

Wav-BERT: Cooperative Acoustic and Linguistic Representation Learning for Low-Resource Speech Recognition

Guolin Zheng^{1*}, Yubei Xiao^{1*}, Ke Gong², Pan Zhou³, Xiaodan Liang¹, Liang Lin^{1,2†}

¹Sun Yat-sen University, China

²Dark Matter AI Research, ³Sea AI Lab

{zhengglin, xiaoyb5}@mail2.sysu.edu.cn

kegong936@gmail.com, zhoupan@sea.com,

xdliang328@gmail.com, linliang@ieee.org

Abstract

Unifying acoustic and linguistic representation learning has become increasingly crucial to transfer the knowledge learned on the abundance of high-resource language data for low-resource speech recognition. Existing approaches simply cascade pre-trained acoustic and language models to learn the transfer from speech to text. However, how to solve the representation discrepancy of speech and text is unexplored, which hinders the utilization of acoustic and linguistic information. Moreover, previous works simply replace the embedding layer of the pre-trained language model with the acoustic features, which may cause the catastrophic forgetting problem. In this work, we introduce Wav-BERT, a cooperative acoustic and linguistic representation learning method to fuse and utilize the contextual information of speech and text. Specifically, we unify a pre-trained acoustic model (wav2vec 2.0) and a language model (BERT) into an end-to-end trainable framework. A Representation Aggregation Module is designed to aggregate acoustic and linguistic representation, and an Embedding Attention Module is introduced to incorporate acoustic information into BERT, which can effectively facilitate the cooperation of two pre-trained models and thus boost the representation learning. Extensive experiments show that our Wav-BERT significantly outperforms the existing approaches and achieves state-of-the-art performance on low-resource speech recognition.

1 Introduction

Recently, Automatic Speech Recognition (ASR) has achieved remarkable success, which can be attributed to two complementary aspects: 1) designing more effective and larger deep neural networks for ASR, and 2) training on a large amount of data (Chan et al., 2016; Watanabe et al., 2017b;

Amodei et al., 2016). However, in practice, unlike the commonly used languages (e.g. English and Chinese) with sufficient training data, many other languages (e.g. Swahili, Tamil) have only low-resource data due to the scarcity of audios and the huge labor resources consumed in transcription. In this way, the aforementioned data-driven mechanism is impractical for low-resource languages and thus suffers from unsatisfactory performance.

To resolve this learning difficulty in the low-resource domain, many efforts have been devoted to leveraging unlabeled data. One mainstream research paradigm is unsupervised pre-training, or representation learning, which has achieved great success in natural language processing (Devlin et al., 2018; Peters et al., 2018) and received increasing attention in speech recognition (Oord et al., 2018; Schneider et al., 2019a). As a representation in this line, wav2vec (Schneider et al., 2019a) and wav2vec 2.0 (Baevski et al., 2020) apply unsupervised contrastive pre-training and show promising results. To utilize linguistic information, some works (Chiu and Chen, 2021; Shin et al., 2019) also aim to build language models to rescore the N -best hypotheses generated by acoustic models. The most recent approach (Yi et al., 2021) even cascaded the pre-trained wav2vec 2.0 and BERT into a single model for low-resource ASR.

However, there leave two critical challenges on how to integrate the acoustic model and language model to utilize the contextual information of speech and text. 1) Representation discrepancy: the acoustic model focuses more on local dependencies of the speech sequence, while the language model aims at capturing long-term semantic information of texts. It is desired to explore an effective model to fuse and leverage the two kinds of representation. 2) Embedding inconsistency: The language model applies a token embedding layer during pre-training but previous methods (Yi et al., 2021) simply replace the embedding layer

* Both authors contributed equally to this work.

† Liang Lin is the corresponding author of this work.

with the features generated by the acoustic model, which may result in the catastrophic forgetting problem (Goodfellow et al., 2013).

To tackle the above challenges, in this work, we make the first attempt to successfully integrate the well-trained acoustic model and language model for low-resource speech recognition. Towards this end, we introduce a new framework that incorporates the two kinds of pre-trained models for cooperative acoustic and linguistic representation learning by exploiting complementary contextual information of both speech and text.

First, to solve representation discrepancy, unlike the previous works (Yi et al., 2021; Yu and Chen, 2021) that simply connect the acoustic model and the language model by treating them as an encoder and a decoder, we consider them as two encoders that provide two different representations. Specifically, we propose a Representation Aggregation Module, a plug-in component to better exploit and fuse the acoustic and linguistic information. We design and evaluate several representation aggregation mechanisms, including Gated Acoustic-Guided Attention, Gated Linguistic-Guided Attention, and Gated Cross-Modal Attention. The experimental results show the proposed Gated Cross-Modal Attention is the most effective method for representation aggregation.

Second, to fill the gap of embedding inconsistency, we introduce an Embedding Attention Module to incorporate the acoustic features into BERT by a gated attention process, which not only preserves the capability of BERT but also takes advantage of acoustic information. Moreover, as BERT requires audio transcripts as input to create word embedding, it may be easy to overfit when using ground truth transcripts. On the other hand, it is also hard to converge when using transcripts predicted by the acoustic model. To facilitate the cooperation of the two encoders, we propose a sampling strategy with decay to randomly select the ground truth and generated transcripts for smooth training.

We adopt pre-trained wav2vec 2.0 (Baevski et al., 2020) and BERT (Devlin et al., 2018) as the encoders to provide acoustic and linguistic representations respectively for their flexible pre-training then fine-tuning paradigm as well as excellent local contextual modeling ability. Accordingly, we denominate our method as Wav-BERT.

We evaluate our method on several datasets with diverse languages from the public IARPA BABEL

dataset (Gales et al., 2014) and AISHELL-1 corpus (Bu et al., 2017). The experimental results demonstrate that our Wav-BERT significantly outperforms the existing approaches on low-resource ASR. Furthermore, our exhaustive ablation studies demonstrate the effectiveness of the proposed mechanisms for cooperative acoustic and linguistic representations learning. We hope this work will be useful for the community on the way to explore different pre-trained models for low-resource ASR.

