\cos$  is cosine similarity. The aim of this loss function is to minimize  $L_1$  distance while maximizing the cosine similarity between the feature vectors. Note that

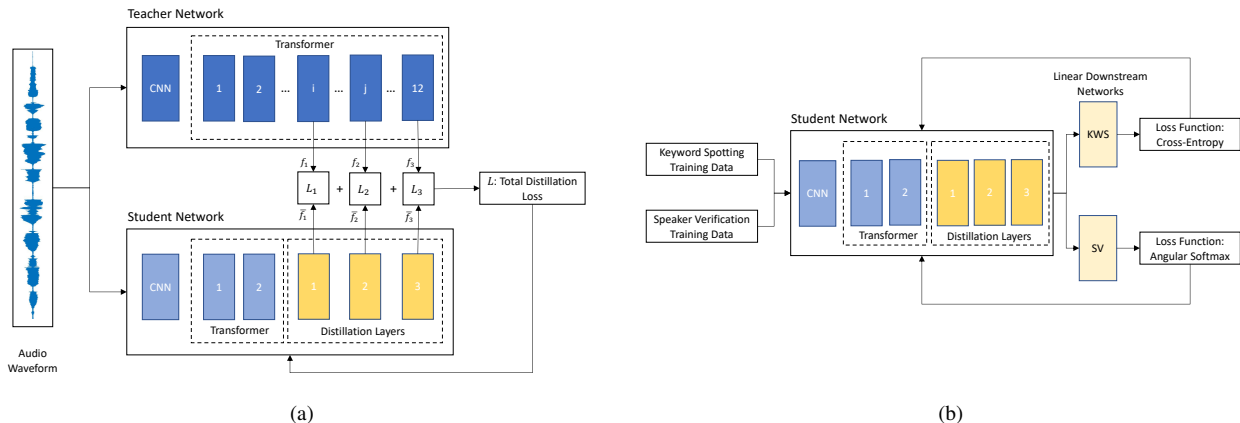


Figure 2: Flowchart of (a) the distillation process where the student model with CNN and Transformer layers is trained via teacher network layers, (b) the fine-tuning step where the complete distilled network is trained with a multi-task training scheme for keyword spotting and speaker verification tasks.

this method can be applied to the networks in various forms by changing the loss functions accordingly.

### 2.3. Multi-task Training

We combine the pre-trained student network with downstream networks to perform multi-task training. For this, we use separate downstream heads for KWS and SV tasks. Note that the choice of downstream heads can vary based on the applications. In our study, we use a linear layer that maps the output of student network to task specific vectors as in Figure 2b to keep the network size small.

In our work, we apply a special multi-task training method where we fine-tune all network parameters (both distilled student network and downstream heads) for all tasks simultaneously [14]. In each iteration, we first train all parameters based on the cross-entropy loss function with the KWS dataset and then use the angular softmax loss with the SV dataset. We alternate between the loss functions and batches of the tasks at each iteration. This alteration between both tasks continues until the defined maximum number of iterations is reached.

Note that the concept of multi-task learning in the literature is commonly considered as fine-tuning multiple downstream heads on their own datasets after freezing the SRL module that they are sharing [22, 23]. However, in this study, we discuss the multi-task training scheme where the SRL module is also trained with the datasets of all downstream tasks one-by-one during fine-tuning. This training scheme allows the SRL module to adapt its learning capability to the downstream tasks more while increasing its performance on all trained tasks. In our experiments, we provide the performance results of both training approaches and show that applying this multi-task training scheme without freezing the SRL module achieve significant performance gain in all investigated setups and tasks.

## 3. EXPERIMENTS

### 3.1. Experimental Setup

In our experiments, we evaluate the performance of the student networks after distillation on single task and multi-task frameworks. We perform our experiments in two settings where we use different models as teacher networks. For knowledge distil-

Table 1: The models constructed with the proposed approach and evaluated in the experiments where their number of parameters and fine-tuned network parts are presented.

