

# Multilingual Machine Translation: Closing the Gap between Shared and Language-specific Encoder-Decoders

Carlos Escolano, Marta R. Costa-jussà, José A. R. Fonollosa, Mikel Artetxe<sup>†</sup>

TALP Research Center, Universitat Politècnica de Catalunya, Barcelona

<sup>†</sup>IXA NLP Group, University of the Basque Country (UPV/EHU)

{carlos.escolano, marta.ruiz, jose.fonollosa}@upc.edu

mikel.artetxe@ehu.eus

## Abstract

State-of-the-art multilingual machine translation relies on a universal encoder-decoder, which requires retraining the entire system to add new languages. In this paper, we propose an alternative approach that is based on language-specific encoder-decoders, and can thus be more easily extended to new languages by learning their corresponding modules. So as to encourage a common interlingua representation, we simultaneously train the  $N$  initial languages. Our experiments show that the proposed approach outperforms the universal encoder-decoder by 3.28 BLEU points on average, while allowing to add new languages without the need to retrain the rest of the modules. All in all, our work closes the gap between shared and language-specific encoder-decoders, advancing toward modular multilingual machine translation systems that can be flexibly extended in lifelong learning settings.

## 1 Introduction

Multilingual machine translation is the ability to generate translations automatically across a (large) number of languages. Research in this area has attracted a lot of attention in recent times both from the scientific and industrial community. Under the neural machine translation paradigm (Bahdanau et al., 2015), the opportunities for improving this area have dramatically expanded. Thanks to the encoder-decoder architecture, there are viable alternatives to expensive pairwise translation based on classic paradigms<sup>1</sup>.

The main proposal in this direction is the universal encoder-decoder (Joh, 2017) with massive multilingual enhancements (Arivazhagan et al., 2019). While this approach enables zero-shot translation and is beneficial for low-resource languages, it has multiple drawbacks: (i) the entire

system has to be retrained when adding new languages or data; (ii) the quality of translation drops when adding too many languages or for those with the most resources (Arivazhagan et al., 2019); and (iii) the shared vocabulary grows dramatically when adding a large number of languages (especially when they do not share alphabets). Other limitations include the incompatibility of adding multiple modalities such as image or speech.

In this paper, we propose a new framework that can be incrementally extended to new languages without the aforementioned limitations (§3). Our proposal is based on language-specific encoders and decoders that rely on a common intermediate representation space. For that purpose, we simultaneously train the initial  $N$  languages in all translation directions. New languages are naturally added to the system by training a new module coupled with any of the existing ones, while new data can be easily added by retraining only the module for the corresponding language.

We evaluate our proposal on three experimental configurations: translation for the initial languages, translation when adding a new language, and zero-shot translation (§4). Our results show that the proposed method is better in the first configurations by improving the universal system: by 0.40 BLEU points on average in the initial training and by 3.28 BLEU points on average when adding new languages. However, our proposed system is still lagging behind universal encoder-decoder in zero-shot translation.

## 2 Related Work

Multilingual neural machine translation can refer to translating from one-to-many languages (Dong et al., 2015), from many-to-one (Zoph and Knight, 2016) and many-to-many (Joh, 2017). Within the many-to-many paradigm, existing approaches can

<sup>1</sup><http://www.euromatrixplus.net>

be further divided into shared or language-specific encoder-decoders.

**Shared Encoder-Decoder.** [joh \(2017\)](#) feed a single encoder and decoder with multiple input and output languages. Given a set of languages, a shared architecture has a universal encoder and a universal decoder that are trained on all initial language pairs at once. The model shares parameters, vocabulary and tokenization among languages to ensure that no additional ambiguity is introduced in the representation. This architecture provides a simple framework to develop multilingual systems because it does not require modifications of a standard neural machine translation model, and information is easily shared among the different languages through common parameters. Despite the model’s advantages in transfer learning, the use of a shared vocabulary and embedding representation forces the model to employ a vocabulary that includes tokens from all the alphabets used. Additionally, recent work ([Arivazhagan et al., 2019](#)), that imposes representational invariance across language, shows increasing the number of languages varies the quality of the languages already in the system (generally enhancing low-resource pairs but being detrimental for high-resource pairs).

