

Continual Learning in Multilingual NMT via Language-Specific Embeddings

Alexandre Bérard

NAVER LABS Europe

alexandre.berard@naverlabs.com

Abstract

This paper proposes a technique for adding a new source or target language to an existing multilingual NMT model without re-training it on the initial set of languages. It consists in replacing the shared vocabulary with a small language-specific vocabulary and fine-tuning the new embeddings on the new language’s parallel data. Some additional language-specific components may be trained to improve performance (e.g., Transformer layers or adapter modules). Because the parameters of the original model are not modified, its performance on the initial languages does not degrade. We show on two sets of experiments (small-scale on TED Talks, and large-scale on ParaCrawl) that this approach performs as well or better as the more costly alternatives; and that it has excellent zero-shot performance: training on English-centric data is enough to translate between the new language and any of the initial languages.

1 Introduction

Multilingual Neural Machine Translation models are trained on multilingual data to translate from and/or into multiple languages (Firat et al., 2016; Johnson et al., 2017). Multilingual NMT is a compelling approach in production, as one only needs to train, deploy and maintain one model (instead of $2 \times N$ ones, where N is the number of languages). It has also been shown to improve MT quality for low-resource languages (at the cost of a slight degradation for high-resource languages) and it can allow translation between languages that have no aligned data (“zero-shot translation”).

However, such models can be costly to train, as they usually involve larger architectures and large datasets. Moreover, because they are trained jointly on all the languages, they require to know in advance the full set of languages. Adding a new language to an existing model usually means re-training the model on the full multilingual dataset.

Naively fine-tuning the original model on the new language’s data is not an option because of vocabulary mismatch (the shared vocabulary needs to be modified to include the new language’s tokens) and catastrophic forgetting (the model will quickly forget how to translate in the other languages).

In this paper, we study the problem of multilingual NMT *incremental training* or *continual learning* and propose a novel way to efficiently add a new source or target language.

Some desirable properties of an incremental training method are:

- No degradation on the existing language pairs;
- Efficient training (e.g., no re-training on the existing language pairs);
- Minimal amount of added parameters: the approach should scale to many languages and the model fit on a single GPU;
- Minimal degradation in inference speed;
- Good zero-shot performance: when training with X-EN (or EN-X) data, where X is a new language, we would like the model to be able to translate from X to any known language Y (resp. from Y to X).

We propose a novel technique for incrementally adding a new source or target language, which consists in substituting the shared embedding matrix with a language-specific embedding matrix, which is fine-tuned on the new language’s data only while freezing the other parameters of the model. In some cases (e.g., when the new language is on the target size), a small number of additional parameters (e.g., adapter modules) have to be trained to match the performance of the re-training baseline. We perform two sets of experiments, with a 20-language Transformer Base trained on TED Talks, and a 20-language Transformer Big (with deep encoder and shallow decoder) trained on ParaCrawl; and show that this approach is fast and parameter-efficient and that it performs as well or better as the more costly alternatives.

2 Related work

Some previous works study how to adapt a multilingual MT model to unseen low-resource languages, but without seeking to maintain good performance in the initial languages (Neubig and Hu, 2018; Lakew et al., 2019). Garcia et al. (2021) introduce a “vocabulary substitution” approach for adding new languages to a multilingual NMT model. They create a new shared BPE vocabulary that includes the new language and initialize the embeddings of the overlapping tokens with their previous values. Then they fine-tune the entire model on the initial model’s training data combined with parallel or monolingual data in the new language. Contrary to ours, their approach assumes access to the initial model’s training data and results in a small performance drop in the existing languages.

Lyu et al. (2020) and Escolano et al. (2020, 2019, 2021) propose multi-decoder / multi-encoder architectures which they show to be compatible with incremental training. To add a new target (resp. source) language, one just has to freeze the model’s encoder (resp. decoder) and train a new language-specific decoder (resp. encoder). However, this results in an enormous number of parameters.

Artetxe et al. (2020); Pfeiffer et al. (2021) incrementally train language-specific embeddings for cross-lingual transfer of BERT classification models. This approach consists of four stages: 1) train a monolingual BERT on language L_1 ; 2) train embeddings on language L_2 using the masked LM objective while freezing the other parameters; 3) fine-tune the L_1 BERT model on the desired classification task using labeled data in language L_1 ; 4) substitute the L_1 embeddings with the L_2 embeddings in the classification model and use it for L_2 -language classification. Artetxe et al. (2020) also combine their approach with L_2 -specific adapter layers and position embeddings. While this algorithm is close to ours, it is used on encoder-only Transformers for classification tasks. Our work extends this algorithm to encoder-decoder Transformers for multilingual MT.

Also similar to our technique, Thompson et al. (2018) do domain adaptation by freezing most of the NMT parameters and only fine-tuning one component (e.g., the source embeddings). Philip et al. (2020) show that adapter modules can be used to adapt an English-centric multilingual model to unseen language pairs, but whose source and target languages are known. We wanted to go further and

use adapter layers to adapt a multilingual model to unseen languages. However, we obtained the surprising result that adapting the embedding matrix is sometimes enough. In the other cases, adapter modules can be used sparingly to match baseline performance. Üstün et al. (2021) introduce “denoising adapters” which they show can be used to incrementally adapt a multilingual MT model to new languages using monolingual data only.

3 Techniques

Figures 1 and 2 illustrate our technique for a new source and a new target language respectively.

The initial model is a many-to-many model with a shared vocabulary and source-side language codes (to indicate the target language).

3.1 New source language

To add a new source language (e.g., Greek), we build a new (smaller) vocabulary for this language only and replace the source embedding matrix with a new embedding matrix corresponding to that vocabulary. Note that some tokens may appear in both vocabularies. Similarly to Pfeiffer et al. (2021); Garcia et al. (2021), we initialize the new embeddings for those tokens with the existing embedding values. We train this new embedding matrix on Greek-English parallel data while freezing all the other parameters. There is no loss in performance in the existing languages as we do not modify the original parameters. At inference, to translate from the initial set of languages, we use the initial shared vocabulary and embeddings. To translate from Greek, we use the Greek embeddings and vocab.

To better adapt to the new source language, we also try combining this language-specific embedding matrix with other language-specific components in the encoder. We either fine-tune the first encoder layer while freezing the other layers, train the full encoder, or plug in adapter modules after encoder layers (Bapna and Firat, 2019) and train these while freezing the Transformer parameters.

Data augmentation As we will show in the experiments, source lang-specific parameters tend to give poor zero-shot results, i.e., when training them on Greek-English data, the resulting model might have trouble translating into other languages than English. For this reason, we try training such models on additional data. One solution is to use a multi-aligned Greek corpus (i.e., Greek paired with all the initial languages), but this might not always

be possible. We experiment with tiny amounts of such data (e.g., 1000 line pairs per initial language); and with synthetic data: translate the English side of the Greek-English corpus into the other languages with the initial model, then use the resulting fake line pairs for training. We call this approach “back-translation” (BT) even though it is arguably closer to distillation than back-translation because the synthetic text will be on the target side.

3.2 New target language

The same incremental training techniques can be used to learn a new target language (e.g., Greek) with some modifications. The decoder has a target embedding matrix and vocabulary projection matrix, which are usually tied and shared with the source embeddings (i.e., the same parameters are used for all 3 purposes). We need to adapt both the target embeddings and output projection to the new Greek vocabulary. Like in the initial model, we tie these two parameters. Additionally, the initial model does not have a “translate into Greek” language code. We add this language code to the source embedding matrix and freeze all source embeddings but this one. It is initialized with the “to English” language code embedding of the initial model. We combine this approach with language-specific parameters (adapter modules or fine-tuned Transformer layers) in the decoder and/or encoder.

3.3 New source and target languages

To translate between two new languages (e.g., Greek to Ukrainian), we train language-specific parameters for each of these languages separately, as described previously. Then, at inference time, we combine these parameters. This is done by taking the new source Greek embedding matrix and target Ukrainian embedding matrix (and vocabulary projection). The “translate into Ukrainian” language code embedding is concatenated to the Greek embedding matrix. Similarly, the combined model includes language-specific layers and adapters from both models. When both models have adapter modules at the same layers (e.g., last encoder layer), we stack them: the target-language adapters are plugged in after the source-language adapters.

3.4 Baselines

We compare our incremental training techniques with two types of baselines: bilingual models trained from scratch with only the new language’s parallel data; and re-training, i.e., training a new




Figure 1: Adding a new source language with our incremental training technique. The source embedding matrix is replaced with the new language’s embeddings and fine-tuned on the new language’s data, while the other parameters are frozen.




Figure 2: Adding a new target language with our incremental training technique. The tied target embedding matrix and output projection are replaced with the new language’s embeddings. Some language-specific parameters can be added in the decoder or encoder, and a new language code is added in the source embedding matrix. Everything is kept frozen except for these new parameters.

multilingual model that includes the new language. To save computation time, similarly to Garcia et al. (2021), we start from the initial model and substitute its vocabulary with a new vocabulary trained with the same settings and data as before plus text in the new language. This ensures a large overlap between the old and new vocabularies. Then, we initialize the embeddings of the overlapping tokens with their previous values and fine-tune the full model on the entire dataset.

