

Confidence-driven Rewriting for Improved Translation

Shachar Mirkin, Sriram Venkatapathy, Marc Dymetman

Xerox Research Centre Europe

6 Chemin de Maupertuis

38240 Meylan, France

firstname.lastname@xrce.xerox.com

Abstract

Some source texts are more difficult to translate than others. One way to handle such texts is to modify them prior to translation. Yet, a prominent factor that is often overlooked is the source translatability with respect to the specific translation system and the specific model that are being used. We present an approach, and an interactive tool implementing it, where source sentences are rewritten in order to maximize confidence estimates with respect to the translation model. The automatically-generated rewrites are then proposed for the user's approval. Such an approach can reduce post-editing effort, replacing it by cost-effective pre-editing that can be done by monolinguals.

1 Introduction

Machine Translation (MT) systems still fall short of being able to compete with expert human translators. The advent of Statistical Machine Translation (SMT) systems has made it often possible to produce understandable “gists” of source texts in a target language spoken by the reader (so-called “assimilation” mode). However, they are not yet capable of producing target texts that are reliable enough to be externally distributed as trusted documents (“dissemination” mode); on the other hand, these target texts are very effective for human post-edition, where they can reduce significantly the cost of producing high-quality translations.

In the past, and typically within rule-based translation systems, one popular approach to im-

proving MT output has been the use of “controlled languages”. These typically consist of sets of more or less formal guidelines for encouraging authors to write clear and simple texts, and it has been shown that such conditions tend to improve the performance of MT systems. However, as (O'Brien, 2006) notes, such generic guidelines may perform differently depending on the specific language pair, technical domain, and even MT model that is being used.

Building on the insight that the quality of translation may depend on an appropriate formulation of the source, but in a way that needs to be geared to the specific MT model in use, we propose the following approach. First, we assume that the original source text in English (say) is not necessarily under the user's control, but may be given to her. While she is a fluent English speaker, she does not know the target language at all, but uses an MT system. Crucially, this system is able to provide estimates of the quality of its translations, or have access to such a quality estimation¹ component. The translation quality of each source sentence is estimated; for those sentences whose estimation is low, automatic rewrites are generated, geared towards improved translation with respect to the translation model in use. Two operational modes may be considered. The first is fully automatic, where the rewriting with the highest estimate replaces the original sentence and gets translated; the second, interactive, lets the user select a rewriting, verifying that it does not distort the original meaning. In this work we focus on the interactive mode and implement it in a prototype tool called SORT.²

¹Also known as *confidence estimation*.

²The tool is presented in more detail in the companion

One advantage of this framework is that the proposed rewritings are implicitly “aware” of the underlying strengths and limitations of the specific MT model. A good quality estimator, for instance, will feel more confident about the translation of an unambiguous word like *weapon* than about that of an ambiguous one such as *arm*, or about translating a known term in its domain than a term not seen during training.

This approach is relevant in many situations, for instance in the context of customer support, where subject matter experts who only speak a certain (source) language, need to provide reliable answers to customers speaking a different (target) language. There, it makes sense to put additional effort at pre-editing the source text if such effort comes with a higher confidence that the translation will be valuable to the customer, or if costly post-editing can be reduced.

The rest of the paper is structured as follows. In Section 2 we explain our approach to source rewriting; in Section 3 we provide details about the confidence estimation technique that we used; in Section 4 we describe SORT; Section 5 provides details about an evaluation of the approach and the analysis of results; in Section 6 we compare our approach to related work; finally, we provide perspectives and conclude in Section 7.

2 Source rewriting

One way to categorize text rewritings methods is by the semantic relation between the resulting text and the original one. The source text can be *paraphrased*, i.e. have its meaning expressed in a different way, but it can also be *generalized*, e.g. by having some of its details omitted.

In principle, paraphrasing is preferable, as it preserves the meaning of the source text. However, in some cases, it is preferable to have a more accurate translation with fewer details than a poor translation of the exact meaning of the original text. This was shown empirically, with human evaluators, in (Mirkin et al., 2009) and (Aziz et al., 2010) which dealt with the problem of unknown words by allowing the generation of an *entailed* version of the source text, and not only paraphrases of it.

On a different dimension, *text simplification* is motivated by the need to provide easy-to-read texts

system-demo paper (Mirkin et al., 2013).

to poorly literate people or language learners. Simplification operations include substitution of words by simpler ones, removal of complicated syntactic structures, shortening of sentences or removal of details not necessary for understanding the core idea of the text (Feng, 2008). Semantically, text simplification techniques may fall between paraphrasing (e.g. when replacing a rare word by a more common synonym) and generalization (e.g. when dropping a dispensable modifier).