**Language-specific Encoder-Decoders** which may or may not share parameters at some point.

*Sharing parameters.* [Firat et al. \(2016b\)](#) proposed extending the bilingual recurrent neural machine translation architecture ([Bahdanau et al., 2015](#)) to the multilingual case ([Vázquez et al., 2019](#); [Lu et al., 2018](#)) by designing a shared attention-based mechanism between the language-specific encoders and decoders to create a language independent representation. As the language specific components rely on the shared modules, modifying those components to add a new language or add further data to the system would require retraining the whole system (similarly to the previous shared approach). ([Lakew et al., 2018](#)) proposes a model based on the addition of new languages to an already trained system by vocabulary adaptation and transfer learning. While limited, it requires some retraining to adapt the model to the new task.

*No sharing.* The system proposed by [Escolano et al. \(2019\)](#) is trained on language-specific encoders and decoders based on joint training with

out parameter or vocabulary-sharing and on enforcing a compatible representation between the jointly trained languages. The advantage of the approach is that it does not require retraining to add new languages and increasing the number of languages does not vary the quality of the languages already in the system. However, the system has to be trained on a multi-parallel corpus and it does not scale well when there is a large number of languages in the initial system, since all encoders and decoders have to be trained simultaneously.

### 3 Proposed Method

Our proposed approach trains a separate encoder and decoder for each of the  $N$  languages available without requiring multi-parallel corpus. We do not share any parameter across these modules, which allows to add new languages incrementally without retraining the entire system.

#### 3.1 Definitions

We next define the notation that we will be using when describing our approach. We denote the encoder and the decoder for the  $i$ th language in the system as  $e_i$  and  $d_i$ , respectively. For language-specific scenarios, both the encoder and decoder are considered independent modules that can be freely interchanged to work in all translation directions.

#### 3.2 Language-Specific Proposed Procedure

In what follows, we describe the proposed training procedure in two steps: joint training and adding new languages.

**Joint Training** The straightforward approach is to train independent encoders and decoders for each language. The main difference from standard pairwise training is that, in this case, there is only one encoder and one decoder for each language, which will be used for all translation directions involving that language. The training algorithm for this language-specific procedure is described in Algorithm 1.

For each translation direction  $s_{i,j}$  in the training schedule  $S$  with language  $i$  as source and language  $j$  as target, the system is trained using the language-specific encoder  $e_i$  and decoder  $d_j$ .

**Adding New Languages** Since parameters are not shared between the independent encoders and decoders, the joint training enables the addition of

---

**Algorithm 1** Multilingual training step

---

```
1: procedure MULTILINGUALTRAININGSTEP
2:    $N \leftarrow$  Number of languages in the system
3:    $S = \{s_{0,0}, \dots, s_{N,N}\} \leftarrow \{(e_i, d_j)\}$ 
4:    $E = \{e_0, \dots, e_N\} \leftarrow$  Language-specific encs.
5:    $D = \{d_0, \dots, d_N\} \leftarrow$  Language-specific decs.
6:   for  $i \leftarrow 0$  to  $N$  do
7:     for  $j \leftarrow 0$  to  $N$  do
8:       if  $s_{i,j} \in S$  then
9:          $l_i, l_j = \text{get\_parallel\_batch}(i, j)$ 
10:         $\text{train}(s_{i,j}(e_i, d_j), l_i, l_j)$ 
```

---

new languages without the need to retrain the existing modules. Let us say we want to add language  $N + 1$ . To do so, we must have parallel data between  $N + 1$  and any language in the system. For illustration, let us assume that we have  $L_{N+1} - L_i$  parallel data. Then, we can set up a new bilingual system with language  $L_{N+1}$  as source and language  $L_i$  as target. To ensure that the representation produced by this new pair is compatible with the previously jointly trained system, we use the previous  $L_i$  decoder ( $d_{l_i}$ ) as the decoder of the new  $L_{N+1} - L_i$  system and we freeze it. During training, we optimize the cross-entropy between the generated tokens and  $L_i$  reference data but update only the parameters of to the  $L_{N+1}$  encoder ( $e_{l_{N+1}}$ ). By doing so, we train  $e_{l_{N+1}}$  not only to produce good quality translations but also to produce similar representations to the already trained languages. Following the same principles, the  $L_{N+1}$  decoder can also be trained as a bilingual system by freezing the  $L_i$  encoder and training the decoder of the  $L_i - L_{N+1}$  system by optimizing the cross-entropy with the  $L_{N+1}$  reference data.

## 4 Experiments in Multilingual Machine Translation

In this section we report machine translation experiments in different settings. Since the main difference between the shared and the language-specific encoders-decoders lies in whether they retrain the entire system when adding new languages, we accordingly design our experiments to compare the systems under this condition.

