



Towards Rehearsal-Free Multilingual ASR: A LoRA-based Case Study on Whisper

Tianyi Xu¹, Kaixun Huang¹, Pengcheng Guo¹, Yu Zhou², Longtao Huang², Hui Xue², Lei Xie^{1†}

¹ Audio, Speech and Language Processing Group (ASLP@NPU), School of Computer Science, Northwestern Polytechnical University, Xi'an, China

² Alibaba Group, China

xutianyi@mail.nwpu.edu.cn, lxie@nwpu.edu.cn

Abstract

Pre-trained multilingual speech foundation models, like Whisper, have shown impressive performance across different languages. However, adapting these models to new or specific languages is computationally extensive and faces catastrophic forgetting problems. Addressing these issues, our study investigates strategies to enhance the model on new languages in the absence of original training data, while also preserving the established performance on the original languages. Specifically, we first compare various LoRA-based methods to find out their vulnerability to forgetting. To mitigate this issue, we propose to leverage the LoRA parameters from the original model for approximate orthogonal gradient descent on the new samples. Additionally, we also introduce a learnable rank coefficient to allocate trainable parameters for more efficient training. Our experiments with a Chinese Whisper model (for Uyghur and Tibetan) yield better results with a more compact parameter set.

Index Terms: Automatic Speech Recognition, Parameter-efficient tuning, Orthogonal gradient, Continual learning

1. Introduction

Recently, a large number of pre-trained foundation models have emerged in the speech areas, either trained in a self-supervised [1, 2] or supervised learning [3, 4] manner. Among these models, OpenAI's Whisper [3] has shown promising results for multilingual automatic speech recognition (ASR) and translation tasks. Despite its robust performance across various languages, Whisper exhibits suboptimal recognition capabilities for low-resource languages, including multilingual Tibetan and Uyghur. The common approach to improve performance in such languages involves fine-tuning the pre-trained model with data from the target language, either with full fine-tuning, or with Lora [5], a popular parameter efficient tuning method. Both approach, however, comes with a significant drawback. The fine-tuned model often experiences a noticeable decline in its ability to recognize the languages it was originally trained in, leading to a phenomenon known as catastrophic forgetting.

To counteract catastrophic forgetting, many continual learning methods have been proposed in the image classification community. Existing studies about continual learning can be mainly classified into rehearsal-based, regularization-based, and architecture-based approaches. Rehearsal-based approaches [6, 7] involve storing samples from previous tasks and combining them with the new data to jointly train the model. However, this method may not always be practical when original samples are unavailable or privacy concerns arise. For large-scale foundation models, the significant memory needed

to store these samples also poses a challenge. Regularization-based approaches [8, 9, 10] add a regularization term to the loss function to discourage large changes to the weights trained on previous tasks. While this method helps maintain training stability, it may limit the model's flexibility to learn new tasks. Architecture-based approaches add task-specific adapters for the new tasks, which prevents modifications to the original parameters [11, 12]. However, this requires the model to infer with the new task-id to decide which part of the model should be used, limiting its applicability to foundation models that lack access to task-ids.

In ASR, research on continual learning is relatively scarce [13, 14]. Existing studies typically perform direct gradient updates on the model's hidden layers within a common vector space for all tasks. However, recent research in the field of Natural Language Processing (NLP) has proposed a promising alternative. When training on new tasks, the loss gradients for new samples are projected orthogonally to the gradient subspace of past samples [15]. Given that neural networks are often over-parameterized, this strategy avoids conflicts with previous loss functions, thereby mitigating the problem of catastrophic forgetting. O-LoRA [16] has indicated that the LoRA parameters not only capture the gradient subspace of previous tasks, but can also be used to approximate the gradients of previous tasks to guide continual learning.

Our main contribution is investigating the catastrophic forgetting phenomenon observed in Whisper fine-tuning and utilizing orthogonal gradient descent to alleviate it. Our experimental result shows that the LoRA parameters of the original models are good representatives of the original data samples, and can be used to instruct the continual learning process. Furthermore, inspired by the AdaLoRA [17], we utilize a dynamic approach to allocate the ranks of matrices more judiciously, thereby improving model performance and further reducing the parameter count necessary for training.

