

The Unreasonable Volatility of Neural Machine Translation Models

Marzieh Fadaee[†]
Informatics Institute
University of Amsterdam
marzieh.f@gmail.com

Christof Monz
Informatics Institute
University of Amsterdam
c.monz@uva.nl

Abstract

Recent works have shown that while Neural Machine Translation (NMT) models achieve impressive performance, questions about understanding the behaviour of these models remain unanswered. We investigate the unexpected volatility of NMT models where the input is semantically and syntactically correct. We discover that with trivial modifications of source sentences, we can identify cases where *unexpected changes* happen in the translation and in the worst case lead to mistranslations. This volatile behaviour of translating extremely similar sentences in surprisingly different ways highlights the underlying generalization problem of current NMT models. We find that both RNN and Transformer models display volatile behaviour in 26% and 19% of sentence variations, respectively.

1 Introduction

The performance of Neural Machine Translation (NMT) models has dramatically improved in recent years, and with sufficient and clean data these models outperform more traditional models. Challenges when sufficient data is not available include translations of rare words (Pham et al., 2018) and idiomatic phrases (Fadaee et al., 2018) as well as domain mismatches between training and testing (Koehn and Knowles, 2017; Khayrallah and Koehn, 2018).

Recently, several approaches investigated NMT models when encountering noisy input and how *worst-case examples* of noisy input can ‘break’ state-of-the-art NMT models (Goodfellow et al., 2015; Michel and Neubig, 2018). Belinkov and Bisk (2018) show that character-level noise in the input leads to poor translation performance. Lee et al. (2018) randomly insert words in different positions in the source sentence and observe that in

[†]Marzieh is now affiliated with Zeta Alpha Vector.

Source: *Ich bin [1] erleichtert und [2] bescheiden.*

[1]	[2]	NMT output
ϕ	ϕ	I am <u>easier</u> and modest.
ϕ	<i>sehr</i> [‡]	I am relieved and very modest.
<i>sehr</i>	ϕ	I am very much <u>easier</u> and modest.
<i>sehr</i>	<i>sehr</i>	I am very <u>easy</u> and very modest.
Reference		
ϕ	ϕ	<i>I am relieved and humble.</i>
<i>sehr</i>	<i>sehr</i>	<i>I am very relieved and very humble.</i>

Table 1: Insertion of the German word *sehr* (English: *very*) in different positions in the source sentence results in substantially different translations. [‡] indicates the original sentence from WMT 2017.

some cases the translations are completely unrelated to the input. While it is to some extent expected that the performance of NMT models that are trained on predominantly clean but tested on noisy data deteriorates, other changes are more unexpected.

In this paper, we explore unexpected and erroneous changes in the output of NMT models. Consider the simple example in Table 1 where the Transformer model (Vaswani et al., 2017) is used to translate very similar sentences. Surprisingly, we observe that by simply altering one word in the source sentence—inserting the German word *sehr* (English: *very*)—an unrelated change occurs in the translation. In principle, an NMT model that generates the translation of the word *erleichtert* (English: *relieved*) in one context, should also be able to generalize and translate it correctly in a very similar context. Note that there are no infrequent words in the source sentence and after each modification, the input is still syntactically correct and semantically plausible. We call a model *volatile* if it displays inconsistent behaviour across similar input sentences during inference.

We investigate to what extent well-established

Modification	Sentence variations
DEL	Some 500 years after the Reformation, Rome [now \ ϕ] has a Martin Luther Square.
SUBNUM	I'm very pleased for it to have happened at Newmarket because this is where I landed [30 \ 31] years ago.
INS	I loved Amy and she is [ϕ \b also] the only person who ever loved me.
SUBGEN	[He \ She] received considerable appreciation and praise for this.

Table 2: Examples of different variations from WMT. [w_i \ w_j] indicates that w_i in the original sentence is replaced by w_j . ϕ is an empty string.

NMT models are volatile during inference. Specifically, we locally modify sentence pairs in the test set and identify examples where a trivial modification in the source sentence causes an ‘unexpected change’ in the translation. These modifications are generated conservatively to avoid insertion of any noise or rare words in the data (Section 2.2). Our goal is not to *fool* NMT models, but instead identify common cases where the models exhibit unexpected behaviour and in the worst cases result in incorrect translations.

