
Paying Attention to Multi-Word Expressions in Neural Machine Translation

Matīss Rikters
Faculty of Computing, University of Latvia

matiss@lielakeda.lv

Ondřej Bojar
Charles University, Faculty of Mathematics and Physics,
Institute of Formal and Applied Linguistics

bojar@ufal.mff.cuni.cz

Abstract

Processing of multi-word expressions (MWEs) is a known problem for any natural language processing task. Even neural machine translation (NMT) struggles to overcome it. This paper presents results of experiments on investigating NMT attention allocation to the MWEs and improving automated translation of sentences that contain MWEs in English→Latvian and English→Czech NMT systems. Two improvement strategies were explored—(1) bilingual pairs of automatically extracted MWE candidates were added to the parallel corpus used to train the NMT system, and (2) full sentences containing the automatically extracted MWE candidates were added to the parallel corpus. Both approaches allowed to increase automated evaluation results. The best result—0.99 BLEU point increase—has been reached with the first approach, while with the second approach minimal improvements achieved. We also provide open-source software and tools used for MWE extraction and alignment inspection.

1 Introduction

It is well known that neural machine translation (NMT) has defined the new state of the art in the last few years (Sennrich et al., 2016a; Wu et al., 2016), but the many specific aspects of NMT outputs are not yet explored. One of which is translation of multi-word units or multi-word expressions (MWEs). MWEs are defined by Baldwin and Kim (2010) as “lexical items that: (a) can be decomposed into multiple lexemes; and (b) display lexical, syntactic, semantic, pragmatic and/or statistical idiomaticity”. MWEs have been a challenge for statistical machine translation (SMT). Even if standard phrase-based models can copy MWEs verbatim, they suffer in grammaticality. NMT, on the other hand, may struggle in memorizing and reproducing MWEs, because it represents the whole sentence in a high-dimensional vector, which can lose the specific meanings of the MWEs even in the more fine-grained attention model (Bahdanau et al., 2015), because MWEs may not appear frequently enough in the training data.

The goal of this research is to examine how MWEs are treated by NMT systems, compare that with related work in SMT, and find ways to improve MWE translation in NMT. We aimed to compare how NMT pays attention to MWEs during translation, using a test set particularly targeted at handling of MWEs, and if that can be improved by populating the training data for the NMT systems with parallel corpora of MWEs.

The objective was to obtain a comparison of how NMT with regular training data and NMT with synthetic MWE data pays attention to MWEs during the translation process as well

as to improve the final NMT output. To achieve this objective, it needed to be broken down into smaller sub-objectives:

- Train baseline NMT systems,
- Extract parallel MWE corpora from the training data,
- Train the NMT systems with synthetic MWE data, and
- Inspect alignments produced by the NMT.

The structure of this paper is as follows: Section 2 summarizes related work in translating MWEs with SMT and NMT. Section 3 describes the architecture of the baseline system and outlines the process of extracting parallel MWE corpora from the training data. Section 4 provides the experiment setup and results. Finally, conclusions and aims for further directions of work are summarized in Section 5.

2 Related Work

There have been several experiments with incorporating separate processing of MWEs in rule-based (Deksne et al., 2008) and statistical machine translation tasks (Bouamor et al., 2012; Skadiņa, 2016). However, there is little literature about similar integrations in NMT workflows so far.

Skadiņa (2016) performed a series of experiments on extracting MWE candidates and integrating them in SMT. The author experimented with several different methods for both the extraction of MWEs and integration of the extracted MWEs into the MT system. In terms of automatic MT evaluation, this allowed to achieve an increase of 0.5 BLEU points (Papineni et al., 2002) for an English→Latvian SMT system.

Tang et al. (2016) introduce an NMT approach that uses a stored phrase memory in symbolic form. The main difference from traditional NMT is tagging candidate phrases in the representation of the source sentence and forcing the decoder to generate multiple words all at once for the target phrase. Although they do mention MWEs, no identification or extraction of MWEs is performed and the phrases they mainly focus on are dates, names, numbers, locations, and organizations, that are collected from multiple dictionaries. For Chinese→English they report a 3.45 BLEU point increase over baseline NMT.

Cohn et al. (2016) describe an extension of the traditional attentional NMT model with the inclusion of structural biases from word-based alignment models, such as positional bias, Markov conditioning, fertility and agreement over translation directions. They perform experiments translating between English, Romanian, Estonian, Russian and Chinese and analyze the attention matrices of the output translations produced by running experiments using the different biases. Specific experiments targeting MWEs are not performed, but they do point out that using fertility, especially global fertility, can be useful for dealing with multi-word expressions. They report a statistically significant improvement of BLEU scores in almost all involved language pairs.

