



# Language-Aware Prompt Tuning for Parameter-Efficient Seamless Language Expansion in Multilingual ASR

Hongli Yang<sup>1</sup>, Sheng Li<sup>2</sup>, Hao Huang<sup>1,3,4,\*</sup>, Ayiduosu Tuohan<sup>1</sup>, Yizhou Peng<sup>5</sup>

<sup>1</sup>School of Computer Science and Technology, Xinjiang University, China

<sup>2</sup>Institute of Science Tokyo, Japan

<sup>3</sup>Joint International Research Laboratory of Silk Road Multilingual Cognitive Computing, China

<sup>4</sup>Xinjiang Key Laboratory of Multi-lingual Information Technology, China

<sup>5</sup>College of Computing and Data Science, Nanyang Technological University, Singapore

hongli@stu.xju.edu.cn, huanghao@xju.edu.cn

## Abstract

Recent advancements in multilingual automatic speech recognition (ASR) have been driven by large-scale end-to-end models like Whisper. However, challenges such as language interference and expanding to unseen languages (language expansion) without degrading performance persist. This paper addresses these with three contributions: 1) Entire Soft Prompt Tuning (Entire SPT), which applies soft prompts to both the encoder and decoder, enhancing feature extraction and decoding; 2) Language-Aware Prompt Tuning (LAPT), which leverages cross-lingual similarities to encode shared and language-specific features using lightweight prompt matrices; 3) SPT-Whisper, a toolkit that integrates SPT into Whisper and enables efficient continual learning. Experiments across three languages from FLEURS demonstrate that Entire SPT and LAPT outperform Decoder SPT by 5.0% and 16.0% in language expansion tasks, respectively, providing an efficient solution for dynamic, multilingual ASR models with minimal computational overhead.

**Index Terms:** multilingual speech recognition, language expansion, Whisper, soft prompt tuning

## 1. Introduction

Recent years have witnessed significant advancements in Automatic Speech Recognition (ASR) models, especially for high-resource languages, achieving remarkable performance [1–4]. As a crucial element of speech interaction in multilingual AI assistants, multilingual speech recognition has attracted considerable research interest, fueling efforts to broaden its applications [5–7].

The development of large-scale multilingual ASR models, such as Whisper [8], Google USM [9], and MMS [10], has revolutionized the field. These foundational models enable the creation of customized multilingual speech recognition models tailored to specific languages, unlocking new possibilities for cross-lingual communication and application development [11–13]. However, two significant challenges still persist in multilingual ASR: (1) language interference, arising from confusion in all-in-one architectures, and (2) the difficulty of expanding to unseen languages, especially low-resource ones, without degrading the performance of existing ones.

To address language interference, prior works have explored various strategies, such as leveraging language ID infor-

mation or designing language-specific modules [14–16], including dedicated encoders for each language. Other approaches employ pruning strategies to create submodels for individual languages [17, 18] or propose novel sampling methods to mitigate data imbalance. While these methods alleviate language interference to some extent, they often introduce complexities in design and fail to accommodate seamless language expansion.

The integration of new low-resource languages into multilingual ASR models has emerged as a critical research area. Traditional approaches, such as Full Fine-Tuning (FFT), update all parameters of a pre-trained model to adapt to a new task or dataset. While effective, FFT is computationally expensive and prone to catastrophic forgetting (CF) of previously learned languages [19, 20]. To mitigate CF, Continual Learning (CL) methods have been proposed, including prototype-based [21], regularization-based [22], replay-based [23], and optimization-based [24] approaches. In response, Parameter-Efficient Fine-Tuning (PEFT) methods such as Low-Rank Adaptation (LoRA) [25–27] and Soft Prompt Tuning (SPT) [28–30] focus on adjusting only a subset of model parameters or adding new ones, keeping most of the original parameters fixed. This approach is more computationally efficient than full fine-tuning, enabling language adaptation without compromising performance on previously learned languages.

Existing approaches like LoRA-Whisper [26] enhance multilingual ASR by attaching language-specific low-rank matrices, leveraging cross-lingual similarities to improve new language adaptation. Similarly, Soft Language Code Tuning (SLCT) and Soft Prompt Tuning for the Whisper decoder (Decoder SPT) were proposed as methods to add unseen languages to multilingual ASR models [29]. SLCT trains only language-specific embeddings, while Decoder SPT uses soft prompts to guide language recognition. However, these methods do not explicitly use modules to store language-specific information, leading to difficulties in achieving effective language discrimination or providing language bias, especially when parameters are shared in a single trainable model.