Model	# Parameters	Fine-tuned Parameters
wav2vec 2.0	96 M	SRL Module + Downstream Heads
wav2vec 2.0 (frozen)	96 M	Downstream Heads
DistilledWav2vec	27 M	SRL Module + Downstream Heads
DistilledWav2vec (frozen)	27 M	Downstream Heads
DistilledHuBERT	27 M	SRL Module + Downstream Heads
DistilledHuBERT (frozen)	27 M	Downstream Heads

lation, we use BASE models of wav2vec 2.0 [6] and HuBERT [7], where they have the same architecture of CNN and 12 transformer layers resulting in 95 million parameters. Note that they are both pre-trained with 960-hour LibriSpeech data [25] to learn speech representations in their own training settings.

The student networks are smaller and simplified versions of the teacher architectures. They consist of CNN and 2 transformer layers as in Figure 2 where their parameters are initialized with the parameters of CNN and 1<sup>st</sup> and 2<sup>nd</sup> transformer layers of the teacher network. Then, we perform the knowledge distillation from the layers 4, 8 and 12 with the same 960-hour LibriSpeech data. The resulting student models with distillation and downstream layers contain 27 million parameters. We refer to these distilled student networks with downstream heads as “DistilledWav2vec” and “DistilledHuBERT”. Note that we keep the layers after distillation, and we fine-tune all network parameters including CNN and transformer layers on downstream tasks. For comparison, we also include the performance of the models when all layers are frozen except the downstream heads during fine-tuning step which is called as partial fine-tuning. We refer to these versions as “DistilledWav2vec (frozen)” and “DistilledHuBERT (frozen)”. The properties of the constructed models regarding the number of parameters and their fine-tuned parts are provided in Table 1 with their model names as referred in the experimental results.

We perform our experiments in both single task and multi-task frameworks for a comprehensive analysis. In these experiments, we use keyword spotting and speaker verification tasks

Table 2: The results of single task experiments performed on keyword spotting and speaker verification tasks separately.

Model	# Parameters	Training Task	KWS Accuracy (%)	SV EER (%)
wav2vec 2.0 [14]	96 M	KWS	98.10	-
wav2vec 2.0 [14]	96 M	SV	-	3.35
DistilledWav2vec	27 M	KWS	<b>98.01</b>	-
DistilledWav2vec	27 M	SV	-	4.48
DistilledHuBERT	27 M	KWS	97.82	-
DistilledHuBERT	27 M	SV	-	<b>4.26</b>

to train the networks for various voice-activated tasks.

For keyword spotting (KWS), we use the Google Speech Command (GSC) V1 dataset [26] with 22236 train and 3081 test utterances of 12 classes and provide accuracy results on the test set. Also, we use VoxCeleb-1 dataset [27] with 143642 train and 4874 test utterances for speaker verification (SV) with Equal Error Rate (EER) evaluation metric. In the experiments, we use Adam optimizer with a learning rate  $10^{-4}$ . Note that we also provide the performance of teacher wav2vec 2.0 model as a reference since our aim is to achieve the performance of such a powerful and large network with much smaller networks.

### 3.2. Single Task Results

In this section, we provide our experimental results regarding the single task setting. We construct two different student networks as DistilledWav2vec and DistilledHuBERT from different teachers. We train them separately for single tasks of KWS and SV. Note that this is different from the multi-task training scheme described in Section 2.3. Here, the networks are trained for only one task at each time.

We provide the evaluation results on test sets in Table 2. As seen from the accuracy and EER values, the student models perform well on both KWS and SV tasks. Particularly, we observe only 0.1% accuracy degradation on the KWS task with DistilledWav2vec compared to the teacher wav2vec 2.0. Moreover, we obtain only 0.9% EER degradation on SV with DistilledHuBERT even though the student networks' size is only 28% of the large teacher network.

In this set of experiments on single task, we show that the knowledge distillation method can be efficiently applied to various networks to reduce the model size while maintaining the performance. However, these models are trained to work on only one task. In the next section, we aim to achieve similar competitive results on multiple tasks with a single network.

### 3.3. Multi-Task Results

We perform multi-task training on the student networks DistilledWav2vec and DistilledHuBERT. In this case, we train the networks with KWS and SV tasks simultaneously where the resulting networks can be used in devices with multi-task applications efficiently.

We obtain the evaluation results as in Table 3. In this setting, training the complete network (both SRL module and downstream heads) instead of freezing the SRL module provides 8.41% and 6.63% accuracy gains on KWS, respectively for DistilledWav2vec and DistilledHuBERT models. Similarly on the SV task, these models obtain 10.42% and 8.40% EER improvements due to the discussed multi-task training scheme. Based on these accuracy and EER values, we observe that entire fine-tuned models perform better than partial fine-tuned ver-

Table 3: The results of multi-task experiments performed simultaneously for keyword spotting and speaker verification.

