

A Multifaceted Evaluation of Neural versus Phrase-Based Machine Translation for 9 Language Directions

Antonio Toral*
University of Groningen
The Netherlands
a.toral.ruiz@rug.nl

Víctor M. Sánchez-Cartagena
Prompsit Language Engineering
Av. Universitat s/n. Edifici Quorum III
E-03202 Elx, Spain
vmsanchez@prompsit.com

Abstract

We aim to shed light on the strengths and weaknesses of the newly introduced neural machine translation paradigm. To that end, we conduct a multifaceted evaluation in which we compare outputs produced by state-of-the-art neural machine translation and phrase-based machine translation systems for 9 language directions across a number of dimensions. Specifically, we measure the similarity of the outputs, their fluency and amount of reordering, the effect of sentence length and performance across different error categories. We find out that translations produced by neural machine translation systems are considerably different, more fluent and more accurate in terms of word order compared to those produced by phrase-based systems. Neural machine translation systems are also more accurate at producing inflected forms, but they perform poorly when translating very long sentences.

1 Introduction

A new paradigm to statistical machine translation, neural MT (NMT), has emerged very recently and has already surpassed the performance of the mainstream approach in the field, phrase-based MT (PBMT), for a number of language pairs, e.g. (Sennrich et al., 2016b; Luong et al., 2015; Costa-Jussà and Fonollosa, 2016; Chung et al., 2016).

In PBMT (Koehn, 2010) different models (translation, reordering, target language, etc.) are trained independently and combined in a log-linear scheme in which each model is assigned a

different weight by a tuning algorithm. On the contrary, in NMT all the components are jointly trained to maximise translation quality. NMT systems have a strong generalisation power because they encode translation units as numeric vectors that represent concepts, whereas in PBMT translation units are encoded as strings. Moreover, NMT systems are able to model long-distance phenomena thanks to the use of recurrent neural networks, e.g. long short-term memory (LSTM) (Hochreiter and Schmidhuber, 1997) or gated recurrent units (Chung et al., 2014).

The translations produced by NMT systems have been evaluated thus far mostly in terms of overall performance scores, be it by means of automatic or human evaluations. This has been the case of last year's news translation shared task at the First Conference on Machine Translation (WMT16).¹ In this translation task, outputs produced by participant MT systems, the vast majority of which fall under either the phrase-based or neural approaches, were evaluated (i) automatically with the BLEU (Papineni et al., 2002) and TER (Snover et al., 2006) metrics, and (ii) manually by means of ranking translations (Federmann, 2012) and monolingual semantic similarity (Graham et al., 2016). In all these evaluations, the performance of each system is measured by means of an overall score, which, while giving an indication of the general performance of a given system, does not provide any additional information.

In order to understand better the new NMT paradigm and in what respects it provides better (or worse) translation quality than state-of-the-art PBMT, Bentivogli et al. (2016) conducted a detailed analysis for the English-to-German language direction. In a nutshell, they found out that NMT (i) decreases post-editing effort, (ii) de-

*Work partly done at his previous position in Dublin City University, Ireland.

¹<http://www.statmt.org/wmt16/translation-task.html>

grades faster than PBMT with sentence length and (iii) results in a notable improvement regarding re-ordering.

In this paper we delve further in this direction by conducting a multilingual and multifaceted evaluation in order to find answers to the following research questions. Whether, in comparison to PBMT, NMT systems result in:

- considerably different output and higher degree of variability;
- more or less fluent output;
- more or less monotone translations;
- translations with better or worse word order;
- better or worse translations depending on sentence length;
- less or more errors for different error categories: inflectional, reordering and lexical;

Hereunder we specify the main differences and similarities between this work and that of Bentivogli et al. (2016):

- Language directions. They considered 1 while our study comprises 9.
- Content. They dealt with transcribed speeches while we work with news stories. Previous research has shown that these two types of content pose different challenges for MT (Ruiz and Federico, 2014).
- Size of evaluation data. Their test set had 600 sentences while our test sets span from 1 999 to 3 000 depending on the language direction.
- Reference type. Their references were both independent from the MT output and also post-edited, while we have access only to single independent references.
- Analyses. While some analyses overlap, some are novel in our experiments. Namely, output similarity, fluency and degree of re-ordering performed.

