

# Towards Personalised and Document-level Machine Translation of Dialogue

Sebastian T. Vincent

Department of Computer Science, University of Sheffield  
Regent Court, 211 Portobello, Sheffield, S1 4DP, UK  
stvincent1@sheffield.ac.uk

## Abstract

State-of-the-art (SOTA) neural machine translation (NMT) systems translate texts at sentence level, ignoring context: intra-textual information, like the previous sentence, and extra-textual information, like the gender of the speaker. Because of that, some sentences are translated incorrectly. Personalised NMT (PersNMT) and document-level NMT (DocNMT) incorporate this information into the translation process. Both fields are relatively new and previous work within them is limited. Moreover, there are no readily available robust evaluation metrics for them, which makes it difficult to develop better systems, as well as track global progress and compare different methods. This thesis proposal focuses on PersNMT and DocNMT for the domain of dialogue extracted from TV subtitles in five languages: English, Brazilian Portuguese, German, French and Polish. Three main challenges are addressed: (1) incorporating extra-textual information directly into NMT systems; (2) improving the machine translation of cohesion devices; (3) reliable evaluation for PersNMT and DocNMT.

## 1 Introduction

Neural machine translation (NMT) represents state-of-the-art (SOTA) results in many domains (Sutskever et al., 2014; Vaswani et al., 2017; Lample et al., 2020), with some authors claiming human parity (Hassan et al., 2018). However, traditional methods process texts in short units like the utterance or sentence, isolating them from the entire dialogue or document, as well as ignoring extra-textual information (e.g. who is speaking, who they are talking to). This can result in a translation hypothesis’ meaning or function being significantly different from the reference or make the text incohesive or illogical. For instance, the sentence in Polish “*Nie*

*poszłam.*” (“*I didn’t go.*”<sup>1</sup>) incorporates gender information in the word *poszłam* (*went<sub>fem</sub>*) – as opposed to *poszedłem* (*went<sub>masc</sub>*) – while the English verb does not incorporate such information. When translating “*I didn’t go.*” into Polish, the machine translation (MT) model must guess the gender of *I*, as this information is not rendered in the English sentence. Rescigno et al. (2020) show that when commercial MT engines need to “guess” the gender of a word, they do so by making implications based on its co-occurrence with other words in the training data. Since training data is often biased (Stanovsky et al., 2020), MT models will reproduce these biases, further propagating and reinforcing them. Clearly, research on context-aware machine translation is needed.

Sentence-level NMT (SentNMT) is especially harmful in the domain of dialogue, where most utterances rely on previously spoken ones, both in content and in style. The way in which an interlocutor chooses to express themselves depends on what they perceive as the easiest for the other person to understand (Pickering and Garrod, 2004). Dialogue is naturally cohesive (Halliday and Matthiessen, 2013), i.e. rid of redundancies, confusing redefinition of terms and unclear references. Part of what makes a conversation fluent is the links between its elements, which SOTA NMT models fail to capture. For instance, the latter utterance in the following exchange: “*They put something on the roof.*” “*What?*” translates to Polish as “*Co takiego?*” (“*What something?*”). The translation uses information unavailable in the utterance itself, i.e. the fact that *what* refers to the noun *something*. A sentence-level translation of *What?* would just be *Co?*, which is more universal, but also more ambiguous. Simply put, even when SentNMT pro-

<sup>1</sup>All examples throughout the report have been generated using Google Translate <http://translate.google.com/>, accessed 26 Nov 2020.

duces a feasible translation, its context agnosticism may prevent it from producing a far better one.

There are growing appeals for developing NMT systems capable of incorporating additional information into hypothesis production: personalised NMT for extra-textual information (e.g. Sennrich et al., 2016; Elaraby et al., 2018; Vanmassenhove et al., 2018) and document-level NMT for intra-textual information (e.g. Bawden, 2019; Tiedemann and Scherrer, 2017; Zhang et al., 2018; Lopes et al., 2020). Evaluation methods predominant within both areas vary vastly from paper to paper, suggesting that for these applications a robust evaluation metric is not readily available. This view is further strengthened by the fact that Hassan et al. (2018), when assessing their MT for human parity, ignored document-level evaluation completely. Läubli et al. (2018) later disputed this choice, showing that professional annotators still overwhelmingly prefer human translation at the level of the document, and therefore human parity has not yet been achieved. This case study shows how much a robust and widely accepted document-level metric is needed.

Currently, researchers working on PersNMT and DocNMT conduct evaluation primarily by reporting the BLEU score for their systems. But they also commonly assert that the metric cannot reliably judge fine-grained translation improvements coming from context inclusion. As a way out, some of them report accuracy on specialised test suites (e.g. Kuang et al., 2018; Bawden, 2019; Voita et al., 2020) or manual evaluation. Although both have limited potential for generalisation, their attention to detail makes them superior tactics of evaluation for applications such as PersNMT and DocNMT.

In this work we utilise TV subtitles, a context-rich domain, in order to investigate whether MT of dialogue can be improved: directly, by **enhancing document coherence and cohesion** through incorporation of intra- and extra-textual information into translation, and indirectly, by **designing suitable evaluation methods for PersNMT and DocNMT**. Dialogue extracted from TV content is an attractive domain for two reasons: (1) there is an abundance of parallel dialogue corpora extracted purely from subtitles, and (2) the data is rich in or could potentially be annotated for a range of meta information such as the gender of the speaker.

