



Emphasizing unseen words: New vocabulary acquisition for end-to-end speech recognition

Leyuan Qu^{a,b,*}, Cornelius Weber^a, Stefan Wermter^a

^a Knowledge Technology, Department of Informatics, University of Hamburg, Hamburg, Germany

^b Department of Artificial Intelligence, Zhejiang Laboratory, Hangzhou, China

ARTICLE INFO

Article history:

Received 19 April 2022

Received in revised form 29 December 2022

Accepted 19 January 2023

Available online 10 February 2023

Keywords:

Automatic speech recognition

Continual learning

Out-of-vocabulary word recognition

End-to-end learning

Loss rescaling

ABSTRACT

Due to the dynamic nature of human language, automatic speech recognition (ASR) systems need to continuously acquire new vocabulary. Out-Of-Vocabulary (OOV) words, such as trending words and new named entities, pose problems to modern ASR systems that require long training times to adapt their large numbers of parameters. Different from most previous research focusing on language model post-processing, we tackle this problem on an earlier processing level and eliminate the bias in acoustic modeling to recognize OOV words acoustically. We propose to generate OOV words using text-to-speech systems and to rescale losses to encourage neural networks to pay more attention to OOV words. Specifically, we enlarge the classification loss used for training neural networks' parameters of utterances containing OOV words (sentence-level), or rescale the gradient used for back-propagation for OOV words (word-level), when fine-tuning a previously trained model on synthetic audio. To overcome catastrophic forgetting, we also explore the combination of loss rescaling and model regularization, i.e. L2 regularization and elastic weight consolidation (EWC). Compared with previous methods that just fine-tune synthetic audio with EWC, the experimental results on the LibriSpeech benchmark reveal that our proposed loss rescaling approach can achieve significant improvement on the recall rate with only a slight decrease on word error rate. Moreover, word-level rescaling is more stable than utterance-level rescaling and leads to higher recall rates and precision rates on OOV word recognition. Furthermore, our proposed combined loss rescaling and weight consolidation methods can support continual learning of an ASR system.

© 2023 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Recently, end-to-end ASR models have been receiving a lot of attention and achieving impressive performance (Bahdanau, Cho, & Bengio, 2015; Graves, 2012; Graves, Fernández, Gomez, & Schmidhuber, 2006). These models significantly simplify the training process to directly map acoustic inputs to characters or words. Additionally, limited domain-specific knowledge is required, which dramatically boosts model development and deployment. However, end-to-end models need a lot of training data and perform poorly on words out-of-vocabulary (OOV) or rarely existing in the training data, for example, trending words and new named entities.

Since it takes substantial efforts to collect labeled OOV speech data for ASR model training, current approaches to tackle the

OOV problem mainly involve a language model (LM) or post-processing, for instance, user-dependent language models (Brown, Della Pietra, Desouza, Lai, & Mercer, 1992; Maskey, Bacchiani, Roark, & Sproat, 2004), LM rescoring (Guo, Sainath, & Weiss, 2019) and finite-state transducer lattice extension (Zhao et al., 2019). However, the post-processing techniques only obtain limited improvement as they do not tackle the root causes at the acoustic level.

Alternatively, fine-tuning end-to-end ASR models with synthetic audio containing OOV words can efficiently improve the recall rate of unseen vocabulary, which usually leverages advanced text-to-speech (TTS) systems to generate audio-text pairs required for ASR model training. However, the catastrophic forgetting problem substantially degrades the overall performance of ASR systems, especially on non-OOV words. Elastic weight consolidation (EWC) (Kirkpatrick et al., 2017) is adapted to tackle this problem but leads to a limited recall rate improvement for OOV word recognition. In this paper, we take this method a step further and propose loss rescaling to encourage models to pay more attention to unknown words. Instead of just fine-tuning ASR models where all words are treated equally, enlarging the loss of

* Corresponding author at: Knowledge Technology, Department of Informatics, University of Hamburg, Hamburg, Germany.

E-mail addresses: qu@informatik.uni-hamburg.de (L. Qu), cornelius.weber@uni-hamburg.de (C. Weber), stefan.wermter@uni-hamburg.de (S. Wermter).

URL: <http://www.informatik.uni-hamburg.de/WTM/> (L. Qu).

utterances containing OOV words (sentence-level) or increasing the gradient of unseen words (word-level) can efficiently incline the model to update the weights related to OOV words. We choose 100 OOV words appearing in LRS3-TED dataset but not existing in LibriSpeech dataset. Then, we crawl texts including the new words from the Internet and synthesize audio with TTS systems. The experimental results of fine-tuning audio-text pairs on a hybrid CTC¹/attention ASR model show a significant improvement on recall. When combining EWC with the word-level loss rescaling, we achieve 45.81% of recall on the ROOV test with only 7.8% and 4.6% of relative WER increase on the LibriSpeech test-clean and test-other datasets respectively. As a result, we have improved the recognition of OOV words while maintaining the accuracy for non-OOV words.

2. Related work

2.1. The recognition of OOV words in end-to-end ASR models

Since the OOV problem has occurred, a number of approaches have been proposed for conventional GMM-HMM² models (Miao, Metzger, & Rawat, 2013; Sheikh, Fohr, Illina, & Linares, 2017) and hybrid DNN-HMM³ models (Hwang & Sung, 2016; Knill et al., 2013). In this section, we only review methods towards end-to-end ASR architectures which have been the most promising methods in speech recognition in recent years.

Aleksic et al. (2015) extend class-based LMs (Brown et al., 1992; Maskey et al., 2004) by creating a user-dependent small LM for contact name recognition on voice commands, which is compiled dynamically based on the contact names on users' devices. Moreover, contact insertion reward is proposed to avoid excessive bias and to balance the information between user-dependent and user-independent cases. Hori, Watanabe, and Hershey (2017b) combine word-level with character-level language modeling in end-to-end architectures. With the word-level LM, the model can achieve better performance by learning stronger and longer context information, while character-level LM is used to overcome the OOV issue that the word-level LM suffers from. A similar idea is investigated by Li, Ye, Zhao, Droppo, and Gong (2017) on acoustic-to-word mapping employing character-level modeling units to tackle OOV issues. Williams, Kannan, Aleksic, Rybach, and Sainath (2018) leverage contextual information, for instance, users' locations, users' favorite songs, and calendar events, to partially rescore the output likelihood from sequence-to-sequence models during beam search instead of bringing an additional LM in. Since previous work does not consider errors generated by speech recognition systems when combining them with external LMs, Guo et al. (2019) incorporate a spelling correction model into the speech recognizer training, that directly maps speech recognizer outputs to ground-truth texts. The experimental results suggest that the proposed spelling correction model outperforms n-best LM rescoring and TTS data fine-tuning.

To enable on-device end-to-end speech recognition models to individually recognize new named entities, such as the contact names on mobile phones, Sim et al. (2019) compare LM biasing and acoustic model fine-tuning methods. Furthermore, several techniques, such as layer freezing, early stopping, and EWC, are investigated to suppress model overfitting during fine-tuning. Instead of using a word-level LM in ASR in which a pre-defined lexicon is required, Likhomanenko, Synnaeve, and Collobert (2019) attempt to decode acoustic models with a character-level LM

which is not constrained by lexicons. The lexicon-free decoder achieves better results on OOV experiments since the character-level LM is naturally able to handle unseen words. Different from traditional hybrid ASR models, for instance DNN-HMM, end-to-end ASR architectures aim to learn LMs and acoustic models into one module, which leads to no clear division between language and acoustic models in the end-to-end fashion. To combine LMs trained from a different domain, Variani, Rybach, Allauzen, and Riley (2020) propose hybrid autoregressive transducer (HAT) to separately model internal LM used for training and an external LM from a different domain used for inference. Consequently, Meng et al. (2021) improve HAT by estimating and subtracting the internal LM scores, and properly integrate an external LM into end-to-end ASR architectures without any further training. To further improve the model performance on proper nouns, Zhao et al. (2019) optimize the shallow-fusion method (Kannan et al., 2018) (integrate an external LM into the inference of a sequence-to-sequence model) by building the LM at subword level instead of at word level. Additionally, they propose the early contextual finite state transducer (FST) to avoid the proper noun candidates being pruned during the Viterbi beam search. Moreover, a common set of prefixes is utilized to avoid the contextual biasing always being active and prevent models from degrading on cases not containing OOV words.

Different from most of the previous work focusing on LM post-processing which requires candidate units existing in n-best lists or decoding lattices, in this paper, we tackle the OOV problem on an earlier processing level by eliminating the bias in acoustic modeling to recognize OOV words acoustically.

