

# Are Ambiguous Conjunctions Problematic for Machine Translation?

Maja Popović and Sheila Castilho

ADAPT Centre,

School of Computing, DCU

Ireland

{firstname.lastname}@adaptcentre.ie

## Abstract

The translation of ambiguous words still poses challenges for machine translation. In this work, we carry out a systematic quantitative analysis regarding the ability of different machine translation systems to disambiguate the source language conjunctions “but” and “and”. We evaluate specialised test sets focused on the translation of these two conjunctions. The test sets contain source languages that do not distinguish different variants of the given conjunction, whereas the target languages do. In total, we evaluate the conjunction “but” on 20 translation outputs, and the conjunction “and” on 10. All machine translation systems almost perfectly recognise one variant of the target conjunction, especially for the source conjunction “but”. The other target variant, however, represents a challenge for machine translation systems, with accuracy varying from 50% to 95% for “but” and from 20% to 57% for “and”. The major error for all systems is replacing the correct target variant with the opposite one.

## 1 Introduction

Ambiguous words are often difficult to translate automatically, even by the state-of-the-art neural machine (NMT) systems. Whereas the NMT approach significantly improved fluency (grammar) of MT outputs compared to the previous state-of-the-art statistical phrase-based (PBMT) models, adequacy (meaning preservation) is still often problematic (Castilho et al., 2017; Klubička et al., 2018). Adequacy is even more problematic for ambiguous words which have two or more meanings depending on the context.

Therefore, ambiguity of nouns, verbs and pronouns has been investigated extensively in recent years (Guillou et al., 2018; Müller et al., 2018; Rios Gonzales et al., 2017, 2018). However, to the best of our knowledge, no results for ambiguity of conjunctions have been reported so far. The only work dealing with conjunctions and machine translation (Huang, 1983) explores conjunction scope for rule-based MT systems and does not address the ambiguity. It should be noted, though, that the conjunction ambiguity is more structural than lexical: it is mainly related to certain aspects of grammar involving the arrangement of words and word types. Therefore, the conjunction ambiguity is related more to fluency than it is to adequacy.

In this work, we present the results of quantitative analysis addressing machine translation of two potentially ambiguous conjunctions, “but” and “and”. The analysis of the conjunction “but” is carried out for {English,French}-into-{Spanish,German,Serbian,Croatian} translation directions, and the conjunction “and” is analysed on {English,Portuguese}-into-{Serbian,Croatian} outputs. Evaluation is carried out on specialised test sets<sup>1</sup> designed for evaluating translation of these ambiguous conjunctions. Instead of comparing the translation output with a reference human translation, our evaluation is based on presence or absence of the correct target language conjunction variant in the translation output. For a small number of sentences with both or with none of the target variants (about 1-2%), manual inspection is carried out.

	<i>but</i> <sub>1</sub>	<i>but</i> <sub>2</sub>
en	We wanted to go to the beach, <b>but</b> we went back to the hotel.	We didn't want to go to the hotel, <b>but</b> to the beach.
de	Wir wollten zum Strand gehen, <b>aber</b> wir sind zurück zum Hotel gegangen.	Wir wollten nicht zum Hotel, <b>sondern</b> zum Strand.
es	Queríamos ir a la playa, <b>pero</b> hemos vuelto al hotel.	No queríamos ir al hotel, <b>sino</b> a la playa.
sr (hr)	Hteli smo da idemo na plažu <b>ali</b> smo se vratili u hotel.	Nismo hteli da idemo u hotel <b>nego</b> na plažu.
en	She will not come <b>but</b> she can call.	She will not come <b>but</b> call.
es	No va a venir, <b>pero</b> puede llamar.	No va a venir, <b>sino</b> a llamar.
de	Sie kommt nicht, <b>aber</b> sie kann anrufen.	Sie kommt nicht <b>sondern</b> ruft an.
sr (hr)	Neće doći <b>ali</b> može da zove.	Neće doći <b>nego</b> će zvati.

