University of Groningen

What Level of Quality Can Neural Machine Translation Attain on Literary Text?

Toral, Antonio; Way, Andy

Published in:
Translation Quality Assessment

DOI:
[10.1007/978-3-319-91241-7_12](https://doi.org/10.1007/978-3-319-91241-7_12)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2018

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Toral, A., & Way, A. (2018). What Level of Quality Can Neural Machine Translation Attain on Literary Text? In J. Moorkens, S. Castilho, F. Gaspari, & S. Doherty (Eds.), *Translation Quality Assessment: From Principles to Practice* (pp. 263-287). (Machine Translation: Technologies and Applications; Vol. 1). Springer International Publishing AG. https://doi.org/10.1007/978-3-319-91241-7_12

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

What Level of Quality Can Neural Machine Translation Attain on Literary Text?

Antonio Toral and Andy Way

Abstract Given the rise of the new neural approach to machine translation (NMT) and its promising performance on different text types, we assess the translation quality it can attain on what is perceived to be the greatest challenge for MT: literary text. Specifically, we target novels, arguably the most popular type of literary text. We build a literary-adapted NMT system for the English-to-Catalan translation direction and evaluate it against a system pertaining to the previous dominant paradigm in MT: statistical phrase-based MT (PBSMT). To this end, for the first time we train MT systems, both NMT and PBSMT, on large amounts of literary text (over 100 million words) and evaluate them on a set of 12 widely known novels spanning from the 1920s to the present day. According to the BLEU automatic evaluation metric, NMT is significantly better than PBSMT ($p < 0.01$) on all the novels considered. Overall, NMT results in a 11% relative improvement (3 points absolute) over PBSMT. A complementary human evaluation on three of the books shows that between 17% and 34% of the translations, depending on the book, produced by NMT (versus 8% and 20% with PBSMT) are perceived by native speakers of the target language to be of equivalent quality to translations produced by a professional human translator.

Keywords Translation quality assessment · Principles to practice · Literature translation · Neural machine translation · Pairwise ranking · Phrase-based statistical machine translation

A. Toral (✉)

Faculty of Arts, Center for Language and Cognition, University of Groningen, Groningen, The Netherlands

e-mail: a.toral.ruiz@rug.nl

A. Way

ADAPT Centre/School of Computing, Dublin City University, Dublin, Ireland

e-mail: andy.way@adaptcentre.ie

1 Introduction

Literary text is considered to be the greatest challenge for machine translation (MT). According to perceived wisdom, despite the tremendous progress in the field of statistical MT over the past two decades, there is no prospect of machines being useful in (assisting with) the translation of this type of content.

However, we believe that the recent emergence of two unrelated technologies opens a window of opportunity to explore this topic:

1. The electronic book: the market share of e-books is continuously growing,¹ as a result of which there is a wide availability of books in digital format, including original novels and their translations. Because the main resource required to train statistical MT systems is bilingual parallel text, we are now able to build MT systems tailored to novels. This should result in better performance, as it has been shown in MT research again and again that for a statistical MT engine to perform optimally it should be trained on similar data to the data it is applied to, e.g. Pecina et al. (2014).
2. Neural MT: NMT is a new approach to statistical MT, which, while having been introduced only very recently,² has already shown great potential, as there is evidence that it can attain better translation quality than the dominant approach to date, namely phrase-based statistical MT (PBSMT). This has been shown for a number of language pairs and domains, including transcribed speeches (Luong and Manning 2015), newswire (Sánchez-Cartagena and Toral 2016) and United Nations documents (Junczys-Dowmunt et al. 2016). Beyond its generally positive performance, NMT is of particular interest for literary texts due to the following two findings:
 - Its performance seems to be especially promising for lexically-rich texts (Bentivogli et al. 2016), which is the case with literary texts.
 - There are claims that NMT “can, rather than do a literal translation, find the cultural equivalent in another language”.³

With respect to the last point, literal (but not word-for-word) translations are deemed acceptable for domains for which PBSMT is already widely used in industry, such as technical documentation, as the aim of the translation process here is purely to carry over the meaning of the source sentence to the target language, without necessarily reflecting any stylistic niceties of the target language. In contrast, literal translations are not at all suitable for literary texts because

¹For example, in the US the market share of e-books surpassed that of printed books for fiction in 2014, <http://www.ingenta.com/blog-article/adding-up-the-invisible-ebook-market-analysis-of-author-earnings-january-2015-2/>

²Working models of NMT have only recently been introduced, but from a theoretical perspective, very similar models can be traced back two decades (Forcada and Neco 1997).

³<http://events.technologyreview.com/video/watch/alan-packer-understanding-language/>

the expectations of the reader are considerably higher; it is not sufficient for the translation to merely preserve the meaning, as it should also preserve the reading experience of the original text.

In this chapter we aim to assess the performance that can be offered by state-of-the-art MT for literary texts. To this end we train PBSMT and NMT systems for the first time on large amounts of literary texts (over 100 million words) and evaluate them on a set of 12 widely known novels that span from the 1920s to the beginning of the twenty-first century.

The rest of the chapter is organised as follows. In the following section, we provide an overview of the research carried out in the field of MT targeting literary texts. Next, we outline our experimental set-up (Sect. 3) and provide technical details of the PBSMT and NMT systems built (Sect. 4). Subsequently we evaluate and analyse the translations produced by both MT systems (Sect. 5). Finally, we conclude and outline lines of future work in Sect. 6.

2 State-of-the-Art in MT of Literary Text

There has been recent interest in the Computational Linguistics community regarding the processing of literary text. The best example is the establishment of an annual workshop (Computational Linguistics for Literature) in 2012, which has run ever since. A popular strand of research concerns the automatic identification of text snippets that convey figurative devices, such as metaphor (Shutova et al. 2013), idioms (Li and Sporleder 2010), humour and irony (Reyes 2013), applied to monolingual text. Conversely, there has been a rather limited amount of work on applying MT to literary texts, as we now survey.

Genzel et al. (2010) constrained SMT systems for poetry to produce French-to-English translations that obey length, meter, and rhyming rules. Form is preserved at the price of producing considerably lower-quality translations; the score according to the BLEU automatic evaluation metric (Papineni et al. 2002) (see the papers by Castilho et al. and Way in this volume for more details of this metric) decreases by around 50%, although it should be noted that their evaluation was not on poetry but on news.

