

Trivial Transfer Learning for Low-Resource Neural Machine Translation

Tom Kocmi Ondřej Bojar

Charles University, Faculty of Mathematics and Physics
Institute of Formal and Applied Linguistics
Malostranské náměstí 25, 118 00 Prague, Czech Republic
<surname>@ufal.mff.cuni.cz

Abstract

Transfer learning has been proven as an effective technique for neural machine translation under low-resource conditions. Existing methods require a common target language, language relatedness, or specific training tricks and regimes. We present a simple transfer learning method, where we first train a “parent” model for a high-resource language pair and then continue the training on a low-resource pair only by replacing the training corpus. This “child” model performs significantly better than the baseline trained for low-resource pair only. We are the first to show this for targeting different languages, and we observe the improvements even for unrelated languages with different alphabets.

1 Introduction

Neural machine translation (NMT) has made a big leap in performance and became the unquestionable winning approach in the past few years (Bahdanau et al., 2014; Sutskever et al., 2014; Sennrich et al., 2017; Vaswani et al., 2017). The main reason behind the success of NMT in realistic conditions was the ability to handle large vocabulary (Sennrich et al., 2016b) and to utilize large monolingual data (Sennrich et al., 2016a). However, NMT still struggles if the parallel data is insufficient (e.g. fewer than 1M parallel sentences), producing fluent output unrelated to the source and performing much worse than phrase-based machine translation (Koehn and Knowles, 2017).

Many strategies have been used in MT in the past for employing resources from additional languages, see e.g. Wu and Wang (2007), Nakov and Ng (2012), El Kholy et al. (2013), or Hoang and Bojar (2016). For NMT, a particularly promising approach is transfer learning or “domain adaptation” where the “domains” are the different languages.

For example, Zoph et al. (2016) train a “parent” model in a high-resource language pair, then use some of the trained weights as the initialization for a “child” model and further train it on the low-resource language pair. In Zoph et al. (2016), the parent and child pairs shared the target language (English) and a number of modifications of the training process were needed to achieve an improvement in translation from Hansa, Turkish, and Uzbek into English with the help of French-English data.

Nguyen and Chiang (2017) explore a related scenario where the parent language pair is also low-resource but it is related to the child language pair. They improved the previous approach by using a shared vocabulary of subword units (BPE, Sennrich et al., 2016b). Additionally, they used transliteration to improve their results.

In this paper, we contribute empirical evidence that transfer learning for NMT can be simplified even further. We leave out the restriction on relatedness of the languages and extend the experiments to parent–child pairs where the target language changes. Moreover, we do not utilize any special modifications to the training regime or data pre-processing.

In contrast to previous work, we test the method with the Transformer model (Vaswani et al., 2017), instead of the recurrent approaches (Bahdanau et al., 2014). As documented in e.g. Popel and Bojar (2018) and anticipated in WMT18,¹ the Transformer model seems superior to other NMT approaches.

2 Method Description

The proposed method is extremely simple: We train the parent language pair for a number of iter-

¹<http://www.statmt.org/wmt18/translation-task.html>

ations and switch the training corpus to the child language pair for the rest of the training, without resetting any of the training (hyper)parameters.

As such, this method is similar to the transfer learning proposed by Zoph et al. (2016) but uses the shared vocabulary as in Nguyen and Chiang (2017). The novelty is that we are removing the restriction about relatedness of the language pairs, and in contrast to the previous papers, we show that this simple style of transfer learning can be used on both sides (i.e. either the source or the target language), not only with the target language common to both parent and child model. In fact, the method is effective also for fully unrelated language pairs.

Our method does not need any modification of existing NMT frameworks. The only requirement is to use a shared vocabulary of subword units (we use wordpieces, Johnson et al., 2017) across both language pairs. This is achieved by learning word-piece segmentation from the concatenated source and target sides of both the parent and child language pairs. All other parameters of the model stay the same as for the standard NMT training.

