

Continual Learning of Neural Machine Translation within Low Forgetting Risk Regions

Shuhao Gu^{1,2*}, Bojie Hu³, Yang Feng^{1,2†}

¹ Key Laboratory of Intelligent Information Processing,
Institute of Computing Technology, Chinese Academy of Sciences (ICT/CAS)

² University of Chinese Academy of Sciences

³ Tencent Minority-Mandarin Translation, Beijing, China
{gushuhao19b,fengyang}@ict.ac.cn, bojiehu@tencent.com

Abstract

This paper considers continual learning of large-scale pretrained neural machine translation model without accessing the previous training data or introducing model separation. We argue that the widely used regularization-based methods, which perform multi-objective learning with an auxiliary loss, suffer from the misestimate problem and cannot always achieve a good balance between the previous and new tasks. To solve the problem, we propose a two-stage training method based on the local features of the real loss. We first search low forgetting risk regions, where the model can retain the performance on the previous task as the parameters are updated, to avoid the catastrophic forgetting problem. Then we can continually train the model within this region only with the new training data to fit the new task. Specifically, we propose two methods to search the low forgetting risk regions, which are based on the curvature of loss and the impacts of the parameters on the model output, respectively. We conduct experiments on domain adaptation and more challenging language adaptation tasks, and the experimental results show that our method can achieve significant improvements compared with several strong baselines.

1 Introduction

The current large-scale pretrained neural machine translation (NMT) models, such as GMNMT (Johnson et al., 2017), mBART50-nn (Tang et al., 2020), and M2M-100 (Fan et al., 2021), are generally trained with large amounts of data from different domains and languages, so the model can learn good semantic representation and mapping relationship at the same time. Based on such models, we hope that they can continually acquire new knowledge from new translation tasks, e.g., new domains

*The work was done when the first author was an intern at Tencent.

Corresponding author: Yang Feng.

Reproducible code: <https://github.com/ictnlp/LFR-NMT>.

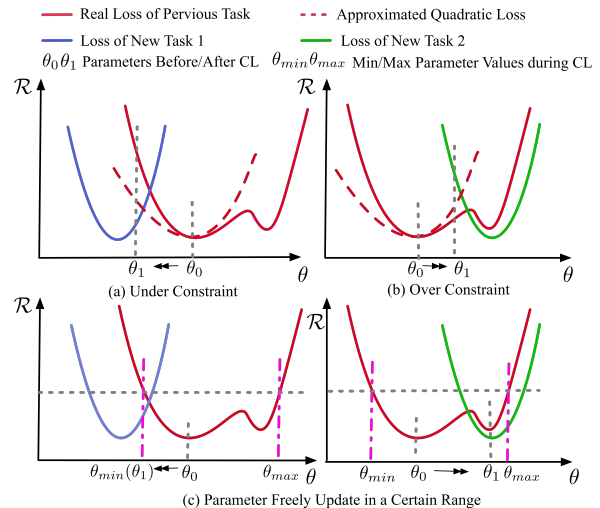


Figure 1: An illustration to indicate the difference between the regularization-based method and our method. \mathcal{R} denotes the task loss. θ denotes the model parameters. \rightarrow and \leftarrow denote the parameter update direction during continual learning. $[\theta_{min}, \theta_{max}]$ denotes the low forgetting risk regions.

and languages, while preserving the previously learned knowledge, i.e., continual learning (CL). However, catastrophic forgetting is the biggest challenge of continual learning, which means the model will forget the previously learned knowledge while learning new knowledge (Gu and Feng, 2020). One solution to avoid forgetting is to mix the previous and the new training data to train the model jointly. However, considering that the training data of the large-scale pre-trained NMT model is usually very large, this method will bring more training consumption and consume more energy. Besides, we sometimes cannot access the previous data due to data privacy or storage limitation.

To solve this problem, many researchers have made their attempts. Some work avoids catastrophic forgetting by constructing pseudo data and mixing them with the new task data for joint training (Kim and Rush, 2016; Ko et al., 2021; Liu et al.,

2021). However, these methods must require that the previous and new tasks are similar, and can not be directly applied to tasks with large differences, such as learning a totally different new language. Some work introduces extra task-specific parameters and only updates these parameters with the new task data (Bapna and Firat, 2019; Escolano et al., 2021). On the one hand, this will increase the size of the model. On the other hand, the task-specific parameters lead to model separation, which makes the model must know which task the input sentence belongs to, otherwise, the translation quality will degrade significantly (Aharoni and Goldberg, 2020). In contrast, the regularization-based methods (Khayrallah et al., 2018; Thompson et al., 2019) don't have such limitations and are more flexible in practice. These methods perform multi-objective learning with an extra penalty item on the parameters, which aims to approximate the real loss on the previous task, and is usually in quadratic form. However, as illustrated in Figure 1 (a) and (b), the approximated loss (the red dashed line) is convex and symmetric, but the real loss (the red solid line) is not necessarily so in most cases, which may lead to the under or over constraint problems. Besides, the multi-objective learning can only make the parameters converge to a region, where the gradients produced by the two losses are equal in size and opposite in direction, i.e., gradients reach balance. However, this does not guarantee that the values of the two losses are small in this region, which means the model may still suffer from the forgetting problem.

