

Facebook FAIR’s WMT19 News Translation Task Submission

Nathan Ng, Kyra Yee, Alexei Baevski, Myle Ott, Michael Auli, Sergey Edunov

Facebook AI Research,
Menlo Park, CA & New York, NY.

Abstract

This paper describes Facebook FAIR’s submission to the WMT19 shared news translation task. We participate in four language directions, English \leftrightarrow German and English \leftrightarrow Russian in both directions. Following our submission from last year, our baseline systems are large BPE-based transformer models trained with the FAIRSEQ sequence modeling toolkit. This year we experiment with different bitext data filtering schemes, as well as with adding filtered back-translated data. We also ensemble and fine-tune our models on domain-specific data, then decode using noisy channel model reranking. Our system improves on our previous system’s performance by 4.5 BLEU points and achieves the best case-sensitive BLEU score for the translation direction English \rightarrow Russian.

1 Introduction

We participate in the WMT19 shared news translation task in two language pairs and four language directions, English \rightarrow German (En \rightarrow De), German \rightarrow English (De \rightarrow En), English \rightarrow Russian (En \rightarrow Ru), and Russian \rightarrow English (Ru \rightarrow En). Our methods are based on techniques and approaches used in our submission from last year (Edunov et al., 2018), including the use of subword models, (Sennrich et al., 2016), large-scale back-translation, and model ensembling. We train all models using the FAIRSEQ sequence modeling toolkit (Ott et al., 2019). Although document level context for En \rightarrow De is now available, all our systems are pure sentence level systems. In the future, we expect better results from leveraging this additional context information.

Compared to our WMT18 submission, we also decide to compete in the En \leftrightarrow Ru and De \rightarrow En translation directions. Although all four directions are considered high resource settings where lar-

ge amounts of bitext data is available, we demonstrate that leveraging high quality monolingual data through back-translation is still very important. For all language directions, we back-translate the Newscrawl dataset using a reverse direction bitext system. In addition to back-translating the relatively clean Newscrawl dataset, we also experiment with back-translating portions of the much larger and noisier Commoncrawl dataset. For our final models, we apply a domain-specific fine-tuning process and decode using noisy channel model reranking (Anonymous, 2019).

Compared to our WMT18 submission in the En \rightarrow De direction, we observe substantial improvements of 4.5 BLEU. Some of these gains can be attributed to differences in dataset quality, but we believe most of the improvement comes from larger models, larger scale back-translation, and noisy channel model reranking with strong channel and language models.

2 Data

For the En \leftrightarrow De language pair we use all available bitext data including the bicleaner version of Paracrawl. For our monolingual data we use English and German Newscrawl. Although our language models were trained on document level data, we did not use document level boundaries in our final decoding step, so all our systems are purely sentence level systems.

For the En \leftrightarrow Ru language pair we also use all available bitext data. For our monolingual data we use English and Russian Newscrawl as well as a filtered portion of Russian Commoncrawl. We choose to use Russian Commoncrawl to augment our monolingual data due to the relatively small size of Russian Newscrawl compared to English and German.

2.1 Data Preprocessing

Similar to last year’s submission for En→De, we normalize punctuation and tokenize all data with the Moses tokenizer (Koehn et al., 2007). For En↔De we use joint byte pair encodings (BPE) with 32K split operations for subword segmentation (Sennrich et al., 2016). For En↔Ru, we learn separate BPE encodings with 24K split operations for each language. Systems trained with this separate BPE encoding performed significantly better than those trained with joint BPE.

2.2 Data Filtering

2.2.1 Bitext

Large datasets crawled from the internet are naturally very noisy and can potentially decrease the performance of a system if they are used in their raw form. Cleaning these datasets is an important step to achieving good performance on any downstream tasks.

We apply language identification filtering (`langid`; Lui et al., 2012), keeping only sentence pairs with correct languages on both sides. Although not the most accurate method of language identification (Joulin et al., 2016), one side effect of using `langid` is the removal of very noisy sentences consisting of mostly garbage tokens, which are classified incorrectly and filtered out.

We also remove sentences longer than 250 tokens as well as sentence pairs with a source/target length ratio exceeding 1.5. In total, we filter out about 30% of the original bitext data. See Table 1 for details on the bitext dataset sizes.

2.2.2 Monolingual

For monolingual Newscrawl data we also apply `langid` filtering. Since the monolingual Newscrawl corpus for Russian is significantly smaller than that of German or English, we augment our monolingual Russian data with data from the commoncrawl corpus. Commoncrawl is the largest monolingual corpus available for training but is also very noisy. In order to select a limited amount of high quality, in-domain sentences from the larger corpus, we adopt the method of Moore and Lewis (2010) for selecting in-domain data (§3.2.1).

