



# Learn and Don't Forget: Adding a New Language to ASR Foundation Models

Mengjie Qian, Siyuan Tang, Rao Ma, Kate M. Knill, Mark J. F. Gales

ALTA Institute, Machine Intelligence Lab, Department of Engineering, Cambridge University, UK

{mq227, st941, rm2114, kmk1001, mjfg100}@cam.ac.uk

## Abstract

Foundation ASR models often support many languages, e.g. 100 languages in Whisper. However, there has been limited work on integrating an additional, typically low-resource, language, while maintaining performance on the original language set. Fine-tuning, while simple, may degrade the accuracy of the original set. We compare three approaches that exploit adaptation parameters: soft language code tuning, train only the language code; soft prompt tuning, train prepended tokens; and LoRA where a small set of additional parameters are optimised. Elastic Weight Consolidation (EWC) offers an alternative compromise with the potential to maintain performance in specific target languages. Results show that direct fine-tuning yields the best performance for the new language but degrades existing language capabilities. EWC can address this issue for specific languages. If only adaptation parameters are used, the language capabilities are maintained but at the cost of performance in the new language.

**Index Terms:** under-represented, low-resource, speech recognition, foundation models, elastic weight consolidation

## 1. Introduction

Recent years have seen impressive advancements in Automatic Speech Recognition (ASR) systems, particularly for languages with abundant linguistic resources, leading to high performance. However, the situation is different for ASR systems tailored to low-resource languages, the performance of which still falls short when compared to their high-resource counterparts. Low-resource languages typically lack a strong online presence, linguistic expertise, sufficient speech and text data, and pronunciation lexicons. These characteristics pose challenges for developing effective ASR systems in such languages.

In the era of hybrid DNN-HMM models, extensive research on ASR for low-resource languages focused on building propriety lexicons for modelling the acoustic units attracts. [1, 2] proposed a global lexicon, which captures dialect variant forms with relatively abstract representations, and a multi-dialect lexicon containing all dialect variants for Irish. Another approach was to develop graphemic lexicons where the “pronunciation” for a word is defined by the letters forming the word e.g. [3]. End-to-end ASR models don't require a lexicon, so this doesn't affect performance. However, the lack of transcribed training data, as for the hybrid models, is a big problem for low-resources languages.

<sup>1</sup>This paper reports on research supported by EPSRC Project EP/V006223/1 (Multimodal Video Search by Examples) and Cambridge University Press & Assessment, a department of The Chancellor, Masters, and Scholars of the University of Cambridge.

To handle the lack of training data, a common method is using self-supervised or semi-supervised training [4, 5, 6], where a model initially trained with limited transcribed data is used to generate transcriptions for unlabeled data which are then added to the training data to update the model. Data augmentation can also be used to tackle this problem where additional data is synthesised from existing data e.g. [7, 8, 9, 10].

With the development and release of pre-trained multilingual end-to-end foundation ASR models [11, 12, 13], utilising a multilingual model is another option for low-resource languages. Zhang et al. [14] proposed a smaller universal monolingual output layer shared across languages for high-quality and high-efficiency multilingual ASR. Chen et al. [15] used hierarchical CTC to leverage language identity throughout the entire encoder-decoder network, aiming to improve ASR performance by correctly identifying languages. Qin et al. [16] utilised multilingual and multilevel unit modeling to improve ASR performance on low-resource Tibetan languages.

The challenge in multilingual models lies in expanding them to new languages without compromising the performance of existing ones, particularly when the training data for new languages is limited or unavailable. In NLP, approaches like language-specific adapters and output heads, as seen in [17, 18], enable parameter-efficient fine-tuning but necessitate prior knowledge of the input language. Alternatively, continual lifelong learning, as proposed by [19], combines data from both existing and new languages for ongoing training. Demonstrating its effectiveness in ASR, [20, 21] have expanded language coverage from 15 to 32 and to 66 languages, respectively.