To overcome these limitations, we extend Decoder SPT to the entire model for Whisper (Entire SPT) and propose Language-Aware Prompt Tuning (LAPT), a novel approach that leverages Soft Prompt Tuning and cross-lingual similarities for seamless language expansion. By capitalizing on the linguistic relationships between new and base languages, LAPT stores shared information within the Whisper model while capturing language-specific details in dedicated soft prompt matrices. When incorporating a new language, a unique soft prompt matrix is assigned, ensuring minimal impact on the performance of existing languages. Specifically, we introduce two variants

\*Corresponding author: Hao Huang

This work was supported by National Natural Science Foundation of China (62466055) and the National Key R&D Program of China (2020AAA0107902).

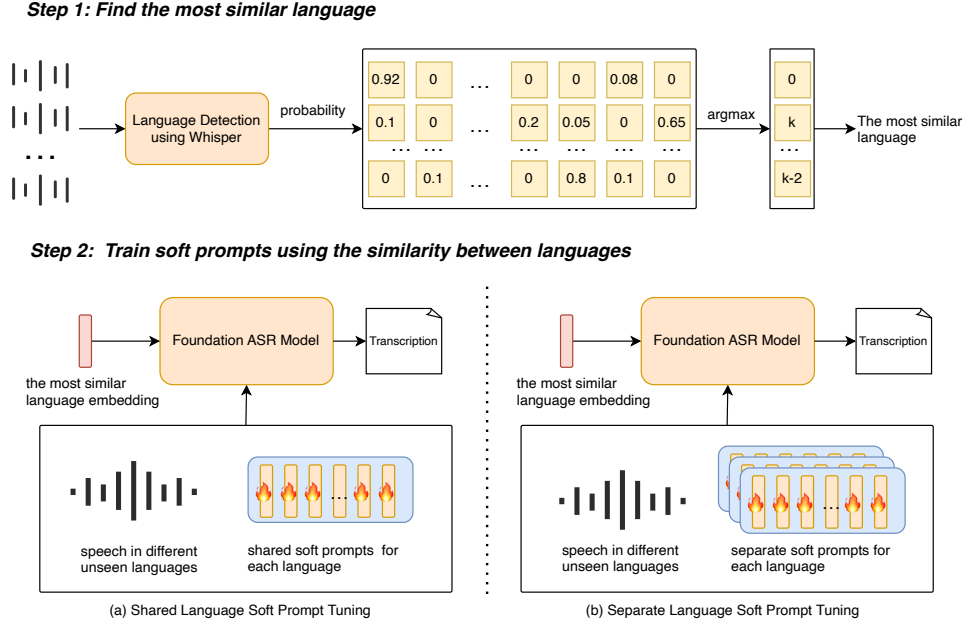


Figure 1: Architecture of LAPT in language expansion. Left: Shared language prompt tuning, Right: Separate language prompt tuning.

of LAPT: shared language prompt tuning and separate language prompt tuning, enabling flexible and efficient language adaptation.

In addition to the methodological contributions, we develop SPT-Whisper<sup>1</sup>, an open-source toolkit that integrates P-Tuning v2 [31], ResMLP [32], LoPT [33], and advanced SPT techniques into Whisper. SPT-Whisper supports continual learning for multilingual ASR, allowing users to customize models for new languages with minimal computational overhead. Although Whisper serves as the exemplar in this work, the foundational model is not limited to Whisper and can be extended to other open-source speech recognition models.

## 2. Methods

### 2.1. Multilingual ASR and Entire SPT

Applying SPT in multilingual ASR expands language support while mitigating catastrophic forgetting (CF). For seen languages, soft prompts are not needed, as the model’s parameters remain fixed during SPT, allowing standard decoding without degrading performance. For new languages, a unique soft prompt matrix is appended to the Whisper model’s input. The original model retains shared language knowledge, while language-specific nuances are encoded within the corresponding embeddings and soft prompts.

