



# Reducing Barriers to Self-Supervised Learning: HuBERT Pre-training with Academic Compute

William Chen<sup>1</sup>, Xuankai Chang<sup>1</sup>, Yifan Peng<sup>2</sup>, Zhaoheng Ni<sup>3</sup>, Soumi Maiti<sup>1</sup>, Shinji Watanabe<sup>1</sup>

<sup>1</sup>Language Technologies Institute, Carnegie Mellon University, USA

<sup>2</sup>Electrical and Computer Engineering, Carnegie Mellon University, USA

<sup>3</sup>Meta, USA

{wc4, xuankaic}@andrew.cmu.edu

## Abstract

Self-supervised learning (SSL) has led to great strides in speech processing. However, the resources needed to train these models has become prohibitively large as they continue to scale. Currently, only a few groups with substantial resources are capable of creating SSL models, which harms reproducibility. In this work, we optimize HuBERT SSL to fit in academic constraints. We reproduce HuBERT independently from the original implementation, with no performance loss. Our code and training optimizations make SSL feasible with only 8 GPUs, instead of the 32 used in the original work. We also explore a semi-supervised route, using an ASR model to skip the first pre-training iteration. Within one iteration of pre-training, our models improve over HuBERT on several tasks. Furthermore, our HuBERT Large variant requires only 8 GPUs, achieving similar performance to the original trained on 128. As our contribution to the community, all models, configurations, and code are made open-source in ESPnet.

**Index Terms:** Self-supervised learning, HuBERT

## 1. Introduction

Self-supervised pre-training has reached new epochs in recent years by scaling Transformer-based models [1]. By increasing both the model size and training data, impressive performance has been achieved across a wide variety of tasks [2–8]. In speech processing, self-supervision has been used at scale to learn powerful representations of acoustic features [9–13]. One promising approach to self-supervised learning (SSL) in speech is masked prediction-based training. Inspired by deep clustering [14] and masked language modelling [15], HuBERT [11] learns to predict the cluster assignment of masked regions of the acoustic input. Importantly, HuBERT requires multiple iterations of pre-training, using clusters derived from the features of previous iterations.

The computational expenses involved with pre-training current SSL models like HuBERT presents a major hurdle to reproducibility and the advancement of open science. Recent approaches are limited to using compression-based techniques such as distillation [16–18] and pruning [19, 20], which require pre-trained SSL weights and cannot create a new standalone model. This makes SSL pre-training *prohibitively expensive* to research groups without large-scale computing resources, like the 128 GPUs needed to train the original HuBERT Large. Unfortunately, the trend towards scaling up these models shows no sign of slowing down, which means that the ability to train them will become even more concentrated in the hands of a select few, further hampering the democratization of knowledge.

In the field of Natural Language Language Processing (NLP), where Large Language Models (LLMs) with billions

of parameters [21, 22] are more common, attempts have been made to address this problem. For example, GPT-NeoX-20B [23], and BLOOM [24] were the results of valuable community-driven efforts to build open-source LLMs. The benefits of this are manifold: researchers are able to share important optimization tricks and ablation logs, therefore reducing the amount of compute required for future studies. However, all of these works are limited to the NLP domain. To the best of our knowledge, such an effort has not been realized in speech.

Our goal is to build a pre-trained SSL model for speech processing without access to large-scale computing resources. Specifically, we want to reproduce HuBERT-style pre-training within a resource constraint that is feasible for academic settings. In doing so, we hope to share valuable insights for training SSL models with the broader research community, and lift the computational barriers that have recently surrounded large-scale self-supervised training.

We approach this challenge from two main directions. The first is to reproduce HuBERT pre-training such that SSL is feasible on a limited compute budget. This includes a multitude of factors, such as storage space, pre-processing time, training time, and most critically, the number of GPUs required for training. The other angle is to reduce the compute necessary to train a standalone SSL model. We explore a semi-supervised SSL (semi-SSL) approach, where the first HuBERT iteration is replaced with a supervised model trained on Automatic Speech Recognition (ASR). In doing so, the costs for SSL pre-training can be drastically lowered, even when considering the costs of training the supervised model.