	En-De	En-Ru
No filter	38.8M	38.5M
+ length filter	35.7M	33.4M
+ <code>langid</code> filter	27.7M	26.0M

Table 1: Number of sentences in bitext datasets for different filtering schemes

3 System Overview

3.1 Base System

Our base system is based on the big Transformer architecture (Vaswani et al., 2017) as implemented in FAIRSEQ. We experiment with increasing network capacity by increasing embed dimension, FFN size, number of heads, and number of layers. We find that using a larger FFN size (8192) gives a reasonable improvement in performance while maintaining a manageable network size. All subsequent models, including ensembles, use this larger FFN Transformer architecture.

We trained all our models using FAIRSEQ (Ott et al., 2019) on 128 Volta GPUs, following the setup described in Ott et al. (2018)

3.2 Large-scale Back-translation

Back-translation is an effective and commonly used data augmentation technique to incorporate monolingual data into a translation system. Back-translation first trains an intermediate target-to-source system that is used to translate monolingual target data into additional synthetic parallel data. This data is used in conjunction with human translated bitext data to train the desired source-to-target system.

In this work we used back-translations obtained by sampling (Edunov et al., 2018) from an ensemble of three target-to-source models. We found that models trained on data back-translated using an ensemble instead of a single model performed better (Table 2). Previous work also found that upsampling the bitext data can improve back-translation (Edunov et al., 2018). We adopt this method to tune the amount of bitext and synthetic data the model is trained on. We find a ratio of 1:1 synthetic to bitext data to perform the best.

3.2.1 Back-translating Commoncrawl

The amount of monolingual Russian data available in the Newscrawl dataset is significantly smaller than that of English and German (Table 3). In

	En→Ru	
	Single Model	Ensemble
newstest15	35.98	36.32
newstest16	32.78	33.28
newstest17	36.57	36.77
newstest18	34.72	34.72

Table 2: SacreBLEU for English-Russian models trained with data back-translated using a single model vs. an ensemble of two models

	En	De	Ru
News crawl	434M	559M	80M
+ langid filter	424M	521M	76M
Common crawl	-	-	1.2B
+ KenLM filter	-	-	60M
Total	424M	521M	136M

Table 3: Number of sentences in monolingual datasets available for back-translation

order to increase the amount of monolingual Russian data for back-translation, we experiment with incorporating Commoncrawl data. Commoncrawl is a much larger and noisier dataset compared to News crawl, and is also non-domain specific. We experiment with methods to identify a subset of Commoncrawl that is most similar to News crawl. Specifically, we use the in-domain filtering method described in Moore and Lewis (2010).

Given an in domain corpus I , in this case News crawl, and a non-domain specific corpus N , in this case Commoncrawl, we would like to find the subcorpus N_I that is drawn from the same distribution as I . For any given sentence s , we can calculate, using Bayes’ rule, the probability a sentence s in N is drawn from N_I

$$P(N_I|s, N) = \frac{P(s|N_I)P(N_I|N)}{P(s|N)} \quad (1)$$

We ignore the $P(N_I|N)$ term, since it will be constant for any given I and N , and use $P(s|I)$ instead of $P(s|N_I)$, since I and N_I are drawn from the same distribution. Moving into the log domain, we can calculate the probability score for a sentence s by $\log P(N_I|s, N) = \log P(s|I) - \log P(s|N)$, or after normalizing for length, $H_I(s) - H_N(s)$, where $H_I(s)$ and $H_N(s)$ are the word-normalized cross entropy scores for a sentence s according to language models L_I and

	En-De	De-En	En-Ru	Ru-En
newstest12	26.7	28.0	-	-
newstest13	27.8	27.6	42.7	27.6
newstest14	21.4	24.0	32.3	22.4
newstest15	25.1	24.6	34.7	21.8
newstest16	24.5	22.0	35.5	19.4
newstest17	25.0	21.9	37.9	19.5
newstest18	25.1	26.0	39.3	20.0

Table 4: Perplexity scores for language models on bolded target languages in all translation directions

L_N trained on I and N respectively.

Our corpora are very large and we therefore use an n -gram model (Heafield, 2011) rather than a neural language model which would be much slower to train and evaluate. We train two language models L_I and L_N on News crawl and Common crawl respectively, then score every sentence s in Common crawl by $H_I(s) - H_N(s)$. We select a cutoff of 0.01, and use all sentences that score higher than this value for back-translation, or about 5% of the entire dataset.

