



MT4SSL: Boosting Self-Supervised Speech Representation Learning by Integrating Multiple Targets

Ziyang Ma¹, Zhisheng Zheng¹, Changli Tang², Yujin Wang², Xie Chen^{1*}

¹ MoE Key Lab of Artificial Intelligence, AI Institute,
X-LANCE Lab, Department of Computer Science and Engineering,
Shanghai Jiao Tong University, Shanghai, China

² Department of Electronic Engineering, Tsinghua University, Beijing, China

{zym.22, chenxie95}@sjtu.edu.cn

Abstract

In this paper, we provide a new perspective on self-supervised speech models from how the training targets are obtained. We generalize the targets extractor into Offline Targets Extractor (Off-TE) and Online Targets Extractor (On-TE). Based on this, we propose a new multi-tasking learning framework for self-supervised learning, MT4SSL, which stands for Boosting Self-Supervised Speech Representation Learning by Integrating Multiple Targets. MT4SSL uses the K-means algorithm as an Off-TE and a teacher network without gradients as an On-TE, respectively. Our model outperforms previous SSL methods by nontrivial margins on the LibriSpeech benchmark, and is comparable to or even better than the best-performing models with fewer data. Furthermore, we find that using both Off-TE and On-TE results in better convergence in the pre-training phase. With both effectiveness and efficiency, we think doing multi-task learning on self-supervised speech models from our perspective is a promising trend. Code is available at <https://github.com/ddlBoJack/MT4SSL>.

Index Terms: self-supervised learning, representation learning, multi-task learning, speech recognition

1. Introduction

Self-supervised learning (SSL) has achieved remarkable success in the field of representation learning, applied in computer vision [1, 2], natural language processing [3, 4], as well as speech processing [5, 6]. For speech representation learning, SSL methods are often used in the pre-training phase to obtain supervisory signals from massive unlabeled audio data.

A core challenge for SSL is to obtain high-quality self-learning targets. Early works directly use input audio as training targets, either contrasting positive samples with negative ones [7, 8], or reconstructing the raw waves [9] and acoustic features [10, 11]. Contemporary works explore many other ways to accomplish this goal, including quantization by quantizers [12, 5], clustering by the K-means algorithm [6, 13], and generation by models [14].

In this work, we uniformly refer to the targets extraction module as the **Targets Extractor (TE)**. Intuitively, we find that all TEs used in SSL models can be easily divided into two categories:

- Offline Targets Extractor (Off-TE). Off-TE is trained in advance, or is an off-the-shelf algorithm. Off-TE will not be updated in the self-supervised pre-training phase.
- Online Targets Extractor (On-TE). On-TE can be trained in advance, or randomly initialized. On-TE will be continuously updated during the pre-training process.

It is obvious that targets from Off-TEs are more coarse-grained, and models using Off-TEs are easier to train. While On-TEs provide finer-and-finer-grained targets during training like curriculum learning does. Models with On-TEs might get better performance [14]. Another observation is that there is a clear complementarity on SUPERB¹ between HuBERT using Off-TE and data2vec using On-TE. Data2vec excels at content-related tasks, while HuBERT works better on speaker-related tasks. Therefore, we hope that both TEs can work together in a multi-task learning framework.

We present **MT4SSL**, short for **Multiple Targets for Self-Supervised Learning**, to optimize the model with targets extracted from Offline Targets Extractor (Off-TE) and targets extracted from Online Targets Extractor (On-TE) simultaneously. Our design is similar to adopt the state-of-the-art models HuBERT [6] and data2vec [14]. Actually, any two or more models using Off-TEs and On-TEs can be integrated into this framework. We conduct experiments on the LibriSpeech benchmark [15]. With 360 hours of unlabeled data, Our model achieves an average of 10% relative WER reduction over the best-performing HuBERT and data2vec on the 1-hour, 10-hour, and 100-hour fine-tuning subsets, and is comparable to or even better than wav2vec 2.0, HuBERT, and WavLM pre-trained with more data on the LibriSpeech 960h dataset. Furthermore, we find that MT4SSL can significantly speed up the convergence in the pre-training phase. The main contributions of this paper can be summarized as three-fold:

1. We provide a new perspective on self-supervised speech models from how the self-training targets are obtained, and generalize the Targets Extractor (TE) into Offline Targets Extractor (Off-TE) and Online Targets Extractor (On-TE).
2. We propose MT4SSL, a new multi-task learning framework that equips the model with both Off-TE and On-TE. We pre-train the model with targets extracted by both TEs and achieve better results than either alone on the LibriSpeech benchmark.
3. We find that MT4SSL has good convergence. Compared with other models, relatively low WER on the speech recognition task can be obtained with only a few pre-training steps. We give a possible explanation for the good performance and convergence. We hope our findings could inspire researchers to develop more powerful self-supervised methods in the speech community.

