

Seamless Language Expansion: Enhancing Multilingual Mastery in Self-Supervised Models

Jing Xu¹, Minglin Wu¹, Xixin Wu^{1,2}, Helen Meng^{1,2}

¹Department of Systems Engineering and Engineering Management,
The Chinese University of Hong Kong, Hong Kong SAR, China

²Centre for Perceptual and Interactive Intelligence (CPII) Limited, HKSAR, China

{1155170427, minglinwu}@link.cuhk.edu.hk, {wuxx, hmmeng}@se.cuhk.edu.hk

Abstract

Self-supervised (SSL) models have shown great performance in various downstream tasks. However, they are typically developed for limited languages, and may encounter new languages in real-world. Developing a SSL model for each new language is costly. Thus, it is vital to figure out how to efficiently adapt existed SSL models to a new language without impairing its original abilities. We propose adaptation methods which integrate LoRA to existed SSL models to extend new language. We also develop preservation strategies which include data combination and re-clustering to retain abilities on existed languages. Applied to mHuBERT, we investigate their effectiveness on speech re-synthesis task. Experiments show that our adaptation methods enable mHuBERT to be applied to a new language (Mandarin) with MOS value increased about 1.6 and the relative value of WER reduced up to 61.72%. Also, our preservation strategies ensure that the performance on both existed and new languages remains intact¹.

Index Terms: adaptation, parameter efficient, multi-lingual, language preservation, speech re-synthesis

1. Introduction

Self-supervised pre-trained models have significantly advanced the field of speech technology, offering robust solutions for learning general representations from vast quantities of unlabeled data[1, 2, 3, 4]. By harnessing the abundant unlabeled data, these models diminish the dependency on expensive and labor-intensive labeled data. Notably, models like Wav2vec 2.0[3] and HuBERT[4] can be easily fine-tuned for a range of downstream tasks, such as automatic speech recognition (ASR), speech synthesis and speech tokenization.

SSL models are typically developed for a limited set of languages[5], but they may encounter new languages when applied in real-world scenarios. Besides, developing a new SSL model for each new language is at high cost due to its large parameters. Hence, it is vital to adapt existed SSL models to a new language. Adapting SSL models to a new language is very challenging, as phonetic and morphological structures differ largely across languages. The situation is more severe when incorporated languages and the new language are from disparate linguistic families and have completely different attributes. For example, languages encompassed in pre-training are non-tonal such as English and the target language is a tonal language like Mandarin. The difference between tonal language and non-tonal language is huge because tonal languages distinguish lexical or grammatical meanings through pitch variations,

while non-tonal languages do not. To deal with language mismatch problem, strategies such as self-supervised adaptive pre-training and language adaptors have been proposed[6, 7]. In [6], unlabeled data collected from the target language of downstream tasks are leveraged to adapt the pre-trained model prior to the last fine-tuning stage. Meta-Adaptable-Adapter is proposed in [7] to learn task-specific adapters for feature extraction and task-independent adapters for feature combination. However, the relationship between parameter efficiency and adaptation performance remains unclear.

When adapting an existed model to new data or new tasks, there is a phenomenon called catastrophic forgetting[8] which indicates that the model tends to forget previously acquired knowledge. The underlying cause is attributed to the model's weights update during adaptation, which inadvertently overwrite antecedent learned representations. This phenomenon can markedly impede the model's performance and capacity for generalization. To mitigate catastrophic forgetting, researchers in speech area have investigated various techniques[9, 10] to strike a balance between adapting to new tasks or data and preserving acquired knowledge. A multi-task learning framework is advocated in [9] where reserving original knowledge and learning new knowledge are treated as two independent tasks. Alternative methodologies involve the orthogonal modification of weights[10, 11, 12]. For instance, [10] introduced a continual learning algorithm known as Regularized Adaptive Weight Modification (RAWM) for fake audio detection. Nevertheless, the strategies for sustaining the knowledge of existed languages while adapting to new ones remain insufficiently explored.