Note that these baselines do not meet our criteria for a good incremental training technique. Bilingual models are parameter-inefficient and cannot do zero-shot translation (except via pivot translation, which is twice as slow). Re-training assumes access to the initial model’s training data and can be very slow. It could also result in a drop in performance in the initial languages.

4 TED Talks Experiments

We adapt a 20-language model trained on TED Talks to Greek (EL), either on the source side or target side. We pick Greek as the new language as it is from an unseen language family and uses an unseen alphabet. We also do experiments with Ukrainian

(UK), Indonesian (ID), or Swedish (SV) as the new language,¹ which are shown in Appendix.

4.1 Data and hyper-parameters

We use the TED Talks corpus (Qi et al., 2018) with the same set of 20 languages as Philip et al. (2020); Bérard et al. (2021).² This corpus is multi-parallel, i.e., it has training data for all 380 (20×19) language pairs. It also includes official valid and test splits for all these language pairs. Table 8 in Appendix shows the training data size per language.

The initial model is the “multi-parallel” baseline from Bérard et al. (2021), a Transformer Base (Vaswani et al., 2017) trained in two stages: English-centric training (38 directions) for 120 epochs; then multi-parallel fine-tuning (380 directions) for 10 epochs.³ More hyper-parameters are given in Appendix (Table 23).

The shared vocabulary is created using BPE (Sennrich et al., 2016) with 64k merge operations and inline casing (Bérard et al., 2019). Both BPE and NMT training use temperature sampling with $T = 5$ (Arivazhagan et al., 2019). Single characters with a total frequency higher than 10 are added to the vocabulary. The Greek vocabulary is obtained with the same BPE settings but on Greek monolingual data with 4k merge operations. The bilingual baselines use a joint BPE model of size 8k and the same settings as in Philip et al. (2020). Our re-training baselines are obtained by creating a new shared BPE model of size 64k including all 20 initial languages plus Greek and fine-tuning the multi-parallel model for 10 more epochs with this vocabulary. Note that there is a vocabulary mismatch with the initial model (which did not have Greek). We initialize the known embeddings with their previous values and the new ones at random and reset the learning rate scheduler and optimizer. We also do a re-training baseline that includes all 4 new languages. Note that contrary to our incremental training approach, those models are trained with the new language(s) on both sides and use multi-aligned parallel data.

Finally, we train a model that follows more closely Garcia et al. (2021): we fine-tune the multi-parallel model for 10 epochs, by replacing the ini-

¹They all use a known script (Latin or Cyrillic). Indonesian is from an unseen language family.

²{en, ar, he, ru, ko, it, ja, zh_cn, es, fr, pt_br, nl, tr, ro, pl, bg, vi, de, fa, hu}

³Note that an “epoch” when using multi-parallel data corresponds to approximately 9 English-centric epochs in terms of updates.

	Model	→EN	←EN	/ EN
1	SOTA – bilingual	32.4	24.4	15.0
2	SOTA – multilingual	30.9	22.3	14.8
3	English-centric	31.8	24.2	13.5
4	(3) + multi-parallel	32.8	24.3	16.3
5	(4) + EL	33.3	24.3	16.6
6	(4) + {EL,UK,SV,ID}	33.2	24.0	16.5

Table 1: BLEU scores (average to English, from English, and between non-English languages) of the baseline models on TED test. “SOTA” corresponds to the bilingual and multi-parallel baselines of Philip et al. (2020). **(3)** and **(4)** are from Bérard et al. (2021).

tial vocabulary with a vocabulary of the exact same size that includes Greek, and whose new tokens are initialized with the outdated embeddings from the old model. Like Garcia et al. (2021), we upscale the new data’s sampling frequency by a factor of 5.

4.2 Evaluation settings

The TED Talks models are evaluated on the provided multi-parallel validation and test sets. Since those are already word-tokenized, we run SacreBLEU with the `--tok none` option.⁴

We report BLEU scores from/into English and average BLEU from/into the 19 other languages than English (which correspond to a zero-shot setting when the incremental training is done on Greek-English only data). We also report chrF scores obtained with SacreBLEU on the test and validation sets in Appendix.⁵

4.3 Results and analysis

Table 29 in Appendix details the notations used in this paper and the tables.

Baselines. Table 1 compares our initial models and re-training baselines against the state of the art on the initial set of 20 languages. In this instance, fine-tuning the initial model with more languages (**5**, **6**) does not degrade BLEU. Appendix Table 9 shows valid and test chrF on more baselines, including our implementation of the vocabulary substitution approach of Garcia et al. (2021).

New source language. Table 2 shows the test BLEU scores of several incrementally-trained models with Greek as a new source language. More re-

⁴SacreBLEU signature: BLEU+c.mixed+#.1+s.exp+tok.none+v.1.5.1

⁵chrF2+numchars.6+space.false+version.1.5.1

ID	Model	Params	EL \rightarrow EN	EL \rightarrow / EN
1	Bilingual baselines	35.7M	38.4	17.1
5	Re-training + EL	–	40.4	19.4
6	Re-training + {EL,UK,SV,ID}	–	40.2	19.4
7	Only embed	2.13M	38.9	18.8
8	(7) + random embed init	2.13M	38.8	18.7*
9	(7) + enc-norm + enc-biases	2.17M	39.5	17.7
10	(7) + enc-adapters-first (dim=64)	2.19M	39.3	0.6
11	(7) + enc-adapters-all (d=64)	2.53M	40.3	0.6
12	(7) + enc-adapters-all (d=512)	5.28M	41.0	0.6
13	(7) + enc-adapters-{1,2,3} (d=1024)	5.28M	41.0	0.6
14	(7) + enc-first-layer	5.28M	40.1	0.7
15	(7) + all-enc-layers	21.0M	40.4	0.6
16	(12) + EL multi-aligned	5.28M	40.5	19.5
17	(12) + EL multi-aligned (BT)	5.28M	40.2	19.0
18	(12) + 1k lines per lang	5.28M	40.9	18.2
19	(12) + 1k lines per lang (BT)	5.28M	41.2	17.8
20	(14) + 1k lines per lang (BT)	5.28M	40.3	18.9
21	(12) + 100 lines per lang (BT)	5.28M	40.9	17.4
22	(7) + {EL,UK,SV,ID}	8.30M	39.3	18.9
23	(14) + {EL,UK,SV,ID}	11.5M	39.9	2.9

Table 2: TED test BLEU scores of incremental training with Greek on the source side. “EL \rightarrow / EN” corresponds to an average BLEU from Greek into all 19 non-English languages. “Params” gives the number of new parameters introduced by each approach. The initial model (4) has 80.2M parameters in total. (*) obtained by using the “translate into X” lang code embeddings from the initial model. The table is divided in 4 parts: baselines trained with multi-aligned data; Greek-English incremental training; incremental training with multi-aligned data (i.e., line pairs between Greek and all 20 languages); and multilingual English-centric incremental training (i.e., on 4 new source languages at once).

sults on Greek, Ukrainian, Indonesian and Swedish are given in Appendix (Tables 10, 11, 12, and 13).

Training the source embeddings only (7) outperforms the bilingual baselines (1) and comes close to the costly re-training baselines (5, 6). In particular, it nearly matches the performance of the latter in the zero-shot EL \rightarrow / EN directions, even though the baselines have training data for those directions. Initializing the known tokens in the new vocabulary with their old embeddings does not improve final performance (7 vs 8). But using language code embeddings from the initial model is necessary to be able to translate into non-English languages. Figure 4 shows that such initialization improves final performance under low-resource settings. Figure 7 in Appendix also shows that it speeds up training.

Training additional components in the encoder, like adapter modules (11, 12, 13) or the first encoder layer (14) helps improve EL \rightarrow EN performance and outperform all baselines, though it is less useful when the new language is from a known family (see Ukrainian and Swedish scores in Ap-

pendix Tables 11 and 13). However, this results in abysmal zero-shot performance (EL \rightarrow / EN). As they only encounter the “to English” language code during training, those models quickly forget how to interpret the other lang codes. This catastrophic forgetting is illustrated by Figure 3, where we see a plunge in EL \rightarrow FR performance after just a few epochs of training. Only tuning the encoder layer norm parameters and biases (9) gives slightly higher EL \rightarrow EN performance without suffering from catastrophic forgetting in the other languages. Note that language code forgetting is less pronounced when the initial model is English-centric (see Table 18 in Appendix). In this setting, adapter modules do not hurt zero-shot translation.