Our analyses are conducted on the best PBMT and NMT systems submitted to the WMT16 translation task for each language direction. This (i) guarantees the reproducibility of our results as all the MT outputs are publicly available, (ii) ensures

that the systems evaluated are state-of-the-art, as they are the result of the latest developments at top MT research groups worldwide, and (iii) allows the conclusions that will be drawn to be rather general, as 6 languages from 4 different families (Germanic, Slavic, Romance and Finno-Ugric) are covered in the experiments.

The rest of the paper is organised as follows. Section 2 describes the experimental setup. Subsequent sections cover the experiments carried out in which we measured different aspects of NMT, namely: output similarity (Section 3), fluency (Section 4), degree of reordering and quality of word order (Section 5), sentence length (Section 6), and amount of errors for different error categories (Section 7). Finally, Section 8 holds the conclusions and proposals for future work.

2 Experimental Setup

The experiments are run on the best PBMT and NMT constrained systems submitted to the news translation task of WMT16. We selected such systems according to the human evaluation (Bojar et al., 2016, Sec. 3.4).² We noted that many of the PBMT systems contain neural features, mainly in the form of language models. If the best PBMT submission contains any neural features we use this as the PBMT system in our analyses as long as none of these features is a fully-fledged NMT system. This was the case of the best submission in terms of BLEU for RU→EN (Junczys-Dowmunt et al., 2016).

Out of the 12 language directions at the translation task, we conduct experiments on 9.³ These are the language pairs between English (EN) and Czech (CS), German (DE), Finnish (FI), Romanian (RO) and Russian (RU) in both directions (except for Finnish, where only the EN→FI direction is covered as no NMT system was submitted for the opposite direction, FI→EN). Finally, there was an additional language at the shared task, Turkish, that is not considered here, as either none of the systems submitted was neural (Turkish→EN), or there was one such system but its performance

²When there are not statistically significant differences between two or more NMT or PBMT systems (i.e. they belong to the same equivalence class), we pick the one with the highest BLEU score. If two NMT or PBMT systems were the best according to BLEU (draw), we pick the one with the best TER score.

³Some experiments are run on a subset of these languages due to the lack of required tools for some of the languages involved.

Language Pair	MT Paradigm	System details
EN→CS	PBMT NMT	Phrase-based, word clusters (Ding et al., 2016) Unsupervised word segmentation and backtranslated monolingual corpora (Sennrich et al., 2016a)
EN→DE	hierarchical PBMT NMT	String-to-tree, neural and dependency language models (Williams et al., 2016) Same as for EN→CS
EN→FI	PBMT NMT	Phrase-based, rule-based and unsupervised word segmentation, operation sequence model (Durrani et al., 2011), bilingual neural language model (Devlin et al., 2014), re-ranked with a recurrent neural language model (Sánchez-Cartagena and Toral, 2016) Rule-based word segmentation, backtranslated monolingual corpora (Sánchez-Cartagena and Toral, 2016)
EN→RO	PBMT NMT	Phrased-based, operation sequence model, monolingual and bilingual neural language models (Williams et al., 2016) Same as for EN→CS
EN→RU	PBMT NMT	Phrase-based, word clusters, bilingual neural language model (Ding et al., 2016) Same as for EN→CS
CS→EN	PBMT NMT	Same as for EN→CS Same as for EN→CS
DE→EN	PBMT NMT	Phrase-based, pre-reordering, compound splitting (Williams et al., 2016) Same as for EN→CS plus reranked with a right-to-left model
RO→EN	PBMT NMT	Phrase-based, operation sequence model, monolingual neural language model (Williams et al., 2016) Same as for EN→CS
RU→EN	PBMT NMT	Phrase-based, lemmas in word alignment, sparse features, bilingual neural language model and transliteration (Lo et al., 2016) Same as for EN→CS

Table 1: Details of the best systems pertaining to the PBMT and NMT paradigms submitted to the WMT16 news translation task for each language direction.

was extremely low (EN→Turkish) and hence most probably not representative of the state-of-the-art in NMT.

Table 1 shows the main characteristics of the best PBMT and NMT systems submitted to the WMT16 news translation task. It should be noted that all the NMT systems listed in the table fall under the encoder-decoder architecture with attention (Bahdanau et al., 2015) and operate on subword units. Word segmentation is carried out with the help of a lexicon in the EN→FI direction (Sánchez-Cartagena and Toral, 2016) and in an unsupervised way in the remaining directions (Sennrich et al., 2016a).

