

# Continuously Learning New Words in Automatic Speech Recognition

Christian Huber<sup>1</sup>

Alexander Waibel<sup>1,2</sup>

<sup>1</sup>Interactive Systems Lab, Karlsruhe Institute of Technology, Karlsruhe, Germany

<sup>2</sup>Carnegie Mellon University, Pittsburgh PA, USA

christian.huber@kit.edu, waibel@cs.cmu.edu

**Abstract**—Despite recent advances, Automatic Speech Recognition (ASR) systems are still far from perfect. Typical errors include acronyms, named entities, and domain-specific special words for which little or no labeled data is available. To address the problem of recognizing these words, we propose a self-supervised continual learning approach: Given the audio of a lecture talk with the corresponding slides, we bias the model towards decoding new words from the slides by using a memory-enhanced ASR model from the literature. Then, we perform inference on the talk, collecting utterances that contain detected new words into an adaptation data set. Continual learning is then performed by training adaptation weights added to the model on this data set. The whole procedure is iterated for many talks. We show that with this approach, we obtain increasing performance on the new words when they occur more frequently (>80 % recall) while preserving the general performance of the model.

**Index Terms**—automatic speech recognition, new-word learning, continual learning, self-supervised learning.

## I. INTRODUCTION

In the last decade, ASR systems improved tremendously in terms of word error rate (WER) due to more data, more computing power, and better architectures [1], [2]. However, these systems are still far from perfect. Although in principal end-to-end systems are open-vocabulary systems, when using appropriate modeling units, such as byte-pair encoded (BPE) characters, in practice, words not seen during training are often not reliably recognized. Typical errors are observed in the categories of cross-lingual words (e.g. 'upgeloaded' is mixing German and English), numbers (e.g. 1945 vs. 19:45 vs. \$19.45), acronyms (e.g. ICASSP), named entities and domain-specific special words (as they occur in specialized meetings or lectures). The word error rate (WER) is only slightly affected by these errors because they are rather infrequent. However, these words are important for understanding the content as they carry a lot of information that is lost when they are misrecognized. Furthermore, proper interpretation of these words is critical for downstream processing such as in speech translation [3], [4]. To measure this, we must not only evaluate WER but also recall, precision, and F1 score for these words, i.e., how often are they recognized if they occur and how often do they produce false positives.

In this work, we tackle the problem of learning such acronyms, named entities, and domain-specific special words from scarce data in a self-supervised manner. For this, we

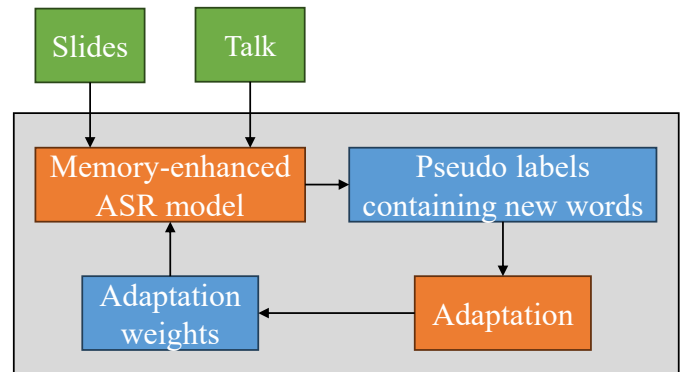


Fig. 1. Illustration of the continual learning: In each learning cycle the model is biased towards the new words from the slides of the current talk, inference is performed and pseudo-labels containing the new words are collected; then the model is adapted.

1) show that we can adapt an ASR model to detect new words with little given labeled data using a factorization-based approach, and 2) use this factorization-based approach in combination with a memory-enhanced ASR model and slides of lecture talks to perform self-supervised continual learning (see Figure 1 and Section III-B). We empirically show that this approach does not lead to catastrophic forgetting even for a large number of learning cycles (66), while improving the recall of new words to more than 80% as new words occur more frequently.

## II. RELATED WORK

An approach to new-word learning of ASR systems is context biasing [5]–[10]. The model is biased with a list of words. Depending on the approach, this list is used in different ways. The problem is that the scale of such approaches for continual learning is limited.

Other approaches study the continual learning of ASR systems [11]–[15]. The access to old data is restricted and multiple regularization-based and data-based methods [16] are used. The problem is that many hours of labeled adaptation data are used for each learning cycle, without considering where to get such data, and/or only a few learning cycles are considered.