Instead of performing multi-objective learning with the approximated loss, we propose a two-stage continual learning method based on the local features of the real loss around the initial parameters. In the first stage, we aim to search the low forgetting risk regions as the parameter update regions, where the performance of the previous task will not degrade significantly when the parameters are updated. Therefore, we can constrain the parameters to such regions to avoid catastrophic forgetting. In the second stage, the parameters will be freely updated in the searched regions only guided by the gradients of the new task, while the update outside these regions will be forbidden, as illustrated in Figure 1 (c). To achieve this, we can use some data related to the previous task to help us find such regions, e.g., the validation set data of the previous translation task, which is usually small-scale and

easy to obtain. We propose two methods to search the low forgetting risk regions. The first method is to find out which parameters have the least impact on the previous task loss. We use the curvature of the loss curve as an indicator and the parameter with small curvature indicates that the loss curve is relatively flat, so the parameters can be updated more greedily. The second method is based on the impacts of parameters on the model output of the previous task. We define an objective function, which tries to maximize the size of update regions as much as possible while restricting the change of model output on the previous task, to help the model learn the update regions automatically. We conduct experiments on the domain adaptation and more challenging language adaptation tasks, and the results show that our method can achieve significant improvements compared with several strong baselines.

Our contribution can be summarized as follows:

- We propose two methods to search the low forgetting risk regions of neural machine translation without accessing the previous training data.
- We conduct experiments on the domain adaptation task and the more challenging language adaptation task of continual learning, and our method can bring significant improvements.
- We prove that our method can also achieve good performance when combined with part of the original training data.

2 Background

In this section, we report the background knowledge used in this paper: the Transformer (Vaswani et al., 2017) model, multilingual NMT, Hessian matrix, and Fisher information matrix.

2.1 Transformer

We denote the input sequence of symbols as $\mathbf{x} = (x_1, \dots, x_I)$ and the ground-truth sequence as $\mathbf{y} = (y_1, \dots, y_J)$. The transformer model is based on the encoder-decoder architecture. The encoder is composed of N identical layers. Each layer has two sublayers. The first is a multi-head self-attention sublayer, and the second is a fully connected feed-forward network. Both of the sublayers are followed by a residual connection operation and a layer normalization operation. The input sequence

\mathbf{x} will be first converted to a sequence of vectors $\mathbf{E}_x = [E_x[x_1]; \dots; E_x[x_J]]$ where $E_x[x_j]$ is the sum of word embedding and position embedding of the source word x_j . Then, this sequence of vectors will be fed into the encoder, and the output of the N -th layer will be taken as source state sequences. We denote it as \mathbf{H}_x . The decoder is also composed of N identical layers. In addition to the same kind of two sublayers in each encoder layer, the cross-attention sublayer is inserted between them, which performs multi-head attention over the output of the encoder. The final output of the N -th layer gives the target hidden states $\mathbf{S} = [\mathbf{s}_1; \dots; \mathbf{s}_{K^*}]$, where \mathbf{s}_k is the hidden states of y_k . We can get the predicted probability of the j -th target word conditioned by the source sentence and the $j - 1$ previous target words. The model is optimized by minimizing a cross-entropy loss of the ground-truth sequence with teacher forcing:

$$\mathcal{L}_{CE} = -\frac{1}{J} \sum_{j=1}^J \log p(y_j | \mathbf{y}_{<j}, \mathbf{x}; \theta), \quad (1)$$

where J is the length of the target sentence and θ denotes the model parameters.

2.2 Multilingual Translation

The multilingual neural machine translation (MNMT) model can translate between multiple different languages with a single model (Johnson et al., 2017). Following Liu et al. (2020), we add a particular language id token at the beginning of the source and target sentences, respectively, to indicate the language.

2.3 Hessian and Fisher Information Matrices

The Hessian Matrix For a twice-differentiable loss \mathcal{L} , the Hessian matrix is the matrix of second-order derivatives of the loss function with respect to the weights, mathematically expressed as $\mathbf{H} = \nabla_{\theta}^2 \mathcal{L}$. Intuitively, its role is to express the curvature of the loss around a given point θ . The smaller the value, the "flatter" the loss is around the given point θ , and vice versa. The flatter the region around θ , the smaller the loss change when the value of θ is changed.

The Fisher Information Matrix The Fisher information matrix \mathbf{F} of the model's conditional distribution $P_{\mathbf{y}|\mathbf{x}}$ is defined as:

$$\mathbf{F} = E_{P_{\mathbf{x},\mathbf{y}}} [\nabla_{\theta} \log p_{\theta}(\mathbf{x}, \mathbf{y}) \nabla_{\theta} \log p_{\theta}(\mathbf{x}, \mathbf{y})^{\top}] \quad (2)$$

Intuitively, the role of the Fisher matrix is very similar to that of the Hessian Matrix. If θ is an accurate set of parameters for the model, we can approximate the Hessian matrix at θ with the Fisher information matrix (Ly et al., 2017).

In practice settings, we can approximate the Fisher information matrix by replacing the model distribution $P_{\mathbf{x},\mathbf{y}}$ with the empirical training distribution $\hat{Q}_{\mathbf{x},\mathbf{y}}$:

$$\begin{aligned} \hat{\mathbf{F}} &= E_{\hat{Q}_{\mathbf{x},\mathbf{y}}} [\nabla_{\theta} \log p_{\theta}(\mathbf{x}, \mathbf{y}) \nabla_{\theta} \log p_{\theta}(\mathbf{x}, \mathbf{y})^{\top}] \\ &= \frac{1}{N} \sum_{n=1}^N \nabla \log p_{\theta}(\mathbf{y}|\mathbf{x}) \nabla \log p_{\theta}(\mathbf{y}|\mathbf{x})^{\top} \end{aligned} \quad (3)$$

3 Method

The goal of our method is to preserve the previously learned knowledge while adapting to the new tasks efficiently. To achieve this, we propose a two-stage training method. In the first stage, we try to depict the local features of the large-scale pre-trained NMT model around its initial parameters θ_0 . Specifically, we want to find a region around θ_0 with low forgetting risk, so we can retain the performance of the model on the previous task as the parameters are updated. We can constrain the parameters within this region so that the model will not suffer from severe forgetting problem. In the second stage, the parameters are updated completely guided by the gradients produced by the new training data within this region, which can mitigate the under-fitting problem on the new task.

