

Beyond Error Propagation in Neural Machine Translation: Characteristics of Language Also Matter

Lijun Wu^{1*}, Xu Tan^{2*}, Di He³, Fei Tian², Tao Qin², Jianhuang Lai¹ and Tie-Yan Liu²

¹School of Data and Computer Science, Sun Yat-sen University

²Microsoft Research

³Key Laboratory of Machine Perception, MOE, School of EECS, Peking University

wulijun3@mail2.sysu.edu.cn; stsljh@mail.sysu.edu.cn;

{xuta,fetia,taoqin,tyliu}@microsoft.com; di_he@pku.edu.cn

Abstract

Neural machine translation usually adopts autoregressive models and suffers from exposure bias as well as the consequent error propagation problem. Many previous works have discussed the relationship between error propagation and the *accuracy drop* (i.e., the left part of the translated sentence is often better than its right part in left-to-right decoding models) problem. In this paper, we conduct a series of analyses to deeply understand this problem and get several interesting findings. (1) The role of error propagation on accuracy drop is overstated in the literature, although it indeed contributes to the accuracy drop problem. (2) Characteristics of a language play a more important role in causing the accuracy drop: the left part of the translation result in a right-branching language (e.g., English) is more likely to be more accurate than its right part, while the right part is more accurate for a left-branching language (e.g., Japanese). Our discoveries are confirmed on different model structures including Transformer and RNN, and in other sequence generation tasks such as text summarization.

1 Introduction

Neural machine translation (NMT) has attracted much research attention in recent years (Bahdanau et al., 2014; Shen et al., 2018; Song et al., 2018; Xia et al., 2018; He et al., 2016; Wu et al., 2017, 2018). The major approach to the task typically leverages an encoder-decoder framework (Cho et al., 2014; Sutskever et al., 2014) and the decoder usually generates the target tokens one by one from left to right autoregressively, in which the generation of a target token is conditioned on previously generated target tokens.

It has been observed that for an NMT model with left-to-right decoding, the right part words in

its translation results are usually worse than the left part words in terms of accuracy (Zhang et al., 2018; Bengio et al., 2015; Ranzato et al., 2015; Hassan et al., 2018; Liu et al., 2016b,a). This phenomenon is referred to as *accuracy drop* in this paper. A straightforward explanation to accuracy drop is *error propagation*: If a word is mistakenly predicted during inference, the error will be propagated and the future words conditioned on this one will be impacted. Different methods have been proposed to address the problem of accuracy drop (Liu et al., 2016a,b; Hassan et al., 2018).

Instead of solving the problem, in this paper, we aim to deeply understand the causes of the problem. In particular, we want to answer the following two questions:

- Is error propagation the main cause of accuracy drop?
- Are there any other causes leading to accuracy drop?

To answer these two questions, we conduct a series of experiments to analyze the problem.

First, we train NMT models separately using left-to-right and right-to-left decoding (Sennrich et al., 2016; Liu et al., 2016b; He et al., 2017; Gao et al., 2018) on several language pairs (i.e., German to English, English to German, and English to Chinese). If error propagation is the main cause of accuracy drop, then the right part words in the translation results generated by right-to-left NMT models should be more accurate than the left part words. However, we observe the opposite phenomenon that the accuracy of the right part words of the translated sentences in both left-to-right and right-to-left models is lower than that of the left part, which contradicts with error propagation. This shows that error propagation alone cannot well explain the accuracy drop and even

*Authors contribute equally to this work.

suggests that error propagation may not exist or matter.

Second, to further investigate the influence of error propagation on accuracy drop, we conduct a set of experiments with teacher forcing (Williams and Zipser, 1989) during inference, in which we feed the ground-truth preceding words to predict the next target word. Teacher forcing eliminates exposure bias as well as error propagation in inference. The results verify the existence of error propagation, since the later part (the right part in left-to-right decoding and the left part in right-to-left decoding) of the translation results get more accuracy improvement with teacher forcing, regardless of the decoding direction. Meanwhile, the accuracy of the right part is still lower than that of the left part with teacher forcing, which demonstrates that there must be some other causes apart from error propagation leading to accuracy drop.

