



LoRA-Whisper: Parameter-Efficient and Extensible Multilingual ASR

Zheshu Song¹, Jianheng Zhuo¹, Yifan Yang¹, Ziyang Ma¹, Shixiong Zhang², Xie Chen^{1,†}

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

²Tencent AI Lab, USA

{songzheshu, chenxie95}@sjtu.edu.cn

Abstract

Recent years have witnessed significant progress in multilingual automatic speech recognition (ASR), driven by the emergence of end-to-end (E2E) models and the scaling of multilingual datasets. Despite that, two main challenges persist in multilingual ASR: language interference and the incorporation of new languages without degrading the performance of the existing ones. This paper proposes LoRA-Whisper, which incorporates LoRA matrix into Whisper for multilingual ASR, effectively mitigating language interference. Furthermore, by leveraging LoRA and the similarities between languages, we can achieve better performance on new languages while upholding consistent performance on original ones. Experiments on a real-world task across eight languages demonstrate that our proposed LoRA-Whisper yields a relative gain of 18.5% and 23.0% over the baseline system for multilingual ASR and language expansion respectively.

Index Terms: multilingual speech recognition, language expansion, Whisper, LoRA

1. Introduction

Automatic speech recognition (ASR) has traditionally concentrated on transcribing speech into written text for single languages [1–5]. Nevertheless, as the demand for cross-lingual communication grows and vast multilingual datasets [6–9] become more accessible, attention has recently turned towards the development of massively multilingual ASR models. With the emergence of large-scale multilingual speech recognition models such as Whisper [10], Google USM [11], and MMS [12], individuals now have the opportunity to construct **customized multilingual speech recognition models** tailored to specific languages based on these foundational models.

However, two significant challenges are still yet to be addressed in multilingual ASR. One is language interference, primarily stemming from language overlap, data imbalance, dialectal accents, etc. Another challenge involves incorporating new languages without compromising the performance of existing ones. To resolve the former problem, there are a series of previous works attempting to mitigate this issue by leveraging language ID information [13, 14] or designing language-specific modules [15–21] such as languages-specific encoders to differentiate each language. Besides, some works [22–24] utilize a pruning strategy in multilingual ASR with a dedicated sub-model for each language, while others propose new sampling methods [25] to address the data imbalance issue. Although the methods mentioned above alleviate language interference to some extent, they are somewhat cumbersome in design and

fail to account for language expansion. When new languages need to be integrated into a multilingual ASR system, a naive approach is to fine-tune the ASR model using data from these new languages. Unfortunately, this often results in *catastrophic forgetting*, referring to the phenomenon that the recognition performance of base languages tends to decline. To solve the above problem, Li et al. [26] proposes lifelong learning [27] solution which remedies the language interference problem by mixing base language data and new language data. However, this approach is inefficient and time-consuming. Libera et al. [28] explores various continual learning methods [29–34] to address the issue of catastrophic forgetting. While these approaches have helped alleviate the problem, it still persists.

Towards this end, we introduce LoRA-Whisper, a parameter-efficient and extensible model for multilingual ASR. LoRA [35], originally introduced in natural language processing (NLP), effectively customizes large language models (LLMs) for specific domains. Drawing inspiration from this, it can also be used to tailor speech recognition models for specific languages. In practice, we assign a language-specific LoRA matrix for each language. This approach allows shared information across languages to be stored within the Whisper model, while language-specific information can be captured in the respective LoRA matrices. When incorporating a new language, a new LoRA matrix is assigned for it, ensuring no impact on the performance of existing languages. Furthermore, by capitalizing on the similarities between the new language and base languages, we can enhance performance on the new language through improved initialization of the new LoRA matrix or by employing mixture of experts (MoE) [36]. Note that the foundational model is not restricted to Whisper but can encompass other open-source speech recognition models as well, we are simply utilizing Whisper as an exemplar in this paper. In summary, the contributions of this paper are as follows:

- We propose LoRA-Whisper to mitigate language interference and avoid catastrophic forgetting when incorporating new languages by attaching language-specific LoRA modules to the Whisper model.
- By utilizing the similarity between languages, notable performance improvement can be achieved on new languages via better initialization of the new LoRA matrix or the employment of MoE.