3.1 Search Parameter Update Region

We propose two methods to search the low forgetting risk (LFR) regions around the initial parameters θ_0 . For the first method, we hope to find out which parameters to update have the least impact on the model loss of the previous task. To achieve this, we propose to use curvature as an indicator to help find the parameter update regions. For the second method, we determine the LFR regions based on the impact of parameters on the model output of the previous task. We propose an objective function to help the model learn the regions.

Curvature-based Method Intuitively, the curvature of the loss function measures how fast the loss changes as the parameter changes around the initial parameters θ_0 . Therefore, the parameters with small curvature can be safely updated without causing forgetting. As described in Section 2.3, the

Hessian matrix is used to represent the curvature of the model loss, but it is almost impossible to obtain the exact Hessian matrix in practice. Therefore, we use the Fisher information matrix to approximate the Hessian matrix with Equation 3. Noting that we cannot access the previous training data, we use a small-scale validation set related to the previous task to compute the Fisher information matrix, which will be described in the experimental part.

After getting the Fisher information matrix, we fix the top $\rho\%$ parameters of each module with large values, and then set the update regions for the rest of the parameters as $[\theta_{\min}, \theta_{\max}]$, which are defined as:

$$\begin{aligned}\theta_{\min} &= \theta_0 - \lambda * |\theta_0|; \\ \theta_{\max} &= \theta_0 + \lambda * |\theta_0|,\end{aligned}\quad (4)$$

where θ_0 denotes the parameter values of the pre-trained NMT model, and λ is a hyper-parameter to control the size of the update region.

Output-based Method In this method, we hope that the model can automatically learn the update region of parameters based on the impact of parameters on the model output of the previous task. Intuitively, preserving the previously learned knowledge requires that the model output on the previous task should not change significantly after parameter update. Meanwhile, we also hope that the update regions should be as large as possible because the large regions will give the model more capacity to learn new tasks. Following the above intuition, we can define the learning objective:

$$\begin{aligned}\mathcal{L}(\theta) &= \frac{1}{N} \sum_{n=1}^N KL(p(\mathbf{y}|\mathbf{x}, \theta) || p(\mathbf{y}|\mathbf{x}, \theta_0)) \\ &\quad - \frac{\alpha}{M} \sum_{i=1}^M (\theta_i - \theta_{0,i})^2,\end{aligned}\quad (5)$$

where N denotes the amount of the training example, KL denotes the KL-divergence, α is a hyper-parameter to control the ratio of the two terms, and M denotes the total amount of the model parameters. The first term of the above objective function let the model output on the previous task stay as close as possible to the pre-trained model, which will discourage the parameter to change. While the second term encourages the model parameters to change more greedily, and maximize the size of regions as much as possible. These two items can be regarded as performing adversarial learning during learning the parameter update regions. Similar to

the curvature-based method, we use a small-scale validation set related to the previous task instead of the data of the new task as the training data.

After this learning process, we can get the updated model parameters θ_1 , then we define the update region $[\theta_{\min}, \theta_{\max}]$ as:

$$\begin{aligned}\theta_{\min} &= \min(\theta_0, \theta_1); \\ \theta_{\max} &= \max(\theta_0, \theta_1).\end{aligned}\quad (6)$$

3.2 Hard-Constrained Training

After finding the parameter update regions, we continually train the model parameters within this region to learn the new translation tasks, i.e., to find:

$$\begin{aligned}\theta^* &= \arg \min_{\theta} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}_N} \mathcal{L}_{CE}(\mathbf{y}|\mathbf{x}, \theta), \\ s.t. & \theta_{\min} \leq \theta \leq \theta_{\max},\end{aligned}\quad (7)$$

where \mathcal{D}_N denotes the training data of the new task. One may suspect that updating parameters only in a constrained region may lead to the insufficient ability of the model to fit the new translation tasks. However, considering the over-parameterization of the large-scale NMT model and the fact that many parameters have not learned sufficient knowledge (Hoefler et al., 2021), the hard-constrained training can give model the capability to adapt to new tasks in most cases, even though the update range of parameters is limited. The experimental results will also prove this.

4 Experiments

In this work, we perform continual learning on two tasks: the domain adaptation task and the more challenging language adaptation task. In the domain adaptation task, the language pairs of the new training data are already supported by the pre-trained NMT model, and the goal is to enable the model to support the translation of new domains. In the language adaptation task, the goal is to enable the model to support the translation of new languages, which are not seen during pretraining. Under the above two scenarios, we hope to retain the translation ability of the pre-trained MNMT model on the original translation task.

4.1 The Pre-trained NMT Model

In the experiments, we use the mBART50-nn (Tang et al., 2020) model as our pre-trained NMT model, which is available in the fairseq library. The

Task		Train	Valid	Test
Multilingual Translation	xx↔En	/	997	1012
Domain Adaptation (De→En)	IT	0.22M		
	Law	0.47M		
	Medical	0.25M	2000	2000
	Subtitles	0.5M		
	Koran	18K		
Language Adaptation	El↔En	1M	997	1012
	Sk↔En	1M		

Table 1: The statics of our datasets. The number in Valid/Test columns denotes the amount of sentence pairs in each domain or translation direction.

mBart50-nn is a many-to-many multilingual NMT model which can support the translation between 50 different languages. The encoder and decoder of mBART50-nn have 12 layers, respectively. The dimensionality of the model is set as 1024, and the attention module has 16 attention heads. The dimensionality of the embedding layer and hidden states is set as 1024, while the inner-layer of the feed-forward network has dimensionality as 4096. The attention module has 16 attention heads both in the encoder and decoder. Besides, the model has a shared source-target vocabulary of about 250k tokens, and the model uses learned positional embeddings with the max token length set as 1024.