Chen et al. (2016) use a similar approach as we do. Their “bootstrapping” automatically extracts smaller parts of training segment pairs and adds them to the training data for NMT. The main difference is that they rely on automatic word alignment and punctuation in the sentence to identify matching sub-segments.

3 Data Preparation and Systems Used

To measure changes introduced by adding synthetic MWE data to the training corpora, first, a baseline NMT system was trained for each language pair. The experiments were conducted on English→Czech and English→Latvian translation directions.



Figure 1: Portions of the final training data set for English→Czech



Figure 2: Portions of the final training data set for English→Latvian

3.1 Baseline NMT System

To be able to compare the results with other MT systems, training and development corpora were used from the WMT shared tasks: data from the News Translation Task¹ for English→Latvian and data from the Neural MT Training Task² (Bojar et al., 2017) for English→Czech. The English→Czech data consists of about 49 million parallel sentence pairs and the English→Latvian of about 4.5 million. The development corpora consist of 2003 sentences for English→Latvian and 6000 for English→Czech.

Neural Monkey (Helcl and Libovický, 2017), an open-source tool for sequence learning, was used to train the baseline NMT systems. Using the configuration provided by the WMT Neural MT Training Task organizers, the baseline reached 11.29 BLEU points for English→Latvian after having seen 23 million sentences in about 5 days and 13.71 BLEU points for English→Czech after having seen 18 million sentences in about 7 days.

3.2 Extraction of Parallel MWEs

To extract MWEs, the corpora were first tagged with morphological taggers: UDPipe (Ramisch, 2012) for English and Czech, LV Tagger (Paikens et al., 2013) for Latvian. After that, the tagged corpora were processed with the Multi-word Expressions toolkit (Ramisch, 2012), and finally aligned with the MPAligner (Pinnis, 2013), intermittently pre-processing and post-processing with a set of custom tools. To extract MWEs from the corpora with the MWE Toolkit, patterns were required for each of the involved languages. Patterns from Skadiņa (2016) were used for Latvian (210 patterns) and English (57 patterns) languages and patterns from Majchráková et al. (2012) and Pecina (2008) for Czech (23 patterns).

This workflow allowed to extract a parallel corpus of about 400 000 multi-word expressions for English→Czech and about 60 000 for English→Latvian. For an extension of this experiment, all sentences containing these MWEs were also extracted from the training corpus, serving as a separate parallel corpus.

4 Experiments

We experiment with two forms of the presentation of MWEs to the NMT system: (1) we add only the parallel MWEs themselves, each pair forming a new “sentence pair” in the parallel corpus, and (2) we use full sentences containing the MWEs. We denote the approaches “MWE phrases” and “MWE sents.” in the following.

4.1 Training Corpus Layout

In both cases, we use the same corpus training corpus layout: we mix the baseline parallel corpus with synthetic data so that MWEs get more exposure to the neural network in training

¹<http://www.statmt.org/wmt17/translation-task.html>

²<http://www.statmt.org/wmt17/nmt-training-task/>

Languages	En→Cs		En→Lv	
Dataset	Dev	MWE	Dev	MWE
Baseline	13.71	10.25	11.29	9.32
+MWE phrases	-	-	11.94	10.31
+MWE sents.	13.99	10.44	-	-

Table 1: Experiment results.

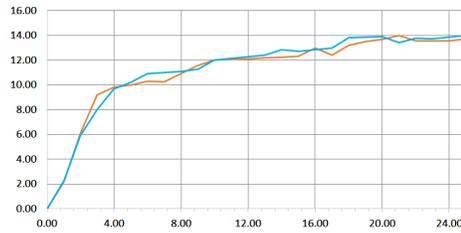


Figure 3: Automatic evaluation progression of En→Cs experiments on validation data. Orange – baseline; blue — baseline with added MWEs.

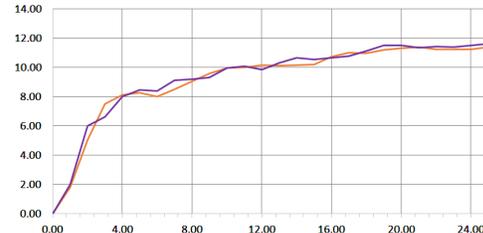


Figure 4: Automatic evaluation progression of En→Lv experiments on validation data. Orange – baseline; purple — baseline with added MWE sentences.

and hopefully allow NMT to learn to translate them better.