During the training we first train the NMT model for the high-resource language pair until convergence. This model is called “parent”. After that, we train the child model without any restart, i.e. only by changing the training corpora to the low-resource language pair.

2.1 Details on Shared Vocabulary

Current NMT systems use vocabularies of subword units instead of whole words. Using subword units gives a balance between the flexibility of separate characters and efficiency of whole words. It solves the out-of-vocabulary words problem and reduces the vocabulary size. The majority of NMT systems use either the byte pair encoding (Sennrich et al., 2016b) or wordpieces (Wu et al., 2016). Given a training corpus and the desired maximal vocabulary size, either method produces deterministic rules for word segmentation to achieve the fewest possible splits.

Our method requires the vocabulary shared across both the parent (translating from language XX to YY) and the child model (translating from AA to BB). This is obtained by concatenating both training corpora into one corpus of sentences in languages AA, BB, XX and YY.²

²Having separate vocabularies for the parent and child and

Due to our focus on low-resource language pairs, we decided to generate the vocabulary in a balanced way by selecting the same amount of sentences from both language pairs. We thus use the same number of sentence pairs of the parent corpus as there are in the child corpus.

We did not experiment with any other balancing of the vocabulary. Future research could also investigate the impact of using only the child corpus for vocabulary generation or various amounts of used sentences.

We generated vocabularies aiming at 32k subword types. The exact size of the vocabulary varies from 26.1k to 34.8k. All experiments of a given language set use the same vocabulary. Vocabulary overlap in each language set is further studied in Section 6.1.

3 Model Description

We use the Transformer sequence-to-sequence model (Vaswani et al., 2017) as implemented in Tensor2Tensor (Vaswani et al., 2018) version 1.4.2. Our models are based on the “big single GPU” configuration as defined in the paper. To fit the model to our GPUs (NVIDIA GeForce GTX 1080 Ti with 11 GB RAM), we set the batch size to 2300 tokens and limit sentence length to 100 wordpieces.

We use exponential learning rate decay with the starting learning rate of 0.2 and 32000 warm up steps and Adam optimized. In our experiments, we find that it is undesirable to reset learning rate as it leads to the loss of the performance from the parent model. Therefore the transfer learning is handled only by changing the training corpora and nothing else.

Decoding uses the beam size of 8 and the length normalization penalty is set to 1.

The models were trained for 1M steps (approx. 140 hours), which was sufficient for models to converge to the best performance. We selected the model with the best performance on the development test for the final evaluation on the testset.

4 Datasets

In our experiments, we compare low-resource and high-resource language pairs spanning two orders

switching from the XX-YY to AA-BB vocabulary when we switch the training corpus leads on an expected drop in performance. Independent vocabularies use different IDs even for identical subwords and the network cannot rely on any of its weights from the parent training.

Lang. pair	Sent. pairs	Words		Vocabulary	
		First	Second	First	Second
ET,EN	0.8 M	14 M	20 M	631 k	220 k
FI,EN	2.8 M	44 M	64 M	1697 k	545 k
SK,EN	4.3 M	82 M	95 M	1059 k	610 k
RU,EN	12.6 M	297 M	321 M	2202 k	3161 k
CS,EN	40.1 M	491 M	563 M	6253 k	4130 k
AR,RU	10.2 M	243 M	252 M	2299 k	2099 k
FR,RU	10.0 M	295 M	238 M	1339 k	2045 k
ES,FR	10.0 M	297 M	288 M	1426 k	1323 k
ES,RU	10.0 M	300 M	235 M	1433 k	2032 k

Table 1: Datasets sizes overview. We consider Estonian and Slovak low-resource languages in our paper. Word counts and vocabulary sizes are from the original corpus, tokenizing only at whitespace and preserving the case.

of magnitude of training data sizes. We consider Estonian (ET) and Slovak (SK) as low-resource languages compared to the Finnish (FI) and Czech (CS) counterparts.

