

# Robust Domain Adaptation for Pre-trained Multilingual Neural Machine Translation Models

Mathieu Grosso, Pirashanth Ratnamogan, Alexis Mathey,  
William Vanhuffel, Michael Fotso Fotso

BNP Paribas

(mathieu.grosso, pirashanth.ratnamogan, alexis.mathey, william.vanhuffel,michael.fotsofotso)@bnpparibas.com

## Abstract

Recent literature has demonstrated the potential of multilingual Neural Machine Translation (mNMT) models. However, the most efficient models are not well suited to specialized industries. In these cases, internal data is scarce and expensive to find in all language pairs. Therefore, fine-tuning a mNMT model on a specialized domain is hard. In this context, we decided to focus on a new task: *Domain Adaptation of a pre-trained mNMT model on a single pair of language* while trying to maintain model quality on generic domain data for all language pairs. The risk of loss on generic domain and on other pairs is high. This task is key for mNMT model adoption in the industry and is at the border of many others. We propose a fine-tuning procedure for the generic mNMT that combines embeddings freezing and adversarial loss. Our experiments demonstrated that the procedure improves performances on specialized data with a minimal loss in initial performances on generic domain for all languages pairs, compared to a naive standard approach (+10.0 BLEU score on specialized data, -0.01 to -0.5 BLEU on WMT and Tatoeba datasets on the other pairs with M2M100).

## 1 Introduction

Building a NMT model supporting multiple language pairs is an active and emerging area of research (NLLB Team et al., 2022; Fan et al., 2020; Tang et al., 2020). Multilingual NMT(mNMT) uses a single model that supports translation in multiple language pairs. Multilingual models have several advantages over their bilingual counterparts (Ari-vazhagan et al., 2019b). This modeling proves to be both efficient and effective as it reduces the operational cost (a single model is deployed for all language pairs) and improves translation performances, especially for low-resource languages.

All these advantages make mNMT models interesting for real-world applications. However, they




Figure 1: Domain Adaptation of a Pre-trained mNMT

are not suitable for specialized industries that require domain-specific translation. Training a model from scratch or fine-tuning all the pairs of a pre-trained mNMT model is almost impossible for most companies as it requires access to a large number of resources and specialized data. That said, fine-tuning a single pair of a pre-trained mNMT model in a specialized domain seems possible. Ideally this domain adaptation could be learned while sharing parameters from old ones, without suffering from catastrophic forgetting (McCloskey and Cohen, 1989). This is rarely the case. The risk of degrading performance on old pairs is high due to the limited available data from the target domain and to the extremely high complexity of the pre-trained model. **In our case, overfitting on fine-tuning data means that the model might not even be multilingual anymore**

In this context, this article focuses on a new real-world oriented task **fine-tuning a pre-trained mNMT model in a single pair of language on a specific domain without losing initial performances on the other pairs and generic data**. Our

research focuses on fine-tuning two state-of-the-art pre-trained multilingual mNMT freely available: M2M100 (Fan et al., 2020) and mBART50 (Tang et al., 2020) which both provide high performing BLEU scores and translate up to 100 languages.

**We explored multiple approaches for this domain adaptation.** Our experiments were made on English to French data from medical domain<sup>1</sup>. This paper shows that fine-tuning a pre-trained model with initial layers freezing, for a few steps and with a small learning rate is the best performing approach.

It is organized as follows : firstly, we introduce standard components of modern NMT, secondly we describe related works, thirdly we present our methods. We finally systematically study the impact of some state-of-the-art fine-tuning methods and present our results.

Our main contributions can be separated into 2 parts:

- Defining a new real-world oriented task that focuses on domain adaptation and catastrophic forgetting on multilingual NMT models
- Defining a procedure that allows to finetune a pre-trained generic model on a specific domain

## 2 Background

### 2.1 Neural Machine Translation

Neural Machine Translation (NMT) has become the dominant field of machine translation. It studies how to automatically translate from one language to another using neural networks.

Most NMT systems are trained using Seq2Seq architectures (Sutskever et al., 2014; Cho et al., 2014) by maximizing the prediction of the target sequence  $V_T = (v_1, \dots, v_T)$ , given the source sentence  $W_S = (w_1, \dots, w_S)$ :

$$P(v_1, \dots, v_T \mid w_1, \dots, w_S)$$

Today the best performing Seq2Seq architecture for NMT is based on Transformers (Vaswani et al., 2017) architecture. They are built on different layers among which the multi-head attention and the

feed-forward layer. These are applied sequentially and are both followed by a residual connection (He et al., 2015) and layer normalization (Ba et al., 2016).