2.2. Data augmentation with synthetic audio for ASR

Proper speech data augmentation does not only boost the model performance, but can also significantly improve system robustness and generalization (Ko, Peddinti, Povey, & Khudanpur, 2015). There are many strategies used in ASR training, for example, noise addition, pitch shifting, speed perturbation, back-translation (Hayashi et al., 2018) and room impulse response injection with real or simulated data (Ko, Peddinti, Povey, Seltzer, & Khudanpur, 2017). More recently, a simple yet effective approach, SpecAugment (Park et al., 2019), has been proposed and achieves state-of-the-art results on the LibriSpeech benchmark corpus. The basic idea for SpecAugment is randomly masking or cropping a fixed area on spectrograms in the time- or frequency-domain, which effectively prevents model overfitting, especially for noisy conditions. Another well-established method is mixing synthetic audio with real data by leveraging advanced TTS models, like Tacotron2 (Shen et al., 2018), DeepVoice3 (Ping et al., 2018) and FlowTron (Valle, Shih, Prenger, & Catanzaro, 2020).

Rossenbach, Zeyer, Schlüter, and Ney (2020) compare commonly used data augmentation strategies with the TTS audio. The results reveal the effectiveness of TTS data in ASR system training. Laptev et al. (2020) investigate the effect of augmenting data with TTS audio for low-resource speech recognition. The resulting models outperform other systems with the same setting and semi-supervised learning methods. Furthermore, other authors explore the influence of the audio quality with different vocoders, i.e. Griffin-Lim and LPCNet (Valin & Skoglund, 2019). Instead of just mixing the synthetic audio data during training, Rosenberg et al. (2019) exploit the impact of the TTS model's effectiveness and diversity on ASR results. Moreover, lexical diversity is also investigated on domain adaptation experiments.

Inspired by the benefit brought by TTS data, we synthesize audio with text crawled from the Internet containing OOV words as the training set for new vocabulary acquisition and model adaptation.

¹ CTC is the abbreviation for Connectionist Temporal Classification.

² GMM and HMM are short for Gaussian Mixture Model and Hidden Markov Model respectively.

³ DNN is short for Deep Neural Network.

2.3. End-to-end ASR architectures

End-to-end ASR architectures aim to directly map acoustic observations to text transcripts, such as characters and words. Different from conventional systems which separately train the acoustic model and LM with different criteria and corpora, end-to-end ASR systems train all modules jointly. Consequently, the training progress is significantly simplified and does not require too much domain-specific knowledge. There are four main categories in end-to-end learning, i.e. CTC, Recurrent Neural Network Transducer (RNN-T), encoder–decoder with attention and hybrid CTC/attention architectures. They mainly differ in their alignment of input acoustic features and output label sequences.

2.3.1. CTC-based end-to-end models

The basic idea of CTC-based approaches is to bring in a special token, ‘blank’, which is dynamically filled in the places between modeling units. Consequently, the exact boundary information required by conventional methods is not needed anymore. Then, a carefully designed dynamic programming algorithm is used to search optimal paths and convert the frame-level token sequences to meaningful utterances by removing blank tokens and merging repeated labels.

Graves et al. (2006) firstly adopt the ground-breaking CTC approach to overcome the problems faced in conventional ASR systems, i.e. the requirement of frame-level segmentation and the mapping from model outputs to ground-truth labels. Hanun, Maas, Jurafsky, and Ng (2014) simplify the CTC building process and propose a modified prefix-search decoding algorithm to completely discard the cumbersome decoding strategies used in HMM-based systems. Amodei et al. (2016) conduct a comprehensive evaluation on model architecture, system optimization, and model deployment. Furthermore, parallel training on model level (Coates et al., 2013) and data level (Dean et al., 2012) is utilized to heavily speed up the training process.

Furthermore, substituting RNNs with Convolutional Neural Networks (CNNs) has been a growing trend since pure CNN-based architectures can dramatically reduce the training time and inference latency, which is critical for speech recognition tasks running in real time. Jasper (Li et al., 2019) achieves state-of-the-art results on the LibriSpeech dataset by stacking the CNN-only blocks. Inspired by Jasper, Kriman et al. (2020) propose QuartzNet which uses a similar fully convolutional architecture but with significantly fewer parameters. QuartzNet enables the CTC-based end-to-end ASR model to run locally on mobile devices.

2.3.2. RNN-T end-to-end models

CTC assumes that every output node is conditionally independent of other outputs, which makes it difficult to model the dependency between adjacent frames. Consequently, RNN-T (Graves, 2012) was proposed to overcome the unreasonable conditionally independence assumptions. Different from CTC which only focuses on acoustic sequence modeling, RNN-T uses a separate module, called prediction network, to model the context information, which can be treated as an intrinsic language model, followed by a joint network to classify the concatenation of encoder and prediction network outputs.

Rao, Sak, and Prabhavalkar (2017) exploit some strategies for RNN-T architecture training and find that CTC-based encoder pre-training and language model-based prediction network pre-training improve the model performance. Zhang, Lu, et al. (2020) present a Transformer-based (Vaswani et al., 2017) RNN-T model, in which the Transformer modules are used to learn representations from speech signals and text sequences. Inspired by Li et al. (2019) and QuartzNet (Kriman et al., 2020), Han et al. (2020) propose ContextNet by introducing a fully convolutional encoder

into the RNN-T architecture. Since the CNN has a weaker receptivity for long context than Long Short-Term Memory (LSTM) and Transformers, the squeeze-and-excitation (SE) layer is integrated into ContextNet to enhance long-distance dependence. Gulati et al. (2020) propose Conformer which combines CNNs with self-attention (Vaswani et al., 2017) to concurrently learn local and global features, and show significant improvement on the LibriSpeech test and test-other set.

2.3.3. Attention-based encoder–decoder models

Another branch of end-to-end systems is the attention-based encoder–decoder architecture. Different from CTC or RNN-T architectures, the attention mechanism dynamically aligns the encoder and decoder time steps to temporally align the input and output sequences. Chorowski, Bahdanau, Serdyuk, Cho, and Bengio (2015) transfer the attention-based recurrent networks to speech recognition and achieve comparable results on the TIMIT benchmark that compare well with conventional methods. Chan, Jaitly, Le, and Vinyals (2016) propose LAS (listen, attend and spell) to directly map acoustic inputs to transcription. The results show that the attention mechanism prevents model overfitting on the training set. Meanwhile, Bahdanau, Chorowski, Serdyuk, Brakel, and Bengio (2016) show that attention-based models can implicitly learn better context information than CTC and conventional models. In addition, local monotonic attention (Tjandra, Sakti, & Nakamura, 2017), full-sequence attention (Prabhavalkar, Sainath, Li, Rao, & Jaitly, 2017), time-restricted self-attention (Povey, Hadrian, Ghahremani, Li, & Khudanpur, 2018), multi-channel attention (Braun, Neil, Anumula, Ceolini, & Liu, 2018) and online attention (Fan, Zhou, Chen, Jia, & Liu, 2018) are proposed to reduce the complexity of attention computation and learn more robust alignments.

2.3.4. Hybrid CTC/attention architectures

To fully incorporate the merits of CTC and attention models, Kim, Hori, and Watanabe (2017) propose to jointly train CTC and attention-based approaches in a multi-task learning fashion by sharing one encoder. The evaluation on WSJ and CHiME-4 noisy speech shows the hybrid architecture can efficiently speed up convergence and learn more robust alignment between input frames and output sequences. Hori, Watanabe, and Hershey (2017a) extend the hybrid CTC/attention method with a joint decoding algorithm by rescaling or combining the probabilities from both objective functions. Then, a monotonic chunk-wise attention (Miao, Cheng, Zhang, Li, & Yan, 2019) and transformer-based encoder (Miao, Cheng, Gao, Zhang, & Yan, 2020) are utilized to enable the hybrid CTC/attention model to work in online streaming tasks. Zhang, Wu, et al. (2020) propose a new two-pass approach (U2) which unifies the streaming and non-streaming ASR models into one architecture. The hybrid CTC/attention architecture is becoming more and more popular. There has been a trend to unify the streaming and non-streaming architecture into one model with this architecture. However, limited research has focused on the recognition of OOV words in end-to-end architectures. In this paper, we tackle the problem of OOV words with the U2 ASR model⁴ which will be described in more detail in 4.1.

3. Methodology

In this section, we demonstrate the proposed loss rescaling approaches at sentence level and word level. Furthermore, we introduce the techniques of L2 regularization and EWC used to overcome catastrophic forgetting problems.

⁴ <https://github.com/wenet-e2e/wenet>




Fig. 1. (a) Utterance loss distribution in one mini-batch. (b) Utterance loss distribution after loss rescaling.

3.1. Loss rescaling at sentence level

During training, the CTC function returns one loss per utterance, and the mean of all utterance losses in the same mini-batch would be used for back-propagation. As shown in Fig. 1(a), each bar in the figure means one utterance loss in a randomly selected mini-batch. We observe that the utterance losses in one mini-batch are evenly distributed. Sometimes, the loss of utterances containing OOV words can be slightly higher or even lower than other utterances without OOV words. Consequently, the model pays equal attention to each utterance or word, which leads to the final model performance heavily relying on the frequency of words in training sets. Sometimes, the model attention is even biased by non-OOV words.