In this paper, to deal with language mismatch problem, we propose parameter efficient adaptation methods which incorporate LoRA into existed SSL models instead of developing a new one. Additionally, we develop preservation strategies that include data combination and re-clustering to mitigate catastrophic forgetting problem. They enable SSL models to be applied to new languages parameter efficiently while retaining capabilities on existed languages. Specifically, the SSL model under study is mHuBERT[13], while the existed language and new language to adapt are English and Mandarin respectively. Speech re-synthesis task is adopted to evaluate the representational ability of adapted SSL model and the efficiency of adaptation. The discrete units of SSL models can naturally serve as the input of speech re-synthesis and open the potential for speech language model (SLM). Speech re-synthesis is a more difficult task than ASR since it requires both representation comprehension and waveform reconstruction. Experiments reveal that the MOS value can improve from 2.26 to 3.80 using our adaptation methods. Also, a 61.72% relative word error rate (WER) reduction compared to un-adapted model in new language speech re-synthesis is observed. These results verify the effectiveness

¹Demo page: <https://jingxu96.github.io/multilingual-ssl-demo/index.html>




Figure 1: (a) Architecture of self-supervised pre-trained models. (b) Architecture of Transformer blocks in self-supervised pre-trained models. (c) Architecture of LoRA which is integrated to the self-attention module of Transformer blocks.

of our adaptation methods on adapting SSL model to a new language. Moreover, results indicate that our preservation strategies substantially mitigate the performance degradation in original languages while maintaining the newly acquired language capabilities, thereby indicating a significant mitigation of catastrophic forgetting.

2. Method

In this section, we first present the framework of our proposed adaptation method in subsection 2.1. Next, in subsection 2.2, we illustrate two adaptation strategies while adapting self-supervised model. In subsection 2.3, we propose two preservation strategies to mitigate catastrophic forgetting caused by adaptation. The evaluation system - speech re-synthesis system is introduced in subsection 2.4.

2.1. Parameter efficient Adaptation

As shown in Figure 1(a), Transformer-based SSL models are generally composed of two principal components: a feature encoder and a Transformer network. The feature encoder, typically a deep neural network, is responsible for transforming raw audio signals into latent representations. Subsequently, these representations are input to the Transformer network, which further process them to contextualized representations. The learning paradigm of these models is multifaceted, encompassing various tasks such as generative task[14], contrastive task[3] and predictive task[2, 4, 15].

Transformer network constitutes the cornerstone of SSL models' success, facilitating the capturing of intricate relationships within speech sequences. Transformer network is stacked by multiple Transformer blocks, whose architecture is shown in Figure 1(b). Each block is equipped with a multi-head self-attention module, a multi-layer perceptron (MLP) layer, and layer normalization layers. Multi-head self-attention module, pivotal to the Transformer block's functionality, enables each temporal frame to compute an attention score with respect to all other frames, thereby capturing their relevance. In pursuit of efficiently enhancing multilingual mastery of SSL models, we advocate the integration of Low Rank Adaptation (LoRA)[16] into the multi-head attention module, thereby augmenting the adaptation efficiency.

As illustrated in Figure 1(c), the update of each pre-trained weight matrix $W_0 \in \mathbb{R}^{d \times k}$, is constrained by representing it via

a low-rank decomposition:

$$W_0 + \Delta W = W_0 + BA \quad (1)$$

where $B \in \mathbb{R}^{d \times r}$ and $A \in \mathbb{R}^{r \times k}$, and the rank $r \ll \min(d, k)$.

During adaptation, only parameters of A and B are updated, W_0 are frozen, exempt from gradient updates. The inputs which multiplied with W_0 and $\Delta W = BA$ are identical, culminating the summation of their resultant output vectors on a coordinate-wise basis. Given $h_0 = W_0 x$, the modified forward pass is:

$$h = W_0 x + \Delta W x = W_0 x + BAx \quad (2)$$

Same as [16], we adopt a random Gaussian initialization for A and zero for B , so at the beginning of training $\Delta W = BA$ is zero. Matrix $\Delta W x$ is scaled by α/r , wherein α is a constant in r . If the initialization is set appropriately, α is a hyperparameter similar to the learning rate during training. Owing to the exclusive update of A and B , the adaptation achieves great parameter efficiency. Furthermore, the stasis of matrix W_0 implies the preservation of the majority of the SSL parameters, thereby conferring the advantage of sustained performance across existed languages.