In this paper, we focus on adding a new language to an ASR foundation model, using Whisper as a case study. Firstly, the zero-shot ability of the foundation model is examined. Secondly, we compare efficient fine-tuning methods such as Low Rank Adaptation (LoRA) with standard fine-tuning. Moreover, we propose Soft Language Code Tuning (SLCT) to train a language-specific embedding vector for the new language. Inspired by techniques from NLP tasks, we implement Soft Prompt Tuning (SPT), introduced in [22] and utilised in a language assessment task [23], to effectively expand language coverage. For fine-tuning, we analysed catastrophic forgetting and mitigated it with Elastic Weight Consolidation (EWC). Initially, we assess the zero-shot ability of Whisper on six languages not originally supported in the model, then focus on three languages to evaluate the performance of various approaches.

## 2. Methods

### 2.1. Fine-tuning

If the parameters of the ASR model are denoted as  $\theta_{ASR}$ , given a set of acoustic features, the goal of the ASR model is to pro-

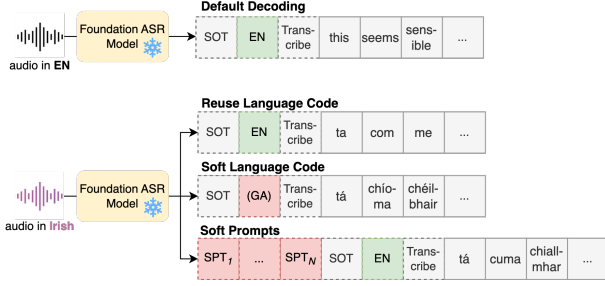


Figure 1: Illustration of the decoding process for existing and new languages using the default decoding and three methods: (a) reusing an existing language code, (b) using a trained soft language code, and (c) employing prepended soft prompts to capture new domain knowledge. Reference text for Irish audio: “ta cuma chiallmhar air seo”.

duce the transcription  $\hat{Y} = \{\hat{y}_1, \hat{y}_2, \dots, \hat{y}_t\}$ , which is determined by the equation  $\hat{Y} = \arg \max_Y P(Y|X; \theta_{ASR})$ . Model fine-tuning (FT) is a standard approach for domain adaptation, where all model parameters are tuned using domain-specific data to minimise the training loss. In this work, we adopt the fine-tuning method to integrate a new target language into Whisper. Benefiting from pre-training on an extensive set of labelled audio-transcription data, Whisper generally requires minimal additional fine-tuning to achieve an effective ASR model tailored for the target domain.

## 2.2. Efficient Fine-tuning with LoRA

Low Rank Adaptation (LoRA) introduces an innovative and efficient fine-tuning strategy by leveraging low-rank approximation to adapt models to new data distributions [24]. We denote the parameters of the ASR model as  $\theta_{ASR} = [\theta_1, \theta_2, \dots, \theta_L]$ , where each  $\theta_l$  represents a weight matrix. Instead of directly updating, LoRA constrains the update of a weight matrix  $\theta_l \in R^{d \times k}$  with a low-rank decomposition, such that  $\theta_l + \Delta\theta_l = \theta_l + B_l \times A_l$ , where  $B_l \in R^{d \times r}$ ,  $A_l \in R^{r \times k}$  and the rank  $r \ll \min(d, k)$ . This introduces new weight matrices  $B$  and  $A$  specific to the adapted domain, while preserving a set of original model parameters  $\theta_{ASR}$  and ensuring the performance on all original tasks is unchanged.

## 2.3. Soft Language Code Tuning

Whisper identifies each language by a unique token during its training phase [11]. For an unseen language, however, this language code does not exist. To tackle this challenge, we introduce the concept of Soft Language Code Tuning (SLCT), depicted in Figure 1. Rather than using the pre-trained embeddings for a specific language code, SLCT trains a new embedding vector for the target language code with the training set of the target language, while keeping other model parameters fixed. The soft language embeddings are exclusively employed for the target language, effectively incorporating a new language into the model while preserving the model’s performance on prior tasks.