Table 1: Examples of difference between the two variants of the English conjunction “but”.

domain	lang.	<i>but</i> <sub>1</sub>	<i>but</i> <sub>2</sub>
News	En-De	65.2	34.8
	En-Sr	79.7	20.3
	En-Hr	78.3	21.7
Subtitles	En-De	97.2	2.8
	Fr-De	96.6	3.4
	En-Sr	97.1	2.9

Table 2: Distribution of sentences requiring each of the two target language variants of the source conjunction “but” in different publicly available parallel corpora.

## 2 Related Work

Lexical ambiguity as a challenge for machine translation has received a lot of attention in recent years. Rios Gonzales et al. (2017) and Rios Gonzales et al. (2018) focus on ambiguous German nouns, while Guillou et al. (2018) and Müller et al. (2018) investigate ambiguous English pronouns. Broader linguistic evaluations presented in Burchardt et al. (2017) and Klubička et al. (2018) also include ambiguity, but conjunctions are not mentioned in any context.

Syntactic ambiguity for rule-based English-to-Bulgarian machine translation is investigated in Pericliev (1984), but these ambiguities are not related to conjunctions. Ambiguity of conjunctions “and” and “or” is investigated for requirement specifications in Sharma et al. (2014) and for legal texts in Adams and Kaye (2006), however without any relation to (either human or machine) translation.

<sup>1</sup><https://github.com/m-popovic/evaluating-ambiguous-conjunctions-MT>

The first work dealing with conjunctions and machine translation is described in Huang (1983). It explores conjunction scope for English parser to be used in rule-based MT systems, but it does not address the ambiguity. Another work related to conjunctions and machine translation is the work of Xu et al. (2014), who proposes using conjunctions for Chinese sentence segmentation in order to achieve better translation quality. Some problems with translating conjunctions by phrase-based machine translation systems involving South Slavic languages are mentioned in Popović and Arčan (2015), but without any systematic quantitative analysis.

To the best of our knowledge, our work represents the first experiments related to ambiguous conjunctions and machine translation. We report the results of an extensive evaluation showing that certain conjunction ambiguities pose a challenge to the state-of-the-art machine translation systems.

## 3 Ambiguity of “But” and “And”

### 3.1 Conjunction “But”

In some languages, there are two possible variants of the conjunction “but”. One variant, *but*<sub>1</sub>, can be used after either a positive or a negative clause. The other variant, *but*<sub>2</sub>, is used after a negative clause when expressing a contradiction. The first clause in the sentence must contain a negation marker, and the second part of the sentence must contradict the first part of the sentence.

Three examples can be seen in Table 1. The sentences on the left have the same context, same or similar meaning, and contain similar words as

	<i>and</i> <sub>1</sub>	<i>and</i> <sub>2</sub>
en	The walls <b>and</b> the door are white.	The walls are white <b>and</b> the door is black.
hr (sr)	Židovi <b>i</b> vrata su bijeli.	Židovi su bijeli <b>a</b> vrata su crna.
en	I studied for the whole day <b>and</b> I learned a lot.	I studied for the whole day <b>and</b> I didn't learn anything.
hr (sr)	Učio sam cijeli dan <b>i</b> svašta naučio.	Učio sam cijeli dan <b>a</b> ništa nisam naučio.
en	Years passed, and he came back.	Years passed, and he still hadn't come back.
hr (sr)	Prošle su godine <b>i</b> on se vratio.	Prošle su godine <b>a</b> on se još nije vratio.
en	Who is this and what is he doing here?	And what is he doing here?
hr (sr)	Tko je to <b>i</b> što on radi ovdje?	<b>A</b> što on radi ovdje?

Table 3: Examples of difference between the two variants of the English conjunction “and”.

domain	lang.	<i>and</i> <sub>1</sub>	<i>and</i> <sub>2</sub>
News	En-Sr	60.0	40.0
	En-Hr	59.4	40.6
Subtitles	En-Sr	62.2	37.8
	Pt-Hr	60.4	39.6

Table 4: Distribution of sentences with two types of conjunctions “and” in different publicly available parallel corpora.

the sentences on the right. Nevertheless, the conjunction “but” in all sentences on the left should be translated into *but*<sub>1</sub> and in those on the right as *but*<sub>2</sub>. This also illustrates the previously mentioned structural nature of conjunction ambiguity.

