

# Automatically Extracting Challenge Sets for Non-local Phenomena in Neural Machine Translation

Leshem Choshen<sup>1</sup> and Omri Abend<sup>1,2</sup>

<sup>1</sup>School of Computer Science and Engineering, <sup>2</sup> Department of Cognitive Sciences  
The Hebrew University of Jerusalem

leshem.choshen@mail.huji.ac.il, oabend@cs.huji.ac.il

## Abstract

We show that the state-of-the-art Transformer MT model is not biased towards monotonic reordering (unlike previous recurrent neural network models), but that nevertheless, long-distance dependencies remain a challenge for the model. Since most dependencies are short-distance, common evaluation metrics will be little influenced by how well systems perform on them. We therefore propose an automatic approach for extracting challenge sets replete with long-distance dependencies, and argue that evaluation using this methodology provides a complementary perspective on system performance. To support our claim, we compile challenge sets for English-German and German-English, which are much larger than any previously released challenge set for MT. The extracted sets are large enough to allow reliable automatic evaluation, which makes the proposed approach a scalable and practical solution for evaluating MT performance on the long-tail of syntactic phenomena.<sup>1</sup>

## 1 Introduction

The assumption that proximate source words are more likely to correspond to proximate target words has often been introduced as a bias (henceforth, *locality bias*) into statistical MT systems (Brown et al., 1993; Koehn et al., 2003; Chiang, 2005). While reordering phenomena, abundant for some language pairs, violate this simplifying assumption, it has often proved to be a useful inductive bias in practice, especially when complemented with targeted techniques for addressing non-monotonic translation (e.g., Och, 2002; Chiang, 2005). For example, if an adjective precedes a noun in one language and modifies it syntactically, it is likely that their corresponding words

<sup>1</sup>Our extracted challenge sets and codebase are found in [https://github.com/borggr/auto\\_challenge\\_sets](https://github.com/borggr/auto_challenge_sets).

will appear close to each other in the translation — i.e., they may not be immediately adjacent or even in the same order in the translation, but it is unlikely that they will be arbitrarily distant from one another.

In the era of Neural Machine Translation (NMT), such biases are implicitly introduced by the sequential nature of the LSTM architecture (Bahdanau et al., 2015, see §2). The influential Transformer model (Vaswani et al., 2017) replaces the sequential LSTMs with self-attention, which does not seem to possess this bias. We show that the default implementation of the Transformer does retain some bias, but that it can be relieved by using learned positional embeddings (§3).

Long-distance dependencies (LDD) between words and phrases present a long-standing problem for MT (Sennrich, 2016), as they are generally more difficult to detect (indeed, they pose an ongoing challenge for parsing as well (Xu et al., 2009)), and often result in non-monotonic translation if the target differs from the source in terms of its word order and lexicalization patterns. The Transformer’s indifference to the absolute position of the tokens raises the question of whether long-distance dependencies are still an open problem.

We address this question by proposing an automatic method to compile challenge sets for evaluating system performance on LDD (§4). We distinguish between two main LDD types: (1) reordering LDD, namely cases where source and target words largely correspond to one another but are ordered differently; (2) lexical LDD, where the way a word or a contiguous expression on the target side is translated is dependent on non-adjacent words on the source side.

We define a methodology for extracting both LDD types. For reordering LDD, we build on Birch (2011), whereas for lexical LDD we compile a list of linguistic phenomena that yield LDD,

and use a dependency parser to find instances of these phenomena in the source side of a parallel corpus. As a test case, we apply this method to construct challenge sets (§4.2) for German-English and English-German. The approach can be easily scaled to other languages for which a good enough parser exists.

Experimenting both with RNN and self-attention NMT architectures, we find that although the latter presents no locality bias, LDD remain challenging. Moreover, lexical LDD become increasingly challenging with their distance, suggesting that syntactic distance remains an important determinant of performance in state-of-the-art (SoTA) NMT.

We conclude that evaluating LDD using targeted challenge sets gives a detailed picture of MT performance, and underscores challenges the field has yet to fully address. As particular types of LDD are not frequent enough to significantly affect coarse-grained measures, such as BLEU (Papineni et al., 2002) or TER (Snover et al., 2006), our evaluation approach provides a complementary perspective on system performance.

## 2 Related Work

### 2.1 Long-distance Dependencies in MT

A common architecture for text-to-text generation tasks is the (Bi)LSTM encoder-decoder (Bahdanau et al., 2015). This architecture consists of several LSTM layers for the encoder and the decoder and a thin attention layer connecting them. LSTM is a recurrent network with a state vector it updates. At every step, it discards some of the current and past information and aggregates the rest into the state. Any information about the past comes from this state, which is a learned “summary” of the previous states (cf. Greff et al., 2017). Hence, for information to reach a certain prediction step, it should be stored and then kept throughout the intermediate steps (tokens). While theoretically information could be kept indefinitely (Hochreiter and Schmidhuber, 1997), practical evidence shows that LSTMs performance decreases with the distance between the trigger and the prediction (Linzen et al., 2016; Liu et al., 2018), and that they have difficulties generalizing over sequence lengths (Suzgun et al., 2018).

Despite being affected by absolute distances between syntactically dependent tokens (Linzen et al., 2016), LSTMs tend to learn to a certain

extent structural information even without being instructed to do so explicitly (Gulordava et al., 2018). Futrell and Levy (2018) discuss similar linguistic phenomena to what we discuss in §4.2, and show that LSTM encoder-decoder systems handle them better than previous N-gram based systems, despite being profoundly affected by distance.

