



DiceHuBERT: Distilling HuBERT with a Self-Supervised Learning Objective

Hyung Gun Chi¹, Zakaria Aldeneh¹, Tatiana Likhomanenko¹, Oggi Rudovic¹, Takuya Higuchi¹,
Li-Wei Chen², Shinji Watanabe², Ahmed Hussen Abdelaziz^{3,†}

¹Apple, ²Carnegie Mellon University, ³Meta

hchi23@apple.com

Abstract

We introduce DiceHuBERT, a knowledge distillation framework for compressing HuBERT, a widely used self-supervised learning (SSL)-based speech foundation model. Unlike existing distillation methods that rely on layer-wise and feature-wise mapping between teacher and student models, DiceHuBERT leverages HuBERT’s iterative self-distillation mechanism by directly replacing the original model with a student model. This replacement allows the student to be trained using the same SSL objective used when pre-training HuBERT, eliminating the need for additional modules or architectural constraints. Experimental results on SUPERB show that DiceHuBERT consistently outperforms existing distillation methods, improving phoneme recognition performance by over 21% and ASR performance by more than 14%. Furthermore, DiceHuBERT demonstrates competitive performance across multiple tasks, highlighting its clear advantage.

Index Terms: Knowledge Distillation, Speech Foundation Model, Self-supervised Learning

1. Introduction

Self-supervised learning (SSL)-based speech foundation models [1–8] have received significant attention due to their ability to generate general-purpose speech representations for a wide range of speech tasks. However, the substantial size of the SSL-based foundation models makes them unsuitable for deployment in mobile applications [9, 10]. Knowledge distillation [11–13] is a widely adopted technique in various domains [14–19] that effectively compresses the model size while maintaining performance. This compression is achieved by transferring knowledge from a larger “teacher” model to a smaller “student” model.

Recent efforts [20–25] in knowledge distillation for SSL-based speech foundation models, particularly HuBERT [2], have focused on compressing architectures and introducing new loss functions. These methods typically distill knowledge by aligning layer-wise features between the teacher and student models. However, they often overlook the original SSL objective and impose architectural constraints on the student model.

In this paper, we introduce DiceHuBERT, Distilling from cluster Centroids HuBERT, to address the aforementioned limitations. DiceHuBERT is trained using the same SSL objective as the teacher model, with knowledge distillation occurring through targets generated from the k -means [26] centroids of the teacher model’s features, as illustrated in Figure 1 (a). DiceHuBERT is a simple yet effective distillation framework. Unlike existing methods that primarily focus on feature distil-

lation, see Figure 1 (b), DiceHuBERT leverages the iterative self-distillation of HuBERT’s pre-training, where knowledge is progressively distilled from each iteration to the next. Building upon this inherent self-distillation mechanism, *DiceHuBERT is a natural extension of HuBERT that enables efficient compression within the HuBERT framework itself, without the need for additional modules.*

In our approach, the student model architecture is unrestricted because distillation relies solely on the target labels produced by the teacher model. In contrast, feature distillation methods typically require layer-wise or feature-wise mapping between teacher and student representations, which imposes design constraints on the student model. This architectural flexibility, combined with the reuse of the original HuBERT self-distillation framework, makes DiceHuBERT more straightforward, practical, and easier to adopt than previous methods.

We compare our SSL-based distillation with feature distillation methods from existing works on four tasks from SUPERB [27]: phoneme recognition (PR), automatic speech recognition (ASR), speaker identification (SID), and automatic speaker verification (ASV). Our results show that SSL-based distillation consistently outperforms feature distillation across various SUPERB downstream tasks. Specifically, our proposed distillation approach outperforms prior HuBERT distillation methods, achieving over 21% relative improvement in PR and more than 14% relative improvement in ASR. Improvements are also observed in ASV and SID. We further investigate different architectures and label choices for the SSL objective during distillation, offering insights to inform future research in this domain.

2. Preliminaries

2.1. Prior HuBERT Distillation

Most prior works on HuBERT distillation [18, 21, 22, 24] focus on transferring knowledge from the pre-trained second-iteration HuBERT model (used as the teacher) to the student model through layer-wise feature distillation, as shown in Figure 1 (b). The student model usually has the same number of layers as the teacher, enabling a direct correspondence between their layer outputs for distillation, while reducing the feature dimensions. The distillation objective is commonly defined as the mean squared error (MSE) between the teacher’s and student’s features. To align the student (S) feature dimension D^S with the teacher (T) feature dimension D^T , a linear projection layer

Note that Figure 1 (b) illustrates general layer-wise feature distillation, which differs among models. e.g., MaskHuBERT [22] employs masked inputs, DistilHuBERT [23] incorporates prediction heads.