3.3 Fine-tuning

Fine-tuning with domain-specific data is a common and effective method to improve translation quality for a downstream task. After completing training on the bitext and back-translated data, we train for an additional epoch on a smaller in-domain corpus. For De→En, we fine-tune on test sets from previous years, including newstest2012, newstest2013, newstest2015, and newstest2017. For En→De, we fine-tune on previous test sets as well as the News-Commentary dataset. For En↔Ru we fine-tune on a combination of News-Commentary, newstest2013, newstest2015, and newstest2017. The other test sets are held out for other tuning procedures and evaluation metrics.

3.4 Noisy Channel Model Reranking

N -best reranking is a method of improving translation quality by scoring and selecting a candidate hypothesis from a list of n -best hypotheses generated by a source-to-target, or forward model. For our submissions, we rerank using a noisy channel model approach (Anonymous, 2019).

Given a target sequence y and a source sequence x , the noisy channel approach applies Bayes’ rule

to model

$$P(y|x) = \frac{P(x|y)P(y)}{P(x)} \quad (2)$$

Since $P(x)$ is constant for a given source sequence x , we can ignore it. We refer to the remaining terms $P(y|x)$, $P(x|y)$, and $P(y)$, as the forward model, channel model, and language model respectively. In order to combine these scores for reranking, we calculate for every one of our n -best hypotheses:

$$\log P(y|x) + \lambda_1 \log P(x|y) + \lambda_2 \log P(y) \quad (3)$$

The weights λ_1 and λ_2 are determined by tuning them with a random search on a validation set and selecting the weights that give the best performance. In addition, we also tune a length penalty.

For all translation directions, our forward models are ensembles of fine-tuned and back-translated models. Since we compete in both directions for both language pairs, for any given translation direction we can use the forward model for the reverse direction as the channel model. Our language models for each of the target languages English, German, and Russian, are big Transformer decoder models with FFN 8192. We train the language models on the monolingual Newscrawl dataset, and use document level context for the English and German models. Perplexity scores for the language models on the bolded target language of each translation direction are shown in table 4. With a smaller amount of monolingual Russian data available, we observe that our Russian language model performs worse than the German and English language models.

To select the length penalty and weights, λ_1 and λ_2 , for decoding, we use random search, choosing values in the range $[0, 2)$ for the weights and values in the range $[0, 1)$ for the length penalty. For all language directions, we choose the weights that give the highest BLEU score on a combined dataset of newstest2014 and newstest2016.

To run our final decoding step, we first use the forward model with beam size 50 to generate an n -best list. We then use the channel and language models to score each of these hypotheses, using the weights and length penalty tuned previously. Finally, we select the hypothesis with the highest score as our output.

System	En→De	
	news2017	news2018
baseline	30.90	45.40
+ langid filtering	30.78	46.43
+ ffn 8192	31.15	46.28
+ ensemble	31.55	47.09
+ BT	33.62	46.66
+ fine tuning	-	47.61
+ ensemble	-	49.27
+ reranking	-	50.63
WMT’18 submission	-	46.10
WMT’19 submission	42.7	

Table 5: SacreBLEU scores on English→German.

3.5 Postprocessing

For En→De and En→Ru, we also change the standard English quotation marks (“...”) to German-style quotation marks („...“).

4 Results

Results and ablations for En→De are shown in Table 5, De→En in Table 6, En→Ru in Table 7 and Ru→En in Table 8. We report case-sensitive SacreBLEU scores using SacreBLEU (Post, 2018)¹, using international tokenization for En→Ru. In the final row of each table we also report the case-sensitive BLEU score of our submitted system on this year’s test set. All single models and individual models within ensembles are averages of the last 10 checkpoints of training. Our baseline systems are big Transformers as described in (Vaswani et al., 2017). The baselines were trained with minimally filtered data, removing only those sentences longer than 250 words and exceeding a source/target length ratio of 1.5 This setup gave us a reasonable baseline to evaluate data filtering.

4.1 English→German

For En→De, langid filtering, larger FFN, and ensembling improve our baseline performance on news2018 by about 1.5 BLEU. Note that our best

¹SacreBLEU signatures:

BLEU+case.mixed+lang.en-de+numrefs.1+smooth.exp+test.wmt{17/18}+tok.13a+version.1.2.11,
 BLEU+case.mixed+lang.de-en+numrefs.1+smooth.exp+test.wmt{17/18}+tok.13a+version.1.2.11,
 BLEU+case.mixed+lang.ru-en+numrefs.1+smooth.exp+test.wmt{17/18}+tok.13a+version.1.2.11,
 BLEU+case.mixed+lang.en-ru+numrefs.1+smooth.exp+test.wmt{17/18}+tok.intl+version.1.2.11

