

Interactive-Predictive Machine Translation based on Syntactic Constraints of Prefix

Na Ye, Guiping Zhang and Dongfeng Cai

Human-Computer Intelligence Research Center

Shenyang Aerospace University, Shenyang 110136, China

yn.yena@hotmail.com, zgp@ge-soft.com, caidf@vip.163.com

Abstract

Interactive-predictive machine translation (IPMT) is a translation mode which combines machine translation technology and human behaviours. In the IPMT system, the utilization of the prefix greatly affects the interaction efficiency. However, state-of-the-art methods filter translation hypotheses mainly according to their matching results with the prefix on character level, and the advantage of the prefix is not fully developed. Focusing on this problem, this paper mines the deep constraints of prefix on syntactic level to improve the performance of IPMT systems. Two syntactic subtree matching rules based on phrase structure grammar are proposed to filter the translation hypotheses more strictly. Experimental results on LDC Chinese-English corpora show that the proposed method outperforms state-of-the-art phrase-based IPMT system while keeping comparable decoding speed.

1 Introduction

In recent years, the machine translation (MT) technology has achieved great progress. However, up till now the MT output still cannot meet the practical requirements on translation quality in many scenarios, and need to be post-edited by human translators before actually put to use. In order to help MT systems collaborate with human translators more effectively, researchers carried out the study on computer-assisted translation (CAT), in which the main goal of MT is supporting professional translators in improving their productivity. Therefore, the ability to interact with human becomes an important research topic in CAT.

The first interactive machine translation systems (Kay and Martins, 1973; Zajac, 1988; Yamron et al., 1993) focus on having human translators disambiguate the source texts through answering questions. However, this question-answering process remains a laborious one for human translators. Under such circumstances, the interactive-predictive machine translation (IPMT) method is proposed (Foster et al., 1997). In the IPMT mode, first the system generates one or more raw suggestions, and then the human translator validates the longest correct prefix in the suggestions and revises the first character in the corresponding suffix, next the new prefix is used to help the system predict the optimal suffix. This process is repeated until the correct translation is acquired. The IPMT technology enables human translators to avoid the burden of explaining the source text, and directly control the final translation generation, so it attracted widespread attention.

It can be seen that the essential difference between IPMT and MT is the introduction of a constraint, namely the target sentence must start with the human validated prefix. To achieve this goal, researchers attempted various MT models (Och et al., 2003; Civera et al., 2004; Tomás and Casacuberta, 2006; Barrachina et al., 2009; González-Rubio et al., 2013). In these methods, prefix is mainly used for performing character-level matching on translation hypotheses to reduce the search space. However, the hypotheses that only match the prefix on character level are perhaps not correct. Figure 1 gives an ex-

This work is licenced under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>

ample.




Figure 1: Example of translation hypothesis selection on character level.

In Figure 1, the validated prefix is “the designated p”. Two hypotheses that start with “the designated person” and “the designated programme” both meet the requirement. If we use “person” to extend the prefix (see the alignment by solid line), then the hypothesis will form a complete subtree translation (circled by dashed line). But if we use “programme” to extend the prefix, then an error will occur. The circled subtree has only been partly translated before the hypothesis turns to translate another subtree (see the alignment by dashed line). In the future interactions, no matter how to extend the hypothesis, no reasonable syntactic alignment will be generated. Therefore, other than character-level matching results with the prefix, we should also take syntactic-level matching results into account for hypothesis selection. Only when a reasonable syntactic alignment is formed between the hypothesis and the source sentence, can we decide this is a good hypothesis.

In this paper we present a mathematical IPMT framework based on the syntactic alignment between the source sentence and the prefix. On the basis of the framework, we proposed two syntactic subtree constraints based on phrase structure grammar for selecting translation hypotheses. Experimental results on LDC corpora show that our method reduced the human-computer interaction times under comparable decoding speed.

2 Related Work

The task of IPMT (Barrachina et al., 2009) is to find the optimal suffix under the condition of a given source sentence s and a validated prefix t_p :

$$\hat{t}_s = \underset{t_s}{\operatorname{argmax}} P(t_s / s, t_p) = \underset{t_s}{\operatorname{argmax}} P(t_p, t_s / s) \quad (1)$$

where $(t_p, t_s)=t$, indicating that the prefix t_p and the predicted suffix t_s concatenate to form a complete translation t .