In Section 2, we discuss relevant contextual phenomena. We then present the research on PersNMT

and DocNMT, and the applicability of MT evaluation metrics to both. In Section 3 we delineate the research questions, the work conducted so far and our plans. Section 4 concludes the paper.

## 2 Background

### 2.1 Contextual phenomena

Two types of contextual phenomena relevant for MT of dialogue are explored: **cohesion** phenomena (related to information that can be found in the text) and **coherence** phenomena (related to the context of situation, which we consider to be external to the text). We emphasise that the phenomena explored below represent a subset of cohesion and coherence constituents, and that our interest in them arises from the difficulties they pose for MT of dialogue.

**Cohesion phenomena** Humans introduce cohesion into speech or written text in three ways: by choosing words related to those that were used before (*lexical cohesion*), by omitting parts of or whole phrases which can be unambiguously recovered by the addressee (*ellipsis and substitution*) and by referring to elements with pronouns or synonyms that the speaker judges recoverable from somewhere else in text (*reference*) (Halliday and Matthiessen, 2013). Cohesion phenomena effectively constitute links in text, whether within one utterance or across several. Figure 1 shows examples of how they can be violated by MT.

Cohesion-related tasks such as coreference or ellipsis resolution have attracted great interest in the recent years (e.g. Rønning et al., 2018; Jwalapuram et al., 2020). Previous research on cohesion within DocNMT has revealed that verb phrase ellipsis, coreference and reiteration (a type of lexical cohesion) may be particularly erroneous in MT (e.g. Tiedemann and Scherrer, 2017; Bawden et al., 2018; Voita et al., 2020).

**Coherence phenomena** Coherence is consistency of text with the context of situation (Halliday and Hasan, 1976). MT of dialogue may be erroneous due to models not having access to extra-textual information<sup>2</sup>, e.g.: (a) speaker gender and number, (b) interlocutor gender and number, (c) social addressing, and (d) discourse situation. Different languages may render such phenomena differently, e.g. formality in German is expressed

<sup>2</sup>Note: the focus here is on *sentence-level translation utilising extra-textual context*.

EN	“It’s just a <u>social call</u> .” “A <u>social call</u> ?”
PL <sub>MT</sub>	“To tylko <b>spotkanie towarzyskie</b> .” “ <b>Połączenie towarzyskie</b> ?” (“It’s just a social gathering.” “A social call?”)
PL <sub>ref</sub>	“To tylko <b>spotkanie towarzyskie</b> .” “ <b>Spotkanie towarzyskie</b> ?” (“It’s just a social gathering.” “A social gathering?”)
<hr/>	
EN	I love it. We all <u>do</u> [=love it].
PL <sub>MT</sub>	<b>Kocham</b> to. Wszyscy <b>to robimy</b> . (“We all do it.”)
PL <sub>ref</sub>	<b>Kocham</b> to. Wszyscy <b>to kochamy</b> . (“We all love it.”)

Figure 1: Mistranslations of cohesion phenomena in translations. In the top example, *social call* is reiterated in source and reference, while MT opts for two different phrases, thereby decreasing lexical cohesion. The bottom example is verb phrase ellipsis, which does not exist in Polish and hence requires that the antecedent verb is repeated.

through the formal pronoun *Sie* (e.g. “Are you hungry?” becomes “*Bist du hungrig*?” when informal and “*Sind Sie hungrig*?” when formal), while in Polish inflections of the pronoun *Pan/Pani/Państwo* (“Mr/Mrs/Mr and Mrs”), the formal equivalent of *ty/wy* (“you”) are used. Then, as observed by [Kranich \(2014\)](#), some languages (such as English) prefer to express formality through politeness via word choices (e.g. *pleased* is a more formal *happy*)<sup>3</sup>.

## 2.2 Personalised Neural Machine Translation

In PersNMT, the aim is to develop a system  $F$  capable of executing the following operation:

$$F(x_{SL}, e, TL) = x_{TL,e}$$

where  $x$  is the source sentence,  $p$  is the extra-textual information (e.g. speaker gender) and  $SL, TL$  are source and target language, respectively;  $x_{TL,e}$  is then a contextual translation of  $x_{SL}$ .

This formulation is inspired by previous work within the area. [Sennrich et al. \(2016\)](#) control the formality of a sentence translated from English to German by using a side constraint. The model is trained on pairs of sentences  $(x_i, y_i)$ , where  $y_i$  is either formal or informal, and a corresponding tag is prepended to the source sentence. At test time, the model relies on the tag to guide the formality

of the translation hypothesis. A similar method has been used in [Vanmassenhove et al. \(2018\)](#) and in [Elaraby et al. \(2018\)](#) to address the problem of speaker gender morphological agreement. [Moryossef et al. \(2019\)](#) address the issue by modifying the source sentence during inference. They prepend the source with a minimal phrase implicitly containing all the relevant information; for example, for a female speaker and a plural audience, the augmented source yields “*She said to them: <src. sent.>*”. Their method improves on multiple phenomena simultaneously (speaker gender and number, interlocutor gender and number) and requires little annotated data, but its performance relies entirely on the MT system’s ability to utilise the added information. Furthermore, there are some side effects, e.g. the authors find the model’s predictions to be often unintentionally influenced by the token *said*.