2. Related Works

In this section, we introduce the progress of two key technologies in MT4SSL, including self-supervised learning (SSL) and

Corresponding author*.

¹<https://superbenchmark.org/leaderboard>

multi-task learning (MTL).

2.1. SSL on Speech Representation Learning

Self-supervised pre-training followed by supervised fine-tuning becomes a mainstream approach for speech representation learning. This training paradigm has been shown to obtain universal representations on a wide variety of speech downstream tasks [16]. Traditional categorization is to divide self-supervised methods into **1) contrastive learning**, **2) predictive learning** and **3) generative learning** based on the pretext tasks. Contrastive learning aims to distinguish positive samples from negative ones. CPC [7] is the first successful representation learning approach for speech using contrastive learning, maximizing the mutual information between the input signal and the learned latent variables. Works that follow this paradigm include wav2vec [8], vq-wav2vec [12] and wav2vec 2.0 [5]. Predictive learning aims to predict the pre-clustered or model-generated targets with the input representations. HuBERT [6] predicts the discrete targets clustered by the K-means algorithm of the masked regions with a BERT-like method. To learn better representations, WavLM [13] polishes HuBERT by employing gated relative position bias and utterance mixing training strategy. ILS-SSL [17] promotes HuBERT by adding an additional loss on the intermediate layers. PBERT [18] and HuBERT-AP [19] refine HuBERT by refining target quality. Another predictive method is data2vec [14], which generates high-quality targets for masked positions with a teacher model fed the same utterance. Generative learning aims to reconstruct the whole speech from latent variables with an auto-encoding model. These works focus on recovering discrete [20] or continuous [21] speech signals using a variational autoencoder (VAE) or Seq2Seq autoencoder (SA) model to obtain representations of speech. Works that follow this paradigm include autoregressive models [9, 22] and non-autoregressive models [10, 11].

2.2. MTL for Speech Representation Learning

There have been some works on MTL for speech representation learning. Early works enumerate as many self-supervised tasks as possible empirically to conduct MTL. PASE [23] and PASE+ [24] solve many self-supervised tasks, such as waveform generation and prosody regression. Some concurrent works solve the problem with different pretext tasks. W2v-BERT [25] combines contrastive learning and mask language modeling to improve speech representation ability. They use the same targets extractor (quantizer) to obtain targets for different pretext tasks. UniSpeech [26] combines self-supervised learning and supervised learning to improve ASR performance. TESSP [27] and SpeechLM [28] enhance the representation ability of speech by introducing the mask language modeling task with paired or unpaired speech-text data.

3. Method

The overall architecture of *MT4SSL* is shown in Figure 1. For the backbone network, we simply use the same architecture as mainstream baseline models [5, 6, 14], which contains an *encoder network* and a *context network*. For the Target Extractor, we use the K-means algorithm as an *Off-TE* as HuBERT does, and use a teacher network without gradients as an *On-TE* as data2vec does.

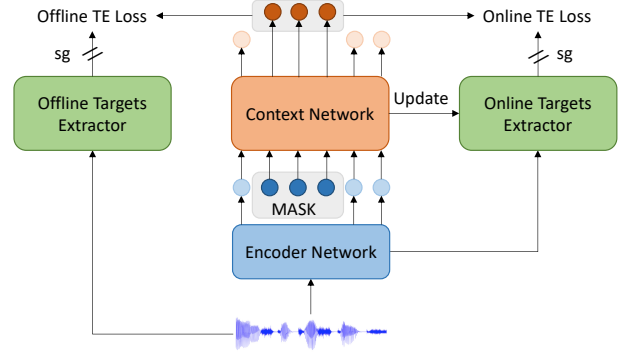


Figure 1: The overall framework of *MT4SSL*. The backbone of the model consists of an encoder network and a context network. The core of the framework is to optimize the model with multiple types of targets extracted from *Off-TE* and *On-TE*. *sg* means stop-gradient.