Although powerful, traditional NMT only translates from one language to another with a high computational cost compared to its statistical predecessor. It has been shown that a simple language token can condition the network to translate a sentence in any target language from any source language (Johnson et al., 2017). It allows to create multilingual models that can translate between multiple languages. Using previous notation the multilingual model adds the condition on target language in the previous modeling

$$P(v_1, \dots, v_T \mid w_1, \dots, w_S, \ell)$$

where  $\ell$  is the target language.

### 2.2 Transfer Learning

Transfer learning is a key topic in Natural Language Processing (Devlin et al., 2018; Liu et al., 2019). It is based on the assumption that pre-training a model on a large set of data in various tasks will help initialize a network trained on another task where data is scarce.

It is already a key area of research in NMT where large set of generic data are freely available (news, common crawl, ...). However, real-world applications require specialized models. In-domain data is rare and more costly to gather for industries (finance, legal, medical, ...) making specialized models harder to train. It is even more true for multilingual model.

In our work, we study how we can adapt a mNMT model on a specific domain by fine-tuning on only one language pair, without losing too much generality for all language pairs.

## 3 Related works

### 3.1 Multilingual Neural Machine Translation

While initial research on NMT started with bilingual translation systems (Sutskever et al., 2014; Cho et al., 2014; Luong et al., 2015; Yang et al., 2020), it has been shown that the NMT framework is extendable to multilingual models (Dong et al., 2015; Firat et al., 2016; Johnson et al., 2017; Dabre et al., 2020) mNMT has seen a sharp increase in the number of publications, since it is easily extendable and it allows both end-to-end modeling and cross

<sup>1</sup><https://opus.npl.eu/EMEA-v3.php>

lingual language representation (Conneau et al., 2017; Linger and Hajaiej, 2020; Conneau et al., 2019).

Competitive multilingual models have been released and open sourced. mBART (Liu et al., 2019) first, was trained following the BART (Lewis et al., 2019) objective before being finetuned on an English-centric multilingual dataset (Tang et al., 2020). M2M100 (Fan et al., 2020) scaled large transformer layers (Vaswani et al., 2017) with a lot of mined data in order to create a mNMT without using English as pivot, that can perform translation between any pairs among 100 languages. More recently, NLLB was released (NLLB Team et al., 2022), extending the M2M100 framework to 200 languages. Those models are extremely competitive as they have similar performance to their bilingual counterpart while allowing a pooling of training and resources.

Our experiments will rely on M2M100 and mBART but it can be generalized to any new pre-trained multilingual model (NLLB Team et al., 2022).

### 3.2 Domain Adaptation

Domain Adaptation in the field of NMT is a key real-world oriented task. It aims at maximizing model performances on a certain in-domain data distribution. Dominant approaches are based on fine-tuning a generic model using either in-domain data only or a mixture of out-of-domain and in-domain data to reduce overfitting (Servan et al., 2016a; Van Der Wees et al., 2017). Many works have extended domain adaptation to multi-domain, where model is finetuned on multiple and different domains (Sajjad et al., 2017; Zeng et al., 2018; Mghabbar and Ratnamogan, 2020).

However, to the best of our knowledge, our work is the first exploring domain adaptation in the context of recent pre-trained multilingual neural machine translation systems, while focusing on keeping the model performant in out-of-domain data in all languages.

### 3.3 Learning without forgetting

Training on a new task or new data without losing past performances is a generic machine learning task, named Learning without forgetting (Li and Hoiem, 2016).

Limiting pre-trained weights updates using either trust regions or adversarial loss is a recent idea that has been used to improve training stability

in both natural language processing and computer vision (Zhu et al., 2019; Jiang et al., 2020; Aghajanyan et al., 2020). These methods haven't been explored in the context of NMT but are key assets that demonstrated their capabilities on other NLP tasks (Natural Language Inference in particular). Our work will apply a combination of those methods to our task.

### 3.4 Zero Shot Translation

MNMT has shown the capability of direct translation between language pairs unseen in training: a mNMT system can automatically translate between unseen pairs without any direct supervision, as long as both source and target languages were included in the training data (Johnson et al., 2017). However, prior works (Johnson et al., 2017; Firat et al., 2016; Arivazhagan et al., 2019a) showed that the quality of zero-shot NMT significantly lags behind pivot-based translation (Gu et al., 2019). Based on these ideas, some paper (Liu et al., 2021) have focused on training a mNMT model supporting the addition of new languages by relaxing the correspondence between input tokens and encoder representations, therefore improving its zero-shot capacity. We were interested in using this method as learning less specific input tokens during the finetuning procedure could help our model not to overfit the training pairs. Indeed, generalizing to a new domain can be seen as a task that includes generalizing to an unseen language.

## 4 Methods

Our new real-world oriented task being at the cross-board of many existing task, we applied ideas from current literature and tried to combine different approaches to achieve the best results.