As previously discussed, when modelling $P(t_p, t_s | s)$, current methods mainly make use of the character-level matching results with the prefix. To enhance the guiding effect of the prefix, some work also considered other factors.

Sanchis-Trilles et al. (2008) integrates the user’s mouse actions into the IPMT system. The method is based on an assumption that the first character of the predicted suffix must be different from the current suffix when the user clicks the mouse on the translation. In this way, the clues hidden in the prefix are further exploited. However, this method still restricted to character or word level matching.

Some researchers investigated the word alignment between the prefix and the source sentence. Nepveu et al. (2004) proposed a cache-based adaptive prediction model. For the predicted translation w of each active word a , once the user accepts w , a word pair (a, w) will be stored in the cache. Higher language model probability and translation probability will be assigned to w if a new source sentence

contains the word a . Ortiz-Martínez et al. (2009) performs word alignment between the prefix and the source sentence. The generation of suffix is limited to the translation of the unaligned parts in the source sentence. These methods deepened the prefix matching level, but still did not reach the syntactic level.

González-Rubio et al. (2013) adopts hierarchical phrase-based model (HPBM) to IPMT. Because HPBMs are on the basis of synchronous grammar, the prefix can guide decoding on a deeper level. However, the synchronous grammar in hierarchical phrases has no linguistic meaning, and cannot evaluate the reasonableness of the hypotheses from the view of syntactic structure.

Compared with the above research, our method made use of the prefix on syntactic level. The hypotheses for which it is impossible to generate reasonable syntactic structure alignment with the source sentence are identified and filtered.

3 IPMT Mathematical Model

In this paper, we examine the syntactic alignment between the source sentence s and t_p when we make use of the prefix. A hidden variable $T(s)$ is introduced to represent the parse tree of s . Consequently, $P(t_p, t_s | s)$ is converted to:

$$P(t_p, t_s | s) = \sum_{T(s)} P(t_p, t_s, T(s) | s) \quad (2)$$

The item on the right side is further deduced with:

$$P(t_p, t_s, T(s) | s) = P(T(s) | s) \times P(t_p | T(s), s) \times P(t_s | t_p, T(s), s) \quad (3)$$

The first factor on the right side is the syntactic parsing model of the source language, which is provided by the parser. The second factor corresponds to the transformation from the source sentence to the prefix, which is the machine translation model. So the two models need not be discussed. The third factor is the key model of this paper, which corresponds to the prediction of suffix t_s . We will emphasize on the modelling of this factor.

There are two ways to model $P(t_s | t_p, T(s), s)$. One is completely adopting the syntax-based MT framework and performing prefix matching during decoding. The other is adding syntactic information into the PBM-based SMT framework as rules. In comparison, the former way is more straightforward, but it is prone to be influenced by the performance of parsing and translation rule extraction algorithms. Incorrect parse tree of the source sentence and incorrect translation rules can both lead to the consequence that the prediction will never succeed. Although some researchers proposed forest-based translation approaches (Mi and Huang, 2008; Zhang et al., 2009) to avoid relying on 1-best parse tree, the computation costs of these approaches are too high for the IPMT systems which have strict speed requirements. The latter way allows the existence of non-syntactic phrases and has larger search space. To alleviate the negative effect of wrong parsing results, we can filter the hypotheses only when there is more than one candidate. In other words, the syntactic structure can play the role as soft constraints if we adopt relatively tolerant syntactic tree matching rules. Once the prefix of a hypothesis has the possibility of being correctly aligned to the source-language parse tree, the hypothesis will be kept. In this way, the system can be more error-tolerant. Therefore, we built the model in the latter way and adopted the phrase-based model (PBM) as in (Barrachina et al., 2009).

