Procesamiento del Lenguaje Natural, Revista nº 53, septiembre de 2014, pp 61-68 recibido 17-04-14 revisado 06-06-14 aceptado 10-06-14

Translating sentences from 'original' to 'simplified' Spanish

Traducción de frases del español 'original' al español 'simplificado'

Sanja Štajner
Research Group in Computational Linguistics
Research Institute in Information and Language Processing
University of Wolverhampton, UK
sanjastajner@wlv.ac.uk

Resumen: La Simplificación de Textos (ST) tiene como objetivo la conversión de oraciones complejas en variantes más sencillas, que serían más accesibles para un público más amplio. Algunos estudios recientes han abordado este problema como un problema de traducción automática (TA) monolingüe (traducir de lengua 'original' a 'simplificada' en lugar de traducir de un idioma a otro), utilizando el modelo estándar de traducción automática basado en frases. En este estudio, investigamos si el mismo enfoque tendría el mismo éxito independientemente del tipo de simplificación que se quiera estudiar, dado que cada público meta require diferentes niveles de simplificación. Nuestros resultados preliminares indican que el modelo estándar podría no ser capaz de aprender las fuertes simplificaciones que se necesitan para algunos usuarios, e.g. gente con el síndrome de Down. Además, mostramos que las tablas de traducción obtenidas durante el proceso de traducción parecen ser capaces de capturar algunas simplificaciones léxicas adecuadas.
Palabras clave: simplificación de textos, traducción automática estadística

Abstract: Text Simplification (TS) aims to convert complex sentences into their simpler variants, which are more accessible to wider audiences. Several recent studies addressed this problem as a monolingual machine translation (MT) problem (translating from 'original' to 'simplified' language instead of translating from one language into another) using the standard phrase-based statistical machine translation (PB-SMT) model. We investigate whether the same approach would be equally successful regardless of the type of simplification we wish to learn (given that different target audiences require different levels of simplification). Our preliminary results indicate that the standard PB-SMT model might not be able to learn the strong simplifications which are needed for certain users, e.g. people with Down's syndrome. Additionally, we show that the phrase-tables obtained during the translation process seem to be able to capture some adequate lexical simplifications.
Keywords: text simplification, phrase-based statistical machine translation

1 Introduction

Since the late nineties, several initiatives raised awareness of the complexity of the vast majority of written documents and the difficulties they pose to people with any kind of reading or learning impairments. These initiatives proposed various guidelines for writing in a simple and easy-to-read language which would be equally accessible to everyone. However, manual adaptation of existing documents could not keep up with the everyday production of material written in a 'complex' language. This motivated the need for automatic Text Simplification (TS), which aims to convert complex sentences into their simpler variants, while preserving the original meaning.

The first TS systems were traditionally rule-based, e.g. (Devlin, 1999; Canning et al., 2000), requiring a great number of hand-crafted simplification rules produced by highly specialised people. They consisted of

ISSN 1135-5948

two main simplification modules – lexical and syntactic. The lexical simplification module replaces long and uncommon words with their shorter and more commonly used synonyms. The syntactic simplification module recursively applies a set of handcrafted rules to each sentence as long as there are any rules which can be applied. The main drawbacks of those systems are that such rules cannot be easily adapted to different languages or genres, and that they lead to TS systems with high precision and low recall. With the emergence of Simple English Wikipedia (SEW)[1], which together with the 'original' English Wikipedia (EW)[2] provided a large parallel corpus for TS, some new machine learning oriented approaches have appeared. Several recent studies addressed text simplification as a monolingual machine translation (MT) problem. Instead of translating from one language to another, they tried to translate from the 'original' to the 'simplified' language.