2 Related Work

2.1 Low resource speech recognition

To tackle the low-resource ASR task, transfer learning ASR (Kunze et al., 2017) and multilingual transfer learning ASR (Dalmia et al., 2018; Watanabe et al., 2017a; Toshniwal et al., 2018) are explored via using different source languages to improve the performance of low-resource languages. Meta-learning approaches (Finn et al., 2017; Nichol et al., 2018) are also adopted for low-resource ASR (Hsu et al., 2020; Xiao et al., 2021) to obtain fast adaptation ability to new tasks with only a few data through meta-learning a model initialization from training tasks. In addition, recent works utilize unsupervised pre-training (Schneider et al., 2019b; Chung and Glass, 2020) and semi-supervised learning (Kahn et al., 2020; Li et al., 2019) to exploit a large amount of unlabeled data to learn general representations for low-resource adaptation. Among them, Wav2vec 2.0 (Baevski et al., 2020) achieved excellent results through self-supervised learning, which learns powerful and contextual acoustic representations of a large speech audio corpus by solving contrastive tasks that require identifying the true quantized latent speech representations for masked time steps. Then it shows strong feasibility of ultra-low resource speech recognition with even only 10 minutes of labeled data.

2.2 Speech recognition with BERT

To use the linguistic information from BERT (Devlin et al., 2018) for improving ASR performance, some works (Chiu and Chen, 2021; Shin et al., 2019; Wang and Cho, 2019) use BERT to re-rank the N-best hypotheses generated by the ASR model. Besides, knowledge distillation (Futami et al., 2020) is explored to use BERT as a teacher model to guide ASR model training. Moreover, some recent works (Yi et al., 2021; Yu and Chen,




Figure 1: Comparison of the architectures of different approaches to fuse BERT into the ASR model. (a) Rescoring methods use BERT to rescore N -best hypotheses generated by wav2vec 2.0 ASR (Shin et al., 2019). (b) Cascade methods directly cascade the BERT decoder on the top of the wav2vec 2.0 encoder through Length Alignment module (Yi et al., 2021). (c) Adapter-BERT inserts adapter modules in each BERT layer (Guo et al., 2020). (d) Our Wav-BERT introduces a Representation Aggregation Module for aggregate acoustic and linguistic representation and an Embedding Attention Module to incorporate acoustic information into text embedding.

2021; Winata et al., 2020) further combine BERT with the ASR model into a unified model and train the model in an end-to-end way. But Yi et al. and Yu et al. both simply connect BERT and the ASR model in series without considering the contextual information of speech and text (Yi et al., 2021; Yu and Chen, 2021). Winata et al. (Winata et al., 2020) modified mBERT model into an auto-regressive decoder and insert a cross-attention layer in each mBERT layer, but the deep bidirectional information of pre-trained BERT cannot be fully utilized in the auto-regressive mode.

3 Preliminaries

Here we briefly introduce the architectures of acoustic and linguistic encoders in our framework. **Wav2vec 2.0.** We adopt wav2vec 2.0 (Baevski et al., 2020) as our acoustic encoder because of its effectiveness and efficiency. It has two stages: (i) contrastive pre-training to learn representations of speech and (ii) fine-tuning to adapt the learned representations on labeled data with connectionist temporal classification (CTC) loss (Graves et al., 2006b) for downstream speech recognition tasks. In this work, we aim to utilize the public pre-trained model and mainly focus on the fine-tuning stage. The architecture of wav2vec 2.0 contains a feature encoder, a context network with a transformer and a quantization module. During fine-tuning, the quantization module is removed and a randomly initialized linear projection layer is attached on top of the context network.

BERT. BERT (Devlin et al., 2018) is employed as our linguistic encoder since it is one of the most popular text pre-training approaches and has shown remarkable performance in many downstream natural language processing tasks. It also consists of

two steps: (i) self-supervised pre-training to learn deep bidirectional linguistic representations from a large text corpus and (ii) fine-tuning to adapt to downstream tasks using labeled data. BERT consists of an embedding table, a multi-layer bidirectional Transformer encoder, and an additional output layer for fine-tuning.

4 Wav-BERT

4.1 Motivation

To transfer the knowledge learned on the abundance of high-resource language data for low-resource speech recognition, many efforts have been devoted to unifying acoustic and linguistic representation learning. We first categorize previous methods and then introduce our solution.

As shown in Figure 1 (a), one simplest way to fuse BERT into an acoustic model in speech recognition is rescoring (Chiu and Chen, 2021; Shin et al., 2019). It uses BERT as a language model to calculate the pseudo-log-likelihood scores of text sentences for reranking the N -best hypotheses generated by the acoustic model. However, this process is time-consuming as it needs to iteratively mask each word in the sentence for inference and then sum up the scores of all masked words. It also requires tuning many hyper-parameters by repetitive experiments, e.g. beam size, balanced weights of the language and acoustic models.

Recently, some works (Yi et al., 2021; Yu and Chen, 2021) directly cascade the decoder BERT on the top of the acoustic encoder, as illustrated by Figure 1 (b). However, such a simple cascade often cannot well fuse the contextual information of speech and text.

Inspired by AB-Net (Guo et al., 2020), we design Adapter-BERT that inserts cross-attention adapters

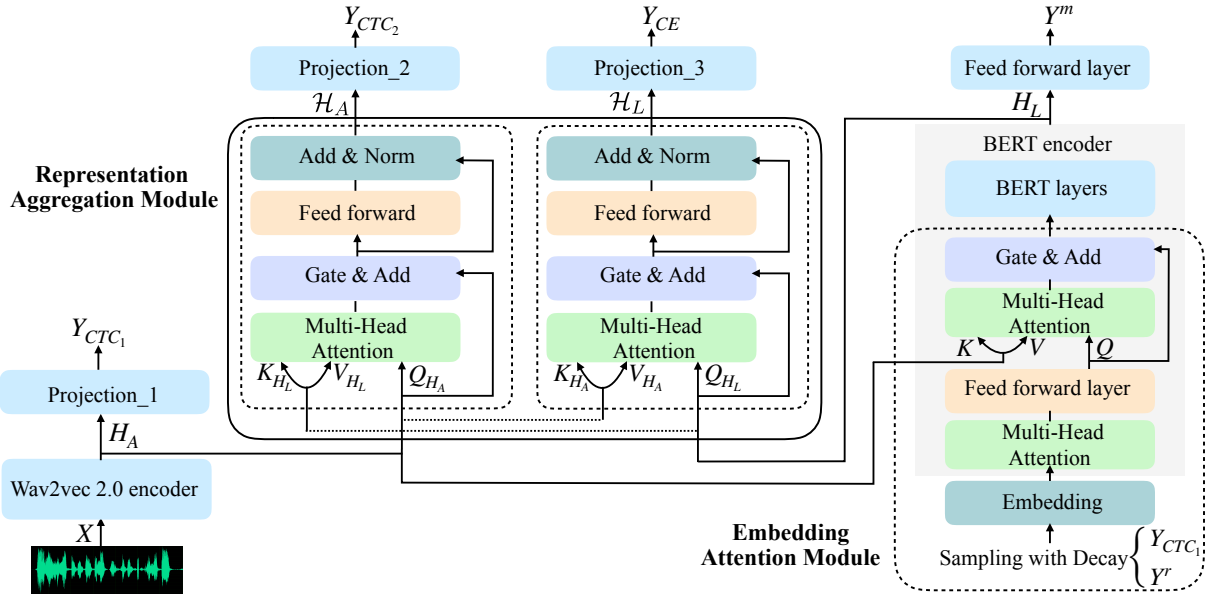


Figure 2: Our Wav-BERT framework, which is composed of two main parts: 1) Representation Aggregation Module that combines a Gated Acoustic-Guided Attention (Left) and a Gated Linguistic-Guided Attention (Right) to construct a Gated Cross-Modal Attention. 2) Embedding Attention Module that includes a Gated Attention and a "Sampling with Decay" mechanism.

