



# TranUSR: Phoneme-to-word Transcoder Based Unified Speech Representation Learning for Cross-lingual Speech Recognition

Hongfei Xue<sup>1</sup>, Qijie Shao<sup>1</sup>, Peikun Chen<sup>1</sup>, Pengcheng Guo<sup>1</sup>, Lei Xie<sup>1\*</sup>, Jie Liu<sup>2</sup>

<sup>1</sup>Audio, Speech and Language Processing Group (ASLP@NPU),  
School of Computer Science, Northwestern Polytechnical University, Xi'an, China  
<sup>2</sup>Huawei Cloud

hfxue@mail.nwpu.edu.cn, lxie@nwpu.edu.cn

## Abstract

UniSpeech has achieved superior performance in cross-lingual automatic speech recognition (ASR) by explicitly aligning latent representations to phoneme units using multi-task self-supervised learning. While the learned representations transfer well from high-resource to low-resource languages, predicting words directly from these phonetic representations in downstream ASR is challenging. In this paper, we propose TranUSR, a two-stage model comprising a pre-trained UniData2vec and a phoneme-to-word Transcoder. Different from UniSpeech, UniData2vec replaces the quantized discrete representations with continuous and contextual representations from a teacher model for phonetically-aware pre-training. Then, Transcoder learns to translate phonemes to words with the aid of extra texts, enabling direct word generation. Experiments on Common Voice show that UniData2vec reduces PER by 5.3% compared to UniSpeech, while Transcoder yields a 14.4% WER reduction compared to grapheme fine-tuning.

**Index Terms:** Cross-lingual speech recognition, unified representation learning, phoneme-to-word

## 1. Introduction

End-to-end ASR systems typically require a large amount of labeled data for training [1], which is difficult to obtain in many languages. This problem makes training of ASR in most low-resource languages a challenging task [2, 3]. Cross-lingual learning is an approach widely used to improve ASR performance in low-resource languages [4, 5, 6], which could learn common pronunciation information from non-target high-resource languages to assist ASR for target low-resource languages. Studies [7, 8, 9, 10] have shown that pre-training on multilingual data, including labeled and unlabeled data from available language, and subsequent fine-tuning on the target low-resource language with labeled data, is an effective cross-lingual learning method.

Most recent research on pre-training for cross-lingual ASR can be classified into two categories. The first category only includes unsupervised loss in the pre-training stage, such as [7, 8, 9, 11, 12, 13]. These methods use much inexpensive, unlabeled multilingual data to learn cross-lingual features. In XLSR [9], a notable representation in the field is presented, which involves pre-training a model using self-supervised loss on 53 languages. The second category involves incorporating a supervised loss during unsupervised pre-training [10, 14, 15]. During the pre-training stage, cross-lingual representations is enhanced using labeled data from high-resource languages. The UniSpeech [10] method combines a supervised CTC loss [16] and

a self-supervised contrastive loss [17] to enhance the quality of the learned representations. Similarly, the JUST method [15] integrates a supervised RNN-T loss with two unsupervised losses. For ASR tasks, studies such as UniSpeech [10] and JUST [15] have demonstrated that the above-mentioned second category typically achieves higher performance compared to the first category such as XLSR [9].

The choice of modeling units in supervised loss during the pre-training stage is typically between multilingual grapheme sets and phoneme sets. Some methods, like JUST [15], use grapheme units, while others, like UniSpeech [10], opt phoneme units. To the best of our knowledge, the UniSpeech method is currently state-of-the-art (SOTA) on the cross-lingual Common Voice dataset [18]. During the pre-training stage, the methods using grapheme units face challenges in learning shared cross-lingual representations due to a lack of shared graphemes among different languages. Conversely, using phoneme units could make models more easily learn shared phonetic representations. International Phonetic Alphabet (IPA) [19] is a typical cross-lingual phoneme set that is suited for this situation. During the fine-tuning stage, a phoneme-based model cannot directly generate words, and requires additional fine-tuning with grapheme units. However, it is also challenging to map phonetic representations to grapheme units accurately. In addition, the quantizer in UniSpeech is required to generate phoneme representations, which is difficult because the quantizer lacks sufficient trainable parameters.