In this paper, we explore the influence of the level of simplification in the training dataset on the performance of a phrase-based statistical machine translation (PB-SMT) model which tries to translate from 'original' to 'simplified' Spanish. Our preliminary results indicate that PB-SMT systems might not be appropriate when the training set contains a great number of 'strong' simplifications (which are needed for some target populations such as people with Down's syndrome for example), while they might work reasonably well when trained on the datasets which contain only the 'weak' simplifications (which are sufficient for some other target populations such as non-native speakers or people with low literacy levels). Additionally, we show that the phrase-based tables produced during the translation process contain a great number of adequate lexical paraphrases which could be used to build a separate lexical simplification module if necessary.

The remainder of the paper is structured as follows: Section 2 presents the related work on text simplification with a special emphasis on previous uses of PB-SMT systems in TS; Section 3 describes the corpora which were used and the experiments conducted; Section 4 presents the performances of the two PB-SMT systems trained on different corpora, and discusses the possibilities for using the phrase-based tables produced during the translation process; Section 5 lists the main findings and gives directions for future work.

2 Related Work

Due to the lack of large parallel corpora of original and simplified texts, many of the recent TS systems are still rule-based, e.g. (Saggion et al., 2011; Drndarević et al., 2013; Orasan, Evans, and Dornescu, 2013). However, the number of machine learning (ML) approaches to TS has increased in the last few years. This increase is especially pronounced in English TS, due to the large and freely available parallel corpus of original and simplified texts – English Wikipedia (EW) and Simple English Wikipedia (SEW). Napoles and Drezde (2010) built a statistical classification system that can distinguish which version of English Wikipedia a text belongs to, thus confirming the possibility of using those corpora in TS. Yatskar et al. (2010) used edit histories in SEW to extract lexical simplifications, and Biran et al. (2011) applied an unsupervised method for learning pairs of complex and simple synonyms from the EW and SEW. Zhu et al. (2010) proposed a tree-based simplification model, while Woodsend and Lapata (2011) used quasi-synchronous grammar to learn a wide range of rewriting transformations for TS.

Several recent studies addressed the TS as a monolingual MT problem. Instead of translating from one language to another, they tried to translate from the 'original' to the 'simplified' language. Coster and Kauchak (2011) applied the standard PB-SMT model implemented in Moses toolkit to 137,000 sentence pairs from the EW and SEW. They also suggested an extension of that model, which adds phrasal deletion to the probabilistic translation model in order to better cover deletion, which is a frequent phenomenon in TS. The obtained results (BLEU = 59.87 on the standard model without phrasal deletion, and BLEU = 60.46 on the extended model) were promising, although not far from the baseline (no translation performed), thus suggesting that the system is overcautious in performing simplifications. In order to overcome this issue, Wubben et al. (2012) performed post-hoc re-ranking on the output

[1]http://simple.wikipedia.org/wiki/Main_Page
[2]http://en.wikipedia.org/wiki/Main_Page

Version	Example
Original	Ahora se amplía, aunque siempre según el parecer del juez, a conducir con un exceso de velocidad superior en 60 kilómetros por hora en vía urbana o en 80 kilómetros por hora en vía interurbana, o conducir bajo la influencia de las drogas o con una tasa de alcohol superior a 1,2 gramos por litro en sangre.
Weak	*Esta medida* se amplía, *dependiendo del juez*, a conducir con un exceso de velocidad mayor de 60 kilómetros por hora en vía urbana o de 80 kilómetros por hora en vía interurbana, o conducir drogado o con una tasa de alcohol mayor a 1,2 gramos por litro en sangre.
Strong	Ahora *los jueces también podrán quitar el coche a las personas condenadas por otras causas. Algunas causas son conducir muy rápido dentro de las ciudades o beber alcohol o tomar drogas antes de conducir.*
Original	El fallo definitivo con la ciudad ganadora del concurso se conocerá el próximo 3 de diciembre de 2010, fecha en la que se celebra el Día Internacional y Europeo de las Personas con Discapacidad.
Weak	*La decisión definitiva* con la ciudad ganadora del concursó se *sabrá* el próximo 3 de diciembre de 2010. *El 3 de Diciembre es* el Día Internacional y Europeo de las Personas con Discapacidad.
Strong	*El premio se entregará el* 3 de diciembre de 2010. *El 3 de diciembre es* el Día Internacional y Europeo de las Personas con Discapacidad.