in each BERT layer with the Mask-Predict algorithm (Ghazvininejad et al., 2019) to fully utilize the bidirectional information of the input sequence, as shown in Figure 1 (c). Nevertheless, the adapters in each layer of BERT will affect the pre-trained parameters of BERT, causing catastrophic forgetting. Moreover, the Mask-Predict decoding suffers from low inference speed.

To solve the representation discrepancy and embedding inconsistency between speech and text, in this work, we introduce Wav-BERT, a cooperative acoustic and linguistic learning framework that fuses and leverages the contextual information of speech and text from the representation level to the embedding level, as shown in Figure 1 (d). We first present an independent Representation Aggregation Fusion Module for acoustic and linguistic representation aggregation, without inserting it in any pre-trained model to avoid destroying the parameters of pre-trained models. Then, an Embedding Attention Module is introduced to better combine acoustic and linguistic embedding instead of simply replacement.

4.2 Our Wav-BERT

The architecture of our Wav-BERT is illustrated in Figure 2. Specifically, wav2vec 2.0 encoder takes raw waveform X as input and outputs acoustic representation H_A , which is then fed into a lin-

ear projection layer with CTC loss (Graves et al., 2006b) (L_{ctc_1}) and the Representation Aggregation Module respectively. For the input of BERT encoder, we employ "Sampling with Decay" mechanism to sample from the masked ground truth Y^r or wav2vec 2.0 CTC output Y_{CTC_1} with probability p and $1 - p$, so as to narrow the gap between training and inference. Next, word embedding E and acoustic embedding H_A are fed into the Gate Attention to model the conditional information from the wav2vec 2.0 encoder side. Through the subsequent BERT transformer layers, we get the linguistic representation H_L . Finally, the Representation Aggregation Module takes linguistic representation H_L as well as acoustic representation H_A as input, generating the CTC output Y_{CTC_2} and cross-entropy (CE) output Y_{CE} , supervised by the CTC (L_{ctc_2}) and CE (L_{ce}) criterion respectively. Simultaneously, the conditional masked language model (CMLM) objective (L_{cmlm}) (Guo et al., 2020) is also attached on BERT encoder followed by a feed-forward layer to supervise the BERT output Y^m . Overall, the objective of our framework is defined as:

where μ_1, μ_2, μ_3 and μ_4 are the corresponding loss weights.

4.2.1 Representation Aggregation Module

To solve representation discrepancy, we first design several representation aggregation mechanisms,

such as Gated Acoustic-Guided Attention, Gated Linguistic-Guided Attention. In our Representation Aggregation Module, we combine a Gated Acoustic-Guided Attention (Left) and a Gated Linguistic-Guided Attention (Right) to construct a Gated Cross-Modal Attention for better exploiting and aggregating the acoustic and linguistic representations.

Specifically, Gated Cross-Modal Attention Module takes acoustic representation H_A generated by wav2vec 2.0 as well as linguistic representation H_L generated by BERT as input and feeds them as the query, key, and value vector respectively to a multi-head attention, which can be formulated as:

$$C_A = \text{ATT}(Q_{H_A}, K_{H_L}, V_{H_L}), \quad (1)$$

$$C_L = \text{ATT}(Q_{H_L}, K_{H_A}, V_{H_A}), \quad (2)$$

where Q_{H_A} means passing H_A as query vector, K_{H_L} as well as V_{H_L} means passing H_L as key and value vector respectively. C_A is the acoustic guided context feature generated by attention which tend to focus on the values in the linguistic representation H_L related to acoustic representation H_A . Vice versa, C_L is the linguistic guided context feature to focus on the values in the H_A related to H_L .

Next, the context feature C_A and acoustic representation H_A are fed into a gated weighting layer to automatically capture the most important information between context and acoustic representation, and generating acoustic-guided linguistic representation H_{AGL} , which can be formulated as:

$$\Phi_A = \text{sigmoid}(W_1[C_A; H_A] + B_1), \quad (3)$$

$$H_{AGL} = H_A + \Phi_A C_A, \quad (4)$$

where W_1 as well as B_1 are model parameters and Φ_A is the gated weight.

Similarly, the context feature C_L and linguistic representation H_L are fed into another gated weighting layer to weigh the expected importance Φ_L and generate linguistic-guided acoustic representation H_{LGA} , which can be formulated as:

$$\Phi_L = \text{sigmoid}(W_2[C_L; H_L] + B_2), \quad (5)$$

$$H_{LGA} = H_L + \Phi_L C_L, \quad (6)$$

where W_2 as well as B_2 are model parameters and Φ_L is the gated weight.

We then feed H_{AGL} and H_{LGA} to a feed-forward layer followed by residual connection respectively and get aggregation representation \mathcal{H}_A as well as \mathcal{H}_L . Finally, two linear projection layers

are attached on the top of Representation Aggregation Module to get the Y_{CTC_2} and Y_{CE} . As the sequence length of Y_{CTC_2} is determined by acoustic representation H_A , we use CTC criterion to align the acoustic frames of Y_{CTC_2} to the ground truth tokens. On the other hand, the sequence length of Y_{CE} is determined by linguistic representation H_L , so we use CE criterion to align the text sequence of Y_{CE} to the ground truth transcript.

The different aggregation mechanisms including Gated Acoustic Guided Attention, Gated Linguistic-Guided Attention and Gated Cross-Modal Attention are evaluated and compared in Table 3.

4.2.2 Embedding Attention Module

Recent works (Yi et al., 2021; Yu and Chen, 2021) directly connect the BERT on the top of the acoustic encoder and simply replace the embedding layer with the acoustic features generated by the acoustic encoder, causing the catastrophic forgetting problem.