The third quarter of Table 2 shows how multi-aligned Greek data can be used to achieve excellent performance in both EL \rightarrow EN and EL \rightarrow / EN directions. The best tradeoff between EL \rightarrow EN and EL \rightarrow / EN performance is achieved by incrementally training with the entire Greek dataset of 2.41M line pairs. However, such data might not

ID	Model	Params	EN \rightarrow EL	/ EN \rightarrow EL
1	Bilingual baselines	35.7M	32.2	18.3
5	Re-training + EL	–	32.5	21.1
6	Re-training + {EL,UK,SV,ID}	–	32.1	21.1
24	Only embed	2.13M	25.7	16.7
25	(24) + non-tied	4.25M	27.1	17.8
26	(24) + dec-adapters-all (dim=64)	2.53M	29.8	19.3
27	(24) + adapters-all (d=64)	2.93M	32.7	19.6
28	(24) + enc-adapters-last (d=1024)	3.18M	32.0	20.0
29	(26) + enc-adapters-last (d=1024)	3.58M	33.5	20.6
30	(24) + dec-last-layer	6.33M	32.6	20.5
31	(30) + enc-adapters-last (d=1024)	7.38M	34.0	20.8
32	(24) + adapters-all (d=430)	7.43M	34.0	18.2
33	(24) + dec-adapters-all (d=690) + enc-adapters-last (d=1024)	7.43M	33.8	20.8
34	(30) + adapters-all (d=90)	7.36M	34.2	19.8
35	(30) + enc-adapters-all (d=170)	7.38M	34.1	19.0
36	(31) + EL multi-aligned	7.35M	32.9	21.1
37	(31) + {EL,UK,SV,ID}	13.5M	33.0	20.4

Table 3: TED test BLEU scores of incremental training with Greek on the target side. “/ EN \rightarrow EL” corresponds to an average BLEU from the 19 non-English languages to Greek. “Params” gives the number of new parameters introduced by each approach.

always be accessible for the new language. Close performance can be reached by training with the same amounts of synthetic data instead (**17**). And more interestingly, only a tiny amount of real (**18**) or back-translated data (**19**, **20**, **21**) in the other 19 languages is needed to obtain good zero-shot results without any loss in EL \rightarrow EN performance.

New target language. Table 3 shows test BLEU scores when incrementally adding Greek on the target side. Additional results on Greek, Ukrainian, Indonesian and Swedish are provided in Appendix (Tables 14, 15, 16, and 17).

With new target languages, only adapting the embedding matrix (tied with vocabulary projection) is not enough and strongly underperforms the baselines (**24** vs **1**, **5** and **6**). Training decoder-side adapter modules (**26**) gets us closer to baseline performance; and tuning the last decoder layer (**30**) bridges the gap with the baselines. However, the most effective strategy is to train some components in both the encoder and decoder (**27**, **29**, **31**, **32**, **33**, **34**, **35**). We observed that it was important for the model to have a way to modify the output of the encoder before it is read by frozen decoder components. Interestingly, only having a large adapter module after the last encoder layer (**28**) is enough to match baseline performance. Adding small adapters after each decoder layer (**29**) fur-

ther improves BLEU and brings the best parameter count / performance tradeoff.

At the same parameter budget, training adapter modules after every encoder layer (**32**, **34**, **35**) gives worse / EN \rightarrow EL performance than an adapter at the last encoder layer combined with decoder-side parameters (**31**, **33**), which is likely caused by the encoder overfitting to English.

In this setting, there is no clear advantage to incremental training with multi-aligned Greek data (**36**), as this hurts EN \rightarrow EL performance, without any notable improvement for / EN \rightarrow EL. Finally, multilingual incremental training (with 4 new target languages at once) is entirely possible (**37**) and gives competitive results to the baselines.

Table 25 in Appendix analyzes the usefulness of learning a new language code, by comparing with three other strategies: incremental training without any language code; with the “to English” language code; or with the language code of a similar language. Interestingly, the more new parameters are learned (esp. encoder-side), the less useful it is to learn a new language code. Moreover, adapting to Swedish by using a fixed English language code gives reasonable performance as the two languages are from the same family. And the proxy “to Russian” language code gives the same results as learning a new language code when adapting to




Figure 3: TED validation BLEU on EL-EN and EL-FR while incrementally training with EL-EN data only (7, 9, 10, 11, 12, 14).

Source model		Target model		BLEU
1		Bilingual		14.9
1		Bilingual (pivot through English)		18.5
6		Re-training + {EL,UK,SV,ID}		22.0
7	Only embed	30	Dec-last-layer	21.1
		31	Dec-last-layer + enc-adapters-last (d=1024)	21.0
14	Enc-first-layer	30	Dec-last-layer	20.7
		31	Dec-last-layer + enc-adapters-last (d=1024)	21.3
		31	Pivot through English*	21.6
20	Enc-first-layer + 1k (BT)	30	Dec-last-layer	21.2
		31	Dec-last-layer + enc-adapters-last (d=1024)	21.3
19	Enc-adapters-all (d=512) + 1k (BT)	30	Dec-last-layer	20.4
		31	Dec-last-layer + enc-adapters-last (d=1024)	20.7

Table 4: TED test BLEU scores on {EL,UK,SV,ID}→{EL,UK,SV,ID} (average over 12 directions) by combining source-language and target-language incrementally-trained parameters. (*) instead of combining model parameters, translate to English with (14), then to the target language with (31).

Ukrainian because both languages are very similar.

New source and target languages Table 4 combines incrementally-trained parameters at inference time to translate between two new languages. Interestingly, combining target-language parameters with source-language parameters that had very poor zero-shot performance (14) gives excellent results. We hypothesize that the language code forgetting issue is less pronounced here because solely activating some language-specific parameters will make the model translate into that language.

Despite showing the best EL → EN performance, source-language encoder adapters (19) tend to perform more poorly when combined with target-language parameters. While better performance is obtained by pivot translation through English with two incrementally trained models (14 and 31), combining the parameters of these two models gives

close results at a faster inference speed.

Discussion Figure 4 shows final BLEU scores of our techniques when training with smaller amounts of data. We observe that incremental training is more data-efficient than bilingual models and can achieve decent performance even with tiny amounts of training data, making it a good fit for adaptation to low-resource languages.

Figure 7 in Appendix illustrates the training speed of our approach compared to our implementation of Garcia et al. (2021). In addition to maintaining the exact same performance on the previous languages and needing only English-centric data, our approach reaches higher performance in much fewer updates than the alternatives. Note that re-training might be an efficient solution if one wants to add several languages at once and on both sides.

Finally, Tables 18 and 19 in Appendix show that




Figure 4: TED validation BLEU from Greek (left) and to Greek (right) by training corpus size, with incremental training (7, 8, 24, 27) versus bilingual baselines (1).

ID	Model	→EN	←EN	/ EN
38	M2M-124	32.4	31.9	25.7
39	Big 6-6 EN-centric	38.8	36.4	18.5
40	Big 12-2 EN-centric	39.6	37.1	21.1
41	(40) + pivot (EN)	–	–	27.6
42	(40) + multi-para.	39.0	36.2	27.6
43	(40) + {AR,RU,ZH}	39.6	36.6	21.0
44	(42) + {AR,RU,ZH}	39.0	35.8	27.5

Table 5: FLORES devtest spBLEU scores of the ParaCrawl/UNPC baselines. Average to English (19 directions), from English (19 directions) and between non-English languages (342 directions). (39, 40, 41, 42) are the same models as in (Bérard et al., 2021).

our incremental training approach can be applied to English-centric initial models with similar results.

5 ParaCrawl Experiments

In this section, we test our approach in a more realistic, large-scale setting: a 20-language Transformer Big initial model trained on ParaCrawl (Bañón et al., 2020). The incremental training experiments are done in three languages: Arabic (AR), Russian (RU), and Chinese (ZH). Arabic and Chinese are both from unseen language families and use unseen scripts. Russian is close to a known language (Bulgarian) and uses the same script. For training on those languages, we use UNPC (Ziems et al., 2016).

5.1 Data and hyper-parameters

We download ParaCrawl v7.1 in the 19 highest-resource languages paired with English.⁶ Then,

⁶{fr, de, es, it, pt, nl, nb, cs, pl, sv, da, el, fi, hr, hu, bg, ro, sk, lt}

like Freitag and Firat (2020), we build a multi-parallel corpus by aligning all pairs of languages through their English side (effectively removing any English duplicate). See Appendix Table 21 for training data statistics. We create a shared BPE vocab with 64k merge operations and inline casing, by sampling from this data with $T = 5$ and include all characters whose frequency is higher than 100.

Our initial model is the Transformer Big 12-2 (i.e., with a deep encoder and shallow decoder) multi-parallel model of Bérard et al. (2021). Like in the previous section, it was trained in two stages: English-centric training (with $T = 5$) for 1M steps; then multi-parallel fine-tuning (with $T = 2$) for 200k more steps. More hyper-parameters are given in Appendix (Table 24).

Incremental training is done for 120k steps with English-centric data from UNPC v1.0 (see Table 22 in Appendix for statistics), which we clean by removing any line pairs where either side is detected as being in the wrong language by `langid.py` (Lui and Baldwin, 2012). We use monolingual BPE models of size 8k. The Chinese data is tokenized with Jieba⁷ before learning the BPE model.

