

# Automated Text Simplification as a Preprocessing Step for Machine Translation into an Under-resourced Language

**Sanja Štajner**

Symanto Research  
Nuremberg, Germany

sanja.stajner@symanto.net

**Maja Popović**

ADAPT Centre, DCU  
Dublin, Ireland

maja.popovic@adaptcentre.ie

## Abstract

In this work, we investigate the possibility of using fully automatic text simplification system on the English source in machine translation (MT) for improving its translation into an under-resourced language. We use the state-of-the-art automatic text simplification (ATS) system for lexically and syntactically simplifying source sentences, which are then translated with two state-of-the-art English-to-Serbian MT systems, the phrase-based MT (PBMT) and the neural MT (NMT). We explore three different scenarios for using the ATS in MT: (1) using the raw output of the ATS; (2) automatically filtering out the sentences with low grammaticality and meaning preservation scores; and (3) performing a minimal manual correction of the ATS output. Our results show improvement in fluency of the translation regardless of the chosen scenario, and difference in success of the three scenarios depending on the MT approach used (PBMT or NMT) with regards to improving translation fluency and post-editing effort.

## 1 Introduction

In spite of recent advances in machine translation (MT), the MT into under-resourced languages is still facing a number of problems. First, there is not enough parallel data to build robust phrase-based and neural systems. Second, the majority of those languages (including Serbian) have a very rich morphology and suffer from data sparsity when it comes to less frequently used cases, tenses, etc. Third, there is a number of syntactic differences which are difficult to capture. For English-to-Serbian phrase-based system, a num-

ber of language-related problems has been identified so far (Popović and Arčan, 2015). Most of them are related to syntactic differences, e.g. missing verb parts due to distinct structure of certain verb tenses, incorrect prepositions, or incorrect translations of English sequences of nouns. Although the neural approach better handles some grammatical aspects, it still often fails to generate correct inflections, prepositions and translations of the English noun phrases (Popović, 2017).

Text simplification (TS) has the goal of transforming given text or sentence into its simpler variant, while preserving the original meaning. What is considered to be a simpler variant depends on the target application, or the target reader in mind. In the case of simplifying texts for humans, a simpler variant is the one that requires a shorter reading time and leads to better text comprehension scores. In the case of text or sentence simplification used as a preprocessing step for a given natural language processing (NLP) task, e.g. machine translation (MT), information extraction (IE), summarization, and semantic role labeling (SRL), a simpler variant is the one that leads to better performances of that NLP system.

Text simplification was originally proposed as a pre-processing step for machine translation (Chandrasekar, 1994) and later for information extraction and parsing (Chandrasekar et al., 1996). At those early stages, automated text simplification (ATS) was not mature enough to help improving performances of those systems. Instead, the idea was explored only hypothetically, using manual text simplification (Chandrasekar, 1994; Vickrey and Koller, 2008). Evans (2011) later showed that an automated simplification of coordinate structures can improve IE systems.

Later, the focus of the ATS shifted towards text accessibility and better social inclusion, having the main goal of making texts easier to understand by

various target readers, e.g. people with low literacy levels (Aluísio and Gasperin, 2010), or people with some kind of reading or cognitive impairments, such as aphasia (Devlin and Unthank, 2006), autism (Orásan et al., 2018), Down’s syndrome (Saggion et al., 2015), or dyslexia (Rello, 2012).

In this study, we want to return to the original motivation for text simplification and explore whether the state-of-the-art ‘general purpose’ ATS system can be used to improve machine translation from English to some under-resourced language, or not. Unlike the previous works, we focus on using fully automated TS output (without any manual corrections), and on filtering out simplifications that are not grammatical and do not preserve the original meaning. Furthermore, we perform experiments using two state-of-the-art MT systems with different architectures, a PBMT and a NMT system. We focus on English-to-Serbian machine translation, taking Serbian as an example of under-resourced languages.

