



Deep versus Wide: An Analysis of Student Architectures for Task-Agnostic Knowledge Distillation of Self-Supervised Speech Models

Takanori Ashihara, Takafumi Moriya, Kohei Matsuura, Tomohiro Tanaka

NTT Corporation, Japan

takanori.ashihara.vk@hco.ntt.co.jp

Abstract

Self-supervised learning (SSL) is seen as a very promising approach with high performance for several speech downstream tasks. Since the parameters of SSL models are generally so large that training and inference require a lot of memory and computational cost, it is desirable to produce compact SSL models without a significant performance degradation by applying compression methods such as knowledge distillation (KD). Although the KD approach is able to shrink the depth and/or width of SSL model structures, there has been little research on how varying the depth and width impacts the internal representation of the small-footprint model. This paper provides an empirical study that addresses the question. We investigate the performance on SUPERB while varying the structure and KD methods so as to keep the number of parameters constant; this allows us to analyze the contribution of the representation introduced by varying the model architecture. Experiments demonstrate that a certain depth is essential for solving content-oriented tasks (e.g. automatic speech recognition) accurately, whereas a certain width is necessary for achieving high performance on several speaker-oriented tasks (e.g. speaker identification). Based on these observations, we identify, for SUPERB, a more compressed model with better performance than previous studies.

Index Terms: self-supervised learning, knowledge distillation, speech representation

1. Introduction

Self-supervised learning (SSL) has become the key technique not only for natural language processing (NLP) and computer vision (CV) communities, but also for the speech community. This approach offers a general-purpose representation learned in an unsupervised fashion and achieves state-of-the-art performance in many downstream tasks. Particularly in speech representation learning, many SSL studies have been published [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11]. They detail impressive performance and capability in solving the multiple tasks associated with the Speech processing Universal PERformance Benchmark (SUPERB) [12]. While SSL is successful, the pre-trained models generally have high computation costs for training and inference due to huge over-parameterization, so they cannot be deployed on mobile devices with limited resources.

To tackle the above problem, knowledge distillation (KD) [13], which is a model compression method, has been used to transfer the knowledge from a large speech SSL model into a small, green model by following a teacher-student framework [14, 15, 16]. There are two main approaches for KD with SSL models: task-specific distillation and task-agnostic distillation. The former is based on distilling the knowledge of a teacher model that has been fine-tuned for a specific-task. For example, in [14, 15], an SSL model fine-tuned for the automatic speech recognition (ASR) task has been utilized to reduce the

parameters without significant ASR performance degradation. On the other hand, in task-agnostic distillation, an internal representation of a pre-trained teacher model is distilled to provide a universal small student model such as DistillHuBERT [16]. We focus on the latter approach in this paper because our aim is to benefit from its ability to solve multiple speech tasks with the same feature extraction module [12]. In addition, task-specific distillation requires the preparation of teacher models fine-tuned to each task and so is more costly than task-agnostic distillation.

When applying KD on an internal representation of itself, since the structure of deep neural networks (DNN) has the core property of accepting scaling with respect to its depth and width, the student model can be thinner and shallower than the teacher model which reduces model size. However, there has been little understanding on how varying the depth and width of the small-footprint SSL model affects the internal speech representation and performance in downstream tasks. In other words, do models with different structures distilled from the same teacher learn similar representations and achieve comparable performance? Investigating this fundamental issue could form the basis for developing more resource-efficient networks such as EfficientNet [17] in CV.

In this paper, we empirically explore the question how to shrink student models in terms of depth and width. First, we use SUPERB to examine the performance realized while varying the depth and width of the student networks distilled from HuBERT_{BASE}; the total number of parameters is kept the same as much as possible. According to the result, we find the performance tendency that the larger the model size, the better the performance as reported in [18], and furthermore, that the deep&narrow student is better than the shallow&wide one in content-oriented tasks and vice versa in several of speaker-oriented tasks. We also confirm that setting KD between student's last and teacher's intermediate layers such as DistillHuBERT is suitable for wide networks, whereas KD between intermediate layers is suitable for deep networks such as [19, 20]. Additionally, we show a similar trend in the student model distilled from the HuBERT_{LARGE} model. This analysis contributes to practical requirements of speech applications because student model size remains the same even if the teacher size is increased. Beyond this analysis, the smaller model that has an intermediate number of layers shows better performance than the previous method through the linear interpolation of prediction-layer and intermediate-layer losses.