[†]Work done while at Apple.

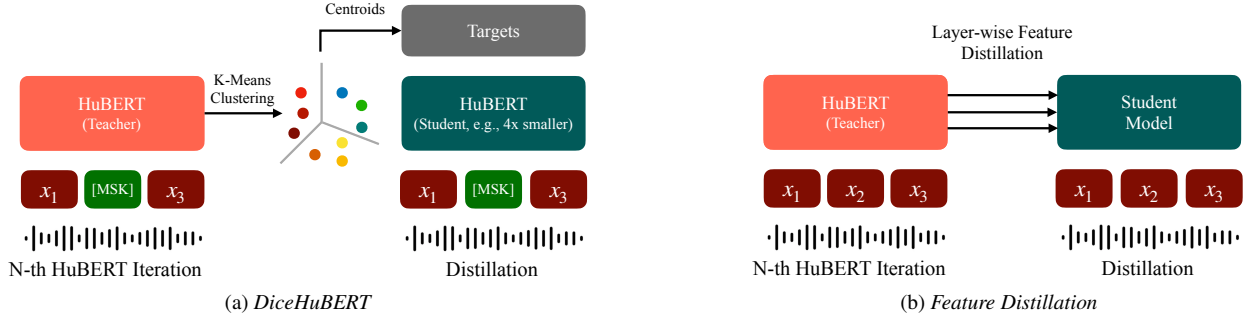


Figure 1: *DiceHuBERT* differs from feature distillation-based methods in that it distills knowledge from target labels generated by a teacher model using *k*-means clustering. This process allows *DiceHuBERT* to be trained using an SSL objective that closely resembles the training objective of the teacher model.

$W_{(l)} \in \mathbb{R}^{D^S \times D^T}$ is employed:

$$\mathcal{L}_{\text{feat}} = \sum_{l \in \{1, \dots, L\}} \alpha_{(l)} \text{MSE}(H_{(l)}^T, H_{(l)}^S W_{(l)}), \quad (1)$$

where $H_{(l)}$ is a feature from l -th layer and $\alpha_{(l)}$ is weight for each layer loss.

Although student model architectures vary, and some works add additional loss functions to enhance alignment, previous work generally follows the above objective for HuBERT distillation.

2.2. HuBERT Pre-training

HuBERT [2] is an SSL speech representation approach that employs a masked token prediction task. The pre-training process involves N iterations, each with distinct targets but the same learning objective predicting the targets of input features corresponding to randomly masked timestamps in the audio. Formally, the SSL loss is defined as the cross-entropy loss computed over the masked timestamps:

$$\mathcal{L}_{\text{SSL}} = - \sum_{t \in M} \log p(z_t | \tilde{X}, t), \quad (2)$$

where M denotes the set of masked timestamps, \tilde{X} is a masked input from a speech utterance $X = [x_1, \dots, x_T]$ using M , and z_t is a target class at timestamp t . The targets are generated through *k*-means clustering: in the first iteration, MFCC features are used; in the following iteration N , typically the 6th layer hidden features from the previous iteration ($N-1$)th model serve as the input for clustering to generate new targets for an N th model. The each following iteration can be viewed as a *self-distillation* process as it uses the pre-trained features from the previous iteration to guide the model’s training, while keeping the model architecture the same with the first iteration.

3. DiceHuBERT

We propose *DiceHuBERT*, a HuBERT distillation framework that offers a simpler and natural distillation approach within HuBERT itself compared to previous methods. *DiceHuBERT* leverages HuBERT’s iterative self-distillation process by reducing the model size when transitioning from iteration N (teacher) to iteration $N+1$ (student), all while maintaining the original SSL objective, see Figure 1 (a). *DiceHuBERT* offers two main advantages over existing HuBERT distillation methods. Firstly, it eliminates the need for an additional distillation framework,

allowing us to directly reuse the pre-training framework of HuBERT without implementing any new modules. As a result, *DiceHuBERT* is architecture-agnostic, providing flexibility in choosing the student model architecture. The feature distillation approach’s reliance on layer-wise feature matching hinders flexibility in student model design. Secondly, our approach utilizes the same SSL objective as the teacher model, which has been shown to be more effective than MSE loss from Eq. (1) for learning robust representations [28]. In the following sections, we focus on two key aspects that can impact distillation performance of *DiceHuBERT*: the student architecture and the targets used for the SSL objective.