This study proposes a **Transcoder based Unified Speech Representation (TranUSR)** learning for cross-lingual ASR. Our TranUSR is a two-stage approach that includes a pre-trained UniData2vec and a phoneme-to-word (P2W) Transcoder. Firstly, we introduce UniData2vec in UniSpeech [10], which is a Data2vec [20] model. UniData2vec replaces UniSpeech's quantized discrete representations with continuous and contextual representations from a teacher model, enhancing the self-supervised targets to accurately capture phonetic information. Moreover, the teacher model has the same number of parameters as the student model, resulting in complete targets that are closely related to phonemes. This facilitates the learning of a unified speech representation in supervised learning and self-supervised learning. Secondly, we introduce Transcoder, which is based on the phoneme-to-grapheme model [21, 22, 23]. The Transcoder includes various components, such as a probabilistic vector input component. With these components, the Transcoder can convert encoder's hypotheses to target language's words with the help of information at different granulates. More importantly, the Transcoder can be trained with large-scale, low-cost text data, making it an attractive solution for low-resource settings. Experiments on the Common Voice dataset show that UniData2vec outperforms the SOTA method

\*Corresponding author.

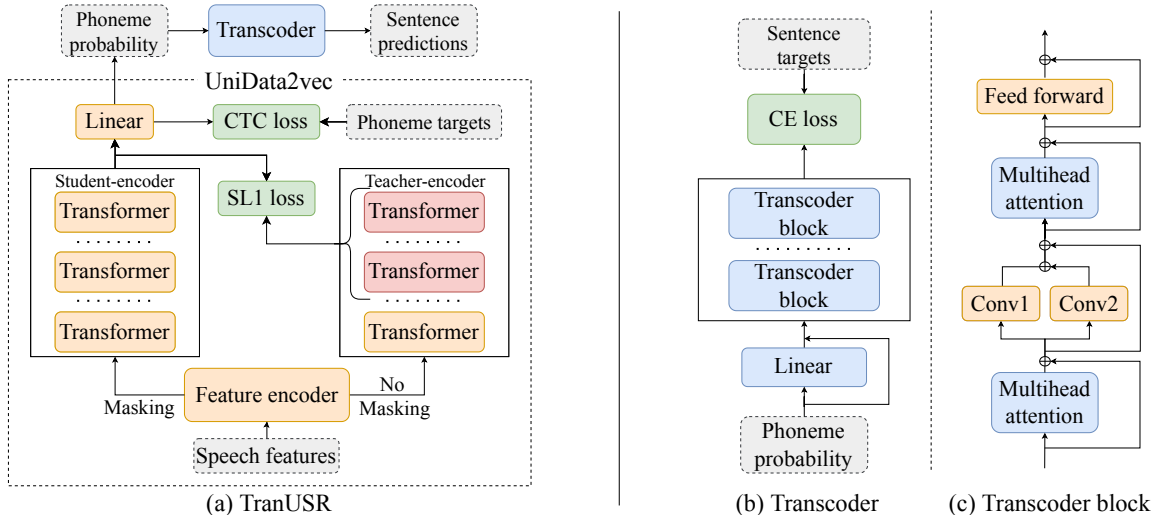


Figure 1: Architecture schematic diagram of TranUSR. Subfigure (a) is the overview of TranUSR, (b) is the Transcoder P2W module, and (c) is the details of the Transcoder blocks.

with an average improvement of relative 5.3% phoneme error rate (PER). Furthermore, Transcoder achieves an average improvement of relative 14.4% word error rate (WER) compared to grapheme fine-tuning.

## 2. Method

As illustrated in Figure 1 (a), the TranUSR framework comprises two modules: UniData2vec and Transcoder. UniData2vec predicts phoneme probabilities from speech features, and Transcoder for each language converts these probabilities into word-level sentences.

### 2.1. Problem statement

In our research, we use four types of high and low-resource data, which can be denoted as  $\mathbb{L}$ ,  $\mathbb{U}$ ,  $\mathbb{F}$  and  $\mathbb{T}$ .  $\mathbb{L} = \{(X_i^L, Y_i^L), i \in (1, N_L)\}$  is labeled non-target high-resource language data consisting of speech audio and transcript labels.  $\mathbb{U} = \{(X_j^U), j \in (1, N_U)\}$  is unlabeled target low-resource language data,  $\mathbb{F} = \{(X_k^F, Y_k^F), k \in (1, N_F)\}$  is labeled target language data, and  $\mathbb{T} = \{(Y_l^T), l \in (1, N_T)\}$  is pure text target language data. While  $\mathbb{L}$ ,  $\mathbb{U}$ , and  $\mathbb{T}$  are easy to obtain,  $\mathbb{F}$  is expensive.

