



Adapting Whisper for Parameter-efficient Code-Switching Speech Recognition via Soft Prompt Tuning

Hongli Yang¹, Yizhou Peng², Hao Huang^{1,3,4,*}, Sheng Li⁵

¹School of Computer Science and Technology, Xinjiang University, China

²College of Computing and Data Science, Nanyang Technological University, Singapore

³Joint International Research Laboratory of Silk Road Multilingual Cognitive Computing, China

⁴Xinjiang Key Laboratory of Multi-lingual Information Technology, China

⁵Institute of Science Tokyo, Japan

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

Abstract

Large-scale multilingual ASR models like Whisper excel in high-resource settings but face challenges in low-resource scenarios, such as rare languages and code-switching (CS), due to computational costs and catastrophic forgetting. We explore Soft Prompt Tuning (SPT), a parameter-efficient method to enhance CS ASR while preserving prior knowledge. We evaluate two strategies: 1) full fine-tuning (FFT) of both soft prompts and the entire Whisper model, demonstrating improved cross-lingual capabilities compared to traditional methods, and 2) adhering to SPT’s original design by freezing model parameters and only training soft prompts. Additionally, we introduce SPT4ASR, a combination of different SPT variants. Experiments on the SEAME and ASRU2019 datasets show that deep prompt tuning is the most effective SPT approach, and our SPT4ASR methods achieve further error reductions in CS ASR, maintaining parameter efficiency similar to LoRA, without degrading performance on existing languages.

Index Terms: Code-switching speech recognition, Soft prompt tuning, Parameter-efficient, Whisper

1. Introduction

Foundation models (FMs) such as LLaMA [1], CLIP [2], and Whisper [3] have demonstrated strong cross-domain capabilities through large-scale pretraining. Among them, Whisper excels in multilingual automatic speech recognition (ASR), benefiting from massive speech-text data and , showcasing the growing capabilities of publicly accessible speech models in handling complex multilingual tasks.

However, a significant challenge that remains in FMs is how to efficiently adapt to low-resource scenarios without compromising the performance of existing ones [4]. To tackle this issue, previous works in the speech signal processing field have utilized transfer learning techniques, including prompt engineering [5–7], full fine-tuning (FFT) [8–10], and parameter-efficient tuning (PEFT) methods such as Low-Rank Adaptation (LoRA) [11–13]. These approaches leverage the pre-trained weights and knowledge of these FMs, adapting them to new low-resource language environments.

Soft Prompt Tuning (SPT) is proposed among PEFT methods for efficiently optimizing a small set of task-specific prompt parameters [14–18], balancing simplicity, adaptability, and resource conservation similar to LoRA. During fine-tuning with




Figure 1: Comparison of (a) Prompt Engineering, (b) FFT, and (c) SPT. The fire and ice emojis in the figure represent trainable and frozen parameters, respectively.

SPT, the parameters of the original model are treated as constants and remain frozen, while the soft prompts are treated as variables and updated during training. Thus, the training process focuses solely on optimizing the soft prompts, enabling efficient and incremental model adaptation. Figure 1 illustrates the key differences between prompt engineering, FFT, and SPT. Recent studies have explored using SPT for cross-lingual transfer, leveraging multilingual language models to transfer knowledge from high-resource to low-resource languages across various NLP tasks [19]. Moreover, some findings suggest that combining soft prompt tuning with partial or full model fine-tuning can further improve performance on downstream tasks [20, 21].

As a low-resource scenario, code-switching (CS) ASR refers to language switching within spontaneous multilingual recordings, including both intra-sentence and inter-sentence cases. Current approaches to CS ASR include joint optimization methods [22–24], conditional factorization methods [25–27], and mixture-of-experts method [28]. However, the recent advancements in Large Language Models (LLMs) offer a new approach to these challenges, potentially outperforming task-specific architectures through their inherent understanding of multiple languages and acoustic patterns.

In this paper, we introduce SPT for cross-lingual ASR transfer learning and validate its effectiveness through CS experiments. Our main contributions include: 1) evaluating various SPT-based approaches, including the combination of different variants called SPT4ASR, to enhance CS ASR capability; 2) introducing SPT-based methods that achieves parameter efficiency while mitigating catastrophic forgetting; and 3) to the best of our knowledge, we are the first to conduct a comprehensive comparison of various SPT-based methods for mitigating catastrophic forgetting in multilingual Whisper models.