2. Background

2.1. Whisper

Whisper [10] is an encoder-decoder Transformer model that is capable of multiple speech tasks, including multilingual speech recognition, speech translation, language identification,

† Corresponding author

and voice activity detection. The input to Whisper is an 80-dimensional log-Mel spectrogram of 30 seconds length $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T]$ where T denotes the context length. The encoder blocks encode the input speech feature into hidden representations \mathbf{H} and the decoder blocks decode the hidden representations into text tokens $\hat{\mathbf{y}}$ recursively conditioned on previous tokens and special prompts \mathbf{p} . In formal terms, this process can be illustrated as follows:

$$\mathbf{H} = \text{AudioEncoder}(\mathbf{X}) \quad (1)$$

$$\hat{\mathbf{y}}_i = \text{TextDecoder}(\mathbf{p}, \hat{\mathbf{y}}_{1:t-1}, \mathbf{H}) \quad (2)$$

2.2. LoRA

LoRA [35] was initially introduced in the field of natural language processing (NLP) as a means to effectively tailor large language models (LLMs) for specific domains or downstream tasks. It was observed that the weights of pre-trained LLMs tend to exist mainly in a low-dimensional space. Taking inspiration from this observation, LoRA reduces the number of trainable parameters by learning pairs of rank decomposition matrices while keeping the original weights fixed. Specifically, consider i -th feed forward layer $f_i(\mathbf{x}) = \mathbf{W}_i\mathbf{x} + \mathbf{b}_i$, where $\mathbf{W}_i \in \mathbb{R}^{d_1 \times d_2}$ and $\mathbf{b}_i \in \mathbb{R}^{d_1}$ denotes the frozen weight and bias. By applying LoRA, the forward process is modified as:

$$f_i(\mathbf{x}) = (\mathbf{W}_i + \Delta\mathbf{W}_i)\mathbf{x} + \mathbf{b}_i; \Delta\mathbf{W}_i = \mathbf{B}_i\mathbf{A}_i \quad (3)$$

where $\mathbf{B}_i \in \mathbb{R}^{d_1 \times r}$ and $\mathbf{A}_i \in \mathbb{R}^{r \times d_2}$ are the two trainable low-rank matrices, with the rank $r \ll \min(d_1, d_2)$.

3. Methods

In this paper, we mainly focus on tackling two challenges in multilingual ASR: one is language interference and the other is new language incorporation. Section 3.1 briefly outlines the main issues to be addressed in this paper. Section 3.2 and Section 3.3 introduce the methods used in multilingual ASR and language expansion in detail.

3.1. Problem statement

In our research, n base languages are employed for the multilingual ASR experiment, alongside an additional m new languages for the language expansion experiment, which can be denoted as $S_1 = \{(\mathbf{X}_i, \mathbf{Y}_i), i \in (1, n)\}$ and $S_2 = \{(\mathbf{X}_j, \mathbf{Y}_j), j \in (n+1, n+m)\}$ where $\mathbf{X}_i, \mathbf{Y}_i$ denotes the speech and transcription of i -th language.

In the multilingual ASR experiment, the aim is to alleviate language interference in S_1 and improve the performance of base languages. The goal of language expansion is to incorporate S_2 into the multilingual model while maintaining the performance of S_1 unaffected, and leverage similarities between S_1 and S_2 to enhance the performance specifically on S_2 .

3.2. Multilingual ASR

Applying LoRA in multilingual ASR is an effective approach to mitigate language interference as shown in Figure 1. For each language, a language-specific LoRA matrix is appended to the encoder and decoder of Whisper. When the input is a piece of speech of k -th language, it will activate k -th LoRA module and pass through Whisper and the corresponding LoRA module in the forward pass.