We introduce a hidden variable A to represent the phrase alignment between t_p and $T(s)$. The third factor of Equation 3 can be transformed to:

$$P(t_s | t_p, T(s), s) = \sum_A P(t_s, A | t_p, T(s), s) \quad (4)$$

This paper estimates the factor on the right side as follows:

$$P(t_s, A | t_p, T(s), s) = \begin{cases} P(t_s | t_p, s) & \text{if } \xi(t_p, A, T(s)) = \text{TRUE} \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

where $\xi(t_p, A, T(s))$ is a function to judge whether the alignment A conforms to the syntactic tree matching rules. The goal is to decide whether there is possibly correct syntactic alignment between the prefix and the source sentence.

4 Hypothesis Selection

Generally, syntactic structure can be represented or labelled by two forms. One is the phrase structure grammar (PSG), and the other is the dependency structure grammar (DSG). In this paper we choose PSG for the extension of prefix. There are two reasons: first, PSG can clearly give the syntactic component borders (subtrees) with complete linguistic meaning; second, the PSG parsers mainly use probabilistic context-free grammar (PCFG), and the produced N-gram translation rules can well model the orders of the syntactic components. Such information can straightforwardly direct the extension of translation hypotheses. But the dependency grammar describes the binary head-modifier structure, so it is inefficient for the PBM model which takes multi-word phrases as the basic processing units.

In order to use the syntactic structure of the source sentence to guide the prefix extension, we analysed the nature of the prefix and propose two subtree matching rules (or constraints) which should be followed while selecting translation hypotheses.

4.1 Complete Subtree Constraint

This constraint requires that the selection of the phrase to extend the prefix needs to consider whether the current subtree¹ has been completely translated. Figure 2 gives an example.




Figure 2: Translation hypothesis extension based on complete subtree constraint.

In Figure 2, the prefix t_p consists of complete words w_1, w_2, w_3 and an incomplete word c_4 . The suffix t_s starts with an incomplete word $\sim c_4$. c_4 and $\sim c_4$ together form a complete word w_4 . The data structure of node n_i records the information whether the word has already been translated (1 for *yes*, 0 for *no*). In this example, n_1, n_2 and n_3 are translated nodes (the word alignments are indicated by solid lines), n_4 and n_5 are untranslated nodes. If the translations of n_4 and n_5 can both be used to extend w_3 (indicated by dashed lines), then we need to examine the subtree (ST_1) to which n_2 belongs. Since there is still a node n_4 left untranslated in the subtree, n_4 should be chosen to perform extension.

The core of the complete subtree constraint is to judge whether the subtree to which the phrase to be extended belongs has been completely translated. Since a subtree may be nested in another subtree, it is possible that the same phrase is contained in multiple subtrees. In this paper we select the subtree with the minimum span to constrain the hypothesis extension.

After deciding the subtree ST to constrain hypothesis extension, the phrase p_1 which is to be extended and the phrase p_2 which is to extend p_1 are examined. Through checking the positions of the words covered by p_1 in the source sentence, the words not covered by subtree ST can be found. Because there

¹ In this paper, since it is meaningless to examine nodes under the complete tree S , the term “subtree” refers to proper subtree and does not include the complete tree.

are re-orderings in the translation, the positions of these words are perhaps not continuous. Therefore, the translation hypotheses should meet the following requirements: a) the positions of the source words covered by p_2 are fully contained in the span of the subtree; or b) the positions of the source words covered by p_2 are continuous and are located at the tail of the source words covered by the subtree, and the source words covered by p_1 fully cover the remained positions of the subtree.

4.2 Subtree Order Constraint

This constraint requires that the selection of the subtree to be extended (translated) needs to consider the overall structure of the syntactic tree and the re-ordering rules. Figure 3 gives an example.




Figure 3: Translation hypothesis extension based on subtree order constraint.

In Figure 3, the prefix consists of a complete word w_1 and an incomplete word c_2 . The suffix t_s starts with an incomplete word $\sim c_2$. If the translations of n_2 and n_3 can both be used to extend w_1 , then we need to examine the subtrees to which n_2 and n_3 belongs. This example possesses an “ NP_1 *de* NP_2 ” structure. According to the re-ordering rules, it should be translated into “ NP_2 *of* NP_1 ”. This means that we should choose the subtree to which n_3 belongs to perform extension.