3.1. Student Architecture

We propose student architectures to examine two key aspects of *DiceHuBERT*: 1) the feature dimension D^S and 2) the number of layers L . To isolate the effects of these factors, we design our student model architecture, keeping all other elements unchanged. We build two types of student models, HuBERT-shallow and HuBERT-narrow, by modifying the original HuBERT-base model. HuBERT-shallow is obtained by reducing **only the number of layers**. HuBERT-narrow, on the other hand, reduces **only the feature dimension and the intermediate feed-forward dimension**. Note that *DiceHuBERT* can be used with any model, including convolutional-only designs.

3.2. Target Labels

We define two distinct target labels for our SSL objective. The first is hard labels, generated through *k*-means clustering of intermediate features from the HuBERT model of N -th iteration. These labels are one-hot encoded, and we can directly apply cross-entropy loss following Eq. (2). The second target is soft labels, which provide richer information than hard labels and enhance knowledge distillation [11]. Soft labels are obtained by calculating the distance of each feature to all centroids generated by *k*-means and further applying softmax to estimate each cluster probability, where hard labels are computed via argmin of these distances. The probability $p(i|H)$ of i -th cluster with centroid H_i for the feature H is defined as:

$$p(i|H) = \frac{\exp(-\rho(H, H_i)/\tau)}{\sum_{j=\{1, \dots, K\}} \exp(-\rho(H, H_j)/\tau)}, \quad (3)$$

where $\rho(\cdot, \cdot)$ is the L_2 distance between two representations, K is the number of clusters, and τ is the temperature parameter. Then, the SSL loss for soft labels is defined as the KL-divergence between the output class distribution of a student model and the soft labels from a teacher model.

Table 1: Comparison with SOTA HuBERT distillation methods. All distillation methods, including our proposed method, use the second iteration of HuBERT-base as the teacher. We exclude LightHuBERT [25] as its compression process requires extensive computational cost.

Methods	# Param	PR	ASR (w/o LM)	ASV	SID
	M ↓	PER ↓	WER ↓	EER ↓	ACC ↑
<i>Baselines</i>					
FBank	-	82.01	23.18	9.56	8.5e-4
Wav2Vec2-base [1]	95	5.74	6.43	6.02	75.18
HuBERT-base [2]	94	5.41	6.42	5.11	81.42
<i>Distillation Methods</i>					
DistilHuBERT [23]	23	16.27	13.37	8.55	73.54
FitHuBERT [21]	22	13.32	12.09	8.00	55.71
DPHuBERT [20]	23	9.67	10.47	5.84	76.83
MaskHuBERT [22]	26	7.30	9.77	5.38	62.83
ArmHuBERT [22]	26	7.72	9.96	5.68	65.03
StarHuBERT-L [24]	26	7.97	8.91	5.45	78.66
<i>Proposed Method</i>					
DiceHuBERT	26	6.23	7.64	5.25	79.23

4. Experiments

4.1. Implementation Details

Since there are no publicly available pre-train weights for the first iteration of HuBERT, we train this model from scratch using the hyper-parameters reported in the HuBERT paper [2]. For our distillation process, we follow exactly the same setup of HuBERT second iteration pre-training, while changing only the model architecture to student following Section 3.1. For example, we set the mask span to 10 and randomly select 8% of the waveform encoder output frames. As training data, we use LibriSpeech 960h [29] following prior works. All our experiments are conducted on 32 NVIDIA A100 40GB GPUs. We used the same setup for all our experiments, unless otherwise stated.

4.2. Evaluation

We evaluate DiceHuBERT on SUPERB [27], focusing our evaluation on four tasks that exhibit the significant performance gaps between the teacher and student models: content-based tasks (PR, ASR) and speaker-based tasks (SID, ASV).

