

Investigating Catastrophic Forgetting During Continual Training for Neural Machine Translation

Shuhao Gu^{1,2}, 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
{gushuhao19b,fengyang}@ict.ac.cn

Abstract

Neural machine translation (NMT) models usually suffer from catastrophic forgetting during continual training where the models tend to gradually forget previously learned knowledge and swing to fit the newly added data which may have a different distribution, e.g. a different domain. Although many methods have been proposed to solve this problem, we cannot get to know what causes this phenomenon yet. Under the background of domain adaptation, we investigate the cause of catastrophic forgetting from the perspectives of modules and parameters (neurons). The investigation on the modules of the NMT model shows that some modules have tight relation with the general-domain knowledge while some other modules are more essential in the domain adaptation. And the investigation on the parameters shows that some parameters are important for both the general-domain and in-domain translation and the great change of them during continual training brings about the performance decline in general-domain. We conduct experiments across different language pairs and domains to ensure the validity and reliability of our findings.

1 Introduction

Neural machine translation (NMT) models (Kalchbrenner and Blunsom, 2013; Cho et al., 2014; Sutskever et al., 2014; Bahdanau et al., 2015; Gehring et al., 2017; Vaswani et al., 2017) have achieved state-of-the-art results and have been widely used in many fields. Due to numerous parameters, NMT models can only play to their advantages based on large-scale training data. However, in practical applications, NMT models often need to perform translation for some specific domain with only a small quantity of in-domain data available. In this situation, continual training (Luong and Manning, 2015), which is also referred to as fine-tuning, is often employed to improve the in-domain translation performance. In this method, the model is first trained with large-scale general-domain training data and then continually trained with the in-domain data. With this method, the in-domain performance can be improved greatly, but unfortunately, the general-domain performance decline significantly, since NMT models tend to overfit to frequent observations (e.g. words, word co-occurrences, translation patterns) in the in-domain data but forget previously learned knowledge. This phenomenon is called catastrophic forgetting. Figure 1 shows the performance trends on the in-domain and general-domain.

Many methods have been proposed to address the catastrophic forgetting problem under the scheme of fine-tuning. Freitag and Al-Onaizan (2016) ensembles the general-domain model and the fine-tuned model together so that the integrated model can consider both domains. Dakwale and Monz (2017) introduces domain-specific output layers for both of the domains and thus the domain-specific features of the two domains can be well preserved. Thompson et al. (2019), Barone et al. (2017), and Khayrallah et al. (2018) propose regularization-based methods that introduce an additional loss to the original objective to help the model trade off between the general-domain and in-domain. All these methods show their effectiveness and have mitigated the performance decline on general-domain, but we still don't know what happened inside the model during continual training and why these methods can alleviate the

*Corresponding author: Yang Feng.

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.




Figure 1: Performance trend for continual training of NMT on the Chinese to English translation task. The model is first trained with the news domain data and then continually trained with the laws domain data (from epoch 0). The y-axis represents the BLEU scores on the test sets of the corresponding domains.

catastrophic forgetting problem. The study on these can help to understand the working mechanism of continual training and inspire more effective solutions to the problem in return.

Given above, in this paper, we focus on the catastrophic forgetting phenomenon and investigate the roles of different model parts during continual training. To this end, we explore the model from the granularities of modules and parameters (neurons). In the module analyzing experiments, we operate the model in two different ways, by freezing one particular module or freezing the whole model except for this module. We find that different modules preserve knowledge for different domains. In the parameter analyzing experiments, we erase parameters according to their importance which is evaluated by the Taylor expansion-based method (Molchanov et al., 2017). According to the experimental results, we find that some parameters are important for both of the general-domain and in-domain and meanwhile they change greatly during domain adaptation which may result in catastrophic forgetting. To ensure the validity and reliability of the findings, we conduct experiments over different language pairs and domains.

Our main contributions are summarized as follows:

- We propose two analyzing methods to explore the model from the perspectives of modules and parameters, which can help us understand the cause of catastrophic forgetting during continual training.
- We find that some modules tend to maintain the general-domain knowledge while some modules are more essential for adapting to the in-domain.
- We find that some parameters are important for both of the general-domain and in-domain, and their over-change in values may result in performance slipping.

2 Background

In our work, we apply our method under the framework of *Transformer* (Vaswani et al., 2017) which will be briefly introduced here. We will also introduce the terminology and symbols used in the rest of the paper. We denote the input sequence of symbols as $\mathbf{x} = (x_1, \dots, x_J)$, the ground-truth sequence as $\mathbf{y}^* = (y_1^*, \dots, y_{I^*}^*)$ and the translation as $\mathbf{y} = (y_1, \dots, y_I)$.

