

Designing the Business Conversation Corpus

Matīss Rikters* and Ryokan Ri* and Tong Li* and Toshiaki Nakazawa

The University of Tokyo

7-3-1 Hongo, Bunkyo-ku, Tokyo, Japan

{matiss, li0123, litong, nakazawa}@logos.t.u-tokyo.ac.jp

Abstract

While the progress of machine translation of written text has come far in the past several years thanks to the increasing availability of parallel corpora and corpora-based training technologies, automatic translation of spoken text and dialogues remains challenging even for modern systems. In this paper, we aim to boost the machine translation quality of conversational texts by introducing a newly constructed Japanese-English business conversation parallel corpus. A detailed analysis of the corpus is provided along with challenging examples for automatic translation. We also experiment with adding the corpus in a machine translation training scenario and show how the resulting system benefits from its use.

1 Introduction

There are a lot of ready-to-use parallel corpora for training machine translation systems, however, most of them are in written languages such as web crawl, news-commentary¹, patents (Goto et al., 2011), scientific papers (Nakazawa et al., 2016) and so on. Even though some of the parallel corpora are in spoken language, they are mostly spoken by only one person (in other words, they are monologues) (Cettolo et al., 2012; Di Gangi et al., 2019) or contain a lot of noise (Tiedemann, 2016; Pryzant et al., 2018). Most of the machine translation evaluation campaigns such as WMT², IWSLT³ and WAT⁴ adopt the written language, monologue or noisy dialogue parallel corpora for their translation tasks. Among them, there is only one clean, dialogue parallel corpus (Salesky et al., 2018) adopted by IWSLT in the conversational

speech translation task, nevertheless, the availability of such kind of corpus is still limited.

The quality of machine translation for written text and monologue has vastly improved due to the increase in the amount of the available parallel corpora and the recent neural network technologies. However, there is much room for improvement in the context of dialogue or conversation translation. One typical case is the translation from pro-drop language to the non-pro-drop language where correct pronouns must be supplemented according to the context. The omission of the pronouns occurs more frequently in spoken language than written language. Recently, context-aware translation models attract attention from many researchers (Tiedemann and Scherrer, 2017; Voita et al., 2018, 2019) to solve this kind of problem, however, there are almost no conversational parallel corpora with context information except noisy OpenSubtitles corpus.

Taking into consideration the factors mentioned above, a document and sentence-aligned conversational parallel corpus should be advantageous to push machine translation research in this field to the next stage. In this paper, we introduce a newly constructed Japanese-English business conversation parallel corpus. This corpus contains 955 scenarios, 30,000 parallel sentences. Table 1 shows an example of the corpus.

We choose the business conversation as the domain of the corpus because 1) the business domain is neither too specific nor too general, and 2) we think that a clean conversational parallel corpus is useful to open new machine translation research directions. We hope that this corpus becomes one of the standard benchmark data sets for machine translation.

What is unique for this corpus is that each scenario is annotated with scene information, as shown in the top of Table 1. In conversations,

*equal contribution

¹<http://www.statmt.org/wmt19/translation-task.html>

²<http://www.statmt.org/wmt19/>

³<http://workshop2019.iwslt.org>

⁴<http://lotus.kuee.kyoto-u.ac.jp/WAT/>

Scene: telephone consultation about intrafirm export

Japanese		English	
Speaker	Content	Speaker	Content
山本	もしもし、山本と申します。	Yamamoto	Hello, this is Yamamoto.
田中	販売部門の田中と申します。	Tanaka	This is Tanaka from the Department of Sales.
田中	輸出に関してご助言いただきたくお電話しました。	Tanaka	I called you to get some advice from you concerning export.
山本	はい、どのようなご用件でしょうか？	Yamamoto	Okay, what's the matter?
田中	イランの会社から遠視カメラの引き合いを受けているのですが、イランに対しては輸出制限があると新聞で読んだことがある気がして。	Tanaka	We got an inquiry from an Iranian company about our far-sight cameras, but I think I read in the newspaper that there are export restrictions against Iran.
田中	うちで売っているようなカメラなら、特に問題にならないのでしょうか？	Tanaka	Is there no problem with cameras like the ones we sell?
山本	恐れ入りますが、イランへの輸出は、かなり制限されているのが事実です。	Yamamoto	I'm afraid that the fact is, exports to Iran are highly restricted.
...