*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).

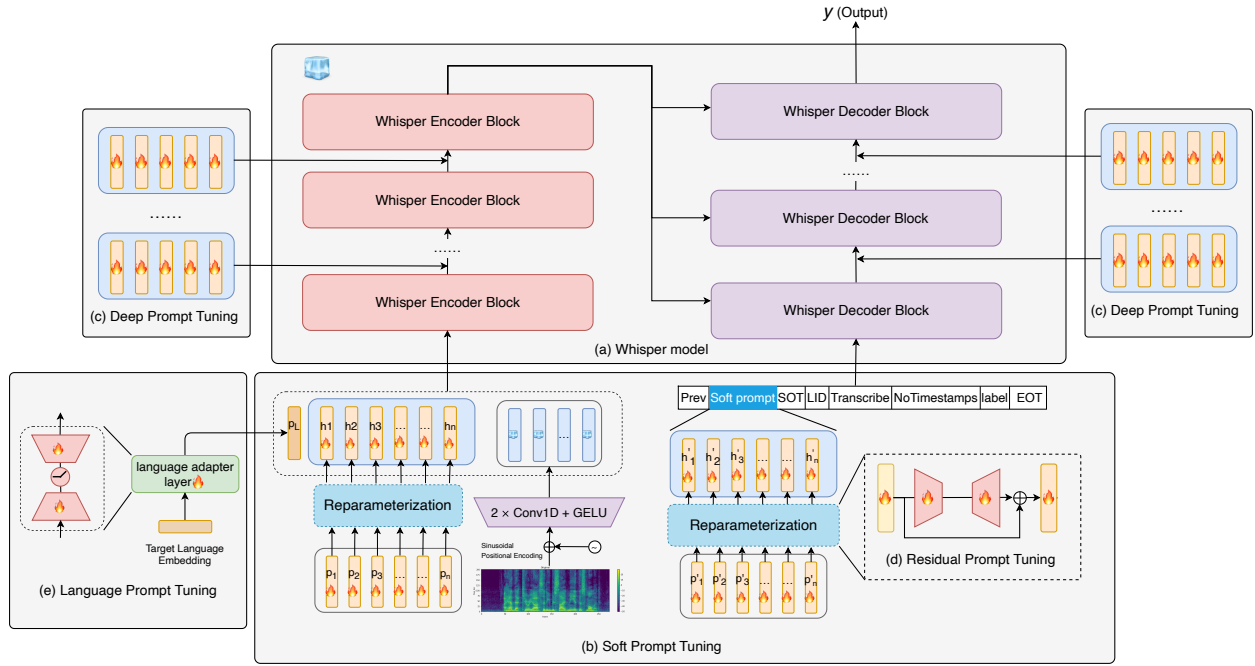


Figure 2: Overview of Whisper model and various Prompt Tuning methods: (a) Whisper is a Transformer-based encoder-decoder model for multilingual ASR and translation tasks, where the entire model is frozen in SPT-based experiments. (b) Vanilla Soft Prompt Tuning involves adding prefix embeddings to the model input features. (c) Deep Prompt Tuning (DPT) extends this by adding prefix embeddings to the output of each block within the Whisper model. (d) Residual Prompt Tuning (ResPT) leverages an autoencoder with reparameterization through residual connections to enhance the model. (e) Language Prompt Tuning (LPT) encodes discrete Language IDs (LIDs) into prompt embeddings to improve language discrimination. (f) SPT4ASR is a combination of the aforementioned methods, specifically designed to improve cross-lingual transfer and language expansion in ASR models.

2. Methods

2.1. Background of Whisper

Fig. 2(a) shows that the Whisper model uses a Transformer-based encoder-decoder architecture. It excels in multi-lingual speech recognition and translation, handling multiple tasks simultaneously through special task tokens as decoder input prompts. For English ASR without predicting timestamps, the decoder input sequence g is structured as:

$$\langle \text{Prev} \rangle, \langle | \text{SOT} | \rangle, \langle | \text{EN} | \rangle, \langle | \text{Transcribe} | \rangle, \langle | \text{NoTimestamps} | \rangle,$$

where $\langle \text{Prev} \rangle$ optionally represents previous text tokens for context during decoding, and succeeding $\langle | \cdot | \rangle$ tokens denote for Start of Transcript, Language ID, Transcription task, and No Timestamps specification.