Generally, sentences with the first variant, *but*<sub>1</sub>, can be found more frequently in the data. Table 2 presents the distribution of the two types of sentences with the conjunction “but” found in publicly available data for several language pairs in two domains: news and subtitles.

### 3.2 Conjunction “And”

Some target languages, such as Serbian and Croatian, distinguish two variants of the conjunction “and”. The first variant, *and*<sub>1</sub>, is used to connect non-contrasting actions or ideas, for example to indicate that one action follows another in the chronological order, or that one idea is the expected result of another. The second variant, *and*<sub>2</sub>, is used to indicate that the two connected facts are different: it introduces a new or different meaning, that is, it introduces an idea that is different or opposite to the idea that is desired, expected or stated previously. Both variants are used to start a new sentence or clause that continues or adds to a previous sentence or clause, however *and*<sub>2</sub> is adding some new, different or unexpected facts.

Four examples can be seen in Table 3. Similarly to the examples for the conjunction “but”, the sentences on the left have similar meaning and contain similar words as the sentences on the right, but all “ands” on the left should be translated into *and*<sub>1</sub> and those on the right into *and*<sub>2</sub>.

Table 4 presents the distribution of the two types of conjunction “and” in publicly available news and subtitles data. Again, the first variant, *and*<sub>1</sub>, is more frequent, although the difference is smaller than between the two variants of the conjunction “but”.

## 4 Experimental Set-Up

### 4.1 Test Sets

In order to estimate a system’s capability to translate ambiguous conjunctions, evaluation is performed on specialised test sets specifically designed for the conjunctions “but” and “and” and their two variants.

The test sets are created semi-automatically using the multilingual subtitles corpora<sup>2</sup> (Tiedemann, 2012). Only short segments (up to 20 words) were included, all noise was removed, and rare named entities which could introduce additional effects were avoided or replaced. Thus, about 1000 source sentences in English and in French were prepared for the conjunction “but”, and 250 source sentences in English and in Portuguese for conjunction “and”. Detailed corpus statistics are presented in Table 5.

It should be noted that although the test sets were created using a bilingual corpus, the resulting test sets do not contain any reference translations. The reason for this is twofold: on the one hand, bilingual manual filtering of noisy and complex

<sup>2</sup><http://www.opensubtitles.org/>

(a) Statistics of the test sets for the source conjunction “but”

source language	target conjunction	number of sentences	number of running words	vocabulary size	average sent. length
English	all	1066	13655	2252	12.8
	<i>but</i> <sub>2</sub>	858	11058	2043	12.9
	<i>but</i> <sub>1</sub>	208	2597	560	12.5
French	all	1010	12963	2162	12.8
	<i>but</i> <sub>2</sub>	806	10478	1823	13.0
	<i>but</i> <sub>1</sub>	204	2485	673	12.2

(b) Statistics of test sets for the source conjunction “and”

source language	target conjunction	number of sentences	number of running words	vocabulary size	average sent. length
English	all	258	3217	769	12.5
	<i>and</i> <sub>2</sub>	206	2651	691	12.9
	<i>and</i> <sub>1</sub>	52	566	248	10.9
Portuguese	all	250	2763	908	11.0
	<i>and</i> <sub>2</sub>	199	2218	767	11.1
	<i>and</i> <sub>1</sub>	51	546	264	10.7

Table 5: Statistics of the test sets for (a) conjunction “but” and (b) conjunction “and”: number of sentences, number of running words, vocabulary size and average sentence length.

content would be very time and resource consuming. On the other hand, reference translations are not really needed, because we are interested only in conjunction disambiguation, therefore, checking the conjunction in the translation hypothesis is sufficient.

In order to encourage and enable future research on the topic, the developed test sets are made publicly available.<sup>3</sup>

## 4.2 MT Outputs

The English and the French test sets for the conjunction “but” were translated into four target languages that distinguish the two variants *but*<sub>1</sub> and *but*<sub>2</sub> in the same way, namely Spanish, German, Serbian and Croatian. The English and the Portuguese test sets for the conjunction “and” were translated into Serbian and Croatian. For all translation directions, two publicly available on-line systems, “Google Translate”<sup>4</sup> and “Bing Translator”<sup>5</sup> are used. All on-line translations were generated between 26th and 30th April 2019.