We use the NiuTrans (Xiao et al., 2012) system to generate tree-to-string translation rules to evaluate whether the order of the subtrees are reasonable. First perform syntactic parsing on the source sentence and get the parse tree T ; then identify all the subtrees ST_i in T and record the child nodes LC_i (left child) and RC_i (right child) of the root node of each subtree; next use these subtree structures to filter the candidate translation rule set RT that can be used by the current sentence. During decoding, the hypothesis with highest score that matches the translation rules is selected for extension.

The key of this constraint is judging whether the phrase p_1 to be extended and the phrase p_2 to extend p_1 are closely adjacent in the translation rule. To achieve this goal, we need to identify the minimum common subtree MCT of the source phrases corresponded to p_1 and p_2 , and then use the translation rules that match this subtree to perform judging.

Since the translation rule base cannot fully cover all the syntactic structure instances, this paper generalized the subtree to acquire the rules that nearly match. If the needed rule is not in the rule base, then the subtrees in MCT will be generalized to their parent nodes from bottom to up. As long as any translation rule can match a generalized subtree, the new parent node will replace the child nodes for order judgement.

4.3 Balancing Strategies

During decoding, the two constraints go forward one by one. First judge whether the two phrases conform to the complete subtree constraint. If so, continue to judge whether they conform to the subtree order constraint. In any of the above two stages, once there is no hypothesis that matches the constraint, the algorithm will go back and accept the result of the previous stage.

However, there exist many mistakes in the parsing results. Therefore, strictly following the subtree constraints according to the parsing results will lead to the loss of some good hypotheses. To achieve balance between analysis depth and searching precision, we adopt the following strategies:

(1) If the hypotheses to extend the current translation contain the candidates that meet the constraints, then searching is limited within the scope of these candidates, otherwise we still search within all the hypotheses. This strategy can better take advantage of the non-syntactic phrases in the PBM-based SMT models.

(2) To reduce the complexity, we do not distinguish the border word of the prefix, but take phrases as the basic processing unit and examine whether the phrase to be extended and the phrase to extend it conform to the subtree constraints.

(3) We do not require that every hypothesis extension conforms to the two constraints. Considering the characteristics of IPMT task, we propose three strategies. a) only consider the subtree constraints when the hypothesis to be extended has not fully covered the prefix; b) only consider the subtree constraints when the hypothesis to be extended just covers the prefix; c) consider the subtree constraints under both the above two situations.

5 Experimental Results

5.1 Data Setup

This paper uses the Chinese-English Hong Kong Laws Parallel Text (LDC2000T47) as the corpora. 200,000 sentence pairs are taken as the training set. 1000 sentence pairs are randomly extracted from the remaining corpora as the development set, and 1558 sentence pairs as the testing set. Table 1 gives a detailed description of the corpora.

Corpus		Chinese	English
Training Set	Sentences	200K	200K
	Words	5.15M	5.11M
	Vocabulary	30K	31K
Development Set	Sentences	1000	1000
	Words	15K	15.7K
	Vocabulary	73.24	48.65
Testing Set	Sentences	1558	1558
	Words	20.6K	21.6K
	Perplexity	72.67	46.75

Table 1: Statistics of the Evaluation Corpora.

The Chinese portions of these data were pre-processed by ICTCLAS word segmenter², and the English portions were tokenized and lowercased. GIZA++ tool was used to perform the bi-directional word alignment of the training data, and the “grow-diag-final” strategy was used to merge the bi-directional results. A 3-gram language model was trained on the English portion of the training corpus with SRILM. For the building of PB SMT models, Moses (Koehn et al., 2007) was used, and the model includes 14 default features. For adjusting feature weights, the MERT (Och, 2003) method was applied, optimizing the BLEU-4 metric obtained on the development corpus. The parse trees were produced by the Berkeley parser, and 1-best tree was used for subtree extraction. During decoding, the size of hypothesis stack is set to 30, and the maximum number of translation options is set to 20 for each source phrase.

We used Key-stroke Ratio (Barrachina et al., 2009) to evaluate the performance of the IPMT systems. The lower the KSR score, the better the system performance. The baseline system is the state-of-the-art PBM-based IPMT approach using multi-stack-decoding algorithm as described in (Barrachina et al., 2009). We follow the user simulation approach for evaluation as previous works in the literature (Barrachina et al., 2009; González-Rubio et al., 2013). Statistical significance test is conducted using the resampling method proposed in (Koehn, 2004; Zhang et al., 2004).