Third, inspired by linguistics, we find that the concept of branching (Berg et al., 2011; Payne, 2006) can help to explain the problem. We conduct the third set of experiments to study the correlation between language branching and accuracy drop. We find that if a target language is right branching such as English, the accuracy of the left part words is usually higher than that of the right part words, no matter for left-to-right or right-to-left NMT models, while for a left-branching target language such as Japanese, the accuracy of the left part words is usually lower than that of the right part, no matter for which models. The intuitive explanation is that a right-branching language has a clearer structure pattern (easier to predict) in the left part of sentence than that in the right part, since the main subject of the sentence is usually put in the left part. We calculate two statistics to verify this assumption: n-gram statistics (including n-gram frequency and conditional probabilities) and dependency parsing statistics. For right-branching languages, we found higher n-gram frequency/conditional probabilities as well as more dependencies in the left part compared with that in the right part. The opposite results are also found in left-branching languages.

We summarize our findings as follows.

- Through empirical analyses, we find that the influence of error propagation is overstated in the literature, which may misguide the future research. Error propagation alone cannot fully explain the accuracy drop in the left or

right part of sentence.

- We find the branching in linguistics well correlates with accuracy drop in the left or right part of sentence and the corresponding analysis on n-gram and dependency parsing statistics well explain this phenomenon.

Our studies show that linguistics can be very helpful to understand existing machine learning models and build better models for language related tasks. We hope that our work can bring some insights to the research on neural machine translation. We believe that our findings can help us to design better translation models. For example, the finding on language branching suggests us to use left-to-right NMT models for right-branching languages such as English and right-to-right NMT models for left-branching languages such as Japanese.

2 Related Work

2.1 Exposure Bias and Error Propagation

Exposure bias and error propagation are two different concepts but often mentioned together in literature (Bengio et al., 2015; Shen et al., 2016; Ranzato et al., 2015; Liu et al., 2016b,a; Zhang et al., 2018; Hassan et al., 2018). Exposure bias refers to the fact that the sequence generation model is usually trained with teacher-forcing while generates the sequence autoaggressively during inference. This discrepancy between training and inference can yield errors that accumulate quickly along the generated sequence, which is known as error propagation (Bengio et al., 2015; Shen et al., 2016; Ranzato et al., 2015).

Bengio et al. (2015) propose the scheduled sampling method to eliminate the exposure bias and the resulting error propagation, which achieves promising performance on sequence generation tasks such as image captioning. Shen et al. (2016); Ranzato et al. (2015) improve the basic maximum likelihood estimation (MLE) with reinforcement learning or minimum risk training and aim to address the limitation of MLE training and exposure bias problem.

2.2 Tackling Accuracy Drop

(Liu et al., 2016b,a; Zhang et al., 2018; Hassan et al., 2018) mainly ascribe accuracy drop (the accuracy of right part words is worse than that in the left part in most cases) to error propagation and

propose different methods to solve this problem. Liu et al. (2016b,a); Hassan et al. (2018) use agreement regularization between the left-to-right and right-to-left models to achieve better performance. Zhang et al. (2018) and (Hassan et al., 2018) propose to use two-pass decoding to refine the generated sequence to yield better quality.

All these works focus on error propagation and accuracy drop. To our knowledge, there is no deep study about other causes of accuracy drop. In this paper, we aim to conduct such a study. Our study shows that accuracy drop is not only caused by error propagation, but also the characteristics of language itself.

3 Error Propagation and Accuracy Drop

3.1 Error Propagation is Not the Only Cause

A left-to-right NMT model feeds target tokens one by one from left to right in training and generate target tokens one by one from left to right during inference, while a right-to-left NMT model trains and generates token in the reverse direction. Intuitively, if error propagation is the root cause of accuracy drop, then a right-to-left NMT model will generate translations with better right half accuracy than the left half. In this section, we study the results of both left-to-right and right-to-left NMT models to analyze the relationship between error propagation and accuracy drop.

We conduct experiments on three translation tasks with different language pairs, which include: IWSLT 2014 German-English (De-En), WMT 2014 English-German (En-De) and WMT 2017 English-Chinese (En-Zh). We choose the state-of-the-art NMT model Transformer (Vaswani et al., 2017) as the basic model structure and train two separate models with left-to-right and right-to-left decoding on each language pair. More details about the datasets and model descriptions can be found in supplementary materials (section A.1 and A.2). We evenly split each generated sentence into the left half and the right half with same number of words¹. Then for both the left and right half, we compute their accuracy with respect to the reference target sentence, in terms of BLEU score (Pa-

¹) For most of the sentences, the last word of the sentence is period which is easy to decode. To make a fair comparison, we simply remove the last period before dividing the translation sentence. 2) For sentence with an odd number of words, we simply remove the word in the middle position to make the left half and right half have the same number of words.

| | De-En | En-De | En-Zh |
|---------------|-------|-------|-------|
| left-to-right | 31.42 | 26.93 | 20.79 |
| right-to-left | 30.00 | 25.35 | 20.23 |