Distinct from existing papers on continual learning research in ASR, this study emphasizes addressing a practical challenge faced by pre-trained speech foundation models: *How to fine-tune large pre-trained models for new tasks in a parameter-efficient way when the original training data is not available.* In real-world applications, new data would accumulate over time and these data could potentially be used to improve the model performance. Furthermore, new requirements would emerge and demand the model to solve new tasks. Our proposed method offers the following advantages: (1) **Rehearsal-free:** In many scenarios, the new data would be much less than the data used to train the original model, therefore it is not efficient to retrain the model using all data. Furthermore, due to privacy concerns, it may not be possible to retain old data. (2) **Parameter-efficient:** Given that the data for new tasks may be

[†]Lei Xie is the corresponding author.




Figure 1: The structure of O-LoRA and O-AdaLoRA.

limited and insufficient for full-finetuning, the proposed method should be parameter-efficient. We use the LoRA method to perform parameter-efficient finetuning and borrow ideas from AdaLoRA to dynamically assign a rank to LoRA parameters to further reduce the additional parameters. (3) **Task-id-free**: For real-world applications, such as multilingual ASR, we aim for the model to infer without explicitly task id, enhancing its generalizability.

To the best of our knowledge, our work is the first to investigate a rehearsal-free method based on orthogonal-gradient decent to solve continual learning challenges encountered by ASR models like Whisper. Although we primarily discuss continual learning for the Whisper model in this study, the proposed methods can be expanded to other speech foundation models with the Transformer architecture.

2. Method

2.1. Problem definition

First, we provide a more formal definition of the domain transfer problem in multilingual ASR models. We begin by assuming an initial model, M_0 , which was trained with LoRA on a dataset D_0 . Subsequent continual learning is built upon M_0 , with an intention to preserve its capabilities, while D_0 is no longer usable. Afterward, we have a series of labeled datasets of various languages, $D = \{D_n\}_{n=1}^N$, intended for subsequent multilingual fine-tuning training. The model trained with D_n is referred to as M_n . N represents the total number of subsets, and the size of each D_n is significantly smaller than D_0 . Utilizing D_0 for review training or retraining with it requires a substantial training cost. For an ideal continual learning model, at each time step n , we desire it to exhibit the following characteristics:

1. **No performance degradation on previous tasks**: M_n achieves the same or better results on D_{n-1} compared to M_{n-1} .
2. **Improved performance on new tasks**: M_n achieves better results on D_n compared to M_{n-1} .
3. **Comparable performance to a model trained on all data**: M_n achieves comparable results to a model trained from scratch using the datasets $D = \{D_n\}_{n=1}^N$.
4. **Parameter efficiency**: The required training expenditure is significantly less than retraining the model using all datasets $D = \{D_n\}_{n=1}^N$.

2.2. LoRA

To achieve parameter efficiency, we use LoRA [5], which was first proposed in NLP, to adapt models to specific domains or downstream tasks. Research has indicated that the weights of pre-trained large language models often occupy a low intrinsic dimensional space. Inspired by this observation, LoRA freezes the original weights and updates only the low-rank in-

cremental weight matrices. Specifically, consider the i -th forward $f_i(\mathbf{x}) = \mathbf{x}\mathbf{W}^\top + \mathbf{b}_i$ where $\mathbf{W}_i \in \mathbb{R}^{d_1 \times d_2}$, and $\mathbf{b}_i \in \mathbb{R}^{d_2}$ denotes the frozen weight and bias, as shown in figure 1 (a). LoRA constrains its updates by representing them as a low-rank decomposition, specifically, the forward is modified as:

$$f_i(\mathbf{x}) = \mathbf{x}(\mathbf{W}_i + \Delta \mathbf{W}_i)^\top + \mathbf{b}_i \Delta \mathbf{W}_i = \mathbf{A}_i \mathbf{B}_i \quad (1)$$

where $\mathbf{A}_i \in \mathbb{R}^{d_1 \times r}$, $\mathbf{B}_i \in \mathbb{R}^{r \times d_2}$, and the rank $r \ll \min\{d_1, d_2\}$, \mathbf{W}_i remains fixed during training and does not receive gradient updates, whereas \mathbf{A}_i and \mathbf{B}_i contain trainable parameters.