In the current prototype, we implemented two rewriting techniques based on text simplification that work either at the full sentence level or at the level of individual words. Our assumption is that simpler sentences are more likely to yield higher quality translations. Clearly, this is not always the case; we leave this decision to the confidence estimation component.

Sentence-level simplification (Specia, 2010) has proposed to model text simplification as an SMT task where the goal is not to translate from one language to another, but to translate texts to their simplified version in the same language. In this approach, a simplification translation model is learnt from a parallel corpus of texts and their simplified versions. Applying this method, we trained an SMT model from English to Simple English, based on the PWKP corpus generated from Wikipedia (Zhu et al., 2010); we used only alignments involving a single sentence on each side. This resulted in a phrase table containing many entries where source and target phrases are identical, but also phrase-pairs that are mapping complex phrases to their simplified counterparts, e.g.:

- *the traditional etymology* → *the name*
- *primarily dry and secondarily cold* → *both cold and dry*
- *the high mountainous alps* → *the alps*

Also, the language model was trained with Simple English sentences to encourage the generation of simpler texts. Given a source text, it is translated to its simpler version, and *n*-best translations are assessed by the confidence estimator.

Lexical simplification One of the primary operations for text simplification is lexical substitution (Table 2 in (Specia, 2010)). Hence, in addition to rewriting a full sentence using the previous technique, we implemented a second method, address-

ing lexical simplification directly, and only modifying local aspects of the source sentence. The approach here was to extract relevant synonyms from our trained SMT model of English to Simple English, and use them as substitutions to simplify new sentences. We extracted all single token mappings from the phrase table of the trained model, removing punctuations, numbers and stop-words. We checked whether their lemmas were synonyms in WordNet (Fellbaum, 1998) (with all possible parts-of-speech as this information was not available in the SMT model). Only those were left as valid substitution pairs. When a match of an English word is found in the source sentence it is replaced with its simpler synonym to generate an alternative for the source.

Below are example lexical simplification rewritings for the source sentence “*Why the Galileo research program superseded rival programs*”:

1. *Why the Galileo research program **replaced** rival programs.*
2. *Why the Galileo research program **super-**
seded competitor programs.*
3. *Why the Galileo research program **replaced**
competitor programs.*

In SORT, both methods were used in conjunction to suggest rewritings for sentences with low confidence estimates.

3 Confidence estimation

MT systems are typically evaluated by comparing their output to *reference* translations. Yet, in many practical scenarios, reference translations are not available. *Quality estimation*, a.k.a. *confidence estimation*, denote methods that estimate the quality of an automated translation without depending on references. Instead, they rely on features extracted from the source, the translation, or from the translation process itself (Blatz et al., 2004; Specia et al., 2009). This is useful, for example, for identifying sentences that are suspected to be poorly translated in order to have them post-edited by human translators.

Our confidence estimator is based on the system and data provided for the 2012 *Quality estimation shared task* (Callison-Burch et al., 2012). In this task, participants were required to estimate the quality of automated translations. Their estimations were compared to human scores of the trans-

lation which referred to the suitability of the translation for post-editing. The scores ranged from 1 to 5, where 1 corresponds to translation that practically needs to be done from scratch, and 5 to translations that requires little to no editing.

The task’s training set consisted of approximately 1800 source sentences in English, their Moses (Koehn et al., 2007) translations to Spanish and the scores given to the translations by the three judges. With this data we trained an SVM regression model using SVM^{light} (Joachims, 1999). Features were extracted with the task’s feature-extraction baseline module. Two types of features are used in this module (i) *black-box* features, that do not assume access to the translation system, such as the length of the source and the target, and language model log probabilities, and (ii) *glass-box* features, which are extracted from the MT model itself, such as the average number of possible translations per source word.

4 The rewriting tool

In this section we shortly describe SORT, an implementation of our confidence-driven rewriting approach (see (Mirkin et al., 2013) for more details about the tool).

With SORT, the process starts when the user uploads a document that needs to be translated. The translation confidence of each sentence is computed and displayed with a color-coded label, to enable quick focus on the sentences that require more attention. Green denotes expected good translations, red labels mark sentences that are estimated to be poorly translated, and orange labels are given to all the remaining sentences.