Table 1: BLEU scores on the test set of the three translation tasks with both left-to-right and right-to-left decoding.

| left-to-right | De-En | En-De | En-Zh |
|----------------------|--------------|-------------|-------------|
| Left | 10.17 | 7.90 | 7.41 |
| Right | 8.39 | 6.60 | 5.91 |
| right-to-left | De-En | En-De | En-Zh |
| Right | 7.83 | 6.45 | 5.77 |
| Left | 9.41 | 7.11 | 7.01 |

Table 2: BLEU scores of the left and right half of left-to-right and right-to-left NMT models. In (Liu et al., 2016a), the authors report the *partial BLEU* score without length penalty, our result is consistent with *partial BLEU* if simply removing length penalty when calculating BLEU.

pineni et al., 2002)².

We first report the BLEU scores of the full translation results (without split) in Table 1. As can be seen, the accuracy of the model is comparable to state-of-the-art results (Vaswani et al., 2017; Wang et al., 2017, 2018). Afterwards we report the BLEU scores of the left half and the right half in Table 2. We have several observations.

- When translating from left-to-right, the BLEU score of the left half is higher than the right half on all the three tasks, which is consistent with previous observation and is able to be explained via error propagation.
- When translating from right-to-left, the accuracy of the left half (in this way it’s the later part of the generated sentence) is still higher than the right half. Such an observation is contradictory to the previous analyses between error propagation and accuracy drop, which regard that accumulated error brought by exposure bias will deteriorate the quality in later part of translation (i.e., the left half).

The inconsistent observation above suggests that error propagation is not the only cause of accuracy drop that there are other factors beyond er-

²We use the *multi-bleu.perl* script <https://github.com/moses-smt/mosesdecoder/scripts/generic/multi-bleu.perl>. When computing BLEU score of the left or right half, the reference is the full reference sentence.

| De-En | left-to-right | | | right-to-left | | |
|-------|---------------|--------------|----------|---------------|--------------|----------|
| | 0 | 1 | Δ | 0 | 1 | Δ |
| Left | 10.17 | 10.71 | 0.54 | 9.41 | 10.41 | 1.00 |
| Right | 8.39 | 9.25 | 0.86 | 7.83 | 8.45 | 0.62 |

| En-De | left-to-right | | | right-to-left | | |
|-------|---------------|-------------|----------|---------------|--------------|----------|
| | 0 | 1 | Δ | 0 | 1 | Δ |
| Left | 7.90 | 9.43 | 1.53 | 7.11 | 10.71 | 3.60 |
| Right | 6.60 | 8.36 | 1.76 | 6.45 | 8.37 | 1.92 |

| En-Zh | left-to-right | | | right-to-left | | |
|-------|---------------|-------------|----------|---------------|-------------|----------|
| | 0 | 1 | Δ | 0 | 1 | Δ |
| Left | 7.41 | 9.11 | 1.70 | 7.01 | 9.83 | 2.82 |
| Right | 5.91 | 8.55 | 2.64 | 5.77 | 7.54 | 1.77 |

Table 3: BLEU scores. "0" represents the translation results without teacher forcing during inference, and "1" represents the translation results with teacher forcing during inference. Δ represents the BLEU score improvement of teacher forcing over normal translation.

ror propagation for accuracy drop. It even challenges the existence of error propagation: does error propagation really exist? In the next section we try to answer this question through teacher forcing experiments.

3.2 The Influence of Error Propagation

Teacher forcing (Williams and Zipser, 1989) in sequence generation means that when training a sequence generation model, we feed the previous ground-truth tokens as inputs to predict the next target word. Here we apply teacher forcing in the inference phase of NMT: to generate the next word \hat{y}_i , we input the preceding ground-truth words $y_{<i}$ rather than the previously generated words $\hat{y}_{<i}$, which largely alleviates the effect of error propagation, since there will be no error propagated from the previously generated words.

Same as last section, we evaluate the quality of the left and right half of the translation results generated by both the left-to-right and right-to-left models. The results are summarized in Table 3. For comparison, we also include the BLEU scores of normal translation (without teacher forcing). We have several findings from Table 3 as follows:

- Exposure bias exists. The accuracy of both left and right half tokens in the normal translation is lower than that in teacher forcing, which feeds the ground-truth tokens as inputs. This demonstrates that feeding the previously generated tokens (which might be in-

correct) in inference indeed hurts translation accuracy.