### 4.1 Hyperparameters search heuristics for efficient fine-tuning

We seek to adapt generic multilingual model to a specific task or domain. (Cettolo et al., 2014; Servan et al., 2016b). Recent works in NMT (Domingo et al., 2019) have proposed methods to adapt incrementally a model to a specific domain. We continue the training of the generic model on specific data, through several iterations (see Algorithm 1). This post-training fine-tuning procedure is done without dropping the previous learning states of the multilingual model. The resulting model is considered as adapted or specialized to a specific

domain. We want to avoid the model to suffer from forgetting on generic domain and pairs. To this end, we include different methods in this fine-tuning, that have been mentioned in the literature. These methods includes in particular choosing a small learning rate (Howard and Ruder, 2018), a triangular learning schedule (Houlsby et al., 2019), reducing the number of steps and freezing some of the layers(Stickland and Murray, 2019).

## 4.2 Smoothness-inducing Adversarial Regularizer

We seek to reduce the loss on generic domain and other pairs. Indeed, due to limited data resources from downstream tasks and the extremely large capacity of pre-trained models, aggressive fine-tuning often causes the adapted model to overfit the data of downstream tasks and forget the knowledge of the pre-trained model. To this end, we added a Smoothness-inducing Adversarial Regularization (SMART) term during the fine-tuning (Jiang et al., 2020). Models fine-tuned on GLUE task with SMART approach outperform even the strongest pre-trained baseline on all 8 tasks. Comparing with BERT (Devlin et al., 2018) and RoBERTa (Liu et al., 2019), BERT<sub>SMART</sub> and RoBERTa<sub>SMART</sub> are performing better by a big margin. This approach gives a smoothness-inducing property to the model  $f$ . This is helpful to prevent overfitting and to improve generalization on low resource target domain for a certain task. Therefore, adding it to our task should avoid overfitting on the new domain.

Given the model  $f(\cdot; \theta)$  and  $n$  data points of the target task denoted by  $\{(x_i, y_i)\}_{i=1}^n$ , where  $x_i$ 's denote the embeddings of the input sentences, given by the first embedding layer of the language model and  $y_i$ 's are the associated labels, SMART is adding a regularization term  $\mathcal{R}_s(\theta)$  to the canonical optimisation loss below:

$$\min_{\theta}(\mathcal{F}(\theta)) = \mathcal{L}(\theta) + \lambda_s \mathcal{R}_s(\theta) \quad (1)$$

where  $\mathcal{L}(\theta)$  is the loss function defined as

$$\mathcal{L}(\theta) = \frac{1}{n} \sum_{i=1}^n \ell(f(x_i; \theta), y_i) \quad (2)$$

and  $\ell(\cdot, \cdot)$  is the loss function depending on the target task,  $\lambda_s > 0$  is a tuning parameter, and  $\mathcal{R}_s(\theta)$  is the smoothness-inducing adversarial regularizer.

Here we define  $\mathcal{R}_s(\theta)$  as

$$\mathcal{R}_s(\theta) = \frac{1}{n} \sum_{i=1}^n \max_{\|\bar{x}_i - x_i\|_p \leq \epsilon} \ell_s(f(\bar{x}_i; \theta), f(x_i; \theta)) \quad (3)$$

where  $\epsilon > 0$  is a tuning parameter. Since NMT is a classification tasks,  $f(\cdot; \theta)$  outputs a probability simplex and  $\ell_s$  is chosen as the symmetrized KL-divergence, i.e.,

$$\ell_s(P, Q) = \mathcal{D}_{\text{KL}}(P\|Q) + \mathcal{D}_{\text{KL}}(Q\|P)$$

## 4.3 Enabling the model to learn less aggressive input tokens

We seek at reducing the loss of performances on the pairs learned during the pre-training of the model. A factor causing a too important language-specific representation is the positional correspondence to input tokens (Liu et al., 2021). Relaxing it should help the model learn the new domain while not focusing too much on the language representation. Recent advances in mNMT showed that we can reduce the positional correspondence learned from the input tokens seen during training thanks to Positional Disentangling Encoder (PDE) (Liu et al., 2021). PDE corresponds to removing some of the residual connections of the model architecture. PDE is reported to beat by +18.5 BLEU models that do not use it on zero shot translation pairs while retaining quality on supervised directions (Liu et al., 2021). Doing this during the domain adaptation fine-tuning helps to learn less specific input tokens (since we train only from English to French). Therefore, adapting this method to our domain adaptation training is straightforward and could bring gain in BLEU on language pairs seen during pre-training while not sacrificing performances on the new specific domain.

## 5 Experimental Settings

### 5.1 Pre-trained Generic Models used

We have worked with two pre-trained mNMT models: M2M100 and mBART50 large.