Greene et al. (2010) translated poetry, choosing target output realisations that conform to the desired rhythmic patterns. Specifically, they translated Dante's *Divine Comedy* from Italian sonnets into English iambic pentameter. Instead of constraining the SMT system, as done by Genzel et al. (2010), they passed its output lattice through a device that maps words to sequences of stressed and unstressed syllables. These sequences were finally filtered with an iambic pentameter acceptor.

Voigt and Jurafsky (2012) examined the role of referential cohesion in translation and found that literary texts have more dense reference chains. They concluded that incorporating discourse features beyond the level of the sentence (Hardmeier 2014) is an important research focus for applying MT to literary texts.

Jones and Irvine (2013) used general-domain MT systems to translate samples of French literature (prose and poetry) into English. They then used qualitative analysis grounded in translation theory on the MT output to assess the potential of MT in literary translation and to address what makes literary translation particularly difficult.

Besacier (2014) used MT followed by post-editing (by non-professional translators) to translate a short story from English into French. Such a workflow was deemed a useful low-cost alternative for translating literary works, albeit at the expense of lower translation quality.

Our recent work (Toral and Way 2015b) contributed to the state-of-the-art in two dimensions. First, we conducted a comparative analysis on the translatability of literary text according to narrowness of the domain and freedom of translation, which is more general and complementary to the analysis by Voigt and Jurafsky (2012). Second, related to Besacier (2014), we evaluated MT for literary text. There were two differences though; first, Besacier (2014) translated a short story, while we translated a novel; second, their MT systems were evaluated against a post-edited reference produced by non-professional translators, while we evaluated our MT systems against the translation produced by a professional translator.

This work builds upon our previous study (Toral and Way 2015b), the following being the main differences between the two: we now train a literary-adapted MT system under the NMT paradigm (while previously we used PBSMT), the translation direction considered is more challenging as the languages are more distant (English-to-Catalan versus Spanish-to-Catalan), we conduct a considerably broader evaluation (12 books now versus just one in the previous work), and we analyse the results with respect to a set of textual features of each novel.

3 Experimental Set-Up

This section covers the experimental settings. We explain the motivation for the language pair chosen for this chapter (Sect. 3.1), describe the data sets used in our experiments (Sect. 3.2) and finally the tools that were utilised (Sect. 3.3).

3.1 Language Pair

In general, it is widely accepted that the quality attainable by MT correlates with the level of relatedness between the pair of languages involved. This is because translations between related languages should be more literal, and complex phenomena (such as metaphorical expressions) might simply transfer rather straightforwardly to the target language, while they are more likely to require complex translations between unrelated languages.

In our previous work (Toral and Way 2015a,b), we considered a closely-related language pair (Spanish-to-Catalan), where both languages belong to the same family (Romance). We built a literary-adapted PBSMT system and used it to translate a novel from an internationally renowned author, Ruiz Zafón. We concluded that our system could be useful to assist with the translation of this kind of text due to the following two findings.

1. For a random subset of sentences from the novel, we asked native speakers to rank the translations coming from the MT system against those from a professional translator (i.e. taken from the published novel in the target language), although they did not know which were which. For over 60% of the sentences, native speakers found both translations to be of the same quality (Toral and Way 2015b).
2. The previous evaluation was carried out at the sentence level, so it might be argued that this is somewhat limited as it does not take context beyond the sentence into account. Accordingly, we subsequently analysed 3 representative passages (up to 10 consecutive sentences): one of average MT quality (i.e. the quality of this passage is similar to the quality obtained by MT on the whole novel, as measured with BLEU), another of high quality (i.e. its BLEU score is similar to the average BLEU score of the 20% highest-scoring passages), and finally, one of low quality (i.e. its BLEU score is similar to the average BLEU score of the 20% lowest-scoring passages). For the passages of high and average quality, we showed that the MT output requires only a few character edits to match the professional translation (Toral and Way 2015a).

Encouraged by the positive results obtained on a closely-related language pair, we have now decided to explore the potential for a less-related pair, correspondingly a more challenging task. The language pair in this study is English-to-Catalan, where the two languages involved belong to different families (Germanic and Romance, respectively).

We choose Catalan as the target language as an example of a mid-size European language.⁴ These are languages into which a significant number of novels have been translated; we have easily identified over 200 English e-books available in Catalan. Nonetheless, this number is very low compared to the amount of books translated into ‘major’ European languages (such as German, French, Italian, or Spanish). Concerning mid-size European languages, because there is (i) a reasonable amount of data available to train literary-adapted MT systems and also (ii) room to have more novels translated if the output translations produced by MT are deemed useful to assist translators, we believe this is a sensible choice of target language type for this line of research.

⁴With this term we refer to European languages with around 5–10 million speakers, as is the case of many other languages in Europe, such as Danish, Serbian, Czech, etc.

Dataset	# sentences	# tokens	
		English	Catalan
Training parallel (in-domain)	1,086,623	16,876,830	18,302,284
Training parallel (OpenSubs)	402,775	3,577,109	3,381,241
Training monolingual (in-domain)	5,306,055	—	100,426,922
Training monolingual (in-domain)	13,841,542	210,337,379	—
Training monolingual (web)	16,516,799	—	486,961,317
Development	2,000	34,562	38,114

Table 1 Number of sentences and tokens (source and target sides) in the training and development data sets

3.2 Data Sets

3.2.1 Training and Development Data

We use parallel and monolingual in-domain data for training. The parallel data comprises 133 parallel novels (over one million sentence pairs), while the monolingual data consists of around 1,000 books written in Catalan (over five million sentences) and around 1,600 books in English⁵ (over 13 million sentences). In addition, we use out-of-domain datasets, namely OpenSubtitles⁶ as parallel data (around 400,000 sentence pairs) and monolingual Catalan data (around 16 million sentences) crawled from the web (Ljubešić and Toral 2014). The development data consists of 2,000 sentence pairs randomly selected from the in-domain parallel training data and removed from the latter data set. Quantitative details of the training and development data sets are shown in Table 1.

3.2.2 Test Data

We test our systems on 12 English novels and their professional translations into Catalan. In so doing we aim to build up a representative sample of literary fiction, encompassing novels from different periods (from the 1920s to the present day) and genres and targeted at different audiences. Details are provided in Table 2. For each novel, aside from the number of sentences and tokens (i.e. words) that it contains, we also show the portion of the source book (percentage of sentences) that was evaluated.⁷

⁵While our experiments are for the English-to-Catalan language pair, we also use English monolingual data to generate synthetic data for our NMT system (see Sect. 4.2).