- Error propagation does exist. We find the error is accumulated along the sequential generation of the sentence. Taking En-Zh and the left-to-right NMT model as an example, the BLEU score improvement of the right half (the second half of the generation) of teacher forcing over normal translation is 2.64, which is much larger than the accuracy improvement of the left half (the first half of the generation): 1.70. Similarly, for En-Zh with the right-to-left NMT model, the BLEU score improvement of the left half (the second half of the generation) of teacher forcing over normal translation is 2.82, which is much larger than the accuracy improvement of the right half (the first half of the generation): 1.77.
- Other causes exist. Taking En-De translation with the left-to-right model as an example, the accuracy of the left half (9.43) is higher than that of the right half (8.36) when there is no error propagation with teacher forcing. Similar results can be found in other language pairs and models. This suggests that there must be some other causes leading to accuracy drop, which will be studied in the next section.

4 Language Branching Matters

Section 3.1 and 3.2 together show that error propagation has influence on but is not the only cause of accuracy drop. We hypothesize that the language itself, i.e., its characteristics, may explain the phenomenon of accuracy drop.

Watanabe and Sumita (2002) finds that left-to-right decoding performs better for Japanese-English translation while right-to-left decoding performs better for English-Japanese translation. We conduct the same analysis settings as in Section 3.1 and 3.2 on English-Japanese (En-Jp) translation dataset. More details about this dataset and model descriptions can be found in supplementary materials (section A.1 and A.2).

Table 4 shows the BLEU score on the En-Jp test set. It can be observed that regardless of decoding direction (i.e., from left-to-right or from right-to-left) and with or without teacher forcing, the accuracy of the right half is always higher than that in the left half. This observation on Japanese is

| | left-to-right | | right-to-left | |
|-------|---------------|--------------|---------------|--------------|
| | 0 | 1 | 0 | 1 |
| left | 7.90 | 9.91 | 7.45 | 8.95 |
| right | 8.70 | 11.52 | 9.24 | 10.59 |

Table 4: BLEU scores on En-Jp test set. "0" represents the normal translation results, and "1" represents the teacher-forcing translation results.

opposite to English, German and Chinese in Section 3.1 and 3.2, and motivates us to investigate the differences between these languages.

We find that a linguistics concept, the branching, can differentiate Japanese from other languages such as English/German. Branching refers to the shape of the parse trees that represent the structure of sentences (Berg et al., 2011; Payne, 2006). Usually, right-branching sentences are head-initial, which means the main subject of the sentence is described first, and is followed by a sequence of modifiers that provide additional information about the subject. On the contrary, left-branching sentences are head-final that putting such modifiers to the left of the sentence (Payne, 2006).

English is a typical right-branching language, while Japanese is almost fully left-branching (Wikipedia, 2018). The two languages demonstrate the opposite phenomenon of accuracy drop as shown in previous studies. When we say a language is typical left/right-branching, we mean most of the sentences in this language follows the left/right-branching structure. While being predominantly right-branching, German is less conclusively so than English. Chinese features a mixture of head-final and head-initial structures, with the noun phrases are head-final while the strict head/complement ordering sentences are head-initial as right-branching (Wikipedia, 2018), but less conclusively than German.

We believe the language branching is a main cause of accuracy drop. Intuitively, the main subject of a right-branching sentence is described first (in the left part) and is followed by additional modifiers (in the right part) (Berg et al., 2011). Therefore, the left half of a right-branching sentence is more likely to possess a clearer structure pattern and thus lead to higher generation accuracy than in the right part, since the main subject is usually simpler and clearer than the modifiers that providing additional information about the subject.

t. In next section, we will verify this intuition this assumption from a statistical perspective.

5 Correlation between Language Branching and Accuracy Drop

As previous work (Arpit et al., 2017) shows, neural networks are easy to learn and memorize simple patterns but difficult to make a correct prediction on noise examples. In this section, we study different branching languages from two aspects, including the n-gram statistics of a target language, which has been used as a kind of characterization of hardness of learning (Bengio et al., 2009), and the dependency statistics in parse trees. We show that these statistics well correlate with the accuracy drop between the left half and the right half of translation results.

5.1 N-gram Statistics

Intuitively speaking, if a pattern occurs frequently and deterministically, it is easy to be learned by neural networks. By comparing the general statistics on the n-gram frequency and n-gram conditional probability of the left and right half tokens, we link the language branching to accuracy drop.

