



Continual Learning Optimizations for Auto-regressive Decoder of Multilingual ASR systems

Chin Yuen Kwok¹, Jia Qi Yip², Eng Siong Chng²

¹Digital Trust Centre, Nanyang Technological University, Singapore

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

kwok0062@e.ntu.edu.sg, jiaqi006@e.ntu.edu.sg, aseschn@ntu.edu.sg

Abstract

Continual Learning (CL) involves fine-tuning pre-trained models with new data while maintaining the performance on the pre-trained data. This is particularly relevant for expanding multilingual ASR (MASR) capabilities. However, existing CL methods, mainly designed for computer vision and reinforcement learning tasks, often yield sub-optimal results when directly applied to MASR. We hypothesise that this is because CL of the auto-regressive decoder in the MASR model is difficult. To verify this, we propose four optimizations on the decoder. They include decoder-layer gradient surgery, freezing unused token embeddings, suppressing output of newly added tokens, and learning rate re-scaling. Our experiments on adapting Whisper to 10 unseen languages from the Common Voice dataset demonstrate that these optimizations reduce the Average Word Error Rate (AWER) of pretrained languages from 14.2% to 12.4% compared with Experience Replay, without compromising the AWER of new languages.

Index Terms: continual learning, task-agnostic, language-agnostic, massively multilingual, speech recognition, human-computer interaction, computational paralinguistics

1. Introduction

Continual Learning (CL) involves adapting pre-trained models to new data without catastrophic forgetting (CF), where the model forgets about previously learnt knowledge and leads to degraded performance on the pre-trained data. This is particularly relevant for expanding multilingual ASR (MASR) capabilities. However, CL for MASR is challenging.

To adapt a model to new data without CF, numerous CL methods are developed with their designs focusing on different aspects: Replay-based methods like ER [1] focus on managing the training data; Regularization-based methods like EWC [2] and MAS [3] reduce model weights mismatch; Optimization-based methods like GEM [4] and A-GEM [5] constrain gradients used to optimize models; LwF [6] and DER [7] reduce model outputs mismatch through distillation [8]; Dynamic-architecture-based [9, 10] methods add new components to the model.

As different CL methods have very different design focuses, previous works [11–13] have compared their effectiveness in the audio domain. Among these, some work specifically study CL methods to adapt MASR models to new languages. They include [14], which incrementally adds new languages and trains with all available data to mitigate forgetting, [15], which trains a mapping matrix for task-specific weight selection, and [16], which uses weight averaging (AVG). Furthermore, the work by [17] offers an extensive evaluation of baseline MASR CL methods, indicating potential areas for improvement.

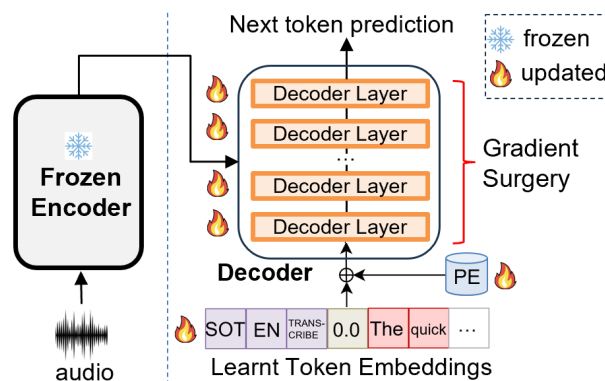


Figure 1: Overview of decoder-layer gradient surgery. For an MASR model that consists of an encoder (left) and decoder (right), we freeze the encoder and only adapt the decoder side. Also, we propose to apply gradient surgery only to the decoder layers, but not to the learnt token embeddings and learnt positional embeddings (PE).

Notably, a significant research gap is that previous CL methods rarely focus their designs on the auto-regressive decoder model architecture of an MASR model. This is because most CL methods are primarily designed for computer vision (CV) and reinforcement learning (RL) tasks [1–7] that make use of a different model architecture like the Convolutional Neural Network, and as a consequence, we show that directly applying these CL methods to MASR yield suboptimal results. We hypothesise that this is because the unique structure of the auto-regressive decoder presents unique CL issues.

They include the following: 1) The decoder uses token embeddings [18] to perform next token prediction, and the result is sensitive to the embeddings' norm and angular information [19]. Gradient-based CL methods which crudely modify these information by removing part of the gradients [20] used to update the embeddings will cause sub-optimal results. 2) Most MASR models make use of the text embedding layer at the decoder which CV and RL models do not, although our experiments show that the embedding layer contributes significantly to CF. 3) The propagation of token output errors during auto-regressive decoding necessitates a meticulously designed strategy for the output of newly added tokens for new languages, ensuring that these errors do not adversely affect the model's performance. 4) The learning rate of the decoder needs to be meticulously tuned to prevent CF.