A similar method of tag-managed tuning has been used to train multilingual NMT systems ([Johnson et al., 2017](#)) and approximately control sequence length in NMT ([Lakew et al., 2019](#)). Outside MT, this method has been the driving force behind large pretrained controllable language models ([Devlin et al., 2019](#); [Keskar et al., 2019](#); [Dathathri et al., 2019](#); [Krause et al., 2020](#); [Mai et al., 2020](#)).

## 2.3 Document-level Neural Machine Translation (DocNMT)

Traditionally, NMT is a sentence-level (**Sent2Sent**) task, where models process each sentence of a document independently. Another way to do it would be to process the entire document at once (**Doc2Doc**), but it is much harder to train a reliable NMT model on document-long sequences. A compromise between the two is a **Doc2Sent** approach which produces the translation sentence by sentence but considers the document-level information as *context* when doing so ([Sun et al., 2020](#)).

**Doc2Doc** [Tiedemann and Scherrer \(2017\)](#) conduct the first Doc2Doc pilot study: they translate documents two sentences at once, each time discarding the first translated sentence and keeping the latter. They find that there is some benefit from doing so, albeit such benefit is difficult to measure. A larger setting was explored in ([Junczys-Dowmunt, 2019](#)): a 12-layer Transformer-Big ([Vaswani et al., 2017](#)) was trained to translate documents of up to 1000 subword units, with performance optimised by noisy back-translation, fine tuning and second-

<sup>3</sup>More examples can be found in the Appendix

pass post editing described in (Junczys-Dowmunt and Grundkiewicz, 2018). Finally, Sun et al. (2020) propose a fully Doc2Doc approach applicable to documents of arbitrary length. They split each document into  $k \in 1, 2, 4, 8, \dots$  parts and treat them as input data to the model, in what they call a **multi-relational** training, as opposed to **single relational** where only the whole document would be fed as input. Despite good results, the last two methods require enormous computational resources, and this limits their commercial application.

**Doc2Sent** When translating a sentence  $s_i$  a Doc2Sent model is granted access to document-level information  $S \subseteq \{s_0 \dots s_{i-1}, s_{i+1} \dots s_n\}$  and/or  $T \subseteq \{t_0 \dots t_{i-1}\}$  where  $n$  is the length of the document. The context information is either concatenated with the source sentence yielding a *uni-encoder* model (Tiedemann and Scherrer, 2017; Ma et al., 2020), or is supplied in an extra encoder yielding a *dual-encoder*<sup>4</sup> model (Zhang et al., 2018; Voita et al., 2020). In most approaches, the performance is optimised when shorter context (1-3 sentences) is used, though Kim et al. (2019) find that applying a simple rule-based *context filter* can stabilise performance for longer contexts. Ma et al. (2020) offer an improvement to uni-decoder which limits the sequence length in the top blocks of the Transformer encoder in the uni-encoder architecture, and Kang et al. (2020) introduce a reinforcement-learning-based *context scorer* which dynamically selects the context best suited for translating the critical sentence.

Jauregi Unanue et al. (2020) challenge the idea that DocNMT can implicitly learn document-level features, and instead propose that the models be *rewarded* when it preserves them. They focus on lexical cohesion and coherence and use respective metrics (Wong and Kit, 2012; Gong et al., 2015) to measure rewards. This method may be successful provided that suitable specialised evaluation metrics are proposed in the future. Nevertheless, more interest has been expressed in literature in achieving high performance w.r.t. such features as a by-product of an efficient architecture, as is the case with SOTA Sent2Sent architectures.

**Other architectures** DocRepair (Voita et al., 2019) is a monolingual post-editing model trained to repair cohesion in a document translated with SentNMT. Kuang et al. (2018) use two cache struc-

tures to influence the model’s token predictions: a dynamic cache  $c_d$  of past token hypotheses with stopword removal and a topic cache  $c_t$  of most probable topic-related words. Finally, Lopes et al. (2020) compress the entire document into a vector and supply it as context during translation.

## 2.4 Evaluation of Machine Translation

Many machine translation evaluation (MTE) metrics have been proposed over the years, much owing to the yearly WMT Metrics task (Mathur et al., 2020). They typically measure similarity between reference  $r$ , hypothesis  $h$  and source  $s$ , expressed in e.g. n-gram overlap (e.g. Papineni et al., 2002), cosine distance of embeddings (e.g. Zhang et al., 2020), translation edit rate (Snover et al., 2006) or trained on human judgements (Shimanaka et al., 2018), with the SOTA represented by COMET which combines the ideas of Zhang et al. and Shimanaka et al.: several distances between  $h, r$  and  $s$  are computed based on contextual embeddings from BERT.