4.2 Data Preparation

Multilingual Translation Task We test the performance of the model on the original translation task before and after continual learning, to verify whether the methods can preserve the previously learned knowledge. To achieve this, we use the FLORES-101 test sets (Goyal et al., 2021), which are extracted from English Wikipedia and cover a variety of different topics and domains. We test the translation performance of other 49 languages to and from English. For our method, we also use the FLORES-101 validation sets to compute the empirical Fisher information matrix and search the parameter update regions.

Domain Adaptation Task For the domain adaptation task, we use the data set proposed by Koehn and Knowles (2017) to simulate a diverse multi-domain setting. The data set includes parallel text in German and English, both of which have already been supported by the mBART50-nn model. The text is mainly from five diverse domains: IT, Law,

Medical, Subtitles, and Koran, which are available in OPUS (Aulamo and Tiedemann, 2019). We use the new split released by Aharoni and Goldberg (2020), and perform German to English translation (De→En) for this task. It should be noted that the De→En translation task has already been supported by the mBART50-nn model.

Language Adaptation Task For the language adaptation task, we adapt the model to support the Greek↔English (El↔En) and Slovak↔English (Sk↔En) translation directions. Greek is from an unseen language family and uses an unseen alphabet, which is very different from all the languages supported by mBART50-nn. In contrast, Slovak is from the same language family as Czech, which is already supported by the mBART50-nn model, so it is more familiar to the model. Therefore, we use them as the new languages because this can simulate most cases where we need language adaptation. We use the training data from OPUS-100 (Zhang et al., 2020) to train the model, and the validation/test sets from FLORES-101 to choose the checkpoint and test the performance.

Following Tang et al. (2020), we use the sentencepiece (Kudo and Richardson, 2018) model, which was trained using monolingual data for 100 languages from XLMR, to process all the above data.

4.3 Systems

We use the open-source toolkit called *Fairseq-py* (Ott et al., 2019) as our Transformer system. The following systems can be divided into two categories. The first category only focuses on either the previous task or the new task.

- **Scratch:** This system is trained from scratch only with the training data from the new translation task.
- **mBART50-nn** (Tang et al., 2020): The large scale pretrained NMT model. All the following systems are implemented based on this model.
- **mBART50-nn + Language-Specific Embedding (LSE)** (Berard, 2021): We insert a new language-specific embedding layer for the new languages and fine-tune these parameters with the new training data for 20k steps. The original parameters are kept fixed during training. For the language adaptation task, we use this system as the baseline and other methods are implemented based on this system.
- **Fine-tuning** (Luong and Manning, 2015): This system is trained based on the pretrained model only with the new training data.

Systems	Multilingual Translation			Domain Adaptation						Avg
	xx→En	En→xx	Avg1	IT	Law	Medical	Subtitles	Koran	Avg2	
Scratch	/	/	/	39.87	53.96	53.88	27.71	18.80	38.84	/
mBART50-nn	25.83	21.48	23.66	35.65	41.81	37.21	27.14	16.41	31.64	27.65
Fine-Tuning	19.27	1.64	10.46	45.99	57.02	54.71	31.98	21.58	42.26	26.36
Mixed-FT	19.33	2.05	10.69	45.8	56.93	53.69	31.5	21.59	41.90	26.30
KD	23.85	19.55	21.7	39.71	49.13	46.64	30.58	20.2	37.25	29.48
L2-Reg	24.03	19.85	21.94	41.03	50.88	49.19	30.1	20.5	38.34	30.14
EWC	24.19	20.29	22.24	41.02	50.25	49.2	30.59	19.68	38.15	30.19
FT-xattn	23.35	18.44	20.90	41.44	51.43	50.03	30.64	19.95	38.70	29.80
LFR-CM	25.09	20.36	22.73	41.73	51.24	50.28	30.96	20.98	39.04	30.89
LFR-OM	24.78	19.48	22.13	43.18	52.72	51.44	31.33	21.51	40.04	31.09

Table 2: The overall BLEU scores of the domain adaptation task. "xx" denotes the languages already supported by mBART50-nn. "Avg1" and "Avg2" denote the average BLEU scores on the multilingual translation task and domain adaptation task, respectively. "Avg" is computed by (Avg1+Avg2)/2 to indicate the balance between the previous and new tasks. The highest scores among all the continual learning methods are marked in bold.

Systems	Multilingual Translation			Language Adaptation					Avg
	xx→En	En→xx	Avg1	El→En	En→El	Sk→En	En→Sk	Avg2	
Scratch	/	/	/	24.93	25.39	28.17	26.59	26.27	/
mBART50-nn+LSE	25.83	21.48	23.66	26.57	16.06	35.82	28.6	26.76	25.21
Fine-Tuning	18.37	1.15	9.76	30.59	26.67	34.96	34.06	31.57	20.67
Mixed-FT	19.55	3.61	11.58	30.33	25.98	33.1	33.89	30.83	21.20
L2-Reg	25.87	18.34	22.11	26.67	18.67	35.41	30.62	27.84	24.97
EWC	25.99	18.2	22.10	26.88	18.5	35.31	30.41	27.78	24.94
FT-xattn	23.15	16.82	19.99	27.27	18.69	35.84	31.07	28.22	24.10
LFR-CM	26.42	18.06	22.24	28.55	20.81	36.1	31.41	29.22	25.73
LFR-OM	26.66	18.68	22.67	28.05	19.76	36.05	30.76	28.66	25.67

Table 3: The overall BLEU scores of the language adaptation task.