Applying SPT to Whisper involves designing separate soft prompt sequences for both the encoder and the decoder. Given a sequence of acoustic features  $X \in \mathbb{R}^{l \times e}$  and a decoder embedding matrix  $E \in \mathbb{R}^{|v| \times e}$ , we determine the optimal soft prompts  $P \in \mathbb{R}^{n \times e}$  for the encoder input and  $P' \in \mathbb{R}^{n \times e}$  for the decoder input. Here,  $l$  is the input speech length,  $|v|$  is the vocabulary size,  $n$  is the length of the soft prompt sequence, and  $e$  is the embedding dimension.

<sup>1</sup><https://github.com/paper-submit-code/SPT-Whisper>

We experimented with different implementations of SPT by applying soft prompts to various components of the Whisper model: specifically, we tested adding soft prompts to only the encoder, only the decoder, or both. For encoder adaptation, the matrix  $[P, X] \in \mathbb{R}^{(n+l) \times e}$  is constructed by concatenating the soft prompt matrix  $P$  with the input features  $X$ . This allows the encoder to attend to both the original acoustic features and the learned prompt information. For decoder adaptation, the soft prompts  $P'$  are inserted into the decoder input at the position of the  $\langle \text{Prev} \rangle$  special token. The decoder’s input sequence includes special tokens such as the start-of-sequence, language ID, and task tokens. The resulting modified input is represented as  $[P', \text{embed}(g)] \in \mathbb{R}^{(n+4) \times e}$ , where  $g$  denotes the sequence of special tokens. These inserted soft prompts provide additional contextual guidance for the decoding process.

### 2.2. Language-Aware Prompt Tuning

A common approach to incorporating language information in multilingual speech recognition models is to represent it as a one-hot vector or embedding, which is then concatenated with the acoustic features to form the input. However, in our proposed approach, we aim to leverage the pre-existing language information in a more efficient and adaptable manner. Consequently, we introduce two effective methods for language expansion, namely Shared Language Prompt Tuning and Separate Language Prompt Tuning as depicted in Fig 1. These methods involve two steps.

**Step 1: Find the most similar language** To integrate a new language into an existing model, the first step is to find the language most similar to the new one from a predefined set of base languages. We begin by randomly selecting  $M$  audio segments from the new language’s dataset. These segments are then processed using the Whisper model to perform language identification, generating a probability distribution across all languages.

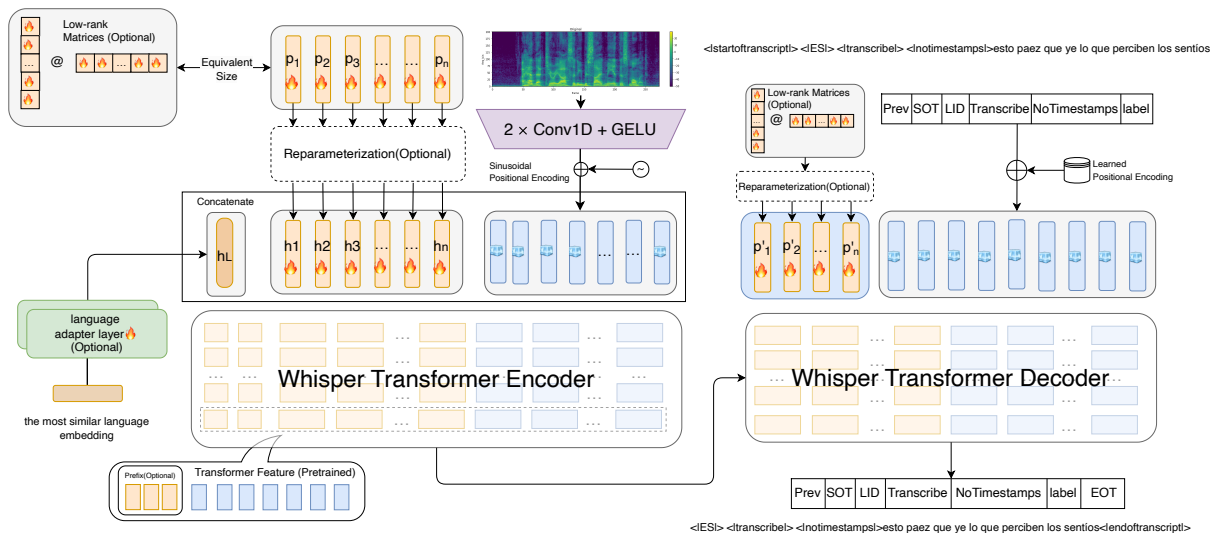


Figure 2: Overview of the SPT-Whisper architecture proposed for parameter-efficient continual learning.