Transformer (Vaswani et al., 2017) models are also encoder-decoder, but instead of LSTMs, they use self-attention. Self-attention is based on gating all outputs of the previous layer as inputs for the current one; put differently, it aggregates all the input in one step. This approach makes information from all parts of the input sequence equally reachable. While this is not the only architecture with such attributes (van den Oord et al., 2016), we focus on it due to its SoTA results for MT (Lakew et al., 2018). The Transformer’s use of self-attention inspired other works in related fields (Devlin et al., 2018), some of which attributed their performance gains to the model’s ability to capture long-range context (Müller et al., 2018).

As the Transformer does not aggregate input sequentially, token positions must be represented through other means. For that purpose, the embedding of each input token  $W$  is concatenated with an embedding of its position in the source sentence  $P$ . While positional embeddings can generally be any vectors, two implementations are commonly used (Tebbifakhr et al., 2018; Guo et al., 2018): learned positional embeddings (learnedPEs;  $P$  is randomly initialized), and sine positional embeddings (SinePEs) defined as:

$$P_{(pos,2i)} = \sin(pos/10,000^{2i/dim})$$

$$P_{(pos,2i+1)} = \cos(pos/10,000^{2i/dim})$$

where  $dim$  is the dimension of the embedding. Vaswani et al. (2017) report that they see no benefit in learnedPEs, and hence use SinePEs, which have much fewer parameters.

Most of the dependencies between words are short. Short-distance linguistic dependencies include some of the most common phenomena in language, such as determination, modification by an adjective and compounding. For example, 62% of the dependencies in the standard UD EWT training set (Silveira et al., 2014) are between tokens that are up to one word apart. It stands to reason that the locality bias is useful in these cases. Nevertheless, as system quality improves, rarer, more challenging dependencies become a priority,

and languages present a countless number of long-distance reordering phenomena (Deng and Xue, 2017). One example is subject-verb agreement, where a correct translation requires that the verb is inflected according to the headword of the subject (e.g., in English “dogs that ..., bark”, while “a dog that ..., barks”). When translating such cases, a locality bias may impede performance, by biasing the model not to attend to both the subject’s head and the main verb (which may be arbitrarily distant), thereby disallowing it to correctly inflect the main verb.

Due to the benefits of the locality bias, it featured prominently in statistical MT, including in the IBM models, where alignments are constrained not to cross too much (Brown et al., 1993), and in predicting probabilities of reorderings (Koehn et al., 2003; Chiang, 2005). Difficulties in handling LDD have motivated the development of syntax-based MT (Yamada and Knight, 2001), that can effectively represent reordering at the phrase level, such as when translating between VSO and SOV languages. However, syntax-based MT models remain limited in their ability to map between arbitrarily different word orders (Sun et al., 2009; Xiong et al., 2012). For example, reorderings that violate the assumption that the trees form contiguous phrases would be difficult for most such models to capture. In the next section (§3) we show that the Transformer, when implemented with learnedPEs, presents no locality bias, and hence can, in principle, learn dependencies between any two positions of the source, and use them at any step during decoding.

## 2.2 MT Evaluation

With major improvements in system performance, crude assessments of performance are becoming less satisfying, i.e., evaluation metrics do not give an indication on the performance of MT systems on important challenges for the field (Isabelle and Kuhn, 2018). String-similarity metrics against a reference are known to be partial and coarse-grained aspects of the task (Callison-Burch et al., 2006), but are still the common practice in various text generation tasks. However, their opaqueness and difficulty to interpret have led to efforts to improve evaluation measures so that they will better reflect the requirements of the task (Anderson et al., 2016; Sulem et al., 2018; Choshen and Abend, 2018b), and to increased interest in defin-

ing more interpretable and telling measures (Lo and Wu, 2011; Hodosh et al., 2013; Birch et al., 2016; Choshen and Abend, 2018a).

A promising path forward is complementing string-similarity evaluation with linguistically meaningful challenge sets. Such sets have the advantage of being interpretable: they test for specific phenomena that are important for humans and are crucial for language understanding. Interpretability also means that evaluation artefacts are more likely to be detected earlier. So far, such challenge sets were constructed for French-English (Isabelle et al., 2017; Isabelle and Kuhn, 2018) and English-Swedish (Ahrenberg, 2018)<sup>2</sup>. Previous challenge sets were compiled by manually searching corpora for specific phenomena of interest (e.g., yes-no questions which are formulated differently in English and French). These corpora are carefully made but are small in size (ten examples per phenomenon), which means that evaluation must be done manually as well.

As our methodology extracts sentences automatically based on parser output, we are able to compile much larger challenge sets, which allows us to apply standard MT measures to each sub-corpus corresponding to a specific phenomenon. The methodology is, therefore, more flexible, and can be straightforwardly adapted to accommodate future advances in MT evaluation.

## 3 Locality in SoTA NMT

In this section we show that encoder-decoder models based on BiLSTM with attention (see §2), do exhibit a locality bias, but that the Transformer, whose encoder is based on self-attention, and in which token position is encoded only through learnedPEs, does not present any such bias.

### 3.1 Methodology

In order to test whether an NMT system presents a locality bias in a controlled environment, we examine a setting of arbitrary absolute order of the source-side tokens. In this case, systems that are predisposed towards monotonic decoding are likely to present lower performance, while systems that have no predisposition as to the order of the target side tokens relative to the source-side tokens are not expected to show any change in per-

<sup>2</sup>In WMT 2019 English-German phenomena were tested with a new corpus, using both human and automatic evaluation. It is not possible, however, to use this evaluation outside the competition (Avramidis et al., 2019).

formance. In order to create a controlled setting, where source-side token order is arbitrary, we extract fixed length sentences, and apply the same permutation to all of them. We then train systems with the permuted source-side data (and the same target-side data), and compare results to a control condition where no permutation is applied.