We show that, regardless of the MT architecture (PBMT or NMT) and the strategy for using ATS system as a pre-processing step (without any manual correction, with filtering for its grammaticality and meaning preservation, and with minimal manual correction of the output), the fluency of the translation can be improved. With regards to improving translation adequacy and post-editing effort, our experiments show that the type of the translation architecture (PBMT or NMT), and the strategy for using ATS system, both play a significant role.

## 2 Related Work

For many language pairs (e.g. English-French, English-Spanish, English-Hindu), attempts were made at rewriting input sentences using paraphrasing or textual entailment to improve the performance of MT systems (Callison-Burch et al., 2006; Mirkin et al., 2009; Aziz et al., 2010; Tyagi et al., 2015; Mirkin et al., 2013a,b). However, they all focus only on out-of-vocabulary words, or difficult to translate shorter  $n$ -grams.

Štajner and Popović (2016) went one step further, using lexico-syntactic automatic text simplification systems as a pre-processing step for English-to-Serbian machine translation. In this way, they covered both lexical and syntactic transformations on the source side. The ATS outputs

were manually inspected by human editors who were also allowed to do minor revisions (correcting the tense, gender, article, etc.) in order to preserve grammaticality and the original meaning on the source side. For both ATS systems used, it was found that (with this minimal human correction of the simplified output) such a pre-processing step improves fluency of the translations, and reduce the post-editing effort. However, the authors only considered manually post-edited ATS output and made no experiments with the raw (uncorrected) ATS output. Nor did they explore how the grammaticality and meaning preservation of the ATS output might influence the results.

In this work, we use one of the ATS systems used by Štajner and Popović (2016), only the system which does not remove any original information. Unlike Štajner and Popović (2016), who used only a phrase-based MT system for English-to-Serbian translation, we also use the current state-of-the-art neural MT system for that language pair (see Section 3.1). We explore three different scenarios for using ATS as the pre-processing step, in search for fully automatic use of ATS in MT, without human correction of the ATS output (see Section 3). We find that success of the ATS used as a pre-processing step heavily depends on the type of the MT system used (PBMT or NMT).

## 3 Experimental Setup

We randomly selected 10 original articles from the 100 news articles automatically simplified by the state-of-the-art lexico-syntactic ATS system (Siddharthan and Angrosh, 2014) in the work of Štajner and Glavaš (2017). The ATS system that consists of a rule-based syntactic simplification module and a supervised lexical simplification module built upon the English Wikipedia - Simple English Wikipedia corpus (Coster and Kauchak, 2011).

We further explored three possible scenarios in which ATS can be used as a pre-processing step for MT (Figure 1):

- **Scenario 1 (Corrected):** Automatically simplified sentences are manually corrected before being used as the source sentences for MT, to ensure the preservation of the original meaning and the grammaticality of the MT input;
- **Scenario 2 (Filtered):** Automatically simpli-

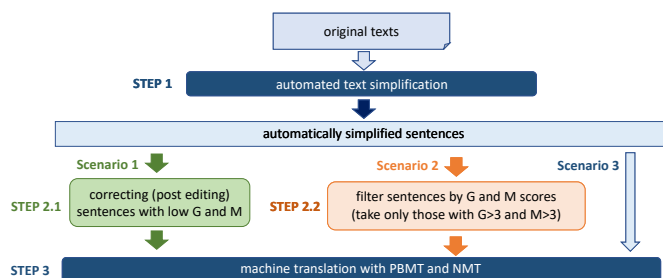


Figure 1: Workflow. The fully-automated steps are shown in dark blue (steps 1 and 3); fully-manual in green (step 2.1); and those that can be used as either automated or manual, in orange (step 2.2).

fied sentences which did not preserve well the original meaning and/or are ungrammatical are filtered out, and in those cases, the original sentences are used instead of them as the MT input;

- **Scenario 3 (Automatic):** Automatically simplified sentences are used as source sentences for MT without any manual correction or filtering beforehand.