Practically all of these metrics are developed to optimise performance at sentence level, an issue which until recently was not brought up often enough within the community. In the latest edition of the Metrics task at WMT (Mathur et al., 2020), a track for document-level evaluation was introduced. However, the organisers approached document-level evaluation as the average of human judgements on sentences in documents. This is not a reliable assessment, since the quality of a text is more than the sum or average of the quality of its sentences. This approach risks “averaging out” the severity of potential inter-sentential errors. Currently, DocNMT models are typically evaluated in terms of BLEU, showing modest improvements over a baseline (e.g. Voita et al., 2018, report 0.7 BLEU improvement). Several authors have argued that BLEU is not well suited to evaluating performance with respect to preserving cross-sentential discourse phenomena (Voita et al., 2020; Lopes et al., 2020). When applied to methods which improve only a certain aspect of translation, BLEU can indicate very little about the accuracy of these improvements. Furthermore, Kim et al. (2019) and Li et al. (2020) argue that even the reported BLEU gains in DocNMT models may not come from document-level quality improvements. Li et al. (2020) show that feeding the incorrect context can improve the metric by a similar amount.

<sup>4</sup>Notation adopted from Ma et al. (2020).

To decide whether DocNMT yield any improvements, a more sophisticated evaluation method is needed. Following the observation that DocNMT improves on individual aspects of translation w.r.t. SentNMT, **test suites** grew in popularity among researchers (Bawden, 2019; Voita et al., 2020; Lopes et al., 2020). In particular, contrastive test suites (Müller et al., 2018) measure whether a model can repeatedly identify and correctly translate a certain phenomenon. They can be seen as robust collections of fine-grained multiple choice questions, yielding for each phenomenon an accuracy score indicative of performance. Producing these suites is time consuming and often requires expertise, but they are of extreme benefit to NMT. A sufficiently rich bed of test suites can evaluate the general robustness of a model, expressed as the average accuracy on these suites.

### 3 Addressing Research Questions

Within this PhD, we seek to answer three research questions (RQs):

- RQ1 Can machine translation of dialogue be personalised by supplying it with extra-textual information?
- RQ2 Is ellipsis problematic for MT, and can MT make use of marking of ellipsis and other cohesion devices to increase cohesion in translation of dialogue?
- RQ3 How can automatic evaluation methods of MT be developed which confidently and reliably reward successful translations of contextual phenomena and, likewise, punish incorrect translations of the same phenomena?

#### 3.1 Modelling Extra-Textual Information in Machine Translation

We hypothesise that supplying the MT model with extra-textual information might help it make better dialogue translation choices. Our hypothesis is motivated by two facts: (1) that human translators base their choices of individual utterances on the understanding of the discourse situation and ensure that each utterance preserves its original function and meaning, and (2) that many instances of utterances and phrases are impossible to interpret unambiguously in isolation from their context.

#### Tuning MT output with external information

Previous works on supplying context via constraints or tags have been narrow in scope, predominantly employing tag controlling (see subsection 2.2).

Following their success we plan to experiment with alternative neural model architectures which allow the incorporation of extra data into sequence-to-sequence transduction and assess whether they are fit for translation. If successful, we see many potential applications of such models in NMT, ranging from those explored in this thesis to limiting the length of the translation, fine-grained personalisation (e.g. on speaker characteristics) and more.

**Per scene domain adaptation** Neural machine translation models can be fine-tuned to a particular domain (e.g. medical transcripts) via domain adaptation (Cuong and Sima'an, 2017). Effective as it is, domain adaptation requires domain-specific data and that the model is trained on it (a time-consuming process). This technique is then inapplicable in scenarios where domains are fine-grained and the adaptation needs to be instantaneous. Per scene adaptation appears to be a promising solution to the problem of wrong lexical choices made by MT models when translating dialogue. The environment or scene in which dialogue occurs is often crucial to interpreting its meaning; a scene-unaware model may misinterpret the function of an utterance and produce an incorrect translation.

Within TV dialogue we define a scene as continuous action which sets boundaries for exchanges. Its characteristics can be expressed in natural language (e.g. extracts from plot synopsis), as tags (e.g. *school, student, sunny, exam*) or as individual categories (e.g. *battle*). Since scene context is document-level, this task can also be seen as a use case for combining PersNMT and DocNMT, and will be explored in this PhD.

#### 3.2 Improving Cohesion for Machine Translation of Dialogue

Work within MT so far has only limitedly explored whether ellipsis poses a significant problem for translation (see Voita et al., 2020). We hypothesise that this is indeed the case: for some language pairs, the quality of machine-translated texts depend on the system's understanding of the ellipsis, when it is present in the source text. Since in dialogue ellipsis typically spans more than one utterance, it is poorly understood by SentNMT and the resulting MT quality is low (Figure 2).

To test our hypothesis, we will analyse ellipsis occurrences in dialogue data. We will use automatic methods to identify 1,000 occurrences of ellipsis in source text and mark spans of their occurrence

EN	“I’m sorry, Dad, but you <u>wouldn’t understand.</u> ” “Oh, sure, I <u>would [understand],</u> princess.”
PL <sub>MT</sub>	“Przepraszam tato, ale <u>nie zrozumiałbyś.</u> ” “Och, <u>oczywiście,</u> księżniczko.”
PL <sub>ref</sub>	“Przykro mi, tato, ale <u>nie zrozumiałbyś.</u> ” “ <u>Pewnie, że zrozumiałbym,</u> księżniczko.”