### 4.1 Data and Implementation

We used 2 million sentences from the *EuroParl* corpus (Koehn, 2005) in German, French, Spanish and English as training data, with parallel sentences among all combinations of these four

languages (without being multi-parallel). For Russian-English, we used 1 million training sentences from the *Yandex* corpus<sup>2</sup>. As validation and test set, we used *newstest2012* and *newstest2013* from WMT<sup>3</sup>, which is multi-parallel across all the above languages. All data were preprocessed using standard Moses scripts (Koehn et al., 2007)

We evaluate our approach in 3 different settings: (i) the *initial* training, covering all combinations of German, French, Spanish and English; (ii) *adding* new languages, tested with Russian-English in both directions; and (iii) *zero-shot* translation, covering all combinations between Russian and the rest of the languages. Additionally we compare two configurations which consists in using non-tied or tied embeddings. In the language-specific approach tied embeddings consists in using language-wise word embeddings: for one language, we use the same word embeddings. Whereas, in the case of non-tied, the encoder and the decoder of each language have different word embeddings. Tied embeddings in the shared system means that both encoder and decoder share the same word embeddings.

All experiments were done using the Transformer provided by Fairseq(Ott et al., 2019)<sup>4</sup>. We used 6 layers, each with 8 attention heads, an embedding size of 512 dimensions, and a vocabulary size of 32k subword tokens with Byte Pair Encoding (Sennrich et al., 2016) (in total for the shared encoders/decoders and per pair for language-specific encoder-decoders). Dropout was 0.1 for the shared approach and 0.3 for language-specific encoders/decoders. Both approaches were trained with an effective batch size of 32k tokens for approximately 200k updates, using the validation loss for early stopping. We used Adam (Kingma and Ba, 2015) as the optimizer, with learning rate of 0.001 and 4000 warmup steps.

### 4.2 Shared vs Language-specific

Table 1 and 2 show comparisons between the shared and language-specific encoders-decoders.

In contrast with our proposed approach, the shared system requires retraining from scratch to

---

<sup>2</sup><https://translate.yandex.ru/corpus?lang=en>

<sup>3</sup><http://www.statmt.org>

<sup>4</sup>Release v0.6.0 available at <https://github.com/pytorch/fairseq>

	Shared		LangSpec	
	¬Tied	Tied	¬Tied	Tied
de-en	24,40	<b>25,04</b>	22,04	24,54
de-es	24,04	25,01	22,38	<b>25,02</b>
de-fr	24,78	25,14	22,57	<b>25,49</b>
en-de	21,39	21,51	19,44	<b>22,01</b>
en-es	28,08	28,19	26,79	<b>29,53</b>
en-fr	28,43	28,67	26,94	<b>29,74</b>
es-de	19,51	20,21	17,7	<b>20,31</b>
es-en	26,66	26,93	24,9	<b>27,75</b>
es-fr	29,47	29,59	27,31	<b>30,08</b>
fr-de	19,22	19,81	16,88	<b>19,97</b>
fr-en	25,78	26,29	23,5	<b>26,55</b>
fr-es	28,15	29,03	26,78	<b>29,07</b>

Table 1: Initial training. In bold, best global results.

add a new language. For that reason, we experiment with two variants of this system: one trained without Russian-English (*Shared*) and another one including this pair (*Shared<sup>RU</sup>*). Note that, to make experiments comparable, we use the *Shared* version when comparing to our *initial* system in Table 1, and the *Shared<sup>RU</sup>* version when *adding* new languages and performing *zero-shot* translation.

**Initial Training** Table 1 shows that the language-specific encoder-decoders outperforms the shared approach in all cases. On average, our proposed approach is better than the shared approach with a difference of 0.40 BLEU points.

**Adding New Languages** Table 2 shows that, when adding a new language into the system, the language-specific encoder-decoders outperform the shared architecture by 2.92 BLEU points for Russian-to-English and by 3.64 BLEU in the opposite direction. It is also worth mentioning that the Russian data is from a different domain than the frozen English modules used for training (*Yandex* corpus and *EuroParl*, respectively). As such, the language specific encoder-decoders are able to outperform the shared architecture when adding a new language and a new domain by learning from the previous information in the frozen modules. Note that additionally, retraining the shared encoder-decoder to add a new language took an entire week, whereas the incremental training with the language-specific encoder-decoders was performed in only one day.