## 2.4. Soft Prompt Tuning

Prompting [25, 26] has become a popular method to adapt large language models (LLM) [27, 28, 29]. Beyond the traditional discrete prompts, soft prompt tuning (SPT) employs continu-

ous vectors that are inserted into the model’s input and optimised via gradient descent on training data [22]. This approach enhances model performance and requires minimal human expertise. [23] adopted soft prompt tuning for the ASR domain and demonstrated its effectiveness on the spoken language assessment task. Inspired by this, our work extends the concept of SPT to incorporate new languages into a foundation ASR model. Specifically,  $m$  token embeddings  $V = \{v_0, \dots, v_{m-1}\}$  were prepended to the decoder input embeddings. During training, only these embeddings  $V$  are updated based on the loss function while keeping the original model parameters,  $\theta_{ASR}$ , unchanged. In this way, extra information is introduced for the model to condition on during the generation process, where  $\hat{Y} = \arg \max_Y P(Y|X; \theta_{ASR}, \theta_V)$ . An illustration of the modified model input to the ASR decoder during training and decoding is shown in Figure 1.

## 2.5. Forgetting and Elastic Weight Consolidation

LoRA, SLCT, and SPT all have the ability to maintain performance on prior domains or tasks as they do not alter the original model parameters. However, model fine-tuning carries the risk of catastrophic forgetting, a common challenge faced by large foundation models, which occurs when the model’s proficiency in prior tasks declines as it undergoes continuous updates or fine-tuning for new tasks or domains [30, 31, 32, 33]. Various strategies to mitigate this issue comprises Pseudo-rehearsal [34, 35], Memory Aware Synapses [36] and Elastic Weight Consolidation (EWC) [37]. In this study, we adopt EWC both as an analytic tool and as a regularisation approach to address the problem of catastrophic forgetting when adding a new language to the foundation ASR model.

EWC employs the observed Fisher information matrix  $F$  to identify important parameters relative to the previous tasks during the pre-training phase. Fine-tuning is then modified to prevent large deviations in these important parameters. For a previous task  $T_A$  (e.g. ASR for an existing language) and a new task  $T_B$  (e.g. ASR for a new language), the loss function  $L$  that we minimise in EWC is defined as follows:

$$L(\theta) = L_B(\theta) + \lambda \sum_i F_i (\theta_i - \theta_{A,i}^*)^2 \quad (1)$$

where  $L_B(\theta)$  is the loss for task  $T_B$  only,  $\lambda$  sets the importance of the previous task  $T_A$  relative to the new task  $T_B$ , for each parameter  $i$ . Here,  $F_i$  represents the  $i$ -th diagonal element of the Fisher information matrix,  $\theta_i$  is the  $i$ -th parameter of the model, and  $\theta_{A,i}^*$  is the optimal value of the  $i$ -th parameter after training on the previous task  $T_A$ .

In addition to using EWC as a training constraint, we also utilise Fisher overlap [37], a measure of the degree of overlap between the two tasks’ Fisher matrices, as an analytic tool to examine potential forgetting of previous languages when integrating a new language into Whisper. The observations derived from this analysis can be tested through EWC fine-tuning. Precisely, we computed the Fisher matrices associated with two languages,  $F_1$  and  $F_2$ , and normalised these to each have unit trace,  $\hat{F}_1$  and  $\hat{F}_2$ , then computed their Fréchet distance, a metric on the space of positive-semidefinite matrices:

$$d^2(\hat{F}_1, \hat{F}_2) = \frac{1}{2} \times \|\hat{F}_1^{1/2} - \hat{F}_2^{1/2}\|_F \quad (2)$$

which is bounded between 0 and 1. We then define the overlap as  $1 - d^2$ , where a value of zero indicates that the two tasks depend on non-overlapping sets of weights, while a value of one indicates that  $F_1 = \alpha F_2$  for some  $\alpha > 0$ .