Figure 2: A wrongly translated exchange with ellipsis. In the source, the word *would* is a negation to *wouldn’t* in the previous utterance. The MT system ignores *I would*: the backtranslation of PL<sub>MT</sub> reads “*Oh, sure, princess.*”

in the corresponding machine and reference translations. All cases will then be manually analysed from the following angles: (i) Is the ellipsis correctly translated? (ii) Is the resulting translation of ellipsis natural/unnatural? (iii) Does the reference translation make use of the elided content? (iv) If the model generates an acceptable translation, could the elided content nevertheless have been used to disambiguate it or make it more cohesive?

Next, we aim to build a DocNMT system which utilises marking of cohesion phenomena to make more cohesive translation choices<sup>5</sup> (Figure 3). We apply the insights from previous research, namely that the Transformer model may track cohesion phenomena when given enough context (Voita et al., 2018), that context preprocessing stabilises performance of contextual MT models (Kim et al., 2019), solutions to the problem of long inputs in DocNMT (e.g. Ma et al., 2020; Sun et al., 2020), and finally our own analysis of the problem.




Figure 3: A draft of our DocNMT pipeline architecture. We preprocess the document to mark cohesion features. Then we use the output as the data for our model.

### 3.3 Applying Evaluation Metrics to Cohesion and Speaker Phenomena

Addressing RQ3 will involve testing the hypothesis that current common and SOTA automatic evaluation metrics fail to successfully reward translations which preserve contextual phenomena and, similarly, fail to punish those which do not.

We will develop a document-level test set of dialogue utterances in five languages, annotated for contextual phenomena. For each phenomenon, we will modify the reference translations to prepare se-

<sup>5</sup>Including elliptical structures in this step will depend on the result of the first experiment.

veral variations: one where all marked phenomena are translated correctly, another one where only 90% is translated correctly, then 80% etc. up to 0%. We will prepare a set of common and SOTA MT evaluation metrics and use them to produce scores for all variants, for all phenomena. If there exists a metric which gives a consistently lower score the more a phenomenon is violated, for all phenomena, then our hypothesis is incorrect and we will use that metric for evaluation in experiments. Otherwise, we will develop our own metric.

The aforementioned test set will also be converted to a contrastive test suite (Müller et al., 2018) and submitted as an evaluation method to WMT News Translation task. The data to be used here is a combination of the Serial Speakers dataset (Bost et al., 2020) and exports from OpenSubtitles (Lison and Tiedemann, 2016), yielding 5.6k utterances total, split into scenes and parallel in five languages.

We hope that this work will substantiate the flaws of sentence-level evaluation and prompt the community to work on context-inclusive methods.

## 4 Conclusions

This work is the proposal of a PhD addressing PersNMT and DocNMT in the dialogue domain. We have presented evidence that sentence-level MT models make cohesion- and coherence-related errors and offered several approaches via which we aim to tackle this problem. We plan to conduct extensive experiments to analyse the problem of ellipsis translation and of the use of sentence-level evaluation metrics to evaluate contextual phenomena. The outcome of this work will also include publicly available test suites, a document-level translation model, a personalised translation model and a context-aware evaluation metric.

## Acknowledgements

This work was supported by the Centre for Doctoral Training in Speech and Language Technologies (SLT) and their Applications funded by UK Research and Innovation [grant number EP/S023062/1].