Table 1: Weak vs strong simplification (deviations from the original sentence are shown in italics)

(simplification hypotheses) based on their dissimilarity to the input (original sentences), i.e. they selected the output that is as different as possible from the original sentence.

Specia (2010) used the standard PB-SMT model implemented in Moses toolkit to try to learn how to simplify sentences in Brazilian Portuguese. She used 4,483 original sentences and their corresponding 'natural' simplifications obtained under the PorSimples project (Gasperin et al., 2009). The project was aimed at people with low literacy levels and the newswire texts were simplified manually by a trained human editor, offering two levels of simplification: 'natural' and 'strong'. Specia (2010) used only the sentence pairs obtained by 'natural' simplification (where the most common simplification operation was lexical substitution), which would correspond to our 'weak' simplification in Spanish. The performance of the translation model was reasonably good – BLEU score of 60.75 – especially taking into account the relatively small size of the corpora (4,483 sentence pairs).

3 Methodology

The main goal of this study was to investigate how far the level of simplification present in the training dataset influences the performance of a PB-SMT system which tries to learn how to translate from 'original' to 'simplified' language. Therefore, we trained the standard PB-SMT system on two TS corpora in Spanish, which contained different levels of simplification (were targeted to different users and were thus compiled using different simplification guidelines). The corpora and the experimental settings are described in the next two sub-sections.

3.1 Corpora

The first corpus (*Strong* simplification) was compiled under the Simplext project[3], following detailed easy-to-read guidelines prepared especially for simplifying texts for readers with Down's syndrome. The 200 newswire texts were simplified manually by the trained human editors. Many sentences required a very high level of simplification (given the specific needs of the target population), as can be observed in Table 1.

The second corpus (*Weak* simplification) was created by three native speakers of Spanish, following the given guidelines with no concrete target population in mind. The guidelines consisted of the same main simplification rules (e.g. use simple sentences, use common words, remove redundant words, use a simpler paraphrase if applicable) as those present in the Simplext guidelines. This time, the editors were explicitly instructed not to use strong paraphrases, i.e. to limit the use of the 'use simpler paraphrase, if applicable' rule to a minimum and not to apply it to the whole sentence but rather only to a specific (short) part of the sentence.

[3]www.simplext.es

The differences in the simplifications obtained by the aforementioned two simplification strategies (*strong* and *weak*) are presented in Table 1. The corpora characteristics: the average number of words per sentence in both original and simplified corpora, and the average sentence-wise BLEU score (S-BLEU) of the sentence pairs (original sentence and its corresponding manually simplified version) for each corpus are presented in Table 2.

Corpus	ASL-O	ASL-S	S-BLEU
Strong	31.82	14.30	0.17
Weak	25.98	16.91	0.60

Table 2: Corpora characteristics: the average number of words per sentence in the original (*ASL-O*) and the simplified corpora (*ASL-S*), and the average sentence-wise BLEU score (*S-BLEU*)

BLEU (Papineni et al., 2002) evaluates MT output by using exact n-gram matching between the hypothesis and the reference translation. Additionally, it applies the brevity penalty which penalises the hypothesis (automatically simplified sentences, in our case) which are shorter than the reference translations (original sentences, in our case). As BLEU is designed to evaluate output on a document level, it is not ideal for sentence-level scoring. Instead, we use S-BLEU (sentence-level BLEU) to evaluate the sentence pairs. Unlike BLEU, S-BLEU will still positively score segments that do not have higher n-gram matching. The low average S-BLEU score on the training dataset (Table 2) suggests that there are many string transformations and strong paraphrases to be learnt, and thus the standard phrase-based translation model might not be the most suitable for the task.