Figure 1 and Figure 2 illustrate how the training data was divided into portions. The block 1xMWE corresponds to the full set of extracted MWEs (400K for En→Cs, 60K for En→Lv) and 2xMWE corresponds to two copies of the set (800K for En→Cs, 120K for En→Lv). For En→Lv the full corpus was used. For En→Cs we used only the first 15M sentences to be able to train multiple epochs on the available hardware. The MWEs get repeated five times in both language pairs. By doing this, the En→Cs data set was reduced from 49M to 17M and the En→Lv data set increased to 4.8M parallel sentences for one epoch of training.

While the experiments were running, early stopping of the training was executed and snapshots of the models for evaluation were taken in stages where the models already were starting to converge. For En→Lv this was after the networks had been trained on 25M sentences (i.e. 5.2 epochs of the mixed corpus), for En→Cs 27M sentences (i.e. 1.6 epochs).

Neural Monkey does not shuffle the training corpus between epochs. This is not a problem if the corpus is properly shuffled and the number of epochs is not very large compared to the size of the epochs. We shuffled only the baseline corpus and the interleaved it with (shuffled) sections for MWEs. This worked well when MWEs were provided in full sentences, but not with MWEs presented as expressions. In the latter case, the NMT started to produce only very short output, losing very much of its performance. We, therefore, shuffle the whole composed corpus for the “MWE phrases” runs, effectively discarding the interleaved composition of the training data.

4.2 Results

Table 1 shows the results for both approaches and both language pairs. Due to hardware constraints, we were not able to try out both approaches on both language pairs.

We evaluate all setups with BLEU (Papineni et al., 2002) on the full development set (distinct from the training set), as shown in the column “Dev”, and on a subset of 611 (En→Lv) and 112 (En→Cs) sentences containing the identified MWEs (column “MWE”).

Sentence	BLEU	Length ratio	Text
Source	-	-	Just like in a city bus or a tram .
Target	100.00	1.00	Stejně jako v městském autobuse či tramvaji .
Baseline	13.54	0.88	jako ve městě autobuse nebo tramvaji .
Improved NMT	41.11	1.00	jen jako v městském autobuse nebo tramvaji .

Figure 5: Differences in translation between baseline and improved NMT system. Improving n-grams are highlighted in green and worsening n-grams — in red.

Sentence	BLEU	Length ratio	Text
Source	-	-	He steps toward the electronic wall map depicting Australia and the surrounding sea areas .
Target	100.00	1.00	Přistoupí k nástěnné elektronické mapě , na níž je znázorněna Austrálie a přilehlé mořské oblasti .
Baseline	10.74	0.75	ukazuje na mapu elektronické stěny zobrazuje Austrálii a okolní mořské oblasti .
Improved NMT	13.87	0.81	vykročil směrem k elektronické mapě , které depicting Austrálie a okolní oblasti .

Figure 6: Differences in translation of a Czech sentence using baseline and improved NMT systems. Improving n-grams are highlighted in green and worsening n-grams — in red.

Figures 3 and 4 illustrate the learning curves in terms of millions of sentences, as evaluated on the full development set.

We see that the difference on the whole development set is not very big for either of the languages, and that it fluctuates as the training progresses.

The improvement is more apparent when evaluated on the dedicated devset of sentences containing multi-word expressions. The improvement for Latvian is even 0.99 BLEU, but arguably, the baseline performance of our system is not very high. Also, more runs should be carried out for a full confidence, but this was unfortunately out of our limits on computing resources.

4.3 Manual Inspection

To find out whether changes in the results are due to the synthetic MWE corpora added, a subset of output sentences from the ones containing MWEs were selected for closer examination. For this task, we used the iBLEU (Madnani, 2011) tool.

In Figure 5, an improvement in the modified NMT translation is visible due to the treatment of the compound nominal “city bus” as a single expression. It seems that the baseline system translates “city” into “městě” and “bus” into “autobuse” individually, resulting in the wrong form of “city” in Czech (a noun used instead of an adjective). On the other hand, the improved NMT translates “city” into “městském” just like the target human translation. Attention alignments will be examined in the following section.