⁶<http://opus.lingfil.uu.se/OpenSubtitles.php>

⁷In order to build the test sets we sentence-align the source and target versions of the books. We keep the subset of sentence pairs whose alignment score is above a certain threshold. See Sect. 3.3.1 for further details.

Author, book and year	% sentences	# sentences	# tokens	
			English	Catalan
Auster's <i>Sunset Park</i> (2010)	75.43%	2,167	70,285	73,541
Collins' <i>Hunger Games #3</i> (2010)	73.36%	7,287	103,306	112,255
Golding's <i>Lord of the Flies</i> (1954)	82.93%	5,195	64,634	69,807
Hemingway's <i>The Old Man and the Sea</i> (1952)	76.01%	1,461	24,233	25,765
Highsmith's <i>Ripley Under Water</i> (1991)	65.86%	5,981	84,339	94,565
Hosseini's <i>A Thousand Splendid Suns</i> (2007)	67.54%	6,619	97,728	105,989
Joyce's <i>Ulysses</i> (1922)	46.65%	11,182	136,250	159,460
Kerouac's <i>On the Road</i> (1957)	76.35%	5,944	106,409	111,562
Orwell's <i>1984</i> (1949)	68.23%	4,852	84,062	90,545
Rowling's <i>Harry Potter #7</i> (2007)	69.61%	10,958	186,624	209,524
Salinger's <i>The Catcher in the Rye</i> (1951)	76.57%	5,591	77,717	77,371
Tolkien's <i>The Lord of the Rings #3</i> (1955)	66.60%	6,209	114,847	129,671

Table 2 Percentage of sentences used from the original data set and number of sentences and tokens in the novels that make up the test set

Author	# books	# sentence pairs	# tokens (English)
Auster	2	6,831	145,195
Collins	2	15,315	216,658
Golding	0	0	0
Hemingway	0	0	0
Highsmith	4	27,024	382,565
Hosseini	1	7,672	105,040
Joyce	2	8,762	146,525
Kerouac	0	0	0
Orwell	2	4,068	88,372
Rowling	6	50,000	836,942
Salinger	4	8,350	141,389
Tolkien	3	23,713	397,328

Table 3 Number of books in the training set, together with their overall number of sentence pairs and source-side tokens for each writer that is also represented in the test set

Whilst obviously none of the novels in the test set is included in the training data, the latter dataset may contain other novels from writers represented in the test set. For example, the test set contains the 7th book in the *Harry Potter* series from Rowling, while the training set contains the previous six books in that series. Table 3 shows, for each writer represented in the test set, how many books appear in the training set from this writer, and how many sentence pairs and tokens (source side) these books amount to.

3.3 Tools

We have leveraged state-of-the-art techniques in the field through the pervasive use of open-source tools throughout the different stages of our experimentation, namely preprocessing, MT experimentation and evaluation, as detailed in the remainder of this section.

3.3.1 Preprocessing

The datasets (see Sect. 3.2) are preprocessed in order to make them suitable for MT. In-domain data is extracted from e-books and converted to plain text with Calibre support tools,⁸ then sentence-split with NLTK (Bird 2006) and Freeling (Padró and Stanilovsky 2012) for English and Catalan, respectively, subsequently tokenised with Moses' scripts (Koehn et al. 2007) and Freeling, for English and Catalan, respectively, and finally sentence-aligned with Hunalign (Varga et al. 2005). Sentence alignment is carried out on lowercased text, in order to reduce data sparsity, with the assistance of a bilingual dictionary extracted from the Catalan–English Apertium rule-based MT system.⁹ Following empirical observations, we keep aligned sentences with confidence scores higher than 0.3 and 0.5 for the training and test sets, respectively.

Subsequently, all datasets are truecased and normalised in terms of punctuation with Moses' scripts. Finally, in the parallel training data we discard sentence pairs where either of the sides has fewer than 1 or more than 80 tokens.

3.3.2 MT Toolkits and Evaluation

PBSMT systems are trained with version 3 of the Moses toolkit, while NMT systems are trained with Nematus (Sennrich et al. 2017).¹⁰ For both paradigms default settings are used, unless mentioned otherwise in the description of the experiments (see Sects. 4.1 and 4.2 for PBSMT and NMT, respectively).

Automatic evaluation is carried out with the BLEU metric and is case-insensitive. Multi-bleu as implemented in Moses 3.0 is used for evaluating the development set while mteval (13a) is used to evaluate the test set. Statistical significance of the difference between systems is computed with paired bootstrap resampling (Koehn 2004) ($p \leq 0.01$, 1 000 iterations).¹¹ Human evaluation is rank-based and is performed with the Appraise tool (Federmann 2012).¹²

⁸<https://calibre-ebook.com/>

⁹<http://sourceforge.net/projects/apertium/files/apertium-en-ca/0.9.3/>

¹⁰<https://github.com/rsennrich/nematus>

¹¹http://www.cs.cmu.edu/~ark/MT/paired_bootstrap_v13a.tar.gz

¹²<https://github.com/cfedermann/Appraise>

4 MT Systems

4.1 PBSMT System

The PBSMT system is trained on both the in-domain and out-of-domain parallel datasets by means of linear interpolation (Sennrich 2012) and uses three reordering models (lexical- and phrase-based as well as hierarchical). In addition, the system makes use of additional feature functions based on the operation sequence model (OSM) (Durrani et al. 2011) and language models based not only on surface n -grams but also on continuous space n -grams (NPLM) (Vaswani et al. 2013). The OSM and NPLM models are built on the in-domain parallel data (both sides in the case of OSM and only the target side for NPLM). The vocabulary size for NPLM is set to 100,000. Surface-form n -gram language models are built on the in-domain and out-of-domain datasets with KenLM (Heafield 2011) and then linearly interpolated with SRILM (Stolcke 2002). Tuning is carried out with batch MIRA (Cherry and Foster 2012).

During development we tuned PBSMT systems using different subsets of the components previously introduced in order to assess their effect on translation quality as measured by the BLEU evaluation metric. Table 4 shows the results, where we start with a baseline trained on in-domain data (in) both for the translation model (TM) and the language model (LM) and we measure the effect of the following:

- Adding NPLM, both using 4- and 5-g, which results in absolute improvements of 0.57 and 0.75 BLEU points, respectively.
- Adding OSM (+0.4).
- Adding linearly interpolated out-domain data both for the TM and the LM (+0.14).

