

Tilde at WMT 2020: News Task Systems

Rihards Krišlauks^{†‡} and Mārcis Pinnis^{†‡}

[†]Tilde / Vienības gatve 75A, Rīga, Latvia

[‡]Faculty of Computing, University of Latvia / Raiņa bulv. 19, Rīga, Latvia

{firstname.lastname}@tilde.lv

Abstract

This paper describes Tilde’s submission to the WMT2020 shared task on news translation for both directions of the English↔Polish language pair in both the constrained and the unconstrained tracks. We follow our submissions from the previous years and build our baseline systems to be morphologically motivated sub-word unit-based Transformer base models that we train using the Marian machine translation toolkit. Additionally, we experiment with different parallel and monolingual data selection schemes, as well as sampled back-translation. Our final models are ensembles of Transformer base and Transformer big models which feature right-to-left re-ranking.

1 Introduction

This year, we developed both constrained and unconstrained NMT systems for the English↔Polish language pair. We base our methods on the submissions of the previous years (Pinnis et al., 2017b, 2018, 2019) including methods for parallel data filtering from Pinnis (2018). Specifically, we lean on Pinnis (2018) and Junczys-Dowmunt (2018) for data selection and filtering, (Pinnis et al., 2017b) for morphologically motivated sub-word units and synthetic data generation, Edunov et al. (2018) for sampled back-translation and finally Morishita et al. (2018) for re-ranking with right-to-left models. We use the Marian toolkit (Junczys-Dowmunt et al., 2018) to train models of Transformer architecture (Vaswani et al., 2017).

Although document level NMT as showcased by (Junczys-Dowmunt, 2019) have yielded promising results for the English-German language pair, we were not able to collect sufficient document level data for the English-Polish language pair. As a result, all our systems this year translate individual sentences.

The paper is further structured as follows: Section 2 describes the data used to train our NMT systems, Section 3 describes our efforts to identify the best-performing recipes for training of our final systems, Section 5 summarises the results of our final systems, and Section 6 concludes the paper.

2 Data

For training of the constrained NMT systems, we used data from the WMT 2020 shared task on news translation¹. For unconstrained systems, we used data from the Tilde Data Library². The 10 largest publicly available datasets that were used to train the unconstrained systems were Open Subtitles from the Opus corpus (Tiedemann, 2016), ParaCrawl (Banón et al., 2020) (although it was discarded due to noise found in the corpus), DGT Translation Memories (Steinberger et al., 2012), Microsoft Translation and User Interface Strings Glossaries³ from multiple releases up to 2018, the Tilde MODEL corpus (Rozis and Skadiņš, 2017), WikiMatrix (Schwenk et al., 2019), Digital Corpus of the European Parliament (Hajlaoui et al., 2014), JRC-Acquis (Steinberger et al., 2006), Europarl (Koehn, 2005), and the QCRI Educational Domain Corpus (Abdelali et al., 2014).

2.1 Data Filtering and Pre-Processing

First, we filtered data using Tilde’s parallel data filtering methods (Pinnis, 2018) that allow discarding sentence pairs that are corrupted, have low content overlap, feature wrong language content, feature too high non-letter ratio, etc. The exact filter configuration is defined in the paper by (Pinnis, 2018).

Then, we pre-processed all data using Tilde’s parallel data pre-processing workflow that nor-

¹<http://www.statmt.org/wmt20/translation-task.html>

²<https://www.tilde.com/products-and-services/data-library>

³<https://www.microsoft.com/en-us/language/translations>

Scenario	Lang. pair	Raw	Tilde	Filtered +DCCEF
(c)	En → Pl	10.8M	6.5M	4.3M
	Pl → En			4.3M
(u)	En → Pl	55.4M	31.5M	23.3M
	Pl → En			24.1M
(u) w/o PC	En → Pl	48.8M	27.0M	21.6M
	Pl → En			21.3M

Table 1: Parallel data statistics before and after filtering. (c) - constrained, (u) - unconstrained, “w/o PC” - “without ParaCrawl”.

malizes punctuation (quotation marks, apostrophes, decodes HTML entities, etc.), identifies non-translatable entities and replaces them with placeholders (e.g., e-mail addresses, Web site addresses, XML tags, etc.), tokenises the text using Tilde’s regular expression-based tokeniser, and applies true-casing.

In preliminary experiments, we identified also that morphology-driven word splitting (Pinnis et al., 2017a) for English↔Polish allowed us to increase translation quality by approximately 1 BLEU point. The finding complies with our findings from previous years (Pinnis et al., 2018, 2017b). Therefore, we applied morphology-driven word splitting also for this year’s experiments.