Under the LoRA-Whisper model, shared information

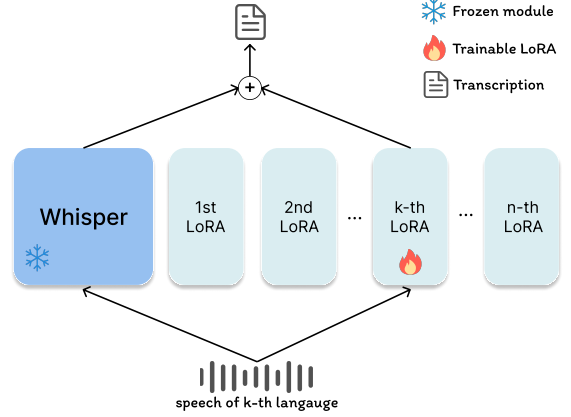


Figure 1: Architecture of LoRA-Whisper in multilingual ASR

across languages resides within the original Whisper model, while language-specific information is stored in the respective LoRA module. As a result, not only is the language interference problem skillfully avoided, but the performance of Whisper on specific languages is also significantly enhanced.

3.3. Language expansion

Apart from mitigating language interference, LoRA can also be naturally extended for language expansion, preventing catastrophic forgetting. Moreover, harnessing the similarities across languages can facilitate more effective training for new languages. Consequently, we introduce two effective methods for language expansion, namely LoRA warm start and LoRA MoE as depicted in Figure 2. These methods involve two steps.

Step 1: Find the most similar language When incorporating a new language into the existing model, we first randomly sample M audio segments from the new language data. These audio are then processed using the Whisper model for language detection. The output provides a probability distribution over all languages. In our experiment, the focus lies solely on the languages employed in the aforementioned multilingual ASR experiment. Specifically, we extract the probabilities associated with these languages and normalize them, which can be denoted as $\mathbf{p}_i = [p_{i1}, p_{i2}, \dots, p_{in}]$; $i = 1, \dots, M$. When incorporating a new language, the most similar language to it can be found by measuring the similarity between the new language and base languages, which is defined as follows:

$$\text{sim}_k = \frac{\sum_{i=1}^M \mathbb{I}(k = \arg \max_j p_{ij})}{M} \quad \text{for } k = 1, \dots, n \quad (4)$$

where \mathbb{I} is indicator function and sim_k is the defined similarity between new language and k -th base language.

Step 2: Continual training on new languages After finding the most similar one, we can leverage the information from base languages to facilitate the training of the new language. In LoRA warm start, the new LoRA matrix is initialized from the LoRA matrix of its most similar language. In LoRA MoE, two LoRA modules are selected in the forward pass to assist in the training of the new language.

In general, aside from avoiding the issue of catastrophic forgetting, LoRA-Whisper enables the exploitation of similarities between new language and base languages, leading to better performance on the new language.