Concretely, we experiment on a German-English setting, extracting all sentences of the most common length (18) from the WMT2015 (Bojar et al., 2015) training data. This results in 130,983 sentences, of which we hold out 1,000 sentences for testing. It is comparable in training set size to a low-resource language setting.

We set a fixed permutation  $\sigma : [18] \rightarrow [18]$  and train systems on three versions of the training data (settings): (1) REGULAR, to be used for control; (2) PERMUTED source-side, in which we apply  $\sigma$  over all source-side tokens; (3) PERPOSEMB where the positional embeddings of the source-side tokens are permuted;<sup>3</sup> and (4) REVERSED, where tokens are input in a reverse order.

We apply the following permutation,  $\sigma$ , to the source-side tokens:

$$\begin{pmatrix} 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 & 13 & 14 & 15 & 16 & 17 \\ 11 & 5 & 9 & 15 & 8 & 14 & 10 & 1 & 3 & 16 & 12 & 2 & 0 & 6 & 17 & 4 & 13 & 7 \end{pmatrix}$$

We did not find any property that would deem this permutation special (examining, e.g., its decomposition into cycles). We therefore assume that similar results will hold for other  $\sigma$ s as well.

We train a Transformer model, optimizing using Adam (Kingma and Ba, 2015). We set the embedding size to 512, dropout rate of 0.1, 6 stack layers in both the encoder and the decoder and 8 attention heads. We use tokenization, truecasing and BPE (Sennrich et al., 2016) as preprocessing, following the same protocol as (Yang et al., 2018).

We experiment both with learnedPEs, and with SinePEs. We train the BiLSTM model using the Nematus implementation (Sennrich et al., 2017b), and use their supplied scripts for preprocessing, training and testing, changing only the datasets used. For all models, we report the highest BLEU score on the test data for any epoch during training, and perform early stopping after 10 consecutive epochs without improvement.

In the Transformer with learnedPEs, 5 repetitions were done in the REGULAR setting, and 5 for

<sup>3</sup> Formally, if the source sentence is  $(t_1, \dots, t_{18})$ , then the input to the Transformer is  $([W(t_1); P(t_{\sigma(1)})], \dots, [W(t_{18}); P(t_{\sigma(18)})])$ .

Model	Positional	Setting	BLEU
Transformer	LearntPE	REGULAR	24.81
	LearntPE	PERMUTED	24.87 (+0.06)
	LearntPE	REVERSE	24.84 (+0.03)
	LearntPE	PERPOSEMB	24.82 (+0.01)
Transformer	SinePEs	REGULAR	25.08
	SinePEs	PERMUTED	23.90 (-1.18)
Nematus		REGULAR	22.32
		PERMUTED	19.67 (-2.65)

Table 1: BLEU score for various Transformer settings on regular and permuted data. In brackets are the differences from REGULAR. Nematus and Transformer using SinePEs show decreased performance when permuting the input. Transformer with learnedPEs does not. Rows correspond to the different models used (Model), which positional embeddings are fed to the Transformer (Positional), and the order of the input tokens (Setting; see text).

the other settings: 5 repetitions for PERMUTED, 1 for PERPOSEMB and 1 for REVERSED. In addition, we trained the BiLSTM model and the Transformer with SinePEs both in the REGULAR condition and in PERMUTED, each was trained once.

## 3.2 Results

Table 1 presents our results. We find that Nematus BiLSTM suffers substantially from permuting the source-side tokens, but that the Transformer does not exhibit a locality bias. Indeed, for learnedPEs in all settings (REGULAR, PERMUTED, REVERSED and PERPOSEMB), BLEU scores are essentially the same. We also find that the common practice of using fixed SinePEs does introduce some bias, as attested by the small performance drop between REGULAR and PERMUTED. Like Vaswani et al. (2017), we find that in the REGULAR settings, learnedPEs are not superior in performance to SinePEs, despite having more expressive power. However, our results suggest that the decision between learnedPEs and SinePEs is not without consequences: learnedPEs are preferable if a locality bias is undesired (this is potentially the case for highly divergent language pairs).

## 3.3 Discussion

Finding that Transformers do not present a locality bias has implications on how to construct their input in MT settings, as well as in other tasks that use self-attention encoders, such as image captioning (You et al., 2016). It is common practice to augment the source-side with globally-applicable

information, e.g., the target language in multi-lingual MT (Johnson et al., 2017). Having no locality bias implies this additional information can be added at any fixed point in the sequence fed to a Transformer, provided that the positional embeddings do not themselves introduce such a bias. This is not the case with BiLSTMs, which often require introducing the same information at each input token to allow them to be effectively used by the system (Yao et al., 2017; Rennie et al., 2017).

## 4 LDD Challenge Sets

One of the stated motivations of the Transformer model is to effectively tackle long-distance dependencies, which are “a key challenge in many sequence transduction tasks” (Vaswani et al., 2017). Our results from the previous section show that indeed fixed reordering patterns are completely transparent for Transformers. This, however, still leaves the question of how Transformers handle linguistic reordering patterns, which may involve varying distances between dependent tokens.

### 4.1 Methodology

We propose a method for scalably compiling challenge sets to support fine-grained MT evaluation for different types of LDD. We address two main types:

**Reordering LDD** are cases where the words on the two sides of the parallel corpus largely correspond to one another, but are ordered differently. These cases may require attending to source words in a highly non-monotonic order, but the generation of each target word is localized to a specific region in the source sentence. For example, in English-German, the verb in a subordinated clause appears in a final position, while the verb in the English source appears right after the subject. Consider “The man that is sitting on the chair”, and the corresponding German “Der Mann, der auf dem Stuhl sitzt” (lit. *the man, that on the chair sits*) — while the verb is placed at different clause positions in the two cases, the words mostly have direct correspondents. Our methodology follows Birch (2011) in detecting such phenomena based on alignment. Concretely, we extract a word alignment between corresponding sentences, and collect all sentences that include a pair of aligned words in the source and target sides, whose indices have a difference of at least  $d \in \mathbb{N}$ .