2.2. FFT

FFT is the mainstream method for adapting Whisper to new languages [8–10]. In this approach, all parameters of Whisper are optimized using the training data from the new language to enhance its performance in that specific language. Given a pre-trained Whisper model with parameters θ and a new dataset Y , the FFT objective is:

$$\max_{\theta} \sum_{(X,y) \in Y} \log P_{\theta}(y | X, \theta), \quad (1)$$

where (X, y) are speech-label pairs from the dataset Y .

2.3. LoRA

LoRA is a prominent PEFT technique that efficiently adapts models to new tasks or languages [12, 13]. It adjusts only a subset of weights by adding low-rank decomposition matrices to the model, while keeping the original parameters θ frozen. This approach significantly reduces computational costs and memory requirements compared to FFT, while maintaining comparable effectiveness. The optimization objective for LoRA can be formulated as:

$$\max_{\theta'} \sum_{(X,y) \in Y} \log P_{\theta}(y | X, \theta, \theta'), \quad (2)$$

where θ' represents the subset of trainable or newly added parameters.

2.4. SPT

SPT adapts Whisper by introducing learnable embeddings P , namely soft prompts, which are concatenated to the model's inputs while freezing Whisper model parameters θ . Specifically, SPT optimizes a trainable embedding matrix $P \in \mathbb{R}^{n \times e}$, where n represents the number of soft prompts and e is the dimension of each embedding. The optimization problem is formulated as follows:

$$\max_P \sum_{(X,y) \in Y} \log P_{\theta}(y | [P, X], \theta). \quad (3)$$

Applying SPT to Whisper involves designing separate soft prompt sequences for both the encoder and decoder. Fig. 2 (b)

illustrates this process. 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, and n is the length of soft prompts.

We experimented with different SPT implementations by adding soft prompts to the encoder, decoder, or both. For encoder adaptation, the model concatenates $[P, X] \in \mathbb{R}^{(n+l) \times e}$, where P is the soft prompt matrix of length n , X is the input features of length l , and e is the embedding dimension. This allows the encoder to process both the original input features and the soft prompts together. For decoder adaptation, soft prompts $P' \in \mathbb{R}^{n \times e}$ are inserted into the $\langle \text{Prev} \rangle$ token of the decoder input sequence g , which includes other special tokens such as start-of-sequence, language ID, and task tokens. The resulting input is $[P', \text{embed}(g)] \in \mathbb{R}^{(n+4) \times e}$, where $n+4$ represents the total number of tokens in the modified sequence. This modified input provides additional context to the decoder, guiding the decoding process effectively.

2.4.1. Deep Prompt Tuning

In the above vanilla SPT approach, where soft prompts are only concatenated with input features before entering the encoder and decoder, we note some limitations: the prompt length is constrained by computational costs, limiting instruction capacity, and the static prompts are too distant from the output layers to directly influence final predictions.

Fig. 2 (c) shows that we introduce Deep Prompt Tuning (DPT) [29], a novel approach that inserts soft prompts at various intermediate blocks throughout the Whisper model. By strategically placing prompts across different levels of the network, we enable a more direct and nuanced influence on the model’s predictions without altering the original parameters.

2.4.2. Residual Prompt Tuning

To overcome limitations of traditional SPT, such as slow convergence and reliance on pre-training prompts, we introduce Residual Prompt Tuning (ResPT) [30], inspired by reparameterization. As illustrated in Fig. 2 (d), ResPT employs a shared multi-layer perceptron (sharedMLP) for reparameterization, enabling efficient and consistent prompt generation. The sharedMLP reparameterizes all prompts, improving training efficiency and robustness through parameter sharing.