² <http://ictclas.nlp.ir.org/>

5.2 Results and Analysis

Table 2 gives the comparative experimental results after using the complete subtree constraint (the distortion distance is limited to 10). ST_CM_BCP represents considering the rule when the hypothesis to be extended has not fully covered the prefix, ST_CM_ACP represents considering the rule when the hypothesis to be extended just covers the prefix, and ST_CM represents considering the rule under both the above two situations. We evaluated the systems on different K -best lists. The best results are displayed in bold fonts and have a statistically significant difference with respect to the baseline (95% confidence).

Method	1-best	5-best	10-best	20-best
baseline	48.66	47.93	47.76	47.55
ST_CM_BCP	48.05	47.41	47.18	46.97
ST_CM_ACP	48.44	47.75	47.48	47.21
ST_CM	47.85	47.16	46.88	46.69

Table 2: KSR scores after using the complete subtree constraint.

From Table 2 we can see that adding source-language syntactic structure constraint can effectively reduce the interaction times. And using the complete subtree constraint in the whole process of prefix generation (ST_CM) achieved more improvement than only using it in a certain stage (ST_CM_BCP and ST_CM_ACP). The reason is that the prefix is a fragment validated by the human, and it can guide the hypothesis extension at any stage. The following experiments in the paper are all based on this setting (consider the rules under both situations). With the increase of the number of translations, the complete subtree constraint also plays a positive role. On the testing corpus, the underling MT engine achieves a BLEU score of 0.2678.

Table 3 gives the KSR scores after adding the subtree order constraint. Experiments are conducted under different distortion distances. ST_CM represents only using the complete subtree constraint, ST_RD represents only using the subtree order constraint, and ST_IMT represents using both constraints.

Method	Distortion distance limitation = 10				Distortion distance limitation = 15			
	1-best	5-best	10-best	20-best	1-best	5-best	10-best	20-best
baseline	48.66	47.93	47.76	47.55	47.61	46.53	46.29	46.09
ST_CM	47.85	47.16	46.88	46.69	47.18	46.11	45.86	45.68
ST_RD	48.21	47.58	47.13	46.94	47.39	46.32	46.05	45.89
ST_IMT	47.74	47.07	46.73	46.41	46.88	45.83	45.59	45.41

Table 3: KSR scores after using the subtree order constraint.

Table 3 shows that after using the subtree order constraint, the interaction times decrease. But the improvement is not as obvious as the complete subtree constraint. The reason is that the translation rules consider the inner structure of the subtrees and are more fine-grained. So the matching difficulty increases. It also can be seen that using both constraints leads to larger improvement in performance, indicating that they are complementary constraints which guide the hypothesis selection from different aspects. These results verify that the user-validated prefixes are effective on multiple levels. In fact, when a user gives a prefix, he/she is making a comprehensive decision after analysing the overall structure and orders of the translation. And the proposed method exploits the syntactic structure information hidden in the prefix. In addition, although there are parsing mistakes, we can still get positive results; this shows that the method is error-tolerant.

Through comparing the IPMT process on specific sentences, we find that the decoding of many sentences will fail in the baseline system. In other words, there is no hypothesis that can match the prefix in the stack. However, after adding the subtree constraints, some hypotheses that do not have reasonable syntactic structures will be filtered during extension, thus some lower-ranked hypotheses have the opportunities to be pushed into the stack and finally match the prefix. In such cases, the efficiency of human-computer interaction greatly improves. In our experiment, 1.8% sentences which cannot find

the correct translation with the baseline method succeeded in finding the correct translation with the new method.

Table 4 shows the decoding speeds of different systems when using 1-best result for user reference (the distortion distance is limited to 10). The speed is the number of sentences decoded per second on the testing corpus. The hardware setting is 4G memory, 500G hard disk, Core i5 3.2GHz processor.

Method	Speed (sent/sec)
baseline	3.43
ST CM BCP	3.07
ST CM ACP	3.23
ST CM	2.86
ST RD	2.94
ST IMT	2.78

Table 4: Decoding speed using different subtree constraints.