2. Related Work

Distilling the representation from large to small task-agnostic models has attracted much attention in the NLP community, especially. DistilBERT [21] set soft target distillation and cosine similarity losses between last layers. Other works such as MobileBERT [22] and TinyBERT [23] transferred knowledge by mapping intermediate layers. More recently, MiniLM [24, 25]

was proposed to attain a shallower architecture by carefully-designed self-attention KD. Unfortunately, to the best of our knowledge, only DistilHuBERT has been proposed as a task-agnostic KD of speech representation. These works focused on developing new effective KD methods, and little attention was paid to the effects of structures on the representations of students, especially for speech representation.

Some previous works have offered analyses of SSL representations [26, 27, 28, 29, 30, 31]. For example, [26] provided a similarity analysis between last-layer output of SSL methods/architectures. In [29, 30, 31], the contribution of features was measured for various probing tasks across layers. Their analysis provided sophisticated insights, but not much was elucidated about how the variations of fundamental network architecture yielded affected outputs.

There have been several empirical studies related to deep versus wide network structures. In the context of speech, shallow networks achieve similar performance to deep networks when mimicking the teacher’s posterior in phoneme recognition [32], but deep structures trained by KD between intermediate-layers are crucial for performance in ASR tasks [19]. Furthermore, [33] has shown that representation differences based on the number of layers seems to be related to robustness against noise. In the CV community, [34] has recently found that deep and wide models yield different characteristic error patterns but similar recognition accuracies. Our study is inspired by these results as regards the distinctive representation between deep and wide architectures, and hence, we attempt to unveil the impact of varying the fundamental network structures (i.e. depth and width) on speech representation/performance.

Other work related to this paper includes [18], which investigated the relationship between the number of parameters and training loss of Mockingjay [6] as well as between the number of parameters and performance on speaker recognition and phoneme classification tasks. They demonstrated positive correlations between model size and loss/performance and exhibited power-law scaling nature on loss. We follow this work by investigating the relationship between the variation of core model architecture and its performance. In addition, we also focus on KD as a more efficient and practical way to train tiny models and also conduct more extensive experiments based on SUPERB which is one of the general benchmarks for speech representation.

3. Method

In this paper, for the purpose of analyzing how the representations of task-agnostic KD models vary with depth and width empirically, we apply two simple KD approaches: prediction-layer distillation and layer-to-layer (L2L) distillation methods. While the former is based on DistilHuBERT [16], the latter is based on representation mapping layer-to-layer such as FitNets [19] to alleviate the difficulty of training deeper networks.

For the whole network architecture, the teacher and student models are composed of 7-layers of convolutional neural networks (CNNs) with input of a raw-waveform, followed by multiple self-attention layers, as is commonly used in SSL studies. For student models, we alter the depth and width of self-attention layers only while fixing the size of the CNNs.

3.1. Prediction-layer distillation

As illustrated in Figure 1, prediction-layer distillation transfers the representation by mapping between the last-layer of student and intermediate-layers of teacher models. Since the effective teacher’s representations for each downstream task are distributed layer-wise, DistilHuBERT [16] utilizes the hidden

states of empirically-selected multiple self-attention layers as teacher knowledge, and attains better performance than KD by utilizing only the last-layer of teacher models. To project into the same dimensional space collected from multiple hidden states from the teacher, the prediction heads are added on the top layer of the student model and hence, the training objective is multi-task learning. When we evaluate the student models, the heads are removed so as to follow DistilHuBERT. In contrast to the original paper, we do not copy the parameters of the teacher model into the student model for initialization, except for convolutional layers, because we investigate not only shallow but also narrow models in which the hidden dimension size differs from those of larger teacher models.




Figure 1: Illustration of student model trained by KD between student’s last and teacher’s intermediate layers based on DistilHuBERT [16].

3.2. Layer-to-layer (L2L) distillation

To train the deeper lightweight model efficiently, we also perform simple L2L distillation in addition to the former distillation as shown in Figure 2. This approach, inspired by FitNets, chooses the intermediate-layers from the teacher model as in DistilHuBERT, but transfers the knowledge to the student model by layer-to-layer mapping. In the case that the intermediate layers from the narrow student yield smaller dimensional output than the teacher model, L2L KD also connects the projection layers to each intermediate layer. The projection layers are removed when evaluation is conducted. In L2L KD, we initialize only the student CNN block from the teacher model in the same way as the former KD and map the student to the teacher self-attention layers.