The last scenario (*Automatic*) is troublesome in the context of ATS used as pre-processing step for MT, as one cannot be sure that the original meaning was preserved during automatic simplification. To account for possible changes of meaning, we slightly modify the common procedures for assessing fluency and adequacy of MT output in order to allow those scores to penalize such text simplification errors (see Section 3.3).

### 3.1 MT Systems

For English-to-Serbian phrase-based and neural machine translation, we use ASISTENT<sup>1</sup> (Arčan et al., 2016), a publicly available web-based MT system offering translation between three South Slavic languages (Croatian, Serbian and Slovenian) and English in both translation directions. Both NMT and PBMT variants were trained on the publicly available data originating from the OPUS website<sup>2</sup> (Tiedemann, 2009) where three domains were available for the Serbian-English language pair: the enhanced version of the SETimes corpus<sup>3</sup> (Tyers and Alperen, 2010) containing “news

and views from South-East Europe”, OpenSubtitles<sup>4</sup> as well as KDE localisation documents and manuals, i.e. technical domain. In total, about 20.7M sentences were used for training (20.5M subtitles, 200k news, 30k technical), and 2k sentences from each of the three domains were used for tuning the systems.

The texts which we are translating in this study are external, i.e. they cannot be found in any of the above described corpora.

### 3.2 Assessment of Quality of ATS Output

The quality of the ATS output was assessed at the sentence level by:

- human assessment of grammaticality (G) and meaning preservation (M) on a 1–5 Likert scale (1 – very bad; 5 – very good).
- measuring the time needed to correct grammaticality and ensure that original meaning is preserved.

The first assessment was used in Scenario 2 (*Filtered*) for filtering simplified sentences according to their G and M scores, while the second assessment was used in Scenario 1 (*Corrected*). The schema of the workflow is presented in Figure 1.

We asked three native English speakers to rate our 130 sentences using the same guidelines as Štajner and Glavaš (2017). The annotators were also provided with several examples for each score. To obtain the final G and M marks used in Scenario 2 (*Filtered*), we averaged the three marks and rounded the results to the closest integers. The average G and M scores were 3.98 and 3.75, respectively. The average pairwise inter-annotator

<sup>1</sup><http://server1.nlp.insight-centre.org/asistent/>

<sup>2</sup><http://opus.lingfil.uu.se>

<sup>3</sup><http://nlp.ffzg.hr/resources/corpora/setimes/>

<sup>4</sup><http://www.opensubtitles.org/>

Version	Sentence
Original	Ex-Soviet leader Mikhail Gorbachev says Russian authorities must annul the parliamentary vote results and hold a new election.
Simplification (uncorrected)	Ex-Soviet leader Mikhail Gorbachev says. Russian authorities must annul the parliamentary vote results. These authorities hold a new election.
Simplification (corrected)	Ex-Soviet leader Mikhail Gorbachev says <b>that</b> Russian authorities must annul the parliamentary vote results. These authorities <b>must</b> hold a new election.
Original	A 21-year-old man was arrested on April 30, on suspicion of murder and was released on bail until May 29 pending further enquiries.
Simplification (uncorrected)	A 21-year-old man was arrested on April 30, on suspicion of murder. This man was <b>followed</b> until May 29 pending further enquiries.
Simplification (corrected)	A 21-year-old man was arrested on April 30, on suspicion of murder. This man was <b>released</b> until May 29 pending further enquiries.

Table 1: Two examples of manual corrections performed on the simplified sentences. Differences between the automatically simplified sentences and their manually corrected versions are shown in bold.

<i>Adequacy</i>	
5	perfectly understandable (regardless of potential poor grammar)
4	understandable with minor ambiguities/differences
3	main gist is preserved but some things are unclear/different from the source
2	difficult to understand and different from the source meaning
1	very bad (regardless to potential grammaticality)
<i>Fluency</i>	
5	perfectly grammatical (regardless of potential meaning loss/change)
4	almost correct – a small number of minor errors
3	a number of grammatical errors although not very heavy
2	many grammatical errors
1	very bad (regardless to potential meaning preservation)

Table 2: Guidelines for assigning adequacy and fluency scores.

agreement was 0.72 and 0.61 (weighted Cohen’s  $\kappa$ ), for the G and M scores respectively.