The Encoder The encoder is composed of N identical layers. Each layer has two sublayers. The first is a multi-head *self-attention sublayer* (abbreviated as SA) and the second is a fully connected feed-forward network, named *FFN sublayer* (abbreviated as FFN). Both of the sublayers are followed by a residual connection operation and a layer normalization operation. The input sequence \mathbf{x} will be first fed into the *embedding layer* (abbreviated as Emb) and 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 input sequence of vectors will be fed into the encoder and the output of the N -th

layer is taken as source hidden states and we denote it as \mathbf{H} .

The Decoder The decoder is also composed of N identical layers. In addition to the same kinds of two sublayers in each encoder layer, a third sublayer is inserted between them, named *cross-attention sublayer* (abbreviated as CA), which performs multi-head attention over the output of the encoder stack. The final output of the N -th layer gives the target hidden states, denoted as $\mathbf{S} = [\mathbf{s}_1; \dots; \mathbf{s}_I]$, where \mathbf{s}_i is the hidden states of y_i .

The Objective The output of the decoder will be fed into the *output layer* (abbreviated as Out) and we can get the predicted probability of the i -th target word over the target vocabulary by performing a linear transformation and a softmax operation to the target hidden states:

$$p(y_i | \mathbf{y}_{<i}, \mathbf{x}) \propto \exp(\mathbf{W}_o \mathbf{s}_i + \mathbf{b}_o) \quad (1)$$

where $\mathbf{W}_o \in \mathbb{R}^{d_{model} \times |V_t|}$ is the parameter matrix of the *output layer*, and $|V_t|$ is the size of target vocabulary. The model is optimized by minimizing a cross-entropy loss which maximizes the probability of the ground-truth sequence with the teacher-forcing training:

$$\mathcal{L} = -\frac{1}{I} \sum_{i=1}^I \log p(y_i^* | \mathbf{y}_{<i}, \mathbf{x}) \quad (2)$$

where I is the length of the target sentence. A detailed description can be found in Vaswani et al. (2017).

3 Module Analysis

In this work, we will investigate the cause of the catastrophic forgetting phenomenon from the perspectives of modules and parameters. As we know, the structure and the various kinds of modules in it has a large impact on translation. In this section, therefore, we will study the function of different modules by isolating them during continual training. The study on the parameters will be discussed in Section 4.

3.1 Analyzing Strategies

We propose two training strategies during the continual training process. The first is to freeze the target module but update the rest of the model, called module-frozen training; the second, in contrast, is to update the target module but freeze the rest of the model, called module-updated training. In this way, we can estimate the ability of each module for preserving the general-domain knowledge and for adapting to the in-domain. In addition, we group the target module based on two criteria. The first is based on its position, e.g., we freeze or update all the sublayers in the *first two layers* of the encoder or the *last two layers* of the decoder; the second is based on its type, e.g., we freeze or update the *self-attention sublayers* or the *cross-attention sublayers* in all the decoder layers.

3.2 Experiments

3.2.1 Data Preparing

We conduct experiments on the following data sets across different languages and domains.

Chinese→English. For this task, general-domain data is from the LDC corpus¹ that contains 1.25M sentence pairs. The LDC data is mainly related to the **News** domain. MT06 and MT02 are chosen as the development and test data, respectively. We choose the parallel sentences with the domain label **Laws** from the UM-Corpus (Tian et al., 2014) as our in-domain data. We filter out repeated sentences and chose 206K, 2K, and 2K sentences randomly as our training, development, and test data, respectively. We tokenize and lowercase the English sentences with Moses² scripts. For the Chinese data, we perform word segmentation by using Stanford Segmenter³.

¹The corpora include LDC2002E18, LDC2003E07, LDC2003E14, Hansards portion of LDC2004T07, LDC2004T08, and LDC2005T06.