### 3. Experiments

#### 3.1. Datasets

Experiments were conducted on the FLEURS benchmark dataset [38], which comprises 102 languages, of which 19 languages are not supported by the Whisper large-v3 model. Each language has about 12 hours of read speech data, with each audio within 30 seconds. Initial tests were performed on six languages randomly selected from the 19 not supported ones: Irish, Asturian, Sorani Kurdish, Cebuano, Kabuverdianu, and Kyrgyz. We then focused on the former three languages (Irish, Asturian, and Sorani Kurdish) for further investigation, selected for their diverse range of zero-shot ASR performance in terms of Character Error Rates (CER). The number of files and hours of the datasets of these three languages are presented in Table 1. Additionally, our proposed approaches are also tested on languages in Whisper’s supported list using the FLEURS dataset, including English (en), Spanish (es), Welsh (cy), and Occitan (oc). Both the original raw transcript and normalisation version with cases lowering and punctuation removing are provided in the dataset, the latter was used as an evaluation reference. Three out-of-domain datasets are utilised to evaluate models’ forgetting behavior, namely (1) Librispeech (LB), an audio read speech corpus [39], (2) Artie Bias Corpus (Artie), a subset of Common Voice Corpus [40], and (3) TED-LIUM (Ted) test set which contains audio from TED conference videos [41].

Table 1: FLEURS datasets used in training.

Language	Train		Dev		Test	
	#	hrs	#	hrs	#	hrs
Irish	2,845	12.12	369	1.49	842	3.46
Asturian	2,511	7.54	398	0.99	946	2.44
Sorani Kurd	3,040	10.46	386	1.23	922	2.99

Whisper has various model sizes. We use the latest version large-v3 as the foundation model because it shows improved performance over a wide range of languages compared with other models<sup>1</sup>. Fine-tuning (FT) employs an initial learning rate of 1e-5 and decays linearly, while SPT and SLCT use initial rates of 1e-4 and 1e-1, respectively. SPT results are presented with 20 soft prompts. LoRA was applied to attention layers with a rank of 8, with varying rank sizes showing negligible differences in results. EWC experiments were conducted with  $\lambda$  values (in Equation 1) selected from the range of [1e-0, 1e-1, 1e-2, 1e-3, 1e-4, 1e-5]. The optimal  $\lambda$  was determined on the development set.

#### 3.2. Zero-shot ASR and Speech Translation

The Whisper large-v3 model has been pre-trained on 100 languages for ASR, Speech Translation (ST) and Language Identification (LID) tasks; exposed to extensive acoustic and language data, it prompts an intriguing question regarding its capability for zero-shot ASR or ST for an unseen language. To investigate this, we performed zero-shot ASR and ST on six languages not supported by Whisper with the FLEURS test sets on the large-v3 model. Results, presented in Table 2, reveal a diverse range of CER, yet consistently poor Word Error Rates (WER) across all languages, underscoring the challenges of zero-shot ASR for unseen languages.

<sup>1</sup><https://github.com/openai/whisper/discussions/1762>

For ST evaluation, we use BLEU and COMET scores, specifically using the wmt22-comet-da [42] model to calculate the COMET score. Although wmt22-comet-da was trained on numerous languages, it lacks Asturian, Cebuano, Kabuverdianu, and Sorani Kurdish, potentially affecting COMET scores for these languages. Whisper achieved notably high BLEU scores, especially for Asturian and Kabuverdianu, despite not being trained on these languages, demonstrating its capability for zero-shot speech translation on unseen languages. Since ASR appears to be more challenging for under-represented languages, we focus on the ASR task for the rest of the paper.

Table 2: Whisper large-v3 model zero-shot performance for under-represented languages on the ASR and ST tasks.