Figure 6 shows an example where the improved NMT scores higher in BLEU points and translates the MWE closer to the human, but loses a part of it in the process. While translating the noun phrase “electronic wall map” the improved system generates a closer match to the human translation “elektronické mapě”, it does not translate the word “wall” that was translated into “stěny” by the baseline system. Upon closer inspection, we discovered that this error was caused by the MWE extractor and aligner because the identified English phrase “electronic wall map” was aligned to an identified Czech phrase “elektronické mapě” and the whole phrase

Sentence	BLEU	Length ratio	Text
Source	-	-	It should be noted that this is not the first time that Facebook has been actively involved in determining what network users see in their news feeds.
Target	100.00	1.00	Jāteic, ka šī nav pirmā reize, kad Facebook aktīvi iesaistās, nosakot, ko tīkla lietotāji redz savās jaunumu plūsmās.
Baseline	10.41	1.04	Jāatzīmē, ka šis nav pirmajā reizē, kad Facebook ir aktīvi iesaistīta, nosakot to, ko tīklā izmanto viņu ziņu pārraides.
Improved NMT	27.41	1.13	Ir jāatzīmē, ka šis ir pirmā reize, kad Facebook ir aktīvi iesaistījies, nosakot to, ko tīkla lietotāji dara viņu ziņu formātā.

- Source:** It should be noted that this is not the first time that Facebook has been actively involved in determining what network users see in their news feeds.
- Baseline:** Jāatzīmē, ka šis nav pirmajā reizē, kad Facebook ir aktīvi iesaistīta, nosakot to, ko tīklā izmanto viņu ziņu pārraides.
- Improved NMT:** Ir jāatzīmē, ka šis ir pirmā reize, kad Facebook ir aktīvi iesaistījies, nosakot to, ko tīkla lietotāji dara viņu ziņu formātā.
- Reference:** Jāteic, ka šī nav pirmā reize, kad Facebook aktīvi iesaistās, nosakot, ko tīkla lietotāji redz savās jaunumu plūsmās.

Figure 7: Differences in translation between baseline and improved NMT system. Improving n-grams are highlighted in green and worsening n-grams — in red.

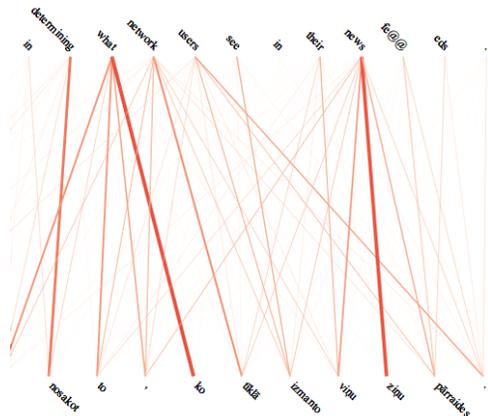


Figure 9: Fragment of soft alignments of the example sentence from the baseline NMT system.

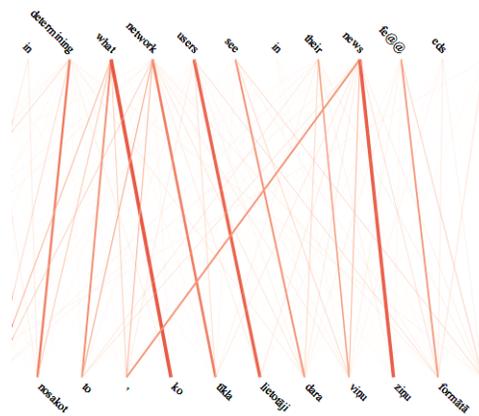


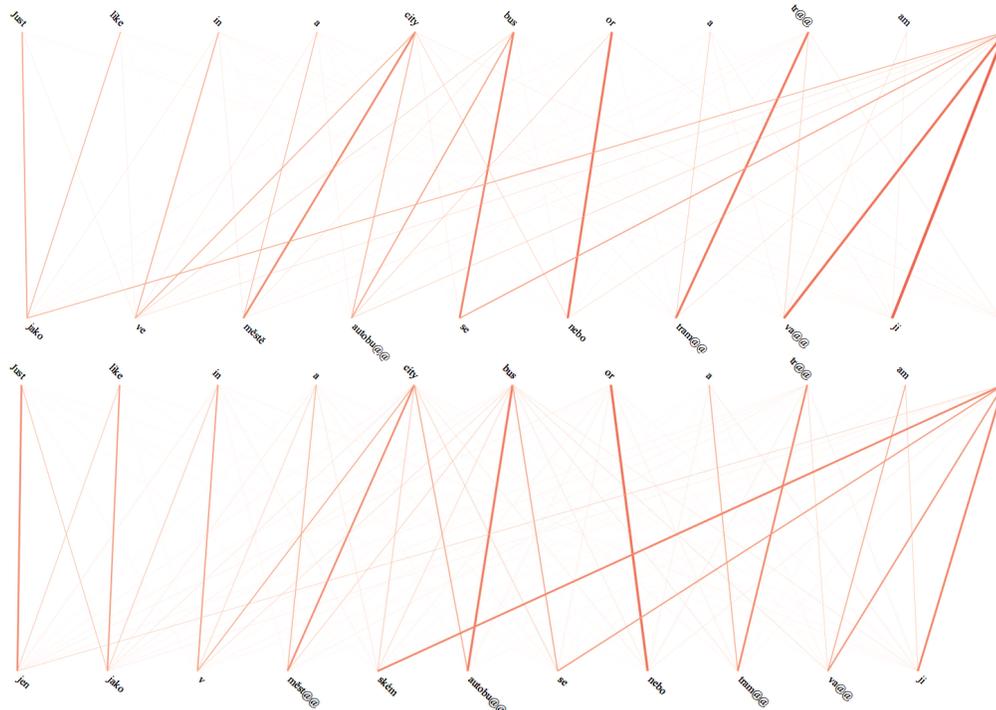
Figure 10: Fragment of soft alignments of the example sentence from the improved NMT system.