To address the four issues above, we propose four CL optimizations tailored for the auto-regressive decoder, which includes decoder-layer gradient surgery, freezing unused token

embeddings, suppressing output of newly added tokens, and learning rate re-scaling to solve them respectively. Our contributions are four-fold. We show that 1) Averaged-GEM (A-GEM) [5], which is a gradient-based CL method, can consistently outperform replay-based methods on MASR CL tasks if we only perform gradient surgery in the decoder layers and not on the token embeddings, 2) selectively freezing part of the text embedding layer at the decoder during training can reduce forgetting, 3) suppressing the output of newly added special tokens in the middle of the transcription can reduce error propagation, and 4) it is crucial to reduce the LR of the decoder rapidly during training to reduce forgetting.

2. Methodology

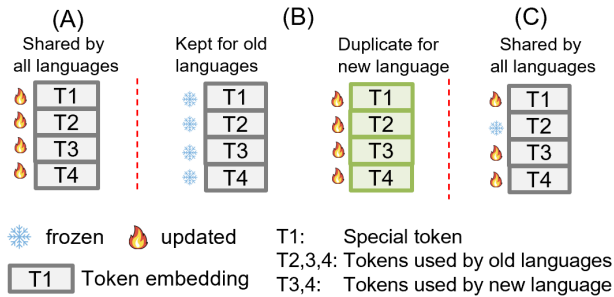


Figure 2: Strategy to adapt the token embeddings at the decoder to new tasks. All the embeddings are initialized from pre-trained weights. A) All the embeddings are shared, and updated for the old and new languages. B) A copy of the embeddings are adapted for the new language, and the original embeddings are kept for the old language. C) All the embeddings are shared, but only special tokens and tokens used by the new language are adapted.

To learn unseen languages, we adapt the Whisper decoder and freeze the encoder [17]. We train the model with A-GEM [5] with four optimizations for the decoder. First, as shown in Figure 1, the gradient surgery from A-GEM is only applied to the decoder layers and not to the token embeddings at the decoder. Second, as shown in the right part of Figure 2, we only update the embeddings of the special tokens and the tokens that appeared in the new languages. Thirdly, as shown in Figure 3, the output of special tokens at specific positions are suppressed. Lastly, we tune the Learning Rate Scheduler (LRS) to have a more rapidly decreasing LR for the decoder during adaptation.

2.1. Whisper Model for Language-agnostic MASR

We use Whisper [21] as the base model for adaptation. This model supports MASR across 75 languages, including Language Identification (LID) functionality, making it language-agnostic [22] as the user does not need to manually specify the language to transcribe. It utilizes a Transformer-based [18] attention encoder-decoder architecture with auto-regressive decoding. This model is selected because while comparative analysis reveals that Whisper’s performance in Word Error Rate (WER) averaged across both seen and unseen languages is comparable to XLS-R [23] as reported by [24], [17] indicate that Whisper outperforms WavLM [25] in overall MASR competency, and [26] demonstrates Whisper’s superiority in LID over models based on wav2vec 2.0 [27] and WavLM. Also, the Whisper decoder is preferred for MASR applications over Large Lan-

A. The	caravan	began to travel	day and night.
B. The	<[eo]><[transcribe]><[notimestamps]>Curven	began to tour the	day and night.
C. The	curve	began to travel	day and night.

Figure 3: Example of Language ID Token Suppression (LID TS). A) Reference. B) Hypothesis before LID TS. C) Hypothesis after LID TS.

guage Model (LLM) based decoders [28], particularly because Whisper decoder supports LID, a feature not available in the LLM-based decoder to the best of our knowledge.

2.2. Experience Replay and A-GEM

To mitigate CF, Experience Replay (ER) aims to remind a model of previously learnt tasks by incorporating a limited subset of data from previous tasks [1] into the new task’s dataset during adaptation. The goal is to minimize:

$$\mathbb{E}_{(x', y') \sim D_B} [\ell(y', f_\theta(x'))] + \beta * \mathbb{E}_{(x'', y'') \sim \hat{D}_A} [\ell(y'', f_\theta(x''))] \quad (1)$$

where f_θ is the model function parameterized by θ , ℓ is the classification loss, \hat{D}_A is a subset of the old task A dataset, D_B is the new task B dataset, and β is a hyper-parameter balancing the trade-off between the terms.