The English-centric bilingual baselines are Transformer Big 6-6 models trained on UNPC for 120k steps with joint BPE vocabularies of size 16k. We do two “re-training” baselines, by fine-tuning either the English-centric or multi-parallel model on their initial ParaCrawl data plus UNPC data in all three new languages. We sample UNPC line pairs in each of the new directions with probability 0.05. The remaining 0.7 probability mass is distributed among the initial ParaCrawl directions with $T = 5$. The new BPE vocabulary is trained

⁷<https://github.com/fxsjy/jieba>

ID	Model	Params	AR	RU	ZH	AR	RU	ZH
			→ EN			→ / EN		
38	M2M-124 (Goyal et al., 2021)	–	25.5	27.5	20.9	19.6	24.0	18.7
45	Bilingual baselines (pivot through English)	193M	32.1	29.9	23.7	23.2	23.9	19.6
43	English-centric (40) + {AR,RU,ZH}	–	30.0	31.7	22.6	14.6	16.9	11.8
44	Multi-parallel (42) + {AR,RU,ZH}	–	31.0	31.6	23.1	19.2	23.5	15.9
50	(44) + pivot through English	–	–	–	–	22.2	24.5	19.1
46	Only embed	8.6M	24.2	30.8	20.5	16.5	23.8	15.6
47	(46) + enc-adapters-all (dim=512)	21.3M	31.7	31.3	23.5	1.0	1.5	1.5
48	(46) + enc-first-layer	21.2M	30.3	31.2	23.2	1.0	2.6	1.5
49	(48) + 20k lines per lang (BT)	21.2M	30.1	31.4	22.8	20.8	24.6	17.8
ID	Model	Params	EN →			/ EN →		
			AR	RU	ZH	AR	RU	ZH
38	M2M-124 (Goyal et al., 2021)	–	17.9	27.1	19.3	13.8	23.0	16.6
45	Bilingual baselines (pivot through English)	193M	29.1	27.5	22.9	21.2	22.5	18.2
43	English-centric (40) + {AR,RU,ZH}	–	26.5	26.8	18.4	15.7	19.3	11.3
44	Multi-parallel (42) + {AR,RU,ZH}	–	28.3	27.3	20.6	16.6	21.6	13.1
50	(44) + pivot through English	–	–	–	–	20.2	21.8	16.3
51	Only embed	8.6M	11.6	19.7	14.0	9.3	15.8	11.5
52	(51) + enc-adapt-last + dec-adapt-all (1024)	14.9M	27.0	26.2	20.9	18.8	20.8	16.8
53	(51) + dec-last-layer	25.4M	26.5	26.9	21.5	19.1	20.9	17.4
54	(53) + enc-adapters-last (dim=1024)	27.5M	28.2	28.0	22.5	20.0	21.8	18.0
55	(54) without lang ID filtering	27.5M	18.0	19.0	13.6	5.2	10.3	5.0
56	(51) + all-dec-layers	42.2M	28.6	28.1	22.4	20.1	21.9	18.1

Table 6: FLORES devtest spBLEU scores of the ParaCrawl/UNPC incrementally-trained models. The top half of each table corresponds to the baselines (SOTA, bilingual or re-training). “Params” gives the number of new parameters introduced by each approach. The incremental training is always done on one language only (i.e., one row can correspond to 3 different models). Note that the parameter counts given in this table are for Arabic (8.63M embedding parameters). Russian and Chinese embeddings have respectively 8.51M and 13.60M parameters.

with the monolingual ParaCrawl/UNPC data in all 23 languages (with $T = 5$). The new embeddings for the known tokens are initialized with their old values and the other embeddings at random.⁸ Note that contrary to the TED Talks experiments, we do not have multi-aligned data for the new languages.

5.2 Evaluation settings

For validation, we use our own split from TED2020 (Reimers and Gurevych, 2020): 3000 random line pairs for each translation direction. We report chrF scores⁹ computed on these valid sets in Appendix.

As test sets, we use FLORES devtest (Goyal et al., 2021). We report scores computed with their new “spBLEU” metric,¹⁰ which runs BLEU on top of a standardized multilingual BPE tokenization.

⁸78% of the new tokens were in the initial vocabulary, and 84% of the old tokens are in the new vocabulary.

⁹chrF2+n.6+s.false+v.1.5.1

¹⁰BLEU+c.mixed+#.1+s.exp+tok.spm+v.1.5.1 (<https://github.com/ngoyal2707/sacrebleu>)

5.3 Results and analysis

Baselines. Table 5 compares our initial model (42) with other baselines. Our multi-parallel model beats the M2M-124 model of Goyal et al. (2021) in all three settings. This is not so surprising, as their model only has 615M parameters for 124 languages, compared to 255M parameters for our 20-language model. Last, we can observe that our “re-training” baselines (43 and 44) perform almost as well as the initial 20-language models (40, 42).

New source or target language. Training only source embeddings (46) is a good strategy for Russian, but underperforms the baselines in the more linguistically distant Arabic and Chinese. Learning more parameters (+8% per source language) can match baseline performance in all 3 languages (47 and 48), but gives poor zero-shot performance. Adding small amounts of “back-translated” data (49) achieves close non-English performance

Source model		Target model	spBLEU
38		M2M-124 (Goyal et al., 2021)	15.0
45		Bilingual baselines (pivot through English)	18.3
44		Multi-parallel (42) + {AR,RU,ZH}	13.0
50		(44) + pivot through English	16.8
46	Only embed	53 Dec-last-layer	13.5
		54 Dec-last-layer + enc-adapters-last (d=1024)	13.9
47	Enc-adapters-all (dim=512)	53 Dec-last-layer	13.0
		54 Dec-last-layer + enc-adapters-last (d=1024)	13.8
		54 Pivot through English	17.7
48	Enc-first-layer	53 Dec-last-layer	10.2
		54 Dec-last-layer + enc-adapters-last (d=1024)	10.6
49	Enc-first-layer + 20k (BT)	53 Dec-last-layer	12.4
		54 Dec-last-layer + enc-adapters-last (d=1024)	12.7

Table 7: FLORES devtest spBLEU scores of the ParaCrawl/UNPC models on {AR,RU,ZH}→{AR,RU,ZH} (average over 6 directions) by combining source-language and target-language incrementally-trained parameters.

to the pivot translation baselines without hurting English-centric scores. For new target languages, the best strategy is to train the last decoder layer with an adapter module at the last encoder layer (54), which matches the re-training baselines in all 3 languages and gets close performance to the parameter-inefficient bilingual baselines. Interestingly, target-side incremental training is very sensitive to training data noise. In a first iteration of our experiments, we trained with unfiltered UNPC data and observed catastrophic performance (55). Simple language ID filtering solved this issue.

New source and target languages. Table 7 combines source-language with target-language incrementally-trained parameters to translate between two new languages. The results are not as good as in our TED Talks experiments. The best combination in this setting (46 with 54) performs considerably worse than pivot translation through English with the baselines. However, it outperforms the “re-training” baseline (44), which has only seen English-centric data for the new languages. And pivot translation with two incrementally-trained models (47 with 54) gives excellent results, close to the bilingual baselines.

6 Conclusion

We propose a new technique for incrementally training multilingual NMT models on a new source or target language. It consists in creating a new monolingual BPE vocabulary for that language, substituting the shared embedding matrix with language-specific embeddings, and training those

while freezing the other model parameters. At inference, translating in any of the initial languages is done by using the initial shared embeddings, and translating in the new language is done by using the newly trained embeddings. This approach does not change performance on the initial languages as the initial parameters are kept aside and not modified. For new source languages, it can achieve close performance to the more costly and less flexible bilingual and re-training baselines. For new target languages, this technique can be combined with language-specific parameters (fine-tuned Transformer layers or adapter modules) to match baseline performance at a small parameter cost. We validate this technique on two sets of experiments: small-scale on TED Talks and large-scale on ParaCrawl; and show that it is compatible with two architectures: Transformer Base 6-6 and Big 12-2. We also show that incremental training on data aligned with English is enough to learn to translate between the new language and any of the initial languages. Translation between a new source and a new target language is also possible by combining their respective parameters at inference. Finally, we provide supplementary material to facilitate reproducibility.¹¹

¹¹<https://europe.naverlabs.com/research/natural-language-processing/efficient-multilingual-machine-translation>