“nástěnné elektronické mapě” was not identified by the MWE extractor at all.

Figure 7 illustrates translations of an example sentence by the En→Lv NMT systems. The MWE, in this case, is “network users” that is translated as “tīkla lietotāji” by the modified system and completely mistranslated by the baseline.

4.4 Alignment Inspection

For inspecting the NMT attention alignments, we developed a tool (Riktors et al., 2017) that takes data produced by Neural Monkey—a 3D array (tensor) filled with the alignment probabilities together with source and target subword units (Sennrich et al., 2016b) or byte pair encodings (BPEs)—as input and produces a soft alignment matrix (Figure 8) of the subword units that highlights all units, that get attention when translating a specific subword unit. The tool includes a web version that was adapted from Nematus (Sennrich et al., 2017) utilities and slightly modified. It allows to output the soft alignments in a different perspective, as connections between BPEs as visible in Figure 9 and Figure 10.



Source: Just like in a city bus or a tram.
Baseline: Jako ve městě autobuse nebo tramvaji.
Improved NMT: Jen jako v městském autobuse nebo tramvaji.
Reference: Stejně jako v městském autobuse či tramvaji.

Figure 11: Soft alignment example visualizations from translating an English sentence into Czech from the baseline (top, hypothesis 1) and improved (bottom, hypothesis 2) NMT systems.



Figure 8: Example of a soft alignment matrix.

In these examples, the attention state of the previously mentioned MWE from En→Lv translations (“network users”) is visible. The alignment inspection tool allows to see that the baseline NMT in Figure 9 has multiple faded alignment lines for both words “network” and “users”, which outlines that the neural network is unsure and looking all around for traces to the correct translation. However, in Figure 10, it is visible that both these words have strong alignment lines to the words “tīkla lietotāji”, that were also identified by the MWE Toolkit as an MWE candidate.

Figure 11 shows one of the previously mentioned En→Cs translation examples. Here it is clear that in the baseline alignment no attention goes to the word “městě” or the subword units “autobu@@” and “se” when translating “city”. In the modified version, on the other hand, some attention from “city” goes into all closely related subword units: “měst@@”, “ském”, “autobu@@”, and “se”. It is also visible that in this example, the translation of “bus” gets attention from not only “autobu@@” and “se” but also the ending subword unit of “city”, i.e. the token “ském”.

5 Conclusion

In this paper, we described the first experiments with handling multi-word expressions in neural machine translation systems. Details on identifying and extracting MWEs from parallel corpora, as well as aligning them and building corpora of parallel MWEs were provided. We explored two methods of integrating MWEs in training data for NMT and examined the output translations of the trained NMT systems with custom built tools for alignment inspection.

In addition to the methods described in this paper, we also released open-source scripts for a complete workflow of identifying, extracting and integrating MWEs into the NMT training and translation workflow.

While the experiments did not show outstanding improvements on the general development data set, an increase of 0.99 BLEU was observed when using an MWE specific test data set. Manual inspection of the output translations confirmed that translations of specific MWEs were improving after populating the training data with synthetic MWE data.

As the next steps, we plan (1) to analyze the obtained results of our experiments in more detail through the help of a larger scale manual human evaluation of the NMT output and (2) to continue experiments to find best ways how to treat different categories of MWEs, i.e. idioms.