Denote a bilingual dataset $D = \{(x_i, y_i), i = 1, \dots, M\}$, where each y_i is a sequence of words $y_i = \{y_i^1, \dots, y_i^{T_i}\}$, T_i is the length of y_i . $F_{i,n}^l$ and $P_{i,n}^l$ denote the average n-gram frequency and n-gram conditional probability of the left half of y_i ³, i.e.,

$$F_{i,n}^l = \frac{1}{T_i/2 - n + 1} \sum_{j=1}^{T_i/2 - n + 1} F(y_i^j, \dots, y_i^{j+n-1}),$$

$$P_{i,n}^l = \frac{1}{T_i/2 - n + 1} \sum_{j=1}^{T_i/2 - n + 1} P(y_i^{j+n-1} | y_i^j, \dots, y_i^{j+n-2}),$$
(1)

where $F(\cdot)$ and $P(\cdot)$ are the n-gram frequency and n-gram conditional probability calculated from the training dataset. Similarly, $F_{i,n}^r$ and $P_{i,n}^r$ denote the n-gram frequency and n-gram conditional probability of the right half.

We calculate the average n-gram frequencies F_n^l and F_n^r of the left half and right half over all the target sentences in the training set. We also calculate the average n-gram conditional probabilities P_n^l and P_n^r over all the training sentences to compare the uncertainty of phrases in the left half and

³Again, we assume T_i is an even number. If not, we simply remove the middle word of y_i , as done in Section 3.1.

right half.

$$F_n^l = \frac{1}{M} \sum_{i=1}^M F_{i,n}^l, F_n^r = \frac{1}{M} \sum_{i=1}^M F_{i,n}^r, \quad (2)$$

$$P_n^l = \frac{1}{M} \sum_{i=1}^M P_{i,n}^l, P_n^r = \frac{1}{M} \sum_{i=1}^M P_{i,n}^r.$$

We also calculate the ratio of the sentences that the frequency/conditional probability of left half is bigger/smaller than that in the right half, denoted as $RF_n^{l>r}/RF_n^{l<r}$ and $RP_n^{l>r}/RP_n^{l<r}$:

$$RF_n^{l>r} = \frac{1}{M} \sum_{i=1}^M \mathbb{1}\{F_{i,n}^l > F_{i,n}^r\},$$

$$RF_n^{l<r} = \frac{1}{M} \sum_{i=1}^M \mathbb{1}\{F_{i,n}^l < F_{i,n}^r\}, \quad (3)$$

$$RP_n^{l>r} = \frac{1}{M} \sum_{i=1}^M \mathbb{1}\{P_{i,n}^l > P_{i,n}^r\},$$

$$RP_n^{l<r} = \frac{1}{M} \sum_{i=1}^M \mathbb{1}\{P_{i,n}^l < P_{i,n}^r\}.$$

We choose $n = 2$ and 3 to calculate the metrics in Equation 2 and 3 on different translation datasets. The numbers are listed in Table 5 and 6.

We can see the 2/3-gram frequency as well as the conditional probability of the left half is higher than that of the right half for right-branching languages including English, German and Chinese in De-En, En-De and En-Zh translation datasets. For left-branching language Japanese, the result is opposite. The n-gram frequency and conditional probability statistics are consistent with our observations on accuracy drop in left/right-branching languages and verify our hypothesis: right-branching languages have clearer patterns in left part (with larger n-gram frequency as well as the conditional probability) and consequently leads to higher translation accuracy in the left part than the right part; left-branching languages are opposite.


We further visualize how the accuracy drop (between the left half and right half of the translation-s) correlates with the gap of n-gram statistics in the left and right part. The accuracy drop (e.g., BLEU score) of left/right half is taken from the teacher-forcing with left-to-right decoding in Table 3, and the n-gram gap is taken from the Δ in the last row of Table 5 and 6. Figure 1 shows strong correlation between accuracy drop and the gap of n-gram statistics: As the gap of n-gram statistics increases from negative values to positive values, the accuracy drop also increases from negative to positive.

| | De-En | | En-De | |
|--------------|----------------|---------------|----------------|---------------|
| | 2-gram | 3-gram | 2-gram | 3-gram |
| F_n^l | 5713.8 | 3122.7 | 13811.8 | 687.1 |
| F_n^r | 3026.5 | 1377.6 | 11692.2 | 419.9 |
| $RF_n^{l>r}$ | 59.6% | 55.8% | 53.8% | 53.6% |
| $RF_n^{l<r}$ | 38.8% | 37.6% | 46.0% | 45.0% |
| Δ | 20.8% | 18.2% | 7.8% | 8.6% |
| | En-Zh | | En-Jp | |
| | 2-gram | 3-gram | 2-gram | 3-gram |
| F_n^l | 17707.0 | 1954.1 | 18910.0 | 1350.0 |
| F_n^r | 16256.4 | 1250.5 | 21076.7 | 1754.0 |
| $RF_n^{l>r}$ | 51.9% | 50.2% | 41.2% | 38.0% |
| $RF_n^{l<r}$ | 46.7% | 43.9% | 51.7% | 52.3% |
| Δ | 5.2% | 6.3% | -10.5% | -14.3% |