Our task involves four steps: (1) training a UniData2vec base model on  $\mathbb{L}$ , (2) initializing with UniData2vec and training a UniData2vec<sup>+</sup> model on  $\mathbb{U}$  to learn cross-lingual phonetically-aware representations, (3) fine-tuning the model on each target low-resource language dataset  $\mathbb{F}$  to improve phoneme recognition ability, and (4) training a P2W Transcoder model for each language using the large text data  $\mathbb{T}$ .

### 2.2. UniData2vec

**Model structure** As shown in Figure 1 (a), our UniData2vec module is composed of a convolutional feature encoder [17], two context Transformer encoders [24], and a linear layer. The convolutional feature encoder is used to map the input speech  $X$  to latent representations  $Z$ . The latent representations  $Z$  are then fed into two Transformer encoders: a masked student encoder and an unmasked teacher encoder, which is similar to [20]. The student encoder outputs context representations  $C$ , while the teacher encoder predicts targets  $Y$ . Finally, the

context representations  $C$  are mapped to phoneme units via the linear layer. In our method, the teacher and student encoders have the same amount of parameters, thus their phonetic representations have the same quality, which is more accurate than the quantizer outputs in UniSpeech [10].

**Multi-task learning** The multi-task learning of unified representation mainly consists of a supervised CTC loss [16] and a self-supervised Smooth L1 (SL1) loss [20]. During this stage, our training object includes three parts: 1) Supervised phoneme recognition on the non-target language, which is achieved by using the CTC loss on  $\mathbb{L}$  to align the context representations  $C$  with phoneme units. 2) Self-supervised representation prediction on the non-target language, which is achieved by using the SL1 loss on  $\mathbb{L}$  to encourage the student encoder to predict the masked samples' targets  $Y$ . 3) Self-supervised representation prediction on the target languages, which is achieved by using the SL1 loss on  $\mathbb{U}$ .

**Self-supervised loss** The self-supervised SL1 loss encourages the student encoder to predict the teacher's contextual representations  $Y$  of masked samples. Specifically, the teacher encoder uses the exponential moving average (EMA) [20] method to parameterize the model parameters. The training targets  $Y$  are constructed based on the output of the top  $K$  layers of the teacher encoder, which represent the time steps masked in the student encoder. Given the context representations  $C$  and the targets  $Y$ , we use the SL1 loss to regress these targets:

$$L_s(Y, C) = \begin{cases} \frac{1}{2}(Y - C)^2 / \beta & |Y - C| \leq \beta \\ (|Y - C| - \frac{1}{2}\beta) & \text{otherwise} \end{cases}, \quad (1)$$

where  $\beta$  controls the transition from square loss to L1 loss.

**Unified representation** Supervised CTC and self-supervised SL1 losses are combined to train a unified representation. During supervised training, the primary loss is CTC, and the context representations  $C$  generated by the student encoder correspond to phoneme units. Since the EMA method parameterizes the teacher encoder, the self-supervised targets  $Y$  generated by the teacher encoder are also related to phoneme units. Therefore, the representations learned by supervised and self-supervised training can be constrained in the same space without the replacement in UniSpeech [10]. The losses can be

expressed separately as:

$$L_{\text{UniData2vec}} = \sum_{(x,y) \in \mathbb{L}} (L_{\text{ctc}} + \alpha L_s), \quad (2)$$

$$L_{\text{UniData2vec}^+} = L_{\text{UniData2vec}} + \sum_{x \in \mathbb{U}} L_s, \quad (3)$$

where  $\alpha$  is a tunable weight of self-supervised loss. In the training of dataset  $\mathbb{L}$ , it could be set to  $(0, 0.5)$ .

**Fine-tuning** During the fine-tuning stage, we use the  $\mathbb{F}$  data to train the CTC loss for phoneme recognition in each target language. This helps the model achieve better phoneme recognition ability on the specific language.

### 2.3. Transcoder

**Model structure** As shown in Figure 1 (b), our Transcoder module consists of a linear layer and multiple Transcoder blocks. Each Transcoder block contains two convolutional layers, two multi-head attention blocks [24], a feed-forward neural network, and multiple residual connections [25], as depicted in Figure 1 (c).