References

- Roei Aharoni, Melvin Johnson, and Orhan Firat. 2019. [Massively multilingual neural machine translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3874–3884, Minneapolis, Minnesota.
- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George Foster, Colin Cherry, et al. 2019. [Massively multilingual neural machine translation in the wild: Findings and challenges](#). *arXiv preprint arXiv:1907.05019*.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. [On the cross-lingual transferability of monolingual representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4623–4637, Online.
- Marta Bañón, Pinzhen Chen, Barry Haddow, Kenneth Heafield, Hieu Hoang, Miquel Esplà-Gomis, Mikel L. Forcada, Amir Kamran, Faheem Kirefu, Philipp Koehn, Sergio Ortiz Rojas, Leopoldo Pla Sempere, Gema Ramírez-Sánchez, Elsa Sarriás, Marek Strelec, Brian Thompson, William Waites, Dion Wiggins, and Jaume Zaragoza. 2020. [ParaCrawl: Web-scale acquisition of parallel corpora](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4555–4567, Online.
- Ankur Bapna and Orhan Firat. 2019. [Simple, scalable 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)*, pages 1538–1548, Hong Kong, China.
- Alexandre Bérard, Ioan Calapodescu, and Claude Roux. 2019. [Naver Labs Europe’s systems for the WMT19 machine translation robustness task](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 526–532, Florence, Italy.
- Alexandre Bérard, Dain Lee, Stéphane Clinchant, Kweonwoo Jung, and Vassilina Nikoulina. 2021. [Efficient inference for multilingual neural machine translation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Punta Cana, Dominican Republic.
- Carlos Escolano, Marta R. Costa-jussà, and José A. R. Fonollosa. 2019. [From bilingual to multilingual neural machine translation by incremental training](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 236–242, Florence, Italy.
- Carlos Escolano, Marta R. Costa-jussà, José A. R. Fonollosa, and Mikel Artetxe. 2021. [Multilingual machine translation: Closing the gap between shared and language-specific encoder-decoders](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 944–948, Online.
- Carlos Escolano, Marta R. Costa-jussà, José A. R. Fonollosa, and Mikel Artetxe. 2020. [Training multilingual machine translation by alternately freezing language-specific encoders-decoders](#). *arXiv preprint arXiv:2006.01594*.
- Orhan Firat, Kyunghyun Cho, and Yoshua Bengio. 2016. [Multi-way, multilingual neural machine translation with a shared attention mechanism](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 866–875, San Diego, California.
- Markus Freitag and Orhan Firat. 2020. [Complete multilingual neural machine translation](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 550–560, Online.
- Xavier Garcia, Noah Constant, Ankur Parikh, and Orhan Firat. 2021. [Towards continual learning for multilingual machine translation via vocabulary substitution](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1184–1192, Online.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’Aurelio Ranzato, Francisco Guzman, and Angela Fan. 2021. [The FLORES-101 evaluation benchmark for low-resource and multilingual machine translation](#). *arXiv preprint arXiv:2106.03193*.
- 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.
- Surafel Melaku Lakew, Alina Karakanta, Matteo Negri, Marcello Federico, and Marco Turchi. 2019. [Adapting multilingual neural machine translation to unseen languages](#). In *16th International Workshop on Spoken Language Translation*.
- Marco Lui and Timothy Baldwin. 2012. [langid.py: An off-the-shelf language identification tool](#). In *Proceedings of the ACL 2012 System Demonstrations*, pages 25–30, Jeju Island, Korea.
- Sungwon Lyu, Bokyoung Son, Kichang Yang, and Jaekyoung Bae. 2020. [Revisiting Modularized Multilingual NMT to Meet Industrial Demands](#). In *Proceedings of the 2020 Conference on Empirical*

- Methods in Natural Language Processing (EMNLP)*, pages 5905–5918, Online.
- Graham Neubig and Junjie Hu. 2018. [Rapid adaptation of neural machine translation to new languages](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 875–880, Brussels, Belgium.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2021. [UNKs everywhere: Adapting multilingual language models to new scripts](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Punta Cana, Dominican Republic.
- Jerin Philip, Alexandre Bérard, Matthias Gallé, and Laurent Besacier. 2020. [Monolingual adapters for zero-shot neural machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4465–4470, Online.
- Ye Qi, Devendra 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, Volume 2 (Short Papers)*, pages 529–535, New Orleans, Louisiana.
- Nils Reimers and Iryna Gurevych. 2020. [Making monolingual sentence embeddings multilingual using knowledge distillation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4512–4525, Online.
- 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 (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany.
- Brian Thompson, Huda Khayrallah, Antonios Anastopoulos, 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*, pages 124–132, Brussels, Belgium.
- Ahmet Üstün, Alexandre Bérard, Laurent Besacier, and Matthias Gallé. 2021. [Multilingual unsupervised neural machine translation with denoising adapters](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Punta Cana, Dominican Republic.
- 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*.
- Michał Ziemski, Marcin Junczys-Dowmunt, and Bruno Pouliquen. 2016. [The United Nations parallel corpus v1.0](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 3530–3534, Portorož, Slovenia.

A Appendix

Language	Code	Family	Script	X-EN lines	X-* lines
English	en	Germanic	Latin	3.56M	3.56M
Arabic	ar	Semitic	Arabic	214.1k	3.43M
Hebrew	he	Semitic	Hebrew	211.8k	3.40M
Russian	ru	Slavic	Cyrillic	208.5k	3.38M
Korean	ko	Koreanic	Hangul	205.6k	3.35M
Italian	it	Romance	Latin	204.5k	3.35M
Japanese	ja	Japonic	Chinese + Kana	204.1k	3.31M
Mandarin Chinese	zh_cn	Sinitic	Chinese	199.9k	3.30M
Spanish	es	Romance	Latin	196.0k	3.23M
French	fr	Romance	Latin	192.3k	3.19M
Brazilian Portuguese	pt_br	Romance	Latin	184.8k	3.11M
Dutch	nl	Germanic	Latin	183.8k	3.05M
Turkish	tr	Turkic	Latin	182.5k	3.02M
Romanian	ro	Romance	Latin	180.5k	3.06M
Polish	pl	Slavic	Latin	176.2k	3.00M
Bulgarian	bg	Slavic	Cyrillic	174.4k	2.95M
Vietnamese	vi	Vietic	Latin	172.0k	2.81M
German	de	Germanic	Latin	167.9k	2.90M
Persian	fa	Iranian	Arabic	151.0k	2.41M
Hungarian	hu	Uralic	Latin	147.2k	2.47M
Greek	el	Hellenic	Greek	134.3k	2.41M
Ukrainian	uk	Slavic	Cyrillic	108.5k	1.81M
Indonesian	id	Malayic	Latin	87.4k	1.61M
Swedish	sv	Germanic	Latin	56.6k	978.0k
Total	all	–	–	7.11M	62.27M

Table 8: Size of the **Top 20 TED Talks** corpus. English has 253.3k unique lines.

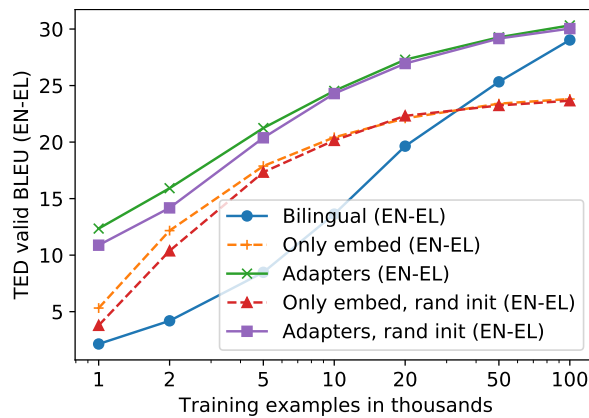


Figure 5: TED validation EN-EL BLEU by training corpus size, with incremental training (24, 27) with or without known embedding initialization, versus bilingual baselines (1).




Figure 6: TED validation BLEU on EL-EN (left) and EL-FR (right) while training with EL-EN data, plus some amount of data (real or back-translated) in the 19 other languages.




Figure 7: TED validation BLEU on EL-EN (left) and EN-EL (right) while training. Comparison of different incremental training approaches with the Garcia et al. (2021) baseline.

ID	Model	Valid chrF			Test chrF		
		→EN	←EN	/ EN	→EN	←EN	/ EN
1	Bilingual (pivot)	.542	.484	.385	.542	.484	.385
3	English-centric	.530	.487	.371	.529	.486	.370
4	(3) + multi-parallel training	.541	.482	.395	.540	.482	.395
57	(3) + EL	.528	.488	.372	.527	.488	.371
58	(3) + {EL, UK, ID, SV}	.526	.485	.370	.526	.485	.370
5	(4) + EL	.545	.481	.398	.545	.482	.398
6	(4) + {EL, UK, ID, SV}	.545	.480	.398	.544	.481	.397
59	(4) + EL (Garcia et al., 2021)	.539	.480	.394	.539	.479	.394
60	(59) @100k steps	.537	.479	.393	.538	.478	.392

Table 9: TED valid and test chrF scores of the baseline models. **(59)** corresponds to the best checkpoint according to validation loss (after 3 epochs, or 320k updates) and **(60)** is after just 100k updates.

ID	Model	Valid chrF		Test chrF	
		EL → EN	EL → / EN	EL → EN	EL → / EN
1	Bilingual baselines	.577	.399	.583	.400
5	Re-training + EL	.591	.426	.596	.425
6	Re-training + {EL,UK,SV,ID}	.590	.425	.594	.424
59	Re-training + EL (Garcia et al., 2021)	.587	.424	.594	.424
60	(59) @100k	.582	.421	.587	.420
7	Only embed	.577	.417	.581	.417
8	(7) + random embed init	.577	.417*	.580	.417*
9	(7) + enc-norm + enc-biases	.582	.407	.587	.407
10	(7) + enc-adapters-first (d=64)	.578	.102	.584	.100
11	(7) + enc-adapters-all (d=64)	.587	.102	.593	.100
12	(7) + enc-adapters-all (d=512)	.593	.102	.602	.100
13	(7) + enc-adapters-{1,2,3} (d=1024)	.595	.103	.603	.101
14	(7) + enc-first-layer	.590	.105	.595	.102
15	(7) + enc-all-layers	.590	.102	.598	.100
16	(12) + EL multi-aligned	.592	.427	.599	.428
18	(12) + 1k lines per lang	.594	.412	.601	.413
19	(12) + 1k lines per lang (BT)	.595	.411	.603	.411
20	(14) + 1k lines per lang (BT)	.589	.422	.596	.422
21	(12) + 100 lines per lang (BT)	.595	.405	.601	.406
22	(7) + {EL,UK,SV,ID}	.582	.419	.585	.419
61	(12) + {EL,UK,SV,ID}	.593	.103	.597	.100
23	(14) + {EL,UK,SV,ID}	.587	.158	.592	.154

Table 10: TED valid and test chrF scores of incremental training with Greek on the source side. (*) obtained by using the “translate into X” lang code embeddings from the initial model.