## References

- Rachel Bawden. 2019. *Going beyond the sentence : Contextual Machine Translation of Dialogue*. Ph.D. thesis, Université Paris-Saclay.
- Rachel Bawden, Rico Sennrich, Alexandra Birch, and Barry Haddow. 2018. [Evaluating discourse phenomena in neural machine translation](#). *NAACL HLT 2018 - 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference*, 1:1304–1313.
- Xavier Bost, Vincent Labatut, and Georges Linarès. 2020. [Serial speakers: A dataset of TV series](#). *arXiv*.
- Hoang Cuong and Khalil Sima'an. 2017. [A survey of domain adaptation for statistical machine translation](#). *Machine Translation*, 31(4):187–224.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2019. [Plug and Play Language Models: A Simple Approach to Controlled Text Generation](#). *arXiv*, pages 1–34.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). In *Proceedings of the 2019 Conference of the North {A}merican Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Mostafa Elaraby, Ahmed Y. Tawfik, Mahmoud Khalid, Hany Hassan, and Aly Osama. 2018. [Gender aware spoken language translation applied to English-Arabic](#). *2nd International Conference on Natural Language and Speech Processing, ICNLSP 2018*, pages 1–6.
- Zhengxian Gong, Min Zhang, and Guodong Zhou. 2015. [Document-Level Machine Translation Evaluation with Gist Consistency and Text Cohesion](#). In *Proceedings of the Second Workshop on Discourse in Machine Translation (DiscoMT)*, pages 33–40. Association for Computational Linguistics.
- M A K Halliday and R Hasan. 1976. *Cohesion in English*. Longman, London.
- M. A.K. Halliday and Christian M.I.M. Matthiessen. 2013. *Halliday's introduction to functional grammar: Fourth edition*. Routledge.
- Hany Hassan, Anthony Aue, Chang Chen, Vishal Chowdhary, Jonathan Clark, Christian Federmann, Xuedong Huang, Marcin Junczys-Dowmunt, William Lewis, Mu Li, Shujie Liu, Tie-Yan Liu, Renqian Luo, Arul Menezes, Tao Qin, Frank Seide, Xu Tan, Fei Tian, Lijun Wu, Shuangzhi Wu, Yingce Xia, Dongdong Zhang, Zhirui Zhang, and Ming Zhou. 2018. [Achieving Human Parity on Automatic Chinese to English News Translation](#). *arXiv*.
- Inigo Jauregi Unanue, Nazanin Esmaili, Gholamreza Haffari, and Massimo Piccardi. 2020. [Leveraging Discourse Rewards for Document-Level Neural Machine Translation](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4467–4482, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. [Google's Multilingual Neural Machine Translation System: Enabling Zero-Shot Translation](#). *Transactions of the Association for Computational Linguistics*, 5:339–351.
- Marcin Junczys-Dowmunt. 2019. [{M}icrosoft Translator at {WMT} 2019: Towards Large-Scale Document-Level Neural Machine Translation](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 225–233, Florence, Italy. Association for Computational Linguistics.
- Marcin Junczys-Dowmunt and Roman Grundkiewicz. 2018. [MS-UEdin Submission to the WMT2018 APE Shared Task: Dual-Source Transformer for Automatic Post-Editing](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 822–826, Belgium, Brussels. Association for Computational Linguistics.
- Prathyusha Jwalapuram, Shafiq Joty, Irina Temnikova, and Preslav Nakov. 2020. [Evaluating pronominal anaphora in machine translation: An evaluation measure and a test suite](#). *EMNLP-IJCNLP 2019 - 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing, Proceedings of the Conference*, pages 2964–2975.
- Xiaomian Kang, Yang Zhao, Jiajun Zhang, and Chengqing Zong. 2020. [Dynamic Context Selection for Document-level Neural Machine Translation via Reinforcement Learning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2242–2254, Online. Association for Computational Linguistics.
- Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, and Richard Socher. 2019. [CTRL: A conditional transformer language model for controllable generation](#). *arXiv*, pages 1–18.
- Yunsu Kim, Duc Thanh Tran, and Hermann Ney. 2019. [When and Why is Document-level Context Useful in Neural Machine Translation?](#) In *Proceedings of the Fourth Workshop on Discourse in Machine Translation (DiscoMT 2019)*, pages 24–34, Hong Kong, China. Association for Computational Linguistics.

- Svenja Kranich. 2014. Translations as a Locus of Language Contact. In Juliane House, editor, *Translation: A Multidisciplinary Approach*, pages 96–115. Palgrave Macmillan UK, London.
- Ben Krause, Akhilesh Deepak Gotmare, Bryan McCann, Nitish Shirish Keskar, Shafiq Joty, Richard Socher, and Nazneen Fatema Rajani. 2020. [GeDi: Generative Discriminator Guided Sequence Generation](#). *arXiv*, pages 1–31.
- Shaohui Kuang, Deyi Xiong, Weihua Luo, and Guodong Zhou. 2018. [Modeling Coherence for Neural Machine Translation with Dynamic and Topic Caches](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 596–606, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Surafel Melaku Lakew, Mattia Di Gangi, and Marcello Federico. 2019. [Controlling the Output Length of Neural Machine Translation](#). *arXiv*.
- Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2020. [Phrase-based & neural unsupervised machine translation](#). *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018*, pages 5039–5049.
- Samuel Läubli, Rico Sennrich, and Martin Volk. 2018. [Has machine translation achieved human parity? A Case for Document-level Evaluation](#). *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018*, pages 4791–4796.
- Bei Li, Hui Liu, Ziyang Wang, Yufan Jiang, Tong Xiao, Jingbo Zhu, Tongran Liu, and Changliang Li. 2020. [Does Multi-Encoder Help? A Case Study on Context-Aware Neural Machine Translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3512–3518, Online. Association for Computational Linguistics.
- Pierre Lison and Jörg Tiedemann. 2016. [OpenSubtitles2016: Extracting Large Parallel Corpora from Movie and TV Subtitles](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 923–929, Portorož, Slovenia. European Language Resources Association (ELRA).
- António Lopes, M Amin Farajian, Rachel Bawden, Michael Zhang, and André F T Martins. 2020. [Document-level Neural MT: A Systematic Comparison](#). In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 225–234, Lisboa, Portugal. European Association for Machine Translation.
- Shuming Ma, Dongdong Zhang, and Ming Zhou. 2020. [A Simple and Effective Unified Encoder for Document-Level Machine Translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3505–3511, Online. Association for Computational Linguistics.
- Florian Mai, Nikolaos Pappas, Ivan Montero, Noah A Smith, and James Henderson. 2020. [Plug and Play Autoencoders for Conditional Text Generation](#). *arXiv*.
- Nitika Mathur, Johnny Wei, Qingsong Ma, and Ondřej Bojar. 2020. [Results of the WMT20 Metrics Shared Task](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 686–723, Online. Association for Computational Linguistics.
- Amit Moryossef, Roei Aharoni, and Yoav Goldberg. 2019. [Filling Gender & Number Gaps in Neural Machine Translation with Black-box Context Injection](#). In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 49–54, Florence, Italy. Association for Computational Linguistics.
- Mathias Müller, Annette Rios, 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 Third Conference on Machine Translation: Research Papers*, pages 61–72, Brussels, Belgium. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a Method for Automatic Evaluation of Machine Translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Martin J. Pickering and Simon Garrod. 2004. [Toward a mechanistic psychology of dialogue](#). *Behavioral and Brain Sciences*, 27(02).
- Argentina Anna Rescigno, Johanna Monti, Andy Way, and Eva Vanmassenhove. 2020. [A Case Study of Natural Gender Phenomena in Translation: A Comparison of Google Translate, Bing Microsoft Translator and DeepL for English to Italian, French and Spanish](#). In *Workshop on the Impact of Machine Translation (iMpaCT 2020)*, pages 62–90, Virtual. Association for Machine Translation in the Americas.
- Ola Rønning, Daniel Hardt, and Anders Søgaard. 2018. [Sluice resolution without hand-crafted features over brittle syntax trees](#). *NAACL HLT 2018 - 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference*, 2:236–241.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Controlling politeness in neural machine translation via side constraints](#). *2016 Conference of*