In Scenario 1 (*Corrected*), we used the manual corrections already provided by Štajner and Glavaš (2017) for the ten selected articles. During manual corrections, only minimal corrections were performed where necessary to restore the original meaning and grammaticality of the sentences. As the goal of those correction was not to make any further simplifications, and the mistakes were easy to notice, the corrections were very fast (15.0 seconds per sentence). They did not even require a native speaker or trained annotator, but rather someone with a proficiency level of English (Štajner and Glavaš, 2017). Several examples of performed corrections are given in Table 1.

### 3.3 Evaluation of MT Output

All translation outputs (translations of the original sentences, their raw (uncorrected) automatic simplifications, and their manually corrected automatic simplifications) obtained by the two MT systems (PBMT and NMT), a total of 390 target

sentences, were evaluated with respect to three aspects:

- adequacy, i.e. how well the sentence preserves the original meaning;
- fluency, i.e. how grammatical the sentence is;
- technical post-editing effort, i.e. the amount of necessary edit operations to correct the output.

These three evaluation aspects concentrate on three distinct things, which are not necessarily correlated. For example, if the reference translation is “We will **not** finish on time.” and the obtained MT output is “We will finish on time.”, the adequacy score will be very low (1), fluency perfect (5), and edit distance very low (only one edit).

Each of the tasks has been carried out separately, i.e. the evaluation of adequacy and fluency were carried out in two separate passes, and post-editing was carried out in the third pass.

The guidelines used for assigning adequacy and fluency scores are presented in Table 2.




Figure 2: Average edit rates (y-axis) for corrected TS output (blue full) and uncorrected/raw TS output (orange pattern), depending on the grammaticality (G) and meaning preservation (M) scores (x-axis) after PBMT and NMT.

Raw edit counts and edit rates (raw counts normalised with the segment length) were calculated using Hjerson (Popović, 2011) for:

- the five classes of edits/errors: inflectional error, reordering error, omission, addition and mistranslation
- the sum of edit counts/rates of all classes

Each translated segment was post-edited, and assigned fluency and adequacy scores, while looking into the corresponding source segment, i.e. the English originals were used for evaluating the translations of the originals, while the corresponding simplified and corrected English sentences were used for evaluating the translations of the simplified sentences. This was done to ensure that the change of meaning during ATS (in scenarios with filtered and fully automatic ATS) is penalised. Reference translations were not available.

## 4 Results

We first explored the influence of grammaticality (G) and meaning preservation (M) scores of the automatically simplified sentences on the post-editing effort needed to correct their translations

(Figure 2).<sup>5</sup>

We see that the differences in effort needed to post-edit the translations of automatically simplified sentences (uncorrected) and corrected simplified sentences (corrected) decrease with the increase of grammaticality (G) and meaning preservation (M) scores of the uncorrected simplifications. This supports our initial idea that, instead of correcting the automatically simplified sentences (Scenario 1 in Figure 1), one could filter out the automatically simplified sentences which did not achieve high enough G and M scores, and instead of those, use the original sentences (Scenario 2 in Figure 1). This filtering should ideally be done automatically. Given that here we just look for the proof of concept, we wanted to ensure that we correctly assign G and M scores, and this was thus done manually. However, it is important to note that several systems for automatic assignment of G and M scores to the ATS outputs have been proposed up to date (Štajner et al., 2016).

The initial exploration (Figure 2) indicated that a good cut-off point for filtering bad simplifications would be around  $G = 3$  and  $M = 3$ .

<sup>5</sup>For the space constraints, we here present graphs only for the average edit rates, but we also analysed the raw counts and found that they follow the same trends as the average edit rates.