3.1. Backbone Network

The encoder network is a 7-layer 1-D convolutional neural network with kernel sizes (5, 2, 2, 2, 2, 2) and strides (10, 3, 3, 3, 3, 2). Given the raw audio input \mathcal{X} at a 16000 Hz sample rate, we downsample the audio with the encoder network denoted with $f : \mathcal{X} \mapsto \mathcal{H}$. The output representations \mathcal{H} are 50 Hz with dimension 512. Then we apply a linear projection for dimension transformation from 512 to 768, followed by the mask matrix to construct the input of the context network, denoted with $m : \mathcal{H} \mapsto \tilde{\mathcal{H}}$. The context network is a 12-layer standard Transformer with learnable convolutional positional encoding. Each Transformer block is set to 768 model dimension, 3072 inner dimension, and 12 attention heads. The context network can be denoted with $g : \tilde{\mathcal{H}} \mapsto \mathcal{Z}$. The final output \mathcal{Z} of the backbone is used for classification and regression, which will be detailed in Section 3.2 and Section 3.3.

3.2. Offline Targets Extractor

The model parameters of Offline Targets Extractor (*Off-TE*) do not update during the pre-training phase. Here we use the K-means algorithm as the *Off-TE*. We train a model with the K-means algorithm before pre-training. The model transforms speech features into C clusters. Thus we use the indices of cluster centers to represent each speech token. Suppose X is a raw audio utterance sampled from \mathcal{X} , the offline targets can be obtained by:

$$Y^f = \mathbf{TE}_f(X), \quad (1)$$

where $\mathbf{TE}_f(\cdot)$ is the extracting operation, and $Y^f = [y_1^f, \dots, y_T^f]^T$ is the self-training offline targets and $y_t^f \in \mathbb{R}^C$ in a one-hot form. Suppose $Z = [z_1, \dots, z_T]^T$ is a masked version obtained through the backbone network from the same raw audio X . $M \subset [T]$ denote the masked indices, and z_t is replaced with a mask token if $t \in M$. We use a projection layer for dimension transformation, which can be written as:

$$Z^f = \mathbf{W}^f Z, \quad (2)$$

where $Z^f = [z_1^f, \dots, z_T^f]^T$ and $z_t^f \in \mathbb{R}^C$. The offline TE loss is defined as:

$$\mathcal{L}^f = \mathbf{CE}(Z^f, Y^f), \quad (3)$$

where $\mathbf{CE}(\cdot)$ computes the cross entropy loss between sources and targets.

3.3. Online Targets Extractor

The model parameters of Online Targets Extractor (On-TE) update continuously during the pre-training phase. We use a teacher network without gradients as the On-TE to obtain online targets. This process can be viewed as a special kind of knowledge distillation [29], or Noisy Student Training (NST) [30]. Suppose H is the hidden representation sampled from \mathcal{H} , which is obtained from \mathcal{X} followed by convolutional subsampling, the online targets can be obtained by:

$$Y^n = \mathbf{TE}_n(H), \quad (4)$$

where $\mathbf{TE}_n(\cdot)$ is the extracting operation, and $Y^n = [y_1^n, \dots, y_T^n]^T$ is the self-training targets. As with offline speech tokens, we use a projection layer for dimension transformation, which can be written as:

$$Z^n = \mathbf{W}^n Z, \quad (5)$$

where Z^n has the same dimension as Y^n . The online TE loss is defined as:

$$\mathcal{L}^n = \text{MSE}(Z^n, Y^n), \quad (6)$$

where $\text{MSE}(\cdot)$ measures the mean squared error between sources and targets. The parameters of the teacher network Δ are initialized with the parameters of the backbone network θ , and the parameters of the teacher network are updated with exponentially moving average (EMA) [29] within each mini-batch, denoted as:

$$\Delta = \tau\Delta + (1 - \tau)\theta, \quad (7)$$

where τ is a parameter that increases linearly during pre-training.

3.4. Loss Function

The core of our model is to integrate multiple targets, thereby enhancing the representational ability of self-supervised learning. The proposed objective is formulated as follows by using both offline TE loss in Eq. 3 and online TE loss in Eq. 6:

$$\mathcal{L} = \mathcal{L}^f + \alpha\mathcal{L}^n, \quad (8)$$

with a tunable parameter α . Note that we only compute loss on the masked parts of the utterance.

4. Experiments

4.1. Dataset

For unsupervised pre-training, we use LibriSpeech [15] corpus with 360-hour unlabeled data (train-clean-360). For supervised fine-tuning, 1-hour, 10-hour splits from Libri-light [31] corpus and 100-hour from LibriSpeech corpus are considered. We conduct model evaluation according to the mainstream test sets: dev-clean/other and test-clean/other from the LibriSpeech corpus.