4.3. Comparison with State-of-the-Art Methods

We compare DiceHuBERT with other state-of-the-art HuBERT distillation methods in Table 1. For this experiment, we use the HuBERT-narrow ($D_{base}/2$) setup as the student architecture, and train our model solely with the SSL loss. For evaluation, we use the default SUPERB recipes but modify the learning rates for the SID task to $5e-3$ following prior works [20, 21, 24] for a fair comparison. All distillation methods in Table 1, including our method, use the second iteration of HuBERT-base as the teacher model. DiceHuBERT outperforms all other distillation methods across all tasks. Notably, DiceHuBERT achieves significant improvements in content-based tasks, such as a 21% relative improvement in ASR and a 14% relative improvement in PR, compared to the previous best models. In contrast, the performance gap in speaker-based tasks is not big but surpasses the other approaches, with improvements of 3.6% in ASV and 0.7% in SID. This comparison demonstrates the advantages of DiceHuBERT over feature distillation-based methods.

Table 2: Performance comparison across different student model architectures. L_{base} and D_{base} represent the number of layers and feature dimension of the HuBERT-base, respectively.

Models	#Param	PR	ASR (w/o LM)	ASV	SID
	M ↓	PER ↓	WER ↓ / UER ↓	EER ↓	ACC ↑
<i>Baselines</i>					
Upper Bound	94	5.41	6.42 / -	5.11	81.42
Lower Bound	26	7.91	8.44 / 2.46	5.32	77.40
<i>HuBERT-Shallow</i>					
$L=L_{base}/6$	23	14.45	13.17 / 3.99	7.37	65.78
$L=L_{base}/4$	30	10.95	11.80 / 3.55	6.55	72.10
<i>HuBERT-Narrow</i>					
$D^S=D_{base}/4$	9	12.55	11.17 / 3.36	6.24	57.64
$D^S=D_{base}/2$	26	7.48	7.97 / 2.35	5.03	76.85

4.4. Ablation Studies

We conduct ablation studies to investigate the impact of different configurations of DiceHuBERT on distillation. We begin by exploring two approaches to reducing model size: decreasing the feature dimension and reducing the number of layers. We then compare our proposed SSL-based distillation with feature distillation while fixing the architecture. Next, we compare different targets for the distillation loss: soft labels and hard labels. Lastly, we investigate the impact of teacher model size by using targets from different teacher sources. Unless otherwise specified, we use a narrow model ($D_{base}/2$) as the student architecture, HuBERT first-iteration model ($N=1$) as the teacher, employing hard labels as our default setup for ablation studies.

Baselines. We establish performance upper and lower bound for our ablation studies in Table 2. To ensure a fair comparison, we select models of second iteration ($N=2$) for our baselines. The upper bound is defined by the performance of the teacher model, which is the second iteration of the HuBERT-base model. To establish a lower bound that is **absent from all prior works**, we train HuBERT-narrow ($D_{base}/2$) from scratch following HuBERT-base pre-training process with two iterations. All the hyperparameters are identical with HuBERT-base pretraining except for the model architecture. Surprisingly, our lower bound outperforms all other distillation methods in Table 1 in terms of ASR and ASV, while achieving comparable results on the other tasks. Previous works have primarily focused on comparing the distilled model performance with that of teacher model, omitting this essential lower bound. The lower bound is vital because it illuminates the extent of improvement achieved through distillation compared to training the student model from scratch.

Student Architectures. In Table 2, we compare distillation performance of our framework on different student architecture that we described in Section 3.1. For simplicity, ‘Shallow’ and ‘Narrow’ represent HuBERT-shallow and HuBERT-narrow, respectively. We first compare shallow ($L_{base}/6$) and narrow ($D_{base}/4$). Narrow ($D_{base}/4$) outperforms shallow ($L_{base}/6$) on PR, ASR, and ASV, even though narrow ($D_{base}/4$) is 2.5 times smaller than shallow ($L_{base}/6$). Next, we compare shallow ($L_{base}/4$) and narrow ($D_{base}/2$). Remarkably, narrow ($D_{base}/2$) demonstrates significantly better performance than shallow ($L_{base}/4$), despite the fact that shallow ($D_{base}/2$) has a smaller model size. These results shows importance of maintaining the number of layers for model compression. In SUPERB, it incorporates a weighted sum of the intermediate features of the models for downstream tasks. Reducing the num-

Table 3: Performance comparison across different distillation methods.