Then, we trained baseline NMT models (see Section 3.2) and language models, which are necessary for dual conditional cross-entropy filtering (DCCEF) (Junczys-Dowmunt, 2018) in order to select parallel data that is more similar to the news domain (for usefulness of DCCEF, refer to Section 3.3). For in-domain (i.e., news) and out-of-domain language model training, we used four monolingual datasets of 3.7M and 10.6M sentences⁴ for the constrained and unconstrained scenarios respectively. Once the models were trained, We filtered parallel data using DCCEF. The parallel data statistics before and after filtering are given in Table 1.

For our final systems, we also generated synthetic data by randomly replacing one to three content words on both source and target sides with unknown token identifiers. This has shown to increase robustness of NMT systems when dealing with rare or unknown phenomena (Pinnis et al., 2017a). This process almost doubles the size of

⁴The sizes correspond to the smallest monolingual in-domain dataset, which in both cases were news in Polish. For other datasets, random sub-sets were selected.

the corpora, therefore, this was not done for the datasets that were used for the experiments documented in Section 3.

For backtranslation experiments, we used all available monolingual data from the WMT shared task on news translation. In order to make use of the Polish CommonCrawl corpus, we scored sentences using the in-domain language models and selected top-scoring sentences as additional monolingual data for back-translation.

Many of the data processing steps were sped up via parallelization with GNU Parallel (Tange, 2011).

3 Experiments

In this section, we describe the details of the methods used and experiments performed to identify the best-performing recipes for training of Tilde’s NMT systems for the WMT 2020 shared task on news translation. All experiments that are described in this section were carried out on the constrained datasets unless specifically indicated that also unconstrained datasets were used.

3.1 NMT architecture

All NMT systems that are described further have the Transformer architecture (Vaswani et al., 2017). We trained the systems using the Marian toolkit (Junczys-Dowmunt et al., 2018). The Transformer *base* model configuration was used throughout the experiments except for the experiments with the *big* model configuration that are described in Section 5. We used gradient accumulation over multiple physical batches (the *--optimizer-delay* parameter in Marian) to increase the effective batch size to around 1600 sentences in the *base* model experiments and 1000 sentences in *big* model experiments. The Adam optimizer with a learning rate of 0.0005 and with 1600 warm-up update steps (i.e., the learning rate linearly rises during warm-up; afterwards decays proportionally to the inverse of the square root of the step number) was used. For language model training, a learning rate of 0.0003 was used.

3.2 Baseline models

We trained baseline models using the Transformer *base* configuration as defined in Section 3.1. The validation results for the baseline NMT systems are provided in Table 2. As we noticed last year that the ParaCrawl corpus contained a large proportion (by our estimates up to 50%) (Pinnis et al., 2019)

System	En \rightarrow Pl	Pl \rightarrow En
Constrained		
Baseline	21.67	32.69
+DCCEF	22.19	33.45
Unconstrained		
Baseline	21.86	33.08
+DCCEF	22.51	30.86
Baseline w/o ParaCrawl	24.29	29.47
+DCCEF	22.60	28.59

Table 2: Comparison of baseline NMT systems trained on data that were prepared with and without DCCEF.

of machine translated content, we trained baseline systems with and without ParaCrawl. It can be seen that when training the En \rightarrow Pl unconstrained system using ParaCrawl, we loose over 2 BLEU points. This is because most machine translated content is on the non-English (in this case Polish) side. For the Pl \rightarrow En direction, the machine-translated content acts as back-translated data and, therefore, does not result in quality degradation. Further, our Pl \rightarrow En systems are trained using ParaCrawl, and En \rightarrow Pl systems – without ParaCrawl.

3.3 Dual Conditional Cross-Entropy Filtering

After the baseline systems, we analysed whether DCCEF allows improving translation quality. The validation results in Table 2 show that translation quality increases for the constrained systems, but degrades for the unconstrained systems. Further, we used DCCEF only for the constrained scenario systems.

3.4 Back-translation

We used monolingual data back-translation to adapt the NMT systems to the news domain. Edunov et al. (2018) has shown that using output sampling instead of beam-search during back-translation yields better-performing NMT systems. Hence, we exclusively used output sampling for monolingual data back-translation. However, due to the abundance of monolingual data for both translation directions, we experimented with different rates of upsampling and back-translated data cutoff to determine whether translation performance doesn't degrade in the presence of a too small proportion of bitext data during training.