To fill the gap of embedding inconsistency, we propose the Embedding Attention Module and insert it behind the embedding layer of BERT to incorporate the acoustic information into the word embedding instead of simply replacing them. We first introduce a Gated Attention operation in this module. As shown in Figure 2, word embedding E generated by embedding layer is fed to a self-attention layer followed by a feed-forward layer to capture higher level linguistic embedding E_L . Then, a multi-head self-attention followed by a gated weighting layer takes E_L as the query vector and acoustic embedding H_A generated by wav2vec 2.0 as the key vector as well as value vector to fuse the linguistic embedding and acoustic embedding. Thus, as a conditional masked language model, BERT can learn to predict the masked word under the conditional acoustic information and provided enhanced linguistic representation.

Furthermore, for the input of the embedding layer of BERT, it is easy to overfit when using ground truth transcripts while it is hard to converge when using transcripts predicted by wav2vec2.0 encoder. To solve this issue, we propose a "Sampling with Decay" mechanism by feeding BERT either the masked ground truth transcript Y^r or the predicted CTC result Y_{CTC_1} with a certain probability during training. The probability p of selecting from Y^r decreases linearly as the number of training steps increases.

Table 1: Results of low resource ASR on IARPA BABEL in terms of CER (%).

Method	Pre-trained	Vi	Sw	Ta	Avg
Mono-BLSTMP (Cho et al., 2018)		54.3	33.1	55.3	47.6
Multi-BLSTMP (Cho et al., 2018)	-	41.0	-	48.5	44.8
Multi-BLSTMP+ VGG (Cho et al., 2018)		37.4	-	45.5	41.5
wav2vec 2.0 (Baevski et al., 2020)		21.8	15.5	29.3	22.2
wav2vec 2.0 w/ 4-gram (Baevski et al., 2020)	wav2vec 2.0 (Base)	21.1	14.9	29.9	22.0
XLSR-Monolingual (Conneau et al., 2020)		25.2	26.8	36.0	29.3
XLSR-10 (Conneau et al., 2020)		21.7	16.6	30.5	22.9
BERT rescoring (Shin et al., 2019)		21.3	15.3	29.1	21.9
Adapter-BERT (Guo et al., 2020)	w/ mBERT	22.5	17.6	29.8	23.3
w2v-cif-bert (Yi et al., 2021)		24.1	21.5	41.9	29.2
our Wav-BERT		19.5	14.8	28.8	21.0
XLSR-10 (Conneau et al., 2020)	wav2vec 2.0 (Large)	19.9	14.9	28.6	21.1
XLSR-53 (Conneau et al., 2020)		21.8	21.3	27.4	23.5
our Wav-BERT w/ XLSR-53	w/ mBERT	19.3	13.8	28.0	20.4

Through the Embedding Attention Module with "Sampling with Decay" mechanism, we further integrate the acoustic and linguistic information from the embedding level to facilitate better fusion between wav2vec 2.0 encoder and BERT encoder. Table 4 verifies the effectiveness of each component of our proposed Embedding Attention Module.

4.2.3 Inference

For inference, we first feed the result Y_{CTC_1} into BERT encoder; then select the one with higher confidence from the two outputs Y_{CTC_2} and Y_{CE} as our final output.

5 Experiments

In this section, we first illustrate the implementation details of our Wav-BERT. Then we introduce two low-resource speech recognition datasets containing several languages as well as the comparison results among our approach and baseline methods. Furthermore, we conduct ablation studies to validate the effectiveness of each main component of our Wav-BERT and present some case studies for perceptual comparison.

Implementation Details. For our proposed Representation Aggregation Module and Embedding Attention Module, the heads and embedding dimensions of all multi-head attention are set to 8 and 768 respectively. Meanwhile, the inner-layer dimension of the position-wise feed-forward is set to 2048. Regarding optimization details, we train our model as well as baselines based on wav2vec 2.0 Base for 200K steps with one GeForce RTX 3090 GPU, setting max tokens and update frequency to 640000 and 4 correspondingly. As for experiments using XLSR-53 (Conneau et al., 2020), three GeForce RTX 3090 GPUs are used with max tokens as 480000 and update frequency as 4. We use the three-stage learning rate policy with the

initial learning rate as $5e-5$, and set each stage ratio to 0.05, 0.45 and 0.5. Besides, we set the weight μ_1, μ_2, μ_3 and μ_4 for each loss to 0.5 for training. Other optimizer settings are the same as wav2vec 2.0 (Baevski et al., 2020). In terms of the "Sampling with Decay" policy, languages in IARPA BABEL start from 100K steps to 200K steps, while in AISHELL-1 it starts from 40k steps to 100k steps, all with p decreasing from 90% to 10%.

Datasets. IARPA BABEL (Gales et al., 2014) is an open-source multilingual corpus of conversational telephone speech. For low resource evaluation, we randomly select 3 kinds of languages with few data: Swahili (Sw), Tamil (Ta) and Vietnamese (Vi). We adopt the same setup as (Conneau et al., 2020) and use the dev folder of the BABEL dataset as our test set since "eval" data are not released. We re-sample audios of all languages to 16kHz. AISHELL-1 (Bu et al., 2017) is an open-source and high-quality Mandarin speech corpus, and is widely used in the speech community, which contains 178 hours of Mandarin speech data. Although the data is in Chinese, a common used language, the quantity is small. Thus, it can also verify our Wav-BERT for low-resource data. Moreover, there are many latest state-of-the-art methods on this dataset to be compared.

For a fair comparison, we use the official wav2vec 2.0 (Base/Large) model, XLSR-53, and mBERT models as the initial encoders. All model checkpoint download links are described in the appendix.

5.1 Results on IARPA BABEL

Table 1 reports the results on IARPA BABEL in terms of character error rate (CER), where our Wav-BERT achieves state-of-the-art performance on all low-resource languages. We find some interesting points comparing the results. First, the performance of the methods without pre-training is quite bad, which indicates that the conventional end-to-end models are impractical for low-resource languages due to the limited data. Second, the pre-training models like wav2vec 2.0 and XLSR largely improve the recognition accuracy thanks to the powerful acoustic representation learned from the huge amount of high-resource language data. Third, in addition to the pre-trained acoustic model, other methods also utilize a pre-trained language model like mBERT while the results change slightly or even become worse. One of the reasons

Table 2: Results of ASR on AISHELL-1 in terms of CER(%).