ResPT’s key advantage is refining soft prompts quickly without requiring extensive pre-training. The use of residual connections allows the model to build on previous prompt information, while the sharedMLP ensures stable learning. Mathematically, the reparameterization is given by:

$$\mathbf{P}' = \text{MLP}(\mathbf{P}) + \mathbf{P}, \quad (4)$$

where \mathbf{P} represents the input soft prompts, $\text{MLP}(\mathbf{P})$ is the output of the shared MLP, and \mathbf{P}' is the refined soft prompt.

2.4.3. Language Prompt Tuning

Building on Whisper’s language identification capabilities and prior work on language token understanding [6–8], we introduce Language Prompt Tuning (LPT) [31] to enhance code-switching ASR performance. Our approach combines this with a Language Encoder, as an adapter module designed to improve language differentiation in code-switching environments. The Language Encoder is a lightweight, task-specific component

that integrates context-specific cues for each language, allowing the model to dynamically adapt to language transitions more efficiently and accurately.

In our implementation of CS ASR, we integrate two LIDs to fully exploit Whisper’s capability to interpret language tokens. Specifically, we utilize Whisper’s pre-trained language embeddings as prompts, guiding the model to more accurately identify distinct languages in mixed-language speech contexts. As shown in Fig. 2 (e), these language-specific embeddings are concatenated with acoustic features before being input into the Whisper model. This method not only improves language differentiation in new CS scenarios but also maintains overall ASR performance across multilingual environments.

Table 1: *Corpus Details with SEAME and ASRU2019*

Corpus	Subset	Duration(Hrs)	Man/Eng
SEAME	Train	93.6	0.83/0.17
	DevMan	7.5	0.91/0.09
	DevSge	3.9	0.55/0.45
ASRU2019	Train	193.0	0.97/0.03
	Test	20.6	0.98/0.02

3. Experiments

3.1. Dataset and Evaluation Metric

Following the approach used in Whisper’s FFT [8], we optimized soft prompts using the SEAME corpus, a conversational Mandarin-English dataset of approximately 100 hours from Southeast Asia, i.e., Malaysia and Singapore, and the ASRU2019 dataset, a 200-hour Mandarin-English CS dataset from Mainland China, released by Datatang for a Mandarin-English ASR challenge. Both the SEAME and ASRU2019 datasets consist of intra-sentential code-switching speech data and detailed information of the datasets of these three languages are presented in Table 1. We adopt the Mix Error Rate (MER) for evaluating CS ASR performance, calculating errors at the character level for Mandarin and at the word level for English. To assess alleviating catastrophic forgetting, we used the Aishell1 dataset for Chinese ASR, the LibriSpeech test set for English ASR, and the Korean language evaluation data from the SPREDS-U1 dataset for Korean ASR.

3.2. Implementation Details

We conducted experiments on Whisper’s small and medium models, training for 10 epochs with a batch size of 8. The AdamW optimizer was used, starting with a learning rate of $1e-3$ for LoRA and SPT, and $1e-6$ for FFT. The training was conducted on a single NVIDIA RTX3090 GPU with 24GB VRAM. For decoding, we selected the best model from all epochs and performed greedy search.

4. Results

4.1. Prompt Length and Position Selection

The performance of SPT is highly sensitive to prompt length. We experimented with Vanilla SPT tuning using lengths $\{16, 32, 64, 128, 256\}$ on the SEAME dataset with the Whisper small model. As shown in Table 2, performance improves with increasing length up to 128, the maximum length the decoder can handle due to context length limitations. Notably, tuning

Table 2: Speech recognition results on SEAME dataset with different prompt lengths and positions (MER %)

Prompt Positions	Prompt Length of 16		Prompt Length of 32		Prompt Length of 64		Prompt Length of 128		Prompt Length of 256	
	devman	devsge	devman	devsge	devman	devsge	devman	devsge	devman	devsge
Vanilla Soft Prompt (Encoder)	27.59	32.11	27.64	35.49	32.12	37.29	24.62	34.81	24.98	34.71
Vanilla Soft Prompt (Decoder)	32.15	45.27	30.23	42.86	27.95	40.56	25.44	39.21	N/A [♣]	N/A [♣]
Vanilla Soft Prompt (Entire)	31.25	36.75	26.39	34.47	23.21	28.92	21.95	27.39	22.31	30.81

Note: ♣ indicates the decoder doesn't support 256 token length due to context limitations. The *Entire* Position uses 256 tokens for the encoder and 128 for the decoder.