The choice of languages was closely related to the languages in this year’s WMT 2018 shared tasks. In particular, Estonian and Finnish (paired with English) were suggested as the main focus for their relatedness. We added Czech and Slovak as another closely related language pair. Russian (RU) for the parent model was chosen for two reasons: (1) written in Cyrillic, there will be hardly any intersection in the shared vocabulary with the child language pairs, and (2) previous work uses transliteration to handle Russian, which is a nice contrast to our work. Finally, we added Arabic (AR), French (FR) and Spanish (ES) for experiments with unrelated languages.

The sizes of the training datasets are in Table 1.

If not specified otherwise we use training, development and test sets from WMT.³ Pairs with training sentences with less than 4 words or more than 75 words on either the source or the target side are removed to allow for a speedup of Transformer by capping the maximal length and allowing a bigger batch size. The reduction of training data is small and based on our experiments, it does not change the performance of the translation model.

We use the Europarl and Rapid corpora for Estonian-English. We disregard Paracrawl due to its noisiness. The development and test sets are from WMT news 2018.

The Finnish-English was prepared as in Östling et al. (2017), removing Wikipedia headlines. The

³<http://www.statmt.org/wmt18/>

dev and test sets are from WMT news 2015.

For English-Czech, we use all parallel data allowed in WMT2018 except Paracrawl. The main resource is CzEng 1.7 (the filtered version, Bojar et al., 2016). The devset is WMT newstest2011 and the testset is WMT newstest2017.

Slovak-English uses corpora from Galuščáková and Bojar (2012), detokenized by Moses.⁴ WMT newstest2011 serves as the devset and testset.

The Russian-English training set was created from News Commentary, Yandex and UN Corpus. As the devset, we use WMT newstest 2012.

The language pairs Arabic-Russian, French-Russian, Spanish-French and Spanish-Russian were selected from UN corpus (Ziemski et al., 2016), which provides over 10 million multi-parallel sentences in 6 languages.

5 Results

In this section, we present results of our approach. Statistical significance of the winner (marked with ‡) is tested by paired bootstrap resampling against the baseline (child-only) setup (1000 samples, conf. level 0.05; Koehn, 2004).

As customary, we label the models with the pair of the source and target language codes, for example the English-to-Estonian translation model is denoted by ENET.

The vocabularies are generated as described in 2.1 separately for each experimented combination of parent and child. The same vocabulary is used whenever the parent and child use the same set of languages, i.e. disregarding the translation direction and model stage (parent or child).

5.1 English as the Common Language

Table 2 summarizes our results for various combinations of high-resource parent and low-resource child language pairs when English is shared between the child and parent either in the encoder or in the decoder.

We confirm that sharing the target language improves performance as previously shown (Zoph et al., 2016; Nguyen and Chiang, 2017). This gains up to 2.44 BLEU absolute for ETEN with the FIEN parent. Using only the parent (FIEN) model to translate the child (ETEN) test set gives a miserable performance, confirming the need for transfer learning or “finetuning”.

⁴<https://github.com/amos-sm/amosdecoder>

Parent - Child	Transfer	Baselines: Only	
		Child	Parent
enFI - enET	19.74‡	17.03	2.32
Flen - ETen	24.18‡	21.74	2.44
enCS - enET	20.41‡	17.03	1.42
enRU - enET	20.09‡	17.03	0.57
RUen - ETen	23.54‡	21.74	0.80
enCS - enSK	17.75‡	16.13	6.51
CSen - SKen	22.42‡	19.19	11.62
enET - enFI	20.07‡	19.50	1.81
ETen - Flen	23.95	24.40	1.78
enSK - enCS	22.99	23.48‡	6.10
SKen - CSen	28.20	29.61‡	4.16

Table 2: Transfer learning with English reused either in source (encoder) or target (decoder). The column “Transfer” is our method, baselines correspond to training on one of the corpora only. Scores (BLEU) are always for the child language pair and they are comparable only within lines or when the child language pair is the same. “Unrelated” language pairs in bold. Upper part: parent larger, lower part: child larger. (“EN” lowercased just to stand out.)