Figure 2: Illustration of student model trained by KD between intermediate-layers such as FitNets [19].

3.3. Distillation objective

In this paper, as the objective function, we employ a combination of maximizing the L1 loss and minimizing the cosine similarity loss explored in DistilHuBERT. The contribution of cosine similarity can be controlled by the weight factor (Equation (1) in [16]), but we set their contribution equal in all experiments.

4. Experimental Setup

The teacher and student models explored in this paper are summarized in Table 2. As teachers, we utilize HuBERT BASE and LARGE, both of which are publicly available from fairseq¹

¹<https://github.com/pytorch/fairseq>

Table 1: Evaluation result for each model distilled from HuBERT *BASE* and each task on SUPERB. The values in the first and second row are taken from [12] and [16], respectively. *Pred.* means the predication-layer distillation and L2L indicates the layer-to-layer distillation in the second column. For clarity, the KD models are indexed from (a) to (h) as shown in the second column.

Model	KD Loss	PR	ASR (w/ LM)	KS	QbE	SID	ASV	SD	IC	SF	ER	Rank↓
		PER↓	WER↓	Acc↑	MTWV↑	Acc↑	EER↓	DER↓	Acc↑	F1↑ / CER↓	Acc↑	
HuBERT <i>BASE</i>	-	5.41	6.42 (4.79)	96.30	0.0736	81.42	5.11	5.88	98.34	88.53 / 25.20	64.92	1.7
DistilHuBERT	Pred.	16.27	13.34 (9.21)	95.98	0.0511	73.54	8.55	6.19	94.99	82.57 / 35.59	63.02	5.8
12-L <i>HALF</i>	(a) Pred.	13.09	11.87 (8.07)	96.97	0.0501	69.11	6.32	6.67	94.91	84.49 / 32.54	62.76	4.6
	(b) L2L	10.67	10.96 (7.68)	97.24	0.0604	69.52	6.13	6.81	96.97	86.11 / 30.93	63.24	2.6
12-L <i>FOURTH</i>	(c) Pred.	18.92	14.02 (9.25)	96.44	0.0495	49.51	6.74	7.12	87.03	81.21 / 37.27	62.82	8.1
	(d) L2L	16.96	13.84 (9.20)	96.40	0.0562	47.67	6.41	7.12	91.62	84.81 / 32.77	61.84	7.0
3-L <i>ONE</i>	(e) Pred.	13.34	12.23 (8.64)	96.69	0.0489	75.71	6.48	6.56	94.15	82.89 / 34.65	63.95	4.6
	(f) L2L	13.96	12.94 (9.11)	96.52	0.0568	47.76	6.18	7.17	96.02	85.99 / 32.38	62.57	5.2
3-L <i>HALF</i>	(g) Pred.	18.62	13.91 (9.27)	96.22	0.0482	62.59	6.86	6.69	91.88	82.78 / 35.75	61.83	8.1
	(h) L2L	18.11	14.48 (9.86)	96.48	0.0502	60.40	6.82	7.31	94.91	81.82 / 37.36	62.78	7.5

[35]. To build lightweight models, for the sake of simplicity, we reduced only the width of self-attention layers by two ratios, specifically one half and one fourth. For the purpose of comparing the depth and width while keeping the total number of parameters constant, 12-L *HALF* and 12-L *FOURTH* were trained to compare 3-L *ONE* and 3-L *HALF*, respectively, which are comparable in terms of the number of parameters. Note that when training the 3-L *ONE* model with L2L KD, we did not add the projection layer because the teacher and student self-attention layers had the same dimension originally. 6-L *HALF* is utilized for further analysis as discussed later.

All student models were trained using the same settings as DistilHuBERT on a single GPU. Specifically, the training steps totaled 200k with 24 batch size, and the models were optimized with Adam; the learning rate was linearly increased to $2e-4$ for the first 7% of updates, and then linearly decayed. As the training dataset, we employed the standard Librispeech 960h [36] without labels. When applying KD, the student models are mapped with the teacher representation by taking all layers for 12-L *HALF* and *FOURTH*, by taking 4th, 8th and 12th layers for 3-L *ONE* and *HALF* and by taking one layer out of two for 6-L *HALF*.