Table 3: ASR performance comparison results of SPT and other methods on SEAME and ASRU2019 datasets (MER %)

ID	Method	SEAME		ASRU	#Params (M)
		devman	devsge	test	
B1	Whisper Small	95.59	75.66	29.74	/
B2	Whisper Medium	76.49	65.36	23.21	/
B3	Whisper Large-v3	61.18	60.14	18.41	/
B4	ESPnet Conformer	16.60	23.30	13.20	47.27M
Whisper-small Tuning Results					
S1	FFT	13.37	19.42	12.92	240.58M
S2	LoRA	13.96	20.59	13.00	1.85M
S3	Vanilla SPT	21.95	27.39	15.60	0.20M
S4	+ DPT	16.57	23.99	13.52	1.39M
S5	+ ResPT	20.23	26.43	15.06	0.79M
S6	+ LPT	18.24	26.68	14.68	0.99M
S7	SPT4ASR	15.48	21.98	13.12	3.74M
S8	Whole model SPT	13.07	18.65	12.84	240.69M
Whisper-medium Tuning Results					
M1	FFT	11.44	17.55	10.18	762.32M
M2	LoRA	11.68	16.99	10.35	4.94M
M3	SPT4ASR	13.12	18.10	11.30	7.48M
M4	Whole model SPT	11.17	16.61	10.18	762.59M

Note: #Params refers to the number of trainable parameters.

the entire model yields the best results. Therefore, in subsequent experiments, we used a prompt length of 128 and applied tuning to the entire model.

4.2. Main Results

Table 3 presents a comparison of tuning methods, including SPT variants, FFT, and LoRA, on the SEAME and ASRU2019 datasets for Whisper-small and Whisper-medium models, along with zero-shot results from three different Whisper model versions (B1-B3) and the ESPnet Conformer results (B4).

4.2.1. Comparison of different SPT variants

As shown in Table 3, among the SPT variants for Whisper-small, DPT (S4) achieves the most significant improvement, significantly outperforming the zero-shot performance of baseline Whisper models (B1-B3), and performing similarly to the ESPnet Conformer results (B4). This enhancement can be attributed to its ability to introduce learnable parameters at multiple layers of the model, thus enabling more granular adaptation across the network hierarchy. ResPT (S5) and LPT (S6) also demonstrate substantial enhancements with fewer parameters. Notably, SPT4ASR, a combination of multiple SPT variants (S7), achieves the best performance among SPT-based methods, with only about 3.74M parameters, highlighting the potential and effectiveness of integrating different SPT approaches.

4.2.2. Comparison of SPT with FFT and LoRA

As shown in Table 3, across both Whisper-small and Whisper-medium models, the SPT4ASR methods (S7 and M3) perform

Table 4: SPT results on AISHELL-1 (CER%), LibriSpeech (WER%) and Korean (WER%) from the SPREDS-U1 datasets in mitigating catastrophic forgetting

ID	Method	Aishell		Librispeech		SPREDS 04kor
		dev	test	test clean	test other	
C1	Whisper Small	12.72	13.91	3.40	7.60	18.90
C2	Whisper Medium	9.38	9.54	2.90	5.90	15.96
C3	+ FFT	15.82	17.16	13.15	22.17	436
C4	+ LoRA	9.37	9.73	3.51	6.97	16.80
C5	+ SPT4AS	9.44	9.70	3.50	7.46	17.03

slightly worse than FFT while utilizing considerably fewer parameters. However, it is worth acknowledging that, in terms of both parameter efficiency and performance, LoRA (S2 and M2) still outperforms SPT4ASR (S7 and M3), whether in the small or medium model. Surprisingly, we found that when applying Vanilla soft prompts to Whisper-small and Whisper-medium models, integrating soft prompts with whole model full tuning resulted in significantly better performance (S8 and M4) compared to standard FFT methods (S1 and M1). For instance, the Whisper-medium model with both whole model FFT and SPT(M4) achieved a devman MER of 11.17% and a devsge MER of 16.61%, outperforming the Medium FFT approach(M1) with only a 0.27M increase in parameters.