ID	Model	Valid chrF		Test chrF	
		UK → EN	UK → / EN	UK → EN	UK → / EN
1	Bilingual baselines	.484	–	.494	–
6	Re-training + {EL,UK,SV,ID}	.522	.402	.534	.402
7	Only embed	.516	.397	.525	.395
9	(7) + enc-norm + enc-biases	.518	.386	.526	.385
10	(7) + enc-adapters-first (d=64)	.519	.100	.525	.099
11	(7) + enc-adapters-all (d=64)	.520	.100	.529	.100
12	(7) + enc-adapters-all (d=512)	.520	.100	.529	.099
13	(7) + enc-adapters-{1,2,3} (d=1024)	.523	.100	.530	.100
14	(7) + enc-first-layer	.521	.136	.529	.134
15	(7) + enc-all-layers	.517	.100	.525	.099
16	(12) + UK multi-aligned	.522	.402	.530	.401
18	(12) + 1k lines per lang	.523	.389	.531	.387
19	(12) + 1k lines per lang (BT)	.520	.384	.528	.382
20	(14) + 1k lines per lang (BT)	.522	.397	.528	.396
21	(12) + 100 lines per lang (BT)	.520	.382	.527	.382
22	(7) + {EL,UK,SV,ID}	.518	.398	.526	.396
61	(12) + {EL,UK,SV,ID}	.524	.101	.532	.100
23	(14) + {EL,UK,SV,ID}	.522	.132	.527	.130

Table 11: TED valid and test chrF scores of incremental training with Ukrainian on the source side.

ID	Model	Valid chrF		Test chrF	
		ID → EN	ID → / EN	ID → EN	ID → / EN
1	Bilingual baselines	.516	–	.533	–
6	Re-training + {EL,UK,SV,ID}	.541	.397	.554	.404
7	Only embed	.533	.390	.548	.399
8	(7) + random embed init	.529	.385*	.547	.395*
9	(7) + enc-norm + enc-biases	.537	.379	.551	.389
10	(7) + enc-adapters-first (d=64)	.535	.100	.551	.101
11	(7) + enc-adapters-all (d=64)	.540	.101	.557	.102
12	(7) + enc-adapters-all (d=512)	.541	.101	.558	.102
13	(7) + enc-adapters-{1,2,3} (d=1024)	.545	.101	.558	.102
14	(7) + enc-first-layer	.540	.101	.554	.102
15	(7) + enc-all-layers	.535	.100	.552	.101
16	(12) + ID multi-aligned	.540	.397	.556	.405
18	(12) + 1k lines per lang	.541	.383	.556	.390
19	(12) + 1k lines per lang (BT)	.542	.382	.558	.388
20	(14) + 1k lines per lang (BT)	.536	.389	.553	.399
21	(12) + 100 lines per lang (BT)	.542	.380	.557	.389
22	(7) + {EL,UK,SV,ID}	.530	.388	.547	.397
61	(12) + {EL,UK,SV,ID}	.543	.101	.560	.102
23	(14) + {EL,UK,SV,ID}	.539	.125	.553	.126

Table 12: TED valid and test chrF scores of incremental training with Indonesian on the source side. (★) obtained by using the “translate into X” lang code embeddings from the initial model.

ID	Model	Valid chrF		Test chrF	
		SV → EN	SV → / EN	SV → EN	SV → / EN
1	Bilingual baselines	.577	–	.579	–
6	Re-training + {EL,UK,SV,ID}	.611	.424	.615	.426
7	Only embed	.601	.417	.607	.420
8	(7) + random embed init	.596	.414*	.605	.417*
9	(7) + enc-norm + enc-biases	.604	.413	.613	.415
10	(7) + enc-adapters-first (d=64)	.606	.100	.613	.102
11	(7) + enc-adapters-all (d=64)	.610	.100	.617	.102
12	(7) + enc-adapters-all (d=512)	.608	.100	.613	.102
13	(7) + enc-adapters-{1,2,3} (d=1024)	.612	.100	.619	.102
14	(7) + enc-first-layer	.608	.102	.617	.104
15	(7) + enc-all-layers	.601	.099	.605	.101
16	(12) + SV multi-aligned	.609	.421	.615	.424
18	(12) + 1k lines per lang	.606	.408	.613	.410
19	(12) + 1k lines per lang (BT)	.607	.406	.611	.409
20	(14) + 1k lines per lang (BT)	.605	.417	.613	.420
21	(12) + 100 lines per lang (BT)	.605	.402	.611	.405
22	(7) + {EL,UK,SV,ID}	.601	.417	.605	.418
61	(12) + {EL,UK,SV,ID}	.616	.100	.620	.102
23	(14) + {EL,UK,SV,ID}	.607	.164	.615	.172

Table 13: TED valid and test chrF scores of incremental training with Swedish on the source side. (★) obtained by using the “translate into X” lang code embeddings from the initial model.

ID	Model	Valid chrF		Test chrF	
		EN → EL	/ EN → EL	EN → EL	/ EN → EL
1	Bilingual baselines	.551	.421	.570	.432
5	Re-training + EL	.551	.452	.569	.460
6	Re-training + {EL,UK,SV,ID}	.550	.451	.568	.460
59	Re-training + EL (Garcia et al., 2021)	.553	.449	.570	.458
60	(59) @100k	.553	.450	.572	.459
24	Only embed	.504	.415	.518	.423
25	(24) + non-tied	.517	.423	.530	.432
26	(24) + dec-adapters-all (d=64)	.533	.435	.551	.444
27	(24) + adapters-all (d=64)	.556	.440	.574	.449
28	(24) + enc-adapters-last (d=1024)	.555	.445	.571	.455
29	(26) + enc-adapters-last (d=1024)	.560	.450	.580	.459
30	(24) + dec-last-layer	.556	.449	.576	.458
31	(30) + enc-adapters-last (d=1024)	.566	.452	.585	.461
32	(24) + adapters-all (d=430)	.566	.426	.585	.436
33	(24) + dec-ad-all (690) + enc-ad-last (1024)	.566	.451	.584	.461
34	(30) + adapters-all (d=90)	.565	.442	.585	.453
35	(30) + enc-adapters-all (d=170)	.567	.435	.586	.444
62	(24) + dec-all-layers	.562	.448	.577	.457
36	(31) + EL multi-aligned	.554	.451	.573	.462
37	(31) + {EL,UK,SV,ID}	.558	.449	.578	.458

Table 14: TED valid and test chrF scores of incremental training with Greek on the target side.

ID	Model	Valid chrF		Test chrF	
		EN → UK	/ EN → UK	EN → UK	/ EN → UK
1	Bilingual baselines	.441	–	.440	–
6	Re-training + {EL,UK,SV,ID}	.460	.401	.459	.394
24	Only embed	.446	.386	.445	.380
25	(24) + non-tied	.452	.390	.449	.384
26	(24) + dec-adapters-all (d=64)	.457	.394	.453	.387
27	(24) + adapters-all (d=64)	.466	.397	.465	.391
28	(24) + enc-adapters-last (d=1024)	.464	.400	.463	.393
29	(26) + enc-adapters-last (d=1024)	.469	.402	.465	.394
30	(24) + dec-last-layer	.468	.402	.464	.394
31	(30) + enc-adapters-last (d=1024)	.471	.403	.467	.395
32	(24) + adapters-all (d=430)	.470	.386	.468	.380
33	(24) + dec-ad-all (690) + enc-ad-last (1024)	.468	.402	.466	.394
34	(30) + adapters-all (d=90)	.469	.400	.467	.393
35	(30) + enc-adapters-all (d=170)	.470	.393	.469	.387
62	(24) + dec-all-layers	.468	.400	.463	.391
36	(31) + UK multi-aligned	.460	.400	.458	.393
37	(31) + {EL,UK,SV,ID}	.467	.402	.465	.395

Table 15: TED valid and test chrF scores of incremental training with Ukrainian on the target side.