4.2. Setup

Our MT4SSL model can be considered as a simplification and fusion of the HuBERT model and the data2vec model, so we maximized the inheritance of the hyperparameters of both to demonstrate the effectiveness of the model. Given the limited computing resources, we simply choose some empirical configuration for the training of MT4SSL model without conducting extensive hyperparameter search.

Model	Language Model	WER%(↓)			
		dev		test	
		clean	other	clean	other
<i>1-hour Labeled Data</i>					
HuBERT	None	29.4	37.1	30.0	37.8
data2vec		24.1	32.6	24.2	33.5
MT4SSL		19.7	27.1	20.1	27.6
HuBERT	4-gram	7.2	16.5	7.9	17.3
data2vec		6.8	15.8	7.2	17.0
MT4SSL		5.5	13.1	5.9	13.9
<i>10-hour Labeled Data</i>					
HuBERT	None	11.4	21.3	11.6	22.5
data2vec		10.8	20.9	10.7	21.9
MT4SSL		9.4	18.0	9.3	18.5
HuBERT	4-gram	4.9	12.7	5.3	13.4
data2vec		5.1	13.8	5.5	14.5
MT4SSL		4.0	11.4	4.5	11.9
<i>100-hour Labeled Data</i>					
HuBERT	None	5.8	15.5	6.0	15.3
data2vec		5.5	15.7	5.6	16.0
MT4SSL		5.1	14.3	5.1	14.3
HuBERT	4-gram	2.9	9.6	3.5	9.9
data2vec		3.0	10.6	3.6	11.2
MT4SSL		2.7	9.6	3.4	9.6

Table 1: Word error rate (WER) on LibriSpeech corpus. We compare the performance on four subsets (dev-clean, dev-other, test-clean, test-other) with (4-gram) and without (None) language model pre-trained on 360 hours of unlabeled data (train-clean-360) and fine-tuned on different amounts of labeled data (1h, 10h, 100h).

Model	Unlabeled Data	WER%(↓)			
		dev		test	
		clean	other	clean	other
<i>1-hour Labeled Data</i>					
wav2vec 2.0[5]	LS-960	24.1	29.6	24.5	29.7
HuBERT[6]	LS-960	24.3	30.2	20.9	27.5
WavLM[13]	LS-960	-	-	24.5	29.2
MT4SSL	LS-360	19.7	27.1	20.1	27.6
<i>10-hour Labeled Data</i>					
wav2vec 2.0[5]	LS-960	10.9	17.4	11.1	17.6
HuBERT[6]	LS-960	12.0	18.1	10.1	16.8
WavLM[13]	LS-960	-	-	9.8	16.0
MT4SSL	LS-360	9.4	18.0	9.3	18.5
<i>100-hour Labeled Data</i>					
wav2vec 2.0[5]	LS-960	6.1	13.5	6.1	13.3
HuBERT[6]	LS-960	5.5	13.0	6.3	13.2
WavLM[13]	LS-960	-	-	5.7	12.0
MT4SSL	LS-360	5.1	14.3	5.1	14.3

Table 2: Word error rate (WER) on LibriSpeech corpus. The results for wav2vec 2.0 and WavLM are obtained from their papers. The results for HuBERT are obtained by fine-tuning their public released model². All results are reported without the language model.

Pre-Training. In the pre-training phase, we train the model with 360 hours LibriSpeech unlabeled data. The training is conducted on NVIDIA GeForce RTX 3090 GPUs, and we simulate 16 GPUs by using k GPUs and setting update frequency with

²<https://github.com/facebookresearch/fairseq/tree/main/examples/hubert>

$16/k$. k is set to 4 in this paper. For the mask strategy, each time step has a probability of $p = 0.065$ to be the starting index and the subsequent $l = 10$ time-steps are masked. This results in the mask embedding covering 49% of all tokens on average. For the optimizing strategy, we use Adam [32] with a learning rate of 0.0005 and a weight decay of 0.01. We train MT4SSL for 800 epoch, with [3%, 90%, 7%] proportion of warm-up, hold-on, and linearly decay. The hyperparameter α that controls the loss weight is set to 1, which means the two losses have the same weight.

For the offline targets, we obtain features from HuBERT model and train an Off-TE before pre-training. Concretely, the targets are obtained by running the K-means clustering with 500 clusters on the 6-th transformer layer output of the first iteration HuBERT model.