²<http://www.statmt.org/moses/>

³<https://nlp.stanford.edu/>

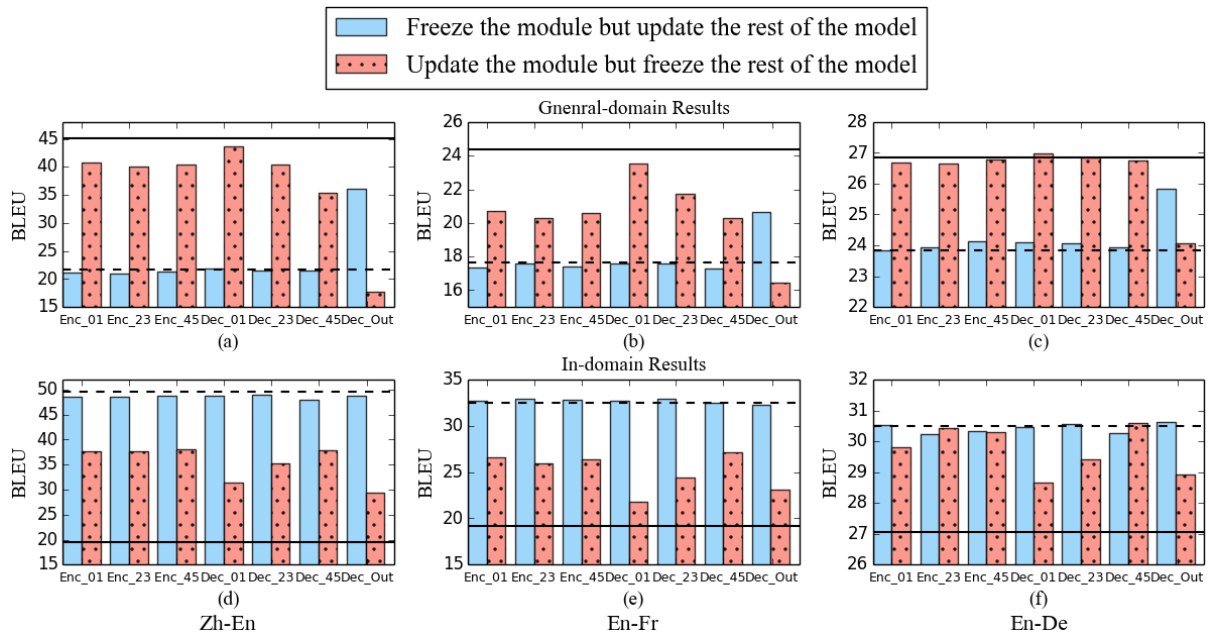


Figure 2: The BLEU of the experiments based on the position of the target module when only freezing the target module (left blue bars) or only updating the target module (right red bars). The first three subfigures (a-c) represent the general-domain BLEU and the last three subfigures (d-f) represent the in-domain BLEU. The horizontal solid line in each subfigure denotes the BLEU before continual training while the dashed line denotes the BLEU after continual training. The x-axis denotes the module at different positions, e.g., 'Enc_01' denotes the first two layers of the encoder, 'Dec_out' denotes the output layer of the decoder.

English→**French**. For this task, we choose 600K sentences randomly from the WMT 2014 corpus as our general-domain data, which are mainly related to the **News** domain. We choose newsdev2013 and newstest2013 as our development and test data, respectively. The in-domain data with 53K sentences is from WMT 2019, and it is mainly related to the **Biomedical** domain. We choose 1K and 1K sentences randomly from the corpora as our development and test data, respectively. We tokenize and truncate the corpora.

English→**German**. For this task, the general-domain data is from the WMT 2016 English to German translation task which is mainly **News** texts. It contains about 4.5M sentence pairs. We choose the news-test 2013 for validation and news-test 2014 for the test. For the in-domain data, we use the parallel training data from the IWSLT 2015 which is mainly from the **Spoken** domain. It contains about 194K sentences. We choose the 2012dev for validation and 2013tst for the test. We tokenize and truncate the corpora.

Besides, integrating operations of 32K, 16K, and 30K are performed to learn BPE (Sennrich et al., 2016) on the general-domain data and then applied to both the general-domain and in-domain data. The dictionaries are also built based on the general-domain data.

3.2.2 Systems

We use the open-source toolkit called *Fairseq-py* (Ott et al., 2019) released by Facebook as our Transformer system. We train the model with two sets of parameters. For the quantitatively analyzing experiments, the system is implemented as the base model configuration in Vaswani et al. (2017) strictly. For the visualizing experiments, we employ a *tiny* setting: the embedding size is set to 32, the FFN size is set to 64, and the rest is the same with the base model.