2.1 Overall Evaluation

First, and in order to contextualise our analyses below, we report the BLEU scores achieved by the best NMT and PBMT systems for each language direction at WMT16’s news translation task in Table 2.⁴ The best NMT system clearly outperforms the best PBMT system for all language directions out of English (relative improvements range from 5.5% for EN→RO to 17.6% for EN→FI) and the human evaluation (Bojar et al., 2016, Sec. 3.4) confirms these results. In the opposite direction, the human evaluation shows that the best NMT system outperforms the best PBMT system for all language directions except when the source language is Russian. This slightly differs from the automatic evaluation, according to which NMT outperforms PBMT for translations from Czech (3.3% relative improvement) and German (9.9%) but underperforms PBMT for translations from Romanian (-3.7%) and Russian (-3.8%).

3 Output Similarity

The aim of this analysis is to assess to which extent translations produced by NMT systems are different from those produced by PBMT systems. We measure this by taking the outputs of the top n ⁵ NMT and PBMT systems submitted to each language direction and checking their pairwise overlap in terms of the chrF1 (Popović, 2015)

⁴We report the official results from <http://matrix.statmt.org/matrix> for the test set *newstest2016* using normalised BLEU (column z *BLEU-cased-norm*).

⁵The number of systems considered is different for each language direction as it depends on the number of systems submitted. Namely, we have considered 2 NMT and 2 PBMT into Czech, 3 NMT and 5 PBMT into German, 2 NMT and 4 PBMT into Finnish, 2 NMT and 4 PBMT into Romanian and 2 NMT and 3 PBMT into Russian.

System	CS	DE	FI	RO	RU
	From EN				
PBMT	23.7	30.6	15.3	27.4	24.3
NMT	25.9	34.2	18.0	28.9	26.0
	Into EN				
PBMT	30.4	35.2	23.7	35.4	29.3
NMT	31.4	38.7	-	34.1	28.2

Table 2: BLEU scores of the best NMT and PBMT systems for each language pair at WMT16’s news translation task. If the difference between them is statistically significant according to paired bootstrap resampling (Koehn, 2004) with $p = 0.05$ and 1 000 iterations, the highest score is shown in bold.

automatic evaluation metric.⁶ In order to make sure that all systems considered are truly different (rather than different runs of the same system) we consider only 1 system per paradigm (NMT and PBMT) submitted by each team for each language direction.

We would consider NMT outputs considerably different (with respect to PBMT) if they resemble each other (i.e. high pairwise overlap between NMT outputs) more than they do to PBMT systems (i.e. low overlap between an output by NMT and another by PBMT). This analysis is carried out only for language directions out of English, as for all the language directions into English there was, at most, 1 NMT submission.

TL	2 NMT	2 PBMT	NMT & PBMT
CS	68.66	77.63	64.34
DE	72.10	72.97	66.80
FI	56.03	57.42	55.55
RO	69.47	75.96	68.77
RU	35.52	43.35	29.87

Table 3: Average of the overlaps between pairs of outputs produced by the top n NMT and PBMT systems for each language direction from English to the target language (TL). The higher the value, the larger is the overlap.

Table 3 shows the results. We can observe the same trends for all the language directions, namely: (i) the highest overlaps are between pairs

⁶Throughout our analyses we use this metric as it has been shown to correlate better with human judgements than the *de facto* standard automatic metric, BLEU, when the target language is a morphologically rich language such as Finnish, while its correlation is on par with BLEU for languages with simpler morphology such as English (Popović, 2015).

of PBMT systems; (ii) next, we have overlaps between NMT systems; (iii) finally, overlaps between PBMT and NMT are the lowest.

We can conclude then that NMT systems lead to considerably different outputs compared to PBMT. The fact that there is higher inter-system variability in NMT than in PBMT (i.e. overlaps between pairs of NMT systems are lower than between pairs of PBMT systems) may surprise the reader, considering the fact that all NMT systems belong to the same paradigm (encoder-decoder with attention) while for some language directions (EN→DE, EN→FI and EN→RO) there are PBMT systems belonging to two different paradigms (pure phrase-based and hierarchical). However, the higher variability among NMT translations can be attributed, we believe, to the fact that NMT systems use numeric vectors that represent concepts instead of strings as translation units.