Method	Pre-trained	AISHELL-1	
		dev	test
Kaldi chain (Yu and Chen, 2021)		-	7.5
Kaldi nnet3 (Yu and Chen, 2021)		-	8.6
LAS (Shan et al., 2019)		-	10.6
ESPnet (Transformer) (Karita et al., 2019)		6.0	6.7
SA-T (Tian et al., 2019)		8.3	9.3
SAN-M (Gao et al., 2020)	-	5.7	6.5
CAT (An et al., 2019)		-	6.3
LFML (Chen et al., 2019)		6.2	6.7
LASO (Bai et al., 2021)		5.9	6.9
NAR-Transformer (Song et al., 2020)		5.6	6.3
Wenet (Zhang et al., 2020)		-	4.7
LASO with BERT (Bai et al., 2021)		5.3	6.1
NAR-BERT-ASR (Yu and Chen, 2021)	BERT	4.9	5.5
wav2vec 2.0 (Baevski et al., 2020)		7.9	8.4
wav2vec 2.0 (cn) (Baevski et al., 2020)	wav2vec 2.0	5.2	5.8
wav2vec 2.0 (cn) w/ 4-gram (Baevski et al., 2020)		4.5	4.9
BERT rescoring (Shin et al., 2019)		4.2	4.5
Adapter-BERT (Guo et al., 2020)	wav2vec 2.0	6.9	7.3
w2v-cif-bert (Yi et al., 2021)	w/ BERT	5.6	6.3
our Wav-BERT w/ wav2vec 2.0		3.8	4.0
our Wav-BERT w/ wav2vec 2.0 (cn)		3.6	3.8

Table 3: Results of different components in Representation Aggregation Module for ASR on IARPA BABEL and AISHELL-1 named CN in terms of CER(%).

Method	Vi	Sw	CN-dev	CN-test	Avg
Gated Cross-Modal Attention	19.5	14.8	3.8	4.0	10.5
w/o Gated Weighting	19.6	14.9	3.9	4.2	10.7
Gated Acoustic-Guided Attention	20.4	15.0	4.4	4.7	11.1
Gated Linguistic-Guided Attention	25.6	18.3	5.7	6.4	14.0

is that the methods that construct adapters in BERT (ADapter-BERT) or simply combine BERT with wav2vec 2.0 (w2v-cif-bert) inevitably suffer from the embedding inconsistency problem and fail to make the best use of pre-trained linguistic representation. As for our Wav-BERT, it effectively facilitates the cooperation of the pre-trained acoustic and language models by the proposed fusion modules from representation level to embedding level. As a result, it can consistently improve the ASR results for different low-resource languages. Moreover, when the pre-trained model (e.g. wav2vec 2.0) becomes larger, the performance of our Wav-BERT will be also improved while it requires more GPU resources to tune the whole model.

Table 4: Results of different components in Embedding Attention Module for ASR on IARPA BABEL and AISHELL-1 named CN in terms of CER(%).

Method	Vi	Sw	CN-dev	CN-test	Avg
Embedding Replacement	21.1	15.4	6.0	6.4	12.2
our Embedding Attention	19.5	14.8	3.8	4.0	10.5
w/o Sampling with Decay	22.0	15.7	5.7	6.2	12.4
w/o Gated Attention	20.7	15.3	4.1	4.3	11.1

5.2 Results on AISHELL-1

Table 2 reports the comparison results on AISHELL-1. In addition to the baselines mentioned above, we also report more latest works for comparison. The data quantity of this dataset is larger than that of IARPA BABEL, so all the methods perform much better. It also accounts for that the performance distance between the methods with pre-trained models and those without pre-trained models becomes small. During the methods without pre-trained models, wenet (Zhang et al., 2020) achieves the best results due to its advanced CTC-Conformer (Graves et al., 2006a; Gulati et al., 2020) architecture, better attention rescoring decoding strategy and larger training epoch number. With the pre-trained language model of BERT, NAR-BERT-ASR (Yu and Chen, 2021) stacked a decoder initialized by a pre-trained BERT model on the top of the transformer encoder and achieves competitive results on AISHELL-1. Regarding methods using the pre-trained acoustic model, the official wav2vec 2.0 Base model that pre-trained on 960 hours of Librispeech corpus achieves great results as the model learned good representations of speech. Furthermore, we also collect and use 1960 hours of public Mandarin speech data to pre-train a wav2vec 2.0 (cn) model, which obtains better performance on AISHELL-1 evaluation. In conclusion, our Wav-BERT not only improves the performance of both wav2vec 2.0 and wav2vec 2.0 (cn) models, but also outperforms other state-of-the-art methods unifying wav2vec 2.0 and BERT. It further demonstrates the generalization of Wav-BERT on different low-resource ASR datasets with different data sizes.

5.3 Comparison of model fusion methods

As illustrate in Section 4.1, there are many different model fusion methods to fuse the pre-trained wav2vec 2.0 and BERT. We compare our Wav-BERT with these methods and report the results in Table 1 and Table 2. First, by using BERT to rescore N -best hypotheses generated by wav2vec 2.0 with CTC beam search, rescoring (Shin et al., 2019) (Figure 1 (a)) is slightly better than wav2vec 2.0, but its inference process is time-consuming. Second, w2v-cif-bert (Yi et al., 2021) uses CIF to connect wav2vec 2.0 and BERT in a cascade way and replace word embedding with acoustic embedding as input for BERT. It is better than wav2vec 2.0 in AISHELL-1 but worse in BABEL for the

Table 5: Predicted examples on AISHELL-1 test set generated by Wav2vec 2.0, BERT rescoring, w2v-cif-bert and our Wav-BERT. The differences words are marked with pronunciation. The wrong words are marked in red. The translations of the sentences are also provided.

Method	Predicted example with translation
wav2vec 2.0 (Baevski et al., 2020)	温州大妈年与(yǔ)半摆(bǎi)冒充女儿成功骗婚多名小鲜肉 Wenzhou aunt Nian and banpai pretended to be their daughter and got married successfully.
BERT rescoring (Shin et al., 2019)	温州大妈年徐(yú)半百(bǎi)冒充女儿成功骗婚多名小鲜肉 More than half of Wenzhou's old aunt pretended to be her daughter and successfully cheated many young people into marriage..
w2v-cif-bert (Yi et al., 2021)	温州大妈年与(yǔ)半百(bǎi)冒充女儿成功骗婚多名小鲜肉 Wenzhou aunt year and half a hundred pretending to be daughters have successfully cheated into marriage, and there are many young people.
our Wav-BERT	温州大妈年逾(yú)半百(bǎi)冒充女儿成功骗婚多名小鲜肉 Wenzhou aunt is more than half a hundred years old, pretending to be her daughter, and has successfully cheated many young people into marriage.

reason that the mBERT is not as well trained as the bert-base-chinese model, resulting in a more severe catastrophic forgetting problem after replacing its input. Third, Adapter-BERT that inserts adapter modules into each BERT layer and tunes it on the training data, has an inconspicuous improvement or even performance degradation since the insertion of adapters affects the pre-trained representation of BERT. Finally, our Wav-BERT significantly surpasses other methods, which indicates that our model can effectively exploit the acoustic and linguistic information through the multi-level hierarchical fusion. Besides, our cooperative learning methods can also help the pre-trained encoders to avoid catastrophic forgetting of pre-training information so that the whole model can converge faster and better.