- the North American Chapter of the Association for Computational Linguistics: *Human Language Technologies, NAACL HLT 2016 - Proceedings of the Conference*, pages 35–40.
- Hiroki Shimanaka, Tomoyuki Kajiwara, and Mamoru Komachi. 2018. **RUSE: Regressor Using Sentence Embeddings for Automatic Machine Translation Evaluation**. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 751–758, Belgium, Brussels. Association for Computational Linguistics.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation. *AMTA 2006 - Proceedings of the 7th Conference of the Association for Machine Translation of the Americas: Visions for the Future of Machine Translation*, pages 223–231.
- Gabriel Stanovsky, Noah A. Smith, and Luke Zettlemoyer. 2020. **Evaluating gender bias in machine translation**. *ACL 2019 - 57th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference*, pages 1679–1684.
- Zewei Sun, Mingxuan Wang, Hao Zhou, Chengqi Zhao, Shujian Huang, Jiajun Chen, and Lei Li. 2020. **Capturing Longer Context for Document-level Neural Machine Translation: A Multi-resolutional Approach**. *arXiv*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. *Advances in Neural Information Processing Systems*, 4(January):3104–3112.
- 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.
- Eva Vanmassenhove, Christian Hardmeier, and Andy Way. 2018. **Getting gender right in neural machine translation**. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018*, pages 3003–3008.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. **Attention is all you need**. *Advances in Neural Information Processing Systems*, pages 5999–6009.
- Elena Voita, Rico Sennrich, and Ivan Titov. 2019. **Context-Aware Monolingual Repair for Neural Machine Translation**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 877–886, Hong Kong, China. Association for Computational Linguistics.
- Elena Voita, Rico Sennrich, and Ivan Titov. 2020. **When a good translation is wrong in context: Context-aware machine translation improves on deixis, ellipsis, and lexical cohesion**. *ACL 2019 - 57th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference*, pages 1198–1212.
- 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.
- Billy T.M. Wong and Chunyu Kit. 2012. Extending machine translation evaluation metrics with lexical cohesion to document level. In *EMNLP-CoNLL 2012 - 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, Proceedings of the Conference*, July, pages 1060–1068.
- Jiacheng Zhang, Huanbo Luan, Maosong Sun, Fei Fei Zhai, Jingfang Xu, Min Zhang, and Yang Liu. 2018. **Improving the transformer translation model with document-level context**. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018*, pages 533–542.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2020. **BERTScore: Evaluating Text Generation with BERT**. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

## A Other examples

In this section we present an extended set of examples supporting our hypotheses stated in the main proposal. All examples in [Figure 4](#), [Figure 5](#) and [Figure 6](#) show examples of mistranslated sentences where the error was related to a specific phenomenon: ellipsis in [Figure 4](#), lexical cohesion in [Figure 5](#) and reference in [Figure 6](#). [Figure 7](#), instead of highlighting translation errors, shows how a sentence in English can have several different translation candidates depending on the extra-textual context embedded in the situation (the corresponding translations are reference translations rather than MT-generated ones).

<b>Context</b>	What would they use it for?
<b>Antecedent</b>	[They would use it for]
<b>EN</b>	Grabbing the balls of a spy.
<b>PL<sub>MT</sub></b>	<b>Łapie</b> szpiega za jaja. ('He/she/they grab(s) the balls of a spy')
<b>PL<sub>ref</sub></b>	<b>Żeby łąpać</b> szpiega za jaja. ('For grabbing the balls of a spy')
<b>Context</b>	A big, dumb, balding, North American ape with no chin.
<b>Antecedent</b>	[with]
<b>EN</b>	And a short temper.
<b>PL<sub>MT</sub></b>	I <b>krótki temperament</b> <sub>nominative</sub> . ('And a short temper.')
<b>PL<sub>ref</sub></b>	I z <b>krótkim temperamentem</b> <sub>instrumental</sub> . ('And with a short temper.')
<b>Context</b>	(...) with a record of zero wins and 48 <u>defeats</u> ...
<b>Antecedent</b>	[a record of zero wins and 48]
<b>EN</b>	Oh, correction. Humiliating <u>defeats</u> , all of them by knockout–
<b>PL<sub>MT</sub></b>	Oh, korekta. <b>Upokarzające porażki, wszystkie</b> <sub>nominative</sub> przez nokautowanie... (‘Oh, correction. Humiliating defeats, all of them by knockout...’)
<b>PL<sub>ref</sub></b>	Oh, korekta. <b>Upokarzających porażek, wszystkich</b> <sub>genitive</sub> przez nokautowanie... (‘Oh, correction. Humiliating defeats, all of them by knockout...’)
<b>Context</b>	“I’ve only got two cupcakes for the three of you.”
<b>Antecedent</b>	[two cupcakes]
<b>EN</b>	“Just take mine [=my cupcake].”
<b>DE<sub>MT</sub></b>	“Nimm einfach <b>meine</b> [=mine <sub>fem</sub> ].”
<b>DE<sub>ref</sub></b>	“Nimm einfach <b>meinen</b> [=mine <sub>masc</sub> ].”