2.3. Continual learning with O-LoRA

We base our method on O-LoRA [16]. O-LoRA incrementally learns new tasks in a direction orthogonal to the LoRA subspace of past tasks while fixing the previous parameters. For each task, we introduce a set of new LoRA parameters dimensionally identical to the old ones denoted as $\mathbf{A}_t, \mathbf{B}_t$, where $\mathbf{A}_t \in \mathbb{R}^{d_1 \times r}$, $\mathbf{B}_t \in \mathbb{R}^{r \times d_2}$ and $r \ll \min\{d_1, d_2\}$. We approximate the parameter update subspace \mathcal{U}_t for the t -th task as the subspace spanned by the column vectors of \mathbf{A}_t , as shown in figure 1 (b). Mathematically, to ensure the orthogonality between the subspace \mathcal{U} and the subspace \mathcal{W} , we need to satisfy:

$$\langle u, w \rangle = 0, \forall u \in \mathcal{U}, \forall w \in \mathcal{W} \quad (2)$$

Therefore, achieving orthogonality between the LoRA subspaces of task i (\mathcal{U}_i) and task t (\mathcal{U}_t) can be expressed as:

$$O_{i,t} = \mathbf{A}_i^\top \mathbf{A}_t = 0 \quad (3)$$

We add an orthogonal loss to enforce this orthogonality:

$$L_{orth}(\mathbf{A}_i, \mathbf{A}_t) = \sum_{j,k} \|O_{i,t}[j, k]\|^2 \quad (4)$$

where $O_{i,t}[j, k]$ denotes the element at the j -th row and k -th column of $O_{i,t}$. Finally our training objective is defined as:

$$L = L_{whisper} + \lambda_1 L_{orth}(\mathbf{A}_i, \mathbf{A}_t) \quad (5)$$

where λ_1 is the weight of the orthogonality loss. During the training process, we fix the previous LoRA parameters $\{\mathbf{A}_i, \mathbf{B}_i | i < t\}$ to mitigate forgetting of past knowledge.

2.4. Orthogonal adaptive low-rank adaptation (O-AdaLoRA)

The importance of weight parameters to the performance varies across different layers and modules. Intuitively, some weight matrices should be prioritized higher in the tuning process. However, the traditional LoRA approach takes a one-size-fits-all strategy by assigning a uniform rank to all weight matrices, which does not account for the possibility that the relative importance of weights may change depending on the context, potentially leading to suboptimal performance. To address

Table 1: *Datasets used in this study.*

Language	Source Dataset	Hours of training audio
Chinese	Wenetspeech, Aishell-1	10000
Uyghur	CommonVoice, Thuyg20, In-house Uyghur Dataset	250
Tibetan	XBMU-AMDO31, TIBMD, In-house Tibetan Dataset	215

this, the AdaLoRA [17] (Adaptive Low-Rank Adaptation) introduces a dynamic rank allocation method that is sensitive to the importance of weights. AdaLoRA makes two key improvements to the original LoRA mechanism: Firstly, it employs a parameterization strategy based on Singular Value Decomposition (SVD) to represent the incremental update term $\Delta \mathbf{W}_j$ as:

$$\Delta \mathbf{W}_j = \mathbf{A}_j \mathbf{\Lambda}_j \mathbf{B}_j; \mathbf{A}_j^\top \mathbf{A}_j = \mathbf{B}_j \mathbf{B}_j^\top = \mathbf{I} \quad (6)$$

where $\mathbf{A}_j \in \mathbb{R}^{d_1 \times r}$ and $\mathbf{B}_j \in \mathbb{R}^{r \times d_2}$ contain the left/right singular vectors of $\Delta \mathbf{W}_j$. \mathbf{I} denotes the identity matrix. and the $\mathbf{A}_j, \mathbf{B}_j$ are enforced to be orthogonal matrices with a regularization term:

$$L_{adaloRA}(\mathbf{A}_j, \mathbf{B}_j) = \|\mathbf{A}_j^\top \mathbf{A}_j - \mathbf{I}\|^2 + \|\mathbf{B}_j \mathbf{B}_j^\top - \mathbf{I}\|^2 \quad (7)$$

The diagonal matrix $\mathbf{\Lambda}_j \in \mathbb{R}^{r \times r}$ contains r singular values, with $r \ll \min\{d_1, d_2\}$. Only $\mathbf{A}_j, \mathbf{\Lambda}_j, \mathbf{B}_j$ are trainable, as shown in figure 1 (c). Second, in adaptation, AdaLoRA dynamically allocates an overall rank budget to its update matrices $\Delta \mathbf{W}_j$. This is achieved by iteratively masking out less important singular values after every gradient update step. A sensitivity-based importance metric is utilized to measure and sort the importance of the k -th triplet $\mathbf{\Lambda}_j^{k,k}, \mathbf{B}_j^{*,k}, \mathbf{A}_j^{k,*}$ of the j -th weight matrix \mathbf{W}_j , which takes account of both singular values and vectors. The non-zero $\mathbf{\Lambda}_j^{k,k}$ acts as the rank coefficient to control the allocated rank budget.