Table 1: An example of the Japanese-English business conversation parallel corpus.

the utterances are often very short and vague, therefore it is possible that they should be translated differently depending on the situations where the conversations are taking place. For example, Japanese expression 「すみません」 can be translated into several English expressions such as “Excuse me.” (when you call a store attendant), “Thank you.” (when you are given some gifts) or “I’m sorry.” (when you need to apologise). By using the scene information, it is possible to discriminate the translations, which is hard to do with only the contextual sentences. Furthermore, it might be possible to connect the scene information to the multi-modal translation, which is also hardly studied recently, such as estimating the scenes by the visual information.

The structure of this paper is as follows: we explain how the corpus is constructed in Section 2, show the fundamental analysis of the corpus in Section 3, report results of machine translation experiments in Section 4, and give a conclusion in Section 5.

2 Description and Statistics of the Corpus

The Japanese-English business conversation corpus, namely Business Scene Dialogue (BSD) corpus, is constructed in 3 steps: 1) selecting business scenes, 2) writing monolingual conversation scenarios according to the selected scenes, and 3) translating the scenarios into the other language. The whole construction process was supervised by a person who satisfies the following conditions to guarantee the conversations to be natural:

- has the experience of being engaged in lan-

Scene	Scenarios	Sentences
JA → EN		
face-to-face	165	5,068
phone call	77	2,329
general chatting	101	3,321
meeting	106	3,561
training	16	608
presentation	4	113
sum	469	15,000
EN → JA		
face-to-face	158	4,876
phone call	99	2,949
general chatting	102	2,988
meeting	103	3,315
training	9	326
presentation	15	546
sum	486	15,000

Table 2: Statistics for the corpus, where JA → EN represents scenarios which are written in Japanese then translated into English and EN → JA represents scenarios constructed in the reverse way.

guage learning programs, especially for business conversations

- is able to smoothly communicate with others in various business scenes both in Japanese and English
- has the experience of being involved in business

2.1 Business Scene Selection

The business scenes were carefully selected to cover a variety of business situations, including meetings and negotiations, as well as so-called water-cooler chats. Details are shown in Table 2. We also paid attention not to select specialised scenes which are suitable only for a limited number of industries. We made sure that all of the se-

lected scenes are generic to a broad range of industries.

2.2 Monolingual Dialogue Scenario Writing

Dialogue scenarios were monolingually written for each of the selected business scenes. Half of the monolingual scenarios were written in Japanese and the other half were written in English (15,000 sentences for each language). This is because we want to cover a wide range of lexicons and expressions for both languages in the corpus. Writing the scenarios only in one language might fail to cover useful, important expressions in the other language when they are translated in the following step.

2.3 Scenario Translation

The monolingual scenarios were translated into the other language by human translators. They were asked to make the translations not only accurate, but also as fluent and natural as a real dialogue at the same time. This principle is adopted to eliminate several common tendencies of human translators when performing Japanese-English translation on a written text. For example, Japanese pronouns are usually omitted in a dialogue, however, when the English sentences are literally translated into Japanese, the translators tend to include unnecessary pronouns. It is acceptable as a written text, but would be rather unusual as a spoken text.

3 Analysis of the Corpus

To understand the difficulty of translating conversations, we conduct an analysis regarding the newly constructed corpus. We choose to use Google Translate⁵, one of the most powerful neural machine translation (NMT) systems which are publicly available, to produce the translations.

Our primary focus is to understand how many sentences require context to be properly translated. We randomly sample 10 scenarios (322 sentences) from the corpus, and check the translations for fatal translation errors, ignoring fluency or minor grammatical mistakes. As a result, 12 sentences have errors due to phrase ambiguity that needs understanding the context, or the real-world situation, and 18 errors of pronouns due to zero anaphora, which is described in the following sec-

⁵<https://translate.google.com/> (May 2019)




Figure 1: The top 10 frequent errors of pronoun translation (fatal errors denoted in boldface)

tion, in the source language (Japanese). Now we focus on the latter errors.