**Lexical LDD** are cases where the translation of a single word or phrase is determined by non-adjacent words on the source side. This requires attending to two or more regions that can be arbitrarily distant from one another. Several phenomena, such as light verbs (Isabelle and Kuhn, 2018), are known from the linguistic and MT literature to yield lexical LDD. Our methodology takes a pre-defined set of such phenomena, and defines rules for detecting each of them over dependency parses of the source-side. See §4.2 for the list of phenomena we experiment on in this paper.

Focusing on LDD, we restrict ourselves to instances where the absolute distance between the word and the dependent is at least  $d \in \mathbb{N}$ . Selecting large enough  $d$  entails that the extracted phenomena are unlikely to be memorized as a phrase with a specific meaning (e.g., encode “*make the whole thing up*” [ $d = 3$ ] as a phrase, rather than as a discontinuous phrase “*make ... up*” with an argument “*the whole thing*”). This increases the probability that such cases, if translated correctly, reflect the MT systems’ ability to recognize that such discontinuous units are likely to be translated as a single piece.

We note, that by extracting the challenge set based on syntactic parses, we by no means assume these representations are internally represented by the MT systems in any way, or assume such a representation is required for succeeding in correctly translating such constructions. The extraction method is merely a way of finding phenomena we have a reason to believe are difficult to translate, and meaningful for language understanding. We use Universal Dependencies (UD; Nivre et al., 2016) as a syntactic representation, due to its cross-lingual consistency (about 90 languages are supported so far), which allows research on difficult LDD phenomena that recur across languages.

Our extraction methods resemble previous challenge set approaches (Isabelle et al., 2017; Isabelle and Kuhn, 2018; Ahrenberg, 2018), in using linguistically motivated sets of sentence pairs to assess translation quality. However, as our extraction method is fully automatic, it allows for the compilation of much larger challenge sets over many language pairs. The challenge sets we extract contain hundreds or thousands of pairs (§4.2). The size of the sets allows using any MT evaluation measures to measure performance, and is thus a much more scalable solution than manual inspection, as

is commonly done in challenge set approaches.

On the other hand, an automatic methodology has the side-effect of being noisier, and not necessarily selecting the most representative sentences for each phenomenon. For instance *befinden sich* (lit. *to determine*) includes a verb and a reflexive pronoun, which do not necessarily appear contiguously in German. However, as *befinden* always appears with the reflexive *sich*, it might not pose a challenge to NMT systems, which can essentially ignore the reflexive pronoun upon translation.

## 4.2 A Test Case on Extracting Sets

Next, we discuss the compilation of German-English and English-German corpora. We select these pairs, as they are among the most studied in MT, and comparatively high results are obtained for them (Bojar et al., 2017). Hence, they are more likely to benefit from a fine-grained analysis.

For the reordering LDD corpus, we align each source and target sentences using FastAlign (Dyer et al., 2013) and collect all sentences with at least one pair of source-side and target-side tokens, whose indices have a difference of at least  $d = 5$ . For example:

**Source:** *Wäre es ein großer Misserfolg, nicht den Titel in der Ligue 1 zu gewinnen, wie dies in der letzten Saison der Fall war?*

**Gloss:** *Would-be it a big failure, not the title in the Ligue 1 to win, as this in the last season the case was?*

**Target:** *In Ligue 1, would not winning the title, like last season, be a big failure?*

We extract lexical LDD using simple rules over source-side parse trees, parsed with UDPipe (Straka and Straková, 2017). For a sentence to be selected, at least one word should separate the detected pair of words. We picked several well-known challenging constructions for translation that involve discontinuous phrases: reflexive-verb, verb-particle constructions and preposition stranding. We note that while these constructions often yield lexical LDDs, and are thus expected to be challenging on average, some of their instances can be translated literally (e.g., *amuse oneself* is translated to *amüsieren sich*).

**Reflexive Verbs.** Prototypically, reflexivity is the case where the subject and object corefer. Reflexive pronouns in English end with *self* or *selves* (e.g., yourselves) and in German include *sich*,

*dich*, *mich* and *uns* among others. However, reflexive pronouns can often change the meaning of a verb unpredictably, and may thus lead to different translations for non-reflexive instances of a verb, compared to reflexive ones. For example, *abheben* in German means taking off (as of a plane), but *sich abheben* means standing out. Similarly, in the example below, *drängte sich* translates to *intrude*, while *drängte* normally translates to *pushed*.

A source sentence is said to include a reflexive verb if one of its tokens is parsed with a reflexive morphological feature (`refl=yes`).

For example:

**Source:** [...] *es ertragen zu müssen, daß eine unsympathische Fremde **sich** unaufhörlich in ihren Familienkreis **drängte**.*

**Target:** [...] *to see an uncongenial alien permanently **intruded** on her own family group.*

**Phrasal Verbs** are verbs that are made up of a verb and a particle (or several particles), which may change the meaning of the verb unpredictably. Examples of English phrasal verbs include *run into* (in the sense of *meet*) and *give in*, and in German they include examples such as *einladen* (*invite*), consisting morphologically of the particle *ein* and the verb *laden* (*load*).

A source sentence is said to include a phrasal verb if a particle dependent (UD labels of `compound:prt` or `prt`) exists in the parse. *trat* in itself means *stepped*, but in the extracted example below, *trat... entgegen* translates to *received*.