We observe that our modifications expose volatilities of both RNN and Transformer translation models in 26% and 19% of sentence variations, respectively. Our findings show how vulnerable current NMT models are to trivial linguistic variations, putting into question the generalizability of these models.

2 Sentence Variations

2.1 Is this another noisy text translation problem?

Noisy input text can cause mistranslations in most MT systems, and there has been growing research interest in studying the behaviour of MT systems when encountering noisy input (Li et al., 2019).

Belinkov and Bisk (2018) propose to swap or randomize letters in a word in the input sentence. For instance, they change the word *noise* in the source sentence into *iones*. Lee et al. (2018) examine how the insertion of a random word in a random position in the source sentence leads to mistranslations. Michel and Neubig (2018) proposes a benchmark dataset for translation of noisy input sentences, consisting of noisy, user-generated comments on Reddit. The types of noisy input text they observe include spelling or typographical errors, word omission/insertion/repetition, and grammatical errors.

In these previous works, the focus of the research is on studying how the MT systems are not robust when handling noisy input text. In these

approaches, the input sentences are semantically or syntactically incorrect which leads to mistranslations.

However, in this paper, our focus is on input text that does *not* contain any types of noise. We modify input sentences in a way that the outcomes are still syntactically and semantically correct. We investigate how the MT systems exhibit volatile behaviour in translating sentences that are extremely similar and only differ in one word without any noise injection.

2.2 Variation generation

While there are various ways to automatically modify sentences, we are interested in simple semantic and syntactic modifications. These trivial linguistic variations should have almost no effect on the translation of the rest of the sentence.

We define a set of rules to slightly modify the source and target sentences in the test data and keep the sentences syntactically correct and semantically plausible.

DEL A conservative approach of modifying a sentence automatically without breaking the grammaticality of a sentence is to remove adverbs. We identify a list of the 50 most frequent adverbs in English and their translations in German*. For every sentence in the WMT test sets, if we find a sentence pair containing both a word and its translation from this list, we remove both words and create a new sentence pair.

SUBNUM Another simple yet effective approach to safely modify sentences is to substitute numbers with other numbers. In this approach, we select every sentence pair from the test sets that contains a number and substitute the number i in both source and target sentences with $i + k$ where $1 \leq k \leq 5$. We choose a small range for change so that the sentences are still semantically correct

*dict.cc




Figure 1: *Levenshtein distance* and *span of change* between translations of sentence variations for RNN and Transformer. The majority of sentence variations falls into the category of *minor* changes between translations (blue area). However, a surprising number of cases have significant changes (red area). RNN exhibits a slightly more unstable pattern i.e., sentence variations with large edit differences and large spans of change.

for the most part and result in few implausible sentences.

INS Randomly inserting words in a sentence has a high chance of producing a syntactically incorrect sentence. To ensure that sentences remain grammatical and semantically plausible after modification, we define a bidirectional n -gram probability for inserting new words as follows:

$$P(w_3|w_1w_2w_4w_5) = \frac{C(w_1w_2w_3w_4w_5)}{C(w_1w_2 \bullet w_4w_5)}$$

w_3 is inserted in the middle of the phrase $w_1w_2w_4w_5$, if the conditional probability is greater than a predefined threshold. The probabilities are computed on the WMT data. This simple approach, instead of using a more complex language model, serves our purposes since we are interested in inserting very common words that are already captured by the n -grams in the training data.

SUBGEN Finally, a local modification is changing the gender of the person in the sentences. The goal of this modification is to investigate the existence and severity of gender bias in our models. This is inspired by recent approaches that have shown that NMT models learn social stereotypes such as gender bias from training data (Escudé Font and Costa-jussà, 2019; Stanovsky et al., 2019).

Note that in a minority of cases these procedures can lead to semantically incorrect sentences, for instance, by substituting numbers we can potentially generate sentences such as “*She was*

born on October 34th“. While this can cause problems for a reasoning task, it barely affects the translation task, as long as the modifications are consistent on the source and target side.

Table 2 shows examples of generated variations. We emphasize that only modifications with local consequences have been selected and we intentionally ignore cases such as *negation* which can result in wider structural changes in the translation of the sentence.