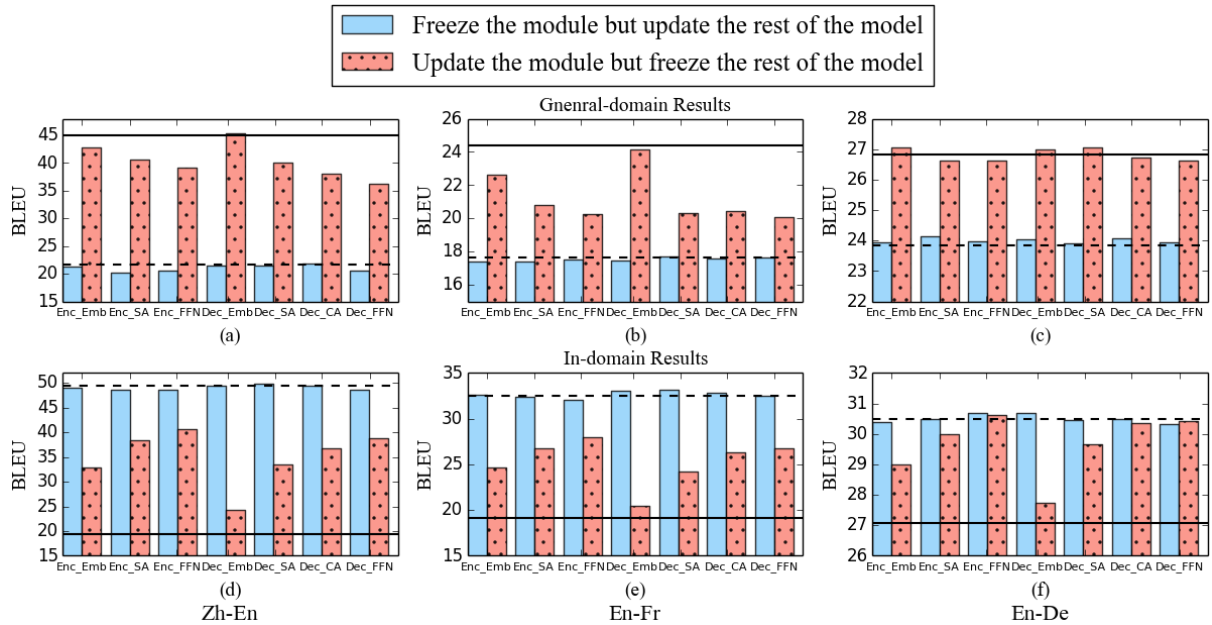


Figure 3: The BLEU of the experiments based on the type of the target module. The x-axis denotes the module of different types, where 'Emb', 'SA', 'CA', and 'FFN' denote the embedding layer, self-attention sublayer, cross-attention sublayer, and FFN sublayer, respectively. The rest of the denotations is just the same as in Figure 2.

3.2.3 The Results Based on the Position of the Target Module

The results of the experiments, which freeze or update the target module based on its position, are shown in Figure 2, where the left blue bars correspond to the results of the module-frozen experiments and the right red bars correspond to the results of module-updated experiments. As can be seen from the results of the module-frozen experiments, for different positions of the encoder and decoder, freezing any single module has a small impact on both the general-domain and in-domain BLEU, when compared with the normal continual training. However, freezing the output layer of the decoder can better alleviate the catastrophic forgetting phenomenon without degrading the in-domain BLEU, which implies that the output layer is more capable of maintaining general-domain knowledge. As for the results of the module-updated experiments, firstly, only updating the encoder can bring larger improvements on the in-domain compared with only updating the decoder, which implies that the encoder is more essential for adapting to the in-domain. Secondly, higher layers of the decoder tend to adapt to the in-domain, which, however, is not the case in the encoder. Lastly, only updating the output layer results in a bad performance on both domains, which indicates that the output layer highly depends on its surrounding modules.

3.2.4 The Results Based on the Type of the Target Module

Figure 3 shows the results of the experiments based on the type. The results of the module-frozen experiments (the left solid bars) show that freezing any type of module has little effect on the final results. As for the results of the module-updated experiments (the right dotted bars), we find that both the encoder embedding layer and the decoder embedding layer tend to preserve the general-domain knowledge, meanwhile, they are bad at adapting to the in-domain. In contrast, the FFN layers are more essential for adapting to the in-domain, though it is also easier for them to cause the catastrophic forgetting problem in general-domain.

3.3 Summary and Inspiration

In this section, we analyze the impacts of different modules on the general-domain and in-domain translation during continual training. we find that some modules, e.g., the output layer, the embedding layers, tend to preserve the general-domain knowledge; some modules, e.g., the encoder layers, the FFN layers,