We attempt to suggest rewritings for sentences that are not estimated to be well translated. When we are able to propose rewriting(s) with higher translation confidence than the original, a magnifying glass icon is displayed next to the sentence. Clicking it displays, on the right side of the screen, an ordered list of the more confident rewritings, along with their corresponding confidence estimations. The differences between the rewriting and the original sentence are highlighted. The first sentence on the list is always the original one, to allow editing it, and to make it easier to view the differences. An example is displayed in Figure 1, which shows a part of SORT’s interface. A rewriting is proposed for the fourth sentence in the document.



Figure 1: Confidence estimations and proposed rewritings for source sentences in SORT.

Here, the suggestion is simply to replace the word *captured* with the word *caught*, a rewriting that is estimated to improve the sentence’s translation.

The user can select one of the suggestions or choose to edit either the original or one of the rewritings; the document is then updated based on her selection. At any time the user (if she speaks the target language) can view the translation of the source or of its rewritten version (with the cogwheel icon). When done, she can save the edited text or its translation. Moses Release 1.0 of an English-Spanish Europarl-trained model³ was used in this work to obtain English-Spanish translations.

5 Analysis

We performed an evaluation of our approach in an English to Spanish translation setting, using a subset of the 2008 News Commentary data.⁴ The main goal is to learn whether the proposed framework leads to improved translations; we further wish to assess its specific components, namely the confidence estimator and the rewiring methods.

First, two annotators who speak English but not Spanish used SORT to rewrite an English text. They reviewed the proposed rewritings and were instructed to “trust the judgment” of the confidence estimator; that is, reviewing the suggestions from the most to the least confident one, they accepted the first rewriting that was fluent and preserved the meaning of the source document as a whole. The original sentence and the selected alternative were then both translated to Spanish and presented as competitors to two native Spanish speakers. The sentences were placed within their context in the original document, taken from the Spanish side of

the corpus. The order of presentation of the two competitors was random. The annotators had to choose which of the two sentences they would prefer having in the given context, or “no clear winner” if they could not choose one over the other.

Our main insights from the evaluation are described next. We start by analyzing the confidence-driven rewritings, and move on to describe their impact on the actual translation.

5.1 Source-side analysis

For this part of the evaluation, 960 English sentences were loaded into SORT. The lexical method offered rewritings for 70% of them; the sentence-level method proposed different rewritings for more than a half of these,⁵ as well as for 116 (12%) other sentences. By construction, the sentence-level method generates at least one rewriting for each sentence;⁶ yet, since this method translates from English to English, sometimes the generated rewriting is identical to the original sentence; thus, for 18% of the sentences, no rewriting was produced. 57% of the sentences with proposed rewritings had higher-confidence suggestions. These were approximately equally divided between the sentence-level and the lexical methods.

Table 1 shows several rewritings that were suggested by the each of the two methods and accepted by the English-speaking annotators. The lexical method replaced a word or two in the sentence by their synonyms; the sentence-level method did that as well, but also generated other rewriting types, including multi-word synonym substitutions, and some other simple paraphrasing.

Example 1 shows sentence-level rewriting,

⁵That is, excluding the cases where the two methods yielded the same suggestion.

⁶In this evaluation we used the single-best translation.

³<http://www.statmt.org/moses/RELEASE-1.0/model/>

⁴Available from <http://www.statmt.org>

Original	Rewriting
or somewhere in that vicinity have welcomed Christmas in this fashion every year reduced the probability of the recession roughly 140,000 drivers	or somewhere in that neighbourhood have welcomed Christmas in this way every year reduced the chance of the recession around 140,000 drivers
exceeds the two-thirds mark (71 percent) to scrap co-payments for doctor visits giant supermarkets located in the shopping triangle it had discharged a patient there's no mistaking the political message not mandatory for either women or men sugar cane from the vast plantations was transported things you can annoy your colleagues with	exceeds the two-thirds mark (71%) to get rid of co-payments for doctor visits giant supermarkets in the shopping triangle it had let go a patient there is no mistaking the political message not a must for either women or men sugarcane from the vast plantations was sent things you can be annoying your colleagues with

Table 1: Snippets from accepted rewriting by the lexical (top) and the sentence-level (bottom) methods. Boldface marks the difference between the original and rewritten text.

where (a) is a segment of the original sentence and (b) is its corresponding rewriting.

- (1) a. *Looking only at the ratios of people who say they definitely **will vote in the referendum**, we find . . .*
- b. *Looking only at the ratios of people who say they definitely **will vote**, we find . . .*

The validity of such simplification is determined by the context. If the context makes it clear that *will vote* refers to *will vote in the referendum*, then — given that a better translation is anticipated — this is a desirable kind of rewriting. Indeed, considering the context, shown below, we see that the removal of the modifier is valid here.