Overall, we showed that the computational overhead of SSL pre-training can be considerably reduced. Due to our optimization of the code, pre-processing, and training settings, we are able to achieve shorter pre-training times with only 8 GPUs than that of prior work [11, 25], which used 32. Our reproduced HuBERT achieves comparable results to the original implementation across the SUPERB benchmark [5]. Furthermore, our semi-SSL models, pre-trained with only 8 GPUs and for a single iteration, improve over comparable state-of-the-art SSL models on 4 different SUPERB tasks. We open source all of our models, training configurations, and code as one of our core contributions, so that different implementations are available to the research community for SSL.

## 2. Model Design and Implementation

In this section, we first provide background on HuBERT-based SSL and discuss our reproduction efforts. We then outline our semi-SSL approach, designed to reduce the amount of HuBERT training iterations. Afterwards, we discuss modifications made to HuBERT pre-processing and training, allowing for it to be done even in more resource constrained settings.




Figure 1: Diagram of our model pipeline. The supervised E-Branchformer is used to extract hidden representations, which are then clustered with k-means. The HuBERT model is pre-trained on these clusters, before being fine-tuned.

## 2.1. Reproducing HuBERT

HuBERT builds off of the wav2vec 2.0 architecture [10], consisting of a convolutional feature extractor, a Transformer encoder, and a codebook output layer. The labels used to train the model are a sequence of cluster assignments. The assignments are obtained using k-means clustering on the acoustic features, where  $k$  is equal to the codebook size. The acoustic features are Mel-Frequency Cepstral Coefficients (MFCCs) or the encoded representations of a pre-trained model. HuBERT is then trained to predict the cluster assignments of masked regions in the audio. We reproduce HuBERT following the original pipeline [11]. Iteration 0 is pre-trained on clusters derived from 39-dimension MFCC features extracted from LibriSpeech 960h [26], which are clustered with a k-means codebook size of 100 for pseudo-labeling. Iteration 1 is then trained on features extracted from the 6th layer of the iteration 0 model clustered with a codebook size of 500. Both iterations use the HuBERT Base architecture. Finally, iteration 2 uses the HuBERT Large architecture and is trained on clusters derived from the 9th layer of the iteration 1 model, using a cluster size of 500.

One issue we found when reproducing HuBERT was that there was little correlation between the pre-training validation curves and fine-tuning performance. For example, a model checkpoint after 400K steps had wildly different ASR results from a checkpoint at 800K steps, despite very similar validation loss/accuracy. This meant that hyperparameter tuning required full pre-training trials, each of which required close to 1000 GPU hours. We conducted 10 trials in total, spending a total of 9600 GPU hours, before finding an optimal setting that matched the original in performance. These settings are shared with the community, so that future efforts can be accelerated.

## 2.2. Semi-Self-Supervised Learning

As mentioned above, SSL training requires significant trial and error due to the lack of direct evaluation metrics. This is most critical for iteration 0 of HuBERT, which serves as the foundation for subsequent iterations. On the other hand, supervised ASR models are much less costly to train and tune, as they can be directly evaluated through fine-tuning metrics, such as validation loss and accuracy<sup>1</sup>. Could we instead replace the entire first iteration of HuBERT pre-training with supervised training? In doing so, we could eliminate the lengthy hyper-parameter tuning process of an entire pre-training iteration.

We explore a semi-supervised SSL (semi-SSL) route, where a supervised model is used to replace iteration 0 of pre-

<sup>1</sup>While these may not be perfectly indicative of final performance, they are still more insightful than HuBERT SSL metrics, such as masked prediction accuracy, which depend on the quality/number of the target clusters. Trivial clusters may lead to high masked prediction accuracy but poor downstream results.

training (Figure 1). Instead of using MFCC features for the initial clusters, we extract the hidden representations of an ASR model. Since the goal is to reduce the amount of training iterations as much as possible, our approach is to use the strongest possible ASR model that can yield better initial clusters. Our choice is a powerful ASR encoder-decoder, built on top of the state-of-the-art E-Branchformer [27] architecture, and trained on the large-scale GigaSpeech corpus [28]. One important decision was determining which layer of the ASR model to extract the features from. We selected features from the middle layer (layer 8) for one pre-training trial, since they had the highest phonetic content when evaluated against frame-level phonetic transcripts [29, 30]. We conducted an additional trial using features from the last layer (layer 16), which was found to encode the most linguistic content [31].