In addition to this, for English-to-German and English-to-Serbian translation, two internal systems trained on much smaller amounts of data

<sup>3</sup><https://github.com/m-popovic/evaluating-ambiguous-conjunctions-MT>

<sup>4</sup><https://translate.google.com/>

<sup>5</sup><https://www.bing.com/translator>

from the news domain were available. Two English-to-German systems, one NMT and one PBMT, are trained on one million parallel sentences from the WMT<sup>6</sup> data. Two English-to-Serbian systems, also one NMT and one PBMT, are trained on the SETimes corpus (Tyers and Alperen, 2010) containing about 200k sentence pairs.

In total, the conjunction “but” is analysed on 20 MT outputs, and the conjunction “and” on 10 MT outputs.

## 4.3 Evaluation

The majority of sentences are checked automatically, however for a small number of sentences a manual inspection is needed. For each sentence, there are four possible outcomes of the automatic evaluation:

- only the correct conjunction is found  
⇒ correct
- only the opposite conjunction is found  
⇒ incorrect
- both conjunctions are found  
⇒ manual inspection

<sup>6</sup><http://www.statmt.org/wmt17/>

- none of the two conjunctions are found  
⇒ manual inspection

Manual inspection is carried out in the following way: if the structure of a sentence with additional or without any conjunctions is correct, then the sentence is considered correct.

All errors which are not related to the conjunction choice are ignored, both by automatic and by manual evaluation.

## 5 Results

The results for both conjunctions and all language pairs are presented in the form of percentage of sentences automatically identified as correct (“aut.”), identified as correct after both automatic check and manual inspection (“full”), and automatically identified as incorrect because the source conjunction is translated into the opposite conjunction (“opposite”).

### 5.1 Conjunction “But”

The results for the source conjunction “but” can be seen in Table 6.

**Target variant  $but_1$ :** Recognising the target conjunction  $but_1$  is generally not problematic: the percentage of correct sentences is almost 100% for all on-line systems, and close to 100% even for the scarce-data *setimes* systems. Apparently, there is no correlation between this accuracy and the amount of the training data, because the two *wmt* systems have lower accuracies than the two *setimes* systems. These lowest accuracies are still very high, i.e. close to 90%.

**Target variant  $but_2$ :** For this variant, the situation is, however, different. On-line systems translate it correctly in 73-95% of cases, and the predominant problem for the rest of the cases is translating it into the opposite variant  $but_1$  (5-25%). The four scarce-data systems are struggling much more with this variant, PBMT systems more than NMT ones. In addition, the influence of the amount of data is obvious, since *setimes* systems are performing much worse than *wmt* systems. For these four systems, replacing  $but_2$  with  $but_1$  is the most frequent problem too, but there were more sentences requiring manual inspection than for the on-line systems.

**Both target variants:** Manual inspection revealed that this is generally not problematic for any of the systems and language pairs: it can happen if “however”, “yet” or similar words that can

be translated as  $but_1$  are present in the source sentence.

**None of the two target variants:** Only a small number of such sentences are generated by on-line systems when translating from English if the English sentence has a structure “not only X, but Y, too”. In these cases, the sentence is paraphrased in the way “not only X, Y too”. For the French source, a number of other sentence structures are paraphrased, and the majority of these paraphrases are correct. An example can be seen in Table 7.

As for the scarce-data systems, manual inspection revealed that the *wmt* NMT system left many sentences untranslated, of which mainly those requiring the conjunction  $but_2$ . However, a number of those requiring the conjunction  $but_1$  were also left untranslated. As for PBMT systems, many of the sentences in the output used a possible translation for  $but_1$  such as “however”, which is of course not correct for sentences requiring  $but_2$ . A number of PBMT sentences had a number of other error types and a low fluency in general. Apart from this, when translating from English into Serbian and Croatian, *Google* and *Bing* often translated “but” as “except”, which is not a possible option for any of the source sentence structures.

### 5.2 Conjunction “And”

The results for the source conjunction “and” can be seen in Table 8.