4.2 NMT System

Due to the lack of established domain-adaptation techniques for NMT at the time when this system was built, our NMT system was trained solely on in-domain data. Specifically, we trained our NMT system on the concatenation of the parallel in-domain training data and a synthetic corpus obtained by machine-translating the Catalan in-domain monolingual training data into English.

TM	LM	OSM	NPLM	BLEU
in	in	–	–	0.3344
in	in	–	4-g	0.3401
in	in	–	5-g	0.3419
in	in	y	5-g	0.3459
inIout	inIout	y	5-g	0.3473

Table 4 Performance of different configurations of the PBSMT system on the development set

We use additional parallel data in which the source side is synthetic (machine-translated from the target language), as this has been reported to be a successful way of integrating target-language monolingual data into NMT (Sennrich et al. 2015) (see also Footnotes 12 and 25 in Way’s chapter in this volume for a discussion on “back-translation”). The in-domain monolingual training data for Catalan is translated into English by means of a Catalan-to-English PBSMT system built for this purpose. This PBSMT system is based on the PBSMT system described in Sect. 4.1. Aside from reversing the translation direction, this PBSMT system is trained on the same datasets and has the same components, except for the following, which are not used: out-of-domain training data (both parallel and monolingual) and NPLM. The reason not to use these components has to do with an efficiency versus translation quality trade-off; this system needs to be fast as it is used to translate over five million sentences (i.e. the in-domain monolingual training data for Catalan), and taking the example of NPLM, this is a rather computationally expensive component to run.

We limit the source and target vocabularies to the 50,000 most frequent tokens in the respective sides of the training data. Training is then run until convergence, with models being saved every 3 h.¹³ Each model is evaluated on the development set using BLEU in order to track performance over training time and find out when the training reaches convergence.

Figure 1 shows the results. We can observe that performance increases very quickly in the first iterations, going from 0.0251 BLEU points for model 1 (i.e. after 3 h of training) to 0.2999 for model 12 (i.e. after 36 h), after which it grows slowly to reach its maximum (0.3356) for model 53 and then plateaus.

We select the four models with the highest BLEU scores. These are, in descending order, 53 (0.3356 points), 76 (0.3333), 74 (0.3322) and 69 (0.3314). We trained these models for 12 h with the embeddings frozen (i.e. the whole network keeps being trained except the first layer (embeddings) which is fixed). We then evaluate ensembles of these four models ‘as is’ as well as with the additional training for 12 h with fixed embeddings. Their BLEU scores are 0.3561 (2.05 points absolute higher than the best individual system, a 6.1% relative improvement) and 0.3555 (1.99 points absolute higher than the best individual system, 5.9% relative), respectively. In other words, ensembling led to a substantial improvement, but fixing embeddings – reported to provide further improvements in several experiments in the literature – did not increase performance in our set-up.

Subsequently, we tried to improve upon this NMT system by implementing the following two functionalities:

1. Using subwords rather than words as translation units. Specifically, we segmented the training data into characters and performed 90,000 operations jointly on both the source and target languages (Sennrich et al. 2016b). These operations iteratively join the most frequent pair of segments. This results in a score of 0.3689 (1.28 points absolute higher than the initial NMT ensemble, a 3.6% relative improvement).

¹³Training is performed on an NVIDIA Tesla K20X GPU.

Fig. 1 BLEU scores obtained by the NMT models on the development set

2. Producing an n -best list and reranking it with a right-to-left NMT system (Senrich et al. 2016a). We trained a so-called right-to-left system, with almost the same settings, the only difference being that the target sentences of the training data are reversed at the word level. We then produced an n -best list containing the top-50 translations with the previous model and re-ranked it with the right-to-left model. This leads to a BLEU score of 0.3948 (2.59 points higher than the previous system, a 7% relative improvement), and almost 6 BLEU points better (a 17.6% relative improvement) than the best individual system.

Due to the fact that we use the same dataset for development in the PBSMT and NMT paradigms, we are able to compare their results. When doing so, however, one should take into account that any such comparison would be unfair to NMT. This is because in the development of PBSMT, the system is optimising its log-linear weights to obtain the highest performance on the development set. Conversely, in the development of NMT we use the development set for validation, i.e. the system is not optimised on the development set. Despite this bias towards PBSMT, we observe that the score obtained by the best NMT system (0.3948, ensemble, using subword units and re-ranked with a right-to-left model) is notably higher (4.75 points, a 13.7% relative improvement) than the score achieved by the best PBSMT system (0.3473, all components, see Table 4).

5 Evaluation

5.1 Automatic Evaluation

As previously mentioned in Sect. 3.3.2, we automatically evaluate the MT systems using the BLEU metric. Table 5 shows the BLEU scores obtained for each novel in the test set with both PBSMT and NMT. The results across the different novels show a very high degree of variability, indicating that fiction is far from a monolithic domain. In fact scores go from a low of 0.1611 (PBSMT for *Ulysses*) to the highest of 0.3892 (NMT for *Harry Potter #7*), which more than doubles the first figure.

As for the performance obtained by the two paradigms that we compare in this chapter, NMT beats PBSMT by a statistically significant margin for all novels. On average, NMT outperforms PBSMT by 10.67% relative and 3 points absolute. The improvement brought about by NMT compared to PBSMT varies widely depending on the book, going from 3.11% (Auster's *Sunset Park*) to 14% (Collins' *Hunger Games #3*).

5.1.1 Analysis

We performed a set of additional analyses in order to obtain further insights from the output translations and, especially to try to find the reason why NMT, while outperforming PBSMT for all the novels, does so by rather diverging margins (from a minimum of 3.11% to a maximum of 14%, see Table 5).