For example:

**Source:** [...] *ich **trat** ihm in wahnsinniger Wut **entgegen**.*

**Target:** [...] *I **received** him in frantic sort.*

**Preposition Stranding** is the case where a preposition does not appear adjacent to the object it refers to. In English, it will often appear at the end of the sentence or a clause. For example, *The banana she stepped **on*** or *The boy I read the book **to***. Preposition stranding is common in English and other languages such as Scandinavian languages or Dutch (Hornstein and Weinberg, 1981). However, in German, it is not a part of standard written language (Beermann and Ik-Han, 2005), although it does (rarely) appear (Fanselow, 1983). We, therefore, extract this challenge set only with English as the source side.

While preposition stranding is often regarded as a syntactic phenomenon, we consider it here a lexical LDD, since the translation of prepositions

	Phenomena	Books	Newstest2013
De↔En	Reorder	7,457	306
	Baseline (full dataset)	51,467	3,000

Table 2: Sizes of reordering and baseline corpora.

		Min Distance				
		All	≥1	≥2	≥3	News
De→En	Phenomena					
	Particle	8,361	7,584	6,261	4,780	232
	Reflexive	13,207	8,122	5,598	4,226	281
En→De	Particle	4,636	786	111	36	17
	Reflexive	3,225	1,188	460	274	11
	Preposition Stranding	682	191	85	40	8

Table 3: Sizes of Lexical LDD corpora. Challenge sets are partitioned (in order of appearance) by the language pairs, the phenomenon type, and the minimal distance between the head and the dependent. Phenomenon appears in the source. Statistics for the Newstest2013 corpora with minimal distance  $\geq 1$  are at the rightmost column, the rest are on Books.

(and in some cases their accompanying verbs) is dependent on the prepositional object, which in the case of preposition stranding, may be distant from the preposition itself. For example, translating *the car we looked for* into German usually uses the verb *suchen* (*search*), while translating *the car we looked at* does not. Translating prepositions is difficult in general (Hashemi and Hwa, 2014), but preposition stranding is especially so, as there is no adjacent object to assist disambiguation.

A source sentence is said to include preposition stranding if it contains two nodes with an edge of the type `obl` (oblique) or a subcategory thereof between them, and the UD POS tag of the dependent is adposition (ADP).

For example,

**Source:** [...] *wherever she wanted to send the hedgehog to* [...]

**Gloss:** [...] *where she the hedgehog rolled-towards wanted* [...]

**Target:** [...] *wo sie den Igel hinrollen wollte* [...]

### 4.3 Experiments

We turn to evaluate SoTA NMT performance on the extracted challenge sets.

**Experimental Setup.** We trained the Transformer on WMT2015 training data (Bojar et al., 2015), for parameters see §3.1. For Nematus we used the non-ensemble pre-trained model from (Sennrich et al., 2017a). Each of the test sets, either a baseline or a challenge sets, for the Transformer and Nematus used a maximum of 10k and

		Transformer		Nematus	
		Books	News	Books	News
De→En	Baseline	9.02	28.23	16.26	26.32
	Reorder	7.16	22.68	13.88	22.73
	Particle	7.52	27.46	15.41	23.98
	Reflexive	8.15	27.84	14.91	27.04
En→De	Baseline	6.33	23.7	12.25	22.03
	Reorder	4.31	19.4	9.02	20.38
	Particle	5.30	17.83	9.55	16.72
	Reflexive	5.07	15.77	9.97	21.81
	Preposition Stranding	5.37	11.82	9.73	6.27

Table 4: BLEU scores on the challenge sets. Minimum distance between head and dependent  $d \geq 1$ . A clear, consistent drop from the Baseline (full corpus) score is observed in all cases. The top part of the table corresponds to German-to-English (De→En) sets, and bottom part to English-to-German (En→De) sets. Within each part, rows correspond to various linguistic phenomena (second column), including reordering LDD (Reorder), Verb-Particle Constructions (Particle), Reflexive Verbs (Reflexive) and Preposition Stranding. Columns correspond to the models (Transformer/Nematus), and the domains (Books/News).

1k sentences per set respectively.<sup>4</sup>

Two parallel corpora were used for extracting the challenge sets. One is newstest2013 (Bojar et al., 2015) from the news domain that is commonly used as a development set for English-German. The other is the relatively unused Books corpus (Tiedemann, 2012) from the more challenging domain of literary translation. The corpora are of sizes 51K and 3K respectively. For lexical LDD, we took the distance ( $d$ ) between the relevant words to be at least 1, meaning there is at least one word separating them. See Tables 2, 3 for the sizes of the extracted corpora.

For evaluation, we use the MOSES implementation of BLEU (Papineni et al., 2002; Koehn et al., 2007), and for reordering LDD, also RIBES (Isozaki et al., 2010), which focuses on reordering. RIBES measures the correlation of n-gram ranks between the output and the reference, where n-gram appears uniquely and in both.

**Manual Validation.** To assess the ability of our procedure to extract relevant LDDs, we manually analyzed over 180 source German sentences ex-

<sup>4</sup>We subsample a smaller test set for Nematus, since the most competitive model for the language pair requires Theano. As Theano is deprecated for two years now, it cannot run on our GPUs, which entails long inference time.