The relevant probabilities are extracted and normalized, yielding the vector  $\mathbf{p}_i = [p_{i1}, p_{i2}, \dots, p_{in}]$ ;  $i = 1, \dots, M$ , where  $p_{ij}$  denotes the probability that the  $i$ -th audio segment belongs to the  $j$ -th base language.

To determine the most similar language, we calculate the similarity between the new language and each base language using the following metric:

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

where  $\mathbb{I}$  is the indicator function and  $\text{sim}_k$  represents the similarity between the new language and the  $k$ -th base language.

**Step 2: Continual training on new languages** Once we have identified the most similar languages, the next step involves training soft prompts using the language similarities. As described in Fig 1, building on Whisper’s pre-existing capabilities in language identification and previous research on language token understanding, we attempt to introduce Whisper’s pre-trained language embeddings as language prompts into SPT-based end-to-end multilingual speech recognition. In addition, we use a prompt encoder to encode prompt embedding with the same dimension as the acoustic features. The prompt embedding serves as a language prompt, guiding the model to accurately recognize specific languages with other soft prompts.

In Shared language prompt tuning, shared soft prompts are shared across each language to encourage efficient adaptation and reduce the parameter overhead. In Separate Language Prompt Tuning, separate soft prompts is trained for each new language, providing more granular control over the adaptation process. When the input is a piece of speech of  $k$ -th language, it will activate  $k$ -th soft prompt and pass through Whisper in the forward pass. Furthermore, language prompts with similarities across languages can be further used to facilitate more effective training for the new languages, thus improving both the accuracy of language identification and the performance of the model on new low-resource languages.

Table 1: FLEURS datasets used in training.

Language	Train		Dev		Test	
	#	hrs	#	hrs	#	hrs
Asturian	2511	7.53	398	0.99	946	2.44
Sorani Kurd	3,040	10.46	386	1.23	922	2.99
Kabuverdianu	2,715	10.51	366	1.32	864	3.27

### 2.3. SPT-Whisper

We implement advanced SPT methods, including LAPT, and provide open-source implementations of all these techniques on SPT-Whisper, as shown in Fig. 2. We aim to provide a fair and practical comparison to advance multilingual ASR.

## 3. Experiments

### 3.1. Dataset and Evaluation Metric

To ensure the reproducibility of our proposed methods, we conducted experiments on FLEURS benchmark multilingual open-source datasets, which comprise 102 languages, of which 19 languages are not supported by the Whisper model. Each language has about 12 hours of read speech data, with each audio within 30 seconds. Due to limited resources, initial tests were performed on three languages randomly selected from the 19 unsupported ones: Asturian, Sorani Kurdish, and Kabuverdianu. We adopt the Character Error Rates (CER) for evaluating ASR performance. Detailed information of the datasets of these three languages are presented in Table 1.

### 3.2. Implementation Details

Considering inference speed and limited computational resources, we evaluate our proposed method on not-large foundation models, such as Whisper-small and Whisper-medium. FFT employs an initial learning rate of  $1e-6$ , which decays linearly, while LoRA and SPT-based methods use initial rates of  $1e-3$ . LoRA was applied to the attention layers with a rank of 8, and varying rank sizes showed negligible differences in re-

Table 2: *Soft Prompt Tuning Configuration Analysis on Whisper-small: Prompt Length and Position (CER %)*

Prompt Length	Encoder SPT	Decoder SPT	Entire SPT
16 Prompt Tokens	12.76	11.61	11.77
32 Prompt Tokens	13.27	11.48	11.61
64 Prompt Tokens	13.04	11.41	10.82
128 Prompt Tokens	12.33	11.91	10.31
256 Prompt Tokens	12.77	N/A <sup>♣</sup>	10.57

Note: ♣ indicates the decoder does not support 256-token length, due to the decoder’s context length limitation. *Entire SPT* with Prompt Length of 256 uses prompts with 256 tokens for the encoder and 128 for the decoder.

sults. We trained the models over 20 epochs with a batch size of 8. The training was conducted on a single NVIDIA RTX3090 GPU with 24GB of VRAM. For decoding, we selected the best model from all epochs based on the dev set and performed a greedy search.