Novel	PBSMT	NMT	Relative improvement (%)
Auster's <i>Sunset Park</i> (2010)	0.3735	0.3851	3.11
Collins' <i>Hunger Games #3</i> (2010)	0.3322	0.3787	14.00
Golding's <i>Lord of the Flies</i> (1954)	0.2196	0.2451	11.61
Hemingway's <i>The Old Man and the Sea</i> (1952)	0.2559	0.2829	10.55
Highsmith's <i>Ripley Under Water</i> (1991)	0.2485	0.2762	11.15
Hosseini's <i>A Thousand Splendid Suns</i> (2007)	0.3422	0.3715	8.56
Joyce's <i>Ulysses</i> (1922)	0.1611	0.1794	11.36
Kerouac's <i>On the Road</i> (1957)	0.3248	0.3572	9.98
Orwell's <i>1984</i> (1949)	0.2978	0.3306	11.01
Rowling's <i>Harry Potter #7</i> (2007)	0.3558	0.3892	9.39
Salinger's <i>The Catcher in the Rye</i> (1951)	0.3255	0.3695	13.52
Tolkien's <i>The Lord of the Rings #3</i> (1955)	0.2537	0.2888	13.84
Average	0.2909	0.3212	10.67

Table 5 BLEU scores obtained by PBSMT and NMT for each of the books that make up the test set. NMT outperforms PBSMT by a statistically significant margin ($p < 0.01$) on all books

More specifically, we considered three characteristics of the source-side of each novel in the test set (lexical richness, novelty with respect to the training data, and average sentence length) and studied whether any of these features correlates to some extent with the performance of the PBSMT and NMT systems and/or with the relative improvement of NMT over PBSMT. The motivation to use these three features is as follows:

- Lexical richness has been already studied in relation to NMT, and there are indications that this MT paradigm has “an edge especially on lexically rich texts” (Bentivogli et al. 2016).
- There is evidence that NMT’s performance degrades with sentence length (Toral and Sánchez-Cartagena 2017).
- Despite, to the best of our knowledge, the lack of empirical evidence, it is still the perceived wisdom that NMT is better at generalising than PBSMT, and so it should perform better than the latter especially on data that is unrelated to the training data.

Lexical Richness

We use type-token ratio (TTR) as a proxy to measure lexical richness. The higher the ratio, the less repetitive the text and hence it can be considered lexically more varied, and thus richer. To measure this we calculate the TTR on the source side of each novel. As they have different sizes, we calculate the TTR for each novel on a random subset of sentences that amount to approximately n words, n being 20,000, a slightly lower number to the number of words contained in the smallest novel in our dataset, *The Old Man and the Sea* with 24,233 words.

Sentence Length

We measure the average sentence length of each novel as the ratio between its total number of tokens and its number of sentences. Both these values were reported in Table 2.

Novelty with Respect to the Training Data

We use n -gram overlap to measure the novelty of a novel with respect to the training data. Concretely, we consider the unique n -grams ($n = 4$) in the parallel training data and in each novel, and calculate the overlap as the ratio between the size of the intersection and the number of unique n -grams in the training set. The higher the overlap, the less novelty that the novel presents with respect to the training data. As in the analysis concerning lexical richness, we consider 20,000 words from randomly selected sentences for each novel.

Novel	TTR	Avg. sentence length	Overlap
Auster's <i>Sunset Park</i> (2010)	0.1865	32.434	0.368
Collins' <i>Hunger Games</i> #3 (2010)	0.1716	14.177	0.393
Golding's <i>Lord of the Flies</i> (1954)	0.1368	12.442	0.370
Hemingway's <i>The Old Man and the Sea</i> (1952)	0.1041	16.587	0.371
Highsmith's <i>Ripley Under Water</i> (1991)	0.1492	14.101	0.404
Hosseini's <i>A Thousand Splendid Suns</i> (2007)	0.1840	14.765	0.377
Joyce's <i>Ulysses</i> (1922)	0.2761	12.185	0.216
Kerouac's <i>On the Road</i> (1957)	0.1765	17.902	0.335
Orwell's <i>1984</i> (1949)	0.1831	17.325	0.343
Rowling's <i>Harry Potter</i> #7 (2007)	0.1665	17.031	0.433
Salinger's <i>The Catcher in the Rye</i> (1951)	0.1040	13.900	0.448
Tolkien's <i>The Lord of the Rings</i> #3 (1955)	0.1436	18.497	0.368

Table 6 TTR, average sentence length and 4-g overlap for the source side of the 12 novels that make up the test set. The highest TTR and average sentence length as well as the lowest *n*-gram overlap values are shown in bold

BLEU	TTR	Avg. sentence length	Overlap
PBSMT	–	–	$r = 0.62, p < 0.05$
NMT	–	–	$r = 0.66, p < 0.01$
Rel. diff	–	$\rho = -0.45^a, p = 0.07$	–

Table 7 Correlations between the BLEU scores for NMT, PBSMT and their relative difference and the other metrics considered (TTR, average sentence length and 4-g overlap) for the 12 novels that make up the test set. Empty cells mean that no significant correlation was found

^aA significant parametric Pearson correlation was found ($r = -0.78, p < 0.01$) but the assumption that both variables come from a bivariate normal distribution was not met, hence the reason why a non-parametric Spearman correlation is shown instead

Results

Table 6 shows the values for each novel and for each of the three features analysed. We can clearly observe an outlier in the data for all the three variables reported. *Ulysses* has the highest TTR by far at 0.276 and is also the novel with the lowest overlap by a wide margin (0.216). As for sentence length, the value for *Sunset Park* (32.434) is over 10 points higher than the value for any other novel.

Table 7 shows the significant correlations between the BLEU scores (for PBSMT, NMT and the relative difference between both, see Table 5) and the three variables analysed (TTR, average sentence length and *n*-gram overlap). Each of the significant correlations is then plotted, including its regression line and its 95% confidence region, in Figs. 2, 3 and 4.

Fig. 2 Spearman correlation between the relative difference between the BLEU scores of the NMT and PBSMT systems and sentence length

Fig. 3 Pearson correlation between the BLEU score of the PBSMT system and 4-g overlap

While Bentivogli et al. (2016) found a moderate correlation ($r = 0.73$) between TTR and the gains by NMT over PBSMT (measured with mTER – multi-reference TER (Snover et al. 2006) – on transcribed speeches), there is no significant correlation in our set-up.

With respect to sentence length (see Fig. 2), we observe a negative correlation ($\rho = -0.45$), meaning that the relative improvement of NMT over PBSMT decreases with sentence length. This corroborates the findings in previous work (Toral and Sánchez-Cartagena 2017) and appears to be the main reason

Fig. 4 Pearson correlation between the BLEU score of the NMT system and 4-g overlap

behind the low relative improvement that NMT achieved for *Sunset Park* (see Table 5), as the average sentence length for this novel is very long compared to all the other novels in our test set (see Table 6).

Finally, we found significant positive correlations between the performance of both PBSMT and NMT and n -gram overlap, the Pearson correlation coefficient being $r = 0.62$ (see Fig. 3) and $r = 0.66$ (see Fig. 4), respectively. This matches the intuition that the performance of a statistical MT system should be better the more the test set resembles the training data. That said, we did not find significant correlations between the relative improvement of NMT over PBSMT and overlap. Thus, the perceived wisdom that the more unrelated the data to be translated from the test set the wider the gap between NMT and PBSMT, does not hold for our set-up.