	Language	Phenomena	BLEU			Spearman	
			All	$\geq 1$	$\geq 2$		$\geq 3$
Transformer	German	Particle	7.56	7.52	7.5	7.49	-0.96
		Reflexive	8.26	8.15	8.04	-	-1
	English	Particle	4.96	5.3	4.96	6.01	0.73
		Reflexive	5.41	5.07	5.25	5.04	-0.63
Nematus	German	Preposition Stranding	5	5.37	4.42	4.64	-0.63
		Particle	15.48	15.41	14.45	12.36	-0.92
		Reflexive	15.27	14.91	15.13	13.14	-0.80
	English	Particle	10.14	9.55	9.36	8.82	-0.98
		Reflexive	9.46	9.97	9.54	9.35	-0.36
		Preposition Stranding	10.01	9.73	9.14	9.04	-0.97

Table 5: The effect of dependency distance for lexical LDDs on SoTA performance. Results are in BLEU over the Books challenge sets. Columns correspond to the minimum distance, where *All* does not restrict distance (control). The rightmost column presents the Spearman correlation of the phenomena’s score with the minimum distance used. All correlations but one are highly negative, implying that distance has a negative effect on performance.

tracted from Books, and 81 English ones including all the instances extracted from News and 45 extracted from Books, where instances are evenly distributed between phenomena and distance of exactly 1,2 or 5. We find that 85% of German sentences, 87% of the English News sentences and 86% of the Books ones indeed contain the target phenomenon. For details of the manual evaluation of the extraction procedure, see Appendix 1.

		News	Books
German	Baseline	0.82	0.57
	Reorder	0.79	0.54
English	Baseline	0.79	0.56
	Reorder	0.77	0.53

Table 6: RIBES scores on the reordering LDD challenge sets. Sentences extracted as being challenging to reorder are harder for the Transformer (lower score). This trend is consistent with our experiments with BLEU. First column indicates the source language.

**Results.** Comparison of the overall BLEU scores of the NMT models (Table 4) against their performance on the challenge sets, shows that the phenomena are challenging for both models. Both in the small development set of newstest2013 and the large set of Books, the challenge subparts are more challenging across the board. For reordering LDD, we further apply RIBES and find a similar trend: RIBES score is lower for the reorder challenge set than the baseline (see Table 6).

In order to confirm that the distance between the head and dependent (the “length” of the depen-

dency) is related to the observed performance drop in the case of lexical LDD, we partition each of the challenge sets according to their length ( $d$ ), and compare the results to a control condition, where all instances of the phenomena listed in §4.2 are extracted, including non-LDD instances, i.e., sentences where the head and the dependent are adjacent. System performance on the sliced challenge sets (Table 5) shows that performance indeed decreases with  $d$ . Results thus indicate that it is not only the presence of the phenomena that make these sets challenging, but that the challenge increases with the distance.

We validate this main finding using manual annotation of German to English cases. Using two annotators (with high agreement between them;  $\kappa=0.79$ ), we find that the decrease in performance with  $d$  is replicated. We measure how many of the detected lexical LDD are correctly translated, ignoring the rest of the source and output, as done in manual challenge set approaches. We find that 60%, 54% and 38% of the cases are translated correctly for  $d \in (1, 2, 5)$ , respectively. This suggests that the extracted phenomena and the distance indeed pose a challenge, and that the automatic metric we use shows the correct trend in these cases. See Appendix 2 for details.

**Discussion.** Interestingly, these results hold true for the Transformer despite its indifference to the absolute word order. Therefore, word distance in itself is not what makes such phenomena challenging, contrary to what one might expect from the definition of LDD. It seems then that these



phenomena are especially challenging due to the non-standard linguistic structure (e.g., syntactic and lexical structure), and the varying distances in which LDD manifest themselves. The models, therefore, seem to be unable to learn the linguistic structure underlying these phenomena, which may motivate more explicit modelling of linguistic biases into NMT models, as proposed by, e.g., [Eriguchi et al. \(2017\)](#) and [Song et al. \(2019\)](#).

We note that our experiments were not designed to compare the performance of BiLSTM and self-attention models. We, therefore, do not see the Transformer’s inferior performance on Books, relative to Nematus as an indication of the general ability of this model in out-of-domain settings. What is evident from the results is that translating Books is a challenge in itself, probably due to the register of the language, and the presence of frequent non-literal translations.

A potential confound is that performance might change with the length of the source in BiLSTMs ([Carpuat et al., 2013](#); [Murray and Chiang, 2018](#)), in Transformers it was reported to increase ([Zhang et al., 2018](#)). Length is generally greater in the challenge set than in the full test set, and generally increases with  $d$ , showing if anything a decrease of performance by length. To assess whether our corpora are challenging due to a length bias, we randomly sample from Books 1,000 corpora with 1,000, 100 and 10 sentences each. The correlation between their corresponding average length and the Transformers’ BLEU score on them was 0.06, 0.09 and 0.03 respectively. While this suggests length is not a strong predictor of performance, to verify that difficulty is not a result of the distribution of lengths in the challenge sets we conduct another experiment.

For each challenge set and each value of  $d$  (0–3), we sample 100 corpora. For each sentence in a given challenge set, we sample a sentence of no more than a difference of 1 in length. This results in a corpus with a similar length distribution, but sampled from the overall population of Books sentences. Results show that the BLEU score of the challenge sets in all German to English cases is lower than any randomly sampled corpus.<sup>5</sup> In the English-German cases, trends are similar, albeit less pronounced. This may be due to the low number of long English sentences, which lead to

<sup>5</sup>Most sampled corpora actually had better scores than the baseline. We believe this is because very short sentences which are mostly noise, are never sampled.

more homogeneous samples. Overall, results suggest that length is extremely unlikely to be the only cause for the observed trends.

## 5 Conclusion

As NMT system performance is constantly improving, more reliable methods for identifying and classifying their failures are needed. Much research effort is therefore devoted to developing more fine-grained and interpretable evaluation methods, including challenge-set approaches. In this paper, we showed that, using a UD parser, it is possible to extract challenge sets that are large enough to allow scalable MT evaluation of important and challenging phenomena.