**Probabilistic vector input** P2W models typically use phoneme one-hot vectors as input, which can result in performance degradation when phoneme sequences contain errors. Our Transcoder is trained on correct texts and fine-tuned on UniData2vec’s hypotheses that may contain errors. To accommodate both types of input, we use soft phoneme posterior probabilities instead of hard one-hot vectors as input. When the phoneme sequence is from the correct text, its probability vector and one-hot vector are equivalent. However, the conventional text embedding layer only supports hard one-hot vectors as input. To overcome this limitation, we use a trainable linear layer to learn the text embedding, which enables correct text pre-training and hypothesis fine-tuning.

**Intra-word and inter-word information extraction** Conventional P2W models typically use convolutional layers to extract intra-word phoneme dependencies. However, a single convolutional layer may not be sufficient to handle words of different lengths since they contain varying numbers of phonemes. To address this issue, our Transcoder module uses two convolutional models with different kernel sizes to extract intra-word information at two granularities, thereby improving accuracy for words of varying lengths. Additionally, we use a multi-head attention block to extract inter-word information.

**Alignment constraint loss** The sequence lengths of phonemes, words can differ between languages, and CTC loss is often used to handle inconsistent output and input sequence lengths. However, P2W mapping has clear alignment boundaries, unlike audio-to-word mapping. To further improve the performance of Transcoder, we use cross-entropy (CE) loss to ensure accurate alignment learning. Since CE loss requires inputs and outputs to have the same length, we align the sequences by prefixing words with ‘\*’ tags. As an example, for the Italian sentence “vostra casa” we align it as shown in Figure 2. With this alignment, our Transcoder can learn to accurately map phonemes to words.

**Phoneme:** v ɔ s t r a k a z a    ①: vostra  
**Word:**    \* \* \* \* \* ① \* \* \* ②    ②: casa

Figure 2: Sequence alignment using ‘\*’ tags as prefixes.

## 3. Experimental setup

**Dataset** This study uses two datasets: Common Voice 6.0<sup>1</sup> [18] and TED2020<sup>2</sup> [26]. To compare with UniSpeech, we use the following eight languages for experiments: English (*en*), Spanish (*es*), French (*fr*), Italian (*it*), Kyrgyz (*ky*), Dutch (*nl*), Russian (*ru*) and Tatar (*tt*). We use *en* as a non-target language  $\mathbb{L}$  with 1350 hours of data. Others are the target languages, and the duration of  $\mathbb{U}$  is in Table 1. For fine-tuning, we use the same settings of UniSpeech [10]. TED2020 is a multilingual text corpus, and we select the following three languages as  $\mathbb{T}$  for evaluation with 400k sentences each: *es*, *fr* and *it*.

**UniData2vec structure** We implement the model within the fairseq [27] framework and refer to the settings in Data2vec2 [20]. The encoder consists of 7 convolution layers, each with 512 channels, the strides are (5, 2, 2, 2, 2, 2, 2), the kernel widths are (10, 3, 3, 3, 3, 2, 2). Both the teacher and student encoders have 12 Transformer layers, whose dimensions are 768, inner dimensions are 3072, and have 8 attention heads. The targets are determined by taking the average of the top  $K = 8$  layers. We generate IPA phoneme transcriptions by an open-source phonemizer tool<sup>3</sup>.

**UniData2vec details** During the pre-training stage, we set the weight  $\alpha$  of the self-supervised loss to 0.15 and the weight  $\beta$  in SL1 to 0.25. We first train a base UniData2vec model on  $\mathbb{L}$  and then transfer it to  $\mathbb{U}$  to form UniData2vec<sup>+</sup>. It is worth noting that both our UniData2vec and UniData2vec<sup>+</sup> models have the same structure. In the fine-tuning stage, we follow the settings in [20]. We use Adam optimization [28] with a peak learning rate of  $5 \times 10^{-4}$  for UniData2vec and  $5 \times 10^{-5}$  for UniData2vec<sup>+</sup>. The UniData2vec model is updated 400K times, while UniData2vec<sup>+</sup> is updated 200K times.

**Transcoder structure** We use three Transcoder blocks with a model dimension of 256 and an inner dimension of 1024, four attention heads, and two convolutional kernels of sizes 3 and 5. For each language, we select the top 50,000 most frequent words as the output vocabulary.

**Transcoder details** Transcoder is first trained on  $\mathbb{T}$  and then optionally fine-tuned using UniData2vec’s hypothesis. We use Adam optimization with a peak learning rate of 5e-4 for 200k updates and a batch size of 256 sentences.