Another dimension of inquiry was with different

strategies for data filtering in the preparation of the back-translated data. Ng et al. (2019) have described a method for domain data extraction from general domain monolingual data using domain and out-of-domain language models. We compared said method with a simpler alternative of using only an in-domain language model for in-domain data scoring. We sorted the monolingual data according to the scores produced by the in-domain language model or by the combination of in-domain and out-of-domain language model scores and experimented with different cutoff points when selecting data for back-translation.

Considering the above, we carried out experiments along two dimensions – 1) monolingual data selection strategy, which was either *combined* or *in-domain*, signifying whether the combined score of both language models or just the score from the in-domain language model was used, respectively, and 2) the bitext and synthetic data mixture selection strategy, which was one of:

- *original ratio* – all available bitext data for the translation direction were combined with all back-translated data having a score ≥ 0 , when using the *combined* selection strategy, or N top-scoring back-translated sentences, when using the *in-domain* selection strategy, where N was selected to match the amount of synthetic data selected in the *combined* case.
- *upsampled 1:1* – the same amount of synthetic data was selected as with the *original ratio* data mixture selection strategy, but bitext was upsampled to match the amount of synthetic data.
- *cutoff* $\{1:1, 1:2, 1:3\}$ – all available bitext data for the translation direction were combined with N top-scoring back-translated sentences where N was chosen so that the ratio of bitext to synthetic data was either 1:1, 1:2 or 1:3.

As a result of the above, we ended up with 96.8M sentences (14% retained) from the English monolingual corpus and 137M (99% retained) sentences from the Polish monolingual corpus after applying the *combined* data selection strategy. Consequently, the same amount of data was selected for the *in-domain* data selection strategy in the case of *original ratio* and *upsampled 1:1* data mixture selection strategies (i.e. when not doing *cutoff*).

	orig. ratio	ups. 1:1	cutoff		
			1:1	1:2	1:3
En → Pl					
combined	23.35	24.01	24.52	24.72	-
in-domain	22.10	22.92	25.02	25.28	25.24
Pl → En					
combined	31.19	33.45	33.29	33.60	-
in-domain	29.67	-	33.40	33.28	-

Table 3: En → Pl back-translation experiment results.

The results for back-translation experiments are presented in Table 3. The systems use the DCCEF-filtered constrained datasets and therefore are directly comparable to the constrained DCCEF systems in Table 2.

For our final systems, we use the *combined* selection strategy for Pl → En and the *in-domain* selection strategy for En → Pl. For unconstrained systems, we identified that there is no significant difference between translation quality; we used the *combined* selection strategy for both language pairs.

3.5 QHAdam optimizer

Last year (Pinnis et al., 2019) we used the QHAdam optimizer (Ma and Yarats, 2018) for model training, however, we hadn’t treated QHAdam and Adam the same in the experimental process, having dedicated substantially more effort to optimizer hyperparameter tuning for QHAdam than Adam. To make an unbiased comparison of the two optimizers, we trained corresponding system variants using QHAdam for the *combined cutoff 1:2*, *in-domain cutoff 1:2* and *in-domain cutoff 1:3* systems from Section 3.4 in the En → Pl translation direction. The BLEU scores for the experiments are found in Table 4. We see that QHAdam performs no better than Adam. We had also done more extensive experiments comparing QHAdam to Adam for a range of learning rate and warm-up step parameter settings on a different dataset, which showed a similar trend, however we do not present those results here. As a result, we didn’t choose QHAdam over Adam in this year’s competition.

We note, however, that we used the recommended safe defaults for the QHAdam’s hyperparameters – $v_1 = 0.8$, $v_2 = 0.7$ – and we haven’t performed a search over these values which could have yielded different results.

	combined cutoff 1:2	in-domain cutoff	
		1:2	1:3
Adam	24.72	25.28	25.24
QHAdam	24.86	24.98	25.00

Table 4: BLEU scores for the QHAdam experiments in the En → Pl translation direction.

3.6 Right-to-Left Re-Ranking

Morishita et al. (2018) report improving the translation performance by using right-to-left (R2L) re-ranking. The method employs a right-to-left model to re-score the n -best list outputs of a regular – left-to-right – model by multiplying both models’ translation probabilities. We implement R2L re-ranking the same as Morishita et al. (2018), but opted to use n -best lists with $n = 12$ (instead of $n = 10$).