	DE-EN			EN-DE		
	2016	2017	2018	2016	2017	2018
RNN	32.5	28.2	35.2	28.1	22.4	34.6
Transformer	36.2	32.1	40.1	33.4	27.9	39.8

Table 3: BLEU scores for different models on the WMT data for translation DE \leftrightarrow EN.

We generate $10k$ sentence variations by applying these modifications to all sentence pairs in WMT test sets 2013–2018 (Bojar et al., 2018). We use RNN and Transformer models to translate sentences and their variations.

2.3 Experimental setup

In the translation experiments, we use the standard EN \leftrightarrow DE WMT-2017 training data (Bojar et al., 2018). We perform NMT experiments with two different architectures: RNN (Luong et al., 2015) and Transformer (Vaswani et al., 2017). We preprocess the training data with Byte-Pair Encoding (BPE) using 32K merge operations (Sennrich et al., 2016). During inference, we use beam search with a beam size of 5. Table 3 shows the case-sensitive BLEU scores as calculated by

Table 4: A random sample of sentences from the WMT test sets and our proposed variations shown with ‘unexpected change’ annotations ($\Delta Translation$). The cases where the unexpected change leads to a change in translation quality are marked in column $\Delta Quality$. $[w_i \setminus w_j]$ indicates that w_i in the original sentence is replaced by w_j . S is the original and modified source sentence, R is the original and modified reference translation, T is the translation of the original sentence, and T_m is the translation of the modified sentence. Differences in translations related to annotations in the original and the modified translations are in **red** and **orange**, respectively. Note that we are interested in *unexpected* changes and do not highlight the changes that are a direct consequence of the modifications.

S	Coes letztes Buch “Chop Suey” handelte von der chinesischen Küche in den USA, während Ziegelman in ihrem Buch “[97\101] Orchard” über das Leben in einem Wohnhaus an der Lower East Side aus der Lebensmittelperspektive erzählt.
R	Mr. Coe’s last book, “Chop Suey,” was about Chinese cuisine in America, while Ms. Ziegelman told the story of life in a Lower East Side tenement through food in her book “[97\101] Orchard.”
T	Coes’s last book, “Chop Suey,” was about Chinese cuisine in the US , while Ziegelman, in her book “97 Orchard” talks about living in a lower East Side.
T_m	Coes last book “Chop Suey” was about Chinese cuisine in the United States , while Ziegelman writes in her book “101 Orchard” about living in a lower East Side.
$\Delta Translation$: [reordered] [paraphrased]	
$\Delta Quality$: No	
S	Man hält [bereits\phi] Ausschau nach Parkbank, Hund und Fußball spielenden Jungs und Mädels.
R	You are [already\phi] on the lookout for a park bench, a dog, and boys and girls playing football.
T	We are already looking for Parkbank, dog and football playing boys and girls.
T_m	Look for Parkbank, dog and football playing boys and girls.
$\Delta Translation$: [word form] [add/remove]	
$\Delta Quality$: Yes	
S	Bei einem Unfall eines Reisebusses mit [43\45] Senioren als Fahrgästen sind am Donnerstag in Krummhörn (Landkreis Aurich) acht Menschen verletzt worden.
R	On Thursday, an accident involving a coach carrying [43\45] elderly people in Krummhörn (district of Aurich) led to eight people being injured.
T	In the event of an accident involving a coach with 43 senior citizens as passengers , eight people were injured on Thursday in Krummaudin (County Aurich) .
T_m	In the event of an accident involving a 45-year-old coach as a passenger , eight people were injured on Thursday in the district of Aurich .
$\Delta Translation$: [word form] [add/remove] [other]	
$\Delta Quality$: Yes	
S	Es ist ein anstrengendes Pensum, aber die Dorfmusiker helfen [normalerweise\phi], das Team motiviert zu halten.
R	It’s a backbreaking pace, but village musicians [usually\phi] help keep the team motivated.
T	It’s a demanding child , but the village musicians usually help keep the team motivated.
T_m	It is a hard-to-use , but the village musician helps to keep the team motivated.
$\Delta Translation$: [word form] [other]	
$\Delta Quality$: Yes	

multi-bleu.perl.