## 4. Results

### 4.1. UniData2vec

**Cross-lingual phoneme recognition** As shown in Table 1, Our re-implementation of the UniSpeech and UniSpeech<sup>+</sup> models shows some improvements and declines compared to the published UniSpeech [10] results, which could be attributed to the updates of the Common Voice dataset and the phonemizer tool. After being fine-tuned on low-resource language  $\mathbb{F}$  for 1 hour, the UniData2vec shows similar competitiveness compared to the re-implemented UniSpeech. But when the fine-tuning dataset is added to 10 hours, the average PER of UniData2vec10 is reduced to relative 5.8% compared with UniSpeech10. We think this is because UniData2vec gives more importance to supervised CTC loss as the primary loss, resulting in better phoneme recognition abilities, which require more data to be transferred to other languages. When further training with the

<sup>1</sup><https://commonvoice.mozilla.org/en/datasets>. Dec 2020 release version is used for training our models. We use the same data size for each language compared to the UniSpeech.

<sup>2</sup><https://opus.nlpl.eu/TED2020.php>

<sup>3</sup><https://github.com/bootphon/phonemizer>

Table 1: *PER (%) results on Common Voice for different methods. The  $\mathbb{L}$  and  $\mathbb{U}$  datasets are described in detail in Section 3.  $\mathbb{U}_{mo}$ : Monolingual unlabelled data. Re-implement: we re-ran the baseline experiment from UniSpeech.*

Model	pre-train	fine-tune	<i>es</i>	<i>fr</i>	<i>it</i>	<i>ky</i>	<i>nl</i>	<i>ru</i>	<i>tt</i>	avg
Unlabeled data ( $\mathbb{U}$ )			168 h	353 h	90 h	17 h	29 h	55 h	17 h	
UniSpeech [10]	$\mathbb{L}$	1 h	10.9	14.8	15.2	11.4	16.2	16.1	9.6	13.5
UniSpeech <sup>+</sup> [10]	$\mathbb{L}, \mathbb{U}_{mo}$	1 h	5.7	7.9	8.1	<b>6.8</b>	9.3	<b>8.6</b>	<b>6.0</b>	7.5
UniSpeech (re-implement)	$\mathbb{L}$	1 h	10.4	12.7	13.4	10.2	14.8	15.7	8.8	12.3
UniData2vec	$\mathbb{L}$	1 h	10.0	13.0	13.3	10.3	15.0	15.1	9.0	12.2
UniSpeech10	$\mathbb{L}$	10 h	7.2	8.1	10.9	4.3	7.3	7.1	3.3	6.9
UniData2vec10	$\mathbb{L}$	10 h	7.0	8.2	10.4	4.1	6.2	6.6	3.2	6.5
UniSpeech <sup>+</sup> (re-implement)	$\mathbb{L}, \mathbb{U}_{mo}$	1 h	6.1	7.7	8.5	7.8	10.8	9.8	7.4	8.3
UniData2vec <sup>+</sup>	$\mathbb{L}, \mathbb{U}_{mo}$	1 h	<b>5.5</b>	<b>6.5</b>	<b>7.1</b>	<b>6.8</b>	<b>8.0</b>	8.8	7.2	<b>7.1</b>

respective monolingual unlabelled data  $\mathbb{U}_{mo}$ , UniData2vec<sup>+</sup> is on average 5.3% lower than UniSpeech<sup>+</sup> in terms of PER. This improvement is more prominent in languages with abundant unlabeled data such as *es*, *fr*, and *it* but weaker in languages with little data such as *ky*, *ru*, and *tt*. Because the weight of self-supervised loss is small during UniData2vec, and the ability to predict masked positions is not enough, thus UniData2vec<sup>+</sup>'s advantages require more data to reflect.

**Unified representation learning** To demonstrate the superiority of UniData2vec's unified representation learning over UniSpeech, we evaluate the quality of the self-supervised phoneme information learned by UniData2vec. Following [29], we use *K*-means to cluster the self-supervised targets' phoneme information and measure purity and mutual information metrics. We extract representations from the top 8 layers in the teacher encoder for UniData2vec and use discrete latent representations for UniSpeech. Figure 3 shows that UniData2vec exhibits a clearer purity than UniSpeech, with the difference becoming more evident in deeper layers. This is because the final layer performs phoneme classification. While UniData2vec generally outperforms UniSpeech in terms of mutual information, there is a performance drop in layer 9, possibly due to the intermediate layers learning less content information.