3.1 Zero Anaphora

As an important preliminary, we briefly introduce a grammatical phenomenon called *zero anaphora*. In Japanese, some arguments of verbs are often omitted from the phrases when they are obvious from the context. When translating them into English, one often has to identify the referent of the omitted argument and recover it in English, as English does not allow omitting the core arguments (i.e., subject, object). In the following Japanese example, the subject of the verb 買った is omitted, but in the English translation a pronoun, for example *he*, has to be recovered. Note that the subject could be anyone, not necessarily *he*, depending on the context. The task of identifying the referent of zero anaphora is called *zero anaphora resolution*, which is one of the most difficult tasks of NLP.

太郎は 買った 牛乳を 飲んだ
Taro-SBJ buy-PST milk-OBJ drink-PST
“Taro drank the milk **he** bought.”

3.2 Quantitative Analysis

To estimate how many sentences need zero anaphora resolution in the business conversation corpus, we counted the number of sentences with the personal pronouns (e.g., 彼, 彼女, 私, あなた in Japanese, *I, you, he, she* in English) in both Japanese and English. As a result, 62% of English sentences contain personal pronouns, while only 11% of Japanese sentences do. This means about 50% of the sentences in the corpus potentially need zero reference resolution when we translate them from Japanese into English.

To reveal what kinds of zero pronouns are hard to translate, we again heuristically count the number of the translation errors of the pronouns for

Previous Source: 支店長はポールをクビにするみたいだよ。
Previous Reference: It seems like the branch manager will be firing **Paul**.
Source: 仕事もあまりしない上に、休み、早退ばかりを希望するから。
Reference: **He** doesn't work much, and **he** takes days off and asks to leave early often.
Google Translate: I do not have much work, and I would like to leave early and leave early.

Figure 2: An example of Japanese to English Google Translate output. The words in boldface are supposed to denote the same referent(Paul).

Previous Source: [Speaker1] 彼の代わりに、優秀な人が入ってくれば、僕の仕事量が減るはずなんだ。
Previous Reference: I think I can work less if there's someone excellent coming in as a replacement for him.
Source: [Speaker2] もう少しの辛抱だよ。
Reference: **You** just need a bit more patience.
Google Translate: I have a little more patience.

Figure 3: An example of Japanese to English Google Translate output. Correct translation needs the speaker information.

the entire corpus. We counted the number of the translated sentences that have pronouns different from their reference sentences. By this heuristic, we detected 3,653 errors (12% of the whole corpus). The top 10 frequent errors are shown in Figure 1.

Some errors such as *we* → *I*, *I* → *me*, might be not fatal, and not be regarded as translation errors. However, there are still many fatal errors among first, second and third-person pronouns (denoted in boldface in the graph).

Looking at the pronouns that the NMT system produced, we can see the tendency of the system to generate frequent pronouns such as *you*, *I*. This suggests that the current system tries to compensate source (Japanese) zero pronouns simply by generating frequent target (English) pronouns. When the referent is denoted in relatively infrequent pronouns in the target language, it is hard to be correctly translated. To deal with this problem, We need to develop more sophisticated systems that take context into account.

3.3 Qualitative Analysis

This section exemplifies some zero-anaphora translation errors and discusses what kind of information is needed to perform correct translation.

A translation that needs world knowledge and inference

In Figure 2, the subjects of the verbs are omitted in the source sentence 「(彼は: he) 仕事もあまりしない上に、(彼は: he) 休み、早退ばかりを希望するから」。This causes the NMT system to incorrectly translate the zero pronouns

into *I*, although they actually refer to *Paul* in the previous sentence and thus have to be translated into *he*.

Resolving these zero pronouns, however, is not straightforward, even if one has access to the information of the previous sentence. For example, to identify the subject of 「仕事もあまりしない」(doesn't work much), one has to know “laziness can lead to being fired” and thereby infer that Paul, who is about to be fired, is the subject. Existing contextual NMT systems (Voita et al., 2018; Bawden et al., 2018; Maruf et al., 2019) still do not seem to be able to handle this complexity.

A translation that needs to know who is talking

In Figure 3, again, the subject is omitted in the source sentence 「(君は: you) もう少しの辛抱だよ。」. The NMT system incorrectly translates the zero pronouns into *I*.