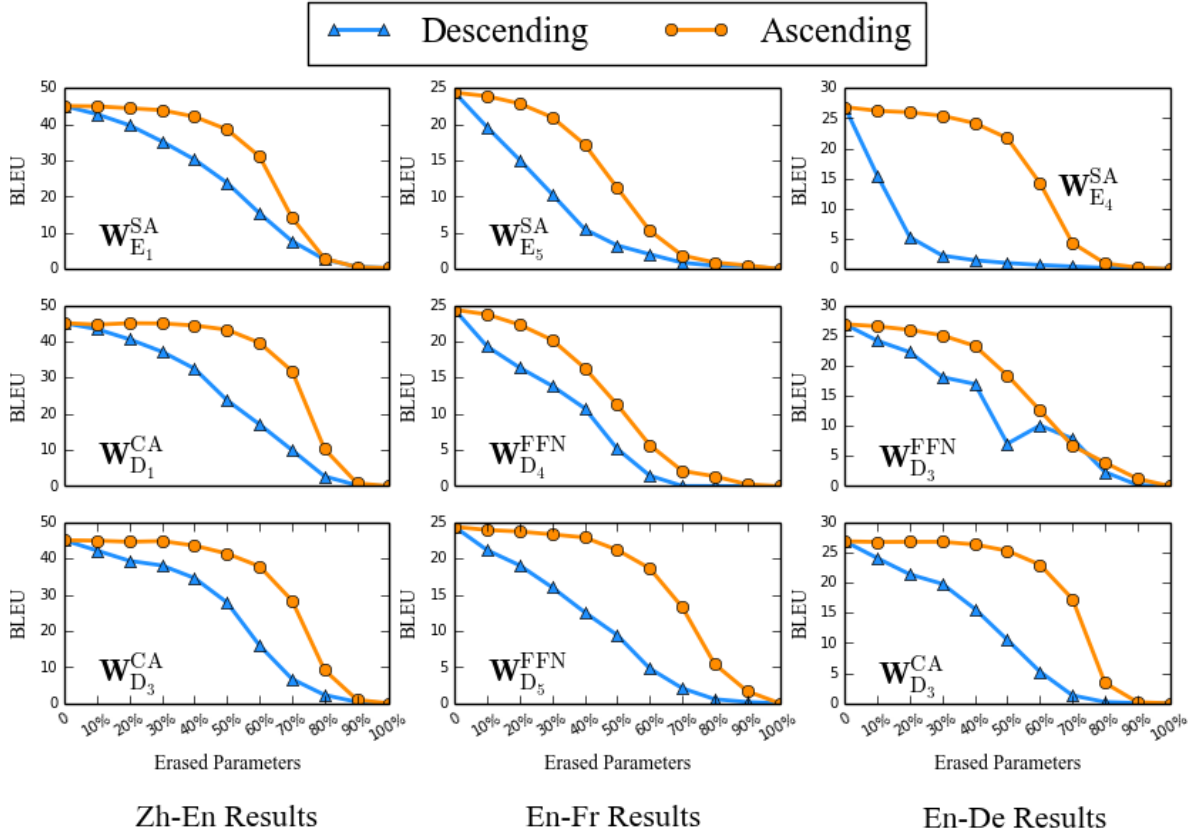


Figure 4: BLEU scores of the parameter erasure experiments with the model trained and tested on the general-domain. The blue curve with triangles stands for erasing parameters in descending order according to the importance while the orange curve with dots stands for erasing parameters in ascending order.

are more essential for adapting to the in-domain. Inspired by our findings, we can freeze those modules which are more important for general-domain during continual training to avoid catastrophic forgetting. To reduce the potential loss on the in-domain translation, we can extend the size of those frozen layers or add domain-specific layers parallel with them, and only update those newly added parameters during continual training.

4 Parameter Analysis

When the structure of the NMT model is fixed, the performance is determined by its parameters. During continual training, the change of the training data distribution makes the distribution of model parameters vary accordingly, which leads to the variation of translation performance on both the general-domain and in-domain. Motivated by this, therefore, we want to investigate the changing trend of model parameters during continual training, aiming to figure out the influence of different parameters. Intuitively, different parameters may have different importance for the NMT model, thus we firstly propose a method for evaluating the importance of parameters in this section. Then, we erase the model parameters increasingly according to the importance to see the change of BLEU scores. Finally, we visualize the parameters and measure their variations in the values to establish the connection between the parameter variation and the catastrophic forgetting phenomenon.

4.1 Importance Evaluation Method

To evaluate the importance of each parameter, we adopt a criterion based on the Taylor expansion (Molchanov et al., 2017), where we directly approximate the change in loss when removing a particular parameter. Let h_i be the output produced from parameter i and H represents the set of other

parameters. Assuming the independence of each parameter in the model, the change of loss when removing a certain parameter can be represented as:

$$|\Delta\mathcal{L}(h_i)| = |\mathcal{L}(H, h_i = 0) - \mathcal{L}(H, h_i)| \quad (3)$$

where $\mathcal{L}(H, h_i = 0)$ is the loss value if the parameter i is pruned and $\mathcal{L}(H, h_i)$ is the loss if it is not pruned. For the function $\mathcal{L}(H, h_i)$, its Taylor expansion at point $h_i = a$ is:

$$\mathcal{L}(H, h_i) = \sum_{n=0}^N \frac{\mathcal{L}^n(H, a)}{n!} (h_i - a)^n + R_N(h_i) \quad (4)$$

where $\mathcal{L}^n(H, a)$ is the n -th derivative of $\mathcal{L}(H, h_i)$ evaluated at point a and $R_n(h_i)$ is n -th remainder. Then, approximating $\mathcal{L}(H, h_i = 0)$ with a first-order Taylor polynomial where h_i equals zero, we get:

$$\mathcal{L}(H, h_i = 0) = \mathcal{L}(H, h_i) - \frac{\partial\mathcal{L}(H, h_i)}{\partial h_i} h_i - R_1(h_i) \quad (5)$$