The approach is often combined with additional optimization constraints. Gradient Episodic Memory (GEM) [4] and Averaged-GEM (A-GEM) [5] use gradient surgery [20] to constrain gradient update to not increase old task loss. Loss gradient vectors of old task g_A and the gradient update on the new task g_B is computed from $D_{\hat{A}}$ and D_B respectively. To not increase old task loss while optimizing new task loss, the gradient \hat{g} is proposed to minimize:

$$\frac{1}{2} \|g_B - \hat{g}\|_2^2 \quad \text{s.t.} \langle \hat{g}, g_A \rangle \geq 0 \quad (2)$$

Equation 1 and 2 can be easily extended to settings with more than one old task by adding extra terms for each old task correspondingly.

2.3. Gradient Surgery Removal

Although the use of gradient surgery for CL has been proven effective in the CV field [4, 5], its effect in the audio field is not better than other CL baseline methods like ER when an auto-regressive decoder is used [13, 17]. We hypothesise that this is because the decoder uses token embeddings to perform next token prediction, and the result is sensitive to the embeddings’ norm and angular information [19]. Gradient-based CL methods which crudely modify these information by removing part of the gradients [20] used to update the embeddings will cause sub-optimal results.

Therefore, to address the issue, we propose to remove A-GEM’s gradient surgery from the token embeddings [18]. We find that this modified version of A-GEM can consistently outperform ER in our experiments.

2.4. Partial Embedding Update

Furthermore, we propose another optimization for the decoder, which is to selectively freeze part of the token embeddings to mitigate CF. This is because MASR models employ sub-word units for word representation [29], and the sub-words can be shared across languages. If the sub-word token embeddings

[30] are adapted to new languages, it can inadvertently overwrite the semantic information of previously learned languages, leading to CF.

To address this issue, as shown in the right part of Figure 2, we propose to update only the embeddings used by all the new languages and freeze the other embeddings.

2.5. Language ID Token Suppression

Moreover, another optimization is applied to suppress the decoder’s output. This is necessary as to adapt Whisper to unseen languages, new language ID tokens are added to the text embedding layer, and the tokens may lead to more errors in the ASR output. Specifically, as shown in Figure 3, we find that these ID tokens may be wrongly transcribed in the middle of the ASR output, and induce more errors in later transcriptions. Therefore, to address this, we propose suppressing the output of the tokens such that these tokens, primarily used for LID, will not interfere with the ASR results.

2.6. Rapid LR Reduction During Adaptation




Figure 4: Change of LR during training as the validation interval changes. split- n refers to validating every $1/n$ epoch.

To further mitigate CF, we explore the effect of lowering the learning rate (LR) to constrain the deviation of the adapted model from its original state, thereby retaining more previously learnt knowledge. One method to modulate LR involves adjusting the Learning Rate Scheduler (LRS). Specifically, as our experiments use an LRS analogous to ReduceLRonPlateau¹, we shorten the validation interval such that the LR can be reduced more promptly by the LRS based on the more frequent validation feedback. Figure 4 shows the effect of shortening the validation interval.

3. Experiments

3.1. Dataset and Model Details

We implement our methods based on the popular SpeechBrain [31] toolkit and CL-MASR [17].

Following previous works [15, 17], we evaluate our method using a subset of the widely used large-scale CommonVoice dataset² [32]. The subsets are extracted following [17]. They consist of 10 languages pre-trained by Whisper and 10 new languages. Each language contains a 10 hours train set, a 1 hour validation set, and a 1 hour test set. whisper-small and whisper-large-v2 are adapted in two CL settings: 1) Adapt to ten pre-trained languages sequentially and test forgetting on ten new

languages 2) Adapt to one pre-trained language and test forgetting on one new language. The former can take up to 6 days utilizing an NVIDIA A40 GPU.

We adapt the models for 2 epochs and set the train batch size to 4. Validation is performed every $1/32$ epoch for all methods. For ER, the replay data size is one hour for every old language. The AdamW optimizer [33] is used with a variant³ of the ReduceLRonPlateau learning rate (LR) scheduler. For CL baselines and our methods, we sweep through the hyper-parameters to optimize average word error rate (AWER). A greedy decoding strategy is used for the experiments.

we refer to our methods as 1) ER-M, which is ER with our decoder optimizations except gradient surgery removal, and 2) A-GEM-M, which is A-GEM with all the decoder optimizations applied. All the experiments are tuned to have a rapid LR reduction during adaptation.