2.2. Adaptation strategies

We propose two adaptation strategies to accomplish adaptation based on the ways of obtaining target labels. In each strategy, model parameters are initialized from an existed SSL model. In our case, the model parameters are all initialized from mHuBERT[13]. We will illustrate the training strategies in mHuBERT setting.

- One-iteration adaptation strategy: In this strategy, the existed SSL model is directly adopted to generate continuous representations of new language. A new K-means model is trained to produce coarse target label for the new language, rendering it a time-efficient training strategy.
- Two-iteration adaptation strategy: In this adaptation strategy, we iterative refine the target label. The coarse target label is derived by applying K-means on MFCC. It is used for model training in the first iteration. By applying K-means on the intermediate layer of the trained model, we can obtain fine-grained target label, which is leveraged for model training in the second iteration. To ensure a fair comparison with one-iteration adaptation strategy, the model parameters

for the second iteration are initialized from publicly released model, rather than model resultant from the first iteration.

2.3. Preservation strategies

After adaptation, we observe the catastrophic forgetting phenomenon wherein the performance drops largely on the previously existing language. To alleviate this problem, we propose two strategies aimed at preserving the ability on existed languages of adapted model.

- Re-clustering using combined data. In this strategy, we introduce both English speech and Mandarin speech into adapted model to produce their respective representations from a designated intermediate layer. These speech samples are randomly selected from a substantial dataset, ensuring parity in the volume of English and Mandarin data. Subsequently, a K-means model is trained utilizing the combined English and Mandarin representations.
- Data combination in adaptation training: In this strategy, we combine English speech with Mandarin speech during adaptation training. To safeguard against a substantial diminution in performance of Mandarin, the ratio of English to Mandarin is maintained below 1. Besides, given the prior integration of English data, we limit its inclusion to a modest quantum. The K-means model is also trained on composite English and Mandarin representations.

2.4. Speech re-synthesis system

Speech re-synthesis offers a more holistic evaluation of the representational capabilities of SSL models. The discrete units derived from SSL models can readily serve as inputs of speech re-synthesis and open the potential for advances of SLM. Therefore, we propose to investigate the performance on speech re-synthesis task. The speech re-synthesis system, as shown in Figure 2, is comprised of a SSL model based discrete unit extractor and a unit-HiFiGAN vocoder[17].

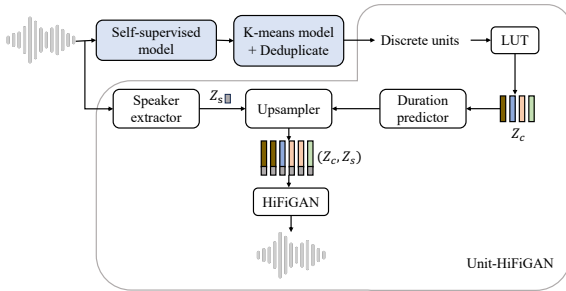


Figure 2: Architecture of speech re-synthesis system.

- Discrete unit extractor. Discrete unit extractor aims to discretize representations extracted by self-supervised model into cluster indices using K-means model. The sequence of cluster indices are de-duplicated and serve as discrete units.
- Unit-HiFiGAN. A neural vocoder is employed to transform speech discrete units back into speech signal. As shown in Figure 2, as an input, the discrete units are converted into continuous representations Z_c via look-up tables (LUT). Speaker embedding Z_s is extracted from wavlm[2], which is trained for speaker verification task, to support multi-speaker speech re-synthesis. The sequences Z_c and Z_s are then up-sampled based on the duration predicted by duration predictor. Speaker embedding Z_s is concatenated to each frame of

up-sampled embedding Z_c . Then, the concatenated sequence is transformed into waveform by the generator of HiFiGAN.

3. Experimental setup

We evaluate the efficacy of our adaptation methods on adapting mHuBERT[13] to Mandarin, a challenging tonal language. HuBERT is known for its ability of discretizing speech signals into discrete units, and have been widely used in speech language models, such as SpeechLM[18], SpeechGPT[19] and SpeechGen[20]. The analysis of WER between transcripts of synthesized speech and ground truth ones serves a robust metric for measuring the performance of our adaptation methods and preservation strategies.