5.2 Human Evaluation

We also conducted a manual evaluation, in order to gain further insights. A common procedure (e.g. conducted in the annual MT shared task at WMT)¹⁴ consists of ranking the translations produced by different MT systems (see also Section 3.4 in Castilho et al. in this volume). Given the source and target sides of the reference (human) translations, and two or more outputs from MT systems, these outputs are ranked according to their quality, e.g. in terms of adequacy and/or fluency.

¹⁴e.g. <http://www.statmt.org/wmt17/translation-task.html>

In our experiment, we are not only interested in comparing two MT systems – PBSMT and NMT – to each other, but also with respect to the human reference translation. Hence, we conduct the rank-based manual evaluation in a slightly modified setting; we do not provide the target of the reference translation as reference but as one of the translations to be ranked. The evaluator thus is provided with the source-side of the reference and three translations, one being the human translation and the other two the translations produced by the PBSMT and NMT systems. The evaluator of course does not know which is which. Moreover, in order to avoid any bias with respect to MT, they are not told whether the translations are human or automatic.

This human evaluation is conducted for three of the books used in the automatic evaluation: Orwell’s *1984*, Rowling’s *Harry Potter #7* and Salinger’s *The Catcher in the Rye*. For each of these books, the sentences in 10 randomly selected passages were ranked. Each passage is made of 10 contiguous sentences, the motivation being to provide the annotator with context beyond the sentence level. Therefore, sentences 1–10 (passage 1) are contiguous in the book, then there is a jump to a second passage contained in sentences 11–20, and so forth.

All the annotations were carried out by two native Catalan speakers with an advanced level of English. They both have a background in linguistics but no in-depth knowledge of statistical MT (again, to avoid any bias with respect to MT). Comprehensive instructions were provided to the evaluators in their native language, in order to minimise ambiguity and thus foster high inter-annotator agreement. Here we reproduce the translation into English of the evaluation instructions:

Given three translations, the task is to rank them:

- Rank a translation *A* higher (rank 1) than a translation *B* (rank 2), if the first translation is better than the second.
- Rank two translations *A* and *B* equally (rank 1 for both *A* and *B*), if both have an equivalent quality.
- Use the highest rank possible, e.g. if there are three translations *A*, *B* and *C*, and the quality of *A* and *B* is equivalent and both are better than *C*, then they should be ranked as follows: *A* = rank 1, *B* = rank 1, *C* = rank 2. Do NOT use lower rankings, e.g.: *A* = rank 2, *B* = rank 2, *C* = rank 3.

Please follow the following guidelines to decide that a translation is better than another:

- Above anything else: the meaning of the original is understood, all the information is preserved and, if possible, the translation sounds natural.
- If all translations preserve the meaning to a similar extent, you might compare the number of errors (e.g. lexical, syntax, etc) in each translation.
- If for a given set of translations, you cannot decide how to rank them, you can skip that set by pressing the button “flag error”.

Figure 5 shows a snapshot of the manual evaluation process. In this example the annotator is asked to rank three translations for the second sentence from the first passage of Salinger’s *The Catcher in the Rye*.

All of a sudden, he said, "For Chrissake, Holden. **This is about a goddam baseball glove.** " Cold as hell.

NA NA NA
— Reference

☐ Rank 1 ☐ Rank 2 ☐ Rank 3

Això és sobre un guant de beisbol.
— Translation 1

☐ Rank 1 ☐ Rank 2 ☐ Rank 3

Això va d'un cony de guant de beisbol.
— Translation 2

☐ Rank 1 ☐ Rank 2 ☐ Rank 3

Es tracta d'un coi de guant de beisbol.
— Translation 3

Fig. 5 Snapshot from the manual evaluation process

5.2.1 Inter-annotator Agreement

The inter-annotator agreement in terms of Fleiss' Kappa (Fleiss 1971) is 0.22 for Orwell, 0.18 for Rowling and 0.38 for Salinger. The values for Orwell and Salinger fall in the band of fair agreement [0.21, 0.4] (Landis and Koch 1977) while that for Rowling is at the higher end of slight agreement [0.01, 0.2]. For the sake of comparison, the average inter-annotator agreement at WMT for the closest language direction to ours (English-to-French) over the last four editions in which that language direction was considered is 0.29, see Table 4 in Bojar et al. (2016).

5.2.2 Pairwise Rankings

From the sets of annotations (rankings between three translations), we extract all pairwise rankings, i.e. the rankings for each pair of translations. Given two translations A and B , the pairwise ranking will be $A > B$ if translation A was ranked higher than B , $A < B$ if A was ranked lower than B , and $A = B$ if both were ranked equally.

It is worth mentioning that while the PBSMT translations consistently cover the source sentences, this is not always the case for the other two translations. NMT has a tendency towards omission errors (Klubička et al. 2017). The human translation sometimes does not cover the source sentence fully either. This may be due to a choice of the translator, e.g. to translate the sentence in a way that diverges notably from the source. There are also some cases where the human translation is

Fig. 6 Pairwise rankings between HT and MT for Orwell’s 1984

misaligned¹⁵ and so it is unrelated to the source sentence. Most cases in which the human translation is ranked lower than MT (PBSMT or NMT) are due to either of these two reasons. It is clearly unjustifiable to rank the human translation lower than MT in these cases, so we remove these pairwise rankings, i.e. $A < B$ where A is the human translation and B corresponds to the translation produced by either MT system.¹⁶

Figures 6, 7 and 8 show the pairwise rankings between each MT system and the human translation (HT) for Orwell’s, Rowling’s and Salinger’s books. In all three books, the percentage of sentences where the annotators perceive the MT translation to be of equivalent quality to the human translation is considerably higher for NMT compared to PBSMT: 16.7% vs. 7.5% for Orwell’s, 31.8% vs. 18.1% for Rowling’s and 34.3% vs. 19.8% for Salinger’s. In other words, if NMT translations were to be used to assist a professional translator (e.g. by means of post-editing), then around one third of the sentences for Rowling’s and Salinger’s and one sixth for Orwell’s would not need any correction.

¹⁵As mentioned in Sect. 3.3.1, the source novels and their human translations were sentence-aligned automatically. The empirically set confidence threshold results in most alignments being correct, but some are erroneous.