Table 5: The n-gram frequency statistics on different translation datasets. F_n^l and F_n^r represent the average of n-gram frequency of left and right half of target sentences. $RF_n^{l>r}$ and $RF_n^{l<r}$ represent the ratio that the n-gram frequency of left half of sentences are bigger/smaller than that of the right half. $\Delta = RF_n^{l>r} - RF_n^{l<r}$. Note that the sum of $RF_n^{l>r}$ and $RF_n^{l<r}$ is less than 1 since sentence with less than 4 words does not contribute to the n-gram statistics.

| | De-En | | En-De | |
|--------------|--------------|--------------|--------------|--------------|
| | 2-gram | 3-gram | 2-gram | 3-gram |
| P_n^l | 0.137 | 0.181 | 0.082 | 0.155 |
| P_n^r | 0.092 | 0.116 | 0.080 | 0.148 |
| $RP_n^{l>r}$ | 59.8% | 56.6% | 50.6% | 51.7% |
| $RP_n^{l<r}$ | 38.7% | 36.4% | 49.2% | 47.0% |
| Δ | 21.2% | 20.2% | 1.4% | 4.7% |
| | En-Zh | | En-Jp | |
| | 2-gram | 3-gram | 2-gram | 3-gram |
| P_n^l | 0.064 | 0.113 | 0.082 | 0.171 |
| P_n^r | 0.055 | 0.108 | 0.086 | 0.191 |
| $RP_n^{l>r}$ | 52.1% | 47.8% | 43.9% | 39.4% |
| $RP_n^{l<r}$ | 46.6% | 47.0% | 49.2% | 50.9% |
| Δ | 5.5% | 0.8% | -5.3% | -11.5% |

Table 6: The n-gram conditional probability statistics on different translation datasets. P_n^l and P_n^r represent the average n-gram conditional probability of left and right half of target sentences. $RP_n^{l>r}$ and $RP_n^{l<r}$ represent the ratio that the n-gram frequency of left half are bigger/smaller than that of the right half. $\Delta = RP_n^{l>r} - RP_n^{l<r}$. Note that the sum of $RP_n^{l>r}$ and $RP_n^{l<r}$ is less than 1 due to two reasons: (1) sentence with less than 4 words does not contribute to the statistics, and (2) we remove the n-gram condition probability with the denominator less than 100 to make probability calculation robust.



(a) Accuracy drop v.s 3-gram frequency gap (%).



(b) Accuracy drop v.s 3-gram conditional probability gap (%).

Figure 1: Accuracy drop (the gap between the left/right BLEU score) with respect to the ΔRF_3 and ΔRP_3 from Table 5 and 6 in the four translation tasks. The x-axis ΔRF_3 and ΔRP_3 represent the gap of between the left and right ratio of the 3-gram frequency/conditional probability defined in Table 5 and 6. The y-axis represents the accuracy drop in terms of BLEU score calculated by the teacher forcing decoding.

5.2 Dependency Statistics

In this subsection, we study language branching from the perspective of dependency structure. We hypothesize that if the left/right half of sentence contains more dependencies between its intra words, this half should be easier to predict, leading to higher accuracy. Here we analyze the English sentence in De-En translation and Japanese sentence in En-Jp translation, since English is fully right-branching and Japanese is fully left-branching as introduced before.

For English parsing, we utilize the well-acknowledged Stanford Parser⁴ to parse the sentences. After obtaining the parsing results, we split the sentence into left and right half, and separately count the numbers of dependencies in each half⁵. For Japanese, we leverage the open-source toolkit J.DepP⁶ to parse the sentence, and then count the number of dependencies of each half.

We provide the results in Table 7. As can be observed, for English sentences, the left-half words depend more on each other than the right-half words, while for the Japanese sentences, the right-half words have more dependencies. This observation is consistent with our observations on accu-

⁴<https://nlp.stanford.edu/software/lex-parser.shtml>

⁵For simplicity, we just count the number of dependency, without considering dependency types. The detailed parsing formats can be found in the supplementary material (Section A.3).