3.2 Experiments

For the translation experiments, we used the standard PB-SMT system implemented in the Moses toolkit (Koehn et al., 2007), the GIZA++ implementation of IBM word alignment model 4 (Och and Ney, 2003), and the refinement and phrase-extraction heuristics described further in (Koehn, Och, and Marcu, 2003). The systems were tuned using minimum error rate training (MERT) (Och, 2003). The Spanish Europarl corpus[4] (portion of 500,000 sentences) was used to build the 3-gram language model with Kneser-Ney smoothing trained with SRILM (Stolcke, 2002). The stack size was limited to 500 hypotheses during decoding.

Both experiments were conducted on exactly the same amount of data: 700 sentence pairs for training and 100 sentence pairs for development. The obtained translation models were evaluated on three test sets: (1) 50 sentence pairs randomly selected from the corpora with strong simplifications (*Test-s*), (2) 50 sentence pairs randomly selected from the corpora with weak simplifications (*Test-w*), and (3) the mixed dataset which contained 100 sentence pairs from the previous two test sets (*Test-m*). In all cases, the sentence pairs used for testing were different from those used for training and development.

4 Results and Discussion

The results of the two translation experiments are presented in Table 3.

Corpus	Test-s	Test-w	Test-m
Strong	0.0937	0.3944	0.2609
Weak	0.0930	0.4638	0.2996
Original	0.0939	0.5282	0.3187

Table 3: Results of the translation experiments (BLEU scores) on the 50 sentences with the strong simplification (*Test-s*), on the 50 sentences with weak simplification (*Test-w*), and on the 100 sentences mixed test set (*Test-m*); the *Original* corpus corresponds to the baseline when no simplification/translation is performed

The BLEU score results on the *Test-s* presented in Table 3 should be taken with caution, given the specificities of the strong simplifications present in that corpus (*Strong*). The 'gold standard' manually simplified versions of the sentences in this corpus are already very different from their corresponding originals (See the BLEU score for the *Original* corpus in Table 3, which corresponds to the baseline when no simplification/translation is performed). Therefore, the same (and very low) BLEU scores obtained by both experiments (*Strong* and *Weak*) on the test set from the corpus with

[4]http://www.statmg.org/europarl/

Ex.	Type	Sentence
1-s	Original	"Amnistía denuncia que el Gobierno egipcio ha detenido a 150 **opositores** antes de las elecciones."
	Automatic-weak	"Amnistía denuncia que el Gobierno egipcio ha detenido a 150 **contrarios** antes de las elecciones."
	Automatic-strong	"Amnistía *Internacional* que el Gobierno egipcio ha detenido a 150 opositores antes de las elecciones."
	Manual-test	"El Gobierno de Egipto detiene a personas de el partido de la oposición antes de las elecciones. La organización Amnistía Internacional ha denunciado que el Gobierno egipcio ha detenido a 150 personas de la oposición."
2-w	Original	"ÁVILA Y BARCELONA, ENTRE LAS CUATRO FINALISTAS DEL PREMIO CIUDAD EUROPEA ACCESIBLE."
	Automatic-weak	"ÁVILA Y BARCELONA **ESTÁN** ENTRE LAS CUATRO FINALISTAS DEL PREMIO CIUDAD EUROPEA ACCESIBLE."
	Automatic-strong	"ÁVILA y BARCELONA, ENTRE las CUATRO FINALISTAS *de el* CIUDAD EUROPEA ACCESIBLE."
	Manual-test	"ÁVILA Y BARCELONA, ENTRE LAS CUATRO ELEGIDAS PARA GANAR EL PREMIO CIUDAD EUROPEA ACCESIBLE."
3-w	Original	"Ahora se amplía, **aunque siempre según el parecer del juez**, a conducir con un exceso de velocidad **superior** en 60 kilómetros por hora en vía urbana o en 80 kilómetros por hora en vía interurbana, o conducir bajo la influencia de las drogas o con una tasa de alcohol superior a 1,2 gramos por litro en sangre."
	Automatic-weak	"Ahora se amplía, **dependiendo del juez** a conducir con un exceso de velocidad **mayor** de 60 kilómetros por hora en vía urbana o en 80 kilómetros por hora en vía interurbana o conducir bajo la influencia de las drogas o con una tasa de alcohol superior a 1,2 gramos por litro en sangre."
	Automatic-strong	"Ahora se amplía, aunque siempre *en* el parecer del juez, a conducir con un exceso de velocidad superior en 60 por *su helicóptero* en vía urbana, en un 80 por *helicóptero en, por vía* interurbana, conducir bajo la influencia de las drogas*,* con *un tipo* de alcohol superior a 1,2 gramos por litro en sangre."
	Manual-test	"Con la reforma del Código Penal la pérdida del vehículo se amplía a conducir con un exceso de velocidad superior en 60 kilómetros por hora en vía urbana o en 80 kilómetros por hora en vía interurbana, o conducir bajo la influencia de las drogas o con una tasa de alcohol superior a 1,2 gramos por litro en sangre."
4-w	Original	"Ana Juan **fue galardonada** con el Premio Nacional de Ilustración correspondiente a 2010, por el conjunto de su obra."
	Automatic-weak	"Ana Juan **recibió** el Premio Nacional de Ilustración correspondiente a 2010, por el conjunto de su obra."
	Automatic-strong	"Ana Juan fue galardonada con el Premio Nacional de Ilustración correspondiente a 2010, por el *que el leído*."
	Manual-test	"Ana Juan ganó el Premio Nacional de Ilustración de 2010 por el conjunto de la obra de Ana Juan."