Figure 3: *Quality of cluster assignments measured by Purity and Mutual Information (MI) obtained through k-means clustering on self-supervised targets.*

## 4.2. Transcoder

**P2W transcoder vs. direct fine-tuning** For a model like UniSpeech or UniData2vec that learns phoneme recognition in the pre-training stage, we can transfer it to predict words of target languages by directly replacing the output units and then fine-tuning. In order to demonstrate the advantages of the Transcoder, we compare it with direct fine-tuning. As shown in Table 2, the Transcoder trained with extra text data results in significantly better results than the baseline BPE [30] model, outperforming the grapheme model WER by about 14.4%. Pure

text data is cheaper and easier to obtain than paired fine-tuning data, which makes the Transcoder scheme more practical.

Table 2: *WER (%) comparison of Transcoder and the direct fine-tuning of UniData2vec<sup>+</sup>. Both the two fine-tuning in the table using 1 hour of data. The BPE model has an output label of 4000 BPE units. The Grapheme model uses grapheme sets from the 400k sentences.*

Model	<i>it</i>	<i>es</i>	<i>fr</i>	avg
UniData2vec <sup>+</sup> + BPE fine-tune	63.5	62.9	69.7	65.4
UniData2vec <sup>+</sup> + Grapheme fine-tune	29.6	29.5	45.3	34.8
UniData2vec <sup>+</sup> + Transcoder	<b>27.4</b>	<b>25.7</b>	<b>36.2</b>	<b>29.8</b>

**Ablation study of Transcoder** We conduct ablation experiments to verify the role of each component in the Transcoder. As shown in Table 3, we evaluate the performance of our proposed Transcoder architecture by ablating its key components, including the use of probabilistic inputs, two convolution modules, and CE loss, using correct text and UniData2vec<sup>+</sup> hypothesis inputs as evaluation metrics. The results demonstrate that removing each of these components leads to weighted average WER increase of 7.3%, 7.0%, and 63.2%, respectively. The Conformer architecture using CE loss compared to our Transcoder resulted in a higher weighted average WER of 6.3%, demonstrating the effectiveness of our proposed architecture.

Table 3: *Ablation study of each component's impact on WER (%) of the Transcoder. Only Italian text is used in experiment.*

Input data	Transcoder			Conformer
	Standard	w/o Prob inputs	w/o Two conv w/o CE loss	
Correct text	<b>3.2</b>	3.2	3.6	3.5
Hypothesis	<b>27.4</b>	29.4	27.8	28.3

## 5. Conclusions

This paper presents an effective approach for low-resource languages that utilizes unified representation learning on phonetically-aware representations, followed by a P2W Transcoder. The proposed UniData2vec model achieves an improvement of over 5.3% compared to the SOTA method, while the Transcoder model yields a 14.4% relative reduction in word error rate compared to grapheme fine-tuning. However, UniData2vec<sup>+</sup> shows weaker performance for low amounts of unlabeled data. In the future, we will unify unlabeled data from multilingual languages to pre-train.