4 Fluency

In this experiment we aim to find out whether the outputs produced by NMT systems are more or less fluent than those produced by PBMT systems. To that end, we take perplexity of the MT outputs on neural language models (LMs) as a proxy for fluency. The LMs are built using TheanoLM (Enarvi and Kurimo, 2016). They contain 100 units in the projection layer, 300 units in the LSTM layer, and 300 units in the *tanh* layer, following the setup described by Enarvi and Kurimo (2016, Sec. 3.2). The training algorithm is Adagrad (Duchi et al., 2011) and we used 1 000 word classes obtained with `mkcls` from the training corpus. Vocabulary is limited to the most frequent 50 000 tokens.

LMs are trained on a random sample of 4 million sentences selected from the News Crawl 2015 monolingual corpora, available for all the languages considered.⁷

Table 4 shows the results. For all the language directions considered but one, perplexity is higher on the PBMT output compared to the NMT output. The only exception is translation into Finnish, in which perplexity on the PBMT output is slightly lower, probably because its fluency was improved by reranking it with a neural LM similar to the one

⁷<http://data.statmt.org/wmt16/translation-task/training-monolingual-news-crawl.tgz>

Language direction	PBMT	NMT	Rel. diff.
EN→CS	202.91	173.33	-14.58%
EN→DE	131.54	107.08	-18.60%
EN→FI	214.10	222.40	3.88%
EN→RO	124.66	116.33	-6.68%
EN→RU	158.18	127.83	-19.19%
CS→EN	110.08	102.36	-7.01%
DE→EN	122.26	104.72	-14.35%
RO→EN	106.08	102.18	-3.68%
RU→EN	123.86	106.75	-13.81%
Average	143.74	129.22	-10.45%

Table 4: Perplexity scores for the outputs of the best NMT and PBMT systems on language models built on 4 million sentences randomly selected from the News Crawl 2015 corpora.

we use in this experiment (Sánchez-Cartagena and Toral, 2016). The average relative difference, i.e. considering all language directions, is notable at -10.45%. Thus, our experiment shows that the outputs produced by NMT systems are, in general, more fluent than those produced by PBMT systems.

One may argue that the perplexity obtained for NMT outputs is lower than that for PBMT outputs because the LMs we used to measure perplexity follow the same model as the decoder of the NMT architecture (Bahdanau et al., 2015) and hence perplexity on a neural LM is not a valid proxy for fluency. However, the following facts support our strategy:

- The manual evaluation of fluency carried out at the WMT16 shared translation task (Bogiar et al., 2016, Sec. 3.5) already confirmed that NMT systems consistently produce more fluent translations than PBMT systems. That manual evaluation only covered language directions into English. In this experiment, we extend that conclusion to language directions out of English.
- Neural LMs consistently outperform *n*-gram based LMs when assessing the fluency of *real* text (Kim et al., 2016; Enarvi and Kurimo, 2016). Thus, we have used the most accurate automatic tool available to measure fluency.

Language direction	Monotone vs.			PBMT vs. Ref.	NMT vs. Ref.
	PBMT	NMT	Ref.		
EN→CS	0.9273	0.9029	0.8295	0.8008	0.7964
EN→DE	0.8006	0.8229	0.7740	0.7560	0.7791
EN→FI	0.8912	0.9172	0.8367	0.7611	0.7819
EN→RO	0.8389	0.8378	0.7937	0.8312	0.8282
EN→RU	0.9342	0.9114	0.8364	0.8249	0.8240
CS→EN	0.7694	0.7589	0.7128	0.8000	0.8015
DE→EN	0.8036	0.7830	0.7409	0.7728	0.7943
RO→EN	0.8693	0.8427	0.8013	0.8187	0.8245
RU→EN	0.8170	0.7891	0.7247	0.7958	0.8069

Table 5: Average Kendall’s tau distance between the word alignments obtained after translating the test set with each MT system being evaluated and a monotone alignment (left); and average Kendall’s tau distance between the word alignments obtained for each MT system’s translation and the word alignments of the reference translation (right). Larger values represent more similar alignments. If the difference between the distances depicted in the two last columns is statistically significant according to paired bootstrap resampling (Koehn, 2004) with $p = 0.05$ and 1 000 iterations, the largest distance is shown in bold.

5 Reordering

In this section we measure the amount of reordering performed by PBMT and NMT systems. Our objective is to empirically determine whether: (i) the recurrent neural networks in NMT systems produce more changes in the word order of a sentence than an PBMT decoder; and whether (ii) these neural networks make the word order of the translations closer to that of the reference.