Model	# Parameters	KWS Accuracy (%)	SV EER (%)
wav2vec 2.0 [14]	96 M	98.19	3.15
wav2vec 2.0 (frozen) [14]	96 M	90.33	20.5
DistilledWav2vec	27 M	97.70	4.59
DistilledWav2vec (frozen)	27 M	89.29	15.01
DistilledHuBERT	27 M	<b>97.90</b>	<b>4.27</b>
DistilledHuBERT (frozen)	27 M	91.27	12.67

sions in all cases. Hence, we conclude that training not only downstream heads but also SRL modules with the multi-task training scheme significantly increases the overall performance.

Moreover, we observe that both DistilledHuBERT and DistilledWav2vec models that are trained in a multi-task framework provide similar results with their single task comparisons. Particularly, when DistilledHuBERT is trained solely on the KWS task, it achieves 97.82% accuracy and when it is trained solely on SV, it has 4.26% EER as in Table 2. On the other hand, when we fine-tune DistilledHuBERT on both KWS and SV at the same time, we achieve 97.90% accuracy on KWS and 4.27% EER on SV as in Table 3. This shows that, instead of using a different network for each task, we can achieve similar competitive performance results on multiple tasks with a single network by applying fine-tuning to both SRL module and downstream heads in a multi-task training scheme. Hence, the resulting network can be efficiently deployed and used for multiple tasks by achieving high performance with low memory and computational costs.

Finally, we conclude that both DistilledHuBERT and DistilledWav2vec student networks after entire fine-tuning in multi-task training setting reach competitive performances. Particularly for KWS task, the performance gap between teacher and student networks is small. This means that we can match the large network performance with a much smaller model even in a multi-task application framework.

## 4. CONCLUSION

In this paper, we apply a knowledge distillation approach to different self-supervised networks to obtain high speech representation learning performance with reduced network size and illustrate its generalization to different settings. We perform single task and multi-task trainings for various speech recognition applications and provide comprehensive performance comparisons. Particularly, we propose to combine knowledge distillation and multi-task training scheme where the complete network of SRL module and downstream heads are fine-tuned with all task datasets simultaneously to obtain a single network that can be efficiently used for multiple tasks. In our experiments, we show that the performance of large models such as wav2vec 2.0 and HuBERT can be achieved with substantially smaller networks (28% of the original model sizes) by using the training approach presented. Moreover, we show that the performance of single task networks on their particular tasks can be achieved with a single network fine-tuned on multiple tasks. Hence, these multi-task models do not only provide high performance, but also reduce memory and computation costs by allowing a single network to be shared among multiple tasks. The resulting networks are thus potential candidates for deployment on edge AI devices while offering competitive performance.