**M2M100** is a multilingual encoder-decoder model, based on large Transformer architecture that can handle 100 languages. It was trained on a non-English-centric dataset of 7.5B sentences from generic domain, as such it is the first true many-to-many NMT model. To ease the fine-tuning process and due to hardware limitations, we worked with the lightest version released (418M parameters).

**mBART50** is a multilingual encoder-decoder model, based on training on an English-centric dataset and on large Transformer architecture that can handle 50 languages. It was trained following the BART objective (Lewis et al., 2019). More formally, the model aims to reconstruct a text that has been previously noised.

We will compare the domain adaptation performance between mBART50 which was trained on English-centric data and M2M100 which was trained on non English-centric data.

## 5.2 Datasets and preprocessing

In order to assess the effectiveness of our different domain adaptation strategies, we focused on the medical domain on the **English to French** using data from the EMEA3<sup>2</sup> dataset (Tiedemann, 2012). We used the same preprocessing as the original publications (BPE joint-tokenization from sentencepiece). We split the dataset into a train and a test dataset. We chose to use the first 5.000 sentences for the testing set and 350.000 sentences for the training set. For the evaluation data on the generic domain, we used generic data from different sources including WMT<sup>3</sup> and Tatoeba<sup>4</sup>. For the evaluation data on the medical domain, we also used EMEA3 dataset in different languages.

## 5.3 Detailed Procedure

We first define a hyperparameters search heuristics procedure. We chose a range of learning rate and trained the model with these values. We set prior threshold between the loss we accept on generic data and the increase we target on medical data. Then apply the procedure in algorithm 1. Having done this, we kept best settings (best learning rate and number of steps for given threshold), and tried freezing first layers to reduce the loss on generic domain. We define  $\epsilon_3$ , a threshold between loss on medical domain and gain on generic domain. We reproduce the same procedure and reports our best results. This allows us to find the optimal model  $\theta_{opt}$ , representing the best compromise between not losing performances on generic data and good adaptation to the medical domain.

<sup>2</sup><https://opus.nlpl.eu/EMEA.php>

<sup>3</sup><https://opus.nlpl.eu/WMT-News.php>

<sup>4</sup><https://github.com/Helsinki-NLP/Tatoeba-Challenge>

---

**Algorithm 1** Hyperparameters search heuristic for domain adaptation using simple fine-tuning Algorithm

---

**Input:**  $T$  : the maximum number of steps;  $L$  : the number of layers we have frozen;  $L_r$ : the learning rate,  $\epsilon_1$ : the threshold for  $\Delta_1$  : the difference of BLEU between baseline and adapted model on EN-FR generic domain data,  $\epsilon_2$  threshold for  $\Delta_2$  : the mean difference of BLEU between baseline and adapted model on all other generic data,  $\theta_0$  is the parameters of the pretrained model,  $\theta_{opt}$ : is the parameters of the model that has optimal value of BLEU on domain and generic.

```

1:  $T \leftarrow 100K$ 
2:  $L \leftarrow 1$ 
3: for  $L_r = 3e - 5, 1e - 5, \dots, 1e - 8$  do
4:    $\theta_s \leftarrow \theta_0$ 
5:   for  $s \leftarrow 1$  to  $T$  do
6:      $\theta_{s+1} \leftarrow \text{AdamUpdate}_B(\theta_s)$ 
7:     Every 2k steps, evaluate model on validation set and compute  $\Delta_1$  and  $\Delta_2$ 
8:     if  $\Delta_1 \leq \epsilon_1 \cup \Delta_2 \leq \epsilon_2$  is true then
9:        $\theta_{opt} \leftarrow \theta_s$ 
10:    else
11:       $\theta_{opt} \leftarrow \theta_s$ 
12:    end For loop
13:    end if
14:  end for
15: end for

```

**Output:**  $\theta_{opt}$

---

**M2M100** We trained M2M100 on the medical EN-FR dataset. We used the adam optimizer ( $\beta_1 = 0.9, \beta_2 = 0.98$ ), label smoothing, a dropout of 0.1 and a weight decay of 0. We applied our hyperparameters search heuristic procedure 1 to find the best model. We set  $\epsilon_1 = 2, \epsilon_2 = 1$ . On this configuration, optimal results were reported with a learning rate of 1e-07, freezing the embeddings at the encoder level, and 60K steps.

**mBART50** We trained mBART50 large on the medical EN-FR dataset. We used the adam optimizer ( $\beta_1 = 0.9, \beta_2 = 0.98$ ), label smoothing, a dropout of 0.3 and a weight decay of 0. Again, we applied our hyperparameters search heuristic procedure to find the best model 1. We had to increase the value of  $\epsilon_1, \epsilon_2$  since mBART50 tends to forget the generic domain quicker than M2M100. We set  $\epsilon_1 = 4, \epsilon_2 = 3$ . On this configuration, optimal results were reported with a learning rate

of  $6e - 07$ , freezing the embeddings at the encoder level and 10K steps.