Language	ASR		ST	
	CER	WER	BLEU	COMET
Asturian	13.5	47.5	27.6	73.1
Cebuano	47.7	72.9	8.8	55.1
Kabuverdianu	35.2	87.8	32.3	73.7
Kyrgyz	50.0	87.4	4.7	56.5
Sorani Kurdish	38.6	110.8	2.2	45.6
Irish	78.8	110.4	1.7	43.1

#### 3.3. Language code selection

Whisper leverages a language code for decoding a supported language. However, there is no corresponding language code when a low-resource language is not supported. One straightforward approach to address this limitation is to find a related source language from the supported list and utilise its language code for inference on the target low-resource language. Taking Irish (‘ga’) as an illustration, we selected four language codes associated with languages exhibiting the highest character and token overlaps with Irish. In addition, we applied Soft Language Code Tuning (SLCT), to train an Irish-specific embedding vector specific. Table 3 presents results on the FLEURS Irish test set. The choice of language code can be seen to impact ASR performance, showing an absolute difference of 33.2% in CER and 20.3% in WER. Similar findings were noted on Asturian and Sorani Kurdish languages. SLCT gives the best performance, setting an upper bound for language code selection.

In contrast to previous findings, fine-tuning proves to be robust to variations in language code, delivering large performance gains regardless of the chosen language code, as per-

Table 3: ASR performance on FLEURS Irish test set with the baseline and the fine-tuned Whisper large-v3 model using different language code (en: English, cy: Welsh, oc: Occitan, es: Spanish, ga: Irish).

Language Code	Baseline		FT	
	CER	WER	CER	WER
en	78.8	110.4	22.3	<b>43.7</b>
cy	62.3	108.0	<b>22.2</b>	44.0
oc	<b>45.6</b>	<b>90.1</b>	22.4	44.2
es	79.2	107.3	22.6	44.7
ga	-	-	22.8	44.5
SLCT	<b>41.5</b>	<b>81.4</b>	-	-

sented in Table 3. With the ‘en’ language code, we observe an improvement of 56% and 66% in CER and WER, respectively, on the test set. Interestingly, when introducing a new language code (‘ga’) for training, we observe slightly worse performance. This is likely due to the fact that the ‘ga’ code is new to the model, suggesting a potential need for more training data to effectively integrate a new language code.

### 3.4. Efficient fine-tuning

Table 4 presents a comparison of different fine-tuning approaches across three target languages: Irish, Asturian, and Sorani Kurdish. Baseline performance is obtained using the same Whisper large-v3 model for all three languages, while other models (FT, SLCT, SPT, and LoRA) are specifically trained for each language using the respective FLEURS training sets. For Irish, FT achieves the best performance improvement, with a 56.5% absolute reduction in CER and a 66.7% reduction in WER compared to other efficient fine-tuning approaches. This trend is consistent with the other two languages, which is unsurprising as the model is tailored to a specific domain for each target language, adjusting all parameters from the original domain. SLCT results in performance gains of 26.3%, 37.3%, and 14.8% on Irish, Asturian, and Sorani Kurdish, respectively. Despite the relatively low performance, it’s worth noting that only 0.00008% of the model’s 1550M parameters are tuned, showcasing noticeable achievement given the tuned parameter size. SPT, which tunes 0.001% of the model parameters, further enhances the performance by 32.2%, 59.2%, and 44.6% relative on Irish, Asturian, and Sorani Kurdish, respectively. LoRA, tuning 0.13% of model parameters, demonstrates improved performance proportionate to the increased size of trained parameters.

Table 4: ASR performance on FLEURS Irish, Asturian, and Sorani Kurdish test sets with different parameter-tuning approaches.