In order to measure the amount of reordering, we used the *Kendall’s tau distance* between word alignments obtained from pairs of sentences (Birch, 2011, Sec. 5.3.2). As the distance needs to be computed from permutations,⁸ we turned word alignments into permutations by means of the algorithm defined by Birch (2011, Sec. 5.2).

For each language direction, we computed word alignments between the source-language side of the test set and the target-language reference, the PBMT output and the NMT output by means of MGIZA++ (Gao and Vogel, 2008). As the test sets are rather small for word alignment (1 999 to 3 000 sentence pairs depending on the language pair), we append bigger parallel corpora to help ensure accurate word alignments and avoid data sparseness. For languages for which in-domain

(news) parallel training data is available (German and Russian), we append that dataset (News Commentary). For the remaining languages (Finnish and Romanian) we use the whole Europarl corpus.

The amount of reordering performed by each system can be estimated as the distance between the word alignments produced by that system and a monotone word alignment. The similarity between the reorderings produced by each MT system and the reorderings in the reference translation can also be estimated as the distance between the corresponding word alignments. Table 5 shows the value of these distances for the language pairs included in our evaluation. The average over all the sentences in the test set of the distance proposed by Birch (2011) is depicted.

It can be observed that the amount of reordering introduced by both types of MT systems is lower than the quantity of reordering in the reference translation. NMT generally produces more changes in the structure of the sentence than PBMT. This is the case for all language pairs but two (EN→DE and EN→FI). A possible explanation for these two exceptions is the following: in the former language pair, the PBMT system is hierarchical (Williams et al., 2016) while in the latter, the output was reranked with neural LMs.

Concerning the similarity between the reorderings produced by both MT systems and those in the reference translation, out of 9 directions, in 5 directions the NMT system performs a reordering

⁸A permutation between a source-language sentence and a target-language sentence is defined as the set of operations that need to be carried out over the words in the source-language sentence to reflect the order of the words in the target-language sentence (Birch, 2011, Sec. 5.2).

closer to the reference, in 1 direction the PBMT system performs a reordering closer to the reference and in the remaining 3 directions the differences are not statistically significant. That is, NMT generally produces reorderings which are closer to the reference translation. The exceptions to this trend, however, do not exactly correspond to the language pairs for which NMT underperformed PBMT.

In summary, NMT systems achieve, in general, a higher degree of reordering than pure, phrase-based PBMT systems, and, overall, this reordering results in translations whose word order is closer to that of the reference translation.

6 Sentence Length

In this experiment we aim to find out whether the performances of NMT and PBMT are somehow sensitive to sentence length. In this regard, Bentivogli et al. (2016) found that, for transcribed speeches, NMT outperformed PBMT regardless of sentence length while also noted that NMT’s performance degraded faster than PBMT’s as sentence length increases. It should be noted, however, that sentences in our content type, news, are considerably longer than sentences in transcribed speeches.⁹ Hence, the current experiment will determine to what extent the findings on transcribed speeches stand also for texts made of longer sentences.

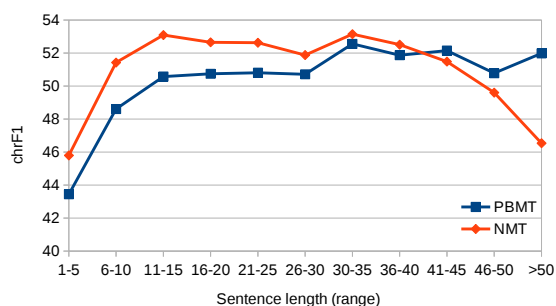


Figure 1: NMT and PBMT chrF1 scores on subsets of different sentence length for the language direction EN→FI.

We split the source side of the test set in subsets of different lengths: 1 to 5 words (1-5), 6 to 10 and so forth up to 46 to 50 and finally longer than 50 words (> 50). We then evaluate the out-

⁹According to Ruiz et al. (2014), sentences of transcribed speeches in English average to 19 words while sentences in news average to 24 words.

puts of the top PBMT and NMT submissions for those subsets with the chrF1 evaluation metric. Figure 1 presents the results for the language direction EN→FI. We can observe that NMT outperforms PBMT up to sentences of length 36-40, while for longer sentences PBMT outperforms NMT, with PBMT’s performance remaining fairly stable while NMT’s clearly decreases with sentence length. The results for the other language directions exhibit similar trends.