As stated, many works in continual learning study the learning of new tasks while limiting the access to the old data. However, access to old data is not always a problem in production, e.g., when publicly available data sets are used

as in our case, but the amount of computation power and particularly the availability of new (labeled) data. It is not feasible to train a new model from scratch each time a new word has to be learned. The same holds for collecting hundreds of labeled utterances for each new word, because of cost and privacy issues. Therefore, we study how we can adapt an ASR model to detect new words with modest computation power and little labeled data or data collected in a self-supervised manner.

### III. EXPERIMENTS

The experiments consist of two parts: First, we look at how to learn new words from little given labeled data. For this, we experimented with a factorization approach [17], [18]: To each weight  $W \in \mathbb{R}^{n \times m}$  of a linear layer of the neural network, a low-rank matrix

$$\overline{W} = \sum_{i=1}^k r_i s_i^T$$

is added, where  $k \in \mathbb{N}$ ,  $r_i \in \mathbb{R}^n$ ,  $s_i \in \mathbb{R}^m$ ,  $i \in \{1, \dots, k\}$ . By changing only the weights  $\overline{W}$ , one can efficiently adapt the model. The hyperparameter  $k$ , typically  $k \ll \min\{n, m\}$ , controls the amount of flexibility of the model, and  $r_i$  and  $s_i$  are denoted as factorization weights.

Second, we perform a continual learning experiment (see Figure 1 and the second part of Section III-B). We use the factorization approach mentioned above together with a memory-enhanced ASR model similar to [5] to first extract pseudo-labels of utterances containing new words and then adapt the model with this new data. The memory-enhanced ASR model consists of a memory

$$Z = (Z_1, \dots, Z_L), \quad L \in \mathbb{N},$$

where each memory entry  $Z_l$ ,  $l \in \{1, \dots, L\}$ , is (the tokenized representation of) a word or a short phrase. The model extracts information from  $Z$  through attention mechanisms, and therefore the prediction of the next token is biased towards the words/phrases in the memory  $Z$ .

#### A. Data

For the experiments we created two data sets. First, we created a new-words data set for the factorization approach experiments based on [19]. We extracted words from the categories named entity (of persons), acronym (abbreviations), and domain-specific special word (products, events, laws, locations, and organizations) from the annotations. Then, we manually filtered those that occurred two or more times. From that, we created a development split (15% of the data, for hyperparameter optimization) and an evaluation split. Both contain a training set (for model adaptation) and a test set (for model evaluation). For each new word, one occurrence is used for the test set and the remaining ones for the training set. For most new words, we have only one training example, for a few examples more.

Second, we downloaded videos of computer science lectures with corresponding slides for the continual learning experiment from YouTube. The data set collected consists of 66

lecture talks. From the slides (given as PDF file) we extracted words that were not in our training data and manually went over the output to filter for obvious errors, e.g., some words next to each other were merged by the text extraction process, but we did not modify the extracted words by content. This resulted in 3891 extracted words (2199 unique words) that were denoted new words.

To evaluate the overall performance of our models, we report the WER on the Tedlium testset (1155 utterances, 2.6 hours).

#### B. Models and Approaches

We use the framework NMTGMinor which is based on PyTorch and uses the Fairseq [20] pre-trained models. For the baseline model, similar to recent work, for example [21], we start with a transformer model [1], [2], where the encoder is initialized with the pretrained Wav2Vec 2.0 model [22] and the decoder is initialized with the decoder part of the pretrained mBART 50 model [23] (using the mBART 50 tokenizer).

We trained the baseline model on ASR data from Common voice, Europarl, Fisher, HOW2, Hub4, Librispeech, Switchboard, Tedlium, and Voxforge (5433 hours in total), denoted baseline data. For some of the mentioned corpora, only lowercase labels without punctuation were available. Therefore, we trained an encoder-decoder model on Wikipedia text to map lowercase text without punctuation to cased text with punctuation and applied that model to the data.

When training the factorization approach, we freeze all weights of the baseline model and only train the added factorization weights until the validation perplexity on the development split does not decrease anymore. To avoid catastrophic forgetting [24], we train the model with a mixture of baseline data and new-words data. To emphasize that the model should learn the new words, we massively upsample utterances containing them [25]. Only training the factorization weights significantly decreases training time compared to training a model from scratch or from a pre-trained model. The model is trained within a few hundred updates (compared to the baseline model which is trained for 60k updates). Each factorization approach was trained in less than four hours on one A6000 GPU. On the other hand, the baseline model needed 60k updates to finish training.