**Target variant  $and_1$ :** Similarly to the translation of “but”, the target variant  $and_1$  is not really problematic for any of the systems, with almost all accuracies close to 100% and all larger than 92%. Also, there is apparently no correlation between this accuracy and the amount of the training data, since the *Bing* system has the lowest accuracy and it is certainly trained on more data than the two *setimes* systems.

**Target variant  $and_2$ :** This target variant is definitely problematic, and much more challenging than  $but_2$ : the highest accuracy is 56.3%. Similarly to the conjunction “but”, the predominant problem is replacing  $and_2$  with  $and_1$ . Also, the accuracies are lower for the scarce-training *setimes* systems. Nevertheless, one curiosity can be noted: the PBMT *setimes* system better disambiguates the conjunction “and” than the NMT *setimes* system. Given that NMT systems are generally more sensitive to the scarcity of training data (Koehn and Knowles, 2017) than PBMT



language pair	system	<i>but</i> <sub>2</sub>			<i>but</i> <sub>1</sub>		
		correct		opposite ( <i>but</i> <sub>1</sub> )	correct		opposite ( <i>but</i> <sub>2</sub> )
		aut.	full		aut.	full	
en-es	google	93.3	93.4	6.2	100	100	0
	bing	76.5	76.6	22.6	100	100	0
en-de	google	88.6	89.1	10.3	99.5	100	0
	bing	94.3	94.4	5.3	99.0	99.0	1.0
	wmt-pbmt	48.9	48.9	46.8	90.9	92.8	4.8
	wmt-nmt	60.5	60.5	28.6	84.6	84.6	0.5
en-sr	google	90.3	90.3	9.3	99.0	99.5	0.5
	bing	79.6	79.6	18.9	100	100	0
	setimes-pbmt	7.0	7.0	88.3	95.7	100	0
	setimes-nmt	53.5	53.6	43.8	95.7	98.6	1.0
en-hr	google	91.8	91.8	7.7	99.5	99.5	0.5
	bing	73.6	73.6	25.4	100	100	0
fr-es	google	92.6	93.8	6.0	100	100	0
	bing	72.0	75.1	24.1	100	100	0
fr-de	google	87.6	89.0	10.4	100	100	0
	bing	88.5	92.9	6.7	100	100	0
fr-sr	google	90.6	90.9	8.5	100	100	0
	bing	75.2	77.3	21.7	100	100	0
fr-hr	google	90.6	90.9	8.4	100	100	0
	bing	72.0	73.9	24.5	100	100	0

Table 6: Percentage of correct target language conjunctions retrieved automatically and by full evaluation, and percentage of opposite target conjunctions for the source language conjunction “but”.

source	Ce n’est pas une étoile <b>mais</b> un cristal.
source (en gloss)	It is not a star <b>but</b> a crystal.
output (es)	No es una estrella, es un cristal.
output (de)	Es ist kein Stern, es ist ein Kristall.
output (en gloss)	It’s not a star, it’s a crystal.

Table 7: Example of correct translation without any of the two conjunction variants (mostly occurring in French-to-Spanish and French-to-German on-line systems).

systems, disambiguating the conjunction “and” is probably more sensitive to the amount of the training data than disambiguating the conjunction “but”. More experiments on systems trained on different training data sizes should be carried out in the future.

**Both target variants:** Manual inspection revealed that all sentences with both “and” variants

are correct: the affected source sentences contain “too”, “as well” or similar, which can be correctly translated as *and*<sub>1</sub>. Three examples can be seen in Table 9.

**None of the two target variants:** A small number of sentences containing “neither” were correctly paraphrased in the target language so that it does not need any of the two “and” variants. These cases are found only in translations generated by on-line systems. An example can be seen in Table 10.