ID	Model	Valid chrF		Test chrF	
		EN → ID	/ EN → ID	EN → ID	/ EN → ID
1	Bilingual baselines	.568	–	.579	–
6	Re-training + {EL,UK,SV,ID}	.579	.498	.591	.504
24	Only embed	.562	.483	.575	.491
25	(24) + non-tied	.569	.487	.582	.496
26	(24) + dec-adapters-all (d=64)	.579	.492	.591	.501
27	(24) + adapters-all (d=64)	.585	.489	.599	.498
28	(24) + enc-adapters-last (d=1024)	.586	.493	.600	.501
29	(26) + enc-adapters-last (d=1024)	.589	.495	.602	.504
30	(24) + dec-last-layer	.588	.496	.598	.503
31	(30) + enc-adapters-last (d=1024)	.587	.496	.601	.503
32	(24) + adapters-all (d=430)	.589	.480	.599	.489
33	(24) + dec-ad-all (690) + enc-ad-last (1024)	.586	.494	.599	.502
34	(30) + adapters-all (d=90)	.588	.493	.603	.502
35	(30) + enc-adapters-all (d=170)	.589	.489	.602	.496
62	(24) + dec-all-layers	.588	.497	.600	.506
36	(31) + ID multi-aligned	.579	.496	.589	.502
37	(31) + {EL,UK,SV,ID}	.584	.493	.597	.501

Table 16: TED valid and test chrF scores of incremental training with Indonesian on the target side.

ID	Model	Valid chrF		Test chrF	
		EN → SV	/ EN → SV	EN → SV	/ EN → SV
1	Bilingual baselines	.557	–	.557	–
6	Re-training + {EL,UK,SV,ID}	.568	.453	.567	.455
24	Only embed	.547	.436	.548	.438
25	(24) + non-tied	.552	.441	.556	.443
26	(24) + dec-adapters-all (d=64)	.569	.451	.572	.453
27	(24) + adapters-all (d=64)	.584	.444	.589	.448
28	(24) + enc-adapters-last (d=1024)	.583	.451	.586	.455
29	(26) + enc-adapters-last (d=1024)	.587	.452	.590	.457
30	(24) + dec-last-layer	.580	.452	.584	.458
31	(30) + enc-adapters-last (d=1024)	.589	.454	.590	.458
32	(24) + adapters-all (d=430)	.587	.437	.597	.443
33	(24) + dec-ad-all (690) + enc-ad-last (1024)	.584	.452	.587	.457
34	(30) + adapters-all (d=90)	.587	.449	.590	.453
35	(30) + enc-adapters-all (d=170)	.588	.443	.591	.449
62	(24) + dec-all-layers	.583	.452	.586	.457
36	(31) + SV multi-aligned	.570	.454	.569	.456
37	(31) + {EL,UK,SV,ID}	.589	.455	.591	.459

Table 17: TED valid and test chrF scores of incremental training with Swedish on the target side.

ID	Model	Valid chrF		Test chrF	
		EL → EN	EL → / EN	EL → EN	EL → / EN
1	Bilingual baselines (pivot)	.577	.407	.583	.407
57	Re-training + EL (pivot)	.575	.408	.578	.408
58	Re-training + {EL,UK,SV,ID} (pivot)	.572	.405	.573	.404
7	Only embed	.571	.402	.579	.403
9	(7) + enc-norm + enc-biases	.574	.403	.581	.403
10	(7) + enc-adapters-first (d=64)	.575	.404	.582	.404
11	(7) + enc-adapters-all (d=64)	.581	.401	.587	.401
12	(7) + enc-adapters-all (d=512)	.584	.386	.589	.387
13	(7) + enc-adapters-{1,2,3} (d=1024)	.584	.403	.591	.404
14	(7) + enc-first-layer	.581	.114	.588	.111
15	(7) + enc-all-layers	.579	.102	.590	.100
18	(12) + 1k lines per lang	.581	.394	.589	.393
19	(12) + 1k lines per lang (BT)	.583	.396	.588	.396
20	(14) + 1k lines per lang (BT)	.579	.404	.587	.404
21	(12) + 100 lines per lang (BT)	.583	.391	.590	.391

Table 18: TED valid and test chrF scores of incremental training with Greek on the source side when the initial model is English-centric (**3**).

ID	Model	Valid chrF		Test chrF	
		EN → EL	/ EN → EL	EN → EL	/ EN → EL
1	Bilingual baselines (pivot)	.551	.435	.570	.444
57	Re-training + EL (pivot)	.557	.430	.577	.440
58	Re-training + {EL,UK,SV,ID} (pivot)	.556	.428	.573	.438
24	Only embed	.518	.401	.534	.409
25	(24) + non-tied	.529	.407	.545	.416
26	(24) + dec-adapters-all (d=64)	.545	.415	.561	.424
27	(24) + adapters-all (d=64)	.564	.414	.584	.423
28	(24) + enc-adapters-last (d=1024)	.562	.420	.581	.429
29	(26) + enc-adapters-last (d=1024)	.568	.424	.586	.432
30	(24) + dec-last-layer	.565	.427	.585	.435
31	(30) + enc-adapters-last (d=1024)	.571	.426	.589	.435
32	(24) + adapters-all (d=430)	.573	.408	.591	.416
33	(24) + dec-ad-all (690) + enc-ad-last (1024)	.568	.425	.589	.434
34	(30) + adapters-all (d=90)	.573	.420	.592	.428
35	(30) + enc-adapters-all (d=170)	.572	.417	.594	.426
62	(24) + dec-all-layers	.566	.426	.583	.434
36	(31) + EL multi-aligned	.553	.441	.570	.450

Table 19: TED valid and test chrF scores of incremental training with Greek on the target side when the initial model is English-centric (**3**).

Source model		Target model		Valid	Test
1		Bilingual		.384	.388
1		Bilingual (pivot through English)		.428	.437
6		Re-training + {EL,UK,SV,ID}		.461	.470
7	Only embed	30	Dec-last-layer	.456	.465
		31	Dec-last-layer + enc-ad-last (d=1024)	.457	.466
14	Enc-first-layer	30	Dec-last-layer	.453	.463
		31	Dec-last-layer + enc-ad-last (d=1024)	.461	.470
		31	Pivot through English*	.460	.469
20	Enc-first-layer + 1k (BT)	30	Dec-last-layer	.459	.468
		31	Dec-last-layer + enc-ad-last (d=1024)	.461	.470
19	Enc-adapters-all (d=512) + 1k (BT)	30	Dec-last-layer	.451	.459
		31	Dec-last-layer + enc-ad-last (d=1024)	.453	.462

Table 20: TED valid and test chrF scores on {EL,UK,SV,ID}→{EL,UK,SV,ID} (average over 12 directions) by combining source-language and target-language incrementally-trained parameters. (*) instead of combining model parameters, translate with (**14**) to English, then to the target language with (**31**).

Language	Code	Family	X-EN lines	X-* lines
English	en	Germanic	450.30M	450.30M
French	fr	Romance	95.43M	215.63M
German	de	Romance	76.49M	192.67M
Spanish	es	Romance	72.97M	191.71M
Italian	it	Romance	38.05M	136.11M
Portuguese	pt	Romance	29.18M	117.68M
Dutch	nl	Germanic	27.36M	104.35M
Norwegian	nb	Germanic	15.38M	65.37M
Czech	cs	Slavic	12.92M	65.55M
Polish	pl	Slavic	12.88M	69.27M
Swedish	sv	Germanic	10.97M	60.16M
Danish	da	Germanic	9.79M	61.28M
Greek*	el	Hellenic	8.92M	48.29M
Finnish	fi	Uralic	6.83M	47.62M
Croatian	hr	Slavic	6.34M	30.47M
Hungarian	hu	Uralic	6.29M	42.53M
Bulgarian*	bg	Slavic	6.10M	36.84M
Romanian	ro	Romance	5.79M	40.52M
Slovak	sk	Slavic	4.56M	36.39M
Lithuanian	lt	Baltic	4.03M	30.21M
Total	all	–	900.60M	2.043B

Table 21: Size of the **Top 20 ParaCrawl** corpus. English has 271.85M unique lines. (*) all languages use the Latin script, except for Greek and Bulgarian (Cyrillic).

Language	Code	Family	X-EN lines
Russian	ru	Slavic	25.17M
Arabic	ar	Semitic	20.04M
Mandarin Chinese	zh	Sinitic	17.45M

Table 22: Size of the UNPC corpus.

Parameter name	Parameter value
share_all_embeddings	True / False ^{5,6}
share_decoder_input_output_embed	True
arch	transformer
lr_scheduler	inverse_sqrt
optimizer	adam
adam_betas	0.9,0.999
fp16	True
clip_norm	0.0
lr	0.0005 / 0.0001 ⁴
warmup_updates	4000
warmup_init_lr	1e-07
criterion	label_smoothed_cross_entropy
label_smoothing	0.1
dropout	0.3 / 0.1 ^{2,3,4}
max_tokens	4000
max_epoch	120 ^{1,5} / 10 ^{2,4} / 20 ^{3,6}
save-interval	1 / 5 ⁵
validate-interval	1 / 5 ⁵
update_freq [*]	4
reset_*	True
lang_temperature [†]	5

Table 23: fairseq v0.10.2 hyper-parameters of the **TED Talks models**. (★) we normalize this value by the number of GPUs to have a constant batch size. For instance, models trained on 4 GPUs use `update_freq=1`. (1) English-centric training stage of the initial model; (2) multi-parallel training stage; (3) our re-training approach; (4) our implementation of Garcia et al. (2021); (5) English-centric incremental training; (6) multi-aligned incremental training. The bilingual baselines use the `transformer_iwslt_de_en` architecture and are trained for 25k steps with validation every 500 steps and patience 3. (†) we implement an on-the-fly data loading pipeline that builds heterogeneous batches by sampling language pair k with probability: $p_k = D_k^{1/T} / (\sum D_i^{1/T})$ where T is the temperature and D_k is the total number of line pairs for that language pair (Aharoni et al., 2019).