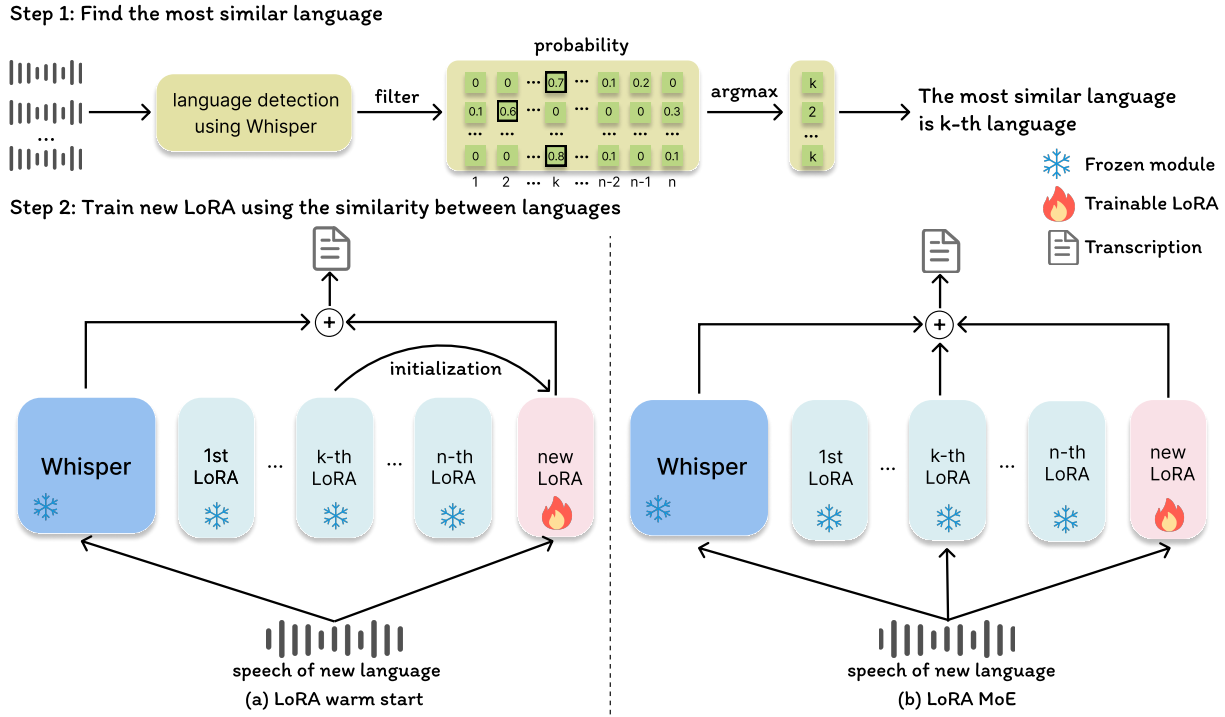


Figure 2: Architecture of LoRA-Whisper in language expansion. Left: LoRA warm start, Right: LoRA MoE

4. Experiments

4.1. Dataset

Our experiments are conducted on MLS dataset [6] and FLEURS [7] dataset as shown in Table 1. Due to limited resources, we only focus on Polish, Portuguese and Italian in MLS dataset. Similarly, five languages are selected from FLEURS dataset, namely Chinese, Danish, Greek, Welsh and Japanese. Four languages are employed in the multilingual ASR experiment and the remaining four languages are used as new languages in the language expansion experiment.

Table 1: Statistics of training and testing data (in hours)

Experiment	Language	Train	Test
Multilingual ASR	Polish(PL)	103.0	2.0
	Portuguese(PT)	160.0	5.0
	Italian(IT)	247.0	5.0
	Chinese(ZH)	9.7	3.1
Language expansion	Danish(DA)	7.5	2.9
	Greek(EL)	10.0	1.9
	Welsh(CY)	12.2	4.3
	Japanese(JA)	7.4	2.4

4.2. Training configuration

We evaluate the performance of our proposed method on Whisper-small. The impact of different LoRA rank on model performance is studied as shown in Table 2. It can be seen that best performance is achieved under $r = 32$. Hence, low-rank matrices where rank $r = 32$ are added to the attention layer

Table 2: WER on MLS Polish under different LoRA rank r

Rank r	8	16	32	48	64
#LoRA param	3.2M	6.5M	13.0M	19.5M	26.0M
WER	9.00	8.54	7.94	8.00	8.00

$\{W_k, W_q, W_v\}$ and fully-connected layer W_{fc} in each transformer layer in both encoder and decoder.

In the training stage, we fix all the parameters of Whisper and optimize the language-specific LoRA modules with AdamW [37] with a peak learning rate of $1e-4$. The number of training epochs is set to 10. All models are trained with 2 NVIDIA RTX 3090 24GB GPUs. In the testing stage, beam search with $beamsize = 5$ is employed to decode the test set.

4.3. Results and analysis

Multilingual ASR The results of multilingual ASR are summarized in Table 3. It can be observed that monolingual full fine-tuning yields optimal outcomes at the cost of training and maintaining four systems, leading to higher maintenance costs. In contrast to monolingual full fine-tuning, multilingual full fine-tuning would lead to significant language interference issues when all training data are mixed together, with the average WER increasing from 9.19% to 11.68% (Table 3: E2 vs E3). Our proposed approach, LoRA-Whisper, effectively eliminates interference between languages by incorporating language-specific LoRA matrices. In comparison to multilingual full fine-tuning, LoRA-Whisper yields superior results with fewer trainable parameters (Table 3: E2 vs E4). Furthermore, when compared to monolingual full fine-tuning, it achieves performance almost on par while requiring training of

Table 3: Comparison of Whisper-small and LoRA-Whisper in WER/CER for base languages.