To emphasize an utterance containing OOV words, we rescale the utterance loss, as indicated in Fig. 1(b), by multiplying it with a hyper-parameter μ in Eq. (1), where x are acoustic inputs, y are the target text references, and O is the set of OOV words.

$$\mathcal{L}_{\text{sentence}}(x, y) = \begin{cases} \mathcal{L}_{\text{CTC}}(x, y), & \text{if } o \text{ not in } y, \forall o \in O, \\ \mu \mathcal{L}_{\text{CTC}}(x, y), & \text{if } o \text{ in } y, \forall o \in O \end{cases} \quad (1)$$

3.2. Loss rescaling at word level

Given an input acoustic vector $\mathbf{x} = (x_0, \dots, x_T)$, and a target label sequence $\mathbf{y} = (y_0, \dots, y_U)$, where $T \gg U$, and T and U are the length of the acoustic vector and target label sequence respectively, the CTC loss aims to maximize the log probability in Eq. (2), where $\tilde{\mathbf{y}}$ is the extended label sequence of \mathbf{y} by inserting blank labels ϕ at the beginning and the end of \mathbf{y} and between every two label tokens, $\tilde{\mathbf{y}} = \mathbf{y} \cup \phi$. When training CTC ASR systems, the last layer of the neural network outputs one $N \times 1$ vector for each acoustic frame x_t at time step t , where N is the number of modeling units. For example, when building on character level, N is 28 (a...z + space + blank).

$$\mathcal{L}_{\text{CTC}} = -P(\tilde{\mathbf{y}}|\mathbf{x}) = - \sum_{\mathbf{a} \in \mathcal{F}^{-1}(\tilde{\mathbf{y}})} P(\mathbf{a}|\mathbf{x}) \quad (2)$$

After processing the last input vector x_T , we can get a $(2U + 1) \times T$ lattice matrix, for instance the one in Fig. 2 for the utterance “News about Brexit” modeling on subword units is shown. Fig. 3 lists some intermediate decoding results of the CTC function. The final CTC loss contains the sum of probabilities of all possible paths that can be converted to be the target labels \mathbf{y} by merging repeat units and removing blank tokens as shown in Eq. (2),

where \mathbf{a} is a possible token path and $\mathcal{F} : \tilde{\mathbf{y}} \rightarrow \mathbf{y}$ is the function that maps the extended label sequence $\tilde{\mathbf{y}}$ back to the true label sequence \mathbf{y} .

We denote $\hat{y}(t, u)$ and $b(t, u)$ as the probability of a label and a blank token at node (t, u) respectively. According to the definition of CTC (Graves, 2008), as shown in Fig. 4, for any blank tokens, there are only two paths from the previous step reaching the current blank token since label tokens cannot be skipped. When the current node is a label token and is same as $u - 2$, there is no connection from $u - 2$ to u . A blank token has to be generated to separate the adjacent same tokens. For instance, there is no path from the first ‘o’ to the second one when decoding the word ‘look’. Lastly, three possible paths can reach the label token when $u \neq u - 2$. Three same cases for the backward procedure are shown in Fig. 4 as well.

According to the three cases shown in Fig. 4, the forward variable $\alpha(t, u)$ can be calculated recursively as follows:

$$\alpha(t, u) = \begin{cases} \hat{y}(t-1, u-1)\alpha(t-1, u-1) & \text{if } (t, u) = \text{blank} \\ +b(t-1, u)\alpha(t-1, u) & \\ b(t-1, u-1)\alpha(t-1, u-1) & \text{if } (t, u) \neq \text{blank} \\ +\hat{y}(t-1, u)\alpha(t-1, u) & \text{and } (t, u) = (t-1, u-2) \\ \hat{y}(t-1, u-2)\alpha(t-1, u-2) & \text{if } (t, u) \neq \text{blank} \\ +b(t-1, u-1)\alpha(t-1, u-1) & \text{and } (t, u) \neq (t-1, u-2) \\ +\hat{y}(t-1, u)\alpha(t-1, u) & \end{cases} \quad (3)$$

It is worth noting that when the two adjacent tokens are same, i.e. $(t, u) = (t-1, u-2)$, there is no direct transition between the two repeated tokens and the second one can only be reached through the blank token or the node of $(t-1, u)$.

Similarly, blue arrows reaching the node (t, u) and the backward variable $\beta(t, u)$ can be represented by Eq. (4) in three cases.

$$\beta(t, u) = \begin{cases} b(t, u)\beta(t+1, u) & \text{if } (t, u) = \text{blank} \\ +\hat{y}(t, u)\beta(t+1, u+1) & \\ \hat{y}(t, u)\beta(t+1, u) & \text{if } (t, u) \neq \text{blank} \\ +b(t, u)\beta(t+1, u+1) & \text{and } (t, u) = (t+1, u+2) \\ \hat{y}(t, u)\beta(t+1, u) & \text{if } (t, u) \neq \text{blank} \\ +b(t, u)\beta(t+1, u+1) & \text{and } (t, u) \neq (t+1, u+2) \\ +\hat{y}(t, u)\beta(t+1, u+2) & \end{cases} \quad (4)$$

$P(\mathcal{A}_{t,u}|\mathbf{x})$, the probability of any candidate paths passing through node (u, t) conditioned on the input sequence \mathbf{x} , can be obtained by multiplying forward (Eq. (3)) and backward probabilities (Eq. (4)) by Eq. (5).

$$P(\mathcal{A}_{t,u}|\mathbf{x}) = \alpha(t, u)\beta(t, u) \quad (5)$$

Thereby, the gradient of the CTC loss function \mathcal{L}_{CTC} w.r.t $\hat{y}(t, u)$ and $b(t, u)$ can be estimated by Eq. (6) and Eq. (7) respectively.

$$\frac{\partial \mathcal{L}_{\text{CTC}}}{\partial \hat{y}(t, u)} \propto \begin{cases} \alpha(t, u)\beta(t+1, u+1) & \text{if } (t, u) = \text{blank} \\ \alpha(t, u)\beta(t+1, u) & \text{if } (t, u) \neq \text{blank} \\ & \text{and } (t, u) = (t+1, u+2) \\ \alpha(t, u)\beta(t+1, u) & \text{if } (t, u) \neq \text{blank} \\ +\beta(t+1, u+2) & \text{and } (t, u) \neq (t+1, u+2) \end{cases} \quad (6)$$

$$\frac{\partial \mathcal{L}_{\text{CTC}}}{\partial b(t, u)} \propto \begin{cases} \alpha(t, u)\beta(t+1, u) & \text{if } (t, u) = \text{blank} \\ \alpha(t, u)\beta(t+1, u+1) & \text{if } (t, u) \neq \text{blank} \end{cases} \quad (7)$$

The CTC function treats all nodes equally and aims to minimize the global loss, which makes models hardly focus on local connections in the decoding lattice. To guide models to pay more attention to the OOV words, we emphasize the OOV words (the nodes

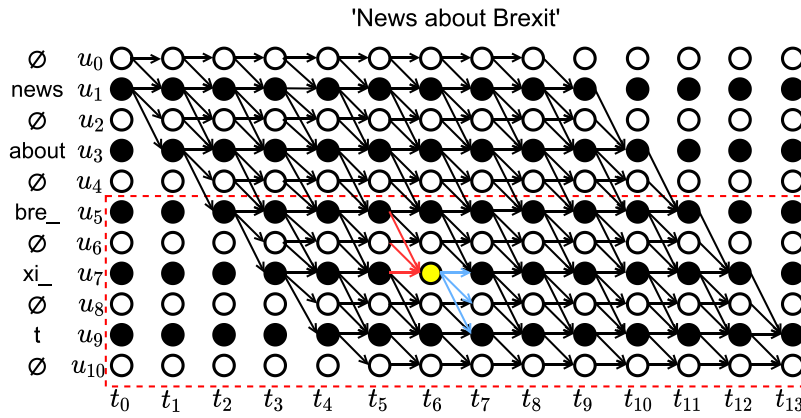


Fig. 2. Illustration of CTC decoding lattices for the example sentence of ‘News about Brexit’, where the modeling unit is subword and “Brexit” is an OOV word. Black nodes are label tokens and white nodes are blank tokens.

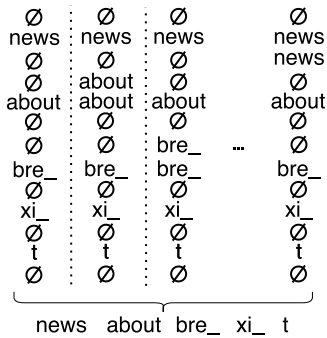


Fig. 3. Possible intermediate decoding sequences of the CTC function.