*When queried about interest **in a referendum**, 60 percent of the sample said it definitely would **vote in one**. . . Given current determination **to vote**, **the referendum** is certain to be valid. . . Looking only at the ratios of people who say they definitely **will vote [in the referendum]**, we find . . .*

Yet, in another case, a modifier removal by the sentence-level method resulted with an invalid rewriting, as shown in Example 2.

- (2) a. *the american **federal** reserve bank , the european central bank . . .*
- b. *the american reserve bank , the european central bank . . .*

This rewriting was mistakenly accepted by an English-speaking annotator. We suspect that the

reason lies in the interface that does highlight differences materialized as additions or substitutions, but does not provide the user with an easy means to spot deletions. We intend to address that in future versions of the tool.

We stress that the sentence-level method is working, well, at the level of a single sentence, as standard SMT systems do. Thus, the removal of the *in a referendum* modifier was not done through consideration of the wider context, but based on rules learned from the training set. It was here successful, but could be mistaken in another case without having the confidence estimator spotting it. Indeed, the user can judge whether a deleted text segment is necessary in context; however, this can also be integrated into the tool, either by using rewriting methods that consider wider context, or by filtering out rewritings that deviate too much from the source sentence. We intend to investigate these directions in future research.

A different kind of error occurred when a wrong synonym was suggested as a replacement (e.g. *Christmas air* for *Christmas atmosphere*). This was somewhat surprising as we had expected the language model features of the confidence estimator to help removing these cases. While they were filtered by the English-speaking users, and thus did not present a problem for translation, they created unnecessary workload. Putting more emphasis on context features in the confidence estimation or explicitly verifying context-suitability of lexical substitutions could help addressing this issue.

Source	Translation
70 thousand Slovaks registered in Britain who moved there for work.	70000 eslovacos registrados en Gran Bretaña que trasladó allí para trabajar
70 thousand Slovaks registered in Britain who went there for work.	70000 eslovacos registrados en Gran Bretaña que fueron allí para trabajar
bought a homestead in San Francisco de Paula	comprado una homestead en San Francisco de Paula
bought a house in San Francisco de Paula	comprado una casa en San Francisco de Paula
quarrels between local residents and Slovak immigrants are beginning to escalate	disputas entre los residentes locales y los inmigrantes eslovaca empiezan a intensificarse
quarrels between local residents and Slovak immigrants are starting to escalate	disputas entre los residentes locales y los inmigrantes eslovaca están empezando a una escalada

Table 2: Examples of rewriting and their corresponding translations to Spanish. The original version and its translations are shown in white; the rewrites and their translations are shaded.

5.2 Impact on translation

440 sentences were used for the Spanish annotation. Approximately a quarter of these had higher-confidence suggestions that were accepted by the English-speaking annotators. 15% of those yielded identical translations to that of the original. These cases almost always originated from the lexical method, when two source synonyms were translated to the same target word. For instance,

- (3) a. *the organizations joining the strike will **begin** an open-ended work stoppage*
b. *the organizations joining the strike will **start** an open-ended work stoppage*

Here, both *will begin* and *will start* were translated to the Spanish word *comenzará*. To save pre-editing effort, we can avoid showing such rewrites to the user, as the result will be the same. Still, this information is useful for updating the confidence in the translation of the original sentence.

~30% of the rewriting translations were preferred by the annotators, in comparison to ~20% where the original was chosen. In the rest of the cases, “no clear winner” was selected.⁷ Among the two methods, the translations of rewrites generated by the sentence-level method were preferred more often by the annotators than those generated by the lexical one.

Table 2 (top) shows two cases where rewriting improved translation, as judged by the Spanish annotators. In general, improvement was manifested

⁷One must consider these figures with caution, as the numbers at this point may be too small to be statistically meaningful.

in better readability, grammaticality or the removal of unknown words from the translation. By contrast, the bottom part of the table shows a case where rewriting resulted in an inferior translation.

The number of preferred original translations indicates that the confidence estimator is not always discriminative enough. By construction, for every rewriting that is displayed, the confidence component estimates the translation of the original to be less accurate than that of the rewriting; yet, this is not always reflected in the preferences of the evaluators. On a different dimension than translation quality, the large number of cases with no clear winner, and the analysis we conducted, indicate that the user’s cognitive effort would be decreased if we displayed only those rewrites associated with a substantial improvement in confidence; due to the nature of our methods, identical or near-identical translations were frequently generated, with only marginal differences in confidence, as in the case of two source synonyms being translated to the same target word.