The remainder R_1 can be represented in the form of Lagrange:

$$R_1(h_i) = \frac{\partial^2\mathcal{L}(H, h_i)}{\partial^2\delta h_i} h_i^2, \quad (6)$$

where $\delta \in (0, 1)$. Considering the use of ReLU activation function in the model, the first derivative of loss function tends to be constant, so the second order term tends to be zero in the end of training. Thus, we can ignore the remainder and get the importance evaluation function as follows:

$$\Theta_{\text{TE}}(h_i) = |\Delta\mathcal{L}(h_i)| = \left| \frac{\partial\mathcal{L}(H, h_i)}{\partial h_i} h_i \right| \quad (7)$$

Intuitively, this criterion disvalues the parameters that have an almost flat gradient of the objective function. In practice, we need to accumulate the product of the activation and the gradient of the objective function w.r.t to the activation, which is easily computed during back-propagation. Finally, the evaluation function is shown as:

$$\Theta_{\text{TE}}(h_i^l) = \frac{1}{T} \sum_t \left| \frac{\delta\mathcal{L}(H, h_i^l)}{\delta h_i^l} h_i^l \right|, \quad (8)$$

where h_i^l is the activation value of the i -th parameter of l -th module and T is the number of the training examples. The criterion is computed on the all training data and averaged over T .

4.2 Parameter Erasure

To prove that different parameters are indeed different in importance and verify the effectiveness of the proposed criterion for evaluating the importance of parameters, we conduct parameter erasure experiments to see the change of BLEU scores. For each parameter matrix \mathbf{W} of the model, we rank all the parameters in it according to the proposed criterion and get the ranked list. Then we erase the model parameters (i.e., masking them to zero) increasingly according to the ranked list in ascending or descending order. If the evaluation results are correct, erasing the parameters in the descending order will hurt the translation quality more than erasing in the ascending order.

4.3 Experiments

The training data and experimental systems are just the same as in section 3.2.

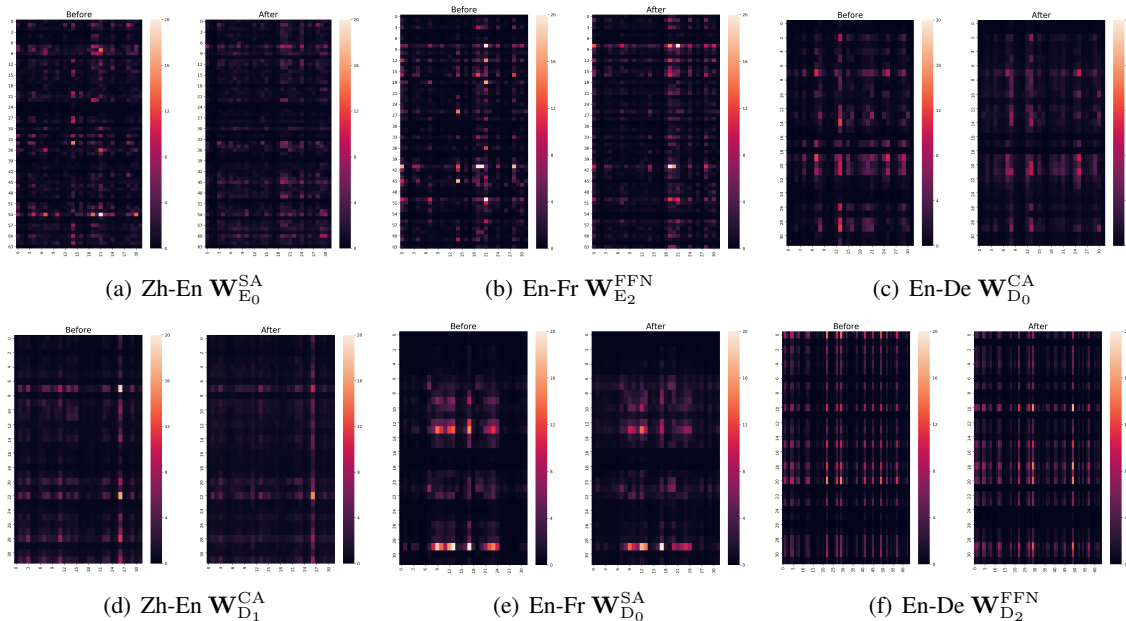


Figure 5: Illustration of parameter importance before and after continual training. Each subfigure represents a parameter matrix of the model, where each square in it represents a parameter. The lighter the square in color, the more important the parameter.

4.3.1 Parameter Erasure for the General-Domain Model

Figure 4 shows part of the results of the parameter erasure experiments for all the three language pairs on the general-domain test sets. In most cases, not surprisingly, both of the two curves decrease monotonically and the curve of erasing the parameters in ascending order according to the importance is higher than the other one. Based on this, we can conclude that some parameters are indeed more important than others for the whole model and have a larger impact on translation quality. In some cases, however, we also found some abnormal results: the two curves intersect or the curve is not monotonous. Although the independence of each parameter has been assumed during the derivation of the proposed method, it isn't always this case in practice, which will lead to these abnormal results. Overall, the proposed criterion can identify those important parameters in most cases, so that we can make use of it to analyze the behavior of the parameters during the continual training process.