It is worth noting that the type of the zero pronoun differs from the one in Figure 2 in that the referent in Figure 3 does not linguistically appear within the text (called *exophora*), while the other does (*endophora*) (Brown and Yule, 1983). The referent of the zero pronoun in Figure 3 is the listener of the utterance (*you*), and it usually does not have another linguistic item (such as the name of the person) that can be referred to. Although some modality expressions and verb types can give constraints to the possible referents (Nakaiwa and Shirai, 1996), essentially, the resolution of *exophora* needs the reference to situation.

In this case, the correct translation depends on who is speaking. In the original conversation, the utterance is from Speaker 2 to Speaker 1, and

Data Set	Devel	Eval	Train
BSD	1000	1000	28,000
AMI	1000	1000	108,499
ON	1000	1000	26,439
Total	162,938		

Table 3: Training, development and evaluation data statistics.

given the context, one can infer that Speaker 2 is speaking to give a consolation to Speaker 1 and thus the subject should be *you* (Speaker 1). However, if the utterance was from Speaker 1, he would then just be complaining about his situation saying “*I just need a bit more patience*”. This example emphasises that the speaker information is essential to translate some utterances in conversation correctly.

4 Machine Translation Experiments

The BSD corpus was created with the intended use of training NMT systems. Thus, we trained NMT models using the corpus in both translation directions. As the BSD corpus is rather small for training reasonable MT systems, we also experimented with combining it with two larger conversational domain corpora. We employed translators to translate the AMI Meeting Corpus (McCowan et al., 2005) (AMI) and the English part of Onto Notes 5.0 (Weischedel et al., 2013) (ON) into Japanese with the same instructions as for translating the BSD corpus. Afterwards, we used them as additional parallel corpora in our experiments.

4.1 Data Preparation

Before training, we split each of the corpora into 3 parts - training, development and evaluation data sets. The sizes of each corpus are shown in Table 3. We used Sentencepiece (Kudo and Richardson, 2018) to create a shared vocabulary of 4000 tokens. We did not perform other tokenisation or truecasing for the training data. We used Mecab (Kudo, 2006) to tokenise the Japanese side of the evaluation data, which we used only for scoring. The English side remained as-is.

4.2 Experiment Setup

We used Sockeye (Hieber et al., 2017) to train transformer architecture models with 6 encoder and decoder layers, 8 transformer attention heads per layer, word embeddings and hidden layers of size 512, dropout of 0.2, maximum sentence length of 128 symbols, and a batch size of 1024

		JA-EN	EN-JA
BSD	SMT	1.90	5.16
	NMT	8.32	8.34
AMI, BSD, ON	SMT	7.27	5.76
	NMT	12.88	13.53
AMI, ON	SMT	2.18	5.74
	NMT	7.08	10.00

Table 4: NMT and SMT experiments using the conversational corpora. Evaluated on the Business Conversation evaluation set.

words, checkpoint frequency of 4000 updates. All models were trained until they reached convergence (no improvement for 10 checkpoints) on development data.

For contrast we also trained statistical MT (SMT) systems using the Moses (Koehn et al., 2007) toolkit and the following parameters: Word alignment using fast-align (Dyer et al., 2013); 7-gram translation models and the ‘wbe-msd-bidirectional-fe-allff’ reordering models; Language model trained with KenLM (Heafield, 2011); Tuned using the improved MERT (Bertoldi et al., 2009).

4.3 Results

Since there are almost no spaces in the Japanese raw texts, we used Mecab to tokenise the Japanese translations and references for scoring. The results in BLEU scores (Papineni et al., 2001) are shown in Table 4 along with several ablation experiments on training NMT and SMT systems using only the BSD data, all 3 conversational corpora, and excluding the BSD corpus from the training data. The results show that adding the BSD to the two larger corpora significantly improves both SMT and NMT performance. For Japanese → English using only BSD as training data achieves a higher BLEU score than using only AMI and ON, while for English → Japanese the opposite is true. Nevertheless, in both translation directions using all 3 corpora outperforms the other results.

We also evaluate the highest-scoring NMT system (trained on all corpora) on all 3 evaluation sets and report BLEU scores and ChrF2 scores (Popović, 2015) in Table 5. We do this to verify that the models are not overfitting on the BSD data, i.e. BLEU and ChrF2 scores are not significantly higher for the BSD evaluation sets when compared to the ON and AMI sets. Results on the ON evaluation set are fairly similar to the BSD results, while results on the AMI evaluation set are

Source:	では、終了する前に、この健康とストレスに関するセルフチェックシートに記入を して頂きたいと思います。
Our Best NMT:	So before we finish, I'd like to fill in the health check-streams with this health and staff check- book.
Google Translate:	I would like you to fill out this health and stress self-check sheet before you finish.
Reference:	Before we finish off, we would like you to fill out this self-check sheet about health and stress.