Table 2: Model settings of teacher and student models. With respect to self-attention blocks, *HALF* and *FOURTH* means the parameter reductions to one half and one fourth, respectively, and *ONE* is the same as the HuBERT *BASE*.

Models	#Params	#Layers	Embed.	FFN	#Head
HuBERT <i>BASE</i> [9]	94.68M	12	768	3072	12
HuBERT <i>LARGE</i> [9]	316.61M	24	1024	4096	16
DistilHuBERT [16]	23.49M	2	768	3072	12
12-L <i>HALF</i>	26.87M	12	384	1536	6
12-L <i>FOURTH</i>	9.93M	12	192	768	3
3-L <i>ONE</i>	30.58M	3	768	3072	12
3-L <i>HALF</i>	10.90M	3	384	1536	6
6-L <i>HALF</i>	16.23M	6	384	1536	6

We evaluated the task-agnostic small KD models on SUPERB [12]. SUPERB consists of 10 downstream tasks: phoneme recognition (PR), automatic speech recognition (ASR), keyword spotting (KS), query by example spoken term detection (QbE), speaker identification (SID), automatic speaker verification (ASV), speaker diarization (SD), intent classification (IC), slot filling (SF) and emotion recognition (ER). These tasks are categorized into content, speaker, semantics and paralinguistics aspects as described in [12]. This bench-

mark freezes the parameters of SSL models to be utilized as a feature extractor and weighted-sums the features collected from hidden layers except for QbE task. The above training and evaluation processes were implemented with S3PRL² [5, 6].

5. Results and Analysis

5.1. How depth and width affects the performance

Here, we analyze the relationship between the variation in student architectures with two KD methods and the resulting performances on SUPERB. Table 1 shows the KD results by using the knowledge of HuBERT *BASE*. The ranking scores in the rightmost column indicate the relative values, which are the averages of each rank on the downstream tasks, as used in [16]. First of all, regarding the variation in KD loss function, prediction-layer loss is suitable for wider architectures such as (e) (the 7th row of the Table 1), whereas L2L loss is effective for deeper architectures such as (b) and (d). By comparing (b) to (e) and (d) to (g, h), which have almost the same model size but different structures, we observe that while deeper networks have higher performance in content-oriented tasks such as PR, ASR and QbE, wider networks have higher performance in speaker-oriented tasks such as SID and SD. Note that 3-L *ONE* and *HALF* were trained by imitating the representation suitable for the ASR task explicitly by heuristically selecting the layers whose representations are weighted for the ASR task (e.g. 8th layer on HuBERT *BASE*) [16]. From the above results, deep models seem to be effective when the task to be solved is based on contents while wide models support tasks related to speakers. This is possible because the context size of self-attention is able to increase with the number of layers, and thus the deep architecture is advantageous for content-oriented tasks that require long context. On the other hand, increasing the representation capacity within lower layers is seemingly effective for speaker-oriented tasks that focus local acoustic features rather than contextual information. As shown in [16], it certainly demonstrates that effective features for SID (ASR) task are concentrated in the lower (upper) layers. We also confirm a clear tradeoff between the number of parameters and overall performance which agrees with [18, 37].

5.2. Is the tendency the same for a larger teacher model?

From Section 5.1, the student models emphasize the characteristic performance in the case of HuBERT *BASE*, do the models distilled from HuBERT *LARGE* show a similar property? These results are important for practical concerns because student size

²<https://github.com/s3prl/s3prl>

Table 3: Evaluation result for each model distilled from HuBERT LARGE. The values of first row are taken from [12]. The values shown from the second row are the results of the KD models trained in our experiment.