A novel result is that the method works also for sharing the source language, improving ENET by up to 2.71 BLEU thanks to ENFI parent.

Furthermore, the improvement is not restricted only to related languages as Estonian and Finnish as shown in previous works. Unrelated language pairs (shown in bold in Table 2) like Czech and Estonian work too and in some cases even better than with the related datasets. We reach an improvement of 3.38 BLEU for ENET when parent model was ENCS, compared to improvement of 2.71 from ENFI parent. This statistically significant improvement contradicts Dabre et al. (2017) who concluded that the more related the languages are, the better transfer learning works. We see it as an indication that the size of the parent training set is more important than relatedness of languages.

The results with Russian parent for Estonian child (both directions) show that transliteration is also not necessary. Because there is no vocabulary sharing between Russian Cyrillic and Estonian Latin (except numbers and punctuation, see Section 6.1 for further details), the improvement could be attributed to a better coverage of English; an effect similar to domain adaptation.

On the other hand, this transfer learning works well only when the parent has more training data than the child. As presented in the bottom part of Table 2, low-resource parents do not generally improve the performance of better-resourced children and sometimes, they even (significantly) decrease

Child Training Sents	Transfer BLEU	Baseline BLEU
800k	19.74	17.03
400k	19.04	14.94
200k	17.95	11.96
100k	17.61	9.39
50k	15.95	5.74
10k	12.46	1.95

Table 3: Maximal score reached by ENET child for decreasing sizes of child training data, trained off an ENFI parent (all ENFI data are used and models are trained for 800k steps). The baselines use only the reduced ENET data.

it. This is another indication, that the most important is the size of the parent corpus compared to the child one.

The baselines are either models trained purely on the child parallel data or only on the parent data. The second baseline only indicates the relatedness of languages because it is only tested but never trained on the child language pair. Also, we do not add any language tag as in Johnson et al. (2017). This also highlights that the improvement of our method cannot be directly attributed to the relatedness of languages: e.g. Czech and Slovak are much more similar than Czech and Estonian (Parent Only BLEU of translation out of English is 6.51 compared to 1.42) and yet the gain from transfer learning is larger for Estonian (+3.38) than from Slovak (+1.62).

5.2 Simulated Very Low Resources

In Table 3, we simulate very low-resource settings by downscaling the data for the child model. It is a common knowledge, that gains from transfer learning are more pronounced for smaller children. The point of Table 3 is to illustrate that our approach is applicable even to extremely small child setups, with as few as 10k sentence pairs. Our transfer learning (“start with a model for whatever parent pair”) may thus resolve the issue of applicability of NMT for low resource languages as pointed out by Koehn and Knowles (2017).

5.3 Parent Convergence

Figure 1 compares the performance of the child model when trained from various training stages of the parent model. The performance of the child clearly correlates with the performance of the parent. Therefore, it is better to use a parent model that already converged and reached its best performance.




Figure 1: Learning curves on dev set for ENFI parent and ENET child where the child model started training after various numbers of the parent’s training steps.

Parent - Child	Transfer	Baseline	Aligned
enFI - ETen	22.75 \ddagger	21.74	24.18
Flen - enET	18.19 \ddagger	17.03	19.74
enRU - ETen	23.12 \ddagger	21.74	23.54
enCS - ETen	22.80 \ddagger	21.74	not run
RUen - enET	18.16 \ddagger	17.03	20.09
enET - ETen	22.04 \ddagger	21.74	21.74
ETen - enET	17.46	17.03	17.03

Table 4: Results of child following a parent with swapped direction. “Baseline” is child-only training. “Aligned” is the more natural setup with English appearing on the “correct” side of the parent, the numbers in this column thus correspond to those in Table 2.