**SMART:** We finetuned the model with the SMART procedure and continue hyperparameters search as in algorithm 1. In Algorithm 2, we note  $\mathcal{R}_s(\theta) = \frac{1}{|B|} \sum_{x_i \in B} \max_{\|\bar{x}_i - x_i\|_p \leq \epsilon} \ell_s(f(\bar{x}_i; \theta), f(x_i; \theta))$  and *AdamUpdate* the ADAM update rule for optimizing equation 1 using the mini-batch  $B$ . Lastly, we set  $T_{\bar{x}} = 1$ . For the perturbation, we set  $\epsilon = 10^{-5}$  and  $\sigma = 10^{-5}$ . The learning rate  $\eta$  is set to  $10^{-3}$ .

---

**Algorithm 2** Adding SMART to procedure

---

**Input:**  $T$  : the total number of iterations;  $\mathcal{X}$ : the dataset;  $\theta_0$ : the parameter of the pre-trained model;  $\sigma^2$ : the variance of the random initialization for  $\bar{x}_i$  's;  $T_{\bar{x}}$ : the number of iterations for updating  $\bar{x}_i$  's;  $\eta$ : the learning rate for updating  $\bar{x}_i$  's;  $\beta$ : clipping value.

- 1:  $\theta_1 \leftarrow \theta_0$
- 2: **for**  $t \leftarrow 1$  to  $T$  **do**
- 3:    $\bar{\theta}_s \leftarrow \theta_t$
- 4:   Sample a mini-batch  $B$  from  $\mathcal{X}$
- 5:   For all  $x_i \in B$ , initialize  $\bar{x}_i \leftarrow x_i + v_i$  with  $v_i \sim \mathcal{N}(0, \sigma^2 I)$
- 6:   **for**  $m = 1, \dots, T_{\bar{x}}$  **do**
- 7:      $\bar{x}_i \leftarrow \bar{x}_i + \eta \mathcal{R}_s(\bar{\theta}_s)$
- 8:   **end for**
- 9:    $\bar{\theta}_{s+1} \leftarrow \text{AdamUpdate}_B(\bar{\theta}_s)$
- 10:    $\theta_{t+1} \leftarrow \text{CLIP}(\bar{\theta}_{s+1}, 1 - \beta, 1 + \beta)$
- 11: **end for**

**Output:**  $\theta_T$

---

**PDE** Finally, we define PDE. It consists in applying Algorithm 1 and then removing first all the residual connection in the penultimate Encoder layers (Chen et al., 2022), then we try removing only the attention layer residual connections (figure 2).




Figure 2: PDE Illustration: Removing Residual Connections on encoder block

## 6 Results and Analysis

### 6.1 Hyperparameters search heuristic

#### 6.1.1 Main Results

**M2M100** As shown in table 1 we reached more than 9.00 increase of BLEU score on the medical dataset without sacrificing performance on generic domain, the loss is not important on most of the pairs (between 0.01 and 0.2). In figure 3, we see that the mean results is rather stable and that the BLEU on generic English to French data does not decrease a lot (around -1.5 BLEU). The model converges after 60K steps so we stop training.

**mBART50** Again we reach more than 9.00 BLEU increase (Figure 4). We observe that after 50K steps mBART50 starts converging around 40.00 BLEU, yet we decided to stop domain adaptation training sooner than with M2M100 as a trade-off between good performance on the EN-FR medical domain and loss of performance on the generic domain. Globally, we achieved better results with M2M100 than mBART50.




Figure 3: Domain Adaptation (Medical Domain EN-FR) of M2M100




Figure 4: Domain Adaptation (Medical Domain EN-FR) of mBART50

### 6.1.2 Catastrophic forgetting with a big learning rate

We tested several learning rate values and we report here our results with a bigger learning rate ( $3e-5$ ). For both models, it led to a catastrophic forgetting on the non-finetuned pairs along with a huge performance increase on the EN-FR Medical dataset, reaching a higher BLEU on the Medical dataset. We decided to focus on a smaller learning rate as a trade-off between loss on generic domain and gain on the medical domain.




Figure 5: Domain Adaptation of M2M100 with big Learning Rate

## 6.2 SMART

We have reported our fine-tuning results for M2M100 and mBART50 with SMART in Table 1.