(a) PBMT								
%	uncorrected		corrected		G, M > 3		G, M ≥ 3	
	Σer	r.e.c.	Σer	r.e.c.	Σer	r.e.c.	Σer	r.e.c.
better	27.7	21.5	<b>46.2</b>	<b>38.5</b>	<b>20.0</b>	<b>16.9</b>	24.6	21.5
worse	50.8	49.2	<b>27.6</b>	<b>26.2</b>	<b>13.8</b>	<b>12.3</b>	27.7	53.9
same	21.5	29.3	26.2	35.3	66.2	70.8	47.7	24.6

(b) NMT								
%	uncorrected		corrected		G, M > 3		G, M ≥ 3	
	Σer	r.e.c.	Σer	r.e.c.	Σer	r.e.c.	Σer	r.e.c.
better	13.8	12.3	<b>46.2</b>	<b>36.9</b>	9.2	6.2	13.8	10.8
worse	67.7	64.6	<b>37.5</b>	<b>32.3</b>	27.7	26.2	41.6	40.0
same	18.5	23.1	32.3	30.8	63.1	67.6	44.6	49.2

Table 3: Percentage of raw simplified sentences (uncorrected), corrected simplified sentences (corrected) and simplified sentences filtered by high grammatical and meaning scores ( $G, M > 3$ , and  $G, M \geq 3$ ) with better/worse/same translations (in terms of edit rate ( $\Sigma\text{er}$ ) and raw edit counts (r.e.c.)) than their original counterparts when translated with phrase-based (a) and neural (b) MT system. Those cases in which more sentences improved than deteriorated are shown in bold.

We further investigated how the distribution of sentences with better, worse, and same post-editing effort differs in the two possible cut-offs: (1) where  $G$  and  $M$  scores are both greater than 3; and (2) where  $G$  and  $M$  scores are both greater or equal to 3. The results are presented in Table 3, together with the corresponding results for the fully automated simplification without any filtering (uncorrected) and for the manually corrected simplification output (corrected).

It seems that filtering automatically simplified sentences according to their grammaticality ( $G$ ) and meaning preservation ( $M$ ) score can substantially decrease the percentage of sentences whose translation is worse (needs more post editing) than the translation of their original counterparts, if the cut-off point is correctly set. This happens in both MT approaches used (PBMT or NMT), but the decrease is more pronounced in PBMT.

Interestingly, in PBMT, filtering sentences according to their  $G$  and  $M$  scores (with  $G, M > 3$ ) results in a lower percentage of sentences with deteriorated translations than correcting the simplification output before translation (corrected). This happens at the cost of increasing the number of sentences with the same MT post-editing effort required, and decreasing the number of sentences with improved translation.

Table 4 shows the differences in fluency and adequacy scores between the translations of the original sentences and the translations of the corrected

simplified texts, uncorrected simplified texts, and the filtered ( $G, M > 3$ ) source sentences.

It can be seen that by using the ATS system in a pre-processing step (marginally) improves the translation adequacy, and only if the simplifications are manually corrected and used in the NMT system.

Using any of the three proposed scenarios (manually corrected simplifications, filtered simplifications, or fully automated simplifications) leads to a higher percentage of sentences with improved rather than those with deteriorated fluency of translation. Interestingly, the percentage of improved sentences is the highest when the uncorrected simplifications are used, but at the cost of the higher percentage of sentences with deteriorated translation adequacy (than in the case of using the manually corrected simplifications). Filtering automatic simplifications according to their grammaticality and meaning preservation substantially decreases the percentage of sentences with both improved and deteriorated translation fluency, but still results in substantially higher percentage of sentences with improved than those with deteriorated fluency.

Several examples of the original English sentences and their simplified versions, together with the scores for the fluency and adequacy of their Serbian translations are presented in Table 5.