- **Mixed-FT**: We mix the small-scaled validation sets related to the previous task, i.e., the FLORES-101 validation sets of the languages supported by mBART50-nn, with the new training data to train the model jointly. We use the temperature-based sampling function to oversample the validation datasets and the temperature is set as 20 (Arivazhagan et al., 2019).

The second category contains several continual learning methods, which aim to get a good balance between previous and new tasks.

- **Knowledge Distillation (KD)** (Dakwale and Monz, 2017): Besides minimizing the training loss of the new task, this method also minimizes the cross-entropy between the output distribution of the mBART50-nn model and the network. The final objective is:

$$\mathcal{L}_{KD}(\theta) = \mathcal{L}_{CE}(\theta) + \alpha KL(p_{\theta_0} || p_{\theta}), \quad (8)$$

where α is the hyper-parameter which controls the contribution of the two parts.

- **L2-Reg** (Barone et al., 2017): This method adds

a L2-norm regularizations on the parameters:

$$\mathcal{L}_{L2}(\theta) = \mathcal{L}_{CE}(\theta) + \frac{\alpha}{M} \sum_{i=0}^M (\theta_i - \theta_{0,i})^2, \quad (9)$$

where i denotes the i -th parameter.

- **EWC** (Thompson et al., 2019): This method models the importance of the parameters with Fisher information matrix and puts more constraints on the important parameters to let them stay close to the original values. The training objective is:

$$\mathcal{L}_{EWC}(\theta) = \mathcal{L}_{CE}(\theta) + \frac{\alpha}{M} \sum_{i=0}^M F_i (\theta_i - \theta_{0,i})^2, \quad (10)$$

where F_i denotes the modeled importance for the i -th parameter.

- **FT-xattn** (Gheini et al., 2021): This method only fine-tunes the cross-attention parameters, which can also mitigate the forgetting problem.

- **LFR-CM**: This system is implemented as the curvature-based method. We set ρ as 75% and λ as 0.1 in the main experiments. We put more results

about the hyperparameters in the next section and Appendix.

- **LFR-OM:** This system is implemented as the output-based method. We set the hyper-parameter as 1. Other training details are listed below.

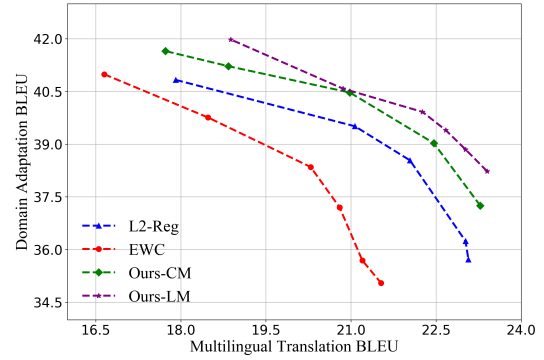
Training Details We set dropout as 0.3 and attention-dropout as 0.1. We employ the Adam optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.98$. We use the inverse square root learning scheduler and set the *warmup_steps* = 4000. For the curvature-based method (LFR-CM), we set $\rho\%$ as 75% and λ as 0.2. For the output-based method (LFR-OM), we set $lr = 2e - 4$ for the domain adaptation task, $lr = 4e - 5$ for the language adaptation, and train the model 5k steps to search the parameter update region. During continual learning, we set $lr = 5e - 4$ and train the model 30k steps for the domain adaptation task; we set $lr = 5e - 5$ and train the model 50k steps for the language adaptation task. We fix all the norm layers of the model in both of the two tasks. In the language adaptation task, we also fix the original and new embedding layers, which we find can also help alleviate the forgetting problem. All the systems are trained on 8 A100 GPUs with the update frequency 2. The max token is 1024 for each GPU. Besides, we use beam search with the size of 4 and length penalty as 1 during decoding.

4.4 Main Results

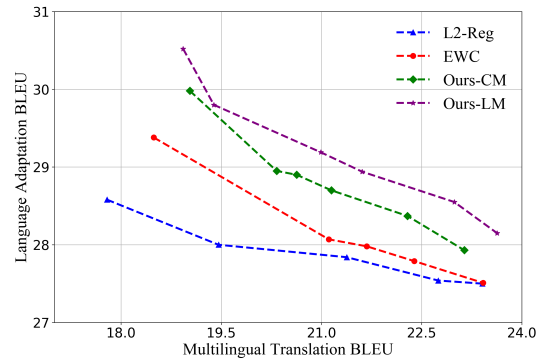
The final translation is evaluated using the 4-gram case-sensitive BLEU (Papineni et al., 2002) with the *SacreBLEU* tool (Post, 2018). Following Goyal et al. (2021), we report the SentencePiece BLEU which uses a SentencePiece tokenizer (SPM) with 256K tokens and then BLEU score is computed on the sentence-piece tokenized text ¹.