Figure 4: An example of Japanese to English NMT output comparing our best NMT to Google Translate.

Previous Source:	ああ、マネージャーはよく休みを取ってるみたいですよ。
Previous Reference:	Well, seems like our manager is taking quite a bit of time off.
Source:	自分が取らないと、他の人が取らないだろと思ってるんでしょうね。
Our Best NMT:	You think other people won't take it if they don't.
Google Translate:	If you don't take it, you might think that other people won't take it.
Reference:	Maybe he thinks if he doesn't take any, then nobody else will.

Figure 5: An example of Japanese to English NMT output, where a context-aware system could be more useful.

	BLEU	ChrF2
ON		
JA → EN	9.08	34.38
EN → JA	14.52	19.73
AMI		
JA → EN	20.88	46.93
EN → JA	23.35	30.25
BSD		
JA → EN	12.88	35.37
EN → JA	13.53	21.97

Table 5: BLEU and ChrF2 scores for all three evaluation data sets using the NMT system trained on all data.

noticeably higher. This can be explained by the fact that the AMI training data set is approximately four times larger than the BSD training data set, and the ON training data set is about the same size as the BSD set.

4.4 Machine Translation Examples

In Figure 4 we can see one of the difficult situations mentioned in Section 3.3, where MT systems find it challenging to generate the correct pronouns in the translation. Of the three pronouns that are in the reference (we, we, you), each system translates one correctly and fails to translate the rest - both systems generate *I* where it should have been *we*, but our system completely omits *you* while Google Translate generates *you* where it should have been *we*.

Figure 5 shows an example where both - our translation and the one from Google Translate are acceptable at the sentence-level, but when looking at the previous source and reference it becomes clear that different personal pronouns should have been used. Our system did generate “they” in

the second part of the sentence, which could be a more casual alternative to “he”, but both systems still failed to find the correct pronoun for the first part by producing “you” instead of “he”. This is an issue that can not be fully resolved by using sentence-level MT and requires a document-level or context-aware solutions.

5 Conclusion

In this paper, we presented a parallel corpus of English-Japanese business conversations. The intended use-cases for the corpus are machine translation system training and evaluation. We describe the corpus in detail and indicate which linguistic phenomena are challenging to translate even for modern MT systems. We also show how adding the BSD corpus to machine translation system training helps to improve translation output of conversational texts.

We point out several examples, where sentence-level MT is unable to produce the correct translation due to lack of context from previous sentences. As the corpus is both - sentence-aligned and document-aligned, we hope that it gets used and inspires new future work such directions as document-level and context-aware neural machine translation, as well as analysing other linguistic phenomena that are relevant to translating conversational texts.

In the near future, we plan to release the full set of business conversational corpora. The set will contain all 3 corpora described in section 4 - an extended version of the Business Scene Dialogue corpus as well as parallel versions of the

AMI Meeting Corpus and Onto Notes 5.0.

Acknowledgements

This work was supported by “Research and Development of Deep Learning Technology for Advanced Multilingual Speech Translation”, the Commissioned Research of National Institute of Information and Communications Technology (NICT), JAPAN.