⁶<http://www.tkl.iis.u-tokyo.ac.jp/~ynaga/jdepp/>

| | English | Japanese |
|-------|--------------|----------------|
| Left | 40242 | 921735 |
| Right | 31509 | 1570630 |

Table 7: Number of dependencies in left and right half of English (De-En) and Japanese (En-Jp) training corpus. The number varies a lot since the two training corpus have different training sentences.

racy drop, and can well explain the high accuracy of left part in English translation and right part in Japanese translation.

6 Extended Analyses and Discussions

We have analyzed the accuracy drop problem from the view of error propagation and language itself in previous sections. In this section, we further provide extended analyses and several discussions to give a more clear understanding of the accuracy drop problem.

6.1 More Languages on Left-Branching

The previous analyses are based on four languages, three right-branching (En, De, Zh) and one left-branching language (Jp). To avoid the experimental bias and randomness, we provide one more translation task, English-Turkish (En-Tr) translation⁷, as Turkish is a left-branching language. We simply calculate the BLEU score of the left/right half in left-to-right and right-to-left

⁷The detailed dataset and model description can be found in supplementary material (section A.1 and section A.2).

| | 0 | 1 |
|-------|-------------|-------------|
| Left | 5.83 | 7.44 |
| Right | 5.27 | 7.96 |

Table 8: BLEU scores on En-Tr test set with left-to-right generation. Normal translation is denoted as “0”, and teacher-forcing translation is denoted as “1”.

| | left-to-right | right-to-left |
|-------|---------------|---------------|
| Full | 27.63 | 25.44 |
| Left | 9.17 | 8.37 |
| Right | 7.51 | 7.25 |

Table 9: BLEU scores on the left-to-right and right-to-left translation sentences on the De-En test set, with RNN-based model. “Full” means the BLEU score of the whole translation sentence.

decodings, as in Section 3.1 and 3.2.

The result is provided in Table 8. For the left-to-right decoding, the accuracy of the left half is higher than that of the right half in the normal translation. However, the accuracy of the right half becomes higher with teacher forcing translation. This demonstrates that English-Turkish translation performs similar to English-Japanese translation as the accuracy of right half is higher than that of the left half. But different from what we observed in Japanese, Turkish shows the opposite phenomenon: the influence of language branching is weaker than error propagation.

6.2 Other Model Structures

One may wonder whether the results in the paper are biased towards a certain model structure as we use Transformer on all the above analyses. To address such concerns, we conduct an additional experiment on De-En translation task with RNN (GRU)-based model⁸. The results are shown in Table 9 and the observations are consistent with what we observed on Transformer. The accuracy of the left half of the De-En translation sentence is always higher than the right half, in both the left-to-right and right-to-left decodings.

6.3 Other Sequence Generation Tasks

We conduct experimental analysis on abstractive summarization, which is also a sequence generation task. The goal of the task is to recap a long news sentence into a short summary. We use Gigaword dataset which contains 3.8M training pairs,

⁸The detailed setting for GRU based RNN model can be found in supplementary material (section A.2).

| | left-to-right | | |
|-------|---------------|-------------|--------------|
| | ROUGE-1 | ROUGE-2 | ROUGE-L |
| Full | 35.55 | 16.66 | 33.01 |
| Left | 24.44 | 9.87 | 23.34 |
| Right | 21.31 | 8.32 | 20.38 |
| | right-to-left | | |
| | ROUGE-1 | ROUGE-2 | ROUGE-L |
| Full | 35.22 | 16.55 | 32.59 |
| Right | 21.62 | 8.41 | 20.48 |
| Left | 23.60 | 9.54 | 22.52 |

Table 10: ROUGE F1 scores for left-to-right and right-to-left generated translation sentences in abstractive summarization task. ROUGE-N stands for N-gram based ROUGE F1 score, ROUGE-L stands for longest common subsequence based ROUGE F1 score. “Full” means the entire translation sentence.

190k validation and 2k test pairs of English sentence, and train an RNN-based model for sentence summarization. The accuracy is measured by the commonly used metric ROUGE F1 score and are reported in Table 10.

We observe the same phenomenon as in translation tasks. The accuracy of the left half is always better than the right half, no matter in left-to-right or right-to-left decoding, since the target language English is a right-branching language.

7 Conclusion

In this work, we studied the problem of accuracy drop between the left half and the right half of the results generated by neural machine translation models. We found the influence of error propagation is overstated in literature and error propagation alone cannot explain accuracy drop. We showed that language branching well correlates to the accuracy drop problem and the evidences on n-gram statistics as well as the dependency statistics well support this correlation. Our discoveries suggest that left-to-right NMT models fit better for right-branching languages (e.g., English) and right-to-left NMT models fit better for left-branching languages (e.g., Japanese).