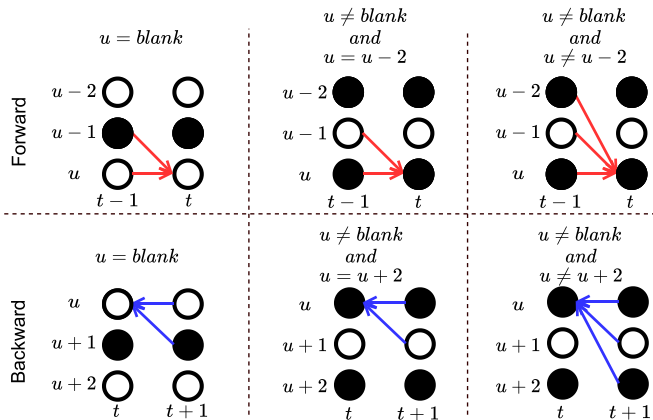


Fig. 4. Diagram of three cases in forward or backward computation, where black nodes are label tokens and white nodes are blank tokens. Red and blue arrows are forward and backward paths respectively.

in the dotted box in Fig. 2) by rescaling the probabilities of OOV nodes in candidate alignments. Thus, the rescaled probability of all alignments passing through OOV nodes is as follows:

$$\tilde{P}(\mathcal{A}_{t,u}|\mathbf{x}) = \begin{cases} \mu P(\mathcal{A}_{t,u}|\mathbf{x}), & \text{if } u \in O \\ P(\mathcal{A}_{t,u}|\mathbf{x}), & \text{otherwise} \end{cases} \quad (8)$$

The regularized loss function at word level is:

$$\mathcal{L}_{word} = - \sum_{\mathcal{A} \in \mathcal{F}^{-1}(\tilde{\gamma})} \tilde{P}(\mathcal{A}_{t,u}|\mathbf{x}) \quad (9)$$

We implement our approach by multiplying the gradients of OOV nodes on the candidate path with μ , as shown in the following equations:

$$\frac{\partial \mathcal{L}_{word}}{\partial \tilde{\gamma}(t, u)} = \begin{cases} \mu \frac{\partial \mathcal{L}_{CTC}}{\partial \tilde{\gamma}(t, u)}, & \text{if } u \in O \\ \frac{\partial \mathcal{L}_{CTC}}{\partial \tilde{\gamma}(t, u)}, & \text{otherwise} \end{cases} \quad (10)$$

$$\frac{\partial \mathcal{L}_{word}}{\partial b(t, u)} = \begin{cases} \mu \frac{\partial \mathcal{L}_{CTC}}{\partial b(t, u)}, & \text{if } u \in O \\ \frac{\partial \mathcal{L}_{CTC}}{\partial b(t, u)}, & \text{otherwise} \end{cases} \quad (11)$$

where O is the set of OOV words tokenized into subwords.

3.3. Overcoming catastrophic forgetting

Directly fine-tuning models on a dataset obeying a different distribution from the original training set may lead to catastrophic forgetting. The updated model may overfit the new dataset but forget the knowledge learned on the original one. To overcome models suffering from catastrophic forgetting, we adopt two approaches during fine-tuning. The first one is mixing partial original audio from LibriSpeech used for baseline model training with synthetic speech, since adding data that obeys the same distribution as the training set can efficiently mitigate the forgetting problem. We explore the effect of different mixing ratios and present the results in Section 5. The other approach is constraining model parameters from updating during fine-tuning with L2 regularization or EWC (Kirkpatrick et al., 2017), and we will introduce the details in the following sections.

3.3.1. L2 regularization

The L2 regularization loss $\mathcal{L}_{L2}(\theta)$ is shown in Eq. (12), where $\mathcal{L}(\theta)$ is the original CTC loss or rescaled loss in Eqs. (1) and (9). θ_i is the i th parameter of the ASR model to be updated during fine-tuning, and θ'_i is the i th parameter in the baseline model which is invariable and saved locally. λ is the coefficient to balance the scale of two parts.

$$\mathcal{L}_{L2}(\theta) = \mathcal{L}_{CTC}(\theta) + \frac{\lambda}{2} \sum_i (\theta_i - \theta'_i)^2 \quad (12)$$

L2 loss takes the difference between the fine-tuned model and the old model into account to ensure the updated model will not stray away too much from the baseline.

3.3.2. Elastic weight consolidation

Different from L2 loss that always refers to a fixed standard and treats all parameters equally, the EWC loss as shown in Eq. (13) uses the diagonal of the Fisher information matrix F to

dynamically weigh the importance of each model parameter for the source task.

$$\mathcal{L}_{EWC}(\theta) = \mathcal{L}_{CTC}(\theta) + \frac{\lambda}{2} \sum_i F_i \cdot (\theta_i - \theta'_i)^2 \quad (13)$$

The Fisher information matrix F can be estimated by the following equation with the gradients of the convergent source model.

$$F_i = \frac{1}{|D|} \sum_{d \in D} \frac{\partial^2 \mathcal{L}_{CTC}(d, \theta'_i)}{\partial \theta_i'^2} \approx \frac{1}{|D|} \sum_{d \in D} \frac{\partial \mathcal{L}_{CTC}(d, \theta'_i)}{\partial \theta_i'^2} \quad (14)$$

where θ' and D are the parameters and the dataset used in the source ASR task respectively. The more important parameters for the source task would have the larger values in the Fisher information matrix, which constrains the change of important parameters and avoids knowledge forgetting. The less important ones would have relatively small values and are encouraged to adapt on the new task. In this paper, the diagonal of the Fisher information matrix is estimated on the LibriSpeech 960 h training set.

4. Experiments

4.1. ASR model architecture

The end-to-end ASR model used in our experiments is the two-pass hybrid CTC/attention architecture, U2 (Zhang, Wu, et al., 2020), as shown in Fig. 5. The shared encoder converts acoustic features x into a latent vector \mathbf{h}^{enc} , then the CTC decoder transforms the latent vector into character/word probability $P(y_t|x_t)$ with the same length as the input frames. Meanwhile, the attention decoder generates one character/word probability $P(y_u|y_{u-1}, \dots, y_0, x)$ per time step by conditioning on the attention content vector \mathbf{c}_u and the decoder output from the last step y_{u-1} . During training, the sum of CTC loss and attention loss is used to do back-propagation, while during inference, the n -best hypotheses produced by the CTC decoder are rescored by the attention decoder to obtain better performance. The candidate with the highest score will be the final output.

We develop our model based on the U2 model published with the WeNet toolkit (Zhang et al., 2021) which unifies streaming and non-streaming ASR models into one architecture by proposing the dynamic chunk-based attention. The encoder consists of 12 conformer blocks, with 4 multi-head attention, 2048 linear units, swish activation, a positional dropout rate of 0.1 and Conv2D kernel size of 31 for each block. The attention decoder contains 6 transformer blocks, and the CTC decoder is composed of 1 linear layer and 1 log softmax function. The U2 model is pre-trained on the 960 h LibriSpeech corpus and achieves 3.18% of WER and 8.72% of WER on the test-clean and the test-other test set respectively when rescoreing with attention decoder. The pre-trained weights are available on the website,⁵ which will be the baseline model for all our experiments.

4.2. OOV set with real audio

We build a 100-OOV-word (see Table 5 in the Appendix) dataset from LRS3-TED (Afouras, Chung, & Zisserman, 2018) corpus since there are no standard OOV corpora published by the community. LRS3-TED is an audio-visual dataset collected from TED and TEDx talks with spontaneous speech and various speaking styles. It is comprised of over 400 h of video by more than

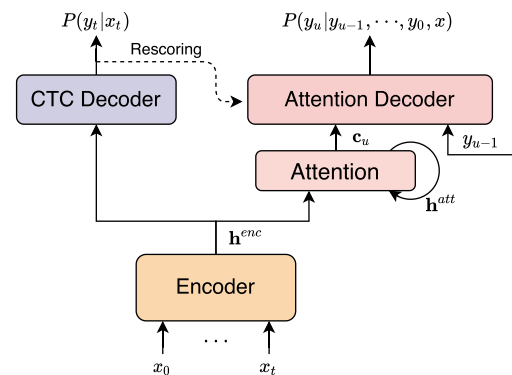


Fig. 5. Two-pass hybrid CTC/attention ASR architecture.

5000 speakers and contains an extensive vocabulary. We filter the vocabulary existing in LRS3-TED but not present in LibriSpeech and select 100 OOV words from more than 100 speakers used for test, where each OOV word contains 50 utterances. We randomly split these utterances into training, validation and test sets with a ratio of 2:1:2. The duration ratio of training, validation and test sets is 3 h : 1.6 h : 2.8 h. In the rest of this paper, the three sets will be referred to as real OOV (ROOV) training, ROOV val and ROOV test set respectively. We report all experimental results on the ROOV test set. More details about the 100 OOV words can be found in Appendix.

4.3. OOV set with synthetic audio

Our goal in this paper is to improve OOV word recognition with synthetic audio and loss rescaling methods. In this subsection, we introduce the synthetic dataset used for model training.

Text Crawling: we crawl 100 sentences for each new word with Scrapy,⁶ where each sentence contains less than 50 words in case of running out of memory. During crawling, we filter those sentences include OOV words out of the selected 100-OOV-word set to ensure that the model performance is only influenced by the chosen ones.