To learn new tasks in a direction orthogonal to the past tasks, we just need to ensure the orthogonality between the subspace of j -th task \mathcal{U}_j and the subspace of t -th task \mathcal{U}_t . Like section 2.3, we add a new set of $\{\mathbf{A}_t, \mathbf{B}_t, \mathbf{\Lambda}_t\}$ as shown in figure 1 (d), and enforce the orthogonality with:

$$\mathbf{O}_{j,t} = \mathbf{A}_j^\top \mathbf{A}_t = 0 \quad (8)$$

Similar to equation (5), our training objective is defined as:

$$L = L_{whisper} + \lambda_1 L_{orth}(\mathbf{A}_j, \mathbf{A}_t) + \lambda_2 L_{adaloRA}(\mathbf{A}_t, \mathbf{B}_t) \quad (9)$$

where λ_1 and λ_2 are regularization coefficients. Since combining all $\{\mathbf{A}_j, \mathbf{B}_j, \mathbf{\Lambda}_j | j < t\}$ before the forward operation is mathematically equivalent to computing each one individually, the additional computational overhead is minimal. The additional memory usage can be mitigated by merging the updates corresponding to the LoRA parameters into the initial ones.

3. Experiments

3.1. Data

In our study, we trained speech recognition models for Chinese, Uyghur, and Tibetan using a variety of datasets. For Chinese, we use the WenetSpeech [18] and Aishell-1 [19] datasets for training. For testing, we use the Aishell-1 test set and an in-house Chinese test set from a voice assistant product with more complex acoustic environments. For Uyghur, we use the Common Voice Dataset 15.0 [20] and THUYG-20 [21], using the test set from THUYG-20 for evaluation. As for Tibetan, we use the TIBMD [22] and XBMU-AMDO31 datasets [23], and we use TIBMD’s test set for testing. We did not differentiate

between the dialects in the TIBMD dataset. For mixed dataset training, the mixture is weighted according to the number of elements in each component while we down-weighted larger datasets. Due to the limited availability of open-source datasets for Tibetan and Uyghur, we add approximately 100 hours of in-house data for each language. A detailed list of our data sources for each language can be found in the referenced table 1.

We noticed that when processing Uyghur and Tibetan, Whisper output words in a Romanized form—written in the Latin alphabet, instead of their original characters. We hypothesize that Whisper’s training included Romanized versions of these languages. To keep our work consistent with Whisper’s training, we use the uroman tool [24] to Romanize our transcripts before training and testing. Also, we stripped out all punctuation from these transcripts, and we discounted any punctuation that appeared in the output when we measured our results. Since Whisper doesn’t include a specific language ID token for Uyghur, we use the language label for Uzbek, which is linguistically similar.

3.2. Experimental setups

All our experiments were performed using the Parameter-Efficient Fine-Tuning (PEFT) library in Hugging Face. To assess the performance with different parameter sizes, we investigated two sizes of the Whisper model: the base model with 74 M parameters and the large model with 1,550 M parameters, we use the large-v2 model unless specified otherwise. For PEFT, we apply low-rank updates to a set of matrices: $W_q, W_k, W_v, W_o, W_{fc1}, W_{fc2}$. Following the parameter config of [25], for LoRA and O-LoRA, we use a rank $r = 32$ for matrix updates. For AdaLoRA and O-AdaLoRA, we set the initial rank $r = 12$ and a target rank of 8. Inspired by [26], we set the learning rate for the PEFT model at $1e-3$ for the initial task and reduced it to $1e-4$ for subsequent tasks. For full-model finetuning, we use a learning rate of $1e-5$. Whisper’s architecture pads audio inputs to 30 seconds, which significantly increases memory usage, leading us to adopt a $batch_size = 2$, with a gradient accumulation of 128 steps. The orthogonal loss coefficient $\lambda_{1,2} = 0.5$ is used. For evaluation metrics, we measure performance in Chinese using the Character Error Rate (CER) and for the Romanized versions of Tibetan and Uyghur, we use the Word Error Rate (WER).