5.4 Direction Swap in Parent and Child

Relaxing the setup in Section 5.1, we now allow a mismatch in translation direction of the parent and child. The parent XX-EN is thus followed by an EN-YY child or vice versa. It is important to note that Transformer shares word embeddings for the source and target side. The gain can be thus due to better English word embeddings, but definitely not due to a better English language model. It would be interesting to study the effect of not sharing the embeddings but we leave it for some future work.

The results in Table 4 document that an improvement can be reached even when none of the involved languages is reused on the same side. This interesting result should be studied in more

Parent - Child	Transfer	Baseline
ARRU - ETEN	22.23	21.74
ESFR - ETEN	22.24 \ddagger	21.74
ESRU - ETEN	22.52 \ddagger	21.74
FRRU - ETEN	22.40 \ddagger	21.74

Table 5: Transfer learning with parent and child not sharing any language.

detail. [Firat et al. \(2016\)](#) hinted possible gains even when both languages are distinct from the low-resource languages but in a multilingual setting. Not surprisingly, the improvements are better when the common language is aligned.

The bottom part of Table 4 shows a particularly interesting trick: the parent is not any high-resource pair but the very same EN-ET corpus with source and target swapped. We see gains in both directions, although not always statistically significant. Future work should investigate if this performance boost is possible even for high-resource languages. Similar behavior has been shown in [Niu et al. \(2018\)](#), where in contrast to our work they mixed the data together and added an artificial token indicating the target language.

5.5 No Language in Common

Our final set of experiments examines the performance of ETEN child trained off parents in totally unrelated language pairs. Without any common language, the gains cannot be attributed, e.g., to the shared English word embeddings. The vocabulary overlap is mostly due to short n-grams or numbers and punctuations.

We see gains from transfer learning in all cases, mostly significant. The only non-significant gain is from Arabic-Russian which does not share the script with the child Latin at all. (Sharing of punctuation and numbers is possible across all the tested scripts.) The gains are quite similar (+0.49–+0.78 BLEU), supporting our assumption that the main factor is the size of the parent (here, all have 10M sentence pairs) rather than language relatedness.

6 Analysis

Here we provide a rather initial analysis of the sources of the gains.

6.1 Vocabulary Overlap

Our method relies on the vocabulary estimated jointly from the child and parent model. In Trans-

ET	EN	RU	% Subwords
✓	-	-	29.93%
-	✓	-	20.69%
-	-	✓	29.03%
✓	✓	-	10.06%
-	✓	✓	1.39%
✓	-	✓	0.00%
✓	✓	✓	8.89%
Total			28.2k (100%)
From parent			41.03%

Table 6: Breakdown of subword vocabulary of experiments involving ET, EN and RU.

former, the vocabulary is even shared across encoder and decoder. With a large overlap, we could expect a lot of “information reuse” between the parent and the child.

Since the subword vocabulary depends on the training corpora, a little clarification is needed. We take the vocabulary of subword units as created e.g. for ENRU-ENET experiments, see Section 2.1. This vocabulary contains 28.2k subwords in total. We then process the training corpora for each of the languages with this shared vocabulary, ignore all subwords that appear less than 10 times in each of the languages (these subwords will have little to no impact on the result of the training) and break down the total 28.2k subwords into classes depending on the languages in which the particular subword was observed, see Table 6.

We see that the vocabulary is reasonably balanced, with each language having 20–30% of subwords unique to it. English and Estonian share 10% subwords not seen in Russian while Russian shares only 0–1.39% of subwords with each of the other languages. Overall 8.89% of subwords are seen in all three languages.

A particularly interesting subset is the one where parent languages help the child model, in other words subwords appearing anywhere in English and also tokens common to Estonian and Russian. For this set of languages, this amounts to $20.69+10.06+1.39+0.0+8.89 = 41.03\%$. We list this number on a separate line in Table 6, “From parent”. These subwords get their embeddings trained better thanks to the parent model.