4.3. Mitigation of Catastrophic Forgetting

To assess the effectiveness of SPT in alleviating catastrophic forgetting, we compared it with FFT and LoRA methods on monolingual ASR tasks using Whisper Medium models. Table 4 shows that while FFT-based models (C3) exhibited significant performance degradation, both SPT4ASR (C5) and LoRA (C4) maintained performance close to the baseline Whisper Medium model (C2). This trend is consistent across Whisper Small and Medium models. This preservation of performance was consistent in both monolingual ASR experiments, suggesting that SPT and LoRA effectively mitigate catastrophic forgetting. The efficacy of SPT and LoRA can be attributed to their plug-and-play nature, allowing for model adaptation without parameter expansion.

5. Conclusion

We introduce FFT with SPT, compare various SPT-based methods, and integrate them into SPT4ASR to assess their effectiveness in low-resource transfer through CS ASR, addressing the gap in comprehensive SPT evaluation for ASR. Our approach reduces errors and mitigates catastrophic forgetting, though it does not outperform FFT, while using fewer task-specific parameters, similar to LoRA. While LoRA remains more effective in terms of performance and parameter efficiency, soft prompt-based methods represent a viable alternative, offering a compelling trade-off between adaptability and resource usage, and warrant continued investigation in multilingual ASR.