## 6 Summary and Outlook

We present a targeted evaluation of 20 translation outputs regarding their performance in lexical choice for the ambiguous source conjunction “but”, and 6 (10) systems regarding the source conjunction “and”. For the source conjunction “but”, we observe that all systems almost perfectly recognise the target conjunction *but*<sub>1</sub>, whereas accuracies for the other target variant *but*<sub>2</sub> are lower, and depend on the model (NMT performs better than PBMT), as well as on the amount of training data. For on-line systems trained on large amounts

language pair	system	<i>and</i> <sub>2</sub>			<i>and</i> <sub>1</sub>		
		correct aut.	full	opposite ( <i>and</i> <sub>1</sub> )	correct aut.	full	opposite ( <i>and</i> <sub>2</sub> )
en-sr	google	38.8	39.8	58.8	98.1	98.1	1.9
	bing	33.5	33.5	62.2	92.3	92.3	5.8
	setimes-pbmt	22.3	22.3	75.4	98.1	98.1	1.9
	setimes-nmt	13.6	13.6	83.5	96.2	98.1	0
en-hr	google	42.7	43.7	54.9	96.2	96.2	3.8
	bing	54.8	56.3	41.9	96.2	96.2	1.9
pt-sr	google	35.7	36.2	59.3	92.2	92.2	5.9
	bing	32.7	32.7	61.8	94.1	94.1	5.9
pt-hr	google	37.7	37.7	57.3	90.2	90.2	3.9
	bing	49.2	49.2	46.2	92.2	92.2	7.8

Table 8: Percentage of correct target language conjunctions retrieved automatically and by full evaluation, and percentage of opposite target conjunctions for the source language conjunction “and”.

source:	I can swim well <b>and so</b> do you.
output (hr) :	ja mogu plivati dobro, <b>a i</b> ti.
output (en gloss):	I can swim well, <i>and</i> <sub>2</sub> <i>and</i> <sub>1</sub> you.
source:	Holly travels a lot, <b>and</b> her sister, <b>too</b> .
output (hr)	Holly puno putuje, <b>a i</b> njena sestra.
output (en gloss):	Holly travels a lot, <i>and</i> <sub>2</sub> <i>and</i> <sub>1</sub> her sister.
source:	John likes burger, <b>and</b> he likes fish, <b>too</b> .
output (hr)	John voli hamburger, <b>a</b> voli <b>i</b> ribu.
output (en gloss):	John likes burger, <i>and</i> <sub>2</sub> likes <i>and</i> <sub>1</sub> fish.

Table 9: Examples of correct translations into Croatian with both target variants of “and”.

of data, accuracies range from 73% to 94%, and for the systems trained on small amounts between 50% and 60%. The errors for all systems are mostly caused by replacing the conjunction *but*<sub>2</sub> with the alternative conjunction *but*<sub>1</sub>.

As for the conjunction “and”, its first variant *and*<sub>1</sub> is also not problematic, even for the systems trained on small amounts of data. The variant *and*<sub>2</sub> is, however, much more challenging than *but*<sub>2</sub>, with the highest accuracy of 56.3%. In addition, disambiguation of this conjunction seems to be more sensitive to the training data scarcity

source	I don’t want you here, <b>and neither</b> does my wife.
output (hr)	Ne želim te ovdje, <b>kao ni</b> moja žena.
output (en gloss)	I don’t want you here, <b>as not</b> my wife.

Table 10: Example of correct translation into Croatian without any of the two conjunction variants.

for NMT than for PBMT. More systematic experiments including both models and different sizes of the training corpus should be carried out in the future to better understand this finding.

In addition to this, there are many other directions for future work. The current study is focused on only two ambiguous conjunctions and only two target language variants for each of them. More conjunctions and ambiguities should be investigated in the future, as well as more source and target languages. Quantitative analysis of correlation between the conjunction disambiguation and overall performance should be a part of future work, too. Also, improving a MT system by, for example, adding more parallel data containing “difficult” target conjunction variants should be investigated as well.

## 7 Acknowledgements

We would like to thank Alberto Poncelas and Pintu Lohar for providing PBMT and NMT translation outputs for English-to-German and English-to-Serbian. The ADAPT Centre for Digital Con-