While our methodology is similar to recent developments in task-oriented pre-training, the core motivations are different. Prior works [32–34] adopted a self-training approach, where supervised models are boosted via unsupervised pre-training. A fine-tuned ASR model is used to initialize SSL, and then fine-tuned again on the same supervised data. Their goal is to improve performance on a *specific* task and domain. In doing so, the performance of the model may be degraded when used for other tasks, such as with the case of speaker verification [32]. However, our motivation is to reduce the cost of pre-training while maintaining performance across a *range* of tasks. Because of this, it is important that our teacher ASR model is trained on a broad range of data, so that the semi-SSL model can be usable across a wide variety of tasks and domains.

## 2.3. Optimizing Pre-processing

Most costs associated with SSL occur during the pre-training stage, but we also found that pre-processing also required a significant investment. The official implementation of HuBERT in fairseq [35] saves the extracted features necessary for k-means to the disk as NumPy [36] arrays. All of the features were loaded into a single node for k-means training. This led to large amounts of memory usage, particularly when using the high dimensional hidden representations. To alleviate this, we implemented a k-means algorithm parallelized across multiple computing nodes. We further optimized the k-means process by swapping NumPy for Kaldiio<sup>2</sup>. This allows the extracted features to be sampled from k-means training prior to loading from disk (rather than after as in the original implementation), reducing the amount of wasted memory and I/O time. On 96 hours of extracted 39-dimension MFCC features, we are able to reduce the maximum memory usage during k-means from 27.3GB to 162MB. The reductions for k-means on the 512-dimension ASR hidden representations are even more significant: memory usage is reduced from 122GB to 4.6 GB.

One major constraint was storage space. The audio used to train the HuBERT-Large model consumed over 7TB of disk space, with the majority from Libri-Light 60k [37]. The feature extraction step in HuBERT only exacerbated the problem due to their high dimensionality. We found that extracting all of the features for Libri-Light would consume 25TB of storage. To avoid this, we made significant modifications to the clustering pipeline. The k-means model was trained using features from LibriSpeech, which only used 500GB of space. The Libri-Light features would then be converted into cluster assignments with the trained k-means model within memory, removing the need

<sup>2</sup><https://github.com/nttclab-sp/kaldiio>

to store the extracted features. Through this, we reduced the storage footprint of these steps from 25TB to 500GB.

### 2.4. Training Optimizations

During training, we only had access to 8 GPUs, one-fourth of the 32 GPUs used in the original implementation. However, the authors of HuBERT also found batch size to be a large determinant in model performance [11]. To make up for the contrast in resources, we use gradient accumulation to increase our effective batch size to similar levels. Furthermore, we utilized mixed precision training with brain float 16 to increase training speed.

## 3. Experimental Setup

For independent reproducibility, we conduct all experiments outside of the original fairseq [35] framework, implementing our full pipeline within ESPnet [38].

### 3.1. Datasets

We use GigaSpeech [28], a 10k hour corpus of English speech across a variety of domains, such as audiobooks and podcasts, to train the ASR model. We use the 960 hours of LibriSpeech [26] audio to pre-train all models using the HuBERT Base architecture. Our models that used the HuBERT Large architecture is trained on 60k hours of audio from Libri-Light [37]. For model fine-tuning, we used the 10 hours labeled subset of Libri-Light. To understand the generalizability of our semi-SSL models, we also evaluate them using the SUPERB Benchmark [5].

### 3.2. HuBERT

We reproduce HuBERT following the original pipeline [11] as described in Section 2.1. Iterations 0 and 1 follow the HuBERT Base architecture, which consists of 12 encoder layers with a hidden size of 768. The models were trained using the Adam optimizer for 880k steps. We used a warmup learning rate scheduler, such that the learning rate was linearly increased for 32k steps until it reached a peak of 0.0005, before exponentially decaying. Mini-batches are constructed using ESPnet’s `numel` sampler, using a batch bins size of 45M to maximize the data in each GPU. Hyper-parameter search was conducted throughout 10 full trials of pre-training. Iteration 2 follows the HuBERT Large architecture, which has 24 encoder layers and a hidden size of 1024. The model was trained with the same optimization settings as the Base-sized models. Due to the size of the model, we reduce the batch size to one audio clip per GPU.