Table 4: Automatic simplification obtained by training the PB-SMT system on two different datasets – the one containing strong simplifications (*Automatic-strong)*, and the other containing weak simplifications (*Automatic-weak*). Differences to the original sentence are shown in italics and bold, where the good replacements are shown in bold and the bad ones in italics. *Manual-test* contains the 'gold standard' manual simplification from the test set.

Ex.	Source phrase	Target phrase	p
1	educar bien resulta cansado	enseñar bien es mucho trabajo	0.50
2	educar bien resulta cansado	educar bien es cansado	0.50
3	sublevación	rebelión	0.50
4	sublevación	sublevación	0.50
5	subrayaron que la edad media de inicio	dijeron que la edad media de inicio	0.67
6	subrayaron que la edad media de inicio	indicaron que la edad media de inicio	0.33
7	sufrieron síndrome de inmersión	Las personas tenían síndrome de inmersión	1.00
8	través del cine	mediante el cine	0.75
9	través del cine	través del cine	0.25
10	través del cine y exponer su heterogeneidad	mediante el cine y mostrar su diversidad	0.50
11	través del cine y exponer su heterogeneidad	través del cine y exponer su heterogeneidad	0.50

Table 5: Examples of *source* and *target* phrases and their 'target given source' probabilities (p) in the phrase-tables produced from the training dataset with weak simplifications

strong simplifications (*Test-s*) does not necessarily mean that both systems are equally unsuccessful. Those results only indicate that the obtained automatic simplifications are very different from the 'gold standard' (which was expected given that no automatic simplification could propose such strong paraphrases as those present in that corpus), but not necessarily bad. However, the manual inspection of the automatically simplified sentences revealed that the output of the system trained on the corpus with *strong* simplifications is barely readable and is not able to learn any adequate simplifications. On the contrary, it only worsens the original sentences by making them ungrammatical and/or changing their meaning (see examples in Table 4). On the other hand, the output of the system trained on the corpus with *weak* simplifications was grammatical and in most of the cases it contained at least one adequate lexical simplification (see examples in Table 4). However, it seems that the system was overcautious in applying any transformations, and thus the output of the system did not differ much from the original sentences. Nevertheless, the automatically simplified sentences obtained by this system were as grammatical and usually less complex than the originals.