The main results of the domain adaptation task are in Table 2. Compared with the Scratch system, Fine-Tuning can greatly improve the performance of domain adaptation tasks, but it also suffers from catastrophic forgetting on the multilingual translation task. Besides, we observe that the forgetting problem in $En \rightarrow xx$ directions is more severer. After analyzing the model output, we find that the model cannot output other languages except for English for all previous translation directions. This may be because the target language of the domain adaptation task is only English. Among all the con-

¹BLEU+case.mixed+numrefs.1+smooth.exp +tok.spm+version.1.5.1



(a) Domain Adaptation Task



(b) Language Adaptation Task

Figure 2: The performance trade-off with different hyper-parameters. The x-axis and y-axis denote the BLEU of the previous and new tasks, respectively. The closer the point is to the upper right corner, the better the performance.

tinual learning methods, our method can get the best overall performance. The main results of the language adaptation task are in Table 3. By adding and updating the new language-specific embeddings (mBART50-nn+LSE), we can achieve good results except for the $En \rightarrow El$ directions. Greek is quite different from the previous languages, so it is more difficult for the model to learn it as the target language only by the existing knowledge. Just like the domain adaptation task, our method outperforms other continual learning methods.

5 Analysis

5.1 Effects of Different Hyper-parameters

For the regularization-based methods, the hyper-parameter α controls the performance trade-off between the previous and new tasks. The larger the hyper-parameters α is, the less decline of the BLEU on the original task will be, and the less improvement of the new task performance will be. As for our method, the proportion of model pa-

System	{Zh,Fr,De} ↔En Avg.	{El, Sk} ↔En Avg.	Zero Avg.
mBart50+LSE	35.24	26.76	5.80
Scratch	26.68	21.82	6.81
Fine-tuning	32.30	28.51	5.07
L2-Reg	30.09	22.53	2.40
EWC	26.32	21.02	3.06
LFR-CM	34.53	29.83	8.55
LFR-OM	34.88	28.94	8.12

Table 4: BLEU of the mixed-training experiments. "Zero Avg." denotes the average BLEU of {Zh, Fr, De}↔{El, Sk}.

rameters to be pruned has a similar effect. erasing more neurons will bring better results in the new task, but will also lead to more degradation in the original task. To better show the full performance trade-off, we conduct experiments with different hyper-parameters. We compare our method with the L2-Reg and EWC systems. For the L2-Reg and EWC method, we vary α from 0.001 to 1. For the curvature-based method, we fix $\rho\%$ as 75% and vary λ from 0.1 to 1. For the learning-based method, we vary the learning rate from $1e - 5$ to $1e - 4$, because we find that adjusting the learning rate is more effective than changing the hyper-parameter. The detailed settings of the hyper-parameters are put in the Appendix. The results are shown in Figure 2. It shows that our method outperforms L2-Reg and EWC at all the operating points significantly. Besides, it also shows that the output-based method is better than the curvature-based method.

5.2 Mixed-Training with Previous Data

In this work, we try to preserve all the previously learned knowledge and consider the situation that performing continual learning without access to any previous training data. But in practice, it is more likely that we just want to preserve some specific knowledge, e.g., for some specific languages, and we have some training data from or related to the previous training task. To prove the effectiveness of our method under this scenario, we conduct further experiments on the language adaptation task. We choose Chinese (Zh↔En), French (Fr↔En), and German (De↔En) from the languages supported by mBART50-nn as our target languages, on which we want to retain the translation performance. Then we collect the corresponding training data from the OPUS-100 dataset,

System	xx→En	En→xx	El↔En	Sk↔En
Fine-tuning	23.91	1.99	14.43	34.5
L2-Reg	25.95	20.53	21.32	32.58
EWC	26.02	20.46	21.21	32.04
LFR-CM	26.64	20.83	21.85	32.84
LFR-OM	26.69	21.39	22.03	33.03

Table 5: BLEU of the sequential language adaptation experiments.

which is much smaller in quantity than the previous training data of mBART50-nn. Last we continually train the model with the mixed data. We test the BLEU scores on the supervised directions and the zero-shot directions between the previous languages (Zh, Fr, De) and the new languages (El, Sk). The results are put in the Table 4, and the detailed hyper-parameter settings are put in the Appendix. We find that the regularization-based method fails to deal with this scenario, and they are even worse than the vanilla Fine-tuning method, which indicates that the soft constraints they put on the parameters are harmful to the model when some previous data is available. Compared with the fine-tuning method, our method can further reduce the forgetting problem with the previous training data and achieve better overall performance.

5.3 Sequential Language Adaptation

In this experiment, we perform the language adaptation task in the sequential training scenario. We first train the model with the EL↔En data and then with the Sk↔En data. For the L2-Reg and EWC methods, we use the model after training with the EL↔En data as the pretrained model for the Sk↔En task to compute the regularization loss. We recompute the Fisher information matrix for the EWC and LFR-CM method after the EL↔En training stage. Besides, we also recompute the update regions after the training of EL↔En for the Sk↔En task. The detailed hyper-parameter settings are put in the Appendix. We report the final results in Table 5. The results show that our method still outperforms other methods under this scenario.

6 Related Work

Recent work on continual learning of NMT can be divided into three categories: data memory based method, task-specific module based method, and regularization based method.

Data Memory Based Method The data memory based methods need to retain part or all of the

training data of the previous task. [Chu et al. \(2017\)](#) fine-tune the pretrained model with the mix of the previous and new data. [Bapna and Firat \(2019\)](#) propose an n-gram level retrieval approach to find the useful context to help translation. [Xu et al. \(2020\)](#) propose to use the similar translations from translation memory to boost the performance. [Liu et al. \(2021\)](#) utilize a bilingual dictionary to generate mixed-language sentences. Compared with these methods, our method doesn't need the previous data and thus is more flexible in practice.