Model	KD Loss	PR	ASR (w/ LM)	KS	QbE	SID	ASV	SD	IC	SF	ER	Rank↓
		PER↓	WER↓	Acc↑	MTWV↑	Acc↑	EER↓	DER↓	Acc↑	F1↑ / CER↓	Acc↑	
HuBERT LARGE	-	3.53	3.62 (2.94)	95.29	0.0353	90.33	5.98	5.75	98.76	89.81 / 21.76	67.62	2.7
12-L HALF	Pred.	9.67	9.59 (6.84)	95.79	0.0507	49.25	5.84	6.20	95.07	84.88 / 31.17	63.59	4.2
	L2L	7.97	9.24 (6.82)	96.24	0.0513	52.42	6.36	6.60	96.92	87.26 / 28.92	64.51	3.2
12-L FOURTH	Pred.	14.10	12.49(8.47)	96.20	0.0482	37.18	6.86	6.93	91.91	83.66/35.11	62.45	6.8
	L2L	12.86	12.91(9.11)	95.34	0.0443	47.51	7.26	7.05	92.86	83.83/34.22	62.20	7.4
3-L ONE	Pred.	12.11	11.35 (8.00)	96.50	0.0474	76.97	7.22	6.61	96.63	85.36 / 31.60	65.80	3.7
	L2L	10.24	12.23 (8.78)	96.40	0.0540	68.90	7.59	7.33	96.97	84.56 / 32.88	65.22	4.3
3-L HALF	Pred.	15.78	13.28(9.34)	96.20	0.0430	60.17	7.17	6.77	94.02	84.67 / 33.82	64.55	5.9
	L2L	15.11	14.31 (9.84)	96.01	0.0532	55.35	7.47	7.81	92.72	84.04 / 34.33	63.40	6.9

Table 4: Evaluation result for each model distilled from HuBERT BASE. The values in the fifth row represent the model trained by the linear interpolation loss (Comb.) between the prediction-layer and L2L losses.

Model	KD Loss	PR	ASR (w/ LM)	KS	QbE	SID	ASV	SD	IC	SF	ER	Rank↓
		PER↓	WER↓	Acc↑	MTWV↑	Acc↑	EER↓	DER↓	Acc↑	F1↑ / CER↓	Acc↑	
HuBERT BASE	-	5.41	6.42 (4.79)	96.30	0.0736	81.42	5.11	5.88	98.34	88.53 / 25.20	64.92	1.3
DistilHuBERT	Pred.	16.27	13.34 (9.21)	95.98	0.0511	73.54	8.55	6.19	94.99	82.57 / 35.59	63.02	4.1
6-L HALF	Pred.	15.14	12.72 (8.68)	96.85	0.0504	67.06	6.36	6.81	93.75	83.65 / 34.35	63.72	3.4
	L2L	13.40	12.66 (8.59)	96.38	0.0545	62.90	6.85	6.95	95.86	83.80 / 33.51	63.09	3.3
	Comb.	14.68	12.43 (8.51)	96.77	0.0516	65.75	6.81	6.83	94.57	84.32 / 33.99	64.78	2.9

is unaffected by the teacher size. Table 3 shows the results, which are similar to Table 1, but different in that the teacher model is HuBERT LARGE, resulting students with better performance on PR, ASR and SF tasks in particular. Note that the 12-L and 3-L models mimic the teacher representations of the even-numbered and multiple of eight layers, respectively. Although we can find a performance tendency that follows that in Section 5.1 on the PR, ASR and SID tasks, the students appear to be inconsistent with the former trends in some tasks such as QbE and KS. However, identifying the reason is difficult because the performance deterioration is apparent in the original application of HuBERT LARGE compared with HuBERT BASE on these tasks, so HuBERT LARGE is considered to have learned a comparatively ineffective representation in such tasks.

Table 5: Result for each shallow model. Pred.-all and L2L-n-of-m are KD method of utilizing all HuBERT BASE layers.

Model	KD Loss	PR	ASR	SID	IC
		PER↓	WER↓	Acc↑	Acc↑
3-L ONE	Pred.	13.34	12.23	75.71	94.15
	Pred.-all	13.85	12.73	71.37	94.33
	L2L	13.96	12.94	47.76	96.02
	L2L-n-of-m	14.64	13.41	44.56	95.57
3-L HALF	Pred.	18.62	13.91	62.59	91.88
	Pred.-all	18.68	14.18	59.73	92.17
	L2L	18.11	14.48	60.40	94.91
	L2L-n-of-m	18.30	14.46	36.57	94.44