We can see that incorporating syntactic information did not lead to much decrease in the predicting/responding speed. This is because the subtree constraints help filtering some bad hypotheses and reduced the search space. This balanced the time cost in rule matching. We should note that in our experiments the testing corpus is parsed in advance. In practice, if the corpus cannot be acquired in advance, then each source sentence should be input to the IPMT procedure after parsing.

5.3 Comparison with Other Work

Some researchers proposed methods (Quirk et al., 2005; Marton and Resnik, 2008; Shen et al., 2008; Gao et al., 2011) that introduce the syntactic information to the phrase-based MT system on the basis of hierarchical phrase-based models. In these methods, hierarchical phrases rather than flat phrases are employed. However, the decoding is not in a left-to-right manner, and has difficulty in applying to the prefix matching process of IPMT systems.

The work more similar with ours are those in (Collins et al., 2005; Wang et al., 2007; Galley and Manning, 2008; Hunter and Resnik, 2010). Galley and Manning (2008) adopts a left-to-right shift-reduce method to build hierarchical structures for flat phrases. But this method is “formally syntactic-based” rather than “linguistically syntactic-based”. Collins et al. (2005) and Wang et al. (2007) perform pre-ordering in the pre-processing stage instead of deciding the phrase orders in the decoding stage. This strategy may remove the translation hypotheses that match the prefix too early. Hunter and Resnik (2010) directly introduce the source-language syntactic constraints into the decoding of phrase-based MT system, which are almost the same as our work. However, this method builds an independent syntactic re-ordering model and scores the hypotheses through features. In fact, in the IPMT systems the user clearly gives the correct prefix, which can act as an explicit constraint that the hypotheses must match. So it is more suitable to use the syntactic constraints in the form of rules.

6 Conclusion

This paper deepened the constraints of prefix in state-of-the-art PBM-based IPMT system. We built the mathematical framework of IPMT model based on the syntactic alignment between the source sentence and the prefix. On the basis of the framework we proposed a method that introduces source-language syntactic information to hypothesis selection in the form of soft constraints, evaluating the hypotheses from the aspects of subtree completion and subtree order. We also proposed the strategies to avoid the matching rules from being too strict, making the system tolerant to the parsing mistakes. Experimental results proved that our method reduced the human-computer interaction times while the decoding speed does not decrease obviously.

Acknowledgements

This work is supported by the National Natural Science Foundation of China (No. 61402299). We would like to thank the anonymous reviewers for their insightful and constructive comments. We also want to thank Yapeng Zhang for help in the preparation of experimental systems in this paper.