3.3. Comparison with other methods

We evaluate O-LoRA and O-AdaLoRA against 5 baseline methods. Since we are investigating methods that are rehearsal-free and task-id-free, rehearsal-based methods like A-GEM [14], and adapter-based methods [13] that require task-ids were not included. Since we found AdaloRA fine-tuning performs the best, we use AdaLoRA for Multilingual fine-tuning and Monolingual finetuning as topline. We compared with the following methods: (1) **SeqFT** (Sequential Finetuning): Full fine-tuning on all tasks in sequence. (2) **SeqLoRA** (Sequential LoRA Finetuning): Sequentially training a fixed LoRA matrix with data from different languages, the LoRA matrix is merged after every language. (3) **LwF** (Learning without Forgetting): Regularization-based method that constrains the shared representation layer to be similar to its original state before learning the new task. LwF was the best memory-free and task-id-free continual learning method from [27, 28, 13]. (4) **Multi** (Multilingual finetuning): Training with a mixed dataset with all three language datasets with AdaLoRA simultaneously. (5) **Mono** (Monolingual finetuning): Finetuning separate whisper models

Table 2: Performance (WER/CER%) after Uyghur finetuning, N/A indicates the model can not transcribe in the correct language. A WER larger than 100 is still in the correct language.

	original-v2		full-ft		LoRA		AdaLoRA		AdaLoRA +Freeze encoder	
	base	large	base	large	base	large	base	large	base	large
Aishell-1	22.20	8.82	N/A	30.87	N/A	24.31	N/A	28.45	N/A	25.40
In-house Mandarin	89.32	47.94	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Thuyg-20	147.72	99.92	51.74	41.17	44.68	39.45	37.76	33.63	38.32	35.89

Table 3: Performance (CER%) of Whisper large-v2 finetuned on Chinese dataset.

	original large-v2	full-ft	LoRA	AdaLoRA
Aishell-1	8.81	5.72	4.23	3.78
In house Chinese	47.94	40.57	36.98	34.43
#Trainable Params	-	100%	3.47%	1.36%




Figure 2: Performance (WER/CER%) of different methods after transfer between languages

with AdaLoRA using datasets from individual languages.

3.4. Limitation of vanilla fine-tuning

We first investigate the possibility of adding a new language (Uyghur) to Whisper across various sizes with baseline approaches: full finetuning, LoRA, AdaLoRA, and freeze encoder approach seen in LLM adaptations. We finetuned Whisper models of different capacities with a dataset composed of Uyghur speech. For this experiment, our language of concern is Chinese and Uyghur. We measure the performance gain in Uyghur with Thuyg-20 while measuring the performance degradation in Chinese with Aishell-1. The findings are detailed in table 2. The smaller Whisper models showed a lower tolerance to catastrophic forgetting, as the base model recognized the Chinese Aishell-1 dataset as Uyghur. A more challenging dataset is found to be more sensitive to catastrophic forgetting, as the Whisper model suffering from catastrophic forgetting can still transcribe Aishell-1 into Chinese, yet the in-house Chinese dataset is being transcribed as Uyghur. While the freeze encoder strategy has the benefit of achieving a similar WER with fewer trainable parameters, it did not have a significant impact on the catastrophic forgetting phenomenon experienced when adding a new language.

3.5. Effectiveness of orthogonal subspace learning

Next, we assessed the effectiveness of orthogonal methods. Unfortunately, since the original Whisper was not trained using

LoRA, it does not retain LoRA matrices, and therefore we cannot directly apply our method to protect all languages in Whisper during subsequent fine-tuning. As a compromise, we procure a Chinese model trained using LoRA. We trained a model using the Chinese dataset discussed in Section 3.1 as our initial model M_0 . The result is shown in table 3 and figure 2. Despite having the most trainable parameters, we found that the full-model fine-tuning approach yielded the poorest performance. We speculate this may be due to that our dataset was significantly smaller than the 680,000 hours of data used to train Whisper thus the model struggled to converge. The LoRA model showed slightly inferior results compared to AdaLoRA. We hypothesize that this could be related to the findings in [25], where it’s suggested that W_v, W_o should have more trainable parameters than W_q, W_k during the finetuning.