Model	Irish		Asturian		Sorani Kurdish	
	CER	WER	CER	WER	CER	WER
Baseline	78.8	110.4	13.5	47.5	38.6	110.8
FT	22.3	43.7	3.7	11.4	7.3	32.9
SLCT	41.5	81.4	7.4	29.8	35.8	95.2
SPT	38.1	74.9	5.1	19.4	13.6	61.4
LoRA	28.3	54.8	4.1	14.0	8.7	41.3

### 3.5. Forgetting and EWC

As discussed in Section 2.5, fine-tuning a model on a specific target language can lead to a performance decline in other languages. To demonstrate this effect, we analysed the Fisher overlap among the test sets for various languages, as depicted in Figure 2. Notably, Irish exhibits a higher Fisher overlap with Welsh (cy) and Occitan (oc) than with English (en) and Spanish (es), indicating a greater parameter overlap between Irish and Welsh (or Occitan). This implies that fine-tuning the model on Irish could adversely affect its performance on Welsh, a finding confirmed by the results in Table 5 where the WER for Welsh (cy) decreases from 27.5% to 96.8%, the biggest degradation compared to other languages. When training is constrained with EWC Fisher parameters extracted from the FLEURS English training set, we note a decline in performance on the Irish test set but improvements across all other languages compared to




Figure 2: Fisher overlap between Fisher parameters extracted for different languages or datasets.

standard fine-tuning. It even manages to preserve the performance on out-of-domain English test sets, halving the WER degradation compared to standard fine-tuning. However, maintaining performance on cy proves more challenging than for en or es due to the higher Fisher overlap between Irish and cy. By constraining the fine-tuning with EWC Fisher parameters extracted with the FLEURS cy training set with a  $\lambda$  value set to 1, we observe better performance preservation on cy but a small decline in performance on Irish. This highlights the difficulty in enhancing performance on two tasks with high Fisher overlap.

Table 5: Evaluation of WER on in-domain and out-of-domain test sets using Whisper large-v3 and its fine-tuned version. The model is fine-tuned with FLEURS Irish training set. Optimal  $\lambda$  for EWC (en):  $1e-1$ ; optimal  $\lambda$  for EWC (cy): 1.0.

Model	FLEURS					Out-of-domain		
	Irish	en	es	cy	oc	LB	Artie	Ted
Baseline	110.4	4.1	2.8	27.5	64.2	3.7	5.6	4.0
FT	43.7	6.1	4.1	96.8	79.8	8.8	13.0	4.2
+ EWC (en)	46.7	4.6	2.8	78.7	71.5	5.4	9.0	4.0
+ EWC (cy)	47.0	4.5	3.0	66.1	71.3	5.4	8.8	4.0

## 4. Conclusion

In this paper, we examine Whisper’s zero-shot ASR and ST capabilities on unseen languages and then focus on adding new languages into Whisper for the ASR task. In addition to the implementation of LoRA, we introduce SLCT and SPT for efficient fine-tuning. While standard fine-tuning achieves the best performance compared to various efficient fine-tuning approaches, it faces the challenge of catastrophic forgetting. To address this, we adopt Fisher overlap as an analytic tool for assessing potential forgetting on different languages. Employing EWC of Fisher parameters, we can constrain forgetting and preserve performance on existing languages. Our experiments reveal that it’s hard to learn a new task and not forget a previous task if two tasks have a high Fisher overlap. In future work, we will extend the investigation into speech translation and explore integrating a new language into foundation models for the translation task while ensuring that performance on previous languages or tasks is preserved.