Our goal with SMART was to reach a higher BLEU score on the generic domain data without sacrificing performances on the medical dataset. In Table 1, we note a good increase in BLEU score. Moreover, we have noted that the BLEU change less when moving learning rate in a reasonable

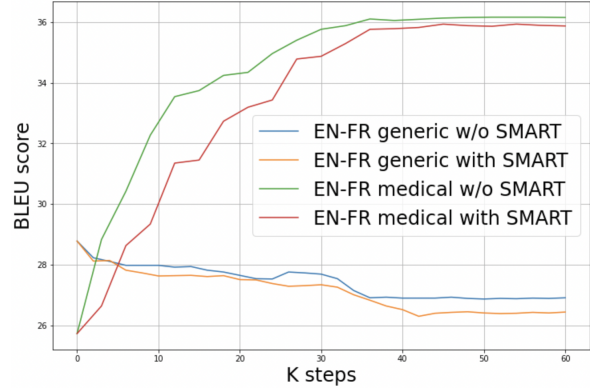


Figure 6: Adding SMART to M2M100 Domain Adaptation training

range compared to the other methods that are extremely sensitive to hyperparameters. In this context, SMART is useful in order to **achieve quick adaptation of a mNMT model to a new domain. It makes domain adaptation procedure more consistent.** Therefore, SMART training procedure allows efficient and robust domain adaptation. However if exploring a large scale of hyper parameters if feasible simple fine-tuning procedure like in Algorithm 1 can provide better results as shown in Table 1.

## 6.3 PDE

We seek at reducing the loss of performances on the pairs learned during pre-training of the model (and that are not used during the post-training domain adaptation). Relaxing the correspondence to the input tokens learned during Domain Adaptation. Fine-tuning was supposed to help learning less specific input tokens and therefore the model would be less likely to forget all the pretrained pairs. As expected, the model learned less aggressive input tokens and do not overfit on English input tokens. However, in practice this does not seem to work well. Indeed, the model is also likelier to forget the pretrained input tokens making this method unfit to our procedure. Using PDE a posteriori (during fine-tuning) seems to be inefficient, since the model is performing worse on all pairs and not only on the English pairs.

We report our results in table 1.

## 6.4 Analysis

### 6.4.1 Zero-shot Domain Adaptation on other pairs

We challenged the approach on domain adaptation on languages unseen during the post-training on the

Table 1: Global results on domain adaptation of M2M100 and mBART50

	M2M100				mBART50			
	Baseline	Finetuned	Finetuned SMART	Finetuned PDE	Baseline	Finetuned	Finetuned SMART	Finetuned PDE
EN-FR medical data	26.94	<b>36.05</b>	35.93	29.71	23.99	<b>33.87</b>	30.3	26.12
EN-FR (WMT)	<b>28.63</b>	26.90	26.41	25.01	<b>31.25</b>	27.10	26.10	17.56
Mean results	24.97	<b>25.38</b>	25.15	21.70	<b>19.83</b>	16.85	15.1	13.63
DE-EN (WMT)	22.42	<b>22.53</b>	22.39	21.15	<b>26.13</b>	21.85	20.64	19.44
EN-DE (WMT)	19.48	<b>19.52</b>	19.21	17.96	<b>22.72</b>	19.84	18.61	16.83
RU-EN (WMT)	<b>26.3</b>	26.27	25.4	24.90	<b>29.72</b>	24.64	23.55	22.47
FR-DE (WMT)	<b>17.82</b>	17.80	17.58	14.69	<b>10.92</b>	8.98	7.1	4.30
EN-FI (WMT)	12.51	<b>12.72</b>	12.51	11.74	<b>13.39</b>	11.23	10.27	9.74
FI-EN (WMT)	<b>23.55</b>	23.18	23.39	21.40	<b>22.10</b>	18.90	17.75	16.33
BG-IT (Tatoeba)	26.65	<b>27.54</b>	27.01	26.20	*	*	*	*
DA-TR (Tatoeba)	20.22	<b>22.27</b>	21.75	20.23	*	*	*	*
PL-RU (Tatoeba)	<b>33.79</b>	33.72	33.68	29.69	<b>14.45</b>	10.49	10.34	9.87
PT-ES (Tatoeba)	51.49	51.98	<b>52.54</b>	50.34	<b>21.57</b>	18.87	16.80	12.54
JA-ES (Tatoeba)	20.55	<b>21.66</b>	21.58	20.30	<b>17.55</b>	15.42	13.27	11.12

medical domain using EMEA3 dataset available on other languages. Table 2 shows that for M2M100 all BLEU scores are increasing, moreover the pairs that implies either English or French are particularly benefiting from this domain adaptation. On mBART50, we also note improvements, first the loss is less important than on generic dataset for the pairs that do not include French as output showing that the model is learning a bit. When French is the output, the domain adaptation is working really fine and we see improvements. Domain-specific data are often hard to gather, especially for low-resource pairs. That’s why being able to improve the performances on a new domain for several pairs using a domain-specific dataset from a single pair is a very interesting propriety from the mNMT models.