(a) PBMT

Difference	Adequacy			Fluency		
	corrected	uncorrected	G, M > 3	corrected	uncorrected	G, M > 3
-3	0	1.5	0	0	0	0
-2	3.1	16.9	9.2	0	0	0
-1	21.5	53.8	30.8	15.4	20.0	10.8
0	60.0	23.1	56.9	61.5	53.8	76.9
1	15.4	4.6	3.1	21.5	34.6	12.3
2	0	0	0	1.5	1.5	0
3	0	0	0	0	0	0
$\Sigma$ positive	15.4	4.6	3.1	<b>23.0</b>	<b>36.1</b>	<b>12.3</b>
$\Sigma$ negative	24.6	81.2	40.0	15.4	20.0	10.8

(b) NMT

Difference	Adequacy			Fluency		
	corrected	uncorrected	G, M > 3	corrected	uncorrected	G, M > 3
-3	1.5	3.1	0	0	0	0
-2	4.6	13.8	6.2	0	3.1	0
-1	20.0	33.8	13.8	13.8	18.5	7.7
0	47.7	38.5	72.3	63.1	52.3	78.5
1	23.1	9.2	6.2	21.5	23.1	12.3
2	3.1	1.5	1.5	0	1.5	0
3	0	0	0	1.5	1.5	1.5
$\Sigma$ positive	<b>26.2</b>	10.7	7.7	<b>23.0</b>	<b>26.1</b>	<b>13.8</b>
$\Sigma$ negative	26.1	50.7	20.0	13.8	21.6	7.7

Table 4: Distribution of differences in adequacy and fluency scores (in terms of percentages) introduced by the manually corrected simplifications, uncorrected simplifications, and by filtering automatic simplifications by their meaning preservation and grammaticality scores.

## 5 Summary and Outlook

Going back to the initial motivation for automatic text simplification, using it as a pre-processing step to improve the performance of machine translation systems, we explored how the current state-of-the-art lexico-syntactic automatic text simplification system behaves in this role. We investigated three possible scenarios: (1) using the output of the ATS system as it is; (2) filtering out the automatic simplifications with low grammaticality and meaning preservation scores and using the original sentences instead; (3) minimally correcting the ATS output in order to preserve the original meaning and the grammaticality of the sentence.

The results of our experiments indicated that:

- The success of the ATS systems depends on the type of MT system used (phrase-based or neural).
- In the case of NMT, only the manually corrected ATS output can reduce the post-editing

effort, and improve the adequacy scores in translation.

- In the case of PBMT, two scenarios (manual correction of ATS output and maintaining only automatic simplifications with high G and M scores) can reduce the post-editing effort in translation, but none of the three investigated scenarios can improve the adequacy of the translation.
- For both MT approaches, PBMT and NMT, all three scenarios for using ATS improve the fluency of the translation noticeably.
- The uncorrected ATS output improves the fluency of the PBMT translation noticeably more than it improves the fluency of the NMT translation.

We also found that even manually corrected ATS output can deteriorate translation adequacy (in about 15% of the cases) and fluency (in about

## (a) PBMT

Ex.	Version	r.e.c/ $\Sigma$ er/A/F	Sentence
1	Original	21/75.0/2/2	The cabinet is also expected to demand banks to set aside a further 35 billion euros (27.9 billion pounds) to cover sound loans in their real estate portfolios.
	Uncorrected (G=4, M=5)	18/58.7/1/2	The cabinet is also expected to demand banks to set aside a further 35 billion euros, to cover sound loans in their real estate portfolios. <b>These</b> is 27.9 billion pounds.
	Corrected	16/50.0/2/2	The cabinet is also expected to demand banks to set aside a further 35 billion euros, to cover sound loans in their real estate portfolios. <b>35 billion</b> euros is 27.9 billion pounds.
2	Original	9/56.7/2/2	The toxic assets now total 184 billion euros, but many fear the hole is even bigger.
	Uncorrected (G=4, M=5)	8/50.8/3/3	The toxic assets now total 184 billion euros. But many fear the hole is even bigger.
	Corrected	7/43.6/4/4	The toxic assets now total 184 billion euros but many fear the hole is even bigger.