## 5. References

- [1] M. Qian, H. Berthelsen, L. Lonergan *et al.*, “Automatic Speech Recognition for Irish: testing lexicons and language models,” in *33rd Irish Signals and Systems Conference (ISSC)*. IEEE, 2022, pp. 1–6.
- [2] L. Lonergan, M. Qian, N. N. Chiarain *et al.*, “Cross-dialect lexicon optimisation for an endangered language ASR system: the case of Irish,” in *Proc. Interspeech 2022*, 2022, pp. 4865–4869.
- [3] M. J. Gales, K. M. Knill, and A. Ragni, “Unicode-based graphemic systems for limited resource languages,” in *Proc. ICASSP*, 2015, pp. 5186–5190.
- [4] H. Zhu, L. Wang, G. Cheng, J. Wang, P. Zhang, and Y. Yan, “Wav2vec-S: Semi-Supervised Pre-Training for Low-Resource ASR,” in *Proc. Interspeech 2022*, 2022, pp. 4870–4874.
- [5] G. Synnaeve, Q. Xu, J. Kahn *et al.*, “End-to-End ASR: from Supervised to Semi-Supervised Learning with Modern Architectures,” in *ICML Workshop on Self-supervision in Audio and Speech*, 2020.
- [6] K. Vesely, “Semi-Supervised DNN Training with Word Selection for ASR,” in *Interspeech*, 2017, pp. 3687–3691.
- [7] J. Xie and J. H. L. Hansen, “MixRep: Hidden Representation Mixup for Low-Resource Speech Recognition,” in *Proc. INTERSPEECH 2023*, 2023, pp. 1304–1308.
- [8] M. Bartelds, N. San, B. McDonnell, D. Jurafsky, and M. Wieling, “Making More of Little Data: Improving Low-Resource Automatic Speech Recognition Using Data Augmentation,” *arXiv preprint arXiv:2305.10951*, 2023.
- [9] D. S. Park, W. Chan, Y. Zhang *et al.*, “SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition,” in *Proc. Interspeech 2019*, 2019, pp. 2613–2617.
- [10] A. Ragni, K. M. Knill, S. P. Rath, and M. J. F. Gales, “Data augmentation for low resource languages,” in *Proc. Interspeech 2014*, 2014, pp. 810–814.
- [11] A. Radford, J. W. Kim, T. Xu *et al.*, “Robust speech recognition via large-scale weak supervision,” in *International Conference on Machine Learning*. PMLR, 2023, pp. 28 492–28 518.
- [12] Y. Zhang, W. Han, J. Qin *et al.*, “Google USM: Scaling Automatic Speech Recognition Beyond 100 Languages,” *arXiv preprint arXiv:2303.01037*, 2023.
- [13] A. Babu, C. Wang, A. Tjandra *et al.*, “XLS-R: Self-supervised cross-lingual speech representation learning at scale,” *arXiv preprint arXiv:2111.09296*, 2021.
- [14] C. Zhang, B. Li, T. Sainath *et al.*, “UML: A Universal Monolingual Output Layer For Multilingual ASR,” in *Proc. ICASSP*, 2023, pp. 1–5.
- [15] W. Chen, B. Yan, J. Shi *et al.*, “Improving massively multilingual asr with auxiliary etc objectives,” in *Proc. ICASSP*, 2023, pp. 1–5.
- [16] S. Qin, L. Wang, S. Li, J. Dang, and L. Pan, “Improving low-resource Tibetan end-to-end ASR by multilingual and multilevel unit modeling,” *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2022, no. 1, pp. 1–10, 2022.
- [17] J. He, C. Zhou, X. Ma *et al.*, “Towards a Unified View of Parameter-Efficient Transfer Learning,” in *International Conference on Learning Representations*, 2021.
- [18] N. Houlsby, A. Giurgiu, S. Jastrzebski *et al.*, “Parameter-efficient transfer learning for NLP,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 2790–2799.
- [19] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter, “Continual lifelong learning with neural networks: A review,” *Neural networks*, vol. 113, pp. 54–71, 2019.
- [20] B. Li, R. Pang, T. N. Sainath *et al.*, “Scaling end-to-end models for large-scale multilingual asr,” in *2021 IEEE Workshop on ASRU*, 2021, pp. 1011–1018.
- [21] B. Li, R. Pang, Y. Zhang *et al.*, “Massively multilingual ASR: A lifelong learning solution,” in *Proc. ICASSP*, 2022, pp. 6397–6401.