	M2M100		mBART50	
	Baseline	Ours	Baseline	Ours
EN-FI medical	12.93	<b>14.83</b>	10.83	10.1
DE-PL medical	12.62	<b>13.6</b>	11.1	7.85
FR-IT medical	23.07	<b>24.62</b>	10.77	8.58
EN-ES medical	32.38	<b>35.15</b>	15.82	17.5
ES-IT medical	25.43	<b>27.06</b>	8.40	7.50
ES-FR medical	24.37	<b>30.64</b>	19.03	25.6
LT-PL medical	12.49	<b>13.8</b>	8.5	7.9
DE-FR medical	18.85	<b>22.20</b>	13.3	18.19
LT-PT medical	17.44	<b>19.26</b>	*	*

Table 2: Zero shot domain adaptation on medical dataset for other pairs

#### 6.4.2 Comparison of initial pre-trained mNMT models (mBART 50 vs M2M100)

We investigated why mBART50 was more likely to forget other pairs compared to M2M100. First, we have worked with the 418M-parameters version of M2M100. This is not the largest M2M100 version

released (and certainly not the most optimized) and this could possibly explain the differences. Then, another hypothesis is the different dataset used during training of both models. Indeed, mBART50 is trained on English-centric data, and M2M100 is not. Non-English centric models are known to achieve higher BLEU especially on low resource data (Fan et al., 2020). Extending this study to domain adaptation, we believe non-English-centric models might be more robust to domain adaptation. We noted that when fine-tuning mBART50 with a bigger learning rate, the first pairs to be forgotten are the non-English ones. Testing this hypothesis on NLLB might be useful.

## 7 Conclusion and Discussion

In this paper, we propose a study of robust domain adaptation approaches on mNMT models where in-domain data is available only for a single language pair. Best performing approach combines embedding freezing and simple fine-tuning with good hyperparameters. This approach shows good improvements with few in-domain data on all language pairs. The framework effectively avoids overfitting and aggressive forgetting on out-of-domain generic data while quickly adapting to in-domain data. We demonstrate that this could be a solution for incremental adaptation of mNMT models. Finally our work is a call for more research in domain adaptation for multilingual models as it is key for real-world applications.

## 8 Limitations

This study was limited by hardware issues. We did not have the possibility to fine-tune on M2M100 large version (12B parameters) that requires 64 GB of VRAM.



Testing our results with a larger version of M2M100 might be interesting.

Also, our study focused on two pre-trained multilingual neural machine translation models. However, many others exist and will be released (NLLB Team et al., 2022). We think that our work is generic enough to be applied on other pre-trained models but extensive experiments on these new models should be carried out.

Finally, the work has been realised on English to French data. We showed domain adaptation is possible for languages with English morphology and tested the impact of this training on many different languages morphology (Japanese, English, Russian, ...). Applying domain adaptation training on other morphology languages and on other domains is also an area to investigate.

## 9 Ethics Statement

The dataset was gathered on OPUS and is largely open-sourced. It was released by (Tiedemann, 2012) and we have downloaded it from OPUS website. We have reviewed the dataset and have not noted any issue with these data. They are very specific to health domain and therefore are not inappropriate. The dataset does not deal with demographic or identity characteristics.

Moreover, these experiments were made using only 2 GPUs and training were relatively short. Given the urgency of addressing climate change, we believe our domain adaptation procedure could help have high-performing mNMT models at small carbon and energy costs. Moreover, SMART framework allows for quicker research of the right hyperparameters, therefore reducing even further the number of experiments and the carbon costs of our method.

## 10 Acknowledgments

The authors would also like to thank Mr. Baoyang Song, Mr. Laurent Lam, Mr. Laatiri Seif Eddine from BNP Paribas for their valuable comments and suggestions. The work is supported by BNP Paribas.

## References

Armen Aghajanyan, Akshat Shrivastava, Ankit Gupta, Naman Goyal, Luke Zettlemoyer, and Sonal Gupta. 2020. Better fine-tuning by reducing representational collapse. *arXiv preprint arXiv:2008.03156*.

Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Roei Aharoni, Melvin Johnson, and Wolfgang Macherey. 2019a. [The missing ingredient in zero-shot neural machine translation](#).

Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George Foster, Colin Cherry, Wolfgang Macherey, Zhifeng Chen, and Yonghui Wu. 2019b. [Massively multilingual neural machine translation in the wild: Findings and challenges](#).

Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. [Layer normalization](#).

Mauro Cettolo, Nicola Bertoldi, Marcello Federico, Holger Schwenk, Loïc Barrault, and Christophe Serivan. 2014. [Translation project adaptation for MT-enhanced computer assisted translation](#). *Machine Translation*, 28:127.