An accumulating body of research is devoted to the ability of modern neural architectures such as LSTMs ([Linzen et al., 2016](#)) and pretrained embeddings ([Hewitt and Manning, 2019](#); [Liu et al., 2019](#); [Jawahar et al., 2019](#)) to represent linguistic features. This paper makes a contribution to this literature in confirming that the Transformer model can indeed be made indifferent to the absolute order of the words, but also shows that this does not entail that the model can overcome the difficulties of LDD in naturalistic data. We may carefully conclude then that despite the remarkable feats of current NMT models, inducing linguistic structure in its more evasive and challenging instances is still beyond the reach of state-of-the-art NMT, which motivates exploring more linguistically-informed models.

## 6 Acknowledgments

This work was supported by the Israel Science Foundation (grant no. 929/17)

## References

- Lars Ahrenberg. 2018. A challenge set for english-swedish machine translation. In *SLTC*.
- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. *Spice: Semantic propositional image caption evaluation*. In *European Conference on Computer Vision*, pages 382–398. Springer.
- Eleftherios Avramidis, Vivien Macketanz, Ursula Strohriegel, and Hans Uszkoreit. 2019. Linguistic evaluation of german-english machine translation using a test suite. In *Proceedings of the Fourth Conference on Machine Translation. Conference on Machine Translation (WMT-2019), located at The*

- 57th Annual Meeting of the Association for Computational Linguistics, August 1-2, Florence, Italy. Association for Computational Linguistics.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473.
- Dorothee Beermann and Lars Ik-Han. 2005. Preposition stranding and locative adverbs in german. *Organizing Grammar*, 86:31.
- Alexandra Birch. 2011. *Reordering metrics for statistical machine translation*. Ph.D. thesis, The University of Edinburgh.
- Alexandra Birch, Omri Abend, Ondřej Bojar, and Barry Haddow. 2016. Hume: Human ucca-based evaluation of machine translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1264–1274.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Shujian Huang, Matthias Huck, Philipp Koehn, Qun Liu, Varvara Logacheva, et al. 2017. Findings of the 2017 conference on machine translation (wmt17). In *Proceedings of the Second Conference on Machine Translation*, pages 169–214.
- Ondrej Bojar, Rajen Chatterjee, Christian Federmann, Barry Haddow, Matthias Huck, Chris Hokamp, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Matt Post, Carolina Scarton, Lucia Specia, and Marco Turchi. 2015. Findings of the 2015 workshop on statistical machine translation. In *WMT@EMNLP*.
- Peter F Brown, Vincent J Della Pietra, Stephen A Della Pietra, and Robert L Mercer. 1993. The mathematics of statistical machine translation: Parameter estimation. *Computational linguistics*, 19(2):263–311.
- Chris Callison-Burch, Miles Osborne, and Philipp Koehn. 2006. Re-evaluation the role of bleu in machine translation research. In *EACL*.
- Marine Carpuat, Lucia Specia, and Dekai Wu. 2013. Proceedings of ssst-7, seventh workshop on syntax, semantics and structure in statistical translation. In *EMNLP 2014*.
- David Chiang. 2005. A hierarchical phrase-based model for statistical machine translation. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, pages 263–270. Association for Computational Linguistics.
- Leshem Choshen and Omri Abend. 2018a. Automatic metric validation for grammatical error correction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Leshem Choshen and Omri Abend. 2018b. Referenceless measure of faithfulness for grammatical error correction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Dun Deng and Nianwen Xue. 2017. Translation divergences in chinese–english machine translation: An empirical investigation. *Computational Linguistics*, 43(3):521–565.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805.
- Chris Dyer, Victor Chahuneau, and Noah A Smith. 2013. A simple, fast, and effective reparameterization of ibm model 2. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–648.
- Akiko Eriguchi, Yoshimasa Tsuruoka, and Kyunghyun Cho. 2017. Learning to parse and translate improves neural machine translation. In *ACL*.
- Gisbert Fanselow. 1983. Zu einigen problemen von kasus, rektion und bindung in der deutschen syntax. *Universität Konstanz*.
- Richard Futrell and Roger P. Levy. 2018. Do rnns learn human-like abstract word order preferences? *CoRR*, abs/1811.01866.
- Klaus Greff, Rupesh K Srivastava, Jan Koutník, Bas R Steunebrink, and Jürgen Schmidhuber. 2017. Lstm: A search space odyssey. *IEEE transactions on neural networks and learning systems*, 28(10):2222–2232.
- Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. 2018. Colorless green recurrent networks dream hierarchically. In *NAACL-HLT*.
- Qipeng Guo, Xipeng Qiu, Pengfei Liu, Yunfan Shao, Xiangyang Xue, and Zheng Zhang. 2018. Star-transformer. *arXiv preprint arXiv:1902.09113*.
- Homa B Hashemi and Rebecca Hwa. 2014. A comparison of mt errors and esl errors. In *LREC*, pages 2696–2700.
- John Hewitt and Christopher D. Manning. 2019. A structural probe for finding syntax in word representations. 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 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.