### 3.3. Prompt Length and Position Selection

The performance of SPT exhibits significant length dependency, as demonstrated by our systematic evaluation on Whisper-small for Asturian recognition. We evaluated SPT with prompt lengths {16, 32, 64, 128, 256}, applied to different components of the model: the encoder (Encoder SPT), decoder (Decoder SPT), and the entire model (Entire SPT). As Table 2 reveals, two critical patterns emerge: (1) Length Threshold: All configurations achieve optimal performance at 128 tokens. (2) Architecture Sensitivity: Decoder SPT show stronger baseline performance, but Entire SPT integration enables greater error reduction. Therefore, we used a prompt length of 128 and applied Entire SPT in subsequent experiments.

Notably, the best performance is achieved when the soft prompt is applied to the entire model, emphasizing the importance of joint optimization between the encoder and decoder for both acoustic feature extraction and language modeling. This finding is significant because, in traditional ASR models, the encoder typically focuses on extracting speech features from the audio signal, while the decoder uses a language model to convert those features into text. Extending the Decoder SPT to the entire model enables better communication between the two, resulting in improved performance in both tasks.

### 3.4. Main Results

Table 3 presents a comparison of various fine-tuning approaches across three target languages: Asturian (ast), Sorani Kurdish (ckb), and Kabuverdianu (kea), using both Whisper-small and Whisper-medium models. The baseline performance is obtained on different Whisper model versions (B1-B3), while FFT, LoRA, and Entire SPT are fine-tuned on both small and medium models for each language using the respective FLEURS training sets. When applying different tuning methods on Whisper-small, FFT provides the best performance, with an average CER of 8.87% on Whisper-small (S1), but at a high computational cost. LoRA (S2) offers a more efficient alternative, yielding a CER of 12.93% on average, but it still underperforms compared to FFT. Entire SPT provides a balanced approach, with both Shared Entire SPT (S3) and Separate Entire SPT (S5) reducing the average CER to 12.43% and 12.54%, respectively, outperforming LoRA. However, the limitation of the Separate SPT approach is that the model size increases as more languages are added, whereas the Shared SPT method addresses

Table 3: *ASR performance on FLEURS Asturian(ast), Sorani Kurdish(ckb) and Kabuverdianu(kea) test sets with different tuning approaches. (CER %)*

ID	Method	Languages				#Params(M)
		ast	ckb	kea	avg	
B1	Whisper Small	14.94	57.14	36.23	36.10	/
B2	Whisper Medium	15.01	65.36	36.20	38.86	/
B3	Whisper Large-v3	14.58	40.05	36.67	30.43	/
Whisper-small Tuning Results						
S1	FFT	7.14	13.16	6.32	8.87	240.58M
S2	LoRA	10.10	18.94	9.75	12.93	1.85M×3
S3	Shared Entire SPT	10.12	17.37	9.81	12.43	0.17M
S4	+ LAPT	<b>9.99</b>	17.35	9.68	12.34	0.96M
S5	Separate Entire SPT	10.31	17.53	9.79	12.54	0.17M×3
S6	+ LAPT	10.01	<b>16.70</b>	<b>9.51</b>	<b>12.07</b>	0.96M×3
Whisper-medium Tuning Results						
M1	FFT	5.94	11.73	5.03	7.57	762.32M
M2	LoRA	7.52	13.76	7.06	9.45	4.94M×3
M3	Shared Entire SPT	7.47	13.34	6.72	9.18	0.44M
M4	+ LAPT	7.45	13.25	6.70	9.13	1.49M
M5	Separate Entire SPT	7.47	13.30	6.76	9.18	0.44M×3
M6	+ LAPT	<b>7.20</b>	<b>13.21</b>	<b>6.67</b>	<b>9.03</b>	1.49M×3

Note: #Params indicates the number of trainable parameters. “×3” denotes the use of three languages with distinct soft prompt groups.

this by sharing soft prompts across multiple similar languages.

Integrating LAPT with Entire SPT further improves performance. The combination of Shared Entire SPT and LAPT (S4) results in a minor improvement, reducing the average CER from 12.43% to 12.34%. However, the biggest improvement comes from the Separate Entire SPT with LAPT (S6), which reduces the average CER to 12.07%. This method achieves the lowest CER for Sorani Kurdish (16.70%), highlighting LAPT’s ability to enhance language-specific recognition.