The R2L re-ranking experiments were performed during the final stages of the competition, hence the baseline systems for those experiments were the final systems that were being prepared for submission to the news translation task. Therefore we present the results in Table 5 in the Results section. We find similar improvements as Morishita et al. (2018), albeit they are slightly smaller.

4 Final Systems

We chose the best performing system variants from Section 3 to serve as a base for the final submission for the news translation task. For the constrained scenario, we trained final systems using parallel data that were filtered with Tilde’s filtering methods and DCCEF, back-translated monolingual data using a ratio of 1:2 (different data selection methods were applied for both translation directions), and synthetic data featuring unknown phenomena. For the unconstrained scenario, we trained final systems using parallel data that were filtered only with Tilde’s filtering methods, back-translated monolingual data that were selected using the *combined* data selection strategy using a ratio of 1:1, and synthetic data featuring unknown phenomena. All models were trained using the Adam optimiser.

When preparing the final systems, we also employed R2L re-ranking (see Section 3.6), ensembling of the best three models, and trained Transformer models using the *big* model configuration.

	Constrained	Unconstrained
Pl → En		
Base	33.48	32.63
+R2L	34.34	33.29
Big	33.79	33.15
+R2L	<u>34.83</u>	33.45
Ensemble of 3	34.19	33.39
+R2L	34.80	33.53
En → Pl		
Base	25.64	26.12
+R2L	26.24	26.52
Big	25.59	26.47
+R2L	26.70	26.78
Ensemble of 3	26.07	26.86
+R2L	<u>26.73</u>	<u>27.12</u>

Table 5: Final system evaluation results (BLEU scores) on validation data (bold marks best scores; submitted systems are underlined).

5 Results

The BLEU scores for the systems that were evaluated for the final submission are shown in Table 5. The results show that right-to-left reranking increased translation quality for all systems. For the En → Pl translation direction, the best results were achieved when using ensembles of three models and better results were achieved by the unconstrained systems. However, for the Pl → En translation direction, the unconstrained systems achieved lower results than the constrained systems. The best results were achieved by the Transformer big model; ensembling did not improve results.

In overall, the results differ from what we have observed in previous years. Back-translation for Pl → En did not improve results, which raises a question of a possible domain mismatch between the monolingual data we back-translated and the development data. Unconstrained systems are only slightly better than constrained systems for En → Pl and even subpar for the Pl → En translation direction, which shows that current NMT methods are not able to benefit from larger datasets. Hence, having in-domain data is more important.

6 Conclusion

In this paper, we described Tilde’s NMT systems for the WMT shared task on news translation. This year, we trained constrained and unconstrained systems for the English↔Polish language pair. We de-

tailed the methods applied and the training recipes.

During our experiments, we identified several avenues of possible further research. We saw that larger datasets even after applying data selection methods did not allow improving translation quality (at least not significantly). We made a similar observation also previous years when participating in WMT. We saw in our results also that back-translation did not yield positive results for En → Pl. We hypothesise that there may be a domain mismatch between the data we used for training and the newsdev2020 dataset. However, this requires further investigation.

Acknowledgements

The research has been supported by the European Regional Development Fund within the research project “Multilingual Artificial Intelligence Based Human Computer Interaction” No. 1.1.1.1/18/A/148. We thank the High Performance Computing Center of Riga Technical University for providing access to their GPU computing infrastructure.

References

- Ahmed Abdelali, Francisco Guzman, Hassan Sajjad, and Stephan Vogel. 2014. The amara corpus: Building parallel language resources for the educational domain. In *LREC*, volume 14, pages 1044–1054.
- Marta Banón, Pinzhen Chen, Barry Haddow, Kenneth Heafield, Hieu Hoang, Miquel Espla-Gomis, Mikel L Forcada, Amir Kamran, Faheem Kirefu, Philipp Koehn, et al. 2020. Paracrawl: Web-scale acquisition of parallel corpora. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4555–4567.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. [Understanding Back-Translation at Scale](#). *arXiv:1808.09381 [cs]*. ArXiv: 1808.09381.
- Najeh Hajlaoui, David Kolovratnik, Jaakko Väyrynen, Ralf Steinberger, and Dániel Varga. 2014. Dcep-digital corpus of the european parliament. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*.
- Marcin Junczys-Dowmunt. 2018. Dual conditional cross-entropy filtering of noisy parallel corpora. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 888–895.
- Marcin Junczys-Dowmunt. 2019. Microsoft translator at wmt 2019: Towards large-scale document-level neural machine translation. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 225–233.