Losses	PR	ASR (w/o LM)	ASV	SID
	PER ↓	WER ↓ / UER ↓	EER ↓	ACC ↑
\mathcal{L}_{SSL}	7.48	7.97 / 2.35	5.03	76.85
$\mathcal{L}_{SSL} + 0.1 \cdot \mathcal{L}_{feat}$	7.64	8.28 / 2.39	5.36	78.01
$\mathcal{L}_{SSL} + 1 \cdot \mathcal{L}_{feat}$	8.97	8.71 / 2.52	5.28	76.98
\mathcal{L}_{feat}	8.19	9.97 / 2.95	5.24	74.00

Table 4: The impact of different target labels for DiceHuBERT.

Labels	PR	ASR (w/o LM)	ASV	SID
	PER ↓	WER ↓ / UER ↓	EER ↓	ACC ↑
Hard Labels	7.48	7.97 / 2.35	5.03	76.85
Soft Labels ($\tau=1$)	7.19	8.01 / 2.31	5.25	76.77
Soft Labels ($\tau=5$)	6.70	8.02 / 2.33	5.51	76.36
Soft Labels ($\tau=10$)	7.43	8.58 / 2.51	5.24	75.39

ber of layers would result in the loss of layer-wise knowledge distilled from the teacher, leading to poor performance for SUPERB. This finding partially explains the inferior performance of DistilHuBERT [18], which shares an identical architecture with a shallow ($D_{base}/6$) model.

Distillation Methods. We compare our proposed SSL-based distillation method with feature distillation discussed in Section 2.1. To isolate the impact of architecture, we fix the model architecture to HuBERT-narrow ($D_{base}/2$) and apply different distillation methods. For feature distillation loss, we use \mathcal{L}_{feat} defined in Eq. (1). We employ the masked input for both \mathcal{L}_{SSL} and \mathcal{L}_{feat} to eliminate the influence of the input. The masked input approach also enables the integration of both losses. As evidenced by MaskHuBERT [22], masking the input does not negatively impact the performance of features distillation. The results are presented in Table 3. We find that model trained with SSL loss consistently outperforms the one with feature distillation across all tasks, highlighting the advantages of SSL for distillation. To explore the potential of combining both approaches, we investigate the effect of adding \mathcal{L}_{SSL} and \mathcal{L}_{feat} for a training objective, varying the weight for \mathcal{L}_{feat} . Combining the losses yields improvements in SID when the weight is 0.1; however, all other task performance were degraded compared to SSL alone, showing that adding feature distillation atop SSL is not always beneficial.

Target Labels. In Table 4, we investigate the impact of the soft labels introduced for the SSL loss in Section 3.2. As HuBERT representations do not lie on the unit sphere, we conduct the experiments by increasing τ in the soft labels definition to alleviate unnormalized feature space (Eq. (3)) and provide reliable targets for a student. We keep the temperature for the student logits as 1, standard default for soft labels distillation [11]. We first see that when the temperature is low ($\tau=1$), the soft labels are close to hard labels, resulting in similar performance to the model trained with hard labels. The PR shows a slight improvement, but all other tasks are degraded. As the temperature increases, some tasks show improvement (e.g., PR at $\tau=5$), but the overall performance generally worsens. Notably, when $\tau=10$, the performance of all tasks was severely degraded, except for a slight improvement in PR. This suggests that over-smoothing the labels is not beneficial for knowledge distillation.

Different Teachers. To assess the influence of teacher models on our distillation framework, we compare the distillation

Table 5: Performance comparison across different teachers. Note that since HuBERT-large is trained on Libri-Light 60K hours [31], direct comparisons with other models are not entirely fair. However, comparisons can still provide insights into the significance of the teacher model in HuBERT distillation.