RNN As the first NMT system, we use a 2-layer bidirectional attention-based LSTM model implemented in OpenNMT (Klein et al., 2017) trained with an embedding size of 512, hidden dimension size of 1024, and batch size of 64 sentences. We use Adam (Kingma and Ba, 2015) for optimization.

Transformer We also experiment with the Transformer model (Vaswani et al., 2017) implemented in OpenNMT. We train a model with 6 layers, the hidden size is set to 512 and the fil-

ter size is set to 2048. The multi-head attention has 8 attention heads. We use Adam (Kingma and Ba, 2015) for optimization. All parameters are set based on the suggestions in Klein et al. (2017) to replicate the results of the original paper.

3 Evaluation of unexpected and erroneous changes

The modifications described above generate sentences that are extremely similar and hence are expected to have a very similar difficulty of translation. We evaluate the NMT models on how robust and consistent they are in translating these sen-




Figure 2: Categories of *unexpected changes* in the translation of sentence variations as provided by annotators. The percentage of sentence variations with *minor* and *major* edit differences, as defined in 3.1, are shown separately. The hatched pattern indicates the ratio of sentence variations for which the translation quality changes. Note that *expected changes* are not plotted here.

tence variations rather than their absolute quality.

3.1 Deviations from Original Translations

The variations are aimed to have minimal effect on changing the meaning of the sentences. Hence, major changes in the translations of these variations can be an indication of volatility in the model. To assess whether the proposed sentence variations result in major changes in the translations, we measure changes in the translations of sentence variations with Levenshtein distance (Levenshtein, 1966). Specifically, Levenshtein distance measures the edit distance between the two translations. We also use the first and last positions of change in the translations, which represents the span of changes.

Ideally, with our simple modifications, we expect a value of zero for the span of change and a value of at most 2 for the Levenshtein distance for a translation pair. This indicates that there is only one token difference between the translation of the original sentence and the modified sentence. We define two types of changes based on these measures: *minor* and *major*. We choose the threshold to distinguish between minor and major changes more conservatively to allow for more variations in the translations. The change in translations is empirically considered *major* if both metrics are greater than 10, and *minor* if both are less than 10. Note that edit distances and spans are based on BPE subword units.

With two very similar source sentences, we expect the Levenshtein distance and span of change between translations of these sentences to be small. Figure 1 shows the results for the RNN and Transformer model. While the majority of sentence variations have minor changes, a substan-

tial number of sentences, 18% of RNN and 13% of Transformer translations, result in translations with major differences. This is surprising and an indication of volatility since these trivial modifications, in principle, should only result in minor and local changes in the translations.

3.2 Oscillations of Variation in Translations

In this section, we look into various sentence-level metrics to further analyze the observed behaviour. In particular, we focus on the SUBNUM modification because with this modification we can generate numerous variations of the same sentence. Having a high number of variations for each sentence gives us the opportunity of observing oscillations of various string matching metrics.

We use sentence-level BLEU, METEOR (Denkowski and Lavie, 2011), TER (Snover et al., 2006), and LengthRatio to quantify changes in the translations. LengthRatio represents the translation length over reference length as a percentage. For a given source sentence, we define the oscillation range as changes in the sentence-level metric for the translations of variations of a given sentence.

While sentence-level metrics are not reliable indicators of translation quality, they do capture fluctuations in translations. With the variations we introduce, in theory there should be no fluctuations in the translations. Table 5 and Figure 3 provide the results. We observe that even though these sentence variations differ by only one number, there are many cases where an insignificant change in the sentence results in unexpectedly large oscillations. Both RNN and Transformer exhibit this behaviour to a certain extent.