Table 7 summarizes this analysis for several language sets, listing what portion of subwords is unique to individual languages in the set, what portion is shared by all the languages and what portion of subwords benefits from the parent training. We see a similar picture across the board, only AR-RU-ET-EN stands out with the very low num-

Languages	Unique in a Lang.	In All	From Parent
ET-EN-FI	24.4-18.2-26.2	19.5	49.4
ET-EN-RU	29.9-20.7-29.0	8.9	41.0
ET-EN-CS	29.6-17.5-21.2	20.3	49.2
AR-RU-ET-EN	28.6-27.7-21.2-9.1	4.6	6.2
ES-FR-ET-EN	15.7-13.0-24.8-8.8	18.4	34.1
ES-RU-ET-EN	14.7-31.1-21.3-9.3	6.0	21.4
FR-RU-ET-EN	12.3-32.0-22.3-8.1	6.3	23.1

Table 7: Summary of vocabulary overlaps for the various language sets. All figures in % of the shared vocabulary.

	BLEU	nPER	nTER	nCDER	chrF3	nCharacTER
Base ENET	16.13	47.13	32.45	36.41	48.38	33.23
ENRU+ENET	19.10	50.87	36.10	39.77	52.12	39.39
ENCS+ENET	19.30	51.51	36.84	40.42	52.71	40.81

Table 8: Various automatic scores on ENET test set. Scores prefixed “n” reported as $(1 - \text{score})$ to make higher numbers better.

ber of subwords (6.2%) available already in the parent. The parent AR-RU thus offered very little word knowledge to the child and yet lead to a gain in BLEU.

6.2 Output Analysis

Since we rely on automatic analysis, we need to prevent some potential overestimations of translation quality due to BLEU. For this, we took a closer look at the baseline ENET model (BLEU of 17.03 in Table 2) and two ENET childs derived from ENCS (BLEU of 20.41) and ENRU parent (BLEU 20.09).

Table 8 confirms the improvements are not an artifact of uncased BLEU. The gains are apparent with several (now cased) automatic scores.

As documented in Table 9, the improved outputs are considerably longer. In the table, we show also individual n -gram precisions and brevity penalty (BP) of BLEU. The longer output clearly helps to reduce the incurred BP but the improvements are also apparent in n -gram precisions. In other words, the observed gain cannot be attributed solely to producing longer outputs.

Table 10 explains the gains in unigram precisions by checking which tokens in the improved outputs (the parent followed by the child) were present also in the baseline (child-only, denoted “b” in Table 10) and/or confirmed by the reference (denoted “r”). We see that about $44+20\%$ of tokens of improved outputs can be seen as “unchanged” compared to the baseline because they appear already in the baseline output (“b”). (The

	Length	BLEU Components	BP
Base ENET	35326	48.1/21.3/11.3/6.4	0.979
ENRU+ENET	35979	51.0/24.2/13.5/8.0	0.998
ENCS+ENET	35921	51.7/24.6/13.7/8.1	0.996

Table 9: Candidate total length, BLEU n -gram precisions and brevity penalty (BP). The reference length in the matching tokenization was 36062.

	ENRU+ENET	ENCS+ENET
rb	15902 (44.2 %)	15924 (44.3 %)
-	9635 (26.8 %)	9485 (26.4 %)
b	7209 (20.0 %)	7034 (19.6 %)
r	3233 (9.0 %)	3478 (9.7 %)
Total	35979 (100.0 %)	35921 (100.0 %)

Table 10: Comparison of improved outputs vs. the baseline and reference.

44% “rb” tokens are actually confirmed by the reference.)