Speech Synthesis: we split the 100 sentences for each new word into training and validation with a ratio of 9:1. The model evaluation will be conducted on the ROOV test set with real audio. Instead of applying a single multi-speaker TTS system, we use several commercial Application Programming Interfaces (APIs), i.e. Baidu TTS API,⁷ Google TTS API,⁸ iFLYTEK TTS API,⁹ Tencent TTS API¹⁰ and Alibaba TTS API¹¹ to synthesize audio and to enable more voices and more variety in speech. In contrast to open-source multi-speaker TTS models, the commercial APIs produce higher quality speech, which is crucial for the following experiments. 8 different speaker voices (4 males and 4 females) are used for training set synthesis and 2 voices (1 male and 2 female) for the validation set. There is no voice overlap between the 2 parts. The duration ratio of synthetic audio used for training and validation is 9.0 h : 2.3 h. In the rest of this paper, the synthetic data will be referred to as synthetic OOV (SOOV) training and SOOV val set respectively.

⁶ <https://scrapy.org/>

⁷ <https://ai.baidu.com/tech/speech/tts>

⁸ <https://cloud.google.com/text-to-speech>

⁹ https://www.xfyun.cn/doc/tts/online_tts/API.html

¹⁰ <https://cloud.tencent.com/document/api>

¹¹ <https://www.alibabacloud.com/help/doc-detail/84435.htm>

⁵ http://mobvoi-speech-public.ufile.ucloud.cn/public/wenet/librispeech/20210216_conformer_exp.tar.gz

Data Augmentation: to avoid overfitting, we perturb speech on speed with the factors of 0.9, 1.0, and 1.1. Furthermore, clean speech is augmented with 5 kinds of room impulse responses.¹² 10 noise sources (Reddy et al., 2019), such as announcements, appliances, and traffic, are added to the reverberated speech with 6 levels of speech-to-noise ratio (0, 4, 8, 12, 16 and 20). Moreover, SpecAugment (Park et al., 2019) with 2 frequency masks (maximum width 50) is utilized on the fly during training.

4.4. Evaluation metrics

We use 3 metrics to evaluate the experimental results of our proposed method:

- **WER:** word error rate is the ratio of error terms, i.e., substitutions, deletions, and insertions, to the total number of words in the reference.
- **Recall:** recall is the number of true positives TP over the sum of the number of true positives and the number of false negatives FN .

$$\text{Recall} = \frac{TP}{TP + FN} \quad (15)$$

- **Precision:** precision is the number of true positives TP over the sum of the number of true positives and the number of false positives FP .

$$\text{Precision} = \frac{TP}{TP + FP} \quad (16)$$

4.5. Training settings

The baseline model is trained on the 960 h LibriSpeech dataset with a batch size of 12, an initial learning rate of $4e-3$ and 25000 warm-up steps. When doing fine-tuning for OOV experiments, we use a bigger batch size of 20 to enable the model to see more utterances not containing OOV words and avoid loss explosion. A tiny initial learning rate of $4e-6$ is utilized for fine-tuning, which is annealed with a value of 1.1 after every 3000 steps, since a tiny learning rate can efficiently ensure stable model learning and retain the previously learned knowledge. In addition, to avoid gradient explosion, we clip all gradients greater than 2, while the threshold used in baseline training is 5.

The validation set during model training is the mixture of the LibriSpeech dev and OOV TTS dev set with the ratio of 1:1. The model checkpoint performing best on the mixture validation set is used for evaluation on test sets with early stopping. It is noteworthy that the attention mechanism and attention decoder are always frozen since we found that fine-tuning the entire network leads to gradient explosion. In addition, it is hard to balance the CTC loss and attention loss since the CTC loss is rescaled in our methods. We only use the attention and the decoder for rescoring.

5. Experimental results

In this section, we report the experimental results from the following perspectives.

5.1. Results of speech mixture from source and target domain

To mitigate catastrophic forgetting, we mix original real speech in the LibriSpeech training set (source domain) with the ROOV or the SOOV training set (target domain). It is still an open question what the best mixing ratio is. We fine-tune the baseline

model with different ratios and report results on the standard LibriSpeech test sets (test-clean and test-other) and the ROOV test set.

As shown in Table 1, when fine-tuning only with ROOV training set (real speech from LRS3-TED), the model shows the inability to retain old knowledge and performs badly on previous LibriSpeech tasks, which leads to a tremendous rise in WER on the test-other set from 8.72% to 41.52%. Context information is crucial for ASR models. In the 0:1 setting, the pre-trained ASR model is destroyed, especially the learned context knowledge. Consequently, the overfitted model is hard to infer a correct context and recognize OOV words.

The forgetting tendency slows down as original data is incorporated into training. When the ratio of audio from LibriSpeech and LRS3-TED is 2:1, the model achieves the highest recall of 32.05%. The more data from the source task (LibriSpeech) is used for training, the more previous knowledge is retained and the better performance is obtained on the LibriSpeech evaluation sets. However, the model tends to focus on the previous LibriSpeech tasks as the ratio increases, which leads to the decrease of recall on the ROOV test.

We can draw the same conclusion when fine-tuning ASR models with synthetic data as shown in Table 2. We prioritize the model performance regarding recall since the goal of this paper is to enable the ASR model to learn new vocabulary, and the catastrophic forgetting issue will be tackled in the next section. Therefore, the 2:1 mixture ratio is used in the following experiments.

5.2. Results of loss rescaling at sentence level

In this section, we explore the effect of loss rescaling at sentence level. We compare the model performance using real speech data with the results using synthetic audio. As shown in Table 3, using L2 and EWC regularization efficiently reduces catastrophic forgetting and improves the recall rate on OOV words while the WER increases only in few cases on the test-clean and test-other test sets. We find $\lambda = 5e7$ is the best weight to balance the L2/EWC loss and the ASR losses.

We reproduce the method proposed by Zheng, Liu, Gunceler, and Willett (2021), in which a RNN-T ASR model is fine-tuned with EWC on mixed real and synthetic audio. However, the dataset is not published. It is noteworthy that the experimental results reported in Tables 3 and 4 are based on our generated data with the method proposed by Zheng et al. (2021). In addition, as shown in Table 3, row “Isolated Words”, we de-emphasize the non-OOV words by fine-tuning ASR models with utterances containing only isolated OOV words which are segmented from real or synthetic continuous speech according to the time alignment information obtained from the Montreal Forced Aligner.¹³ Compared with the method proposed by Zheng et al. (2021), fine-tuning ASR models with only isolated OOV words can effectively improve the recall rate but it leads to much more serious forgetting on non-OOV recognition.

When just fine-tuning the base model with real or synthetic OOV audio, all words are treated equally, which leads the model to hardly focus on the OOV words we concern. Therefore, we propose loss rescaling and encourage the model to pay more attention to OOV words by enlarging the loss of sentences containing unknown words. For the loss rescaling weight μ , we examine the values 1, 10, 100, 1000, and 10,000. As we can see in Table 3, the OOV recall rapidly increases when rescaling the target sentences by 100 times compared to only fine-tuning using

¹² <http://www.iks.rwth-aachen.de/en/research/tools-downloads/databases/aachen-impulse-response-database/>

¹³ <https://github.com/MontrealCorpusTools/Montreal-Forced-Aligner>

Table 1

The influence of the ratio of speech data from LibriSpeech and real speech in ROOV training set (LRS3-TED) on ASR and OOV word recognition.

Model	Ratio	WER	WER	WER	Recall	Precision
		Test-clean	Test-other	ROOV test	ROOV test	ROOV test
Baseline	–	3.18	8.72	15.33	1.37	100
LibriSpeechL + ROOV training	0:1	30.18	41.52	38.21	19.58	97.28
	1:1	23.34	30.41	25.83	24.30	93.18
	2:1	13.26	26.77	23.27	32.05	98.41
	3:1	7.26	15.25	22.31	30.82	98.02
	4:1	4.11	11.05	15.52	26.12	99.04

Table 2

The influence of the ratio of speech data from LibriSpeech and synthetic speech in SOOV training set on ASR and OOV word recognition.

Model	Ratio	WER	WER	WER	Recall	Precision
		Test-clean	Test-other	ROOV test	ROOV test	ROOV test
Baseline	–	3.18	8.72	15.33	1.37	100
LibriSpeechL + SOOV training	0:1	35.54	53.42	42.27	13.51	99.83
	1:1	29.05	39.08	30.22	20.51	91.21
	2:1	20.34	28.72	25.31	27.54	97.67
	3:1	13.37	19.21	23.22	26.25	98.53
	4:1	6.71	16.23	20.31	23.21	98.08

Table 3

Loss rescaling at sentence level with L2/EWC regularization. μ and λ are the loss weight in Eq. (1) and the weight of L2/EWC in Eq. (12)/Eq. (13) respectively. R is short for “Real” and represents using ROOV training set with real audio for fine-tuning. S is short for “Synthetic” and denotes using SOOV training set with synthetic audio for fine-tuning.