## (b) NMT

Ex.	Version	r.e.c/e.r./F/A	Sentence
1	Original	5/28.3/5/4	Jeffrey Burrows was killed at his home in Norfolk Square in Brighton, East Sussex, on April 29.
	Uncorrected (G=5, M=3)	3/14.1/4/4	Jeffrey Burrows was killed at his home in Norfolk Square in Brighton, on April 29. Brighton is East Sussex.
	Corrected	2/9.1/5/4	Jeffrey Burrows was killed at his home in Norfolk Square in Brighton, on April 29. Brighton is <b>in</b> East Sussex.
2	Original	3/14.3/3/3	Spanish stocks were down 2.1 percent on Friday morning, in line with other European markets, after getting a big boost on Thursday from the banking reform plans.
	Uncorrected (G=5, M=5)	3/14.3/4/4	Spanish stocks were down 2.1 percent on Friday morning, in line with other European markets. This happened after getting a big boost on Thursday from the banking reform plans.

Table 5: Examples of original sentences, automatically simplified sentences (uncorrected), and their manually corrected versions, together with corresponding translation scores in terms of edit operations, adequacy and fluency scores.

25% of the cases), and increases the post-editing effort in translation (in about 26-27% of the cases in PBMT, and about 32-37% of the cases in NMT). This indicates that the general-purpose ATS systems (which were not initially developed for improving MT performances) might not be suitable for this task. Depending on the source-target language combination, it might be better to design a MT-oriented text simplification system, which would target only the sentence structures and vo-

cabulary which poses particular difficulties in that language pair and MT approach.

Our results also showed that there are notable differences in how the three scenarios for using ATS as a pre-processing step influence the performances of the PBMT and NMT (at least for the language pair investigated here). This indicates that the two MT approaches require different pre-processing strategies in order to improve their performances.



## 6 Acknowledgements

The ADAPT Centre for Digital Content Technology ([www.adaptcentre.ie](http://www.adaptcentre.ie)) at Dublin City University is funded under the Science Foundation Ireland Research Centres Programme (Grant 13/RC/2106) and is co-funded under the European Regional Development Fund.