The memory-enhanced ASR model [5] consists of the audio encoder and decoder of the baseline model and additionally has a memory encoder and a memory decoder (both with six layers and initialized with parts of the pre-trained mBART 50 model). Each memory entry is encoded independently by the memory encoder. Each memory decoder layer consists of the standard mBART decoder layer followed by a memory-attention layer and a memory-entry-attention layer extracting information from the encoded memory entries (see [5] for more details). During training of the memory-enhanced ASR model, only the memory encoder and memory decoder are trained, the rest of the model is frozen.

In the continual learning experiment (see Figure 1), we start with an empty data set of new-words adaptation data.

We iterate over the given talks and do for each talk the following four steps: **1)** We use the new words extracted from the corresponding slides as memory to bias the model toward the new words, **2)** We perform inference on the talk using the model from the last iteration step (using the baseline model in the first iteration step), **3)** We extract all utterances from the talk containing a word from the current memory or any past memory and add them together with the pseudo-labels created to the new words adaptation data. Using utterances containing a word from a past memory is critical for continual learning, since we want to use the collected data to increase the performance of new words even when they are no longer present in the memory. All utterances in which a new word occurs the second time are used for the new words validation data, and **4)** We use the factorization approach applied to the new-words adaptation data as well as the baseline data to adapt the model. Especially, we train new factorization weights initialized randomly. After the training, we use the best model according to perplexity on the new words validation data. These four steps are iterated for all given talks. We denote this approach by MEM-CL. Note that it is not a severe problem if a new word is not correctly recognized by the memory-enhanced ASR model since an important new word most likely will occur multiple times, and therefore the memory-enhanced ASR model can also detect it later.

For comparison, we report the approach where no factorization weights are learned but instead the memory is always enriched by the new words of the current talk (denoted by MEM-ALL). Furthermore, we vary this approach by keeping words in memory only if they are recognized in the talk in which they are added to the memory (denoted by MEM-FOUND). This is similar to MEM-CL, where only utterances containing found new words are transferred to the new-words adaptation data.

#### IV. RESULTS

For the factorization experiment, we tuned the learning rate and the upsampling factor of the new-words data on the development split of the new-words data set. In Figure 2 the results can be seen for the learning rate  $10^{-4}$  and the upsampling factor  $10^5$  (this corresponds to approximately half new and baseline data in each batch). We compare the addition of factorization weights to the whole model and to only the decoder. We report only the F1 score (evaluated only on the new words) because the precision of all approaches is very high ( $>98\%$ ) and therefore the recall strongly correlates with the F1 score.

On the left of Figure 2, we see the amount of factorization parameters versus the F1 score evaluated on the new words. Until the point  $k = 4$ , the performance increases for both the factorized decoder and the factorized encoder and decoder, and the model is able to use the additional flexibility. Generally, the model with factorized encoder and decoder is better than the model with only factorized decoder. However, adapting only the factorized decoder is faster by approximately a factor of two to five. For comparison, the memory-enhanced ASR




Fig. 2. Results of the factorization experiment: Left: Number of parameters (MB, 16-bit) vs. F1-score after training with the new-words data for a factorized decoder and a factorized encoder+decoder. The baseline model has F1-score 0.402. Right: Number of training samples per new word vs. F1-score for the different categories with a factorized encoder+decoder and  $k = 4$ .

model trained for the continual learning experiment scores recall 0.721, precision 0.964 and F1-score 0.825 when using a memory containing all new words from the eval split. On the right of Figure 2, the number of samples per new word versus the F1 score is shown for each category. We see a similar behavior for all categories and that the performance increases when adapting with more samples per new word. For four and five samples per new word, the performance does not increase anymore, probably since there are not many new words for which that number of samples is available. Furthermore, acronyms seem to work best, possibly due to the small number of characters each part of the acronym can be, followed by named entities and special words.

For the continual learning experiment, we use  $k = 4$ , factorized decoder only to reduce the amount of computational power required and the upsampling factor  $10^5$ . We report the WER on the Tedlium testset over the course of the learning cycles (see Figure 3, left) as well as the forward transfer, that is, if a new word occurred, how likely is it recognized in a later talk (see Figure 3, middle and right). Since the baseline decoder and the memory decoder independently predict the next token (afterward the results are combined), the factorization adaptation of the baseline decoder does not interfere with the memory extension of the model.