Task-specific Module Based Method The task-specific module based methods need to assign the model parameters to different tasks. [Bapna and Firat \(2019\)](#) injects domain-specific adapter modules into each layer of the pretrained model, then they fix the model and only train the new adapters. Based on this, here are also some work ([Zeng et al., 2018, 2019](#); [Gu et al., 2019](#); [Cao et al., 2021](#)) that adds different kinds of task-specific modules to the model. Besides, some work tries to divide the model into different parts based on the parameter importance on each task ([Gu et al., 2021](#); [Liang et al., 2020](#)). Although they don't increase the model, these methods actually divide the model into different task-specific parts, which makes the model need to know which task the input comes from. Compared with these methods, our method doesn't introduce model separation explicitly.

Regularization Based Method The regularization based methods mitigate forgetting by introducing an additional penalty term in the learning objective. [Khayrallah et al. \(2018\)](#) and [Thompson et al. \(2019\)](#) add regularization terms to let the model parameters stay close to their original values. [Dakwale and Monz \(2017\)](#) minimize the cross-entropy between the output distribution of the pretrained model and the fine-tuned model. Different from the above work, our method constrains the model parameters within low forgetting risk regions to avoid the forgetting problem.

Besides, our work is also related to unstructured model pruning ([See et al., 2016](#); [Zhu and Gupta, 2018](#); [Frankle and Carbin, 2019](#)), because we both aim to search and operate on the unimportant parameters. The difference is that our method is to find an update region around the initial parameters, while model pruning should directly remove some parameters, that is, set them to zero.

7 Conclusion

In this work, we propose a continual learning method for NMT by constraining model parameters within low forgetting risk regions. We propose two methods to find such regions, the first is based on the curvature of the loss function and the second is based on the model output of the previous task. Then we can continually train the model parameters within this region. The experimental results on the domain adaptation and language adaptation tasks prove that our method can achieve significant improvements over several strong baselines.

Limitations

Because our method does not introduce additional parameters or use the previous data, the overall performance of our method is weaker than the data memory based methods and task-specific module based methods. It is difficult for our method to achieve the same performance as the Fine-tuning method in the new task without causing catastrophic forgetting.

Acknowledgements

We thank all the anonymous reviewers for their insightful and valuable comments.

References

- Roei Aharoni and Yoav Goldberg. 2020. [Unsupervised domain clusters in pretrained language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7747–7763.
- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George F. Foster, Colin Cherry, Wolfgang Macherey, Zhifeng Chen, and Yonghui Wu. 2019. [Massively multilingual neural machine translation in the wild: Findings and challenges](#). *CoRR*, abs/1907.05019.
- Mikko Aulamo and Jörg Tiedemann. 2019. [The OPUS resource repository: An open package for creating parallel corpora and machine translation services](#). In *Proceedings of the 22nd Nordic Conference on Computational Linguistics, NoDaLiDa 2019, Turku, Finland, September 30 - October 2, 2019*, pages 389–394.
- Ankur Bapna and Orhan Firat. 2019. [Non-parametric adaptation for neural machine translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational*

- Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 1921–1931.
- Antonio Valerio Miceli Barone, Barry Haddow, Ulrich Germann, and Rico Sennrich. 2017. [Regularization techniques for fine-tuning in neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 1489–1494.
- Alexandre Berard. 2021. [Continual learning in multilingual NMT via language-specific embeddings](#). In *Proceedings of the Sixth Conference on Machine Translation, WMT@EMNLP 2021, Online Event, November 10-11, 2021*, pages 542–565.
- Yue Cao, Hao-Ran Wei, Boxing Chen, and Xiaojun Wan. 2021. [Continual learning for neural machine translation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 3964–3974.
- Chenhui Chu, Raj Dabre, and Sadao Kurohashi. 2017. An empirical comparison of simple domain adaptation methods for neural machine translation. *arXiv preprint arXiv:1701.03214*.
- Praveen Dakwale and Christof Monz. 2017. Fine-tuning for neural machine translation with limited degradation across in-and out-of-domain data. *Proceedings of the XVI Machine Translation Summit*, page 117.
- Carlos Escolano, Marta R. Costa-jussà, and José A. R. Fonollosa. 2021. [From bilingual to multilingual neural-based machine translation by incremental training](#). *J. Assoc. Inf. Sci. Technol.*, 72(2):190–203.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Michael Auli, and Armand Joulin. 2021. [Beyond english-centric multilingual machine translation](#). *J. Mach. Learn. Res.*, 22:107:1–107:48.
- Jonathan Frankle and Michael Carbin. 2019. [The lottery ticket hypothesis: Finding sparse, trainable neural networks](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Mozhdeh Gheini, Xiang Ren, and Jonathan May. 2021. [Cross-attention is all you need: Adapting pretrained transformers for machine translation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 1754–1765.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’Aurelio Ranzato, Francisco Guzmán, and Angela Fan. 2021. [The FLORES-101 evaluation benchmark for low-resource and multilingual machine translation](#). *CoRR*, abs/2106.03193.
- Shuhao Gu and Yang Feng. 2020. [Investigating catastrophic forgetting during continual training for neural machine translation](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4315–4326, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Shuhao Gu, Yang Feng, and Qun Liu. 2019. Improving domain adaptation translation with domain invariant and specific information. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT*, pages 3081–3091.
- Shuhao Gu, Yang Feng, and Wanying Xie. 2021. [Pruning-then-expanding model for domain adaptation of neural machine translation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 3942–3952.
- Torsten Hoefler, Dan Alistarh, Tal Ben-Nun, Nikoli Dryden, and Alexandra Peste. 2021. [Sparsity in deep learning: Pruning and growth for efficient inference and training in neural networks](#). *J. Mach. Learn. Res.*, 22:241:1–241:124.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda B. Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. [Google’s multilingual neural machine translation system: Enabling zero-shot translation](#). *Trans. Assoc. Comput. Linguistics*, 5:339–351.
- Huda Khayrallah, Brian Thompson, Kevin Duh, and Philipp Koehn. 2018. [Regularized training objective for continued training for domain adaptation in neural machine translation](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation, NMT@ACL 2018, Melbourne, Australia, July 20, 2018*, pages 36–44.
- Yoon Kim and Alexander M. Rush. 2016. [Sequence-level knowledge distillation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 1317–1327.
- Wei-Jen Ko, Ahmed El-Kishky, Adithya Renduchintala, Vishrav Chaudhary, Naman Goyal, Francisco Guzmán, Pascale Fung, Philipp Koehn, and Mona T. Diab. 2021. [Adapting high-resource NMT models to translate low-resource related languages without parallel data](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and*