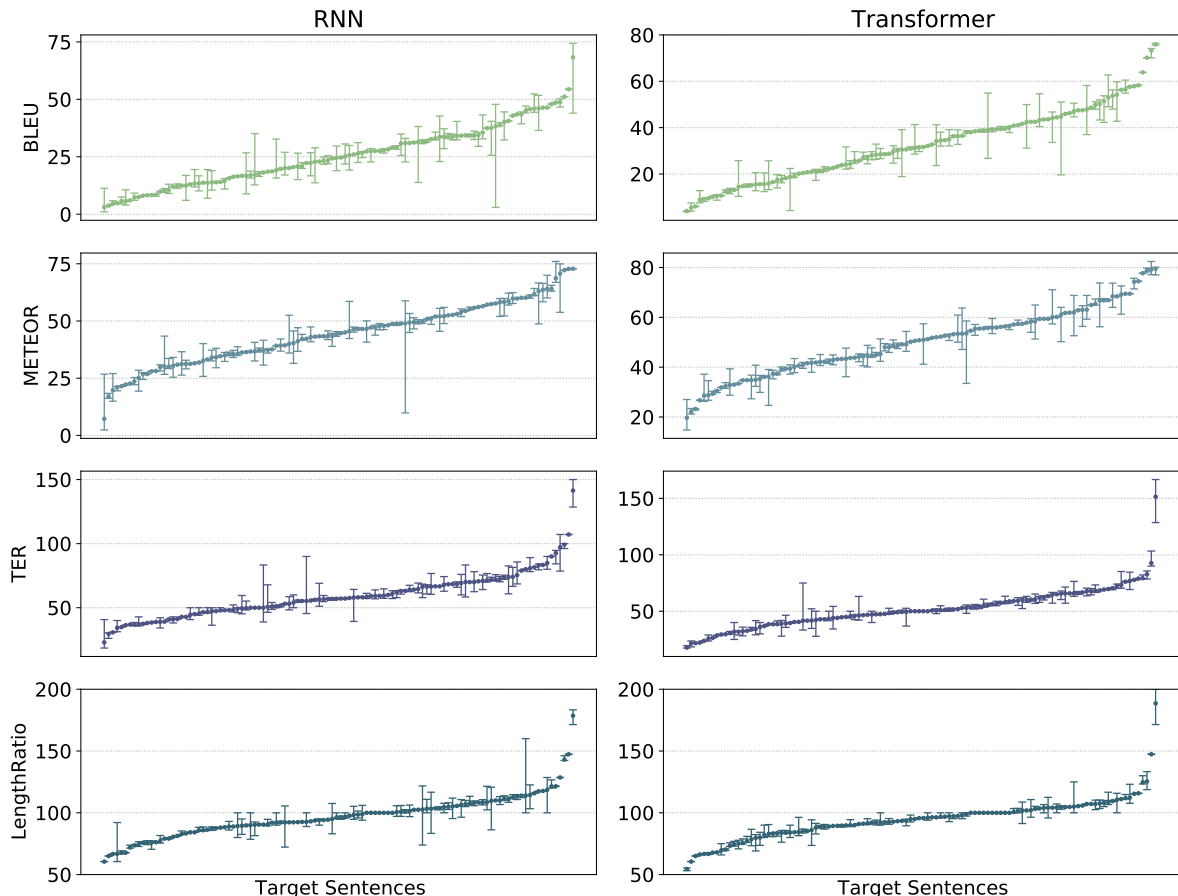


Figure 3: Oscillations of various sentence-level attributes for randomly sampled sentences from our test data and their SUBNUM variations. The data points are the mean values for all variations of each sentence, and the error bars indicate the range of oscillation of the metrics. The x-axis represents test sentence instances, sorted based on the corresponding metric. Ideally each data point should have zero oscillation.

Table 5: Mean oscillations for SUBNUM variations. In theory the variations should result in zero oscillations for every metric.

	BLEU	METEOR	TER	LengthRatio
RNN	4.0	3.8	5.2	5.3
Transformer	3.8	3.3	4.2	3.4

3.3 The Effect of Volatility on Translation Quality

While edit distances and spans of change provide some indication of volatility, they do not capture all aspects of this unexpected behaviour. It is also not entirely clear what effect these unexpected changes have on translation quality. To further investigate this, we also perform manual evaluations.

In the first evaluation, we provide annotators with a pair of sentence variations and their corresponding translations and ask them to identify

the differences between the two sentence pairs. In the second evaluation, we additionally provide the source sentences and reference translations, and ask the annotators to rank the sentence variations based on the translation quality similar to [Bojar et al. \(2016\)](#). In total the annotators evaluated 400 randomly selected sentence quadruplets.