5.3. Is it due to lack of information from the teacher?

The above results could be due to not capturing all teacher representations. In other words, only selected layers are utilized for teacher prediction for KD. Therefore, we adopted an additional training approach of utilizing all teacher layers. Although DistilHuBERT makes it easy to scale the number of teachers layers by adding the prediction heads to student models (called Pred.-all), it is difficult to achieve with L2L KD. To alleviate this problem, some attempts have been made for efficient intermediate layer mapping [38, 39, 40] in the NLP community;

we adopt the simpler method inspired by the RAIL-KD [40] for our analysis. When applying L2L KD, this method randomly selects as many layers as the student model from all layers of the teacher model at each batch (referred to as L2L-n-of-m). Table 5 compares the performance on PR, ASR, SID and IC for two shallow models (3-L ONE and HALF), trained by different KD losses. The results show there is no significant improvement in performance for both methods. Perhaps there is a need to simply advance the KD method which is better-suited for speech representation. However, in our experiment, the severe degradation in PR/ASR performance seems not to be due to omission in teacher layer selection but to structures.

5.4. Can a student that has an intermediate number of layers offers universal performance?

The above results can be summarized as narrow or shallow models have advantages and disadvantages in terms of performance. Therefore, we built a KD model consisting of an intermediate number of layers, specifically 6-layers. Table 4 presents the performance of KD models that were transferred knowledge from HuBERT BASE. To construct the more universal model, we prepared a combination model trained by the linear interpolation loss between prediction-layer and L2L distillation losses, which corresponds to the 5th row in Table 4. In this experiment, we decided to weight the former loss term by 0.8 and the latter term by 0.2 based on our preliminary experiment. The result is that our model offers not only better performance but also greater compression than DistilHuBERT.

6. Conclusions

In this paper, we empirically investigated and analyzed how varying the depth and width of small network structures impacted the speech representation formed by task-agnostic KD from a large SSL model. Our findings indicate that a deep&narrow student is better than the shallow&wide equivalent in content-oriented tasks, while the reverse is true for several speaker-oriented tasks. We also developed and tested a smaller model with better performance than the previous approach. In future work, we will investigate other benchmarks such as SUPERB-SG [41] to confirm these findings.