Test sets	μ	λ	WER↓		WER↓		WER↓		Recall↑		Precision↑	
			Test-clean	Test-other	Test-clean	Test-other	ROOV test	ROOV test	ROOV test	ROOV test	ROOV test	ROOV test
Model	–	–	R	S	R	S	R	S	R	S	R	S
Baseline	1	0	3.18		8.72		15.33		1.37		100	
L2	1	5e7	5.21	5.32	10.22	10.91	15.03	18.84	24.13	15.24	98.12	99.02
	10	5e7	5.53	5.36	10.93	12.02	15.17	18.69	35.12	24.42	98.51	92.03
	100	5e7	6.07	6.38	11.38	11.94	15.01	19.28	50.02	31.26	94.35	95.71
	1000	5e7	6.83	6.93	11.83	12.48	15.74	19.74	53.42	39.72	90.04	89.43
	10000	5e7	7.22	7.79	13.48	13.59	16.33	19.26	51.78	39.77	83.44	78.25
Zheng et al. (2021)	1	5e7	3.18	3.20	8.93	9.02	14.87	18.23	30.57	16.20	98.43	99.12
Isolated Words	1	5e7	5.83	6.32	11.01	12.27	14.65	18.47	35.92	22.48	97.65	98.15
	10	5e7	5.21	5.20	9.91	9.83	14.59	18.34	49.22	30.71	98.11	99.04
EWC	100	5e7	5.29	5.37	9.82	9.98	14.66	18.04	54.28	41.22	97.21	97.55
	1000	5e7	5.55	5.54	10.54	11.37	14.67	18.37	54.06	42.38	88.24	89.41
	10000	5e7	6.07	6.12	11.78	11.84	14.52	18.02	55.31	42.58	80.51	84.28

L2 (50.02% VS 24.13% for real speech and 31.26% VS 24.13% for synthetic audio).

As a bigger μ is used, the recall further rises, but the WER on the test-clean and the test-other test sets is getting worse. We hypothesize that directly rescaling the entire sentence loss may also enhance irrelevant words or noises, which leads to gradient explosion during training and accelerates forgetting previous knowledge. Hence, we have to use a very small learning rate and clip the gradients over 2.0 to ensure the progress of fine-tuning. In contrast to L2 regularization, EWC can provide more stable and resilient protection of the weights important for the previous LibriSpeech tasks but still with a relatively high loss in the ASR performance, as shown in Table 3.

In addition, when fine-tuning only with the synthetic audio, we obtain competitive recall rate compared to utilizing real speech from the LRS3-TED dataset, for example, when using EWC and rescaling the loss by 1000 times larger, we achieve 54.06% VS 42.38% recall rate for real and synthetic speech. Furthermore, compared with the method proposed by Zheng et al. (2021), rescaling the OOV utterance loss can achieve significant improvement on recall with only slight decrease on WER and precision.

5.3. Results of loss rescaling at word level

Instead of enhancing the entire sentence loss, in this section, we report the results of only rescaling unknown words. As shown in Table 4, the λ weight, which is needed to balance the ASR and L2/EWC loss, is smaller (1e7) than the one used at sentence level (5e7) in Table 3, and we do not observe gradient explosion during training unless μ is very large, e.g. 1e4. The results without loss rescaling ($\mu = 1$) is slightly different from the one shown in Table 3, which is caused by different λ values used for L2/EWC normalization. Using 10 times smaller μ can obtain a similar or even higher recall at word-level rescaling, for example, using μ of 100 gets a 45.81% recall rate at word level compared to using μ of 1000 getting a 42.38% recall rate at sentence level when regularizing models with EWC and fine-tuning with only synthetic data.

Rescaling loss on OOV words obtains lower WER on the standard LibriSpeech test sets. Moreover, the bigger μ is used, the higher recall is achieved for the ROOV test. Furthermore, a worse WER is obtained for the LibriSpeech benchmark, which is observed at sentence-level loss rescaling as well. To make a trade-off between WER and Recall, 100 times loss rescaling is performing

Table 4

Loss rescaling at word level with L2/EWC regularization. μ and λ are the loss weight in Eq. (1) and the weight of L2/EWC in Eq. (12)/Eq. (13) respectively. R is short for “Real” and represents using ROOV training set with real audio for fine-tuning. S is short for “Synthetic” and denotes using SOOV training set with synthetic audio for fine-tuning.

Test sets	μ	λ	WER↓ Test-clean		WER↓ Test-other		WER↓ ROOV test		Recall↑ ROOV test		Precision↑ ROOV test	
			R	S	R	S	R	S	R	S	R	S
Baseline	1	0	3.18		8.72		15.33		1.37		100	
L2	1	1e7	4.03	4.04	9.05	9.23	14.95	15.21	24.01	13.77	98.12	99.02
	10	1e7	4.16	4.23	9.31	9.01	15.01	15.03	33.45	24.63	98.51	92.03
	100	1e7	4.53	5.03	9.48	14.83	14.95	18.28	48.41	32.49	94.35	95.71
	1000	1e7	4.71	4.62	10.73	14.45	15.19	19.74	55.28	43.24	90.04	89.43
	10000	1e7	5.33	5.59	12.74	14.15	16.26	19.26	55.83	44.04	83.44	78.25
Zheng et al. (2021)	1	1e7	3.19	3.18	8.83	8.79	14.09	15.44	30.83	18.12	97.25	98.44
Isolated Words	1	1e7	4.22	5.47	10.15	11.19	14.31	16.24	38.07	27.62	95.14	94.33
	10	1e7	3.22	3.31	8.94	8.83	13.96	15.22	49.10	31.43	94.23	94.57
	100	1e7	3.30	3.43	8.97	9.12	13.47	15.17	59.74	45.81	90.45	89.14
EWC	1000	1e7	3.55	3.62	9.99	10.45	14.24	15.54	59.76	46.03	80.61	81.09
	10000	1e7	4.01	5.37	10.28	11.23	14.15	15.87	62.19	46.71	77.17	69.42

best for our experimental results with only a 7.8%/4.6% relative WER increase on the test-clean/test-other test set or a 45.81% recall rate on the ROOV test. In addition, we obtain competitive performance by only using synthetic audio compared with using real speech data for fine-tuning.

5.4. Discussion

New vocabulary emerges all the time due to the evolution of human language. Therefore, it is important to enable a trained ASR system to dynamically acquire unseen vocabulary. The combined loss rescaling and weight consolidation methods proposed in this paper can support continual learning (Parisi, Kemker, Part, Kanan, & Wermter, 2019) of an ASR system. The methods neither require any labeled data nor do they require retraining a new ASR model from scratch.

An interesting finding is that enhancing the gradient of blank tokens within and after OOV words is important as well, which encourages the decoding procedure moving forward, for instance, the rows of u_6 , u_8 and u_{10} in Fig. 2. Otherwise, the decoding progress is cut off and models repeatedly produce one token, such as “news about bre_ bre_ bre_” when only enlarging the gradient of “bre_” in the utterance of “new about bre_ xi_ t”. Sometimes, the fine-tuned ASR system even repeats one token, for example “bre_ bre_ bre_ bre_ bre_”, when μ is very large.

Additionally, we find that the performance and the speed of convergence are affected by the batch size, especially for loss rescaling at sentence level. When the batch size is small, e.g. 5, all utterances in one batch may contain OOV words, which leads to a bigger rescaled loss. Consequently, the model suffers from gradient explosion, and L2 or EWC regularization can hardly constrain the model weights diverging.

6. Conclusion

In this paper, we present the use of synthetic speech to boost an ASR model on the recognition of OOV words. In addition to fine-tuning with audio containing OOV words, we propose to rescale loss at sentence level or word level, which encourages models to pay more attention to unknown words. Experimental results reveal that fine-tuning the baseline ASR model combined with loss rescaling and L2/EWC regularization can significantly improve the OOV word recall rate and efficiently overcome models suffering from catastrophic forgetting. Furthermore, loss rescaling at word level is more stable than at sentence level and results in less ASR performance loss on general non-OOV words and previous LibriSpeech tasks. The combination of proposed

loss rescaling, which updates the new task-related parameters (OOV word recognition), and EWC, which retains the old task-learned weights (speech recognition on the LibriSpeech dataset), can enable continual learning of an ASR system.