- Marcin Junczys-Dowmunt, Roman Grundkiewicz, Tomasz Dwojak, Hieu Hoang, Kenneth Heafield, Tom Neckermann, Frank Seide, Ulrich Germann, Alham Fikri Aji, Nikolay Bogoychev, André F. T. Martins, and Alexandra Birch. 2018. **Marian: Fast Neural Machine Translation in C++**. In *Proceedings of ACL 2018, System Demonstrations*, pages 116–121, Melbourne, Australia. Association for Computational Linguistics.
- Philipp Koehn. 2005. **Europarl : A Parallel Corpus for Statistical Machine Translation**. In *Proceedings of the 10th Machine Translation Summit (MT Summit)*, pages 79–86.
- Jerry Ma and Denis Yarats. 2018. Quasi-Hyperbolic Momentum and Adam for Deep Learning. *arXiv preprint arXiv:1810.06801*.
- Makoto Morishita, Jun Suzuki, and Masaaki Nagata. 2018. **NTT’s Neural Machine Translation Systems for WMT 2018**. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 461–466, Belgium, Brussels. Association for Computational Linguistics.
- Nathan Ng, Kyra Yee, Alexei Baevski, Myle Ott, Michael Auli, and Sergey Edunov. 2019. **Facebook FAIR’s WMT19 News Translation Task Submission**. *arXiv:1907.06616 [cs]*. ArXiv: 1907.06616.
- Mārcis Pinnis. 2018. **Tilde’s Parallel Corpus Filtering Methods for WMT 2018**. In *Proceedings of the Third Conference on Machine Translation*, pages 952–958, Brussels, Belgium. Association for Computational Linguistics.
- Mārcis Pinnis, Rihards Krišlauks, Daiga Deksnē, and Toms Miks. 2017a. **Neural Machine Translation for Morphologically Rich Languages with Improved Sub-word Units and Synthetic Data**. In *Proceedings of the 20th International Conference of Text, Speech and Dialogue (TSD2017)*, volume 10415 LNAI, Prague, Czechia.
- Mārcis Pinnis, Rihards Krišlauks, Toms Miks, Daiga Deksnē, and Valters Šics. 2017b. **Tilde’s Machine Translation Systems for WMT 2017**. In *Proceedings of the Second Conference on Machine Translation (WMT 2017), Volume 2: Shared Task Papers*, pages 374–381, Copenhagen, Denmark. Association for Computational Linguistics.
- Mārcis Pinnis, Rihards Krišlauks, and Matīss Rikters. 2019. **Tilde’s Machine Translation Systems for WMT 2019**. In *Proceedings of the Fourth Conference on Machine Translation*, pages 526–533, Florence, Italy. Association for Computational Linguistics.
- Mārcis Pinnis, Matīss Rikters, and Rihards Krišlauks. 2018. **Tilde’s Machine Translation Systems for WMT 2018**. In *Proceedings of the Third Conference on Machine Translation*, pages 477–485, Brussels, Belgium. Association for Computational Linguistics.
- Roberts Rozis and Raivis Skadiņš. 2017. **Tilde MODEL - Multilingual Open Data for EU Languages**. In *Proceedings of the 21st Nordic Conference on Computational Linguistics*, pages 263–265.
- Holger Schwenk, Vishrav Chaudhary, Shuo Sun, Hongyu Gong, and Francisco Guzmán. 2019. Wikimatrix: Mining 135m parallel sentences in 1620 language pairs from wikipedia. *arXiv*, pages arXiv-1907.
- Ralf Steinberger, Andreas Eisele, Szymon Kłoczek, Spyridon Pilos, and Patrick Schlüter. 2012. Dgtm: A freely available translation memory in 22 languages. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 454–459.
- Ralf Steinberger, Bruno Pouliquen, Anna Widiger, Camelia Ignat, Tomaz Erjavec, Dan Tufi, and Dániel Varga. 2006. The JRC-Acquis: A Multilingual Aligned Parallel Corpus with 20+ Languages. In *Proceedings of the 5th International Conference on Language Resources and Evaluation (LREC’2006)*, volume 4, pages 2142—2147.
- Ole Tange. 2011. **Gnu parallel - the command-line power tool**. *The USENIX Magazine*, 36(1):42–47.
- Jörg Tiedemann. 2016. Finding alternative translations in a large corpus of movie subtitle. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 3518–3522.
- 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 *Advances in Neural Information Processing Systems*, pages 5998–6008.