## 6. References

- [1] J. Li, Y. Wu, Y. Gaur, C. Wang, R. Zhao, and S. Liu, "On the comparison of popular end-to-end models for large scale speech recognition," in *21st Annual Conference of the International Speech Communication Association (Interspeech)*, 2020, pp. 1–5.
- [2] L. Besacier, E. Barnard, A. Karpov, and T. Schultz, "Automatic speech recognition for under-resourced languages: A survey," *Speech Commun.*, pp. 85–100, 2014.
- [3] H. Yadav and S. Sitaram, "A survey of multilingual models for automatic speech recognition," in *Proceedings of the Thirteenth Language Resources and Evaluation Conference (LREC)*, 2022, pp. 5071–5079.
- [4] J. Huang, J. Li, D. Yu, L. Deng, and Y. Gong, "Cross-language knowledge transfer using multilingual deep neural network with shared hidden layers," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2013, pp. 7304–7308.
- [5] S. Toshniwal, T. N. Sainath, R. J. Weiss, B. Li, P. J. Moreno, E. Weinstein, and K. Rao, "Multilingual speech recognition with a single end-to-end model," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 4904–4908.
- [6] A. Kannan, A. Datta, T. N. Sainath, E. Weinstein, B. Ramabhadran, Y. Wu, A. Bapna, Z. Chen, and S. Lee, "Large-scale multilingual speech recognition with a streaming end-to-end model," in *20th Annual Conference of the International Speech Communication Association (Interspeech)*, 2019, pp. 2130–2134.
- [7] A. Conneau and G. Lample, "Cross-lingual language model pre-training," in *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2019, pp. 7057–7067.
- [8] M. Rivière, A. Joulin, P. Mazaré, and E. Dupoux, "Unsupervised pretraining transfers well across languages," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7414–7418.
- [9] A. Conneau, A. Baevski, R. Collobert, A. Mohamed, and M. Auli, "Unsupervised cross-lingual representation learning for speech recognition," in *22nd Annual Conference of the International Speech Communication Association (Interspeech)*, 2021, pp. 2426–2430.
- [10] C. Wang, Y. Wu, Y. Qian, K. Kumatani, S. Liu, F. Wei, M. Zeng, and X. Huang, "Unispeech: Unified speech representation learning with labeled and unlabeled data," in *Proceedings of the 38th International Conference on Machine Learning (ICML)*, 2021, pp. 10937–10947.
- [11] A. Baevski, W. Hsu, A. Conneau, and M. Auli, "Unsupervised speech recognition," in *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2021, pp. 27826–27839.
- [12] A. Babu, C. Wang, A. Tjandra, K. Lakhotia, Q. Xu, N. Goyal, K. Singh, P. von Platen, Y. Saraf, J. Pino, A. Baevski, A. Conneau, and M. Auli, "XLS-R: self-supervised cross-lingual speech representation learning at scale," in *23rd Annual Conference of the International Speech Communication Association (Interspeech)*, 2022, pp. 2278–2282.
- [13] Y. Lu, M. Huang, X. Qu, P. Wei, and Z. Ma, "Language adaptive cross-lingual speech representation learning with sparse sharing sub-networks," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 6882–6886.
- [14] C. Talnikar, T. Likhomanenko, R. Collobert, and G. Synnaeve, "Joint masked CPC and CTC training for ASR," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 3045–3049.
- [15] J. Bai, B. Li, Y. Zhang, A. Bapna, N. Siddhartha, K. C. Sim, and T. N. Sainath, "Joint unsupervised and supervised training for multilingual ASR," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 6402–6406.
- [16] A. Graves, S. Fernández, F. J. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML)*, 2006, pp. 369–376.
- [17] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [18] R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, "Common voice: A massively-multilingual speech corpus," in *Proceedings of The 12th Language Resources and Evaluation Conference (LREC)*, 2020, pp. 4218–4222.
- [19] I. P. Association, *Handbook of the International Phonetic Association: A guide to the use of the International Phonetic Alphabet*. Cambridge University Press, 1999.
- [20] A. Baevski, W. Hsu, Q. Xu, A. Babu, J. Gu, and M. Auli, "data2vec: A general framework for self-supervised learning in speech, vision and language," in *International Conference on Machine Learning (ICML)*, 2022, pp. 1298–1312.
- [21] R. Masumura, N. Makishima, M. Ithori, A. Takashima, T. Tanaka, and S. Orihashi, "Phoneme-to-grapheme conversion based large-scale pre-training for end-to-end automatic speech recognition," in *21st Annual Conference of the International Speech Communication Association (Interspeech)*, 2020, pp. 2822–2826.
- [22] Y. Kubo and M. Bacchiani, "Joint phoneme-grapheme model for end-to-end speech recognition," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6119–6123.
- [23] X. Wang, Z. Yao, X. Shi, and L. Xie, "Cascade RNN-Transducer: Syllable based streaming on-device mandarin speech recognition with a syllable-to-character converter," in *Spoken Language Technology Workshop (SLT)*, 2021, pp. 15–21.
- [24] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2017, pp. 5998–6008.
- [25] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [26] N. Reimers and I. Gurevych, "Making monolingual sentence embeddings multilingual using knowledge distillation," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020, pp. 4512–4525.
- [27] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, "fairseq: A fast, extensible toolkit for sequence modeling," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, 2019, pp. 48–53.
- [28] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations (ICLR)*, 2015.
- [29] W. Hsu, B. Bolte, Y. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE ACM Trans. Audio Speech Lang. Process.*, pp. 3451–3460, 2021.
- [30] R. Sennrich, B. Haddow, and A. Birch, "Neural machine translation of rare words with subword units," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2016.