We conducted further research into the phenomena of catastrophic forgetting using this Chinese-finetuned Whisper-large model with sequential language training. After initially training on Chinese data for 20,000 steps, we subsequently incorporated Uyghur and Tibetan datasets for an additional 5,000 steps each, using various methods. Figure 2 shows the changes in WER/CER.

Although Whisper’s paper [3] suggested that mixed language training may lead to improvements over single-language performance, we found that the three monolingual models performed better than the multi-language model. We suspect this discrepancy arises because our dataset was not as large as Whisper’s and the LoRA matrix did not learn the ability to integrate knowledge across languages. The orthogonal gradient-based approach O-LoRA significantly reduced the catastrophic forgetting issue, with the SeqFT and SeqLoRA having a final Chinese WER after Tibetan fine-tuning of over 30%. And outperformed the SOTA method LwF that did not require review or task-id, with a lower final Chinese WER (22.82% vs 12.35%). Our O-AdaLoRA sped up model convergence and achieved lower WER (53.62% vs 50.45%) in the Tibetan fine-tuned result, with a smaller parameter size (3.47% vs 1.36%) while having a similar performance retaining the previous tasks.

4. Conclusions

This paper takes the first step in investigating fine-tuning Whisper from a continual learning point-of-view and using orthogonal gradient descent to adapt the pre-trained Whisper model for Uyghur and Tibetan languages. We investigated various fine-tuning methods and examined their susceptibility to catastrophic forgetting. We have demonstrated that the LoRA matrices of a model can be used to guide subsequent continual learning in ASR. We proposed to allocate the parameters with a learnable rank coefficient to improve the overall performance. This enables the fine-tuning of foundation models that are entirely deprived of pre-training data while avoiding previously acquired language knowledge being overwritten, greatly enhancing the potential reusability of pre-trained models.