6 Related work

One approach to produce translatable text is by enforcing the use of a “controlled language” during source text authoring. (Mitamura, 1999) discusses the design of such a language in the context of rule-based translation and proposes tools for checking source sentences; (Carbonell et al., 1997) propose an interactive tool to enforce such constraints on the source in a similar context. In the case of restricted semantic domains, (Dymetman et

al., 2000; Power et al., 2003) propose methods that directly generate multilingual target texts based on a semantically-guided authoring process (see also (Ranta, 2011) for a formal paradigm that may be used for similar applications). In the context of SMT, the first two authors of this paper recently presented an authoring tool that consults the MT system itself to propose phrases that can be used during composition to obtain better translations (Venkatapathy and Mirkin, 2012). All the above methods address authoring of source texts from scratch. By contrast, this paper is concerned with modifying an existing text to improve its translatability (an approach that was however mentioned as a future perspective in (Mitamura, 1999)).

For the modification of existing text, another approach is to paraphrase the source or to generate entailed sentences (Callison-Burch et al., 2006; Mirkin et al., 2009; Marton et al., 2009; Aziz et al., 2010). Yet, these works focus on handling out-of-vocabulary (OOV) words, do not assess the translatability of the source sentences and do not offer an interactive solution for the author. Another way to use paraphrases for improved translation has been proposed by (Max, 2010) who paraphrases source texts to increase the number of training examples for the SMT system.

A different approach is pre-ordering, a strategy used to improve translation for language pairs that have different syntactic structures (e.g. Japanese SOV and Chinese SVO). There, the words in the source text are reordered to make their order more similar to that of the target language (Isozaki et al., 2010; Wu et al., 2011). Unlike our approach, pre-ordering is not specific to the MT system; further, the technique does not produce fluent, readable sentences that can be validated by users, and is therefore less suitable for interactive source rewriting. (Choumane et al., 2005) propose an iterative system where the author helps the translation system “understand” a given text by tagging text positions that represent potential syntactic ambiguities. As in most cases mentioned earlier, the rules used by this method are generic and are not tailored for a specific MT system or model.

Monolingual-based editing for translation is proposed in the MonoTrans2 project (Hu et al., 2011); monolingual speakers of the source and the target languages collaborate to improve the translation. Unlike our approach, here both the feed-

back for poorly translated sentences and the actual modification of the source are done by humans.

7 Conclusions

We introduced an approach for interactively rewriting texts for translation under the guidance of a confidence estimator, and implemented a tool integrating two rewriting techniques. We observed that the translations of rewrites were preferred more often than those of the original source.

While we focused on an interactive mode, selections could be made automatically by choosing a rewrite associated not only with a high translation confidence, but also with a high “rewrite confidence” (i.e. the probability of preserving the original meaning). This is a direction we plan to pursue. A precise confidence estimator along these two dimensions would add considerable value to the current interactive framework as well, by supporting more powerful rewriting techniques having the potential to explore a larger space of possible reformulations. Methods based on textual entailment (Dagan et al., 2009), coupled with robust filtering methods, would be a natural framework for such investigations.

Another direction would be to explore rewriting techniques and confidence estimators that are more context-aware. Such methods would facilitate removing redundant information within rewrites and help maintaining discourse fluency. A context-matching component as used in (Mirkin et al., 2009) would be useful for ensuring substitution accuracy.

Finally, the ability to adapt the confidence estimators based on observations of user interactions would lead to an improved usability of the proposed approach.

Acknowledgements

We wish to thank Ioan Calapodescu for his assistance in implementing the SORT interface. We also thank Matthias Gallé, José-Antonio Rodríguez and Matias Tealdi for their help with the Spanish evaluation and analysis.