5.4 Ablation Studies

5.4.1 Representation Aggregation Module

To investigate the effectiveness of our Representation Aggregation Module, we present results for Gated Linguistic-Guided Attention, Gated Acoustic-Guided Attention, removing gated weighting in Table 3. We can find that the effect of gated weighting, while small, is still existent, which can automatically measure the importance of the acoustic and linguistic representation while aggregating those two kinds of representation. Compared with Gated Cross-Modal Attention, Gated Acoustic-Guided Attention and Gated Linguistic-Guided Attention increases the average CER by 0.6% and 3.5% respectively, which indicates that the attention in each direction plays an important role in our Representation Aggregation Module while Gated Acoustic-Guided Attention makes a greater contribution since speech recognition task is more dependent on acoustic information.

5.4.2 Embedding Attention Module

The results in Table 4 further verify the effectiveness of our Embedding Attention Module. First, we report the result of Embedding Replacement that simply replaces the original word embedding with the acoustic embedding as the input of BERT like previous works (Yu and Chen, 2021). As expected, the performance is poor especially on AISHELL-1, which indicates that such simple replacement methods will be affected by the embedding inconsistency problem. In contrast, we solve this challenge by the proposed Embedding Attention Module including the sampling mechanism and Gated Attention, so that the performance is largely improved. Second, when turning off "Sampling with Decay" or Gated Attention, the average CER increased by 1.9% and 0.6% respectively. It demonstrates that the "Sampling with Decay" mechanism effectively alleviates the embedding inconsistency of BERT between inference and training. Mover, the Gated Attention effectively provides additional acoustic information to the input of BERT, facilitating it to capture more reliable linguistic representation.

5.5 Case Studies

We further present some case studies in Table 5, to illustrate the importance of acoustic and linguistic information for speech recognition. We provided some transcript examples obtained from the baseline methods and our Wav-BERT with the same input from AISHELL-1 test set. The pronunciations of the keywords and the English translation of the whole sentence are also provided. As can be observed, all the baseline methods predict one or two wrong words with similar pronunciation as the wrong words, which leads to an unreasonable sentence. On the contrary, thanks to the cooperative learning of acoustic and linguistic information, our Wav-BERT can successfully recognize the whole sentence without any word error.

6 Conclusion

In this work, based on the powerful wav2vec 2.0 and BERT models, we introduce cooperative acoustic and linguistic representation learning for low-resource speech recognition. To solve the representation discrepancy and embedding inconsistency challenges, we design a Representation Aggregation Module and an Embedding Attention Module to facilitate the cooperation of the two pre-trained models and thus boost the representation learning. Extensive experimental results demonstrate that our proposed Wav-BERT can significantly improve low-resource ASR performances in different languages. In future work, we will investigate more effective modules to infuse more types of knowledge, and apply our framework to more pre-trained models to promote the development of low-resource speech tasks.

Acknowledgement

This work was supported in part by National Key R&D Program of China under Grant No. 2020AAA0109700, National Natural Science Foundation of China (NSFC) under Grant No.U19A2073 and No.61976233, Guangdong Province Basic and Applied Basic Research (Regional Joint Fund-Key) Grant No.2019B1515120039, Guangdong Outstanding Youth Fund (Grant No. 2021B1515020061), Shenzhen Fundamental Research Program (Project No. RCYX20200714114642083, No. JCYJ20190807154211365).

References

Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, et al. 2016. Deep speech 2: End-to-end speech recognition in english and mandarin. In *International conference on machine learning*, pages 173–182. PMLR.

Keyu An, Hongyu Xiang, and Zhijian Ou. 2019. Cat: crf-based asr toolkit. *arXiv preprint arXiv:1911.08747*.

Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in Neural Information Processing Systems*, 33.

Ye Bai, Jiangyan Yi, Jianhua Tao, Zhengkun Tian, Zhengqi Wen, and Shuai Zhang. 2021. Fast end-to-

end speech recognition via non-autoregressive models and cross-modal knowledge transferring from bert. *arXiv preprint arXiv:2102.07594*.

- Hui Bu, Jiayu Du, Xingyu Na, Bengu Wu, and Hao Zheng. 2017. Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline. In *2017 20th Conference of the Oriental Chapter of the International Coordinating Committee on Speech Databases and Speech I/O Systems and Assessment (O-COCOSDA)*, pages 1–5. IEEE.
- William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals. 2016. Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4960–4964. IEEE.
- Nanxin Chen, Shinji Watanabe, Jesús Villalba, and Nanjim Dehak. 2019. Listen and fill in the missing letters: Non-autoregressive transformer for speech recognition. *arXiv preprint arXiv:1911.04908*.
- Shih-Hsuan Chiu and Berlin Chen. 2021. Innovative bert-based reranking language models for speech recognition. In *2021 IEEE Spoken Language Technology Workshop (SLT)*, pages 266–271. IEEE.
- Jaejin Cho, Murali Karthick Baskar, Ruizhi Li, Matthew Wiesner, Sri Harish Mallidi, Nelson Yalta, Martin Karafiát, Shinji Watanabe, and Takaaki Hori. 2018. Multilingual sequence-to-sequence speech recognition: Architecture, transfer learning, and language modeling. *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 521–527.
- Yu-An Chung and James Glass. 2020. Generative pre-training for speech with autoregressive predictive coding. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3497–3501. IEEE.
- Alexis Conneau, Alexei Baevski, Ronan Collobert, Abdelrahman Mohamed, and Michael Auli. 2020. Unsupervised cross-lingual representation learning for speech recognition. *arXiv preprint arXiv:2006.13979*.
- Siddharth Dalmia, Ramon Sanabria, Florian Metze, and Alan W Black. 2018. Sequence-based multilingual low resource speech recognition. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4909–4913. IEEE.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*.