### 3.3. Semi-supervised SSL Pre-training

#### 3.3.1. Supervised ASR Teacher

The ASR model is a hybrid CTC/attention [39] encoder-decoder. The encoder is a 17-layer E-Branchformer [27] (an enhanced version of Branchformer [40]), while the decoder is a 6-layer Transformer [1]. We use the same configuration as the existing E-Branchformer recipe for LibriSpeech. Each encoder layer has a hidden size of 512, 8 attention heads, and a kernel size of 31. Each decoder layer has 8 attention heads and a feed-forward dimension of 2048.

#### 3.3.2. Semi-SSL Models

As mentioned in Section 2.2, we first use the ASR teacher to extract representations of audio from LibriSpeech, for both layers 8 and 16. Clusters are then created using the k-means algorithm with a codebook size of 500. Due to the larger memory footprint from the increased output size, the batch bins size is reduced to 40M. All other settings remain the same as the reproduced HuBERT. Our models are denoted as the following:

**GS-8:** Iteration 0 clusters were derived from the middle layer

Table 1: Word Error Rate on Librispeech dev/test sets. SSL models were fine-tuned on 10 hours of data from Librilight, using CTC loss without a language model. Results from [25] were obtained from the official repository<sup>4</sup>.

model	dev		test	
	clean	other	clean	other
<i>Prior Work (Iteration 1)</i>				
HuBERT Base[25]	10.9	17.5	10.9	17.8
<i>Iteration 0</i>				
HuBERT Base	15.5	23.9	15.9	24.1
GS-8	12.2	19.4	12.3	19.9
GS-16	<b>9.6</b>	<b>15.7</b>	<b>9.7</b>	<b>16.3</b>
<i>Iteration 1</i>				
HuBERT Base	10.4	17.4	10.5	17.6
GS-16 (Base)	10.2	16.0	10.3	16.4
GS-16 (Large)	<b>8.4</b>	<b>13.8</b>	<b>8.2</b>	<b>13.8</b>
<i>Iteration 2</i>				
HuBERT Large	9.5	14.9	9.6	14.9

(layer 8) of the ASR model, which contained the most phonetic content. We only train this model for iteration 0, using the HuBERT Base architecture.

**GS-16 Base:** Iteration 0 training targets were derived from the last layer (layer 16) of the ASR teacher, which should contain the most linguistic content [31]. This model was trained for two iterations, both using the HuBERT Base architecture.

**GS-16 Large:** This model was trained on clusters obtained from iteration 0 of GS-16 Base. It uses the same settings as the HuBERT Large described in Section 3.2

### 3.4. Fine-tuning

We follow the fine-tuning setup of [25]<sup>3</sup>, allowing for comparisons in both compute usage and downstream performance. Models are fine-tuned on 10 hours of Libri-Light using CTC loss [41] for 10200 updates, with the pre-trained encoder frozen for the first 4000. We use a character level vocabulary of size 31. The 10 best checkpoints are averaged for inference, which is done with CTC greedy decoding and without a language model.

### 3.5. Hardware

We train the majority of models on publicly accessible HPC clusters with 8xA100 GPUs. To conduct more experimental trials, we also train models using Amazon AWS EC2; we used a single 40GB A100x8 instance and 32GB V100x8 instance.

## 4. Results

### 4.1. ASR Fine-tuning

The results are presented in Table 1, evaluated using Word Error Rate (WER) on the LibriSpeech test sets. Our reproduced HuBERT Base yielded strong results by outperforming prior work [25], demonstrating the feasibility of SSL even with limited compute. Both GS-8 and GS-16 performed better than HuBERT Base in iteration 0, showing the effectiveness of semi-supervised SSL. We observed that GS-16 performed much better than GS-8, likely as a result of its clusters encoding content more relevant to ASR. Furthermore, GS-16 had a lower WER than our reproduced HuBERT Base, which was pre-trained for an additional iteration (0.8/1.3 WER lower on test-clean/other).

<sup>3</sup><https://github.com/pytorch/audio/tree/main/examples/hubert>

<sup>4</sup>Ibid.