The proposed target word loss rescaling method is simple and effective, but there are still some issues left to be improved. Currently, results are evaluated on synthetic audio data which is different from the spontaneous speech recorded in the real world. Future work could focus on real-scenario speech collecting and model evaluation, which enables us to well understand and compare the contribution of using synthetic and real speech containing OOV words. Additionally, the current OOV word set needs to be known, hence, how to automatically detect and optimize OOV words is a potential direction. The trade-off between WER on universal test sets (e.g. LibriSpeech test-clean and test-other sets) and recall rate on the OOV set is another issue. A dynamic L2/EWC weight (Leang, Sistu, Bürger, Bursuc, & Yogamani, 2020) can be adopted to replace the fixed λ weight. Later in the fine-tuning, a fixed regularization weight could influence model updating. Moreover, we are interested in investigating the effectiveness of our proposed method on RNN-T and attention-based encoder-decoder ASR systems. It is also worthwhile to explore our loss rescaling method on some general unbalanced label problems, for example, speaker diarization and voice verification. Since continual learning in sequence processing is a young research field (Ahrens, Abawi, & Wermter, 2021; Cossu, Carta, Lomonaco, & Bacciu, 2021; Ehret et al., 2020), our loss rescaling method may have wider implication for data where novel elements are learnt in temporal or spatial context with known elements.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

This study used only external publicly available data sources, as referenced.

Acknowledgments

The authors gratefully acknowledge the support from the China Scholarship Council (CSC), from the Youth Foundation Project of Zhejiang Lab (No. K2023KH0AA01), and from the German Research Foundation DFG (projects TRR 169 and LeCAREbot).

Table 5

The 100 OOV words selected from LRS3 dataset.

MEANINGFUL	QUANTUM	RESILIENCE	VACCINE	MONITORING
RESEARCHER	GLOBALLY	GLOBALIZATION	EMPOWERMENT	SPACECRAFT
HORMONE	HEALTHCARE	PRESCHOOL	WORKFORCE	ALGORITHM
ISRAELI	NOBEL	EMPATHY	ECOSYSTEM	INTERACTING
SCHIZOPHRENIA	SOFTWARE	INTEGRATE	PROGRAMMED	YOGA
ALZHEIMER	VIETNAM	RETHINK	PAKISTAN	LATVIA
PARKINSON	SOCCER	RACISM	SILICON	MARIJUANA
NEUROSCIENCE	GENETICALLY	MAINSTREAM	DEMOGRAPHIC	UNEMPLOYMENT
PARADIGM	GENOME	CREATIVITY	INSULIN	PROSTHETIC
TRANSGENDER	KENYA	RACIST	STRESSFUL	PERSONALIZED
RAPED	VIRAL	STORYTELLING	ENTREPRENEURSHIP	ADULTHOOD
MICROSOFT	STEREOTYPE	EXPERTISE	LITERACY	RAINFORREST
MINDSET	GOOGLE	WORKPLACE	YOUTUBE	SAUDI
FACEBOOK	IDEOLOGY	INTERACTIVE	AUTISM	CHEMOTHERAPY
SUSTAINABILITY	ACTIVIST	ROBOTIC	SOCIETAL	EBOLA
PRODUCTIVITY	TEENAGER	DOPAMINE	SYNDROME	VIABLE
TRANSFORMATIVE	CORTEX	OXYTOCIN	PARENTING	JIHAD
RECYCLING	DIMENSIONAL	TARGETED	LAPTOP	NIGERIA
SUSTAINABLE	PROGRAMMING	SMARTPHONE	COLLABORATIVE	VEGAN
EINSTEIN	RESEARCHING	VACCINE	WIKIPEDIA	ANTIMATTER

Appendix

See Table 5.

References

- Afouras, T., Chung, J. S., & Zisserman, A. (2018). LRS3-TED: a large-scale dataset for visual speech recognition. arXiv preprint arXiv:1809.00496.
- Ahrens, K., Abawi, F., & Wermter, S. (2021). DRILL: Dynamic representations for imbalanced lifelong learning. In *International conference on artificial neural networks* (pp. 409–420). Springer.
- Aleksic, P., Allauzen, C., Elson, D., Kracun, A., Casado, D. M., & Moreno, P. J. (2015). Improved recognition of contact names in voice commands. In *IEEE international conference on acoustics, speech and signal processing* (pp. 5172–5175). IEEE.
- Amodei, D., Ananthanarayanan, S., Anubhai, R., Bai, J., Battenberg, E., Case, C., et al. (2016). Deep Speech 2: End-to-end speech recognition in English and Mandarin. In *International conference on machine learning* (pp. 173–182). PMLR.
- Bahdanau, D., Cho, K., & Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. In *International conference on learning representations*.
- Bahdanau, D., Chorowski, J., Serdyuk, D., Brakel, P., & Bengio, Y. (2016). End-to-end attention-based large vocabulary speech recognition. In *IEEE international conference on acoustics, speech and signal processing* (pp. 4945–4949). IEEE.
- Braun, S., Neil, D., Anumula, J., Ceolini, E., & Liu, S.-C. (2018). Multi-channel attention for end-to-end speech recognition. In *Proc. INTERSPEECH* (pp. 17–21).
- Brown, P. F., Della Pietra, V. J., Desouza, P. V., Lai, J. C., & Mercer, R. L. (1992). Class-based n-gram models of natural language. *Computational Linguistics*, 18(4), 467–480.
- Chan, W., Jaitly, N., Le, Q. V., & Vinyals, O. (2016). Listen, attend and spell. In *IEEE international conference on acoustics, speech and signal processing*.
- Chorowski, J. K., Bahdanau, D., Serdyuk, D., Cho, K., & Bengio, Y. (2015). Attention-based models for speech recognition. In *Advances in neural information processing systems*, vol. 28.
- Coates, A., Huval, B., Wang, T., Wu, D., Catanzaro, B., & Andrew, N. (2013). Deep learning with COTS HPC systems. In *International conference on machine learning* (pp. 1337–1345). PMLR.
- Cossu, A., Carta, A., Lomonaco, V., & Bacciu, D. (2021). Continual learning for recurrent neural networks: an empirical evaluation. *Neural Networks*, 143, 607–627.
- Dean, J., Corrado, G. S., Monga, R., Chen, K., Devin, M., Le, Q. V., et al. (2012). Large scale distributed deep networks. In *Advances in neural information processing systems* (pp. 1232–1240).
- Ehret, B., Henning, C., Cervera, M. R., Meulemans, A., Von Oswald, J., & Grewe, B. F. (2020). Continual learning in recurrent neural networks. arXiv preprint arXiv:2006.12109.
- Fan, R., Zhou, P., Chen, W., Jia, J., & Liu, G. (2018). An online attention-based model for speech recognition. In *Proc. INTERSPEECH* (pp. 4390–4394).
- Graves, A. (2008). Supervised sequence labelling with recurrent neural networks. In *Studies in computational intelligence*.
- Graves, A. (2012). Sequence transduction with recurrent neural networks. In *Proceedings of the 23rd international conference on machine learning*.
- Graves, A., Fernández, S., Gomez, F., & Schmidhuber, J. (2006). Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on machine learning* (pp. 369–376).
- Gulati, A., Qin, J., Chiu, C.-C., Parmar, N., Zhang, Y., Yu, J., et al. (2020). Conformer: Convolution-augmented transformer for speech recognition. In *Proc. INTERSPEECH* (pp. 5036–5040).
- Guo, J., Sainath, T. N., & Weiss, R. J. (2019). A spelling correction model for end-to-end speech recognition. In *IEEE international conference on acoustics, speech and signal processing* (pp. 5651–5655). IEEE.
- Han, W., Zhang, Z., Zhang, Y., Yu, J., Chiu, C.-C., Qin, J., et al. (2020). ContextNet: Improving convolutional neural networks for automatic speech recognition with global context. In *Proc. INTERSPEECH* (pp. 3610–3614).
- Hannun, A. Y., Maas, A. L., Jurafsky, D., & Ng, A. Y. (2014). First-pass large vocabulary continuous speech recognition using bi-directional recurrent DNNs. arXiv preprint arXiv:1408.2873.
- Hayashi, T., Watanabe, S., Zhang, Y., Toda, T., Hori, T., Astudillo, R., et al. (2018). Back-translation-style data augmentation for end-to-end ASR. In *IEEE spoken language technology workshop* (pp. 426–433). IEEE.
- Hori, T., Watanabe, S., & Hershey, J. R. (2017a). Joint CTC/attention decoding for end-to-end speech recognition. In *Proceedings of the 55th annual meeting of the Association for Computational Linguistics (Volume 1: Long papers)* (pp. 518–529).
- Hori, T., Watanabe, S., & Hershey, J. R. (2017b). Multi-level language modeling and decoding for open vocabulary end-to-end speech recognition. In *2017 IEEE automatic speech recognition and understanding workshop* (pp. 287–293). IEEE.
- Hwang, K., & Sung, W. (2016). Character-level incremental speech recognition with recurrent neural networks. In *IEEE international conference on acoustics, speech and signal processing* (pp. 5335–5339). IEEE.
- Kannan, A., Wu, Y., Nguyen, P., Sainath, T. N., Chen, Z., & Prabhavalkar, R. (2018). An analysis of incorporating an external language model into a sequence-to-sequence model. In *IEEE international conference on acoustics, speech and signal processing* (pp. 1–5828). IEEE.
- Kim, S., Hori, T., & Watanabe, S. (2017). Joint CTC-attention based end-to-end speech recognition using multi-task learning. In *IEEE international conference on acoustics, speech and signal processing* (pp. 4835–4839). IEEE.
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13), 3521–3526.
- Knill, K. M., Gales, M. J., Rath, S. P., Woodland, P. C., Zhang, C., & Zhang, S.-X. (2013). Investigation of multilingual deep neural networks for spoken term detection. In *IEEE workshop on automatic speech recognition and understanding* (pp. 138–143). IEEE.
- Ko, T., Peddinti, V., Povey, D., & Khudanpur, S. (2015). Audio augmentation for speech recognition. In *Sixteenth annual conference of the international speech communication association*.
- Ko, T., Peddinti, V., Povey, D., Seltzer, M. L., & Khudanpur, S. (2017). A study on data augmentation of reverberant speech for robust speech recognition. In *IEEE international conference on acoustics, speech and signal processing* (pp. 5220–5224). IEEE.
- Kriman, S., Beliaev, S., Ginsburg, B., Huang, J., Kuchaiev, O., Lavrukhin, V., et al. (2020). Quartznet: Deep automatic speech recognition with 1D time-channel separable convolutions. In *IEEE international conference on acoustics, speech and signal processing* (pp. 6124–6128). IEEE.