4.1 Additional experiment

Given the notable similarity of our 'weak' simplifications with the 'natural' simplifications used in (Specia, 2010), we performed an additional experiment. We randomly selected only a portion of the corpus used in (Specia, 2010) – 741 sentence pairs for training, 94 for development and 90 for testing – and performed the same translation experiment as for our two Spanish corpora (using the same setup in the Moses toolkit, but this time using the Lácio-Web corpus[5] for the LM). The average S-BLEU score in this portion of Brazilian Portuguese corpora was 0.58, thus very similar to the one obtained on our Spanish corpus with weak simplifications (Table 2). The obtained BLEU score on the test set for Brazilian Portuguese was 0.5143, while the baseline (no simplification) was 0.5747. These results are again comparable to those obtained on our Spanish corpus with weak simplifications (Table 3).

4.2 Phrase-tables

We additionally examined the phrase-tables produced from the training dataset with *weak* simplifications. We observed many examples of identical source and target phrases with high probabilities. However, the phrase-tables contained a great number of adequate lexical simplifications and simple rewritings (Table 5). While the phrase-tables also provided many examples of bad lexical substitutions, most of them had a very low probabilities. These substitutions were thus discarded in the later stages by either the translation model or the language model.

In many cases, the probability score of the phrases which remain unchanged in the source and target was equal to or higher than the probability of the target phrase which is an adequate simplification of the source phrase (see examples 3 and 4, and 10 and 11 in Table 5). This might be one of the main reasons for the system being overcautious in applying any transformations. If this is the

[5]http://www.nilc.icmc.usp.bf/lacioweb/

case, the translation model could be modified in the way that it forces the system to pick the target phrase which is different from the source phrase whenever the probability of such a translation is higher than some carefully selected threshold.

Alternatively, the phrase-tables obtained during the translation process could be used to build an independent lexical simplification module. Such a module would go beyond the one word substitution level, offering lexical simplification for any phrase which consists of up to seven words (the default configuration in the Moses toolkit builds phrases with up to seven tokens). However, given the small size of the training sets, this approach would suffer from the sparseness problem. It would, therefore, need to be combined with a traditional lexical simplification module which would be used in cases when the 'complex' phrase cannot be found in the phrase-table.

5 Conclusions and Future Work

Text simplification has recently been treated as a statistical machine translation problem and addressed by using the standard phrase-based SMT models. Motivated by the fact that different target populations need different types of simplification, we investigated how much the level of simplification present in the training datasets influences the success of such a TS system.

It appears that a PB-SMT model works reasonably well only when the training dataset does not contain a great number of strong simplifications. Our results indicate that such translation models should not be used when we wish to learn strong simplifications which are needed for some specific audiences, e.g. people with Down's syndrome. Given the very small size of the training datasets used in this study, the reported results should only be regarded as preliminary. To the best of our knowledge, there are no other parallel corpora consisting of original and manually simplified texts in Spanish which could be used to enlarge our training datasets. Therefore, we cannot completely rule out the possibility that the PB-SMT systems would not reach some reasonably good performance if trained on much larger datasets.

The phrase-tables produced during the translation process open two possible avenues for future research. First, it would be interesting to explore how much we could improve the performance of the PB-SMT system if we force it to use the target phrases which are different from the source ones whenever the probability of such a translation is higher than some carefully selected threshold. Second, we could build an independent lexical simplification module based on the information contained in the phrase-tables. Such a lexical simplification module would go beyond performing the substitutions on the word level, offering lexical simplifications for phrases which consists of up to seven words.

Acknowledgements

I would like to express my gratitude to Profs. Ruslan Mitkov and Horacio Saggion, my director of studies and co-supervisor, for their input and help with the resources, as well as to the reviewers for their valuable comments and suggestions.

References

Biran, O., S. Brody, and N. Elhadad. 2011. Putting it Simply: a Context-Aware Approach to Lexical Simplification. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 496–501, Portland, Oregon, USA. Association for Computational Linguistics.

Canning, Y., J. Tait, J. Archibald, and R. Crawley. 2000. Cohesive generation of syntactically simplified newspaper text. In *Proceedings of the Third International Workshop on Text, Speech and Dialogue*, TDS '00, pages 145–150, London, UK, UK. Springer-Verlag.