¹⁶While the majority of $HT < MT$ cases are unjustified, not all of them are. By removing these rankings, the results are slightly biased in favour of HT and thus overly conservative with respect to the potential of MT.

Fig. 7 Pairwise rankings between HT and MT for Rowling’s *Harry Potter* #7

Fig. 8 Pairwise rankings between HT and MT for Salinger’s *The Catcher in the Rye*

Having looked at pairwise rankings between MT and human translations, we move our attention now to the pairwise rankings between the two types of MT systems. The results for all three books are depicted in Fig. 9. In all the books

Fig. 9 Pairwise rankings between PBSMT and NMT

the trends are similar. The biggest chunk (41.4%, 54.7%) corresponds to cases where NMT translations are ranked higher than PBSMT’s (PBSMT < NMT). The second relates to translations by both systems which are ranked equally (27.8%, 39.4%), (PBSMT = NMT). Finally, the smallest chunk (less than 20% in all three books) signifies translations for which PBSMT is ranked higher than NMT (PBSMT > NMT).

5.2.3 Overall Human Scores

In addition to the pairwise rankings, we derive an overall score for each translation type (HT, NMT and PBSMT) and novel based on the rankings. To this end we use the TrueSkill method adapted to MT evaluation (Sakaguchi et al. 2014) following its usage at WMT15,¹⁷ i.e. we run 1,000 iterations of the rankings recorded with Appraise followed by clustering ($p < 0.05$).

Figure 10 depicts the results. For all three books considered, all the translation types are put in different clusters, meaning that the differences between every pair of translation types are significant. The ordering of the translation types corroborates that seen in the pairwise analysis (see Sect. 5.2.2), namely human translations come on top, followed by NMT outputs and finally, in third place, PBSMT outputs.

¹⁷<https://github.com/mjpost/wmt15>

Fig. 10 Overall human evaluation scores with TrueSkill

If we consider PBSMT's score as a baseline, the score given to the human translations as a gold standard, and the distance between the two as the potential room for improvement for MT, we could interpret NMT's score as the progress made in our journey towards better translation quality for novels, departing from PBSMT and targeting human translations as the goal to be reached ultimately. Using this analogy, although there is still a long way to go, with NMT we have covered already a considerable part of the journey: 20%, 22% and 18% for Orwell's, Rowling's and Salinger's books, respectively; while it may not yet be a case of *A Thousand Splendid Suns*, it can be said with confidence that we are *On The Road!*

6 Conclusions and Future Work

This chapter has assessed the quality attainable for novels by the two most common paradigms to MT at present, NMT and PBSMT. To this end, we built the first in-domain PBSMT and NMT systems for literary text by training them on large amounts of parallel novels. We then automatically evaluated the translation quality of the resulting systems on a set of 12 widely known novels spanning from the 1920s to the present day. The results proved favourable for NMT, which outperformed PBSMT by a significant margin for all the 12 novels.

We then delved deeper into the results by analysing the effect of three features of each novel: its lexical richness, its degree of novelty with respect to the training data, and its average sentence length. Only for the last feature did we find a meaningful

correlation with NMT relative improvement over PBSMT, which corroborates the tendency for improvements in NMT over PBSMT to decrease with sentence length. This seems to be the main reason behind NMT achieving a relatively low improvement over PBSMT for one of the novels, but we note that particular novel to be a considerable outlier in terms of sentence length.

We have also conducted a human evaluation, where we manually ranked the translations produced by NMT, PBSMT, as well as the human translations for three of the books. Again, NMT outperformed PBSMT. For two out of the three books native speakers perceived NMT translations to be of equivalent quality to those of human translations in around one third of the cases (one fifth for PBSMT).

As for future work, we would like to assess the feasibility of using MT to assist with the translation of literary text. To that end, we plan to carry out an experiment in which we integrate MT into the workflow of professional literary translators by means of post-editing and assess its impact in the translation process (e.g. temporal and technical effort) as well as in the translation result (e.g. quality and reading experience of the resulting translation).

Acknowledgements Carme Armentano and Álvaro Bellón ranked the translations used for the human evaluation. The research leading to these results has received funding from the European Association for Machine Translation through its 2015 sponsorship of activities programme (project PiPeNovel). The second author is supported by the ADAPT Centre for Digital Content Technology, funded under the SFI Research Centres Programme (Grant 13/RC/2106). We would like to thank the Center for Information Technology of the University of Groningen and the Irish Centre for High-End Computing (<http://www.ichec.ie>) for providing computational infrastructure.

References

- Bentivogli L, Bisazza A, Cettolo M, Federico M (2016) Neural versus phrase-based machine translation quality: a case study. In: Proceedings of the 2016 conference on empirical methods in natural language processing, Austin, pp 257–267
- Besacier L (2014) Traduction automatisée d’une oeuvre littéraire: une étude pilote. In: Traitement automatique du langage naturel (TALN), Marseille, <http://hal.inria.fr/hal-01003944>
- Bird S (2006) NLTK: the natural language toolkit. In: Proceedings of the COLING/ACL on interactive presentation sessions, Sydney, pp 69–72
- Bojar O, Chatterjee R, Federmann C, Graham Y, Haddow B, Huck M, Jimeno Yepes A, Koehn P, Logacheva V, Monz C, Negri M, Neveol A, Neves M, Popel M, Post M, Rubino R, Scarton C, Specia L, Turchi M, Verspoor K, Zampieri M (2016) Findings of the 2016 conference on machine translation. In: Proceedings of the first conference on machine translation, Berlin, pp 131–198
- Cherry C, Foster G (2012) Batch tuning strategies for statistical machine translation. In: Proceedings of the 2012 conference of the North American chapter of the Association for Computational Linguistics: human language technologies, Montréal, pp 427–436
- Durrani N, Schmid H, Fraser A (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, vol 1, Portland, pp 1045–1054
- Federmann C (2012) Appraise: an open-source toolkit for manual evaluation of machine translation output. Prague Bull Math Linguist 98:25–35