## References

- Sandra Maria Aluísio and Caroline Gasperin. 2010. Fostering digital inclusion and accessibility: The porsimples project for simplification of portuguese texts. In *Proceedings of YIWICALA Workshop at NAACL HLT 2010*. pages 46–53.
- Mihael Arčan, Maja Popović, and Paul Buitelaar. 2016. Asistent – A Machine Translation System for Slovene, Serbian and Croatian. In *Proceedings of the Conference on Language Technologies and Digital Humanities*. Ljubljana, Slovenia, pages 13–20.
- Wilker Aziz, Marc Dymetman, Shachar Mirkin, Lucia Specia, Nicola Cancedda, and Ido Dagan. 2010. Learning an expert from human annotations in statistical machine translation: the case of out-of-vocabulary words. In *Proceedings of EAMT*.
- Chris Callison-Burch, Philipp Koehn, and Miles Osborne. 2006. Improved statistical machine translation using paraphrases. In *Proceedings of the 2016 Human Language Technology Conference of the North American Chapter of the ACL (HLT-NAACL)*. pages 17–24.
- R. Chandrasekar. 1994. *A Hybrid Approach to Machine Translation using Man Machine Communication*. Ph.D. thesis, Tata Institute of Fundamental Research/University of Bombay, Bombay.
- Raman Chandrasekar, Christine Doran, and B. Srinivas. 1996. Motivations and Methods for Text Simplification. In *Proceedings of COLING 1996*. pages 1041–1044.
- William Coster and David Kauchak. 2011. Simple English Wikipedia: a new text simplification task. In *Proceedings of ACL&HLT*. pages 665–669.
- Siobhan Devlin and Gary Unthank. 2006. Helping aphasic people process online information. In *Proceedings of the 8th international ACM SIGACCESS conference on Computers and accessibility*. ACM, New York, NY, USA, Assets '06, pages 225–226.
- Richard J. Evans. 2011. Comparing methods for the syntactic simplification of sentences in information extraction. *Literary and Linguistic Computing* 26(4):371–388.
- Shachar Mirkin, Lucia Specia, Nicola Cancedda, Ido Dagan, Marc Dymetman, and Idan Szpektor. 2009. Source-language entailment modeling for translating unknown terms. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2 - Volume 2*. Association for Computational Linguistics, Stroudsburg, PA, USA, ACL '09, pages 791–799.
- Shachar Mirkin, Sriram Venkatapathy, and Marc Dymetman. 2013a. Confidence-driven Rewriting for Improved Translation. In *Proceedings of the XIV MT Summit, Nice, France*. pages 257–264.
- Shachar Mirkin, Sriram Venkatapathy, Marc Dymetman, and Ioan Calapodescu. 2013b. SORT: An Interactive Source-Rewriting Tool for Improved Translation. In *Proceedings of ACL, Sofia, Bulgaria*. pages 85–90.
- Constantin Orăsan, Richard Evans, and Ruslan Mitkov. 2018. *Intelligent Natural Language Processing: Trends and Applications*, Springer, chapter Intelligent Text Processing to Help Readers with Autism, pages 287–312.
- Maja Popović. 2011. Hjerson: An Open Source Tool for Automatic Error Classification of Machine Translation Output. *The Prague Bulletin of Mathematical Linguistics* 96:59–68.
- Maja Popović. 2017. Comparing Language Related Issues for NMT and PBMT between German and English. *The Prague Bulletin of Mathematical Linguistics* 108(1):209–220.
- Maja Popović and Mihael Arčan. 2015. Identifying Main Obstacles for Statistical Machine Translation of Morphologically Rich South Slavic languages. In *Proceedings of the 18th Annual Conference of the European Association for Machine Translation (EAMT 2015)*. Antalya, Turkey.
- Luz Rello. 2012. Dyswebxia: a model to improve accessibility of the textual web for dyslexic users. In *SIGACCESS Access. Comput.*, ACM, New York, NY, USA, 102, pages 41–44.
- Horacio Saggion, Sanja Štajner, Stefan Bott, Simon Mille, Luz Rello, and Biljana Drndarevic. 2015. Making It Simplext: Implementation and Evaluation of a Text Simplification System for Spanish. *ACM Transactions on Accessible Computing* 6(4):14:1–14:36.
- Advait Siddharthan and M. A. Angrosh. 2014. Hybrid text simplification using synchronous dependency grammars with hand-written and automatically harvested rules. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*. pages 722–731.
- Jorg Tiedemann. 2009. News from OPUS – A Collection of Multilingual Parallel Corpora with Tools and Interfaces. In *Advances in Natural Language Processing*, Borovets, Bulgaria, volume V, chapter V, pages 237–248.

- Shruti Tyagi, Deepti Chopra, and Iti Mathur. 2015. Classifier based text simplification for improved machine translation. In *Proceedings of International Conference on Advances in Computer Engineering and Applications (ICACEA)*, Ghaziabad, India, pages 46–50.
- Francis M. Tyers and Murat Alperen. 2010. South-East European Times: A parallel corpus of the Balkan languages. In *Proceedings of the LREC Workshop on Exploitation of Multilingual Resources and Tools for Central and (South-) Eastern European Languages*. Valetta, Malta, pages 49–53.
- David Vickrey and Daphne Koller. 2008. Sentence simplification for semantic role labeling. In *Proceedings of ACL&HLT*, volume 344–352.
- Sanja Štajner and Goran Glavaš. 2017. Leveraging event-based semantics for automated text simplification. *Expert Systems with Applications* 82:383–395.
- Sanja Štajner and Maja Popović. 2016. Can text simplification help machine translation? In *Proceedings of the 19th Annual Conference of the European Association for Machine Translation (EAMT 2016)*. Riga, Latvia, pages 230–242.
- Sanja Štajner, Maja Popović, Horacio Saggion, Lucia Specia, and Mark Fishel. 2016. Shared Task on Quality Assessment for Text Simplification. In *Proceedings of the LREC Workshop on Quality Assessment for Text Simplification (QATS)*, pages 22–31.