The annotators identified 71% and 68% of changes in the variation translation as *expected* for the RNN and Transformer model, respectively. The main types of *unexpected* changes identified by the annotators are a change of word form, e.g., verb tense,, reordering of phrases, paraphrasing parts of the sentence, and an ‘other’ category, e.g., preposition. A sentence pair can have multiple labels based on the types of changes. Table 4 provides examples from the test data.

Statistics for each category of unexpected change is shown in Figure 2. Our first observation is that, as to be expected, there are very few

‘unexpected changes’ when two variations lead to translations with *minor* differences. Interestingly, the vast majority of changes are due to paraphrasing and dropping of words. Comparing the performance of the RNN and Transformer model, we see that both RNN and Transformer display inconsistent translation behaviour. While Transformer has slightly fewer sentences with major changes, it has a higher number of sentence variations in the major category that result in a change in translation quality. From the annotators’ assessments, we find that in 26% and 19% of sentence variations, the modification results in a change in *translation quality* for the RNN and Transformer model, respectively.

3.4 Generalization and Compositionality

Because of their ability to generalize beyond their training data, deep learning models achieve exceptional performances in numerous tasks. The generalization ability allows MT systems to generate long sentences not seen before. Recently there has been some interest in understanding whether this performance depends on recognizing shallow patterns, or whether the networks are indeed capturing and generalizing linguistic rules.

In simple terms, compositionality is the ability to construct larger linguistic expressions by combining simpler parts. For instance, if a model understands the correct compositional rules to understand ‘*John loves Mary*’, it must also understand ‘*Mary loves John*’ (Fodor and LePore, 2002). Investigating the compositional behaviour of neural networks in real-world natural language problems is a challenging task. Recently, several works have studied deep learning models’ understanding of compositionality in natural language by using synthetic and simplified languages (Andreas, 2019; Chevalier-Boisvert et al., 2019). Baroni (2019) shows that to a certain extent neural networks can be productive without being compositional.

Although we do not specifically look into the compositional potential of MT systems, we are inspired by compositionality in defining our modifications. We argue that the observed volatile behaviour of the MT systems in this paper is a side effect of current models not being compositional. If an MT system has a good ‘understanding’ of the underlying structures of the sentences ‘*Mary is 10 years old*’ and ‘*Mary is 11 years old*’, it must also translate them very similarly regardless of the ac-

curacy of the translation. While current evaluation metrics capture the accuracy of the NMT models, these volatilities go unnoticed.

Current neural models are successful in generalizing without learning any explicit compositional rules, however, our findings signal that they still lack robustness. We highlight this lack of robustness and suspect that it is associated with these models’ lack of understanding of the compositional nature of language.

4 Conclusion

In this paper, we showed the unexpected volatility of NMT models by using a simple approach to modifying standard test sentences without introducing noise, i.e., by generating semantically and syntactically correct variations. We show that even with trivial linguistic modifications of source sentences we can effectively identify a surprising number of cases where the translations of extremely similar sentences are surprisingly different, see Figure 1. Our manual analyses show that both RNN and Transformer models exhibit volatile behaviour with changes in translation quality for 26% and 19% of sentence variations, respectively. This highlights the problem of generalizability of current NMT models and we hope that our insights will be useful for developing more robust NMT models.

Acknowledgments

We thank Arianna Bisazza for helpful discussions. This research was funded in part by the Netherlands Organization for Scientific Research (NWO) under project numbers 639.022.213 and 612.001.218. We also thank NVIDIA for their hardware support and the anonymous reviewers for their helpful comments.

References

- Jacob Andreas. 2019. [Measuring compositionality in representation learning](#). *CoRR*, abs/1902.07181.
- Marco Baroni. 2019. [Linguistic generalization and compositionality in modern artificial neural networks](#). *CoRR*, abs/1904.00157.
- Yonatan Belinkov and Yonatan Bisk. 2018. [Synthetic and natural noise both break neural machine translation](#). In *Proceedings of the International Conference on Learning Representations (ICLR)*.

- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Aurelie Neveol, Mariana Neves, Martin Popel, Matt Post, Raphael Rubino, Carolina Scarton, Lucia Specia, Marco Turchi, Karin Verspoor, and Marcos Zampieri. 2016. [Findings of the 2016 conference on machine translation](#). In *Proceedings of the First Conference on Machine Translation*, pages 131–198, Berlin, Germany. Association for Computational Linguistics.
- Ondřej Bojar, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, and Christof Monz. 2018. [Findings of the 2018 conference on machine translation \(wmt18\)](#). In *Proceedings of the Third Conference on Machine Translation, Volume 2: Shared Task Papers*, pages 272–307, Belgium, Brussels. Association for Computational Linguistics.
- Maxime Chevalier-Boisvert, Dzmitry Bahdanau, Salem Lahlou, Lucas Willems, Chitwan Saharia, Thien Huu Nguyen, and Yoshua Bengio. 2019. [BabyAI: First steps towards grounded language learning with a human in the loop](#). In *International Conference on Learning Representations*.
- Michael Denkowski and Alon Lavie. 2011. [Meteor 1.3: Automatic metric for reliable optimization and evaluation of machine translation systems](#). In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, pages 85–91, Edinburgh, Scotland. Association for Computational Linguistics.
- Joel Escudé Font and Marta R. Costa-jussà. 2019. [Equalizing gender bias in neural machine translation with word embeddings techniques](#). In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 147–154, Florence, Italy. Association for Computational Linguistics.
- Marzieh Fadaee, Arianna Bisazza, and Christof Monz. 2018. [Examining the tip of the iceberg: A data set for idiom translation](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*. European Language Resource Association.
- J.A. Fodor and E. LePore. 2002. *The Compositionality Papers*. Clarendon Press.
- G. Frege. 1892. "über sinn und bedeutung. In Mark Textor, editor, *Funktion - Begriff - Bedeutung*, volume 4 of *Sammlung Philosophie*. Vandenhoeck & Ruprecht, Göttingen.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. [Explaining and harnessing adversarial examples](#). In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Theo M. V. Janssen. 2001. [Frege, contextuality and compositionality](#). *Journal of Logic, Language and Information*, 10(1):115–136.
- Huda Khayrallah and Philipp Koehn. 2018. [On the impact of various types of noise on neural machine translation](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 74–83. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Lei Ba. 2015. [Adam: A method for stochastic optimization](#). In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. [Opennmt: Open-source toolkit for neural machine translation](#). In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72. Association for Computational Linguistics.
- Philipp Koehn and Rebecca Knowles. 2017. Six challenges for neural machine translation. *arXiv preprint arXiv:1706.03872*.
- Katherine Lee, Orhan Firat, Ashish Agarwal, Clara Fannjiang, and David Sussillo. 2018. [Hallucinations in neural machine translation](#). In *Neural Information Processing Systems (NeurIPS) Workshop on Interpretability and Robustness for Audio, Speech, and Language*. NeurIPS.
- Vladimir I Levenshtein. 1966. [Binary codes capable of correcting deletions, insertions, and reversals](#). *Soviet physics doklady*, 10(8):707–710.
- Xian Li, Paul Michel, Antonios Anastasopoulos, Yonatan Belinkov, Nadir Durrani, Orhan Firat, Philipp Koehn, Graham Neubig, Juan Pino, and Hassan Sajjad. 2019. [Findings of the first shared task on machine translation robustness](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 91–102, Florence, Italy. Association for Computational Linguistics.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. [Effective approaches to attention-based neural machine translation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, Lisbon, Portugal. Association for Computational Linguistics.
- Paul Michel and Graham Neubig. 2018. [Mtnet: A testbed for machine translation of noisy text](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 543–553. Association for Computational Linguistics.
- Richard Montague. 1974. Universal grammar. In *Formal Philosophy: Selected Papers of Richard Montague*, number 222–247 in *Theoria*, New Haven, London. Yale University Press.
- Francis Jeffrey Pelletier. 1994. The principle of semantic compositionality. *Topoi*, 13:11–24.

- Ngoc-Quan Pham, Jan Niehues, and Alexander Waibel. 2018. [Towards one-shot learning for rare-word translation with external experts](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 100–109. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. 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. In *Proceedings of Association for Machine Translation in the Americas*, pages 223–231.
- Gabriel Stanovsky, Noah A. Smith, and Luke Zettlemoyer. 2019. [Evaluating gender bias in machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1679–1684, Florence, Italy. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.