- Micah Hodosh, Peter Young, and Julia Hockenmaier. 2013. Framing image description as a ranking task: Data, models and evaluation metrics. *Journal of Artificial Intelligence Research*, 47:853–899.
- Norbert Hornstein and Amy Weinberg. 1981. Case theory and preposition stranding. *Linguistic inquiry*, 12(1):55–91.
- Pierre Isabelle, Colin Cherry, and George F. Foster. 2017. A challenge set approach to evaluating machine translation. In *EMNLP*.
- Pierre Isabelle and Roland Kuhn. 2018. A challenge set for french→english machine translation. *arXiv preprint arXiv:1806.02725*.
- Hideki Isozaki, Tsutomu Hirao, Kevin Duh, Katsuhito Sudoh, and Hajime Tsukada. 2010. Automatic evaluation of translation quality for distant language pairs. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 944–952. Association for Computational Linguistics.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. [What does BERT learn about the structure of language?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3651–3657, Florence, Italy. Association for Computational Linguistics.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda B. Viégas, Martin Wattenberg, Gregory S. 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.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondrej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180.
- Philipp Koehn, Franz Josef Och, and Daniel Marcu. 2003. Statistical phrase-based translation. In *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology-Volume 1*, pages 48–54. Association for Computational Linguistics.
- Surafel Melaku Lakew, Mauro Cettolo, and Marcello Federico. 2018. A comparison of transformer and recurrent neural networks on multilingual neural machine translation. In *COLING*.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. Assessing the ability of lstms to learn syntax-sensitive dependencies. *Transactions of the Association for Computational Linguistics*, 4:521–535.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019. [Linguistic knowledge and transferability of contextual representations](#). 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 1073–1094, Minneapolis, Minnesota. Association for Computational Linguistics.
- Nelson F. Liu, Omer Levy, Roy Schwartz, Chenhao Tan, and Noah A. Smith. 2018. Lstms exploit linguistic attributes of data. In *Rep4NLP@ACL*.
- Chi-kiu Lo and Dekai Wu. 2011. Meant: an inexpensive, high-accuracy, semi-automatic metric for evaluating translation utility via semantic frames. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 220–229. 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 *WMT*.
- Kenton Murray and David Chiang. 2018. Correcting length bias in neural machine translation. In *WMT*.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D. Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty, and Daniel Zeman. 2016. Universal dependencies v1: A multilingual treebank collection. In *Proc. of LREC*, pages 1659–1666.
- Franz Josef Och. 2002. *Statistical machine translation: from single-word models to alignment templates*. Ph.D. thesis, Bibliothek der RWTH Aachen.
- Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W. Senior, and Koray Kavukcuoglu. 2016. Wavenet: A generative model for raw audio. In *SSW*.
- 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 on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.

- Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2017. Self-critical sequence training for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7008–7024.
- Rico Sennrich. 2016. How grammatical is characterlevel neural machine translation. *Assessing MT quality with contrastive translation pairs. CoRR, abs/1612.04629*.
- Rico Sennrich, Alexandra Birch, Anna Currey, Ulrich Germann, Barry Haddow, Kenneth Heafield, Antonio Valerio Miceli Barone, and Philip Williams. 2017a. The university of edinburgh’s neural mt systems for wmt17. In *WMT*.
- Rico Sennrich, Orhan Firat, Kyunghyun Cho, Alexandra Birch, Barry Haddow, Julian Hirschler, Marcin Junczys-Dowmunt, Samuel Läubli, Antonio Valerio Miceli Barone, Jozef Mokry, and Maria Nadejde. 2017b. [Nematus: a toolkit for neural machine translation](#). In *Proceedings of the Software Demonstrations of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, pages 65–68, Valencia, Spain. 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)*, volume 1, pages 1715–1725.
- Natalia Silveira, Timothy Dozat, Marie-Catherine de Marneffe, Samuel Bowman, Miriam Connor, John Bauer, and Christopher D. Manning. 2014. A gold standard dependency corpus for English. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC-2014)*.
- 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*, volume 200.
- Linfeng Song, Daniel Gildea, Yue Zhang, Zhiguo Wang, and Jinsong Su. 2019. Semantic neural machine translation using amr. *arXiv preprint arXiv:1902.07282*.
- Milan Straka and Jana Straková. 2017. [Tokenizing, pos tagging, lemmatizing and parsing ud 2.0 with udpipeline](#). In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 88–99, Vancouver, Canada. Association for Computational Linguistics.
- Elior Sulem, Omri Abend, and Ari Rappoport. 2018. Semantic structural evaluation for text simplification. In *NAACL-HLT*.
- Jun Sun, Min Zhang, and Chew Lim Tan. 2009. A non-contiguous tree sequence alignment-based model for statistical machine translation. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2-Volume 2*, pages 914–922. Association for Computational Linguistics.
- Mirac Suzgun, Yonatan Belinkov, and Stuart M Shieber. 2018. On evaluating the generalization of lstm models in formal languages. *arXiv preprint arXiv:1811.01001*.
- Amirhossein Tebbifakhr, Ruchit Agrawal, Matteo Negri, and Marco Turchi. 2018. Multi-source transformer with combined losses for automatic post editing. In *WMT*.
- Jörg Tiedemann. 2012. Parallel data, tools and interfaces in opus. In *Lrec*, volume 2012, pages 2214–2218.
- 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 *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Deyi Xiong, Min Zhang, and Haizhou Li. 2012. Modeling the translation of predicate-argument structure for smt. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pages 902–911. Association for Computational Linguistics.
- Peng Xu, Jaeho Kang, Michael Ringgaard, and Franz Och. 2009. Using a dependency parser to improve smt for subject-object-verb languages. In *Proceedings of human language technologies: The 2009 annual conference of the North American chapter of the association for computational linguistics*, pages 245–253. Association for Computational Linguistics.
- Kenji Yamada and Kevin Knight. 2001. A syntax-based statistical translation model. In *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics*.
- Zhen Yang, Wei Chen, Feng Wang, and Bo Xu. 2018. [Improving neural machine translation with conditional sequence generative adversarial nets](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1346–1355. Association for Computational Linguistics.
- Ting Yao, Yingwei Pan, Yehao Li, Zhaofan Qiu, and Tao Mei. 2017. Boosting image captioning with attributes. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 4904–4912.

Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. 2016. Image captioning with semantic attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4651–4659.

Biao Zhang, Deyi Xiong, and Jinsong Su. 2018. Accelerating neural transformer via an average attention network. In *ACL*.