3.2. Results and Discussion

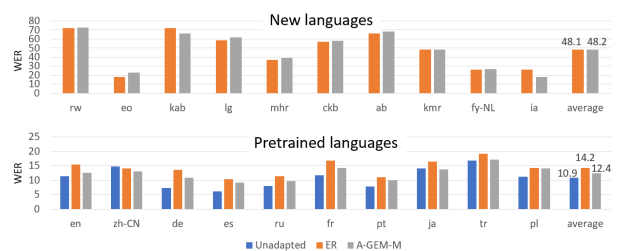


Figure 5: WER of transcribing 10 pre-trained languages and 10 new languages of varying difficulties. Results are obtained without manually specifying the language to transcribe. “Unadapted” means the unadapted model.

We show both the results of adapting to Whisper-small and Whisper-large-v2 models. Table 1 shows the results of adapting Whisper-small to two unseen languages separately. Vanilla full fine-tune (FT) can enable Esperanto and Interlingua ASR for Whisper, but leads to CF. All CL baselines reduce forgetting significantly and improve upon FT result by 27.2% – 66.9% relatively. Among these, ER has the best overall performance, although EWC and MAS has slightly better partial results than ER for English and Esperanto in the language-aware setting. Weight averaging (AVG) performs less well than ER as we did not use a task-specific layer as in [16].

Language-agnostic results are generally worse than language-aware results as the errors in LID may induce more error in transcription. Among the CL baselines, LwF, EWC, MAS and DER do not reduce forgetting as much as A-GEM and ER. We hypothesize that this is because the distillation of token-level logits for LwF and DER has no effect on the newly added language ID tokens. Similarly, the model weights’ importance computed by EWC and MAS does not include the newly added language ID embedding weights, so its update is not regularized. As these methods fail to regularize the language ID tokens, they may cause more LID error, which in turn leads to more transcription errors. Additionally, we observe that when compared with language-aware results, the increase of WER in language-agnostic results mainly occurs in the old languages English and German. This suggests that the adapted Whisper

¹https://pytorch.org/docs/stable/generated/torch.optim.lr_scheduler.ReduceLRonPlateau.html

²<https://commonvoice.mozilla.org/en>

³https://speechbrain.readthedocs.io/en/latest/_modules/speechbrain/nnet/schedulers.html#NewBobScheduler

Table 1: WER of adapting Whisper-small to new language Interlingua (ia) and Esperanto (eo), and test forgetting on the pretrained language German (de) and English (en) respectively. Language-agnostic [22] means the user does not need to specify the language to transcribe, and vice-versa for Language-aware. “None” method refers to not performing adaptation.

Method	Language-aware WER (%)					Language-agnostic WER (%)				
	de	en	ia	eo	avg	de	en	ia	eo	avg
None	14.00	14.57	n/a	n/a	n/a	14.00	14.57	n/a	n/a	n/a
FT	64.83	68.99	12.31	18.04	41.0 _{-00.0%}	88.40	90.55	12.31	18.04	52.3 _{-00.0%}
<i>CL baselines</i>										
AVG [16]	17.94	16.20	16.53	38.96	22.4 _{-45.4%}	26.55	18.37	16.53	40.00	25.4 _{-51.4%}
LwF [6]	16.81	16.04	14.82	22.20	17.5 _{-57.3%}	40.07	74.96	14.82	22.18	38.1 _{-27.2%}
EWC [2]	23.78	15.62	13.31	19.21	18.0 _{-56.1%}	41.12	39.87	13.34	19.21	28.4 _{-45.7%}
MAS [3]	18.29	15.64	12.42	20.14	16.7 _{-59.3%}	21.97	88.91	12.46	20.14	35.9 _{-31.4%}
A-GEM [5]	18.10	18.15	12.42	18.74	16.9 _{-58.9%}	18.68	19.09	12.68	18.91	17.3 _{-66.9%}
DER [7]	16.31	17.13	13.66	19.54	16.7 _{-59.3%}	23.59	56.95	13.66	20.25	28.6 _{-45.3%}
ER [1]	16.65	15.64	12.94	20.34	16.4 _{-60.0%}	17.20	18.26	13.14	20.45	17.3 _{-66.9%}
<i>our methods</i>										
ER-M	14.56	14.96	14.12	19.88	15.9 _{-61.2%}	15.13	16.99	14.45	20.06	16.7 _{-68.1%}
A-GEM-M	15.37	15.14	12.77	18.39	15.4 _{-62.4%}	15.85	15.96	13.18	18.56	15.9 _{-69.6%}

model always bias its LID result of the input audios towards one of the newly adapted languages whenever it is uncertain.