We do not have a reference transcript available for the talks. Therefore, we manually labeled all detected occurrences of new words not present in the baseline hypothesis and all our approaches (4315) if they are correctly recognized or false positives. Based on these labels, we calculate the forward transfer: For each new word  $w$  that occurs somewhere in talk  $i$  and  $j$ ,  $i < j$ , we evaluate if  $w$  is recognized in talk  $j$  using the model produced in iteration step  $j - 1$  (empty memory for MEM-CL). We group according to how many training examples have been available until iteration step  $j - 1$  (false positives not counted), such that each bucket contains more than 300 samples. Then, we calculate precision and recall for




Fig. 3. Results of the continual learning experiment: Left: General performance: Accumulated number of new words from slides versus WER (in %) on the Tedium testset. Middle and right: Forward transfer: Number of training samples per new word versus forward transfer recall and precision.

each bucket. In total, the forward transfer is calculated on 402 unique words.

We see that for MEM-CL, the general performance (see Figure 3, left) is preserved, while for MEM-ALL and MEM-FOUND, the WER increases consistently over the course of the learning cycles (more than 12% and 7% relative respectively). This is because the MEM-ALL and MEM-FOUND models have to distinguish between more and more words in the memory which produces false positives. MEM-ALL and MEM-FOUND have 2199 and 692 memory entries at the end of the last learning cycle, respectively. For comparison, the baseline model scores 4.39% WER.

For forward transfer recall (see Figure 3, middle), we see that for MEM-ALL and MEM-FOUND the recognition performance of new words is at approximately 70% after the new words are added to the memory. The models cannot utilize more occurrences of the same word. Therefore, performance does not increase over the course of the learning cycles. It even degrades a bit because the number of training samples per new word correlates with the total number of words in the memory, and with more words in the memory the performance degrades. In contrast, the performance of MEM-CL increases consistently to over 80% when more training samples of a new word arrive. For the forward transfer precision (see Figure 3, right) we find that MEM-CL has the best performance of all approaches and fewer false alarms are produced if more training samples per new word are available. Words which MEM-CL recognized better than the baselines include back-propagate, CSPs, elementwise, frontend, LRU, MDP, MDPs, MRU, prefetcher, RNN, RNNs, SRAM, SIMD and tanh.

#### A. Limitations

The extraction of words from the slides is not trivial. We took all words from the slides and filtered by the words in the training data. Errors occur when the words on the slides are capitalized but in the transcript they should not be capitalized, and when the extracted words contain spelling errors. A further

problem arises when the morphological variances of known words are extracted, as their presence in memory increases the rate of false positives. Adding other morphological variances of the same word may help here. All these mentioned errors can propagate through the learning cycles. Moreover, training with pseudo-labels could lead to a degradation of the general performance. We did not observe that, however, one could restrict the training for the pseudo-labels to tokens that are part of new words. Other false positives can occur, e.g., when first a word like ISCA is learned followed by a word ISA which is written similarly. In this case, a small supervised intervention could help reduce the number of false positive utterances added to the adaptation data.

## V. CONCLUSION

We demonstrated a self-supervised continuous learning approach for learning new words. This is done by iteratively extracting new words from slides of a given talk, detecting new words by a memory-enhanced ASR model, and using collected data for adapting low-rank matrix weights added to each weight matrix of the model. With this approach, we can increase the performance of new words as they occur more often to more than 80% recall while the general performance of the model is preserved.

## VI. ACKNOWLEDGEMENTS

This research was supported in part by a grant from Zoom Video Communications, Inc. Furthermore, the projects on which this research is based were funded by the Federal Ministry of Education and Research (BMBF) of Germany under the number 01EF1803B (RELATER), the Horizon research and innovation program of the European Union under grant agreement No 101135798 (Meetween), and the KIT Campus Transfer GmbH (KCT) staff in accordance with the collaboration with Carnegie – AI. The authors gratefully acknowledge the support. We thank Ngoc Quan Pham for providing his framework NMTGMinor to train our systems.