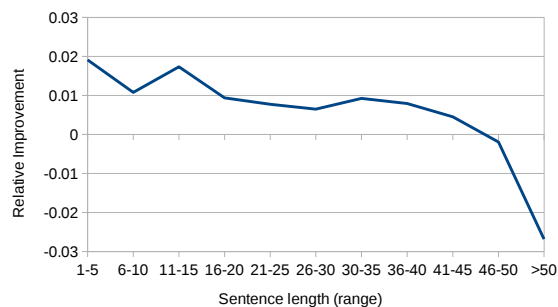


Figure 2: Relative improvement of the best NMT versus the best PBMT submission on chrF1 for different sentence lengths, averaged over all the language pairs considered.

Figure 2 shows the relative improvements of NMT over PBMT for each sentence length subset, averaged over all the 9 language directions considered. We observe a clear trend of this value decreasing with sentence length and in fact we found a strong negative Pearson correlation (-0.79) between sentence length and the relative improvement (chrF1) of the best NMT over the best PBMT system.

The correlations for each language direction are shown in Table 6. We observe negative correlations for all the language directions except for DE→EN.

Direction	CS	DE	FI	RO	RU
From EN	-0.72	-0.26	-0.89	-0.01	-0.74
Into EN	-0.19	0.10	-	-0.36	-0.70

Table 6: Pearson correlations between sentence length and relative improvement (chrF1) of the best NMT over the best PBMT system for each language pair.

7 Error Categories

In this experiment we assess the performance of NMT versus PBMT systems on a set of error

Error type	EN→CS	EN→DE	EN→FI	EN→RO	EN→RU	Average
Inflection	-16.18%	-13.26%	-11.65%	-15.13%	-16.79%	-14.60%
Reordering	-7.97%	-21.92%	-12.12%	-15.91%	-6.18%	-12.82%
Lexical	-0.44%	-3.48%	-1.09%	2.17%	-0.09%	-0.59%

Table 7: Relative improvement of NMT versus PBMT for 3 error categories, for language directions out of English.

Error type	CS→EN	DE→EN	RO→EN	RU→EN	Average
Inflection	-4.38%	-2.47%	-3.65%	-21.12%	-7.91%
Reordering	-8.68%	-21.09%	-9.48%	-8.50%	-11.94%
Lexical	-1.92%	-4.91%	-3.90%	5.32%	-1.35%

Table 8: Relative improvement of NMT versus PBMT for 3 error categories, for language directions into English.

categories that correspond to five word-level error classes: inflection errors, reordering errors, missing words, extra words and incorrect lexical choices. These errors are detected automatically using the edit distance, word error rate (WER), precision-based and recall-based position-independent error rates (hPER and rPER, respectively) as implemented in Hjerison (Popović, 2011). These error classes are then defined as follows:

- Inflection error (hINFer). A word for which its full form is marked as a hPER error while its base form matches the base form in the reference.
- Reordering error (hRer). A word that matches the reference but is marked as a WER error.
- Missing word (MISer). A word that occurs as deletion error in WER, is also a rPER error and does not share the base form with any hypothesis error.
- Extra word (EXTer). A word that occurs as insertion error in WER, is also a hPER error and does not share the base form with any reference error.
- Lexical choice error (hLEXer). A word that belongs neither to inflectional errors nor to missing or extra words.

Due to the fact that it is difficult to disambiguate between three of these categories, namely missing words, extra words and lexical choice errors (Popović and Ney, 2011), we group them in

a unique category, which we refer to as lexical errors.

As input, the tool requires the full forms and base forms of the reference translations and MT outputs. For base forms, we use stems for practical reasons. These are produced with the Snowball stemmer from NLTK¹⁰ for all languages except for Czech, which is not supported. For this language we used the aggressive variant in `czech_stemmer`.¹¹

Tables 7 and 8 show the results for language directions out of English and into English, respectively. For all language directions, we observe that NMT results in a notable decrease of both inflection (-14.6% on average for language directions out of EN and -7.91% for language directions into EN) and reordering (-12.82% from EN and -11.94 into EN) errors. The reduction of reordering errors is compatible with the results of the experiment presented in Section 5.¹²

Differences in performance for the remaining error category, lexical errors, are much smaller. In addition, the results for that category show a mixed picture in terms of which paradigm is better, which makes it difficult to derive conclusions that apply regardless of the language pair. Out of En-

¹⁰<http://www.nltk.org>

¹¹http://research.variancia.com/czech_stemmer/

¹²Although the results depicted both in this section and in Section 5 show that NMT performs better reordering in general, results for particular language pairs are not exactly the same in both sections. This is due to the fact that the quality of the reordering is computed in different ways. In this section, only those words that match the reference are considered when identifying reordering errors, while in Section 5 all the words in the sentence are taken into account. That said, in Section 5 the precision of the results depends on the quality of word alignment.

glish, NMT results in slightly less errors (0.59% decrease on average) for all target languages except for RO (2.17% increase). Similarly, in the opposite language direction, NMT also results in slightly better performance overall (1.35% error reduction on average), and looking at individual language directions NMT outperforms PBMT for all of them except RU→EN.