Additionally, our methods outperform all CL baselines and improve the FT result by 61.2 – 69.6%. Forgetting is further reduced compared to all CL baselines while we can maintain similar WER for the newly adapted languages. Lastly in Figure 5, we plot the results of adapting whisper-large-v2 to the ten unseen language of various difficulties and test forgetting on the ten pretrained languages. Our method A-GEM-M outperforms the best CL baseline ER and reduce the Average Word Error Rate (AWER) of pretrained languages from 14.2% to 12.4% without compromising the AWER of new languages. We did not compare Dynamic-architecture-based methods [9, 10] as they are not language-agnostic.

3.3. Ablation study

We further perform ablation studies and show the results in Table 2. For the experiments, we adapt Whisper-small to the unseen language Esperanto (eo) and test forgetting on the pretrained language English (en). The results show that Gradient Surgery Removal reduces English WER from 18.15% to 16.47%, showing that removing the gradient surgery in the text embedding layer reduces forgetting.

For Partial Embedding Update, it improves English WER from 15.46% to 14.96%, showing that freezing tokens not used by the new language Esperanto reduces forgetting. We further try to naively train a completely separate text embedding layer at the decoder exclusively for Esperanto, and keep the original embedding layer for English. Results show that it further reduces WER slightly. However, the naive approach is not language-agnostic and requires more memory to store a separate embedding layer.

For Language ID Token Suppression, it has no effect on English WER when no adaptation is performed. However, after adapting a model to the unseen language, it reduces English WER from 15.57% to 14.48% as it prevents the model from wrongly producing the newly added language ID token as the output, as shown in Figure 3.

Lastly when we shorten the validation interval during training, it reduces English WER from 19.68% to 15.64% with not

much effect on Esperanto WER. It shows that although tuning the validation interval is not crucial for the adapted language WER, it can reduce forgetting in the pre-trained languages by reducing learning rate more promptly whenever the validation loss is not decreasing, as shown in Figure 4.

Table 2: Ablation study of Language ID Token Suppression, Partial Embedding Update, Gradient Surgery Removal and Shorten Validation Interval. “None” method refers to not performing adaptation.

Method	WER (%)		
	en	eo	avg
<i>Gradient Surgery Removal</i>			
A-GEM	18.15	18.74	18.4
+ rm emb gradient surgery	16.47	18.32	17.4
<i>Partial Embedding Update</i>			
ER	19.68	20.53	20.1
+ suppress token	15.46	20.34	17.9
+ update partial emb	14.96	19.88	17.5
+ task-wise full emb	14.48	20.32	17.4
<i>Language ID Token Suppression</i>			
None	14.57	n/a	n/a
+ suppress token	14.57	n/a	n/a
ER	19.68	20.53	20.1
+ update partial emb	15.57	20.42	18.0
+ suppress token	14.48	20.32	17.4
<i>Shorten Validation Interval</i>			
ER + val. every epoch	19.68	20.53	20.1
+ val. every 1/32 epoch	15.64	20.34	18.0

4. Conclusion

To conclude, this paper provides four optimizations to adapt multilingual auto-regressive decoder ASR models in a continuous learning framework to new languages and ablation study has shown the effectiveness of the proposed methods.

5. Acknowledgement

This research is supported by the National Research Foundation, Singapore and Infocomm Media Development Authority under its Trust Tech Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore and Infocomm Media Development Authority.