- the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021, pages 802–812.
- Philipp Koehn and Rebecca Knowles. 2017. [Six challenges for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation, NMT@ACL 2017, Vancouver, Canada, August 4, 2017*, pages 28–39.
- Taku Kudo and John Richardson. 2018. [Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018: System Demonstrations, Brussels, Belgium, October 31 - November 4, 2018*, pages 66–71.
- Jianze Liang, Chengqi Zhao, Mingxuan Wang, Xipeng Qiu, and Lei Li. 2020. Finding sparse structure for domain specific neural machine translation. *AAAI 2021*.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *Trans. Assoc. Comput. Linguistics*, 8:726–742.
- Zihan Liu, Genta Indra Winata, and Pascale Fung. 2021. [Continual mixed-language pre-training for extremely low-resource neural machine translation](#). In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, pages 2706–2718.
- Minh-Thang Luong and Christopher D Manning. 2015. Stanford neural machine translation systems for spoken language domains. In *Proceedings of the International Workshop on Spoken Language Translation*, pages 76–79.
- Alexander Ly, Maarten Marsman, Josine Verhagen, Raoul PPP Grasman, and Eric-Jan Wagenmakers. 2017. A tutorial on fisher information. *Journal of Mathematical Psychology*, 80:40–55.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Demonstrations*, pages 48–53.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Abigail See, Minh-Thang Luong, and Christopher D. Manning. 2016. [Compression of neural machine translation models via pruning](#). In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, pages 291–301, Berlin, Germany. Association for Computational Linguistics.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. [Multilingual translation with extensible multilingual pretraining and finetuning](#). *CoRR*, abs/2008.00401.
- Brian Thompson, Jeremy Gwinnup, Huda Khayrallah, Kevin Duh, and Philipp Koehn. 2019. Overcoming catastrophic forgetting during domain adaptation of neural machine translation. In *NAACL-HLT 2019*, pages 2062–2068.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems*, pages 5998–6008.
- Jitao Xu, Josep Maria Crego, and Jean Senellart. 2020. [Boosting neural machine translation with similar translations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 1580–1590.
- Jiali Zeng, Yang Liu, Jinsong Su, Yubin Ge, Yaojie Lu, Yongjing Yin, and Jiebo Luo. 2019. [Iterative dual domain adaptation for neural machine translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 845–855.
- Jiali Zeng, Jinsong Su, Huating Wen, Yang Liu, Jun Xie, Yongjing Yin, and Jianqiang Zhao. 2018. [Multi-domain neural machine translation with word-level domain context discrimination](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 447–457.
- Biao Zhang, Philip Williams, Ivan Titov, and Rico Senrich. 2020. [Improving massively multilingual neural machine translation and zero-shot translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 1628–1639.
- Michael Zhu and Suyog Gupta. 2018. [To prune, or not to prune: Exploring the efficacy of pruning for](#)

model compression. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Workshop Track Proceedings*. OpenReview.net.

A Appendix

A.1 The mBART50-nn Model

The mBart50-nn model is a many-to-many multilingual NMT model which can support the translation between 50 different languages. The encoder and decoder of mBART50-nn have 12 layers. The dimensionality of the model is set as 1024, while the inner-layer of the feed-forward network has dimensionality as 4096. The attention module has 16 attention heads both in the encoder and decoder. Besides, the model has a shared source-target vocabulary of about 250k tokens, and the model uses learned positional embeddings with the max token length set as 1024.

A.2 Hyper-parameter Settings for the Analysis

In this section, we report the detailed hyper-parameter settings in our experiments.

For the main experiments, we set α as 1 for the KD method, 0.01 for the L2-Reg method, and 0.05 for the EWC method. We set $\rho\%$ as 75% and λ as 0.1 for the LFR-CM method, and α as 1 for the LFR-OM method.

For the experiments studying the effects of different hyper-parameters (Section 5.1), we tried the following hyper-parameters:

- L2-Reg (α): 0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 1.
- EWC (α): 0.001, 0.01, 0.05, 0.1, 0.5, 1, 5.
- LFR-CM (λ): 0.05, 0.1, 0.2, 0.4, 0.6, 0.8, 1.
- LFR-OM (lr): $1e - 5$, $2e - 5$, $3e - 5$, $4e - 5$, $5e - 5$, $1e - 4$.

For the mixed-training with previous data experiments (Section 5.2), we set α as 0.01 for the L2-Reg and EWC methods, λ as 0.4 for the LFR-CM method, and lr as $1e - 4$ for the LFR-OM method.

For the sequential language adaptation experiments, we set α as 0.5 for the L2-Reg and EWC methods, λ as 0.05 for the LFR-CM method, and lr as $2e - 5$ for the LFR-OM method.