8 Conclusions

We have conducted a multifaceted evaluation to compare NMT versus PBMT outputs across a number of dimensions for 9 language directions. Our aim has been to shed more light on the strengths and weaknesses of the newly introduced NMT paradigm, and to check whether, and to what extent, these generalise to different families of source and target languages. Hereunder we summarise our findings:

- The outputs of NMT systems are considerably different compared to those of PBMT systems. In addition, there is higher inter-system variability in NMT, i.e. outputs by pairs of NMT systems are more different between them than outputs by pairs of PBMT systems.
- NMT outputs are more fluent. We have corroborated the results of the manual evaluation of fluency at WMT16, which was conducted only for language directions into English, and we have shown evidence that this finding is true also for directions out of English.
- NMT systems introduce more changes in word order than pure PBMT systems, but less than hierarchical PBMT systems.¹³ Nevertheless, for most language pairs, including those for which the best PBMT system is hierarchical, NMT's reorderings are closer to the reorderings in the reference than those of PBMT. This corroborates the findings on reordering by Bentivogli et al. (2016).
- We have found negative correlations between sentence length and the improvement brought by NMT over PBMT for the majority of the languages examined. While for most sentence lengths NMT outperforms PBMT, for very long sentences PBMT outperforms

¹³The latter finding applies only to one language direction as only for that one the best PBMT system is hierarchical.

NMT. The latter was not the case in the work by Bentivogli et al. (2016). We believe the reason behind this different finding is twofold. Firstly, the average sentence length in their evaluation dataset was considerably shorter; and secondly, the NMT systems included in our evaluation operate on subword units, which increases the effective sentence length they have to deal with.

- NMT performs better in terms of inflection and reordering consistently across all language directions. We thus confirm that the findings of Bentivogli et al. (2016) regarding these two error types apply to a wide range of language directions. Differences regarding lexical errors are much smaller and inconsistent across language directions; for 7 of them NMT outperforms PBMT while for the remaining 2 the opposite is true.

The results for some of the evaluations, especially error categories (Section 7) have been analysed only superficially, looking at what conclusions can be derived that apply regardless of language direction. Nevertheless, all our data is publicly released,¹⁴ so we encourage interested parties to use this resource to conduct deeper language-specific studies.

Acknowledgments

The research leading to these results is supported by the European Union Seventh Framework Programme FP7/2007-2013 under grant agreement PIAP-GA-2012-324414 (Abu-MaTran) and by Science Foundation Ireland through the CNGL Programme (Grant 12/CE/I2267) in the ADAPT Centre (www.adaptcentre.ie) at Dublin City University.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *Proceedings of ICLR 2015*, San Diego, CA, USA.
- Luisa Bentivogli, Arianna Bisazza, Mauro Cettolo, and Marcello Federico. 2016. Neural versus phrase-based machine translation quality: a case study. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 257–267, Austing, TX, USA, November.