6. References

- [1] D. Brignac, N. Lobo, and A. Mahalanobis, "Improving replay sample selection and storage for less forgetting in continual learning," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 3540–3549.
- [2] 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.
- [3] 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.
- [4] D. Lopez-Paz and M. Ranzato, "Gradient episodic memory for continual learning," *Advances in neural information processing systems*, vol. 30, 2017.
- [5] A. Chaudhry, M. Ranzato, M. Rohrbach, and M. Elhoseiny, "Efficient lifelong learning with a-gem," *arXiv preprint arXiv:1812.00420*, 2018.
- [6] Z. Li and D. Hoiem, "Learning without forgetting," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 12, pp. 2935–2947, 2017.
- [7] P. Buzzega, M. Boschini, A. Porrello, D. Abati, and S. Calderara, "Dark experience for general continual learning: a strong, simple baseline," *Advances in neural information processing systems*, vol. 33, pp. 15 920–15 930, 2020.
- [8] D.-T. Truong, R. Tao, J. Q. Yip, K. A. Lee, and E. S. Chng, "Emphasized non-target speaker knowledge in knowledge distillation for automatic speaker verification," in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 10 336–10 340.
- [9] A. Mallya, D. Davis, and S. Lazebnik, "Piggyback: Adapting a single network to multiple tasks by learning to mask weights," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 67–82.
- [10] A. A. Rusu, N. C. Rabinowitz, G. Desjardins, H. Soyer, J. Kirkpatrick, K. Kavukcuoglu, R. Pascanu, and R. Hadsell, "Progressive neural networks," *arXiv preprint arXiv:1606.04671*, 2016.
- [11] 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.
- [12] 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.
- [13] S. Vander Eeck *et al.*, "Rehearsal-free online continual learning for automatic speech recognition," *arXiv e-prints*, pp. arXiv–2306, 2023.
- [14] B. Li, R. Pang, Y. Zhang, T. N. Sainath, T. Strohmaier, P. Haghani, Y. Zhu, B. Farris, N. Gaur, and M. Prasad, "Massively multilingual asr: A lifelong learning solution," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6397–6401.
- [15] Z. Yu, Y. Zhang, K. Qian, C. Wan, Y. Fu, Y. Zhang, and Y. C. Lin, "Master-asr: achieving multilingual scalability and low-resource adaptation in asr with modular learning," in *International Conference on Machine Learning*. PMLR, 2023, pp. 40 475–40 487.
- [16] S. V. Eeck *et al.*, "Weight averaging: A simple yet effective method to overcome catastrophic forgetting in automatic speech recognition," *arXiv preprint arXiv:2210.15282*, 2022.
- [17] L. Della Libera, P. Mousavi, S. Zaiem, C. Subakan, and M. Ravanelli, "Cl-masr: A continual learning benchmark for multilingual asr," *arXiv preprint arXiv:2310.16931*, 2023.
- [18] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [19] D. Demeter, G. Kimmel, and D. Downey, "Stolen probability: A structural weakness of neural language models," *arXiv preprint arXiv:2005.02433*, 2020.
- [20] T. Yu, S. Kumar, A. Gupta, S. Levine, K. Hausman, and C. Finn, "Gradient surgery for multi-task learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 5824–5836, 2020.
- [21] 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.
- [22] A. Datta, B. Ramabhadran, J. Emond, A. Kannan, and B. Roark, "Language-agnostic multilingual modeling," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 8239–8243.
- [23] A. Conneau, A. Baevski, R. Collobert, A. Mohamed, and M. Auli, "Unsupervised cross-lingual representation learning for speech recognition," *arXiv preprint arXiv:2006.13979*, 2020.
- [24] A. Rouditchenko, S. Khurana, S. Thomas, R. Feris, L. Karlinsky, H. Kuehne, D. Harwath, B. Kingsbury, and J. Glass, "Comparison of multilingual self-supervised and weakly-supervised speech pre-training for adaptation to unseen languages," *arXiv preprint arXiv:2305.12606*, 2023.
- [25] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [26] K. Praveen, B. Radhakrishnan, K. Sabu, A. Pandey, and M. Shaik, "Language identification networks for multilingual everyday recordings," 2023.
- [27] A. Babu, C. Wang, A. Tjandra, K. Lakhotia, Q. Xu, N. Goyal, K. Singh, P. von Platen, Y. Saraf, J. Pino *et al.*, "Xls-r: Self-supervised cross-lingual speech representation learning at scale," *arXiv preprint arXiv:2111.09296*, 2021.
- [28] 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," *arXiv preprint arXiv:2310.13289*, 2023.
- [29] R. Sennrich, B. Haddow, and A. Birch, "Neural machine translation of rare words with subword units," *arXiv preprint arXiv:1508.07909*, 2015.
- [30] K. Ethayarajh, "How contextual are contextualized word representations? comparing the geometry of bert, elmo, and gpt-2 embeddings," *arXiv preprint arXiv:1909.00512*, 2019.
- [31] M. Ravanelli, T. Parcollet, P. Plantinga, A. Rouhe, S. Cornell, L. Lugosch, C. Subakan, N. Dawlatabad, A. Heba, J. Zhong *et al.*, "Speechbrain: A general-purpose speech toolkit," *arXiv preprint arXiv:2106.04624*, 2021.
- [32] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, "Common voice: A massively-multilingual speech corpus," *arXiv preprint arXiv:1912.06670*, 2019.
- [33] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.