In Whisper-medium, similar trends are observed, with FFT (M1) delivering the best performance, though at a higher computational cost. The Shared Entire SPT (M3) provides more balanced results compared to the Separate Entire SPT (M5), particularly in terms of parameter efficiency, despite having slightly lower performance. The Separate Entire SPT with LAPT configuration (M6) is expected to outperform LoRA (M2) and Shared Entire SPT with LAPT (M4), making it a more efficient approach for multilingual ASR models.

Overall, experimental results demonstrate that Entire SPT significantly improves the performance of multilingual ASR by applying soft prompts across the entire model. Additionally, the combination of LAPT with Entire SPT provides a scalable solution for expanding ASR models by leveraging language-specific features while minimizing interference during language expansion, all without the computational overhead of FFT.

## 4. Conclusion

We propose a novel approach to address language interference and enable efficient language expansion in multilingual ASR. Entire SPT applies soft prompts to both the encoder and decoder to enhance language recognition, while combining Entire SPT with LAPT leverages cross-lingual similarities for seamless expansion. Additionally, we introduce SPT-Whisper for parameter-efficient continual learning in multilingual ASR, providing a scalable solution for expanding models without compromising performance.

## 5. References

- [1] A. Graves, S. Fernández, F. J. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *ICML*, ser. ACM International Conference Proceeding Series, vol. 148. ACM, 2006, pp. 369–376.
- [2] W. Chan, N. Jaitly, Q. V. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *ICASSP*. IEEE, 2016, pp. 4960–4964.
- [3] S. Kim, T. Hori, and S. Watanabe, "Joint ctc-attention based end-to-end speech recognition using multi-task learning," in *ICASSP*. IEEE, 2017, pp. 4835–4839.
- [4] 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.
- [5] 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," in *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023.
- [6] C. Tang, W. Yu, G. Sun, X. Chen, T. Tan, W. Li, L. Lu, Z. Ma, and C. Zhang, "SALMONN: towards generic hearing abilities for large language models," in *ICLR*. OpenReview.net, 2024.
- [7] Q. Fang, S. Guo, Y. Zhou, Z. Ma, S. Zhang, and Y. Feng, "LLaMA-omni: Seamless speech interaction with large language models," in *The Thirteenth International Conference on Learning Representations*, 2025.
- [8] 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, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, ser. Proceedings of Machine Learning Research, vol. 202. PMLR, 2023, pp. 28 492–28 518.
- [9] Y. Zhang, W. Han, J. Qin, Y. Wang, A. Bapna, Z. Chen, N. Chen, B. Li, V. Axelrod, G. Wang *et al.*, "Google usm: Scaling automatic speech recognition beyond 100 languages," *arXiv preprint arXiv:2303.01037*, 2023.
- [10] V. Pratap, A. Tjandra, B. Shi, P. Tomasello, A. Babu, S. Kundu, A. Elkahky, Z. Ni, A. Vyas, M. Fazel-Zarandi *et al.*, "Scaling speech technology to 1,000+ languages," *Journal of Machine Learning Research*, vol. 25, no. 97, pp. 1–52, 2024.
- [11] Y. Yang, Y. Peng, H. Huang, E. S. Chng, and X. Zhong, "Adapting openai's whisper for speech recognition on code-switch mandarin-english seame and asru2019 datasets," in *2024 Asia Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2024, pp. 1–6.
- [12] B. Talafha, A. Waheed, and M. Abdul-Mageed, "N-Shot Benchmarking of Whisper on Diverse Arabic Speech Recognition," in *24th Annual Conference of the International Speech Communication Association, Interspeech 2023, Dublin, Ireland, August 20-24, 2023*. ISCA, 2023, pp. 5092–5096.
- [13] Y. Liu, X. Yang, and D. Qu, "Exploration of whisper fine-tuning strategies for low-resource asr," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2024, no. 1, p. 29, 2024.
- [14] S. Li, Y. You, X. Wang, K. Ding, and G. Wan, "Enhancing multilingual speech recognition through language prompt tuning and frame-level language adapter," in *ICASSP*. IEEE, 2024, pp. 10 941–10 945.