## 5. References

- [1] A. T. Liu, S.-w. Yang, P.-H. Chi, P.-c. Hsu, and H.-y. Lee, "Mockingjay: Unsupervised speech representation learning with deep bidirectional transformer encoders," in *IEEE ICASSP*, 2020.
- [2] T. Sercu, G. Saon, J. Cui, X. Cui, B. Ramabhadran, B. Kingsbury, and A. Sethy, "Network architectures for multilingual speech representation learning," in *IEEE ICASSP*, 2017.
- [3] Y.-A. Chung, W.-N. Hsu, H. Tang, and J. Glass, "An Unsupervised Autoregressive Model for Speech Representation Learning," in *Proc. Interspeech*, 2019, pp. 146–150.
- [4] S. Latif, R. Rana, S. Khalifa, R. Jurdak, J. Qadir, and B. W. Schuller, "Survey of deep representation learning for speech emotion recognition," *IEEE Transactions on Affective Computing*, pp. 1–1, 2021.
- [5] G. Trigeorgis, F. Ringeval, R. Brueckner, E. Marchi, M. A. Nicolaou, B. Schuller, and S. Zafeiriou, "Adieu features? End-to-end speech emotion recognition using a deep convolutional recurrent network," in *IEEE ICASSP*, 2016, pp. 5200–5204.
- [6] 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*, vol. 33, 2020.
- [7] 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*, vol. 29, 2021.
- [8] A. T. Liu, S.-W. Li, and H.-y. Lee, "Tera: Self-supervised learning of transformer encoder representation for speech," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 2351–2366, 2021.
- [9] 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 ICASSP*, 2022, pp. 6147–6151.
- [10] A. Baevski, S. Schneider, and M. Auli, "vq-wav2vec: Self-supervised learning of discrete speech representations," in *International Conference on Learning Representations*, 2020.
- [11] L. Pepino, P. Riera, and L. Ferrer, "Emotion Recognition from Speech Using wav2vec 2.0 Embeddings," in *Proc. Interspeech*, 2021, pp. 3400–3404.
- [12] E. Morais, R. Hoory, W. Zhu, I. Gat, M. Damasceno, and H. Aronowitz, "Speech emotion recognition using self-supervised features," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 6922–6926.
- [13] N. Vaessen and D. A. Van Leeuwen, "Fine-tuning wav2vec2 for speaker recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 7967–7971.
- [14] S. Hussain, V. Nguyen, S. Zhang, and E. Visser, "Multi-task voice activated framework using self-supervised learning," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 6137–6141.
- [15] D.-H. Kim, J.-H. Lee, J.-H. Mo, and J.-H. Chang, "W2V2-Light: A Lightweight Version of Wav2vec 2.0 for Automatic Speech Recognition," in *Proc. Interspeech*, 2022.
- [16] J. Gou, B. Yu, S. J. Maybank, and D. Tao, "Knowledge distillation: A survey," *International Journal of Computer Vision*, vol. 129, no. 6, pp. 1789–1819, 2021.
- [17] A. Gresse, M. Quillot, R. Dufour, and J.-F. Bonastre, "Learning Voice Representation Using Knowledge Distillation for Automatic Voice Casting," in *Proc. Interspeech*, 2020, pp. 160–164.
- [18] G. Kurata and G. Saon, "Knowledge Distillation from Offline to Streaming RNN Transducer for End-to-End Speech Recognition," in *Proc. Interspeech*, 2020, pp. 2117–2121.
- [19] Z. Peng, A. Budhkar, I. Tuil, J. Levy, P. Sobhani, R. Cohen, and J. Nassour, "Shrinking bigfoot: Reducing wav2vec 2.0 footprint," in *Proceedings of the Second Workshop on Simple and Efficient Natural Language Processing*. Virtual: Association for Computational Linguistics, Nov. 2021, pp. 134–141.
- [20] R. Wang, Q. Bai, J. Ao, L. Zhou, Z. Xiong, Z. Wei, Y. Zhang, T. Ko, and H. Li, "LightHuBERT: Lightweight and Configurable Speech Representation Learning with Once-for-All Hidden-Unit BERT," in *Proc. Interspeech*, 2022.
- [21] A. Pasad, J.-C. Chou, and K. Livescu, "Layer-wise analysis of a self-supervised speech representation model," in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2021, pp. 914–921.
- [22] H.-J. Chang, S.-w. Yang, and H.-y. Lee, "DistilHuBERT: Speech Representation Learning by Layer-Wise Distillation of Hidden-Unit BERT," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.
- [23] S. wen 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, T.-H. Huang, W.-C. Tseng, K. tik Lee, D.-R. Liu, Z. Huang, S. Dong, S.-W. Li, S. Watanabe, A. Mohamed, and H. yi Lee, "SUPERB: Speech Processing Universal PERFORMANCE Benchmark," in *Proc. Interspeech*, 2021, pp. 1194–1198.
- [24] X. Chang, T. Maekaku, P. Guo, J. Shi, Y.-J. Lu, A. S. Subramanian, T. Wang, S.-w. Yang, Y. Tsao, H.-y. Lee, and S. Watanabe, "An exploration of self-supervised pretrained representations for end-to-end speech recognition," in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2021, pp. 228–235.
- [25] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *IEEE ICASSP*, 2015, pp. 5206–5210.
- [26] P. Warden, "Speech commands: A dataset for limited-vocabulary speech recognition," *arXiv preprint arXiv:1804.03209*, 2018.
- [27] A. Nagrani, J. S. Chung, and A. Zisserman, "VoxCeleb: A Large-Scale Speaker Identification Dataset," in *Proc. Interspeech*, 2017, pp. 2616–2620.