System	De→En	
	news2017	news2018
baseline	37.28	45.32
+ langid and ffn 8192	38.45	46.16
+ ensemble	38.82	46.76
+ BT	41.08	48.78
+ fine tuning	-	49.07
+ ensemble	-	49.60
+ reranking	-	51.13
WMT'19 submission	40.8	

Table 6: SacreBLEU scores on German→English.

bitext only systems already outperforms our system from last year by 1 BLEU point. This is perhaps due to the addition of higher quality bitext data and improved data filtering techniques. The addition of back-translated (BT) data improves single model performance by only 0.3 BLEU, but combining this with fine-tuning and ensembling gives us a total of 3 BLEU. Finally, applying reranking on top of these strong ensembled systems gives another 1.4 BLEU.

4.2 German→English

For De→En, as with En→De, we see similar improvements with langid filtering, larger FFN, and ensembling on the order of 1.4 BLEU. Compared to En→De however, we also observe that the addition of back-translated data is much more significant, improving single model performance by over 2.5 BLEU. Fine-tuning, ensembling, and reranking add an additional 2.4 BLEU, with reranking contributing 1.5 BLEU, a majority of the improvement.

4.3 English→Russian

For En→Ru, we observe large improvements of 2.4 BLEU over a bitext-only model after applying langid filtering, larger FFN, and ensembling. Since we start with a lower quality initial En↔Ru bitext dataset, we observe a large improvement of 3.5 BLEU by adding back-translated data. Augmenting this back-translated data with Common-crawl adds an additional 0.2 BLEU. Finally, applying fine-tuning, ensembling, and reranking adds 2.2 BLEU, with reranking contributing 1 BLEU.

System	En→Ru	
	news2017	news2018
baseline	35.42	31.53
+ langid filtering	35.69	31.77
+ ffn 8192	36.66	33.49
+ ensemble	37.42	33.93
+ BT NC	40.09	37.07
+ BT NC + CC	40.42	37.3
+ fine tuning	-	37.74
+ ensemble	-	38.59
+ reranking	-	39.53
WMT'19 submission	36.3	

Table 7: SacreBLEU scores on English→Russian

System	Ru→En	
	news2017	news2018
baseline	37.07	32.69
+ langid and ffn 8192	37.72	33.44
+ ensemble	38.69	34.29
+ BT	41.68	36.49
+ fine tuning	-	38.54
+ ensemble	-	38.96
+ reranking	-	40.16
WMT'19 submission	40.0	

Table 8: SacreBLEU scores on Russian→English

4.4 Russian→English

For Ru→En, we observe similar trends to En↔De, with langid filtering, larger FFN, and ensembling improving performance of a bitext-only system by 1.6 BLEU. Backtranslation adds 3 BLEU, again most likely due to the lower quality bitext data available. Fine-tuning, ensembling, and reranking add almost 4 BLEU, with reranking contributing 1.2 BLEU.

4.5 Reranking

For every language direction, reranking gives a significant improvement, even when applied on top of an ensemble of very strong back-translated models. We also observe that the biggest improvement of 1.5 BLEU comes in the De→En language direction, and the smallest improvement of 1 BLEU in the En→Ru direction. This is perhaps due to the relatively weak Russian language model, which is trained on significantly less data

compared to English and German. Improving our language models may lead to even greater improvements with reranking.

5 Conclusions

This paper describes Facebook FAIR’s submission to the WMT19 news translation task. For all four translation directions, En↔De and En↔Ru, we use the same strategy of filtering bitext data, back-translating monolingual data, then training strong individual models on a combination of this data. Each of these models is fine-tuned and ensembled into a final system that is used for decoding with noisy channel model reranking. We demonstrate the effectiveness of our reranking approach, even when applied on top of very strong systems, and achieve the best case-sensitive BLEU score for En→Ru and competitive results in all other directions.

References

- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. Understanding back-translation at scale. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 489–500.
- Kenneth Heafield. 2011. Kenlm: faster and smaller language model queries. In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, pages 187–197.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2016. Bag of tricks for efficient text classification. *arXiv preprint arXiv:1607.01759*.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondrej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *ACL Demo Session*.
- 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. Association for Computational Linguistics.
- Robert Moore and William Lewis. 2010. Intelligent selection of language model training data. In *Proceedings of the ACL 2010 Conference Short Papers*, pages 220–224.
- 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*.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. Scaling neural machine translation. In *Proc. of WMT*.
- Matt Post. 2018. A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191. Association for Computational Linguistics.
- 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*, pages 1715–1725.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan Gomez, ukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st Conference on Neural Information Processing Systems*.