ID	Model	Finetune	#Train param	PL	PT	IT	ZH	Avg
E1		No	-	10.90	13.82	20.43	9.25	13.60
E2	Whisper-small	Full(multilingual)	240M	8.30	13.34	10.69	14.37	11.68
E3		Full(monolingual)	240M*4	7.65	10.21	10.27	8.64	9.19
E4	LoRA-Whisper	LoRA	13M*4	7.94	10.81	10.49	8.82	9.52

Table 4: Experimental results of language expansion. E1 denotes the results of original Whisper model and E2 denotes the results on base languages before language expansion. Full means multilingual full fine-tune with new language data and full+ means using both new language and base language data. Seed model serves as an initialization for continual training.

ID	Model	Finetune	Seed model	#Train param	New languages					Base languages				
					DA	EL	CY	JA	Avg	PL	PT	IT	ZH	Avg
E1	Whisper-small	No	E1	-	33.97	31.81	58.62	12.04	34.11	10.90	13.82	20.43	9.25	13.60
E2		No	E2	-	-	-	-	-	-	8.30	13.34	10.69	14.37	11.68
E5		Full	E2	240M	38.88	26.08	32.90	15.46	28.33	17.25	25.56	17.41	68.93	32.29
E6		Full+	E2	240M	41.35	28.89	33.65	16.61	30.13	9.77	15.41	12.16	16.07	13.35
E7	LoRA-Whisper	LoRA	E4	13M*4	28.28	22.84	30.63	10.11	22.97	7.94	10.81	10.49	8.82	9.52
E8		Warm start	E4	13M*4	27.45	21.77	28.05	10.03	21.83	7.94	10.81	10.49	8.82	9.52
E9		LoRA MoE	E4	13M*4	27.56	21.56	28.07	10.03	21.81	7.94	10.81	10.49	8.82	9.52

only 5% of parameters (Table 3: E3 vs E4).

Language expansion As can be seen from Table 4, when adding new languages, fine-tuning the model solely with new language data results in serious catastrophic forgetting, with WER on base languages nearly triples compared to the previous result (Table 4: E2 vs E5). A simple way to mitigate catastrophic forgetting is to mix a portion of original data with new data, and then train the model on this combined dataset, as indicated by E6 outlined in Table 4. For each base language, we extract 5 hours of training data and merge them with new language data. It can be seen that this can alleviate the phenomenon of catastrophic forgetting to a certain degree but at the sacrifice of performance on new languages caused by language interference. The LoRA-Whisper we propose can solve this problem in an elegant way. Without affecting the performance of base languages, the LoRA-Whisper model can make use of the similarity between languages to flexibly expand new languages. Our proposed methods (LoRA warm start and LoRA MoE) yield a relative gain of 23% and 5% over full fine-tuning and LoRA fine-tuning respectively, demonstrating the efficacy of our model (Table 4: E8 vs E5 & E7, E9 vs E5 & E7).

4.4. Ablation study

In language expansion, the most similar language to the new language are chosen to assist in training. To validate the effectiveness of this approach, a series of experiments have been conducted on LoRA warm start.

Utilizing Step 1 as illustrated in Figure 2, we can obtain that Danish and Greek are most similar to Portuguese, while Welsh and Japanese are most similar to Polish and Chinese respectively. From Table 5, it is evident that the model attains optimal performance on new languages when the new LoRA matrix is initialized from the LoRA matrix of the most similar language. Another noteworthy observation is that initializing

it with less relevant language’s LoRA matrix may result in decreased model performance compared to training from scratch, underscoring the significance of selecting the most similar language to the new one.

Table 5: Ablation study on LoRA warm start. - means the new LoRA matrix is trained from scratch.