## REFERENCES

- [1] 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.
- [2] N.-Q. Pham, T.-S. Nguyen, J. Niehues, M. Müller, S. Stüker, and A. Waibel, "Very deep self-attention networks for end-to-end speech recognition," *arXiv preprint arXiv:1904.13377*, 2019.
- [3] C. Fügen, A. Waibel, and M. Kolss, "Simultaneous translation of lectures and speeches," *Machine translation*, vol. 21, pp. 209–252, 2007.
- [4] A. Waibel and C. Fuegen, "Simultaneous translation of open domain lectures and speeches," Jan. 3 2012, uS Patent 8,090,570.
- [5] C. Huber, J. Hussain, S. Stüker, and A. Waibel, "Instant one-shot word-learning for context-specific neural sequence-to-sequence speech recognition," in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 1–7.
- [6] D. Le, M. Jain, G. Keren, S. Kim, Y. Shi, J. Mahadeokar, J. Chan, Y. Shangquan, C. Fuegen, O. Kalinli *et al.*, "Contextualized streaming end-to-end speech recognition with trie-based deep biasing and shallow fusion," *arXiv preprint arXiv:2104.02194*, 2021.
- [7] P. Maergner, A. Waibel, and I. Lane, "Unsupervised vocabulary selection for real-time speech recognition of lectures," in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2012, pp. 4417–4420.
- [8] B. Suhm, M. Woszczyna, and A. Waibel, "Detection and transcription of new words," in *Eurospeech*. Citeseer, 1993, pp. 2179–2182.
- [9] B. Suhm and A. Waibel, "Towards better language models for spontaneous speech," in *Proc. ICSLP'94*, vol. 2, 1994, pp. 831–834.
- [10] A. Waibel and I. R. Lane, "Enhanced speech-to-speech translation system and methods for adding a new word," Mar. 3 2015, uS Patent 8,972,268.
- [11] H.-J. Chang, H.-y. Lee, and L.-s. Lee, "Towards lifelong learning of end-to-end asr," *arXiv preprint arXiv:2104.01616*, 2021.
- [12] L. Fu, X. Li, L. Zi, Z. Zhang, Y. Wu, X. He, and B. Zhou, "Incremental learning for end-to-end automatic speech recognition," in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 320–327.
- [13] M. Yang, I. Lane, and S. Watanabe, "Online continual learning of end-to-end speech recognition models," *arXiv preprint arXiv:2207.05071*, 2022.
- [14] G. Chennupati, M. Rao, G. Chadha, A. Eakin, A. Raju, G. Tiwari, A. K. Sahu, A. Rastrow, J. Droppo, A. Oberlin *et al.*, "Ilasr: privacy-preserving incremental learning for automatic speech recognition at production scale," in *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2022, pp. 2780–2788.
- [15] 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.
- [16] L. Wang, X. Zhang, H. Su, and J. Zhu, "A comprehensive survey of continual learning: theory, method and application," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [17] N.-Q. Pham, T.-N. Nguyen, S. Stüker, and A. Waibel, "Efficient weight factorization for multilingual speech recognition," *arXiv preprint arXiv:2105.03010*, 2021.
- [18] 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," *arXiv preprint arXiv:2106.09685*, 2021.
- [19] M. Del Rio, N. Delworth, R. Westerman, M. Huang, N. Bhandari, J. Palakapilly, Q. McNamara, J. Dong, P. Zelasko, and M. Jetté, "Earnings-21: A practical benchmark for asr in the wild," *arXiv preprint arXiv:2104.11348*, 2021.
- [20] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, "fairseq: A fast, extensible toolkit for sequence modeling," in *Proceedings of NAACL-HLT 2019: Demonstrations*, 2019.
- [21] X. Li, C. Wang, Y. Tang, C. Tran, Y. Tang, J. Pino, A. Baevski, A. Conneau, and M. Auli, "Multilingual speech translation with efficient finetuning of pretrained models," *arXiv preprint arXiv:2010.12829*, 2020.
- [22] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in Neural Information Processing Systems*, vol. 33, pp. 12449–12460, 2020.
- [23] Y. Liu, J. Gu, N. Goyal, X. Li, S. Edunov, M. Ghazvininejad, M. Lewis, and L. Zettlemoyer, "Multilingual denoising pre-training for neural machine translation," *Transactions of the Association for Computational Linguistics*, vol. 8, pp. 726–742, 2020.
- [24] R. M. French, "Catastrophic forgetting in connectionist networks," *Trends in cognitive sciences*, vol. 3, no. 4, pp. 128–135, 1999.
- [25] C. Huber, J. Hussain, T.-N. Nguyen, K. Song, S. Stüker, and A. Waibel, "Supervised adaptation of sequence-to-sequence speech recognition systems using batch-weighting," in *Proceedings of the 2nd Workshop on Life-long Learning for Spoken Language Systems*, 2020, pp. 9–17.