For the online targets, we use the average of the top 8 blocks of the transformer layer outputs from the teacher network as data2vec model designs. For the parameter update of the teacher model, we apply a linearly increasing strategy for τ from $\tau_s = 0.99$ to $\tau_e = 0.999$ for the first 7.5% training steps. The parameter τ is kept constant for the remainder of training steps.

Fine-Tuning. In the fine-tuning phase, we use Connectionist Temporal Classification (CTC) [33] loss to keep consistent with the baseline models. The hyper-parameters of the fine-tuning stage are still kept consistent with the mainstream models [5, 6, 14].

4.3. Results

Table 1 shows results of the MT4SSL on the LibriSpeech benchmark compared to other state-of-the-art models. The models are pre-trained on LibriSpeech 360 hours dataset (train-clean-360), and fine-tuned on Libri-light 1-hour, 10-hour and LibriSpeech 100-hour subsets. We compare the performance on dev-clean/other and test/other with and without the language model, respectively. We adopt the 4-gram language model trained on the official LibriSpeech language modeling data. Given 10 hours of labeled data, MT4SSL can achieve 13.0% (dev-clean), 13.9% (dev-other), 13.1% (test-clean) and 15.5% (test-other) relative WER reduction over the best-performing model without a language model, and 18.4% (dev-clean), 10.2% (dev-other), 15.1% (test-clean) and 11.2% (test-other) relative WER reduction with a 4-gram language model. For the fine-tuning on 1-hour and 100-hour labeled data, the improvements are consistent for the MT4SSL over other models.

Table 2 shows results of MT4SSL trained with 360 hours of audio data, compared to the state-of-the-art models trained with 960 hours. We compare results of MT4SSL with wav2vec 2.0 and WavLM from their papers, and HuBERT from their public released page. Despite using less data, our model is comparable to or even better than the state-of-the-art models. The performance on noisy subsets is less competitive, and one possible explanation is that our pre-training data (train-clean-360) only consists of clean audio.

4.4. Convergence Analysis

In this section, we analyze our MT4SSL in terms of model convergence quantitatively and qualitatively. All experiments are carried out with the following configuration: all the models are pre-trained on 360 hours of LibriSpeech unlabeled data, fine-tuned on 10 hours of Libri-light labeled data, and evaluated on dev-other subset of LibriSpeech corpus.

We find that MT4SSL not only achieves amazing results on




Figure 2: Word error rate (WER) along with pre-training epochs among different SSL models. The models are fine-tuned on 10-hour labeled data and evaluated on LibriSpeech dev-other dataset without adopting the language model.

the benchmark, but also exhibits good convergence. As shown in Figure 2, we plot the WER trend with respect to the number of pre-training epochs. By comparing MT4SSL with data2vec and HuBERT, it can be seen that HuBERT which utilizes offline targets has better convergence than data2vec which utilizes online targets. However, data2vec achieves better performance than HuBERT when fully trained. MT4SSL combines their advantages and converges to a relatively low WER quickly.

One possible explanation is that the fixed offline targets are less difficult to learn for the model to learn than the ever-changing online targets. Hence, the model which adopts offline targets converges faster. However, the online targets have finer granularity than the offline targets, so the model which uses online targets has better representation capabilities. The learning of the two targets may not be conflicting but cooperative, so MT4SSL can achieve both efficiency and effectiveness.

5. Discussion

In this work, a preliminary attempt to combine Off-TE using the K-means algorithm and On-TE using a randomly initialized teacher has yielded good results. There are several interesting aspects to explore:

- Will there be On-TEs and Off-TEs that cooperate better?
- Will using different targets extractors at different stages of pre-training improve efficiency, effectiveness, or both?
- Is MTL in SSL a better way to obtain universal features in various speech-related downstream tasks?

We will research the above problems in the future.

6. Conclusion

In this paper, we present a new framework for speech-based self-supervised learning, which is named MT4SSL, by simultaneously optimizing the model with offline targets and online targets, without caring about specific pretext tasks. The proposed method improves both performance and convergence upon the state-of-the-art models.

7. Acknowledgements

This work was supported by the National Natural Science Foundation of China (No. 62206171), the International Cooperation Project of PCL, and Alibaba Group through Alibaba Innovative Research Program.