New language	Most similar language	Initialization				
		-	PL	PT	IT	ZH
DA	PT	28.28	30.19	27.45	28.26	28.15
EL	PT	22.84	23.1	21.77	21.81	22.19
CY	PL	30.63	28.05	28.42	28.5	29.34
JA	ZH	10.11	11.69	10.34	10.36	10.03

4.5. Limitation

The limitation of LoRA-Whisper lies in the model size will become larger as the number of languages continues to increase. Hence, future research will explore the sharing of LoRA within multiple similar languages and expand to more languages.

5. Conclusions

In this study, we introduce LoRA-Whisper, a parameter-efficient and extensible multilingual ASR model. By attaching language-specific LoRA modules to the Whisper model, our approach effectively solves the problem of language interference and achieves better language expansion via LoRA warm start or LoRA MoE, allowing people build customized multilingual speech recognition models based on these speech foundation models. We hope that our study can facilitate the research on language expansion in multilingual ASR.

6. Acknowledgements

This work was supported by the National Natural Science Foundation of China (No. 62206171 and No. U23B2018), Shanghai Municipal Science and Technology Major Project under Grant 2021SHZDZX0102 and the International Cooperation Project of PCL and Tencent AI Lab Rhino-Bird Focused Research Program.

7. References

- [1] A. Graves, S. Fernández, F. J. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proc. ICML*, Pittsburgh, 2006.
- [2] A. Graves, A. Mohamed, and G. E. Hinton, "Speech recognition with deep recurrent neural networks," in *Proc. ICASSP*, Vancouver, 2013.
- [3] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *Proc. ICASSP*, Shanghai, 2016.
- [4] S. Kim, T. Hori, and S. Watanabe, "Joint CTC-attention based end-to-end speech recognition using multi-task learning," in *Proc. ICASSP*, New Orleans, 2017.
- [5] J. Li *et al.*, "Recent advances in end-to-end automatic speech recognition," *APSIPA Transactions on Signal and Information Processing*, vol. 11, no. 1, 2022.
- [6] V. Pratap, Q. Xu, A. Sriram, G. Synnaeve *et al.*, "MLS: A large-scale multilingual dataset for speech research," in *Proc. Interspeech*, Shanghai, 2020.
- [7] A. Conneau, M. Ma, S. Khanuja, Y. Zhang *et al.*, "FLEURS: Few-shot learning evaluation of universal representations of speech," in *arXiv preprint arXiv:2205.12446*, 2022.
- [8] C. Wang, M. Riviere, A. Lee, A. Wu *et al.*, "VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation," in *Proc. ACL*, Bangkok, 2021.
- [9] R. Ardila, M. Branson, K. Davis, M. Kohler *et al.*, "Common Voice: A massively-multilingual speech corpus," in *Proc. ACL*, Marseille, 2020.
- [10] A. Radford, J. W. Kim, T. Xu, G. Brockman *et al.*, "Robust speech recognition via large-scale weak supervision," in *Proc. ICML*, Hawaii, 2023.
- [11] Y. Zhang, W. Han, J. Qin, Y. Wang *et al.*, "Google USM: Scaling automatic speech recognition beyond 100 languages," in *arXiv preprint arXiv:2303.01037*, 2023.
- [12] V. Pratap, A. Tjandra, B. Shi, P. Tomasello *et al.*, "Scaling speech technology to 1,000+ languages," in *arXiv preprint arXiv:2305.13516*, 2023.
- [13] B. Li, T. N. Sainath, K. C. Sim, M. Bacchiani *et al.*, "Multi-dialect speech recognition with a single sequence-to-sequence model," in *Proc. ICASSP*, Calgary, 2018.
- [14] A. Waters, N. Gaur, P. Haghani, P. Moreno *et al.*, "Leveraging language ID in multilingual end-to-end speech recognition," in *Proc. ASRU*, Sentosa, 2019.
- [15] L. Zhou, J. Li, E. Sun, and S. Liu, "A configurable multilingual model is all you need to recognize all languages," in *Proc. ICASSP*, Singapore, 2022.