In the two-iteration adaptation strategy, the target label of initial iteration is the clustering indices of 39-dimensional MFCC features. Subsequently, the representations extracted from 6th-layer of the model, obtained in first iteration, are clustered to form the target label of second iteration. In one-iteration adaptation strategy, the representations are extracted from 11th layer of released model. The initial iteration in two-iteration adaptation strategy is trained for 250k steps. To have a fair comparison between two adaptation strategies, the second iteration is trained for 400k steps same as the one-iteration strategy.

We use MiniBatchKMeans algorithm to realize K-means clustering. The mini-batch size is set to 10,000 frames. K-means++ with 20 random starts is used for better initialization. We set the rank r in the adaptation methods to be 24 and integrate LoRA to the projection layer of query, key, value and output of self-attention layer in each Transformer block.

During adaptation, AISHELL-2[21] is employed, comprising 1000 hours of clean Mandarin reading-speech data. The English speech data used in preservation strategies is LibriSpeech[22], from which a subset of 100 hours are randomly selected. The speech re-synthesis task is conducted using AISHELL-3[23], an 85-hour high-fidelity multi-speaker Mandarin speech corpus designed for Text-to-Speech (TTS). Regarding the English corpus for speech re-synthesis, VCTK[24] is chosen, encompassing about 44-hour speech uttered by 110 English speakers. From the VCTK corpus, a random selection of 5% of speakers constitutes the test set, while 95% of the remaining speech from other speakers form the training set, with the rest allocated as the validation set. The speech samples are uniformly down-sampled to 16 kHz.

4. Results

4.1. Results on Mandarin adaptation

We conduct an analysis of WER and administer a Mean Opinion Score (MOS) test on the generated speech to evaluate the effectiveness. The transcripts of generated speech are generated by Whisper large-v3[25]. For all MOS tests in this paper, 15 listeners are invited to assess 10 samples randomly selected from each system based on their naturalness. A reduction in WER and an increase in MOS after adaptation would suggest that the adapted model is capable of effectively encoding and tokenizing the new language, thereby serving as a robust indicator of successful adaptation.

As shown in "Before adaptation" of Table 1, training K-means model without adaptation yields in an 18.90% relative reduction in WER, suggesting that improved tokenization enhances performance. Applying one-iteration adaptation strategy and two-iteration adaptation strategy to SSL model can achieve

Table 1: Effectiveness of adaptation methods on Mandarin speech re-synthesis task

System	Size	Trainable param.	WER↓	MOS↑
Ground Truth	-	-	8.02%	4.58±0.07
Mandarin HuBERT	94.8 M	-	26.79%	3.27±0.09
Before adaptation				
Un-adapted mHuBERT + Un-trained K-means	94.8 M	-	50.57%	2.26±0.11
Un-adapted mHuBERT + Trained K-means	94.8 M	-	41.01%	2.55±0.12
After adaptation				
Adapted mHuBERT (One-iteration) + Trained K-means	96.6 M	2.026 M (2.098%)	32.88%	2.77±0.12
Adapted mHuBERT (Two-iteration) + Trained K-means	96.6 M	2.026 M (2.098%)	21.94%	3.80±0.08

relative WER reductions of 34.98% and 56.61% respectively. This implies a potential trade-off between time-efficiency and adaptation performance, as two-iteration strategy gives superior target labels. Furthermore, Table 1 reveals the parameter efficiency of our adaptation methods with less than 3% of model parameters trainable. The results of MOS corroborate the WER findings, thereby validating our methods’ efficacy. Additionally, a widely used open-sourced implementation of Mandarin HuBERT trained on 1W+ hours of Wenetspeech[26] exhibits marginally inferior performance compared to our two-iteration adaptation strategy, which may be due to the differences in data distribution and suggest our adaptation strategies can transfer knowledge from other languages.

4.2. Results on language preservation and adaptation

To assess the extent of catastrophic forgetting, we perform a language maintenance test, depicted in "Before preservation" part of Table 2. In this test, the adapted mHuBERT model is utilized to extract English speech representations whereas the K-means model is trained on representations derived from un-adapted mHuBERT. English unit-HiFiGAN is utilized to transform speech units into waveform. It is observed that one-iteration adaptation strategy exhibits superior performance on existed language and after adaptation, there is a noticeable decline in performance on existed languages.