6. References

- [1] A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Yang, A. Fan *et al.*, “The llama 3 herd of models,” *arXiv preprint arXiv:2407.21783*, 2024.
- [2] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, “Learning transferable visual models from natural language supervision,” in *Proc. 38th Int. Conf. Mach. Learn. (ICML)*, ser. Proceedings of Machine Learning Research, vol. 139. PMLR, 2021, pp. 8748–8763.
- [3] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” in *Proc. 40th Int. Conf. Mach. Learn. (ICML)*, ser. Proc. Mach. Learn. Res., vol. 202. PMLR, 2023, pp. 28492–28518.
- [4] I. J. Goodfellow, M. Mirza, D. Xiao, A. Courville, and Y. Bengio, “An empirical investigation of catastrophic forgetting in gradient-based neural networks,” *arXiv preprint arXiv:1312.6211*, 2013.
- [5] S. Wang, C. H. Yang, J. Wu, and C. Zhang, “Can whisper perform speech-based in-context learning?” in *Proc. Int. Conf. Acoust. Speech Signal Process. (ICASSP)*. IEEE, 2024, pp. 13421–13425.
- [6] P. Peng, B. Yan, S. Watanabe, and D. Harwath, “Prompting the hidden talent of web-scale speech models for zero-shot task generalization,” in *Proc. 24th Annu. Conf. Int. Speech Commun. Assoc. (Interspeech)*. ISCA, 2023, pp. 396–400.
- [7] C. Yang, K. Huang, and H. Lee, “Do Prompts Really Prompt? Exploring the Prompt Understanding Capability of Whisper,” *CoRR*, vol. abs/2406.05806, 2024.
- [8] 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 *Proc. APSIPA Annu. Summit Conf. (APSIPA ASC)*. IEEE, 2024, pp. 1–6.
- [9] R. Jain, A. Barcovschi, M. Y. Yiwere, P. Corcoran, and H. Cucu, “Adaptation of whisper models to child speech recognition,” in *Proc. 24th Annu. Conf. Int. Speech Commun. Assoc. (Interspeech)*. ISCA, 2023, pp. 5242–5246.
- [10] B. Talafha, A. Waheed, and M. Abdul-Mageed, “N-shot benchmarking of whisper on diverse arabic speech recognition,” in *Proc. 24th Annu. Conf. Int. Speech Commun. Assoc. (Interspeech)*. ISCA, 2023, pp. 5092–5096.
- [11] 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,” in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2022.
- [12] Z. Song, J. Zhuo, Y. Yang, Z. Ma, S. Zhang, and X. Chen, “Lora-whisper: Parameter-efficient and extensible multilingual asr,” in *Proc. 25th Annu. Conf. Int. Speech Commun. Assoc. (Interspeech)*, 2024, pp. 3934–3938.
- [13] 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 *Proc. 25th Annu. Conf. Int. Speech Commun. Assoc. (Interspeech)*, 2024, pp. 2534–2538.
- [14] B. Lester, R. Al-Rfou, and N. Constant, “The power of scale for parameter-efficient prompt tuning,” in *Proc. 2021 Conf. Empir. Methods Nat. Lang. Process. (EMNLP)*. Assoc. Comput. Linguistics, 2021, pp. 3045–3059.
- [15] Z. Chen, C. Fu, C. Liu, S. D. Li, and H. Lee, “Exploring efficient-tuning methods in self-supervised speech models,” in *Proc. IEEE Spoken Lang. Technol. Workshop (SLT)*. IEEE, 2022, pp. 1120–1127.
- [16] H. Ma, Z. Peng, M. Shao, J. Li, and J. Liu, “Extending whisper with prompt tuning to target-speaker asr,” in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*. IEEE, 2024, pp. 12516–12520.
- [17] 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 *Proc. 25th Annu. Conf. Int. Speech Commun. Assoc. (Interspeech)*, 2024, pp. 2544–2548.
- [18] D. Ng, C. Zhang, R. Zhang, Y. Ma, F. R. Gutierrez, T. H. Nguyen, C. Ni, S. Zhao, E. S. Chng, and B. Ma, “Are soft prompts good zero-shot learners for speech recognition?” in *ICASSP 2024*. IEEE, 2024, pp. 10366–10370.
- [19] F. Philippy, S. Guo, S. Haddadan, C. Lothritz, J. Klein, and T. F. Bissyandé, “Soft Prompt Tuning for Cross-lingual Transfer: When Less is More,” *CoRR*, vol. abs/2402.03782, 2024.
- [20] M. Zhao and H. Schütze, “Discrete and Soft Prompting for Multilingual Models,” 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. 8547–8555.
- [21] L. Huang, S. Ma, D. Zhang, F. Wei, and H. Wang, “Zero-shot Cross-lingual Transfer of Prompt-based Tuning with a Unified Multilingual Prompt,” in *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*. Association for Computational Linguistics, 2022, pp. 11488–11497.
- [22] Y. Peng, Y. Liu, J. Zhang, H. Xu, Y. He, H. Huang, and E. S. Chng, “Internal language model estimation based language model fusion for cross-domain code-switching speech recognition,” *CoRR*, vol. abs/2207.04176, 2022.
- [23] Y. Peng, J. Zhang, H. Xu, H. Huang, and E. S. Chng, “Minimum word error training for non-autoregressive transformer-based code-switching asr,” in *IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*. IEEE, 2022, pp. 7807–7811.
- [24] H. Liu, H. Xu, L. P. García, A. W. H. Khong, Y. He, and S. Khudanpur, “Reducing language confusion for code-switching speech recognition with token-level language diarization,” in *IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*. IEEE, 2023, pp. 1–5.
- [25] B. Yan, C. Zhang, M. Yu, S.-X. Zhang, S. Dalmia, D. Berrebbi, C. Weng, S. Watanabe, and D. Yu, “Joint modeling of code-switched and monolingual asr via conditional factorization,” in *IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*. IEEE, 2022, pp. 6412–6416.
- [26] Y. Lu, M. Huang, H. Li, J. Guo, and Y. Qian, “Bi-encoder transformer network for mandarin-english code-switching speech recognition using mixture of experts,” in *Proc. 21st Annu. Conf. Int. Speech Commun. Assoc. (Interspeech)*. ISCA, 2020, pp. 4766–4770.
- [27] J. Tian, J. Yu, C. Zhang, Y. Zou, and D. Yu, “Lae: Language-aware encoder for monolingual and multilingual asr,” in *Proc. 23rd Annu. Conf. Int. Speech Commun. Assoc. (Interspeech)*. ISCA, 2022, pp. 3178–3182.
- [28] Z. You, S. Feng, D. Su, and D. Yu, “Speechmoe: Scaling to large acoustic models with dynamic routing mixture of experts,” in *Proc. 22nd Annu. Conf. Int. Speech Commun. Assoc. (Interspeech)*. ISCA, 2021, pp. 2077–2081.
- [29] 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.
- [30] 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 Assoc. Comput. Linguistics: ACL 2023*. Assoc. Comput. Linguistics, 2023, pp. 6740–6757.
- [31] 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 *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*. IEEE, 2024, pp. 10941–10945.