8. References

- [1] H. Bao, L. Dong, S. Piao, and F. Wei, "Beit: Bert pre-training of image transformers," in *Proc. of ICLR*, 2021.
- [2] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proc. of CVPR*, 2022.
- [3] J. D. M.-W. C. Kenton and L. K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *Proc. of NAACL*, 2019.
- [4] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.
- [5] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Proc. of NeurIPS*, 2020.
- [6] 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," *Trans. of TASLP*, 2021.
- [7] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [8] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Unsupervised pre-training for speech recognition," *Proc. Interspeech*, 2019.
- [9] Y.-A. Chung, W.-N. Hsu, H. Tang, and J. Glass, "An unsupervised autoregressive model for speech representation learning," *Proc. Interspeech*, 2019.
- [10] A. T. Liu, S.-W. Li, and H.-y. Lee, "Tera: Self-supervised learning of transformer encoder representation for speech," *Trans. of TASLP*, 2021.
- [11] S. Ling, Y. Liu, J. Salazar, and K. Kirchhoff, "Deep contextualized acoustic representations for semi-supervised speech recognition," in *Proc. of ICASSP*, 2020.
- [12] A. Baevski, S. Schneider, and M. Auli, "vq-wav2vec: Self-supervised learning of discrete speech representations," in *Proc. of ICLR*, 2019.
- [13] 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," *J. of JSTSP*, 2022.
- [14] 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," *Proc. of ICML*, 2022.
- [15] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *Proc. of ICASSP*, 2015.
- [16] 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," *Proc. of Interspeech*, 2021.
- [17] C. Wang, Y. Wu, S. Chen, S. Liu, J. Li, Y. Qian, and Z. Yang, "Improving self-supervised learning for speech recognition with intermediate layer supervision," in *Proc. of ICASSP*, 2022.
- [18] C. Wang, Y. Wang, Y. Wu, S. Chen, J. Li, S. Liu, and F. Wei, "Supervision-guided codebooks for masked prediction in speech pre-training," *Proc. of Interspeech*, 2022.
- [19] S. Ren, S. Liu, Y. Wu, L. Zhou, and F. Wei, "Speech pre-training with acoustic piece," *Proc. of Interspeech*, 2022.
- [20] J. Chorowski, R. J. Weiss, S. Bengio, and A. Van Den Oord, "Unsupervised speech representation learning using wavenet autoencoders," *Trans. of TASLP*, 2019.
- [21] Y.-C. Chen, S.-F. Huang, H.-y. Lee, Y.-H. Wang, and C.-H. Shen, "Audio word2vec: Sequence-to-sequence autoencoding for unsupervised learning of audio segmentation and representation," *Trans. of TASLP*, 2019.
- [22] Y.-A. Chung, H. Tang, and J. Glass, "Vector-quantized autoregressive predictive coding," *Proc. Interspeech*, 2020.
- [23] S. Pascual, M. Ravanelli, J. Serrà, A. Bonafonte, and Y. Bengio, "Learning problem-agnostic speech representations from multiple self-supervised tasks," *Proc. Interspeech*, 2019.
- [24] M. Ravanelli, J. Zhong, S. Pascual, P. Swietojanski, J. Monteiro, J. Trmal, and Y. Bengio, "Multi-task self-supervised learning for robust speech recognition," in *Proc. of ICASSP*, 2020.
- [25] 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 *Proc. of ASRU*, 2021.
- [26] C. Wang, Y. Wu, Y. Qian, K. Kumtani, S. Liu, F. Wei, M. Zeng, and X. Huang, "Unispeech: Unified speech representation learning with labeled and unlabeled data," in *Proc. of ICML*, 2021.
- [27] Z. Yao, S. Ren, S. Chen, Z. Ma, P. Guo, and L. Xie, "Tessp: Text-enhanced self-supervised speech pre-training," *arXiv preprint arXiv:2211.13443*, 2022.
- [28] Z. Zhang, S. Chen, L. Zhou, Y. Wu, S. Ren, S. Liu, Z. Yao, X. Gong, L. Dai, J. Li *et al.*, "Speechlm: Enhanced speech pre-training with unpaired textual data," *arXiv preprint arXiv:2209.15329*, 2022.
- [29] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. Guo, M. Gheshlaghi Azar *et al.*, "Bootstrap your own latent-a new approach to self-supervised learning," *Proc. of NeurIPS*, 2020.
- [30] D. S. Park, Y. Zhang, Y. Jia, W. Han, C.-C. Chiu, B. Li, Y. Wu, and Q. V. Le, "Improved noisy student training for automatic speech recognition," *Proc. Interspeech*, 2020.
- [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 *Proc. of ICASSP*, 2020.
- [32] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. of ICLR*, 2015.
- [33] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proc. of ICML*, 2006.