Table 2: Evaluation on the SUPERB Benchmark. The tasks are grouped into 5 categories: recognition, detection, semantics, speaker, and paralinguistics. Metrics for each task are: phoneme error rate (PR), word error rate (ASR), maximum term weighted value (QbE), F1 and concept error rate (SF), equal error rate (ASV), diarization error rate (SD), and accuracy (KS, IC, and ER)

Method	Params	Score	Recognition		Detection		Semantics			Speaker		Paralinguistics
			PR↓	ASR↓	KS↑	QbE↑	IC↑	SF (F1)↑	SF (CER)↓	ASV↓	SD↓	ER↑
<i>Prior work</i>												
DistillHuBERT [16]	23.5M	76.2	16.3	13.4	95.9	5.1	95.0	82.6	36.6	8.6	6.2	63.0
LightHuBERT [42]	95.0M	81.2	4.2	5.7	<b>96.8</b>	7.4	98.5	88.4	25.3	5.1	<b>5.5</b>	66.3
HuBERT Base [11]	94.7M	80.7	5.4	6.4	96.3	7.4	98.3	88.5	25.2	5.1	5.9	64.9
HuBERT Large [11]	316M	81.4	<b>3.5</b>	<b>3.6</b>	95.3	<b>3.5</b>	<b>98.8</b>	<b>89.8</b>	<b>21.8</b>	6.0	5.8	<b>67.6</b>
<i>Iteration 0</i>												
GS-8	94.7M	80.9	6.3	6.9	<b>96.1</b>	<b>8.1</b>	98.9	89.4	22.5	<b>5.9</b>	<b>6.2</b>	<b>64.1</b>
GS-16	94.7M	80.8	<b>6.1</b>	<b>5.4</b>	95.6	5.6	<b>99.1</b>	<b>90.0</b>	<b>21.5</b>	6.4	6.5	64.0
<i>Iteration 1</i>												
HuBERT Base	94.7M	80.7	5.8	6.9	95.8	10.1	97.6	89.0	24.1	5.6	6.1	62.8
GS-16 (Base)	94.7M	81.3	<b>5.3</b>	6.6	96.3	<b>10.5</b>	<b>99.0</b>	<b>89.9</b>	<b>22.1</b>	5.6	6.2	63.1
GS-16 (Large)	316M	81.4	7.4	<b>5.4</b>	<b>97.4</b>	10.1	98.9	89.0	23.3	6.4	<b>5.6</b>	<b>66.1</b>

## 4.2. Universal Representations Benchmark

SUPERB consists of several speech processing tasks designed to evaluate SSL models. Tasks include Phoneme Recognition (PR), ASR, Keyword Spotting (KS), Query by Example (QbE), Intent Classification (IC), Slot Filling (SF), Automatic Speaker Verification (ASV), Speaker Diarization (SD), and Emotion Recognition (ER). To obtain an overall picture of each model’s performance, we calculate a score in the same way as [12, 42]: QbE is multiplied by 100, error rates are subtracted from 1, and average the scores across all tasks.

The results on SUPERB are contained in Table 2. To compare with prior work, we include results for the original HuBERT implementation. We also include the scores of distillation-based approaches [16, 42], although they cannot be used to train a standalone SSL model. Our reproduced HuBERT performs at a similar level to the original, showing that our training optimizations did not degrade the performance. Our semi-SSL models also achieved comparable scores to the original HuBERT within a single iteration of pre-training (80.9 and 80.8 vs 80.7). These results indicate that features used for the initial cluster assignments had a significant effect on downstream performance. GS-16 had significant gains in semantic-related tasks and ASR. However, this came at the cost of performance in other categories, such as detection. On the other hand, GS-8 obtained scores that were more even across the board, having less degradation on detection-related tasks.

The GS-16 models trained for an additional iteration achieved even stronger results. The overall score improved from 80.9 to 81.3 for the GS-16 Base model, which is close to the original HuBERT Large (81.4). Our GS-16 Large also matches the performance of the original HuBERT Large, while being trained for 1 less iteration. This shows that our semi-SSL technique is effective in reducing the computational expenses of SSL while maintaining performance.

## 4.3. Training Time

Table 3 breaks down the resources and time required to train each model. Each HuBERT Base iteration could be trained in 960 to 1200 GPU hours using 8 Nvidia A100s. We highlight this significant improvement over prior work, which consumed 1344 to 2150 GPU hours [25]. Specifically, they required 16.8 hours for every 100K gradient steps when training a HuBERT Base model on 32 A100 GPUs. Our optimizations resulted in a significant speedup, such that the same amount of steps only

Table 3: Models by compute usage. Our models consume less GPU resources while maintaining similar performance.