- Hayato Futami, Hirofumi Inaguma, Sei Ueno, Masato Mimura, Shinsuke Sakai, and Tatsuya Kawahara. 2020. Distilling the knowledge of bert for sequence-to-sequence asr. *arXiv preprint arXiv:2008.03822*.
- Mark JF Gales, Kate M Knill, Anton Ragni, and Shakti P Rath. 2014. Speech recognition and keyword spotting for low-resource languages: Babel project research at cued. In *Fourth International Workshop on Spoken Language Technologies for Under-Resourced Languages (SLTU-2014)*, pages 16–23. International Speech Communication Association (ISCA).
- Zhifu Gao, Shiliang Zhang, Ming Lei, and Ian McLoughlin. 2020. San-m: Memory equipped self-attention for end-to-end speech recognition. *arXiv preprint arXiv:2006.01713*.
- Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. 2019. Mask-predict: Parallel decoding of conditional masked language models. *arXiv preprint arXiv:1904.09324*.
- Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. 2013. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*.
- Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. 2006a. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, pages 369–376.
- Alex Graves, Santiago Fernández, Faustino J. Gomez, and Jürgen Schmidhuber. 2006b. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *ICML '06*.
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, et al. 2020. Conformer: Convolution-augmented transformer for speech recognition. *arXiv preprint arXiv:2005.08100*.
- Junliang Guo, Zhirui Zhang, Linli Xu, Hao-Ran Wei, Boxing Chen, and Enhong Chen. 2020. Incorporating bert into parallel sequence decoding with adapters. In *Advances in Neural Information Processing Systems*, volume 33, pages 10843–10854. Curran Associates, Inc.
- Tingwei Guo, Cheng Wen, Dongwei Jiang, Ne Luo, Ruixiong Zhang, Shuaijiang Zhao, Wubo Li, Cheng Gong, Wei Zou, Kun Han, et al. 2021. Didispeech: A large scale mandarin speech corpus. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6968–6972. IEEE.
- Kenneth Heafield. 2011. Kenlm: Faster and smaller language model queries. In *Proceedings of the sixth workshop on statistical machine translation*, pages 187–197.
- Jui-Yang Hsu, Yuan-Jui Chen, and Hung-yi Lee. 2020. Meta learning for end-to-end low-resource speech recognition. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7844–7848. IEEE.
- Jacob Kahn, Ann Lee, and Awni Hannun. 2020. Self-training for end-to-end speech recognition. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7084–7088. IEEE.
- Shigeki Karita, Nanxin Chen, Tomoki Hayashi, Takaaki Hori, Hirofumi Inaguma, Ziyang Jiang, Masao Someki, Nelson Enrique Yalta Soplín, Ryuichi Yamamoto, Xiaofei Wang, et al. 2019. A comparative study on transformer vs rnn in speech applications. In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 449–456. IEEE.
- Julius Kunze, Louis Kirsch, Ilia Kurenkov, Andreas Krug, Jens Johannsmeier, and Sebastian Stober. 2017. Transfer learning for speech recognition on a budget. In *Rep4NLPACL*.
- Bo Li, Tara N Sainath, Ruoming Pang, and Zelin Wu. 2019. Semi-supervised training for end-to-end models via weak distillation. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2837–2841. IEEE.
- Alex Nichol, Joshua Achiam, and John Schulman. 2018. On first-order meta-learning algorithms. *ArXiv*, abs/1803.02999.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of NAACL-HLT 2019: Demonstrations*.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli. 2019a. wav2vec: Unsupervised pre-training for speech recognition. In *INTER-SPEECH*.
- Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli. 2019b. wav2vec: Unsupervised pre-training for speech recognition. In *INTER-SPEECH 2019*.

Changhao Shan, Chao Weng, Guangsen Wang, Dan Su, Min Luo, Dong Yu, and Lei Xie. 2019. Component fusion: Learning replaceable language model component for end-to-end speech recognition system. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5361–5635. IEEE.

Yao Shi, Hui Bu, Xin Xu, Shaoji Zhang, and Ming Li. 2020. Aishell-3: A multi-speaker mandarin tts corpus and the baselines. *arXiv preprint arXiv:2010.11567*.

Joonbo Shin, Yoonhyung Lee, and Kyomin Jung. 2019. Effective sentence scoring method using bert for speech recognition. In *Asian Conference on Machine Learning*, pages 1081–1093. PMLR.

Xingchen Song, Zhiyong Wu, Yiheng Huang, Chao Weng, Dan Su, and Helen Meng. 2020. Non-autoregressive transformer asr with ctc-enhanced decoder input. *arXiv preprint arXiv:2010.15025*.

Zhengkun Tian, Jiangyan Yi, Jianhua Tao, Ye Bai, and Zhengqi Wen. 2019. Self-attention transducers for end-to-end speech recognition. *arXiv preprint arXiv:1909.13037*.

Shubham Toshniwal, Tara N Sainath, Ron J Weiss, Bo Li, Pedro Moreno, Eugene Weinstein, and Kanishka Rao. 2018. Multilingual speech recognition with a single end-to-end model. In *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 4904–4908. IEEE.

Alex Wang and Kyunghyun Cho. 2019. Bert has a mouth, and it must speak: Bert as a markov random field language model. *arXiv preprint arXiv:1902.04094*.

Shinji Watanabe, Takaaki Hori, and John R. Hershey. 2017a. Language independent end-to-end architecture for joint language identification and speech recognition. In *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 265–271.

Shinji Watanabe, Takaaki Hori, Suyoun Kim, John R Hershey, and Tomoki Hayashi. 2017b. Hybrid ctc/attention architecture for end-to-end speech recognition. *IEEE Journal of Selected Topics in Signal Processing*, 11(8):1240–1253.

Genta Indra Winata, Guangsen Wang, Caiming Xiong, and Steven Hoi. 2020. Adapt-and-adjust: Overcoming the long-tail problem of multilingual speech recognition. *arXiv preprint arXiv:2012.01687*.

Yubei Xiao, Ke Gong, Pan Zhou, Guolin Zheng, Xi-aodan Liang, and Liang Lin. 2021. Adversarial meta sampling for multilingual low-resource speech recognition. In *AAAI*.

Cheng Yi, Shiyu Zhou, and Bo Xu. 2021. Efficiently fusing pretrained acoustic and linguistic encoders for low-resource speech recognition. *IEEE Signal Processing Letters*.

Fu-Hao Yu and Kuan-Yu Chen. 2021. Non-autoregressive transformer-based end-to-end asr using bert. *arXiv preprint arXiv:2104.04805*.

Binbin Zhang, Di Wu, Zhuoyuan Yao, Xiong Wang, Fan Yu, Chao Yang, Liyong Guo, Yaguang Hu, Lei Xie, and Xin Lei. 2020. Unified streaming and non-streaming two-pass end-to-end model for speech recognition. *arXiv preprint arXiv:2012.05481*.