- [16] E. Sun, J. Li, Y. Hu, Y. Zhu *et al.*, "Building high-accuracy multilingual ASR with gated language experts and curriculum training," in *arXiv preprint arXiv:2303.00786*, 2023.
- [17] W. Wang, G. Ma, Y. Li, and B. Du, "Language-routing mixture of experts for multilingual and code-switching speech recognition," in *Proc. Interspeech*, Dublin, 2023.
- [18] N. Gaur, B. Farris, P. Haghani, I. Leal *et al.*, "Mixture of informed experts for multilingual speech recognition," in *Proc. ICASSP*, Toronto, 2021.
- [19] Y. Zhu, P. Haghani, A. Tripathi, B. Ramabhadran *et al.*, "Multilingual speech recognition with self-attention structured parameterization," in *Proc. Interspeech*, Shanghai, 2020.
- [20] S. Li, Y. You, X. Wang, K. Ding *et al.*, "Enhancing multilingual speech recognition through language prompt tuning and frame-level language adapter," in *arXiv preprint arXiv:2309.09443*, 2023.
- [21] G. I. Winata, G. Wang, C. Xiong, and S. Hoi, "Adapt-and-Adjust: Overcoming the long-tail problem of multilingual speech recognition," in *arXiv preprint arXiv:2012.01687*, 2020.
- [22] J. Xie, K. Li, J. Guo, A. Tjandra *et al.*, "Dynamic ASR pathways: An adaptive masking approach towards efficient pruning of a multilingual ASR model," in *arXiv preprint arXiv:2309.13018*, 2024.
- [23] M. Yang, A. Tjandra, C. Liu, D. Zhang *et al.*, "Learning ASR pathways: A sparse multilingual ASR model," in *Proc. ICASSP*, Rhodes, 2023.
- [24] Y. Lu, M. Huang, X. Qu, P. Wei *et al.*, "Language adaptive cross-lingual speech representation learning with sparse sharing sub-networks," in *Proc. ICASSP*, Singapore, 2022.
- [25] A. Kannan, A. Datta, T. N. Sainath, E. Weinstein *et al.*, "Large-scale multilingual speech recognition with a streaming end-to-end model," in *Proc. Interspeech*, Graz, 2019.
- [26] B. Li, R. Pang, Y. Zhang, T. N. Sainath *et al.*, "Massively multilingual ASR: A lifelong learning solution," in *Proc. ICASSP*, Singapore, 2022.
- [27] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan *et al.*, "Continual lifelong learning with neural networks: A review," *Neural Networks*, vol. 113, pp. 54–71, 2019.
- [28] L. D. Libera, P. Mousavi, S. Zaiem, C. Subakan *et al.*, "CL-MASR: A continual learning benchmark for multilingual ASR," in *arXiv preprint arXiv:2310.16931*, 2023.
- [29] D. Rolnick, A. Ahuja, J. Schwarz, T. Lillicrap *et al.*, "Experience replay for continual learning," in *Proc. NeurIPS*, Vancouver, 2019.
- [30] A. Chaudhry, M. Ranzato, M. Rohrbach, and M. Elhoseiny, "Efficient lifelong learning with A-GEM," in *Proc. ICLR*, New Orleans, 2019.
- [31] P. Buzzega, M. Boschini, A. Porrello, D. Abati *et al.*, "Dark experience for general continual learning: A strong, simple baseline," in *Proc. NeurIPS*, 2020.
- [32] A. Mallya, D. Davis, and S. Lazebnik, "Piggyback: Adapting a single network to multiple tasks by learning to mask weights," in *Proc. ECCV*, Munich, 2018.
- [33] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness *et al.*, "Overcoming catastrophic forgetting in neural networks," *Proceedings of the national academy of sciences*, vol. 114, no. 13, pp. 3521–3526, 2017.
- [34] S. Hou, X. Pan, C. C. Loy, Z. Wang *et al.*, "Learning a unified classifier incrementally via rebalancing," in *Proc. CVPR*, Long Beach, 2019.
- [35] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu *et al.*, "LoRA: Low-rank adaptation of large language models," in *arXiv preprint arXiv:2106.09685*, 2021.
- [36] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis *et al.*, "Outrageously large neural networks: The sparsely-gated mixture-of-experts layer," in *arXiv preprint arXiv:1701.06538*, 2017.
- [37] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *arXiv preprint arXiv:1711.05101*, 2019.