Model	GPU	Steps	GPU Hours	Days
<i>Iteration 0</i>				
HuBERT Base [25]	A100x32	250K	1344	2
HuBERT Base (ours)	A100x8	880K	960	5
GS-8	A100x8	880K	1166	6
GS-16	A100x8	880K	1166	6
<i>Iteration 1</i>				
HuBERT Base [25]	A100x32	400k	2150	3
HuBERT Base (ours)	A100x8	880k	960	5
GS-16 Base	V100x8	880k	1600	8
GS-16 Large	A100x8	3M	2880	15
<i>Iteration 2</i>				
HuBERT Large	A100x16	2.4M	7550	20

required 16.1 hours with 8 GPUs.

## 5. Conclusion

Advances in SSL have focused on scaling larger models to more data. While effective, this constrains SSL research to only those that have considerable computing resources. Our work focuses on training state-of-the-art SSL models within academic resource constraints. We reproduced HuBERT independently from the original implementation, while optimizing its training to fit within only 8 GPUs, and without performance loss. To further reduce costs, we leverage supervised models to initialize the clusters assignments for a semi-SSL HuBERT. Our evaluation on SUPERB shows that semi-SSL models can reach comparable, if not better, results across many tasks within a single iteration of pre-training. Furthermore, we showcase the feasibility of academic SSL by pre-training a 300M parameter semi-SSL model using only 8 GPUs, as compared to the original 128, which yields results comparable to HuBERT Large. All of our models, training configurations, and code are made available open-source as our core contribution to the community.

## 6. Acknowledgements

This work used the PSC Bridges2 and the NCSA Delta systems from the Advanced Cyberinfrastructure Coordination Ecosystem, as part of project cis210027p and supported by NSF award number ACI-1445606.