A Datasets

Both IARPA BABEL dataset (Gales et al., 2014) and AISHELL-1 (Bu et al., 2017) are open-source and high-quality speech datasets, and are widely used in the speech community. Among them, AISHELL-1 can be downloaded for free here¹, For each speaker in it, around 360 utterances (about 26 minutes of speech) are released. Table 6 provides a summary of all subsets in the corpus. As for IARPA BABEL, it can be purchased through LDC² (eg. Vietnamese Language Pack³). Table 7 summarizes the amount of data in hours for the language used in our experiments on the "Full Language Pack" (FLP) condition. Researchers can easily reproduce or compare our results with the same languages.

Table 6: AISHELL-1 dataset statistics.

Subset	Duration(hrs)	Male	Female
Training	150	161	179
Development	10	12	28
Test	5	13	7

Table 7: IARPA BABEL dataset statistics.

Language	Train(hrs)	Eval(hrs)
Vietnamese	87.72	11.00
Swahili	44.39	10.65
Tamil	69.35	11.68

B Ours Wav-BERT Model

Our model checkpoint described in Sec 5 can be downloaded [here](#). With limited storage space, thus we only upload the model using wav2vec 2.0 Base.

¹<https://www.openslr.org/33/>

²<https://www ldc.upenn.edu/>

³<https://catalog ldc.upenn.edu/LDC2017S01>

C Pre-trained Models

We use different pre-trained acoustic and language models in our experiment described in Sec 5. All of them are open-source except the wav2vec 2.0 (Baeovski et al., 2020) pre-trained in Chinese by ourselves. For pre-trained language models, the bert-base-chinese model can be download here ⁴, and the multilingual mBERT can be download here ⁵. For pre-trained acoustic models, the official wav2vec 2.0 pre-trained on English can be download here ⁶, and the XLSR-53 (Conneau et al., 2020) model can be downloaded here ⁷. Besides, though the wav2vec 2.0(cn) pre-trained on 1,960 hours of Chinese data cannot open-source, both the used training code and datasets are open-source, which means researchers still can reproduce our results. In details, we base on the Fairseq framework ⁸ (Ott et al., 2019) to pre-train our model 8 GeForce RTX 3090 GPUs with max tokens and update frequency setting to 1400000 and 8 respectively, consuming about one week to train 400K steps. Besides, the used datasets are DiDiSpeech (Guo et al., 2021), PVTC ⁹, ST-CMDS ¹⁰, aidatantang ¹¹, AISHELL-1, AISHELL-3 (Shi et al., 2020), MAGICDATA ¹², MagicDataSpeech ¹³, Primewords ¹⁴ and Thchs ¹⁵.

D Baselines

We describe some baseline methods below, which are reproduced by ourselves or experimented with the open-source code.

1. Wav2vec 2.0 w/ 4-gram: For each language, results from the trained wav2vec 2.0 model with beam search, are rescored by the 4-gram language model. Specifically, the 4-

⁴<https://s3.amazonaws.com/models.huggingface.co/bert/bert-base-chinese.tar.gz>

⁵<https://s3.amazonaws.com/models.huggingface.co/bert/bert-base-multilingual-uncased.tar.gz>

⁶https://dl.fbaipublicfiles.com/fairseq/wav2vec/wav2vec_small.pt

⁷https://dl.fbaipublicfiles.com/fairseq/wav2vec/xlsr_53_56k.pt

⁸<https://github.com/pytorch/fairseq>

⁹<https://www.pvtc2020.org/index.html>

¹⁰<http://www.openslr.org/38/>

¹¹<http://www.openslr.org/62/>

¹²<http://www.openslr.org/68/>

¹³<https://www.biendata.xyz/competition/magicdata/>

¹⁴<http://www.openslr.org/47/>

¹⁵<http://www.openslr.org/18/>

gram model is trained by transcripts in the training set of each language, using the KenLM (Heafield, 2011) framework. And the beam size for beam search is set to 50.

2. BERT rescoring (Chiu and Chen, 2021; Shin et al., 2019): For each language, results from the trained wav2vec 2.0 model with beam search, are rescored by the fine-tuned language model(mBERT or bert-base-chinese model). Specifically, the linguistic decoder is fine-tuned by transcripts in the training set of each language using masked language model(MLM) objective (Devlin et al., 2018) of BERT. In rescoring stage, we mask each word in the sentence once at a time, then sum all the log-likelihoods of the masked words from each masked input instance. Finally rescoring the sentence with both the likelihoods from acoustic and language model. Besides, considering it is time-consuming, the beam size for beam search is set to 5.
3. Adapter-BERT: This method is inspired by AB-Net (Guo et al., 2020), cross-attention adapters are inserted to each BERT layer to unify the wav2vec 2.0 and BERT model. Output from the feed-forward layer at the last of BERT is supervised by the cross-entropy criterion. In inference, the Mask-Predict algorithm (Ghazvininejad et al., 2019) is adopted.
4. Embedding Replacement: Inspired by previous work (Yu and Chen, 2021), we use similar architecture as it but replace the acoustic encoder with wav2vec 2.0 and keep our Representation Aggregation Module. We use position embeddings as query vector and acoustic representation from wav2vec 2.0 as key and value vector to attention block followed by 3 self-attention block, which is the same as (Yu and Chen, 2021), generating aligned acoustic representation H_{pos} . Then H_{pos} is used as the input of BERT, replacing the word embedding. Finally, Representation Aggregation Module takes both the H_{pos} and linguistic representation from BERT as input, just the same as our Wav-BERT. It is worth mention that the length of the position embedding is set to 60, considering it cost too much GPU memory for a larger value.

E More Implementation Details

Most of the significant experiment details are described in Sec 5. Aiming to let researcher reproduce our result more easily, we describe more details below. About the data augmentation, mask probability and mask channel probability are set to 0.65 and 0.5 respectively the same as setting in wav2vec 2.0 (Baevski et al., 2020) for 100 hour training data. Besides, we use adam optimizer, setting adam betas and adam eps to (0.9,0.98) and 1e-08 individually. In data preprocessing, we use feature normalize for wav2vec 2.0 Base model but not for the XLSR-53 model, keeping consistent with the pre-training setting. Also, we filter some samples whose length of speech shorter than 0.5 seconds as well as number of subwords less than 1 or bigger than 512 in training set. Regarding the training time, training our Wav-BERT model with wav2vec 2.0 Base model spends less than 2 days, and 5 days with the XLSR-53 model. Finally, the number of parameters in our model with wav2vec 2.0 Base is about 380M, and 600M with XLSR-53, which is slightly different with different languages.