- [22] 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*, 2021, pp. 3045–3059.
- [23] R. Ma, M. Qian, M. J. Gales, and K. Knills, “Adapting an ASR Foundation Model for Spoken Language Assessment,” in *9th Workshop on SLATE*, 2023, pp. 104–108.
- [24] E. J. Hu, P. Wallis, Z. Allen-Zhu *et al.*, “LoRA: Low-Rank Adaptation of Large Language Models,” in *International Conference on Learning Representations (ICLR)*, 2021.
- [25] P. Liu, W. Yuan, J. Fu *et al.*, “Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing,” *ACM Computing Surveys*, vol. 55, no. 9, pp. 1–35, 2023.
- [26] T. Shin, Y. Razeghi, R. L. Logan IV *et al.*, “AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts,” in *Proc. EMNLP*, 2020, pp. 4222–4235.
- [27] J. D. M.-W. C. Kenton and L. K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” in *Proceedings of NAACL-HLT*, 2019, pp. 4171–4186.
- [28] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, “Exploring the limits of transfer learning with a unified text-to-text transformer,” *The Journal of Machine Learning Research*, vol. 21, no. 1, pp. 5485–5551, 2020.
- [29] OpenAI, “GPT-4 Technical Report,” *arXiv preprint arXiv:2303.08774*, 2023.
- [30] M. McCloskey and N. J. Cohen, “Catastrophic interference in connectionist networks: The sequential learning problem,” in *Psychology of learning and motivation*. Elsevier, 1989, vol. 24, pp. 109–165.
- [31] R. French, “Catastrophic forgetting in connectionist networks,” *Trends in cognitive sciences*, vol. 3, no. 4, pp. 128–135, 1999.
- [32] R. Ratcliff, “Connectionist models of recognition memory: constraints imposed by learning and forgetting functions,” *Psychological review*, vol. 97, no. 2, p. 285, 1990.
- [33] I. J. Goodfellow, M. Mirza, A. Courville, and Y. Bengio, “An Empirical Investigation of Catastrophic Forgetting in Gradient-Based Neural Networks,” *stat*, vol. 1050, p. 6, 2014.
- [34] A. Robins, “Catastrophic forgetting, rehearsal and pseudorehearsal,” *Connection Science*, vol. 7, no. 2, pp. 123–146, 1995.
- [35] C. Atkinson, B. McCane, L. Szymanski *et al.*, “Pseudo-rehearsal: Achieving deep reinforcement learning without catastrophic forgetting,” *Neurocomputing*, vol. 428, pp. 291–307, 2021.
- [36] R. Aljundi, F. Babiloni, M. Elhoseiny *et al.*, “Memory aware synapses: Learning what (not) to forget,” in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 139–154.
- [37] J. Kirkpatrick, R. Pascanu, N. Rabinowitz *et al.*, “Overcoming catastrophic forgetting in neural networks,” *Proceedings of the national academy of sciences*, vol. 114, no. 13, pp. 3521–3526, 2017.
- [38] A. Conneau, M. Ma, S. Khanuja, Y. Zhang, V. Axelrod, S. Dalmia, J. Riesa, C. Rivera, and A. Bapna, “Fleurs: Few-shot learning evaluation of universal representations of speech,” in *Spoken Language Technology Workshop (SLT)*, 2023, pp. 798–805.
- [39] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an ASR corpus based on public domain audio books,” in *Proc. ICASSP*, 2015, pp. 5206–5210.
- [40] J. Meyer, L. Rauchenstein, J. D. Eisenberg, and N. Howell, “Artie bias corpus: An open dataset for detecting demographic bias in speech applications,” in *Proceedings of the Twelfth Language Resources and Evaluation Conference*, 2020, pp. 6462–6468.
- [41] F. Hernandez, V. Nguyen, S. Ghannay, N. Tomashenko, and Y. Esteve, “TED-LIUM 3: Twice as much data and corpus repartition for experiments on speaker adaptation,” in *Speech and Computer: 20th International Conference*. Springer, 2018, pp. 198–208.
- [42] R. Rei, J. G. De Souza, D. Alves *et al.*, “COMET-22: Unbabel-IST 2022 submission for the metrics shared task,” in *Proc. WMT*, 2022, pp. 578–585.