Parameter name	Parameter value
max_source_positions	256
max_target_positions	256
share_all_embeddings	True / False ⁵
share_decoder_input_output_embed	True
arch	transformer_vaswani_wmt_en_de_big
encoder_layers	12 / 6 ⁴
decoder_layers	2 / 6 ⁴
lr_scheduler	inverse_sqrt
optimizer	adam
adam_betas	0.9,0.98
fp16	True
clip_norm	1.0
lr	0.0005
warmup_updates	4000
warmup_init_lr	1e-07
criterion	label_smoothed_cross_entropy
label_smoothing	0.1
dropout	0.1
max_tokens	8000
max_update	1000000 ¹ / 200000 ² / 360000 ³ / 120000 ^{4,5}
save_interval_updates	20000 / 10000 ⁴ / 5000 ⁵
validate_interval_updates	20000 / 10000 ⁴ / 5000 ⁵
update_freq [*]	32
reset_*	True
lang_temperature [†]	5 / 2 ^{2,3}

Table 24: fairseq hyper-parameters of the **ParaCrawl/UNPC models**. (★) (†) see Table 23. (1) English-centric training stage of the initial model; (2) multi-parallel training stage; (3) our re-training approach; (4) bilingual baselines; (5) incremental training.

ID	Model	Lang code	EN →				/ EN →			
			EL	UK	ID	SV	EL	UK	ID	SV
24	Embed only	None	-17	-18	-19	-10	-15	-18	-18	-8
		EN	-18	-19	-17	-5	-14	-16	-17	-7
		Proxy	-6	-2	-2	-13	-3	+0	-2	-6
30	(24) + dec-last-layer	None	-2	-4	-3	+3	-4	-6	-3	-1
		EN	+0	-2	-1	+5	-3	-5	-3	+0
		Proxy	+3	+1	-1	+4	+0	+1	-1	+2
31	(30) + enc-adapters-last (d=1024)	None	-1	-4	+2	-6	-2	-5	-3	-2
32	Enc-adapters-all (dim=430)	None	+2	-1	-1	+1	-8	+1	+1	-3
33	(24) + enc-adapters-last (d=1024) + dec-adapters-all (d=690)	None	-1	-1	1	-1	-2	-4	-1	-4

Table 25: TED valid chrF delta ($\times 1000$) of target-side incremental learning techniques with fixed language codes, compared to models with learned language codes. “None” corresponds to training and decoding without any language code. “EN” trains and decodes with the pre-trained (and frozen) “to English” language code. “Proxy” uses the closest pre-trained language code (RU for UK, BG for EL, DE for SV and VI for ID). This is an oracle, obtained by computing the Euclidean distance between trained language codes in (6).

ID	Model	→EN	←EN	/ EN
39	Big 6-6 English-centric	.582	.571	.400
40	Big 12-2 English-centric	.587	.577	.435
42	(40) + multi-parallel training	.583	.573	.486
41	(40) + pivot through English	–	–	.488
43	(40) + {AR,RU,ZH}	.585	.574	.433
44	(42) + {AR,RU,ZH}	.580	.569	.486

Table 26: TED2020-valid chrF scores of the ParaCrawl/UNPC baselines.

ID	Model	AR	RU	ZH	AR	RU	ZH
		→ EN			→ / EN		
45	Bilingual baselines (pivot through English)	.499	.460	.430	.429	.423	.381
43	English-centric (40) + {AR,RU,ZH}	.488	.480	.430	.372	.385	.337
44	Multi-parallel (42) + {AR,RU,ZH}	.494	.479	.430	.395	.418	.345
50	(44) + pivot through English	–	–	–	.424	.433	.382
46	Only embed	.447	.469	.416	.378	.425	.365
63	(46) without lang ID filtering	.447	.469	.416	.378	.425	.365
47	(46) + enc-adapters-all (d=512)	.502	.478	.434	.154	.157	.152
48	(46) + enc-first-layer	.491	.474	.428	.154	.168	.152
64	(48) without lang ID filtering	.488	.474	.427	.154	.158	.151
49	(48) + 20k lines per lang (BT)	.492	.474	.427	.417	.433	.376
65	(49) without lang ID filtering	.433	.457	.385	.148	.158	.148

ID	Model	EN →			/ EN →		
		AR	RU	ZH	AR	RU	ZH
45	Bilingual baselines (pivot through English)	.423	.437	.187	.358	.400	.156
43	English-centric (40) + {AR,RU,ZH}	.412	.439	.179	.295	.384	.115
44	Multi-parallel (42) + {AR,RU,ZH}	.423	.443	.182	.300	.402	.126
50	(44) + pivot through English	–	–	–	.356	.402	.153
51	Only embed	.314	.398	.158	.277	.364	.134
66	(51) without lang ID filtering	.282	.395	.130	.224	.301	.084
52	(51) + enc-adapters-last + dec-adapters-all (d=1024)	.417	.437	.187	.348	.397	.153
53	(51) + dec-last-layer	.412	.441	.187	.348	.401	.154
54	(53) + enc-adapters-last (d=1024)	.426	.446	.192	.356	.404	.156
55	(54) without lang ID filtering	.312	.363	.107	.072	.185	.037
56	(51) + dec-all-layers	.426	.446	.193	.357	.404	.159

Table 27: TED2020-valid chrF scores of the ParaCrawl/UNPC incrementally-trained models.

Source model	Target model	chrF
45	Bilingual baselines (pivot through English)	.274
44	Multi-parallel (42) + {AR,RU,ZH}	.237
50	(44) + pivot through English	.271
46	Only embed	53 Dec-last-layer .248
	54 Dec-last-layer + enc-adapters-last (d=1024)	.252
47	Enc-adapters-all (d=512)	53 Dec-last-layer .242
	54 Dec-last-layer + enc-adapters-last (d=1024)	.251
	54 Pivot through English*	.274
48	Enc-first-layer	53 Dec-last-layer .223
	54 Dec-last-layer + enc-adapters-last (d=1024)	.234
49	Enc-first-layer + 20k (BT)	53 Dec-last-layer .243
	54 Dec-last-layer + enc-adapters-last (d=1024)	.251

Table 28: TED2020-valid chrF scores of the ParaCrawl/UNPC models on {AR,RU,ZH}→{AR,RU,ZH} (average over 6 directions) by combining source-language and target-language incrementally-trained parameters. (*) instead of combining model parameters, translate to English with (**48**), then to the target language with (**54**).

English-centric	Denotes a parallel corpus that only has alignments with English (i.e., 38 language pairs in our settings). Also denotes many-to-many models that are trained with such data. In this setting, translation between non-English language pairs is called “zero-shot.”
Multi-parallel	Denotes a parallel corpus that has alignments between all possible language pairs (380 in our case), and by extension, models that are trained with such data (AKA “complete multilingual NMT”, Freitag and Firat, 2020).
Enc-adapters- X ($d = N$)	Train adapter modules of bottleneck dimension N after the encoder layer X .
Enc-adapters-all ($d = N$)	Train adapter modules of bottleneck dimension N after all encoder layers.
Enc-layer- X	Fine-tune the X th Transformer encoder layer.
Enc-norm + enc-biases	Fine-tune the layer norm parameters and all the biases in the Transformer encoder.
+ random embed init	Initialize the new language-specific embeddings at random, instead of initializing the embeddings of the known tokens with their previous values.
+ non-tied	Train separate target embeddings and output projection matrix (by default they are tied, i.e., they correspond to the same parameter).
+ EL multi-aligned	Train with multi-aligned EL data: not just paired with English, but with all of the 20 languages (2.41M lines pairs instead of 134k).
+ EL multi-aligned (BT)	Like above, but the non-EN data is obtained by translating the English side of the EL-EN corpus to the other 19 languages with (3). For better comparison with the above method, we back-translate the same number of lines per language as in the multi-aligned EL corpus.
+ N lines per lang	Append to the EL-EN training data N line pairs for each of the 19 non-English languages.
+ N lines per lang (BT)	Like above, but the N line pairs per language are obtained by translating the English side of the EL-EN corpus with (3).
/ EN	19 non-English languages in the initial model. As a column header, it means an average score over all 342 non-English translation directions.
\rightarrow EN	Average score over all 19 $X \rightarrow$ EN translation directions.
\leftarrow EN	Average score over all 19 EN \rightarrow Y translation directions.
Re-training + $\{L_1, L_2 \dots\}$	Fine-tune the initial model with an updated BPE vocabulary and embedding matrix that include the new languages (L_1, L_2 , etc.), and on the multi-aligned data of all the 20 initial languages plus the new ones.
(K) + $\{L_1, L_2 \dots\}$	Use the incremental training technique K , but on several languages at once (L_1, L_2 , etc.) This means that a single shared BPE is trained for all these languages (whose size is multiplied by the number of languages) and the newly-trained parameters are shared between them.

Table 29: Summary of the notations used in this paper.