Acknowledgement

This study was supported in parts by the grants H2020-ICT-2014-1-645442 (QT21), the ICT COST Action IC1207 *ParseME: Parsing and multi-word expressions. Towards linguistic precision and computational efficiency in natural language processing*, and Charles University Research Programme “Progres” Q18 – Social Sciences: From Multidisciplinarity to Interdisciplinarity.

References

- Bahdanau, D., Cho, K., and Bengio, Y. (2015). Neural Machine Translation by Jointly Learning to Align and Translate. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Baldwin, T. and Kim, S. N. (2010). Multiword expressions. In *Handbook of Natural Language Processing, Second Edition*, pages 267–292. Chapman and Hall/CRC.
- Bojar, O., Helcl, J., Kocmi, T., Libovický, J., and Musil, T. (2017). Results of the WMT17

Neural MT Training Task. In *Proceedings of the Second Conference on Machine Translation (WMT17)*, Copenhagen, Denmark.

Bouamor, D., Semmar, N., and Zweigenbaum, P. (2012). Identifying bilingual multi-word expressions for statistical machine translation. In *LREC*, pages 674–679.

Chen, W., Matusov, E., Khadivi, S., and Peter, J.-T. (2016). Guided alignment training for topic-aware neural machine translation. *AMTA 2016, Vol.*, page 121.

Cohn, T., Hoang, C. D. V., Vymolova, E., Yao, K., Dyer, C., and Haffari, G. (2016). Incorporating structural alignment biases into an attentional neural translation model. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 876–885, San Diego, California. Association for Computational Linguistics.

Deksne, D., Skadins, R., and Skadina, I. (2008). Dictionary of multiword expressions for translation into highly inflected languages. In *LREC*.

Helcl, J. and Libovický, J. (2017). Neural Monkey: An open-source tool for sequence learning. *The Prague Bulletin of Mathematical Linguistics*, (107):5–17.

Madnani, N. (2011). ibleu: Interactively debugging and scoring statistical machine translation systems. In *Semantic Computing (ICSC), 2011 Fifth IEEE International Conference on*, pages 213–214. IEEE.

Majchráková, D., Dušek, O., Hajič, J., Karčová, A., and Garabík, R. (2012). Semi-automatic detection of multiword expressions in the slovak dependency treebank.

Paikens, P., Rituma, L., and Pretkalnina, L. (2013). Morphological analysis with limited resources: Latvian example. In *Proceedings of the 19th Nordic Conference of Computational Linguistics (NODALIDA 2013); May 22-24; 2013; Oslo University; Norway. NEALT Proceedings Series 16*, number 085, pages 267–277. Linköping University Electronic Press.

Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. (2002). Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.

Pecina, P. (2008). Reference data for czech collocation extraction. In *Proc. of the LREC Workshop Towards a Shared Task for MWEs (MWE 2008)*, pages 11–14.

Pinnis, M. (2013). Context independent term mapper for european languages. In *RANLP*, pages 562–570.

Ramisch, C. (2012). A generic framework for multiword expressions treatment: from acquisition to applications. In *Proceedings of ACL 2012 Student Research Workshop*, pages 61–66. Association for Computational Linguistics.

Riktters, M., Fishel, M., and Bojar, O. (2017). Visualizing Neural Machine Translation Attention and Confidence. *The Prague Bulletin of Mathematical Linguistics*, 109.

Sennrich, R., Firat, O., Cho, K., Birch, A., Haddow, B., Hirschler, J., Junczys-Dowmunt, M., Läubli, S., Miceli Barone, A. V., Mokry, J., and Nadejde, M. (2017). Nematus: a toolkit for neural machine translation. In *Proceedings of the Software Demonstrations of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, pages 65–68, Valencia, Spain. Association for Computational Linguistics.

- Sennrich, R., Haddow, B., and Birch, A. (2016a). Edinburgh neural machine translation systems for wmt 16. In *Proceedings of the First Conference on Machine Translation*, pages 371–376, Berlin, Germany. Association for Computational Linguistics.
- Sennrich, R., Haddow, B., and Birch, A. (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.
- Skadiņa, I. (2016). Multi-word expressions in english-latvian. In *Human Language Technologies–The Baltic Perspective: Proceedings of the Seventh International Conference Baltic HLT 2016*, volume 289, page 97. IOS Press.
- Tang, Y., Meng, F., Lu, Z., Li, H., and Yu, P. L. H. (2016). Neural machine translation with external phrase memory. *CoRR*, abs/1606.01792.
- Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., et al. (2016). Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.