- [15] E. Sun, J. Li, Y. Hu, Y. Zhu, L. Zhou, J. Xue, P. Wang, L. Liu, S. Liu, E. Lin *et al.*, "Building high-accuracy multilingual asr with gated language experts and curriculum training," in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–7.
- [16] W. Wang, G. Ma, Y. Li, and B. Du, "Language-routing mixture of experts for multilingual and code-switching speech recognition," in *INTER\_SPEECH*. ISCA, 2023, pp. 1389–1393.
- [17] J. Xie, K. Li, J. Guo, A. Tjandra, Y. Shangguan, L. Sari, C. Wu, J. Jia, J. Mahadeokar, and O. Kalinli, "Dynamic ASR pathways: An adaptive masking approach towards efficient pruning of a multilingual ASR model," in *ICASSP*. IEEE, 2024, pp. 12 201–12 205.
- [18] Y. Lu, M. Huang, X. Qu, P. Wei, and Z. Ma, "Language adaptive cross-lingual speech representation learning with sparse sharing sub-networks," in *ICASSP*. IEEE, 2022, pp. 6882–6886.
- [19] R. Jain, A. Barcovschi, M. Y. Yiwere, P. Corcoran, and H. Cucu, "Adaptation of Whisper models to child speech recognition," in *24th Annual Conference of the International Speech Communication Association, Interspeech 2023, Dublin, Ireland, August 20-24, 2023*. ISCA, 2023, pp. 5242–5246.
- [20] H. Yang, M. Zhang, S. Tao, M. Ma, and Y. Qin, "Chinese ASR and NER Improvement Based on Whisper Fine-tuning," in *25th International Conference on Advanced Communication Technology, ICACT 2023, Pyeongchang, Korea, Republic of, February 19-22, 2023*. IEEE, 2023, pp. 213–217.
- [21] U. Michieli, P. P. Parada, and M. Ozay, "Online continual learning in keyword spotting for low-resource devices via pooling high-order temporal statistics," *arXiv preprint arXiv:2307.12660*, 2023.
- [22] S. Vander Eeckt *et al.*, "Rehearsal-free online continual learning for automatic speech recognition," *Proceedings Interspeech 2023*, pp. 944–948, 2023.
- [23] U. Cappellazzo, D. Falavigna, and A. Brutti, "An investigation of the combination of rehearsal and knowledge distillation in continual learning for spoken language understanding," in *Proceedings of INTERSPEECH 2023*, 2023, pp. 735–739.
- [24] A. Chaudhry, M. Ranzato, M. Rohrbach, and M. Elhoseiny, "Efficient lifelong learning with a-gem," in *International Conference on Learning Representations*.
- [25] E. J. Hu, yelong shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-rank Adaptation of Large Language Models," in *International Conference on Learning Representations*, 2022.
- [26] Z. Song, J. Zhuo, Y. Yang, Z. Ma, S. Zhang, and X. Chen, "Lora-whisper: Parameter-efficient and extensible multilingual asr," in *Interspeech 2024*, 2024, pp. 3934–3938.
- [27] T. Xu, K. Huang, P. Guo, Y. Zhou, L. Huang, H. Xue, and L. Xie, "Towards rehearsal-free multilingual asr: A lora-based case study on whisper," in *Interspeech 2024*, 2024, pp. 2534–2538.
- [28] B. Lester, R. Al-Rfou, and N. Constant, "The Power of Scale for Parameter-efficient Prompt Tuning," in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*. Association for Computational Linguistics, 2021, pp. 3045–3059.
- [29] M. Qian, S. Tang, R. Ma, K. M. Knill, and M. J. Gales, "Learn and don't forget: Adding a new language to asr foundation models," in *Interspeech 2024*, 2024, pp. 2544–2548.
- [30] H. Ma, Z. Peng, M. Shao, J. Li, and J. Liu, "Extending whisper with prompt tuning to target-speaker asr," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 12 516–12 520.
- [31] X. Liu, K. Ji, Y. Fu, Z. Du, Z. Yang, and J. Tang, "P-Tuning v2: Prompt Tuning Can Be Comparable to Fine-tuning Universally Across Scales and Tasks," *CoRR*, vol. abs/2110.07602, 2021.
- [32] A. Razdaibiedina, Y. Mao, M. Khabsa, M. Lewis, R. Hou, J. Ba, and A. Almahairi, "Residual Prompt Tuning: improving prompt tuning with residual reparameterization," in *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*. Association for Computational Linguistics, 2023, pp. 6740–6757.
- [33] S. Guo, S. Damani, and K.-h. Chang, "Lopt: Low-rank prompt tuning for parameter efficient language models," *arXiv preprint arXiv:2406.19486*, 2024.