## 7. References

- [1] A. Vaswani *et al.*, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [2] A. Srivastava *et al.*, “Beyond the imitation game: Quantifying and extrapolating the capabilities of language models,” *arXiv preprint arXiv:2206.04615*, 2022.
- [3] A. Wang *et al.*, “GLUE: A multi-task benchmark and analysis platform for natural language understanding,” in *Proc. EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, 2018, pp. 353–355.
- [4] A. Wang *et al.*, “SuperGLUE: A stickier benchmark for general-purpose language understanding systems,” in *Proc. NeurIPS*, H. Wallach *et al.*, Eds., vol. 32, 2019.
- [5] S.-w. Yang *et al.*, “SUPERB: Speech Processing Universal Performance Benchmark,” in *Proc. Interspeech*, 2021, pp. 1194–1198.
- [6] J. Deng *et al.*, “ImageNet: A large-scale hierarchical image database,” in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 248–255.
- [7] X. Chang *et al.*, “An exploration of self-supervised pretrained representations for end-to-end speech recognition,” in *Proc. ASRU*, 2021, pp. 228–235.
- [8] A. Mohamed *et al.*, “Self-supervised speech representation learning: A review,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1179–1210, 2022.
- [9] S. Schneider *et al.*, “wav2vec: Unsupervised Pre-Training for Speech Recognition,” in *Proc. Interspeech*, 2019, pp. 3465–3469.
- [10] A. Baevski *et al.*, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Proc. NeurIPS*, H. Larochelle *et al.*, Eds., vol. 33, 2020, pp. 12 449–12 460.
- [11] W.-N. Hsu *et al.*, “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.
- [12] S. Chen *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.
- [13] C. Wang *et al.*, “Improving self-supervised learning for speech recognition with intermediate layer supervision,” in *Proc. ICASSP*, 2022, pp. 7092–7096.
- [14] M. Caron *et al.*, “Deep clustering for unsupervised learning of visual features,” in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 132–149.
- [15] J. Devlin *et al.*, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proc. ACL*, 2019, pp. 4171–4186.
- [16] H.-J. Chang, S.-w. Yang, and H.-y. Lee, “DistilHuBERT: Speech representation learning by layer-wise distillation of hidden-unit BERT,” in *Proc. ICASSP*, 2022.
- [17] Y. Lee *et al.*, “FitHuBERT: Going Thinner and Deeper for Knowledge Distillation of Speech Self-Supervised Models,” in *Proc. INTERSPEECH*, 2022.
- [18] T. Ashihara *et al.*, “Deep versus Wide: An Analysis of Student Architectures for Task-Agnostic Knowledge Distillation of Self-Supervised Speech Models,” in *Proc. INTERSPEECH*, 2022.
- [19] C.-I. J. Lai *et al.*, “PARP: Prune, Adjust and Re-Prune for Self-Supervised Speech Recognition,” in *Proc. NeurIPS*, 2021.
- [20] Y. Peng *et al.*, “Structured Pruning of Self-Supervised Pre-trained Models for Speech Recognition and Understanding,” in *Proc. ICASSP*, 2023.
- [21] A. Chowdhery *et al.*, “PaLM: Scaling language modeling with pathways,” *arXiv preprint arXiv:2204.02311*, 2022.
- [22] T. Brown *et al.*, “Language models are few-shot learners,” in *Proc. NeurIPS*, H. Larochelle *et al.*, Eds., vol. 33, 2020, pp. 1877–1901.
- [23] S. Black *et al.*, “GPT-neox-20b: An open-source autoregressive language model,” in *Challenges & Perspectives in Creating Large Language Models*, 2022.
- [24] T. L. Scao *et al.*, “BLOOM: A 176b-parameter open-access multilingual language model,” *arXiv preprint arXiv:2211.05100*, 2022.
- [25] Y.-Y. Yang *et al.*, “Torchaudio: Building blocks for audio and speech processing,” in *Proc. ICASSP*, 2022, pp. 6982–6986.
- [26] V. Panayotov *et al.*, “Librispeech: An asr corpus based on public domain audio books,” in *Proc. ICASSP*, 2015, pp. 5206–5210.
- [27] K. Kim *et al.*, “E-Branchformer: Branchformer with enhanced merging for speech recognition,” in *Proc. SLT*, 2023, pp. 84–91.
- [28] G. Chen *et al.*, “GigaSpeech: An Evolving, Multi-Domain ASR Corpus with 10,000 Hours of Transcribed Audio,” in *Proc. Interspeech*, 2021, pp. 3670–3674.
- [29] M. McAuliffe *et al.*, “Montreal Forced Aligner: Trainable Text-Speech Alignment Using Kaldi,” in *Proc. Interspeech*, 2017, pp. 498–502.
- [30] L. Lugosch *et al.*, “Speech Model Pre-Training for End-to-End Spoken Language Understanding,” in *Proc. Interspeech*, 2019, pp. 814–818.
- [31] K. Shim, J. Choi, and W. Sung, “Understanding the role of self attention for efficient speech recognition,” in *Proc. ICLR*, 2022.
- [32] C. Wang *et al.*, “Supervision-Guided Codebooks for Masked Prediction in Speech Pre-training,” in *Proc. Interspeech*, 2022, pp. 2643–2647.
- [33] H. Y. Kim *et al.*, “ASBERT: Asr-specific self-supervised learning with self-training,” in *Proc. SLT*, 2023, pp. 9–14.
- [34] F. L. Kreyssig *et al.*, “Biased self-supervised learning for ASR,” *arXiv preprint arXiv:2211.02536*, 2022.
- [35] M. Ott *et al.*, “Fairseq: A fast, extensible toolkit for sequence modeling,” in *Proceedings of NAACL-HLT 2019: Demonstrations*, 2019.
- [36] C. R. Harris *et al.*, “Array programming with NumPy,” *Nature*, vol. 585, no. 7825, pp. 357–362, 2020.
- [37] J. Kahn *et al.*, “Libri-Light: A benchmark for asr with limited or no supervision,” in *Proc. ICASSP*, 2020, pp. 7669–7673.
- [38] S. Watanabe *et al.*, “ESPnet: End-to-end speech processing toolkit,” in *Proc. Interspeech*, 2018, pp. 2207–2211.
- [39] S. Watanabe *et al.*, “Hybrid CTC/attention architecture for end-to-end speech recognition,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [40] Y. Peng *et al.*, “Branchformer: Parallel MLP-attention architectures to capture local and global context for speech recognition and understanding,” 2022.
- [41] A. Graves *et al.*, “Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks,” in *Proc. ICML*, 2006, pp. 369–376.
- [42] R. Wang *et al.*, “LightHuBERT: Lightweight and Configurable Speech Representation Learning with Once-for-All Hidden-Unit BERT,” in *Proc. Interspeech*, 2022, pp. 1686–1690.