References

- Rachel Bawden, Rico Sennrich, Alexandra Birch, and Barry Haddow. 2018. Evaluating Discourse Phenomena in Neural Machine Translation. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Nicola Bertoldi, Barry Haddow, and Jean-Baptiste Fouet. 2009. Improved Minimum Error Rate Training in Moses. *The Prague Bulletin of Mathematical Linguistics*, 91(1):7—16.
- Gillian Brown and George Yule. 1983. *Discourse Analysis*. Cambridge University Press, Cambridge.
- Mauro Cettolo, Christian Girardi, and Marcello Federico. 2012. Wit³: Web inventory of transcribed and translated talks. In *Proceedings of the 16th Conference of the European Association for Machine Translation (EAMT)*, pages 261–268, Trento, Italy.
- Mattia Antonino Di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2019. MuST-C: a Multilingual Speech Translation Corpus. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, Minneapolis, MN, USA.
- Chris Dyer, Victor Chahuneau, and Noah A Smith. 2013. A simple, fast, and effective reparameterization of ibm model 2. In *Proceedings of NAACL-HLT 2013*. Association for Computational Linguistics.
- Isao Goto, Bin Lu, Ka Po Chow, Eiichiro Sumita, and Benjamin Tsou. 2011. Overview of the patent machine translation task at the ntcir-9 workshop. In *Proc. of NTCIR-9 Workshop Meeting*, pages 559–578.
- Kenneth Heafield. 2011. Kenlm: Faster and smaller language model queries. In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, pages 187–197. Association for Computational Linguistics.
- Felix Hieber, Tobias Domhan, Michael Denkowski, David Vilar, Artem Sokolov, Ann Clifton, and Matt Post. 2017. *Sockeye: A toolkit for neural machine translation*. *ArXiv e-prints*.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. *Moses: open source toolkit for statistical machine translation*.
- Taku Kudo. 2006. Mecab: Yet another part-of-speech and morphological analyzer. <http://mecab.sourceforge.jp>.
- Taku Kudo and John Richardson. 2018. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71.
- Sameen Maruf, André FT Martins, and Gholamreza Haffari. 2019. Selective attention for context-aware neural machine translation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3092–3102.
- Iain McCowan, Jean Carletta, Wessel Kraaij, Simone Ashby, S Bourban, M Flynn, M Guillemot, Thomas Hain, J Kadlec, Vasilis Karaiskos, et al. 2005. The ami meeting corpus. In *Proceedings of the 5th International Conference on Methods and Techniques in Behavioral Research*, volume 88, page 100.
- Hiromi Nakaiwa and Satoshi Shirai. 1996. Anaphora Resolution of Japanese Zero Pronouns with Deictic Reference. In *Proceedings of International Conference on Computational Linguistics*.
- Toshiaki Nakazawa, Manabu Yaguchi, Kiyotaka Uchimoto, Masao Utiyama, Eiichiro Sumita, Sadao Kurohashi, and Hitoshi Isahara. 2016. Aspec: Asian scientific paper excerpt corpus. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC 2016)*, pages 2204–2208, Portorož, Slovenia. European Language Resources Association (ELRA).
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2001. *BLEU*. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics - ACL '02*, page 311, Morristown, NJ, USA. Association for Computational Linguistics.
- Maja Popović. 2015. chrF: character n-gram f-score for automatic mt evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395.
- R. Pryzant, Y. Chung, D. Jurafsky, and D. Britz. 2018. JESC: Japanese-English Subtitle Corpus. *Language Resources and Evaluation Conference (LREC)*.
- Elizabeth Salesky, Susanne Burger, Jan Niehues, and Alex Waibel. 2018. Towards fluent translations from disfluent speech. In *Proceedings of the IEEE Workshop on Spoken Language Technology (SLT)*, Athens, Greece.

- Jörg Tiedemann. 2016. [Finding alternative translations in a large corpus of movie subtitle](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, pages 3518–3522, Portorož, Slovenia. European Language Resources Association (ELRA).
- Jörg Tiedemann and Yves Scherrer. 2017. [Neural machine translation with extended context](#). In *Proceedings of the Third Workshop on Discourse in Machine Translation*, pages 82–92, Copenhagen, Denmark. Association for Computational Linguistics.
- Elena Voita, Rico Sennrich, and Ivan Titov. 2019. [When a good translation is wrong in context: Context-aware machine translation improves on deixis, ellipsis, and lexical cohesion](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1198–1212, Florence, Italy. Association for Computational Linguistics.
- Elena Voita, Pavel Serdyukov, Rico Sennrich, and Ivan Titov. 2018. [Context-aware neural machine translation learns anaphora resolution](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1264–1274, Melbourne, Australia. Association for Computational Linguistics.
- Ralph Weischedel, Martha Palmer, Mitchell Marcus, Eduard Hovy, Sameer Pradhan, Lance Ramshaw, Nianwen Xue, Ann Taylor, Jeff Kaufman, Michelle Franchini, et al. 2013. Ontonotes release 5.0 ldc2013t19. *Linguistic Data Consortium, Philadelphia, PA*, 23.