- Fleiss JL (1971) Measuring nominal scale agreement among many raters. *Psychol Bull* 76(5): 378–382
- Forcada ML, Neco RP (1997) Recursive hetero-associative memories for translation. In: *Proceedings of the biological and artificial computation: from neuroscience to technology: international work-conference on artificial and natural neural networks, IWANN'97 Lanzarote*, 4–6 June 1997. Springer, Berlin/Heidelberg, pp 453–462
- Genzel D, Uszkoreit J, Och F (2010) “Poetic” statistical machine translation: rhyme and meter. In: *Proceedings of the 2010 conference on empirical methods in natural language processing*, Cambridge, pp 158–166
- Greene E, Bodrumlu T, Knight K (2010) Automatic analysis of rhythmic poetry with applications to generation and translation. In: *Proceedings of the 2010 conference on empirical methods in natural language processing*, Cambridge, pp 524–533
- Hardmeier C (2014) Discourse in statistical machine translation. PhD thesis, University of Uppsala
- Heafield K (2011) Kenlm: faster and smaller language model queries. In: *Proceedings of the sixth workshop on statistical machine translation*, Edinburgh, pp 187–197
- Jones R, Irvine A (2013) The (un)faithful machine translator. In: *Proceedings of the 7th workshop on language technology for cultural heritage, social sciences, and humanities*, Sofia, pp 96–101
- Junczys-Dowmunt M, Dwojak T, Hoang H (2016) Is neural machine translation ready for deployment? A case study on 30 translation directions. *arXiv preprint arXiv:161001108*
- Klubička F, Toral A, Sánchez-Cartagena VM (2017) Fine-grained human evaluation of neural versus phrase-based machine translation. *Prague Bull Math Linguist* 108:121–132
- Koehn P (2004) Statistical significance tests for machine translation evaluation. In: *Proceedings of the conference on empirical methods in natural language processing*, Barcelona, vol 4, pp 388–395
- Koehn P, Hoang H, Birch A, Callison-Burch C, Federico M, Bertoldi N, Cowan B, Shen W, Moran C, Zens R, Dyer C, Bojar O, Constantin A, Herbst E (2007) Moses: open source toolkit for statistical machine translation. In: *Proceedings of the 45th annual meeting of the ACL on interactive poster and demonstration sessions*, Prague, pp 177–180
- Landis RJ, Koch GG (1977) The measurement of observer agreement for categorical data. *Biometrics* 33(1):159–174
- Li L, Sporleder C (2010) Using Gaussian mixture models to detect figurative language in context. In: *Human language technologies: the 2010 annual conference of the North American chapter of the Association for Computational Linguistics*, Los Angeles, pp 297–300
- Ljubešić N, Toral A (2014) caWaC – a Web corpus of Catalan and its application to language modeling and machine translation. In: *Proceedings of the ninth international conference on language resources and evaluation (LREC'14)*, Reykjavik, pp 1728–1732
- Luong MT, Manning CD (2015) Stanford neural machine translation systems for spoken language domain. In: *Proceedings of the international workshop on spoken language translation*, Da Nang, pp 76–79
- Padró L, Stanilovsky E (2012) Freeling 3.0: towards wider multilinguality. In: *Proceedings of the eighth international conference on language resources and evaluation (LREC-2012)*, Istanbul, pp 2473–2479
- Papineni K, Roukos S, Ward T, Zhu WJ (2002) Bleu: a method for automatic evaluation of machine translation. In: *Proceedings of the 40th annual meeting of Association for Computational Linguistics*, Philadelphia, pp 311–318
- Pecina P, Toral A, Papavassiliou V, Prokopidis P, Tamchyna A, Way A, van Genabith J (2014) Domain adaptation of statistical machine translation with domain-focused web crawling. *Lang Resour Eval* 49(1):147–193
- Reyes A (2013) Linguistic-based patterns for figurative language processing: the case of humor recognition and irony detection. *Procesamiento del Lenguaje Natural* 50:107–109
- Sakaguchi K, Post M, Van Durme B (2014) Efficient elicitation of annotations for human evaluation of machine translation. In: *Proceedings of the ninth workshop on statistical machine translation*, Baltimore, pp 1–11

- Sánchez-Cartagena VM, Toral A (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, Berlin, pp 362–370
- Sennrich R (2012) Perplexity minimization for translation model domain adaptation in statistical machine translation. In: Proceedings of the 13th conference of the European chapter of the Association for Computational Linguistics, Avignon, pp 539–549
- Sennrich R, Haddow B, Birch A (2015) Improving neural machine translation models with monolingual data. arXiv preprint arXiv:151106709
- Sennrich R, Haddow B, Birch A (2016a) Edinburgh neural machine translation systems for WMT 16. In: Proceedings of the first conference on machine translation, Berlin, pp 371–376
- Sennrich R, Haddow B, 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, Berlin. Long papers, vol 1, pp 1715–1725
- Sennrich R, Firat O, Cho K, Birch A, Haddow B, Hitschler J, Junczys-Dowmunt M, Läubli S, Miceli Barone AV, Mokry J, 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, Valencia, pp 65–68
- Shutova E, Teufel S, Korhonen A (2013) Statistical metaphor processing. *Comput Linguist* 39(2):301–353
- Snover M, Dorr B, Schwartz R, Micciulla L, Makhoul J (2006) A study of translation edit rate with targeted human annotation. In: AMTA 2006: proceedings of the 7th conference of the association for machine translation in the Americas, “Visions for the Future of Machine Translation”, Cambridge, pp 223–231
- Stolcke A (2002) SRILM—an extensible language modeling toolkit. In: Proceedings of the 7th international conference on spoken language processing (ICSLP 2002), Denver, pp 901–904
- Toral A, Sánchez-Cartagena VM (2017) A multifaceted evaluation of neural versus phrase-based machine translation for 9 language directions. In: Proceedings of the 15th conference of the European chapter of the Association for Computational Linguistics: volume 1, Long Papers, Valencia, pp 1063–1073
- Toral A, Way A (2015a) Machine-assisted translation of literary text: a case study. *Translat Spaces* 4:241–268
- Toral A, Way A (2015b) Translating literary text between related languages using SMT. In: Proceedings of the fourth workshop on computational linguistics for literature, Denver, pp 123–132
- Varga D, Halaácsy P, Kornai A, Nagy V, Németh L, Trón V (2005) Parallel corpora for medium density languages. In: International conference RANLP-2005, recent advances in natural language processing, proceedings, Borovets, pp 590–596
- Vaswani A, Zhao Y, Fossom V, Chiang D (2013) Decoding with large-scale neural language models improves translation. In: Proceedings of the 2013 conference on empirical methods in natural language processing, Seattle, pp 1387–1392
- Voigt R, Jurafsky D (2012) Towards a literary machine translation: the role of referential cohesion. In: Proceedings of the NAACL-HLT 2012 workshop on computational linguistics for literature, Montréal, pp 18–25