The differing tokens are more interesting: “-” denotes the cases when the improved system produced something different from the baseline and also from the reference. Gains in BLEU are due to “r” tokens, i.e. tokens only in the improved outputs and the reference but not the baseline “b”. For both parent setups, there are about 9–9.7 % of such tokens. We looked at these 3.2k and 3.5k tokens and we have to conclude that these are regular *Estonian* words; no Czech or Russian leaks to the output and the gains are *not* due to simple token types common to all the languages (punctuation, numbers or named entities). We see identical BLEU gains even if we remove all such simple tokens from the candidates and references. A better explanation of the gains thus still has to be sought for.

7 Related Work

Firat et al. (2016) propose multi-way multi-lingual systems, with the main goal of reducing the total number of parameters needed to cater multiple source and target languages. To keep all the language pairs “active” in the model, a special training schedule is needed. Otherwise, catastrophic forgetting would remove the ability to translate among the languages trained earlier.

Johnson et al. (2017) is another multi-lingual approach: all translation pairs are simply used at once and the desired target language is indicated with a special token at the end of the source side. The model implicitly learns translation between

many languages and it can even translate among language pairs never seen together.

Lack of parallel data can be tackled by unsupervised translation (Artetxe et al., 2018; Lample et al., 2018). The general idea is to mix monolingual training of autoencoders for the source and target languages with translation trained on data translated by the previous iteration of the system.

When no parallel data are available, the training set of closely related high-resource pair can be used with transliteration approach as described in Karakanta et al. (2018).

Aside from the common back-translation (Sennrich et al., 2016a; Kocmi et al., 2018), simple copying of target monolingual data back to source (Currey et al., 2017) has been also shown to improve translation quality in low-data conditions.

Similar to transfer learning is also curriculum learning (Bengio et al., 2009; Kocmi and Bojar, 2017), where the training data are ordered from foreign out-of-domain to the in-domain training examples.

8 Conclusion

We presented a simple method for transfer learning in neural machine translation based on training a parent high-resource pair followed a low-resource language pair dataset. The method works for shared source or target side as well as for language pairs that do not share any of the translation sides. We observe gains also from totally unrelated language pairs, although not always significant.

One interesting trick we propose for low-resource languages is to start training in the opposite direction and swap to the main one afterwards.

The reasons for the gains are yet to be explained in detail but our observations indicate that the key factor is the size of the parent corpus rather than e.g. vocabulary overlaps.

Acknowledgments

This study was supported in parts by the grants SVV 260 453, GAUK 8502/2016, and 18-24210S of the Czech Science Foundation. This work has been using language resources and tools stored and distributed by the LINDAT/CLARIN project of the Ministry of Education, Youth and Sports of the Czech Republic (projects LM2015071 and OP VVV VI CZ.02.1.01/0.0/0.0/16 013/0001781).