- Laptev, A., Korostik, R., Svischev, A., Andrusenko, A., Medennikov, I., & Rybin, S. (2020). You do not need more data: improving end-to-end speech recognition by text-to-speech data augmentation. In *2020 13th international congress on image and signal processing, biomedical engineering and informatics* (pp. 439–444). IEEE.
- Leang, I., Sistu, G., Bürger, F., Bursuc, A., & Yogamani, S. (2020). Dynamic task weighting methods for multi-task networks in autonomous driving systems. In *2020 IEEE 23rd international conference on intelligent transportation systems* (pp. 1–8). IEEE.
- Li, J., Lavrukhin, V., Ginsburg, B., Leary, R., Kuchaiev, O., Cohen, J. M., et al. (2019). Jasper: An end-to-end convolutional neural acoustic model. In *Proc. INTERSPEECH* (pp. 71–75).
- Li, J., Ye, G., Zhao, R., Droppo, J., & Gong, Y. (2017). Acoustic-to-word model without OOV. In *2017 IEEE automatic speech recognition and understanding workshop* (pp. 111–117). IEEE.
- Likhomanenko, T., Synnaeve, G., & Collobert, R. (2019). Who needs words? Lexicon-free speech recognition. In *Proc. INTERSPEECH* (pp. 3915–3919).
- Maskey, S. R., Bacchiani, M., Roark, B., & Sproat, R. (2004). Improved name recognition with meta-data dependent name networks. In *IEEE international conference on acoustics, speech, and signal processing, vol. 1* (pp. 1–789). IEEE.
- Meng, Z., Parthasarathy, S., Sun, E., Gaur, Y., Kanda, N., Lu, L., et al. (2021). Internal language model estimation for domain-adaptive end-to-end speech recognition. In *2021 IEEE spoken language technology workshop* (pp. 243–250). IEEE.
- Miao, H., Cheng, G., Gao, C., Zhang, P., & Yan, Y. (2020). Transformer-based online CTC/attention end-to-end speech recognition architecture. In *IEEE international conference on acoustics, speech and signal processing* (pp. 6084–6088). IEEE.
- Miao, H., Cheng, G., Zhang, P., Li, T., & Yan, Y. (2019). Online hybrid CTC/Attention architecture for end-to-end speech recognition. In *Proc. INTERSPEECH* (pp. 2623–2627).
- Miao, Y., Metz, F., & Rawat, S. (2013). Deep maxout networks for low-resource speech recognition. In *IEEE workshop on automatic speech recognition and understanding* (pp. 398–403). IEEE.
- Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2019). Continual lifelong learning with neural networks: A review. *Neural Networks*, 113, 54–71.
- Park, D. S., Chan, W., Zhang, Y., Chiu, C.-C., Zoph, B., Cubuk, E. D., et al. (2019). SpecAugment: A simple data augmentation method for automatic speech recognition. In *Proc. INTERSPEECH* (pp. 1613–1617).
- Ping, W., Peng, K., Gibiansky, A., Arik, S. O., Kannan, A., Narang, S., et al. (2018). Deep Voice 3: 2000-speaker neural text-to-speech. In *International conference on learning representations* (pp. 214–217).
- Povey, D., Hadian, H., Ghahremani, P., Li, K., & Khudanpur, S. (2018). A time-restricted self-attention layer for ASR. In *IEEE international conference on acoustics, speech and signal processing* (pp. 5874–5878). IEEE.
- Prabhavalkar, R., Sainath, T. N., Li, B., Rao, K., & Jaitly, N. (2017). An analysis of “Attention” in sequence-to-sequence models. In *Proc. INTERSPEECH* (pp. 3702–3706).
- Rao, K., Sak, H., & Prabhavalkar, R. (2017). Exploring architectures, data and units for streaming end-to-end speech recognition with RNN-Transducer. In *IEEE automatic speech recognition and understanding workshop* (pp. 193–199). IEEE.
- Reddy, C. K., Beyrami, E., Pool, J., Cutler, R., Srinivasan, S., & Gehrke, J. (2019). A scalable noisy speech dataset and online subjective test framework. In *Proc. INTERSPEECH* (pp. 1816–1820).
- Rosenberg, A., Zhang, Y., Ramabhadran, B., Jia, Y., Moreno, P., Wu, Y., et al. (2019). Speech recognition with augmented synthesized speech. In *2019 IEEE automatic speech recognition and understanding workshop* (pp. 996–1002). IEEE.
- Rossenbach, N., Zeyer, A., Schlüter, R., & Ney, H. (2020). Generating synthetic audio data for attention-based speech recognition systems. In *IEEE international conference on acoustics, speech and signal processing* (pp. 7069–7073). IEEE.
- Sheikh, I., Fohr, D., Illina, I., & Linares, G. (2017). Modelling semantic context of OOV words in large vocabulary continuous speech recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 25(3), 598–610.
- Shen, J., Pang, R., Weiss, R. J., Schuster, M., Jaitly, N., Yang, Z., et al. (2018). Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions. In *IEEE international conference on acoustics, speech and signal processing* (pp. 4779–4783). IEEE.
- Sim, K. C., Beaufays, F., Benard, A., Guliani, D., Kabel, A., Khare, N., et al. (2019). Personalization of end-to-end speech recognition on mobile devices for named entities. In *2019 IEEE automatic speech recognition and understanding workshop* (pp. 23–30). IEEE.
- Tjandra, A., Sakti, S., & Nakamura, S. (2017). Local monotonic attention mechanism for end-to-end speech and language processing. In *The international joint conference on natural language processing*.
- Valin, J.-M., & Skoglund, J. (2019). LPCNet: Improving neural speech synthesis through linear prediction. In *IEEE international conference on acoustics, speech and signal processing* (pp. 5891–5895). IEEE.
- Valle, R., Shih, K., Prenger, R., & Catanzaro, B. (2020). FlowTron: an autoregressive flow-based generative network for text-to-speech synthesis. In *International conference on learning representations*.
- Variani, E., Rybach, D., Allauzen, C., & Riley, M. (2020). Hybrid autoregressive transducer (HAT). In *IEEE international conference on acoustics, speech and signal processing* (pp. 6139–6143). IEEE.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). Attention is all you need. In *Advances in neural information processing systems* (pp. 5998–6008).
- Williams, I., Kannan, A., Aleksic, P. S., Rybach, D., & Sainath, T. N. (2018). Contextual speech recognition in end-to-end neural network systems using beam search. In *Proc. INTERSPEECH* (pp. 2227–2231).
- Zhang, Q., Lu, H., Sak, H., Tripathi, A., McDermott, E., Koo, S., et al. (2020). Transformer transducer: A streamable speech recognition model with transformer encoders and RNN-T loss. In *IEEE international conference on acoustics, speech and signal processing* (pp. 7829–7833).
- Zhang, B., Wu, D., Yang, C., Chen, X., Peng, Z., Wang, X., et al. (2021). WeNet: Production first and production ready end-to-end speech recognition toolkit. In *Proc. INTERSPEECH*.
- Zhang, B., Wu, D., Yao, Z., Wang, X., Yu, F., Yang, C., et al. (2020). Unified streaming and non-streaming two-pass end-to-end model for speech recognition. *arXiv preprint arXiv:2012.05481*.
- Zhao, D., Sainath, T. N., Rybach, D., Rondon, P., Bhatia, D., Li, B., et al. (2019). Shallow-fusion end-to-end contextual biasing. In *Proc. INTERSPEECH* (pp. 1418–1422).
- Zheng, X., Liu, Y., Gunceler, D., & Willett, D. (2021). Using synthetic audio to improve the recognition of out-of-vocabulary words in end-to-end ASR systems. In *IEEE international conference on acoustics, speech and signal processing* (pp. 5674–5678). IEEE.