Table 2: Effectiveness of preservation strategies on English (en) and Mandarin (zh). In the system notion, the term preceding the hyphen denotes the language used during adaptation, whereas the term following it indicates the adaptation strategy.

System	en WER↓	zh WER↓	en MOS↑	zh MOS↑
GT	1.44%	8.02%	4.63±0.07	4.58±0.07
English HuBERT	11.48%	-	4.03±0.08	-
Before preservation				
zh-2iter	38.17%	21.94%	2.84±0.10	3.80±0.08
zh-1iter	15.50%	32.88%	3.97±0.08	2.77±0.12
After preservation				
zh-2iter	13.84%	23.71%	4.04±0.08	3.61±0.08
zh-1iter	9.70%	34.36%	4.24±0.07	2.80±0.10
enzh-2iter	7.94%	27.18%	4.21±0.08	3.35±0.10

The effectiveness of preservation strategies are shown in Table 2. Speech units of Mandarin and English are converted using bi-lingual unit-HiFiGAN. It is observed that employing solely Mandarin speech for adaptation, one-iteration adaptation strategy yields better results in English but worse results in Mandarin. This phenomenon underscores the delicate balance between adapting a new language and preserving existed languages. Subsequent to re-clustering with combined data, there is a marked improvement in the performance of existed language, with WER reducing from 38.17% to 13.84% in two-iteration adaptation strategy. Besides, our method behaves even better with the original model (before any adaptation) in ex-

isted language with the performance in new language remains largely unaffected. These results confirmed the performance of our preservation strategies. Discrepancies between WER and MOS are noted, particularly when WER values are close, as discerning differences between two similarly performing systems may be challenging for listeners. Nevertheless, after adaptation, there is a significant increase in MOS for English, rising from 2.84 to 4.04 in two-iteration adaptation strategy. The strategy of combining data during adaptation has been shown to improve performance in existed languages while reducing more performance on new language. The impressive performance of adapted mHuBERT, utilizing speech exclusively from new language and employing re-clustering preservation strategy, facilitates the processing of both Mandarin and English effectively.

4.3. Ablation study

We also enlarge the clustering number of K-means model to 3000 to examine the impact of clustering number. For Adaptation, we only focus on the performance of new language, whereas for preservation, we consider the balance between new languages and existed language. It is observed from Table 3 that increasing the number of clusters yields the most substantial adaptation performance, with a relative WER reduction of up to 61.72% compared to un-adapted model. During preservation, the performance enhancement in both existed and new languages can also be observed, with WER reduced up to around 2% in each language.

Table 3: Ablation study on cluster numbers

System	en WER↓	zh WER↓
Adaptation		
zh-2iter-1000	-	21.94%
zh-2iter-3000	-	19.36%
Preservation		
zh-2iter-1000	13.84%	23.71%
zh-2iter-3000	12.53%	21.5%

5. Discussion

In this paper, we introduce parameter efficient adaptation methods which integrate LoRA into existed SSL models and preservation strategies including re-clustering and data combination to adapt an existed SSL model to a new language without harming its original abilities. We select speech re-synthesis task to evaluate the efficacy of our proposed approaches. Experiments corroborate that our methods can seamlessly improve the multi-lingual mastery in SSL models while preserving capabilities for existed languages. Currently, we only focus on a most-widely adopted and representative SSL model, mHuBERT. In the future, we plan to assess the effectiveness of our methods across a broader range of SSL models and diverse downstream tasks.

6. Acknowledgements

The first author Jing Xu would like to thank the guidance and support from Xiao Chen, Daxin Tan and Yuting Yeung. This work is partially supported by the Human-Computer Communications Laboratory, Department of Systems Engineering & Engineering Management, The Chinese University of Hong Kong.