References

- Sergio Barrachina, Oliver Bender, Francisco Casacuberta, Jorge Civera, Elsa Cubel, Shahram Khadivi, Antonio Lagarda, Hermann Ney, Jesús Tomás, Enrique Vidal, and Juan-Miguel Vilar. 2009. Statistical Approaches to Computer-Assisted Translation. *Computational Linguistics*, 35(1):3–28.
- Jorge Civera, Elsa Cubel, Antonio L. Lagarda, David Picó, Jorge González, Enrique Vidal, Francisco Casacuberta, Juan M. Vilar, and Sergio Barrachina. 2004. From Machine Translation to Computer Assisted Translation using Finite-State Models. In *Proceedings of the EMNLP*, pages 349–356.
- Michael Collins, Philipp Koehn, and Ivona Kucerov. 2005. Clause Restructuring for Statistical Machine Translation. In *Proceedings of the ACL*, pages 531–540.
- George Foster, Pierre Isabelle, and Pierre Plamondon. 1997. Target-text Mediated Interactive Machine Translation. *Machine Translation*, 12(1):175–194.
- Michel Galley and Christopher D. Manning. 2008. A Simple and Effective Hierarchical Phrase Reordering Model. In *Proceedings of the EMNLP*, pages 848–856.
- Yang Gao, Philipp Koehn, and Alexandra Birch. 2011. Soft Dependency Constraints for Reordering in Hierarchical Phrase-based Translation. In *Proceedings of the EMNLP*, pages 857–868.
- Jesús González-Rubio, Daniel Ortiz-Martínez, José-Miguel Benedí, and Francisco Casacuberta. 2013. Interactive Machine Translation using Hierarchical Translation Models. In *Proceedings of the EMNLP*, pages 244–254.
- Tim Hunter and Philip Resnik. 2010. Exploiting Syntactic Relationships in a Phrase-based Decoder: An Exploration. *Machine Translation*, 24(2):123–140.
- Martin Kay and Gary R. Martins. 1973. The MIND System. *Natural Language Processing*.
- Philipp Koehn, 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 *Proceedings of the ACL*, pages 177–180.
- Philipp Koehn. 2004. Statistical Significance Tests for Machine Translation Evaluation. In *Proceedings of the EMNLP*, pages 388–395.
- Yuval Marton and Philip Resnik. 2008. Soft Syntactic Constraints for Hierarchical Phrased-based Translation. In *Proceedings of the ACL*, pages 1003–1011.
- Haitao Mi and Liang Huang. 2008. Forest-based Translation Rule Extraction. In *Proceedings of the EMNLP*, pages 206–214.
- Laurent Nepveu, Guy Lapalme, Philippe Langlais, and George Foster. 2004. Adaptive Language and Translation Models for Interactive Machine Translation. In *Proceedings of the EMNLP*, pages 190–197.
- Franz Josef Och, Richard Zens, and Hermann Ney. 2003. Efficient Search for Interactive Statistical Machine Translation. In *Proceedings of the EACL*, pages 287–293.
- Franz Josef Och. 2003. Minimum Error Rate Training in Statistical Machine Translation. In *Proceedings of the ACL*, pages 160–167.
- Daniel Ortiz-Martínez, Ismael García-Varea, and Francisco Casacuberta. 2009. Interactive Machine Translation based on Partial Statistical Phrase-based Alignments. In *Proceedings of the RANLP*, pages 330–336.
- Chris Quirk, Arul Menezes, and Colin Cherry. 2005. Dependency Treelet Translation: Syntactically-informed Phrasal SMT. In *Proceedings of the ACL*, pages 271–279.
- German Sanchis-Trilles, Daniel Ortiz-Martínez, and Jorge Civera. 2008. Improving Interactive Machine Translation via Mouse Actions. In *Proceedings of the EMNLP*, pages 485–494.
- Libin Shen, Jinxi Xu, and Ralph Weischedel. 2008. A New String-to-Dependency Machine Translation Algorithm with a Target Dependency Language Model. In *Proceedings of the ACL*, pages 577–585.
- Jesús Tomás and Francisco Casacuberta. 2006. Statistical Phrase-based Models for Interactive Computer-Assisted Translation. In *Proceedings of the COLING/ACL*, pages 835–841.
- Chao Wang, Michael Collins, and Philipp Koehn. 2007. Chinese Syntactic Reordering for Statistical Machine Translation. In *Proceedings of the EMNLP*, pages 737–745.

- Tong Xiao, Jingbo Zhu, Hao Zhang, and Qiang Li. 2012. Niutrans: An Open Source Toolkit for Phrase-based and Syntax-based Machine Translation. In *Proceedings of the ACL*, pages 19–24.
- Jonathan Yamron, James Baker, Paul Bamberg, Haakon Chevalier, Taiko Dietzel, John Elder, Frank Kampmann, Mark Mandel, Linda Manganaro, Todd Margolis, and Elizabeth Steele. 1993. Lingstat: An Interactive, Machine-aided Translation System. In *Proceedings of the workshop on Human Language Technology*, pages 191–195.
- Remi Zajac. 1988. Interactive Translation: A New Approach. In *Proceedings of the COLING*, pages 785–790.
- Ying Zhang, Stephan Vogel, and Alex Waibel. 2004. Interpreting BLEU/NIST Scores: How Much Improvement Do We Need to Have a Better System? In *Proceedings of the LREC*, pages 2051–2054.
- Hui Zhang, Min Zhang, Haizhou Li, Aiti Aw, and Chew Lim Tan. 2009. Forest-based Tree Sequence to String Translation Model. In *Proceedings of the ACL*, pages 172–180.