Guanhua Chen, Shuming Ma, Yun Chen, Dongdong Zhang, Jia Pan, Wenping Wang, and Furu Wei. 2022. [Towards making the most of cross-lingual transfer for zero-shot neural machine translation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 142–157, Dublin, Ireland. Association for Computational Linguistics.

Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.

Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2017. Word translation without parallel data. *arXiv preprint arXiv:1710.04087*.

Raj Dabre, Chenhui Chu, and Anoop Kunchukuttan. 2020. A survey of multilingual neural machine translation. *ACM Computing Surveys (CSUR)*, 53(5):1–38.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [Bert: Pre-training of deep bidirectional transformers for language understanding](#).

Miguel Domingo, Mercedes García-Martínez, Álvaro Peris, Alexandre Helle, Amando Estela, Laurent Bié, Francisco Casacuberta, and Manuel Herranz. 2019. [Incremental adaptation of NMT for professional post-editors: A user study](#). *CoRR*, abs/1906.08996.

- Daxiang Dong, Hua Wu, Wei He, Dianhai Yu, and Haifeng Wang. 2015. [Multi-task learning for multiple language translation](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1723–1732, Beijing, China. Association for Computational Linguistics.
- 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, Edouard Grave, Michael Auli, and Armand Joulin. 2020. [Beyond english-centric multilingual machine translation](#). arXiv.
- Orhan Firat, Kyunghyun Cho, and Yoshua Bengio. 2016. [Multi-way, multilingual neural machine translation with a shared attention mechanism](#).
- Jiatao Gu, Yong Wang, Kyunghyun Cho, and Victor O. K. Li. 2019. [Improved zero-shot neural machine translation via ignoring spurious correlations](#).
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. [Deep residual learning for image recognition](#).
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for nlp](#).
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#).
- Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Tuo Zhao. 2020. [SMART: Robust and efficient fine-tuning for pre-trained natural language models through principled regularized optimization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. [Google’s multilingual neural machine translation system: Enabling zero-shot translation](#). *Transactions of the Association for Computational Linguistics*, 5:339–351.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. [Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#).
- Zhizhong Li and Derek Hoiem. 2016. [Learning without forgetting](#).
- Mathis Linger and Mhamed Hajaiej. 2020. [Batch clustering for multilingual news streaming](#). arXiv preprint arXiv:2004.08123.
- Danni Liu, Jan Niehues, James Cross, Francisco Guzmán, and Xian Li. 2021. [Improving zero-shot translation by disentangling positional information](#).
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. [Effective approaches to attention-based neural machine translation](#). arXiv preprint arXiv:1508.04025.
- Michael McCloskey and Neil J. Cohen. 1989. [Catastrophic interference in connectionist networks: The sequential learning problem](#). *The Psychology of Learning and Motivation*, 24:104–169.
- Idriss Mghabbar and Pirashanth Ratnamogan. 2020. [Building a multi-domain neural machine translation model using knowledge distillation](#). arXiv preprint arXiv:2004.07324.
- NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barraud, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. [No language left behind: Scaling human-centered machine translation](#).
- Hassan Sajjad, Nadir Durrani, Fahim Dalvi, Yonatan Belinkov, and Stephan Vogel. 2017. [Neural machine translation training in a multi-domain scenario](#). arXiv preprint arXiv:1708.08712.
- Christophe Servan, Josep Crego, and Jean Senellart. 2016a. [Domain specialization: a post-training domain adaptation for neural machine translation](#). arXiv preprint arXiv:1612.06141.
- Christophe Servan, Josep Maria Crego, and Jean Senellart. 2016b. [Domain specialization: a post-training domain adaptation for neural machine translation](#). *CoRR*, abs/1612.06141.
- Asa Cooper Stickland and Iain Murray. 2019. [Bert and pals: Projected attention layers for efficient adaptation in multi-task learning](#).
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). *Advances in neural information processing systems*, 27.

- 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](#).
- Jörg Tiedemann. 2012. Parallel data, tools and interfaces in opus. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey. European Language Resources Association (ELRA).
- Marlies Van Der Wees, Arianna Bisazza, and Christof Monz. 2017. Dynamic data selection for neural machine translation. *arXiv preprint arXiv:1708.00712*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#).
- Shuoheng Yang, Yuxin Wang, and Xiaowen Chu. 2020. [A survey of deep learning techniques for neural machine translation](#).
- 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*, pages 447–457.
- Chen Zhu, Yu Cheng, Zhe Gan, Siqu Sun, Tom Goldstein, and Jingjing Liu. 2019. Freelib: Enhanced adversarial training for natural language understanding. *arXiv preprint arXiv:1909.11764*.