Models	#Param	PR	ASR (w/o LM)	ASV	SID
	M ↓	PER ↓	WER ↓ / UER ↓	EER ↓	ACC ↑
<i>HuBERT-base ($N=1$)</i>					
Teacher	95	8.56	8.92 / 2.62	5.58	74.33
Student	26	7.48	7.97 / 2.35	5.03	76.85
<i>HuBERT-base ($N=2$)</i>					
Teacher [†]	95	5.41	6.42 / -	5.11	81.42
Student	26	6.23	7.60 / 2.21	5.25	79.23
<i>HuBERT-large</i>					
Teacher [†]	317	3.54	3.62 / -	5.98	90.33
Student	26	6.35	6.90 / 2.02	5.66	76.39

[†]Numbers from the SUPERB paper [27].

results obtained using various teacher models in Table 5. We employ the first iteration HuBERT-base ($N=1$), the second iteration HuBERT-base ($N=2$), and the HuBERT-large model as the teacher models. We pre-trained HuBERT-base ($N=1$) from scratch, but used pre-trained weights from the official HuBERT repository for the others. For target generation using k -means, we select the 6th layer for HuBERT-base ($N=1$) and the 9th layer for HuBERT-base ($N=2$), as used in the HuBERT paper for target generation [2]. For HuBERT-large, the 18th layer out of 24 was used. When HuBERT-base ($N=1$) serves as a teacher, surprisingly, the student model outperforms all tasks. Despite the teacher model’s poor performance, it generates high-quality targets for the student model that outperform the teacher model even with a smaller model size. These experiments demonstrate the advantages of self-distillation in HuBERT. Furthermore, comparing student performance between HuBERT-base ($N=1$) and HuBERT-base ($N=2$) reveals further performance improvement as the teacher’s quality improves. Lastly, when HuBERT-large is used as a teacher, ASR performance improves compared to HuBERT-base ($N=2$), but not in other tasks. This difference might be attributed to the SUPERB model utilizing all 24 layers of HuBERT-large for downstream tasks, which performs well as a teacher. However, during distillation, only a single layer from the teacher model is selected for target generation, potentially leading to knowledge loss for the student. Future work could explore how to define aggregation between different layers’ representations to cluster and provide targets for distillation, e.g. recent work [30] showed that average of representations instead of selected one layer representation improves HuBERT framework.

5. Conclusion

In this paper, we introduce DiceHuBERT, a model compression framework that leverages HuBERT’s self-distillation mechanism by directly replacing the original model with a student model. DiceHuBERT outperforms feature distillation methods across various SUPERB downstream tasks, achieving over 21% improvement in ASR and 14% in PR compared to previous SOTA HuBERT distillation approaches. Additionally, our ablation studies provide valuable insights into designing efficient architectures and selecting optimal distillation targets for future HuBERT compression research.