7. References

- [1] A. van den Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [2] Y.-A. Chung, W.-N. Hsu, H. Tang, and J. Glass, "An unsupervised autoregressive model for speech representation learning," in *Interspeech*, 2019.
- [3] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Unsupervised pre-training for speech recognition," in *Interspeech*, 2019.
- [4] S. Pascual, M. Ravanelli, J. Serrà, A. Bonafonte, and Y. Bengio, "Learning problem-agnostic speech representations from multiple self-supervised tasks," in *Interspeech*, 2019.
- [5] A. T. Liu, S.-W. Li, and H. yi Lee, "TERA: Self-supervised learning of transformer encoder representation for speech," *arXiv preprint arXiv:2007.06028*, 2020.
- [6] 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 *ICASSP*, 2020.
- [7] S. Ling and Y. Liu, "Decoar 2.0: Deep contextualized acoustic representations with vector quantization," *arXiv preprint arXiv:2012.06659*, 2020.
- [8] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *NeurIPS*, 2020.
- [9] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "HuBERT: Self-supervised speech representation learning by masked prediction of hidden units," *TASLP*, 2021.
- [10] 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," *arXiv preprint arXiv:2110.13900*, 2021.
- [11] A. Baevski, W.-N. Hsu, Q. Xu, A. Babu, J. Gu, and M. Auli, "data2vec: A general framework for self-supervised learning in speech, vision and language," *arXiv preprint arXiv:2202.03555*, 2022.
- [12] S. wen Yang, P.-H. Chi, Y.-S. Chuang, C.-I. J. Lai, K. Lakhota, 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 *Interspeech*, 2021.
- [13] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [14] Z. Peng, A. Budhkar, I. Tuil, J. Levy, P. Sobhani, R. Cohen, and J. Nassour, "Shrinking bigfoot: Reducing wav2vec 2.0 footprint," in *SustainNLP*, 2021.
- [15] X. Yang, Q. Li, and P. C. Woodland, "Knowledge distillation for neural transducers from large self-supervised pre-trained models," *arXiv preprint arXiv:2110.03334*, 2021.
- [16] H.-J. Chang, S. wen Yang, and H. yi Lee, "DistilHuBERT: Speech representation learning by layer-wise distillation of hidden-unit BERT," *arXiv preprint arXiv:2110.01900*, 2021.
- [17] M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *ICML*, 2019.
- [18] J. Pu, Y. Yang, R. Li, O. Elibol, and J. Droppo, "Scaling effect of self-supervised speech models," in *Interspeech*, 2021.
- [19] A. Romero, N. Ballas, S. E. Kahou, A. Chassang, C. Gatta, and Y. Bengio, "FitNets: Hints for thin deep nets," in *ICLR*, 2015.
- [20] T. Moriya, H. Kanagawa, K. Matsui, T. Fukutomi, Y. Shinohara, Y. Yamaguchi, M. Okamoto, and Y. Aono, "Efficient building strategy with knowledge distillation for small-footprint acoustic models," in *SLT*, 2018.
- [21] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter," in *EMC2-NeurIPS*, 2019.
- [22] Z. Sun, H. Yu, X. Song, R. Liu, Y. Yang, and D. Zhou, "MobileBERT: a compact task-agnostic BERT for resource-limited devices," in *ACL*, 2020.
- [23] X. Jiao, Y. Yin, L. Shang, X. Jiang, X. Chen, L. Li, F. Wang, and Q. Liu, "TinyBERT: Distilling BERT for natural language understanding," in *EMNLP*, 2020.
- [24] W. Wang, F. Wei, L. Dong, H. Bao, N. Yang, and M. Zhou, "MiniLM: Deep self-attention distillation for task-agnostic compression of pre-trained transformers," in *NeurIPS*, 2020.
- [25] W. Wang, H. Bao, S. Huang, L. Dong, and F. Wei, "MiniLMv2: Multi-head self-attention relation distillation for compressing pre-trained transformers," in *ACL-IJCNLP*, 2021.
- [26] Y.-A. Chung, Y. Belinkov, and J. Glass, "Similarity analysis of self-supervised speech representations," in *ICASSP*, 2021.
- [27] S. wen Yang, A. T. Liu, and H. yi Lee, "Understanding self-attention of self-supervised audio transformers," in *Interspeech*, 2020.
- [28] D. Ma, N. Ryant, and M. Liberman, "Probing acoustic representations for phonetic properties," in *ICASSP*, 2021.
- [29] A. Pasad, J. Chou, and K. Livescu, "Layer-wise analysis of a self-supervised speech representation model," in *ASRU*, 2021.
- [30] J. Shah, Y. K. Singla, C. Chen, and R. R. Shah, "What all do audio transformer models hear? Probing acoustic representations for language delivery and its structure," *arXiv preprint arXiv:2101.00387*, 2021.
- [31] 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 *ASRU*, 2021.
- [32] J. Ba and R. Caruana, "Do deep nets really need to be deep?" in *NeurIPS*, 2014.
- [33] O. Vinyals and N. Morgan, "Deep vs. wide: Depth on a budget for robust speech recognition," in *Interspeech*, 2013.
- [34] T. Nguyen, M. Raghu, and S. Kornblith, "Do wide and deep networks learn the same things? Uncovering how neural network representations vary with width and depth," in *ICLR*, 2021.
- [35] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, "fairseq: A fast, extensible toolkit for sequence modeling," in *NAACL-HLT*, 2019.
- [36] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *ICASSP*, 2015.
- [37] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, "Scaling laws for neural language models," *arXiv preprint arXiv:2001.08361*, 2020.
- [38] P. Passban, Y. Wu, M. Rezagholizadeh, and Q. Liu, "ALP-KD: Attention-based Layer Projection for Knowledge Distillation," in *AAAI*, 2021.
- [39] Y. Wu, P. Passban, M. Rezagholizadeh, and Q. Liu, "Why skip if you can combine: A simple knowledge distillation technique for intermediate layers," in *EMNLP*, 2020.
- [40] M. A. Haidar, N. Anchuri, M. Rezagholizadeh, A. Ghaddar, P. Langlais, and P. Poupart, "RAIL-KD: RANdom Intermediate Layer mapping for Knowledge Distillation," *arXiv preprint arXiv:2109.10164*, 2021.
- [41] H.-S. Tsai, H.-J. Chang, W.-C. Huang, Z. Huang, K. Lakhota, S. wen Yang, S. Dong, A. T. Liu, C.-I. J. Lai, J. Shi, X. Chang, P. Hall, H.-J. Chen, S.-W. Li, S. Watanabe, A. Mohamed, and H. yi Lee, "SUPERB-SG: Enhanced Speech processing Universal PERFORMANCE Benchmark for Semantic and Generative Capabilities," *arXiv preprint arXiv:2203.06849*, 2022.