Coster, W. and D. Kauchak. 2011. Learning to Simplify Sentences Using Wikipedia. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, pages 1–9.

Devlin, S. 1999. *Simplifying natural language text for aphasic readers.* Ph.D. thesis, University of Sunderland, UK.

Drndarević, B., S. Štajner, S. Bott, S. Bautista, and H. Saggion. 2013. Automatic Text Simplication in Spanish: A Comparative Evaluation of Complementing Components. In *Proceedings of the*

12th International Conference on Intelligent Text Processing and Computational Linguistics. Lecture Notes in Computer Science. Samos, Greece, 24-30 March, 2013., pages 488–500.

Gasperin, C., L. Specia, T. Pereira, and S.M. Aluísio. 2009. Learning When to Simplify Sentences for Natural Text Simplification. In *Proceedings of the Encontro Nacional de Inteligência Artificial (ENIA-2009), Bento Gonçalves, Brazil.*, pages 809–818.

Koehn, P., H. Hoang, A. Birch, C. Callison-Burch, M. Federico, N. Bertoldi, B. Cowan, W. Shen, C. Moran, R. Zens, C. Dyer, O. Bojar, A. Constantin, and E. Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *Proceedings of ACL*.

Koehn, P., F. J. Och, and D. Marcu. 2003. Statistical phrase-based translation. In *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology - Volume 1*, NAACL '03, pages 48–54, Stroudsburg, PA, USA. Association for Computational Linguistics.

Napoles, C. and M. Dredze. 2010. Learning simple wikipedia: a cogitation in ascertaining abecedarian language. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Linguistics and Writing: Writing Processes and Authoring Aids*, CL&W '10, pages 42–50, Stroudsburg, PA, USA. Association for Computational Linguistics.

Och, F. 2003. Minimum Error Rate Training in Statistical Machine Translation. In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 160–167.

Och, F. J. and H. Ney. 2003. A systematic comparison of various statistical alignment models. *Computational Linguistics*, 29(1):19–51.

Orasan, C., R. Evans, and I. Dornescu, 2013. *Towards Multilingual Europe 2020: A Romanian Perspective*, chapter Text Simplification for People with Autistic Spectrum Disorders, pages 287–312. Romanian Academy Publishing House, Bucharest.

Papineni, K., S. Roukos, T. Ward, and W. Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of ACL*.

Saggion, H., E. Gómez Martínez, E. Etayo, A. Anula, and L. Bourg. 2011. Text Simplification in Simplext: Making Text More Accessible. *Revista de la Sociedad Española para el Procesamiento del Lenguaje Natural*, 47:341–342.

Specia, L. 2010. Translating from complex to simplified sentences. In *Proceedings of the 9th international conference on Computational Processing of the Portuguese Language*, pages 30–39, Berlin, Heidelberg.

Stolcke, A. 2002. SRILM - an Extensible Language Modeling Toolkit. In *Proceedings of the International Conference on Spoken Language Processing (ICSLP)*, pages 901–904.

Woodsend, K. and M. Lapata. 2011. Learning to Simplify Sentences with Quasi-Synchronous Grammar and Integer Programming. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Wubben, S., A. van den Bosch, and E. Krahmer. 2012. Sentence simplification by monolingual machine translation. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers - Volume 1*, ACL '12, pages 1015–1024, Stroudsburg, PA, USA. Association for Computational Linguistics.

Yatskar, M., B. Pang, C. Danescu-Niculescu-Mizil, and L. Lee. 2010. For the sake of simplicity: unsupervised extraction of lexical simplifications from wikipedia. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, HLT '10, pages 365–368, Stroudsburg, PA, USA. Association for Computational Linguistics.

Zhu, Z., D. Berndard, and I. Gurevych. 2010. A Monolingual Tree-based Translation Model for Sentence Simplification. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 1353–1361.