¹⁴https://github.com/antot/neural_vs_phrasebased_smt_eacl17

- Alexandra Birch. 2011. *Reordering metrics for statistical machine translation*. Ph.D. thesis, The University of Edinburgh.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Aurelie Neveol, Mariana Neves, Martin Popel, Matt Post, Raphael Rubino, Carolina Scarton, Lucia Specia, Marco Turchi, Karin Verspoor, and Marcos Zampieri. 2016. Findings of the 2016 conference on machine translation. In *Proceedings of the First Conference on Machine Translation*, pages 131–198, Berlin, Germany, August.
- Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *Proceedings of the Deep Learning and Representation Learning Workshop, NIPS*.
- Junyoung Chung, Kyunghyun Cho, and Yoshua Bengio. 2016. A character-level decoder without explicit segmentation for neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 1693–1703, Berlin, Germany, August.
- Marta R. Costa-Jussà and José A. R. Fonollosa. 2016. Character-based neural machine translation.
- Jacob Devlin, Rabih Zbib, Zhongqiang Huang, Thomas Lamar, Richard Schwartz, and John Makhoul. 2014. Fast and robust neural network joint models for statistical machine translation. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1370–1380, Baltimore, Maryland, June.
- Shuoyang Ding, Kevin Duh, Huda Khayrallah, Philipp Koehn, and Matt Post. 2016. The JHU Machine Translation Systems for WMT 2016. In *Proceedings of the First Conference on Machine Translation*, pages 272–280, Berlin, Germany, August.
- John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159.
- Nadir Durrani, Helmut Schmid, and Alexander Fraser. 2011. A Joint Sequence Translation Model with Integrated Reordering. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1045–1054, Portland, Oregon, USA, June.
- Seppo Enarvi and Mikko Kurimo. 2016. TheanoLM – An Extensible Toolkit for Neural Network Language Modeling. In *Proceedings of the 17th Annual Conference of the International Speech Communication Association*.
- Christian Federmann. 2012. Appraise: An open-source toolkit for manual evaluation of machine translation output. *The Prague Bulletin of Mathematical Linguistics*, 98:25–35, September.
- Qin Gao and Stephan Vogel. 2008. Parallel implementations of word alignment tool. In *Software Engineering, Testing, and Quality Assurance for Natural Language Processing, SETQA-NLP '08*, pages 49–57, Columbus, Ohio.
- Yvette Graham, Timothy Baldwin, Alistair Moffat, and Justin Zobel. 2016. Can machine translation systems be evaluated by the crowd alone. *Natural Language Engineering*, FirstView:1–28, 1.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Marcin Junczys-Dowmunt, Tomasz Dwojak, and Rico Sennrich. 2016. The AMU-UEDIN Submission to the WMT16 News Translation Task: Attention-based NMT Models as Feature Functions in Phrase-based SMT. In *Proceedings of the First Conference on Machine Translation*, pages 319–325, Berlin, Germany, August.
- Yoon Kim, Yacine Jernite, David Sontag, and Alexander Rush. 2016. Character-aware neural language models. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 2741–2749, Phoenix, Arizona, USA, February.
- Philipp Koehn. 2004. Statistical significance tests for machine translation evaluation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, volume 4, pages 388–395, Barcelona, Spain.
- Philipp Koehn. 2010. *Statistical Machine Translation*. Cambridge University Press, New York, NY, USA, 1st edition.
- Chi-kiu Lo, Colin Cherry, George Foster, Darlene Stewart, Rabib Islam, Anna Kazantseva, and Roland Kuhn. 2016. NRC Russian-English Machine Translation System for WMT 2016. In *Proceedings of the First Conference on Machine Translation*, pages 326–332, Berlin, Germany, August.
- Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, Lisbon, Portugal, September.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA, July.

- Maja Popović and Hermann Ney. 2011. Towards automatic error analysis of machine translation output. *Comput. Linguist.*, 37(4):657–688, December.
- Maja Popović. 2011. Hjerson: An open source tool for automatic error classification of machine translation output. *The Prague Bulletin of Mathematical Linguistics*, 96:59–67.
- Maja Popović. 2015. chrF: character n-gram F-score for automatic MT evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal, September.
- Nicholas Ruiz and Marcello Federico. 2014. Complexity of spoken versus written language for machine translation. In *17th Annual Conference of the European Association for Machine Translation, EAMT*, pages 173–180, Dubrovnik, Croatia, June.
- Víctor M. Sánchez-Cartagena and Antonio Toral. 2016. Abu-matran at wmt 2016 translation task: Deep learning, morphological segmentation and tuning on character sequences. In *Proceedings of the First Conference on Machine Translation*, pages 362–370, Berlin, Germany, August.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. Edinburgh Neural Machine Translation Systems for WMT 16. In *Proceedings of the First Conference on Machine Translation*, pages 371–376, Berlin, Germany, August.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. Improving Neural Machine Translation Models with Monolingual Data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 1693–1703, Berlin, Germany, August.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation. In *Proceedings of AMTA*, pages 223–231.
- Philip Williams, Rico Sennrich, Maria Nadejde, Matthias Huck, Barry Haddow, and Ondřej Bojar. 2016. Edinburgh’s statistical machine translation systems for wmt16. In *Proceedings of the First Conference on Machine Translation*, pages 399–410, Berlin, Germany, August.