5. References

- [1] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [2] 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.
- [3] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *International Conference on Machine Learning*. PMLR, 2023, pp. 28 492–28 518.
- [4] Y. Chu, J. Xu, X. Zhou, Q. Yang, S. Zhang, Z. Yan, C. Zhou, and J. Zhou, "Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models," *arXiv preprint arXiv:2311.07919*, 2023.
- [5] 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 *International Conference on Learning Representations*, 2022. [Online]. Available: <https://openreview.net/forum?id=nZeVKeeFYf9>
- [6] A. Chaudhry, M. Rohrbach, M. Elhoseiny, T. Ajanthan, P. K. Dokania, P. H. Torr, and M. Ranzato, "On tiny episodic memories in continual learning," *arXiv preprint arXiv:1902.10486*, 2019.
- [7] D. Rolnick, A. Ahuja, J. Schwarz, T. Lillicrap, and G. Wayne, "Experience replay for continual learning," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [8] F. Zenke, B. Poole, and S. Ganguli, "Continual learning through synaptic intelligence," in *International conference on machine learning*. PMLR, 2017, pp. 3987–3995.
- [9] R. Aljundi, F. Babiloni, M. Elhoseiny, M. Rohrbach, and T. Tuytelaars, "Memory aware synapses: Learning what (not) to forget," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 139–154.
- [10] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska *et al.*, "Overcoming catastrophic forgetting in neural networks," *Proceedings of the national academy of sciences*, vol. 114, no. 13, pp. 3521–3526, 2017.
- [11] Z. Wang, Y. Liu, T. Ji, X. Wang, Y. Wu, C. Jiang, Y. Chao, Z. Han, L. Wang, X. Shao *et al.*, "Rehearsal-free continual language learning via efficient parameter isolation," in *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2023, pp. 10933–10946.
- [12] A. Razdaibiedina, Y. Mao, R. Hou, M. Khabsa, M. Lewis, and A. Almahairi, "Progressive prompts: Continual learning for language models," in *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*, 2023.
- [13] S. Vander Eeck and H. Van Hamme, "Using adapters to overcome catastrophic forgetting in end-to-end automatic speech recognition," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [14] M. Yang, I. R. Lane, and S. Watanabe, "Online continual learning of end-to-end speech recognition models," in *Interspeech 2022, 23rd Annual Conference of the International Speech Communication Association, Incheon, Korea, 18-22 September 2022*, H. Ko and J. H. L. Hansen, Eds. ISCA, 2022, pp. 2668–2672.
- [15] M. Farajtabar, N. Azizan, A. Mott, and A. Li, "Orthogonal gradient descent for continual learning," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 3762–3773.
- [16] X. Wang, T. Chen, Q. Ge, H. Xia, R. Bao, R. Zheng, Q. Zhang, T. Gui, and X. Huang, "Orthogonal subspace learning for language model continual learning," in *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, H. Bouamor, J. Pino, and K. Bali, Eds. Association for Computational Linguistics, 2023, pp. 10658–10671.
- [17] Q. Zhang, M. Chen, A. Bukharin, P. He, Y. Cheng, W. Chen, and T. Zhao, "Adaptive budget allocation for parameter-efficient fine-tuning," in *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023.
- [18] B. Zhang, H. Lv, P. Guo, Q. Shao, C. Yang, L. Xie, X. Xu, H. Bu, X. Chen, C. Zeng *et al.*, "Wenetspeech: A 10000+ hours multi-domain mandarin corpus for speech recognition," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6182–6186.
- [19] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, "Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline," in *2017 20th conference of the oriental chapter of the international coordinating committee on speech databases and speech I/O systems and assessment (O-COCOSDA)*. IEEE, 2017, pp. 1–5.
- [20] R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, "Common voice: A massively-multilingual speech corpus," in *Proceedings of The 12th Language Resources and Evaluation Conference, LREC 2020, Marseille, France, May 11-16, 2020*, N. Calzolari, F. Béchet, P. Blache, K. Choukri, C. Cieri, T. Declerck, S. Goggi, H. Isahara, B. Maegaard, J. Mariani, H. Mazo, A. Moreno, J. Odiijk, and S. Piperidis, Eds. European Language Resources Association, 2020, pp. 4218–4222. [Online]. Available: <https://aclanthology.org/2020.lrec-1.520/>
- [21] A. Roze, S. Yin, Z. Zhang, D. Wang, and A. Hamdulla, "Thugy20: A free uyghur speech database," in *NCMMSC'15*, 2015.
- [22] Y. Zhao, X. Xu, J. Yue, W. Song, X. Li, L. Wu, and Q. Ji, "An open speech resource for tibetan multi-dialect and multitask recognition," *International Journal of Computational Science and Engineering*, vol. 22, no. 2-3, pp. 297–304, 2020.
- [23] J. N. Senyan Li, Guanyu Li, "XBMU-AMDO31: An open source of Amdo Tibetan speech database and speech recognition baseline system," in *National Conference on Man-Machine Speech Communication, NCMMSC2022*, 2022.
- [24] U. Hermjakob, J. May, and K. Knight, "Out-of-the-box universal Romanization tool uroman," in *Proceedings of ACL 2018, System Demonstrations*, F. Liu and T. Solorio, Eds. Melbourne, Australia: Association for Computational Linguistics, Jul. 2018, pp. 13–18.
- [25] W. Liu, Y. Qin, Z. Peng, and T. Lee, "Sparsely shared lora on whisper for child speech recognition," *arXiv preprint arXiv:2309.11756*, 2023.
- [26] G. I. Winata, L. Xie, K. Radhakrishnan, S. Wu, X. Jin, P. Cheng, M. Kulkarni, and D. Preotiuc-Pietro, "Overcoming catastrophic forgetting in massively multilingual continual learning," in *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, A. Rogers, J. L. Boyd-Graber, and N. Okazaki, Eds. Association for Computational Linguistics, 2023, pp. 768–777.
- [27] H. Chang, H. Lee, and L. Lee, "Towards lifelong learning of end-to-end ASR," in *Interspeech 2021, 22nd Annual Conference of the International Speech Communication Association, Brno, Czechia, 30 August - 3 September 2021*, H. Hermansky, H. Cernocký, L. Burget, L. Lamel, O. Scharenborg, and P. Motlíček, Eds. ISCA, 2021, pp. 2551–2555.
- [28] S. Vander Eeck and H. Van Hamme, "Continual learning for monolingual end-to-end automatic speech recognition," in *2022 30th European Signal Processing Conference (EUSIPCO)*. IEEE, 2022, pp. 459–463.