4.3.2 Parameter Importance Visualization

The distribution of parameter importance can be seen as an important feature, which can imply the inner change of the model. Therefore we try to visualize the importance distribution before and after continual training. To achieve this, we retrain models with the tiny parameter setting, considering the convenience for presentation. The model is first trained with the general-domain data and we get the importance evaluation matrices for all the modules, which are computed using the general-domain data. Then the model is continually trained with the in-domain data and we also get the importance evaluation matrices for the in-domain, which, in contrast, are computed with the in-domain data. These two matrices are then visualized with heatmaps.

Figure 5 shows the results. Firstly, the more important parameters, which are lighter in the figures, lie in certain rows or columns of the target parameter matrix. Considering that the parameters in the same row or column are connected to the same neurons in the former or latter layer, we argue that some neurons are more important for the model, which is consistent with the conclusion of Bau et al. (2019). Secondly, the parameter distribution before and after continual training is very similar; most of the lighter squares in the left pictures are still lighter than other squares in the right pictures. This observation result indicates that the parameters which are important for the general-domain translation still have larger impacts on the in-domain translation after the continual learning process in most cases. Lastly, there




Figure 6: The value change of the model parameters. The x-axis stands for the importance interval of the model before fine-tuning, e.g., [20%, 30%) means the importance of parameters is between top 20% to 30%. The y-axis represents the Euclidean distance between the parameters of the model before and after continual training.

are still lots of less important parameters after continual training, which have limited influence on the in-domain translation.

4.3.3 Parameter Variation across Domains

From the results above, we find that the important parameters for the general-domain still have large impacts on the in-domain translation after continual training. To figure out the variation of parameters with different importance, we compute the average Euclidean distance for all the parameters in the model before and after the continual learning process:

$$\text{distance} = \frac{1}{N} \sum_{i=1}^N \sqrt{(\mathbf{W}_i^G - \mathbf{W}_i^I)^2}, \quad (9)$$

where N denotes the number of different modules in the model; \mathbf{W}_i denotes the parameter matrix for the i -th module; G and I denote the general-domain and in-domain, respectively. Then all the parameters in each module are ranked and divided into ten groups according to their importance. The result of the average Euclidean distance of parameters in each importance interval is shown in Figure 6. We find that the top parameters change more greatly than the less important parameters. Considering their impacts on the translation, we conclude that it is because of the excessive change of the important parameters that causes the catastrophic forgetting phenomenon.

4.4 Summary and Inspiration

In this section, we propose a method for evaluating the parameter importance. Then through the parameter erasure experiments, we find that some parameters are more important and have a greater influence on the output. Next based on the importance distribution visualization results, we find that some parameters are important for both of the domains. Finally, the average Euclidean distance of model parameters before and after continual learning is calculated and we find that the important parameters change more greatly, which causes the catastrophic forgetting problem. Inspired by our findings, we can freeze part of those important parameters during continual training to avoid catastrophic forgetting. Besides, we can retrain those unimportant parameters to further improve the in-domain translation.

5 Related Work

Analyzing Work Recently, much work has been concerned with analyzing and evaluating the NMT model from different perspectives. Shi et al. (2016) investigates how NMT models output target strings of appropriate lengths. Ding et al. (2017) analyzes the contribution of each contextual word to arbitrary hidden states. Qi et al. (2018) analyzes when the pre-trained word embeddings can help in NMT tasks. Voita et al. (2019) analyzes the importance of different attention heads. Bau et al. (2019) investigates the importance and function of different neurons in NMT. Wuebker et al. (2018) finds that a large proportion

of model parameters can be frozen during adaptation with minimal reduction in translation quality by encouraging structured sparsity. Wang and Sennrich (2020) links the exposure bias problem (Ranzato et al., 2016; Shao et al., 2018; Zhang et al., 2019) to the phenomenon of NMT tends to generate hallucinations under domain shift. Gu et al. (2020) finds that the NMT tends to generate more high-frequency tokens and less low-frequency tokens than reference. Compared with them, this work mainly focuses on investigating the functions of the different modules and parameters in the NMT model during continual training. In this sense, the work of Thompson et al. (2018) is most related to ours, which tries to understand the effectiveness of continued training for improving the in-domain performance. Compared with their work, our work also explores the cause of the catastrophic forgetting problem in general-domain. Besides, our work also analyzes the performance of NMT at the neuron (parameter) level.