References

Aziz, Wilker, 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 *Proc. of EAMT*.
- Blatz, John, Erin Fitzgerald, George Foster, Simona Gandrabur, Cyril Goutte, Alex Kulesza, Alberto Sanchez, and Nicola Ueffing. 2004. Confidence estimation for machine translation. In *Proc. of COLING*.
- Callison-Burch, Chris, Philipp Koehn, and Miles Osborne. 2006. Improved statistical machine translation using paraphrases. In *Proc. of HLT-NAACL*.
- Callison-Burch, Chris, Philipp Koehn, Christof Monz, Matt Post, Radu Soricut, and Lucia Specia. 2012. Findings of the 2012 workshop on statistical machine translation. In *Proc. of WMT*.
- Carbonell, Jaime G, Sharlene L Gallup, Timothy J Harris, James W Higdon, Dennis A Hill, David C Hudson, David Nasjleti, Mervin L Rennich, Peggy M Andersen, Michael M Bauer, et al. 1997. Integrated authoring and translation system. US Patent 5,677,835.
- Choumane, Ali, Hervé Blanchon, and Cécile Roisin. 2005. Integrating translation services within a structured editor. In *Proc. of the ACM symposium on Document engineering*.
- Dagan, Ido, Bill Dolan, Bernardo Magnini, and Dan Roth. 2009. Recognizing textual entailment: Rational, evaluation and approaches. *Natural Language Engineering*, pages 15(4):1–17.
- Dymetman, Marc, Veronika Lux, and Aarne Ranta. 2000. Xml and multilingual document authoring: Convergent trends. In *Proc. of COLING*.
- Fellbaum, Christiane, editor. 1998. *WordNet: An Electronic Lexical Database (Language, Speech, and Communication)*. The MIT Press.
- Feng, Lijun. 2008. Text simplification: A survey. Technical report, CUNY.
- Hu, Chang, Philip Resnik, Yakov Kronrod, Vladimir Eidelman, Olivia Buzek, and Benjamin B. Bederson. 2011. The value of monolingual crowdsourcing in a real-world translation scenario: simulation using haitian creole emergency sms messages. In *Proc. of WMT*.
- Isozaki, Hideki, Katsuhito Sudoh, Hajime Tsukada, and Kevin Duh. 2010. Head finalization: a simple re-ordering rule for sov languages. In *Proc. of the Joint Fifth Workshop on SMT and MetricsMATR*.
- Joachims, T. 1999. Making large-scale SVM learning practical. In Schölkopf, B., C. Burges, and A. Smola, editors, *Advances in Kernel Methods - Support Vector Learning*, chapter 11, pages 169–184. MIT Press.
- Koehn, Philipp, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondrej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *Proc. of ACL, Demo and Poster Sessions*.
- Marton, Yuval, Chris Callison-Burch, and Philip Resnik. 2009. Improved statistical machine translation using monolingually-derived paraphrases. In *Proc. of EMNLP*.
- Max, Aurélien. 2010. Example-based paraphrasing for improved phrase-based statistical machine translation. In *Proc. of EMNLP*.
- Mirkin, Shachar, Lucia Specia, Nicola Cancedda, Ido Dagan, Marc Dymetman, and Idan Szpektor. 2009. Source-language entailment modeling for translating unknown terms. In *Proc. of ACL-IJCNLP*.
- Mirkin, Shachar, Sriram Venkatapathy, Marc Dymetman, and Ioan Calapodescu. 2013. Sort: An interactive source-rewriting tool for improved translation. In *ACL System Demonstrations*.
- Mitamura, Teruko. 1999. Controlled language for multilingual machine translation. In *Proc. of Machine Translation Summit VII*.
- O’Brien, Sharon. 2006. Controlled Language and Post-Editing. *Multilingual*, 17(7):17–19.
- Power, Richard, Donia Scott, and Anthony Hartley. 2003. Multilingual generation of controlled languages. In *Proc. of EAMT/CLAW*.
- Ranta, Aarne. 2011. *Grammatical Framework: Programming with Multilingual Grammars*. CSLI Publications, Stanford. ISBN-10: 1-57586-626-9 (Paper), 1-57586-627-7 (Cloth).
- Specia, Lucia, Nicola Cancedda, Marc Dymetman, Marco Turchi, and Nello Cristianini. 2009. Estimating the sentence-level quality of machine translation systems. In *Proc. of EAMT*.
- Specia, Lucia. 2010. Translating from complex to simplified sentences. In *Proc. of PROPOR*.
- Venkatapathy, Sriram and Shachar Mirkin. 2012. An SMT-driven authoring tool. In *Proc. of COLING 2012: Demonstration Papers*.
- Wu, Xianchao, Katsuhito Sudoh, Kevin Duh, Hajime Tsukada, and Masaaki Nagata. 2011. Extracting pre-ordering rules from predicate-argument structures. In *Proc. of IJCNLP*.
- Zhu, Zheming, Delphine Bernhard, and Iryna Gurevych. 2010. A monolingual tree-based translation model for sentence simplification. In *Proc. of COLING*.