7. References

- [1] A. Mohamed, H.-y. Lee, L. Borgholt, J. D. Havtorn, J. Edin, C. Igel, K. Kirchhoff, S.-W. Li, K. Livescu, L. Maaloe, T. N. Sainath, and S. Watanabe, "Self-supervised speech representation learning: A review," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, p. 1179–1210, Oct. 2022. [Online]. Available: <http://dx.doi.org/10.1109/JSTSP.2022.3207050>
- [2] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [3] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [4] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [5] A. Conneau, A. Baevski, R. Collobert, A. Mohamed, and M. Auli, "Unsupervised cross-lingual representation learning for speech recognition," *arXiv preprint arXiv:2006.13979*, 2020.
- [6] M. M. Shaik, D. Klakow, and B. M. Abdullah, "Self-supervised adaptive pre-training of multilingual speech models for language and dialect identification," *arXiv preprint arXiv:2312.07338*, 2023.
- [7] Y. Chen, H. Zhang, X. Yang, W. Zhang, and D. Qu, "Meta-adaptable-adaptor: Efficient adaptation of self-supervised models for low-resource speech recognition," *Available at SSRN 4641394*.
- [8] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska *et al.*, "Overcoming catastrophic forgetting in neural networks," *Proceedings of the national academy of sciences*, vol. 114, no. 13, pp. 3521–3526, 2017.
- [9] J. Xue, J. Han, T. Zheng, X. Gao, and J. Guo, "A multi-task learning framework for overcoming the catastrophic forgetting in automatic speech recognition," *arXiv preprint arXiv:1904.08039*, 2019.
- [10] X. Zhang, J. Yi, J. Tao, C. Wang, and C. Y. Zhang, "Do you remember? overcoming catastrophic forgetting for fake audio detection," 2023.
- [11] S. Wen and L. Itti, "Overcoming catastrophic forgetting problem by weight consolidation and long-term memory," *arXiv preprint arXiv:1805.07441*, 2018.
- [12] X. Liu, M. Masana, L. Herranz, J. Van de Weijer, A. M. Lopez, and A. D. Bagdanov, "Rotate your networks: Better weight consolidation and less catastrophic forgetting," in *2018 24th International Conference on Pattern Recognition (ICPR)*. IEEE, 2018, pp. 2262–2268.
- [13] A. Lee, H. Gong, P.-A. Duquenne, H. Schwenk, P.-J. Chen, C. Wang, S. Popuri, Y. Adi, J. Pino, J. Gu *et al.*, "Text-less speech-to-speech translation on real data," *arXiv preprint arXiv:2112.08352*, 2021.
- [14] A. T. Liu, S.-w. Yang, P.-H. Chi, P.-c. Hsu, and H.-y. Lee, "Mockingjay: Unsupervised speech representation learning with deep bidirectional transformer encoders," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6419–6423.
- [15] 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," in *International Conference on Machine Learning*. PMLR, 2022, pp. 1298–1312.
- [16] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "Lora: Low-rank adaptation of large language models," *arXiv preprint arXiv:2106.09685*, 2021.
- [17] A. Polyak, Y. Adi, J. Copet, E. Kharitonov, K. Lakhotia, W.-N. Hsu, A. Mohamed, and E. Dupoux, "Speech resynthesis from discrete disentangled self-supervised representations," *arXiv preprint arXiv:2104.00355*, 2021.
- [18] 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.
- [19] D. Zhang, S. Li, X. Zhang, J. Zhan, P. Wang, Y. Zhou, and X. Qiu, "Speechgpt: Empowering large language models with intrinsic cross-modal conversational abilities," *arXiv preprint arXiv:2305.11000*, 2023.
- [20] H. Wu, K.-W. Chang, Y.-K. Wu, and H.-y. Lee, "Speechgen: Unlocking the generative power of speech language models with prompts," *arXiv preprint arXiv:2306.02207*, 2023.
- [21] J. Du, X. Na, X. Liu, and H. Bu, "Aishell-2: Transforming mandarin asr research into industrial scale," *arXiv preprint arXiv:1808.10583*, 2018.
- [22] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [23] Y. Shi, H. Bu, X. Xu, S. Zhang, and M. Li, "Aishell-3: A multi-speaker mandarin tts corpus and the baselines," *arXiv preprint arXiv:2010.11567*, 2020.
- [24] J. Yamagishi, C. Veaux, and K. MacDonald, "Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit (version 0.92)," 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:213060286>
- [25] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *International Conference on Machine Learning*. PMLR, 2023, pp. 28 492–28 518.
- [26] B. Zhang, H. Lv, P. Guo, Q. Shao, C. Yang, L. Xie, X. Xu, H. Bu, X. Chen, C. Zeng *et al.*, "Wenetspeech: A 10000+ hours multi-domain mandarin corpus for speech recognition," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6182–6186.