**Zero-shot** The shared encoder-decoder clearly outperforms the language-specific encoder-

	Shared <sup>RU</sup>		LangSpec	
	¬Tied	Tied	¬Tied	Tied
ru-en	24,71	24,62	25,52	<b>27,54</b>
en-ru	19,91	20,03	21,44	<b>23,94</b>
ru-de	15,36	<b>16,52</b>	12,73	13,77
ru-es	21,38	<b>23,12</b>	18,71	21,08
ru-fr	21,38	<b>22,04</b>	18,05	19,85
de-ru	16,23	<b>17,27</b>	14,39	16,99
es-ru	16,98	<b>18,78</b>	15,93	18,46
fr-ru	16,79	<b>17,83</b>	15,16	17,47

Table 2: Adding a new language translation and Zero-shot.

decoders by 1.39 BLEU points on average. This difference in performance suggests that, while limiting the amount of shared information during training can improve our model performance, it may also harm zero-shot translation.

Note that employing tied embeddings has a larger impact in the language-specific architecture than in the shared one. In fact, it has been key for closing the performance gap between language-specific and shared architectures.

Previous works, e.g. (Firat et al., 2016a; Lu et al., 2018), suffered from attention mismatch which was solved by sharing intermediate layers. In our case, we believe that we do not suffer from this problem because, within an initial set of  $N$  languages, we train  $N * N - 1$  systems using pair-wise corpus (without requiring multi-parallel corpus as previous works (Escolano et al., 2019)). Due to our proposed joint training, once we have trained our initial system, we end up with only  $N$  encoders and  $N$  decoders ( $2 * N$ ).

## 5 Conclusions

In this paper, we present a novel method to train language-specific encoders-decoders without sharing any parameters at all. More relevantly, our system allows to incrementally add new languages into the system without having to retrain it and without varying the translation quality of initial languages in the system. When adding a new language, the language-specific encoder-decoders outperform the shared ones by 3.28 BLEU on average and, most importantly, the training of this new language was done in only one day, as opposed to the week taken by the shared system.

## Acknowledgments

This work is supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 947657).

## References

2017. [Google’s multilingual neural machine translation system: Enabling zero-shot translation](#). *Transactions of the Association for Computational Linguistics*, 5:339–351.
- 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. 2019. [Massively multilingual neural machine translation in the wild: Findings and challenges](#). *CoRR*, abs/1907.05019.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Daxiang Dong, Hua Wu, Wei He, Dianhai Yu, and Haifeng Wang. 2015. [Multi-task learning for multiple language translation](#). In *Proceedings of the ACL-IJCNLP*, pages 1723–1732, Beijing.
- 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 ACL: Student Research Workshop*, Florence, Italy.
- Orhan Firat, Kyunghyun Cho, and Yoshua Bengio. 2016a. [Multi-way, multilingual neural machine translation with a shared attention mechanism](#). In *Proceedings of the 2016 Conference of the NAACL-HLT, month = jun*, pages 866–875, San Diego, California.
- Orhan Firat, Baskaran Sankaran, Yaser Al-Onaizan, Fatos T. Yarman Vural, and Kyunghyun Cho. 2016b. [Zero-resource translation with multi-lingual neural machine translation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 268–277, Austin, Texas.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Philipp Koehn. 2005. [Europarl: A parallel corpus for statistical machine translation](#). In *MT summit*, volume 5, pages 79–86. Citeseer.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, et al. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *Proceedings of the ACL: Demo Papers*, pages 177–180.
- Surafel Melaku Lakew, Aliia Erofeeva, Matteo Negri, Marcello Federico, and Marco Turchi. 2018. [Transfer learning in multilingual neural machine translation with dynamic vocabulary](#). *CoRR*, abs/1811.01137.
- Yichao Lu, Phillip Keung, Faisal Ladhak, Vikas Bhardwaj, Shaonan Zhang, and Jason Sun. 2018. [A neural interlingua for multilingual machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 84–92, Belgium, Brussels.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of NAACL-HLT 2019: Demonstrations*.
- 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.
- Raúl Vázquez, Alessandro Raganato, Jörg Tiedemann, and Mathias Creutz. 2019. [Multilingual NMT with a language-independent attention bridge](#). In *Proceedings of the RepL4NLP Workshop*, pages 33–39, Florence, Italy.
- Barret Zoph and Kevin Knight. 2016. [Multi-source neural translation](#). In *Proceedings of the 2016 Conference of the NAACL-HLT*, pages 30–34, San Diego.