References

- Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2018. Unsupervised neural machine translation. In *Proceedings of the Sixth International Conference on Learning Representations*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural Machine Translation by Jointly Learning to Align and Translate. *CoRR*, abs/1409.0473.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48. ACM.
- Ondřej Bojar, Ondřej Dušek, Tom Kocmi, Jindřich Libovický, Michal Novák, Martin Popel, Roman Sudarikov, and Dušan Variš. 2016. Czeng 1.6: enlarged czech-english parallel corpus with processing tools dockered. In *International Conference on Text, Speech, and Dialogue*, pages 231–238. Springer.
- Anna Currey, Antonio Valerio Miceli Barone, and Kenneth Heafield. 2017. Copied monolingual data improves low-resource neural machine translation. In *Proceedings of the Second Conference on Machine Translation*, pages 148–156.
- Raj Dabre, Tetsuji Nakagawa, and Hideto Kazawa. 2017. An empirical study of language relatedness for transfer learning in neural machine translation. In *Proceedings of the 31st Pacific Asia Conference on Language, Information and Computation*, pages 282–286.
- Ahmed El Kholly, Nizar Habash, Gregor Leusch, Evgeny Matusov, and Hassan Sawaf. 2013. Language independent connectivity strength features for phrase pivot statistical machine translation. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 412–418, Sofia, Bulgaria. Association for Computational Linguistics.
- 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. Association for Computational Linguistics.
- Petra Galuščáková and Ondrej Bojar. 2012. Improving smt by using parallel data of a closely related language. In *Proc. of HLT*, pages 58–65.
- Duc Tam Hoang and Ondrej Bojar. 2016. Pivoting methods and data for czech-vietnamese translation via english. *Baltic Journal of Modern Computing*, 4(2):190–202.
- Melvin Johnson, Mike Schuster, Quoc Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernand a Vigas, 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.
- Alina Karakanta, Jon Dehdari, and Josef van Genabith. 2018. Neural machine translation for low-resource languages without parallel corpora. *Machine Translation*, 32(1):167–189.
- Tom Kocmi and Ondřej Bojar. 2017. Curriculum Learning and Minibatch Bucketing in Neural Machine Translation. In *Recent Advances in Natural Language Processing 2017*.
- Tom Kocmi, oman Sudarikov, and Ondřej Bojar. 2018. CUNI Submissions in WMT18. In *Proceedings of the 3rd Conference on Machine Translation (WMT)*, Brussels, Belgium.
- Philipp Koehn. 2004. Statistical significance tests for machine translation evaluation. In *Proceedings of EMNLP*, volume 4, pages 388–395.
- Philipp Koehn and Rebecca Knowles. 2017. Six challenges for neural machine translation. In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.
- Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018. Phrase-based & neural unsupervised machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Preslav Nakov and Hwee Tou Ng. 2012. Improving statistical machine translation for a resource-poor language using related resource-rich languages. *Journal of Artificial Intelligence Research*, 44:179–222.
- Toan Q. Nguyen and David Chiang. 2017. Transfer learning across low-resource, related languages for neural machine translation. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 296–301. Asian Federation of Natural Language Processing.
- Xing Niu, Michael Denkowski, and Marine Carpuat. 2018. Bi-directional neural machine translation with synthetic parallel data. In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 84–91, Melbourne, Australia. Association for Computational Linguistics.
- Robert Östling, Yves Scherrer, Jörg Tiedemann, Gongbo Tang, and Tommi Nieminen. 2017. The helsinki neural machine translation system. In *Proceedings of the Second Conference on Machine Translation*, pages 338–347, Copenhagen, Denmark. Association for Computational Linguistics.

- Martin Popel and Ondej Bojar. 2018. Training Tips for the Transformer Model. *The Prague Bulletin of Mathematical Linguistics*, 110(1):43–70.
- Rico Sennrich, Alexandra Birch, Anna Currey, Ulrich Germann, Barry Haddow, Kenneth Heafield, Antonio Valerio Miceli Barone, and Philip Williams. 2017. The university of edinburgh’s neural mt systems for wmt17. In *Proceedings of the Second Conference on Machine Translation, Volume 2: Shared Task Papers*, pages 389–399, Copenhagen, Denmark. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. Improving Neural Machine Translation Models with Monolingual Data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. 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. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Ashish Vaswani, Samy Bengio, Eugene Brevdo, Francois Chollet, Aidan Gomez, Stephan Gouws, Llion Jones, Lukasz Kaiser, Nal Kalchbrenner, Niki Parmar, Ryan Sepassi, Noam Shazeer, and Jakob Uszkoreit. 2018. Tensor2Tensor for Neural Machine Translation. In *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Papers)*, pages 193–199, Boston, MA. Association for Machine Translation in the Americas.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6000–6010. Curran Associates, Inc.
- Hua Wu and Haifeng Wang. 2007. Pivot language approach for phrase-based statistical machine translation. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 856–863, Prague, Czech Republic. Association for Computational Linguistics.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Michal Ziemski, Marcin Junczys-Dowmunt, and Bruno Pouliquen. 2016. The united nations parallel corpus v1. 0. In *LREC*.
- Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. 2016. Transfer learning for low-resource neural machine translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1568–1575, Austin, Texas. Association for Computational Linguistics.