Figure 4: Examples of translations where resolving ellipsis is crucial to generating a correct translation hypothesis. **Context** is the utterance containing the antecedent, and **Antecedent** is the content which is elided in the current utterance. In the first two examples from the top, the Polish translation requires including part of the antecedent in order to maintain cohesion. In the third example from the top, the antecedent decides the inflection of all the words relating to the word *defeats* which is repeated in the current utterance. Finally, the bottom example contains nominal ellipsis, and the model uses an incorrect inflection of *mein* since it fails to make the connection with the antecedent.

<b>EN</b>	“Sorry, Dad. I <u>know</u> you mean well.” “Thanks for <u>knowing</u> I mean well.”
<b>PL<sub>MT</sub></b>	“Przepraszam tato. <b>Wiem</b> , że chcesz dobrze.” “Dzięki, że <b>wiedziłeś</b> , że chcę dobrze.”
<b>PL<sub>ref</sub></b>	“Przepraszam tato. <b>Wiem</b> , że chcesz dobrze.” “Dzięki, że <b>wiesz</b> , że chcę dobrze.”
<b>EN</b>	“You’re a <u>dimwit</u> .” “Maybe so, but from now on... this <u>dimwit</u> is on easy street.”
<b>PL<sub>MT</sub></b>	“Jesteś <b>głupcem</b> .” (‘You’re a fool.’) “Może i tak, ale od teraz ... ten <b>głupek</b> ( <i>dimwit</i> ) jest na łatwej ulicy.”
<b>PL<sub>ref</sub></b>	“Jesteś <b>głupkiem</b> .” (‘You’re a dimwit.’) “Może i tak, ale od teraz ... ten <b>głupek</b> ( <i>dimwit</i> ) jest na łatwej ulicy.”

Figure 5: Examples of mistranslated lexical cohesion. In the top example, although the MT model managed to translate most of the repeated phrase in the same way, it failed to maintain the verb *know* in the present tense. In the bottom example a different translation of *dimwit* is used in the two utterances. Note that it is okay for a model to give a different hypothesis to a word than the human translator would, as long as it agrees with the source *and* is cohesive with the rest of the text (i.e. all occurrences of the word are translated in the same way).

<b>EN</b>	The <u>grabber</u> . What would they use <u>it</u> for?
<b>DE<sub>MT</sub></b>	Der Grabber <sub>masc.</sub> . Wofür würden sie <b>es</b> <sub>neut</sub> verwenden?
<b>DE<sub>ref</sub></b>	Der Grabber <sub>masc.</sub> . Wofür würden sie <b>ihn</b> <sub>masc</sub> verwenden?
<b>EN</b>	Leave <u>ideology</u> to the armchair generals. <u>It</u> does me no good.
<b>PL<sub>MT</sub></b>	Ideologie <sub>fem</sub> zostawcie generałom foteli. Nic mi <b>to</b> <sub>neut</sub> nie da.
<b>PL<sub>ref</sub></b>	Ideologie <sub>fem</sub> zostawcie generałom foteli. Nic mi <b>ona</b> <sub>fem</sub> nie da.

Figure 6: Examples of mistranslated multi-sentence dialogue where reference is the violated phenomenon. In both examples, the gender of the referent is different in source and target languages, therefore the pronoun which refers to it is mistranslated.

<b>EN</b>	I never expected to be involved in every policy or decision, but I have been completely cut out of everything.
<b>PL (fem)</b>	Nigdy nie oczekiwałam <u>wglądu</u> w każdą decyzję, ale zostałam <u>odcięta</u> od wszystkiego.
<b>PL (masc)</b>	Nigdy nie oczekiwałem <u>wglądu</u> w każdą decyzję, ale zostałem <u>odcięty</u> od wszystkiego.
<b>EN</b>	And who have you called, by the way ?
<b>PL (to masc)</b>	Do kogo już dzwoniłeś?
<b>PL (to fem)</b>	Do kogo już dzwoniłaś?
<b>PL (to Plural)</b>	Do kogo już dzwoniлиście?
<b>PL (to Plural<sub>fem</sub>)</b>	Do kogo już dzwoniłyście?
<b>EN</b>	He was shot previous to your arrival?
<b>PL (formal)</b>	Został postrzelony przed <u>pana</u> przyjazdem?
<b>PL (informal)</b>	Został postrzelony przed <u>Twoim</u> przyjazdem?

Figure 7: Examples of situation phenomena that can occur in text: speaker gender agreement (top), addressee gender agreement (middle), formality (bottom).