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

## References

- Kenneth A. Adams and Alan S. Kaye. 2006. Revisiting the Ambiguity of "And" and "Or" in Legal Drafting. *St. John's Law Review* 90(4).
- Aljoscha Burchardt, Vivien Macketanz, Jonathan Dehdari, Georg Heigold, Jan-Thorsten Peter, and Philip Williams. 2017. A Linguistic Evaluation of Rule-Based, Phrase-Based, and Neural MT Engines. *The Prague Bulletin of Mathematical Linguistics* 108(1):159–170.
- Sheila Castilho, Joss Moorkens, Federico Gaspari, Rico Sennrich, Vilelmini Sasoni, Yota Georgakopoulou, Pintu Lohar, Andy Way, Antonio Miceli Barone, and Maria Gialama. 2017. A comparative quality evaluation of pbsmt and nmt using professional translators. In *Proceedings of MT Summit XVI*. pages 116–131.
- Liane Guillou, Christian Hardmeier, Ekaterina Lapshinova-Koltunski, and Sharid Loiciga. 2018. A Pronoun Test Suite Evaluation of the English–German MT Systems at WMT 2018. In *Proceedings of the 3rd Conference on Machine Translation (WMT 2018)*. Association for Computational Linguistics, Belgium, Brussels, pages 576–583.
- Xiuming Huang. 1983. Dealing with Conjunctions in a Machine Translation Environment. In *Proceedings of the 1st Conference on European Chapter of the Association for Computational Linguistics (EACL 1983)*. Pisa, Italy, pages 81–85.
- Filip Klubička, Antonio Toral, and Víctor M. Sánchez-Cartagena. 2018. Quantitative Fine-grained Human Evaluation of Machine Translation Systems: A Case Study on English to Croatian. *Machine Translation* 32(3):195–215.
- Philipp Koehn and Rebecca Knowles. 2017. Six Challenges for Neural Machine Translation. In *Proceedings of the First Workshop on Neural Machine Translation*. Vancouver, pages 28–39.
- Mathias Müller, Annette Rios Gonzales, Elena Voita, and Rico Sennrich. 2018. A Large-Scale Test Set for the Evaluation of Context-Aware Pronoun Translation in Neural Machine Translation. In *Proceedings of the 3rd Conference on Machine Translation (WMT 2018)*. Association for Computational Linguistics, Belgium, Brussels, pages 61–72.
- Vladimir Pericliev. 1984. Handling syntactical ambiguity in machine translation. In *Proceedings of the 10th International Conference on Computational Linguistics and 22nd Annual Meeting of the Association for Computational Linguistics (ACL 1984)*. Stanford, California, USA, pages 521–524.
- Maja Popović and Mihael Arčan. 2015. Identifying main obstacles for statistical machine translation of morphologically rich south Slavic languages. In *Proceedings of the 18th Annual Conference of the European Association for Machine Translation (EAMT 2015)*. Antalya, Turkey, pages 97–104.
- Annette Rios Gonzales, Laura Mascarell, and Rico Sennrich. 2017. Improving word sense disambiguation in neural machine translation with sense embeddings. In *Proceedings of the 2nd Conference on Machine Translation (WMT 2017)*. Copenhagen, Denmark, pages 11–19.
- Annette Rios Gonzales, Mathias Müller, and Rico Sennrich. 2018. The word sense disambiguation test suite at wmt18. In *Proceedings of the 3rd Conference on Machine Translation (WMT 2018)*. Association for Computational Linguistics, Belgium, Brussels, pages 594–602.
- Richa Sharma, Jaspreet Bhatia, and K. K. Biswas. 2014. Machine Learning for Constituency Test of Coordinating Conjunctions in Requirements Specifications. In *Proceedings of the 3rd International Workshop on Realizing Artificial Intelligence Synergies in Software Engineering (RAISE 2014)*. Hyderabad, India, pages 25–31.
- Jörg Tiedemann. 2012. Parallel data, tools and interfaces in OPUS. In *Proceedings of the 8th International Conference on Language Resources and Evaluation (LREC 2012)*. Istanbul, Turkey, pages 2214–2218.
- Francis M Tyers and Murat Serdar Alperen. 2010. South-east European Times: A parallel corpus of Balkan languages. In *Proceedings of the LREC Workshop on Exploitation of Multilingual Resources and Tools for Central and (South-) Eastern European Languages*. Valetta, Malta, pages 49–53.
- Li Fang Xu, Yun Zhu, Li Jiao Yang, and Yao Hong Jin. 2014. Research on sentence segmentation with conjunctions in patent machine translation. In *Applied Science, Materials Science and Information Technologies in Industry*. Trans Tech Publications, volume 513 of *Applied Mechanics and Materials*, pages 4605–4609.