Continual Training Continual training, which is also referred to as fine-tuning, is widely used in NMT for the domain adaptation task. Luong and Manning (2015) fine tunes the general-domain model with the in-domain data. Chu et al. (2017) fine tunes the model with the mix of the general-domain data and over-sampled in-domain data. 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) minimizes the cross-entropy between the output distribution of the general-domain model and the fine-tuned model. Gu et al. (2019) adds a discriminator to help preserve the domain-shared features and fine tunes the whole model on the mixed training data. Jiang et al. (2020) proposes to obtain the word representations by mixing their embedding in individual domains based on the domain proportions. Compared with them, our work pays attention to exploring the inner change of the model during continual training as well as the cause of the catastrophic forgetting phenomenon.

6 Conclusion

In this work, we focus on the catastrophic forgetting phenomenon of NMT and aim to find the inner reasons for this. Under the background of domain adaptation, we propose two analyzing methods from the perspectives of modules and parameters (neurons) and conduct experiments across different language pairs and domains. We find that some modules tend to maintain the general-domain knowledge while some modules tend to adapt to the in-domain; we also find that some parameters are more important for both the general-domain and in-domain translation and the change of them brings about the performance decline in general-domain. Based on our findings, we have proposed several ideas that may help improve the vanilla continual training method. We will prove the effectiveness of these ideas in future work.

Acknowledgements

We thank all the anonymous reviewers for their insightful and valuable comments. This work was supported by National Key R&D Program of China (NO. 2017YFE0192900).

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR*.
- 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.
- Anthony Bau, Yonatan Belinkov, Hassan Sajjad, Nadir Durrani, Fahim Dalvi, and James R. Glass. 2019. Identifying and controlling important neurons in neural machine translation. In *7th International Conference on Learning Representations, ICLR*.
- Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP*, pages 1724–1734.

- 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.
- Yanzhuo Ding, Yang Liu, Huanbo Luan, and Maosong Sun. 2017. Visualizing and understanding neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1150–1159, July.
- Markus Freitag and Yaser Al-Onaizan. 2016. Fast domain adaptation for neural machine translation. *CoRR*, abs/1612.06897.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. Convolutional sequence to sequence learning. In *Proceedings of the 34th International Conference on Machine Learning, ICML*, pages 1243–1252.
- 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, Jinchao Zhang, Fandong Meng, Yang Feng, Wanying Xie, Jie Zhou, and Dong Yu. 2020. Token-level adaptive training for neural machine translation. *CoRR*, abs/2010.04380.
- Haoming Jiang, Chen Liang, Chong Wang, and Tuo Zhao. 2020. Multi-domain neural machine translation with word-level adaptive layer-wise domain mixing. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 1823–1834.
- Nal Kalchbrenner and Phil Blunsom. 2013. Recurrent continuous translation models. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1700–1709.
- 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, NMTACL 2018, Melbourne, Australia, July 20, 2018*, pages 36–44.
- 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.
- Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. 2017. Pruning convolutional neural networks for resource efficient inference. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, C] onference Track Proceedings*.
- 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 NAACL-HLT 2019: Demonstrations*.
- Ye Qi, Devendra Singh Sachan, Matthieu Felix, Sarguna Padmanabhan, and Graham Neubig. 2018. When and why are pre-trained word embeddings useful for neural machine translation? In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT*, pages 529–535.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. Sequence level training with recurrent neural networks. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL*.
- Chenze Shao, Xilin Chen, and Yang Feng. 2018. Greedy search with probabilistic n-gram matching for neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 4778–4784.
- Xing Shi, Kevin Knight, and Deniz Yuret. 2016. Why neural translations are the right length. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP*, pages 2278–2282.

- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Brian Thompson, Huda Khayrallah, Antonios Anastasopoulos, Arya D. McCarthy, Kevin Duh, Rebecca Marvin, Paul McNamee, Jeremy Gwinnup, Tim Anderson, and Philipp Koehn. 2018. Freezing subnetworks to analyze domain adaptation in neural machine translation. In *Proceedings of the Third Conference on Machine Translation: Research Papers, WMT 2018, Belgium, Brussels, October 31 - November 1, 2018*, pages 124–132.
- Brian Thompson, Jeremy Gwinnup, Huda Khayrallah, Kevin Duh, and Philipp Koehn. 2019. Overcoming catastrophic forgetting during domain adaptation of 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*, pages 2062–2068.
- Liang Tian, Derek F. Wong, Lidia S. Chao, Paulo Quaresma, Francisco Oliveira, and Lu Yi. 2014. Um-corpus: A large english-chinese parallel corpus for statistical machine translation. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation, LREC 2014, Reykjavik, Iceland, May 26-31, 2014*, pages 1837–1842.
- 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.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. 2019. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL*, pages 5797–5808.
- Chaojun Wang and Rico Sennrich. 2020. On exposure bias, hallucination and domain shift in neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 3544–3552.
- Joern Wuebker, Patrick Simianer, and John DeNero. 2018. Compact personalized models for neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 881–886.
- Wen Zhang, Yang Feng, Fandong Meng, Di You, and Qun Liu. 2019. Bridging the gap between training and inference for neural machine translation. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4334–4343.