<https://github.com/facebookresearch/fairseq>

6. References

- [1] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [2] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM transactions on audio, speech, and language processing*, vol. 29, pp. 3451–3460, 2021.
- [3] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [4] A. Babu, C. Wang, A. Tjandra, K. Lakhotia, Q. Xu, N. Goyal, K. Singh, P. Von Platen, Y. Saraf, J. Pino *et al.*, “Xls-r: Self-supervised cross-lingual speech representation learning at scale,” *arXiv preprint arXiv:2111.09296*, 2021.
- [5] X. Song, G. Wang, Y. Huang, Z. Wu, D. Su, and H. Meng, “Speech-xlnet: Unsupervised acoustic model pretraining for self-attention networks,” 2020.
- [6] Z. Huang, S. Watanabe, S.-w. Yang, P. García, and S. Khudanpur, “Investigating self-supervised learning for speech enhancement and separation,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6837–6841.
- [7] X. Chang, T. Maekaku, P. Guo, J. Shi, Y.-J. Lu, A. S. Subramanian, T. Wang, S.-w. Yang, Y. Tsao, H.-y. Lee *et al.*, “An exploration of self-supervised pretrained representations for end-to-end speech recognition,” in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 228–235.
- [8] Y.-A. Chung, Y. Zhang, W. Han, C.-C. Chiu, J. Qin, R. Pang, and Y. Wu, “W2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training,” in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 244–250.
- [9] T. N. Sainath and C. Parada, “Convolutional neural networks for small-footprint keyword spotting,” in *Interspeech*, 2015, pp. 1478–1482.
- [10] M. Wu, S. Panchapagesan, M. Sun, J. Gu, R. Thomas, S. N. P. Vitaladevuni, B. Hoffmeister, and A. Mandal, “Monophone-based background modeling for two-stage on-device wake word detection,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5494–5498.
- [11] G. Hinton, “Distilling the knowledge in a neural network,” *arXiv preprint arXiv:1503.02531*, 2015.
- [12] 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.
- [13] J. Ba and R. Caruana, “Do deep nets really need to be deep?” *Advances in neural information processing systems*, vol. 27, 2014.
- [14] G. Chen, W. Choi, X. Yu, T. Han, and M. Chandraker, “Learning efficient object detection models with knowledge distillation,” *Advances in neural information processing systems*, vol. 30, 2017.
- [15] J. Cui, B. Kingsbury, B. Ramabhadran, G. Saon, T. Sercu, K. Audhkhasi, A. Sethy, M. Nussbaum-Thom, and A. Rosenberg, “Knowledge distillation across ensembles of multilingual models for low-resource languages,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 4825–4829.
- [16] Y. Yang, J. Qiu, M. Song, D. Tao, and X. Wang, “Distilling knowledge from graph convolutional networks,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 7074–7083.
- [17] T. Asami, R. Masumura, Y. Yamaguchi, H. Masataki, and Y. Aono, “Domain adaptation of dnn acoustic models using knowledge distillation,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 5185–5189.
- [18] S. Gandhi, P. von Platen, and A. M. Rush, “Distil-whisper: Robust knowledge distillation via large-scale pseudo labelling,” *arXiv preprint arXiv:2311.00430*, 2023.
- [19] D. Chen, J.-P. Mei, H. Zhang, C. Wang, Y. Feng, and C. Chen, “Knowledge distillation with the reused teacher classifier,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 11 933–11 942.
- [20] Y. Peng, Y. Sudo, S. Muhammad, and S. Watanabe, “Dphubert: Joint distillation and pruning of self-supervised speech models,” in *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, vol. 2023, 2023, pp. 62–66.
- [21] Y. Lee, K. JANG, J. Goo, Y. Jung, and H.-R. Kim, “Fithubert: Going thinner and deeper for knowledge distillation of speech self-supervised learning,” in *23rd Annual Conference of the International Speech Communication Association, INTERSPEECH 2022*. ISCA, 2022, pp. 3588–3592.
- [22] K. Jang, S. Kim, S. Yun, and H.-R. Kim, “Recycle-and-distill: Universal compression strategy for transformer-based speech ssl models with attention map reusing and masking distillation,” in *24th International Speech Communication Association, Interspeech 2023*. International Speech Communication Association, 2023, pp. 316–320.
- [23] H.-J. Chang, S.-w. Yang, and H.-y. Lee, “Distilhubert: Speech representation learning by layer-wise distillation of hidden-unit bert,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7087–7091.
- [24] K. Jang, S. Kim, and H. Kim, “Star: Distilling speech temporal relation for lightweight speech self-supervised learning models,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 10 721–10 725.
- [25] 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,” *Proc. Interspeech 2022*, pp. 1686–1690, 2022.
- [26] D. Arthur and S. Vassilvitskii, “k-means++: The advantages of careful seeding,” *Stanford, Tech. Rep.*, 2006.
- [27] S. W. Yang, P. H. Chi, Y. S. Chuang, C. I. J. Lai, K. Lakhotia, Y. Y. Lin, A. T. Liu, J. Shi, X. Chang, G. T. Lin *et al.*, “Superb: Speech processing universal performance benchmark,” in *22nd Annual Conference of the International Speech Communication Association, INTERSPEECH 2021*. International Speech Communication Association, 2021, pp. 3161–3165.
- [28] A. K. Menon, A. S. Rawat, S. J. Reddi, S. Kim, and S. Kumar, “Why distillation helps: a statistical perspective,” *arXiv preprint arXiv:2005.10419*, 2020.
- [29] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [30] L.-W. Chen, T. Higuchi, H. Bai, A. Hussen Abdelaziz, A. Rudnicky, S. Watanabe, T. Likhomanenko, B.-J. Theobald, and Z. Aldeneh, “Exploring prediction targets in masked pre-training for speech foundation models,” *arXiv e-prints*, pp. arXiv-2409, 2024.
- [31] J. Kahn, M. Rivière, W. Zheng, E. Kharitonov, Q. Xu, P. E. Mazaré, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen, T. Likhomanenko, G. Synnaeve, A. Joulin, A. Mohamed, and E. Dupoux, “Libri-light: A benchmark for asr with limited or no supervision,” in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7669–7673, <https://github.com/facebookresearch/libri-light>.