For future works, we will study more left/right-branching languages as well as other languages that have no obvious branching characteristics. We will also investigate how language branching influences other natural language tasks, especially for neural networks based models.

References

- Devansh Arpit, Stanislaw K. Jastrzebski, Nicolas Balas, David Krueger, Emmanuel Bengio, Maxinder S. Kanwal, Tegan Maharaj, Asja Fischer, Aaron C. Courville, Yoshua Bengio, and Simon Lacoste-Julien. 2017. A closer look at memorization in deep networks. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 233–242. PMLR.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 1171–1179.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48. ACM.
- Thomas Berg et al. 2011. *Structure in language: A dynamic perspective*. Routledge.
- Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar; A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1724–1734.
- Fei Gao, Lijun Wu, Li Zhao, Tao Qin, Xueqi Cheng, and Tie-Yan Liu. 2018. Efficient sequence learning with group recurrent networks. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, volume 1, pages 799–808.
- 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. *CoRR*, abs/1803.05567.
- Di He, Hanqing Lu, Yingce Xia, Tao Qin, Liwei Wang, and Tiejian Liu. 2017. Decoding with value networks for neural machine translation. In *Advances in Neural Information Processing Systems*, pages 178–187.
- Di He, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tiejian Liu, and Wei-Ying Ma. 2016. Dual learning for machine translation. In *Advances in Neural Information Processing Systems*, pages 820–828.
- Lemao Liu, Andrew M. Finch, Masao Utiyama, and Eiichiro Sumita. 2016a. Agreement on target-bidirectional lstms for sequence-to-sequence learning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 2630–2637.
- Lemao Liu, Masao Utiyama, Andrew M. Finch, and Eiichiro Sumita. 2016b. Agreement on target-bidirectional neural machine translation. In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 411–416.
- 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, July 6-12, 2002, Philadelphia, PA, USA.*, pages 311–318.
- Thomas Payne. 2006. *Exploring language structure: a student’s guide*. Cambridge University Press.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. *CoRR*, abs/1511.06732.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Edinburgh neural machine translation systems for WMT 16. In *Proceedings of the First Conference on Machine Translation, WMT 2016, colocated with ACL 2016, August 11-12, Berlin, Germany*, pages 371–376.
- Shiqi Shen, Yong Cheng, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2016. Minimum risk training for neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*.
- Yanyao Shen, Xu Tan, Di He, Tao Qin, and Tie-Yan Liu. 2018. Dense information flow for neural machine translation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 1294–1303.
- Kaitao Song, Xu Tan, Di He, Jianfeng Lu, Tao Qin, and Tie-Yan Liu. 2018. Double path networks for sequence to sequence learning. In *Proceedings of*

- the 27th International Conference on Computational Linguistics, COLING 2018, Santa Fe, New Mexico, USA, August 20-26, 2018, pages 3064–3074.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 3104–3112.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 6000–6010.
- Yijun Wang, Yingce Xia, Li Zhao, Jiang Bian, Tao Qin, Guiquan Liu, and Liu. 2018. Dual transfer learning for neural machine translation with marginal distribution regularization. In *AAAI*.
- Yuguang Wang, Shanbo Cheng, Liyang Jiang, Jiajun Yang, Wei Chen, Muze Li, Lin Shi, Yanfeng Wang, and Hongtao Yang. 2017. Sogou neural machine translation systems for WMT17. In *Proceedings of the Second Conference on Machine Translation, WMT 2017, Copenhagen, Denmark, September 7-8, 2017*, pages 410–415.
- Taro Watanabe and Eiichiro Sumita. 2002. Bidirectional decoding for statistical machine translation. In *19th International Conference on Computational Linguistics, COLING 2002, Howard International House and Academia Sinica, Taipei, Taiwan, August 24 - September 1, 2002*.
- Wikipedia. 2018. Head-directionality parameter — Wikipedia, the free encyclopedia. <http://en.wikipedia.org/w/index.php?title=Head-directionality%20parameter&oldid=835385416>. [Online; accessed 08-May-2018].
- Ronald J. Williams and David Zipser. 1989. A learning algorithm for continually running fully recurrent neural networks. *Neural Computation*, 1(2):270–280.
- Lijun Wu, Fei Tian, Tao Qin, Jianhuang Lai, and Tie-Yan Liu. 2018. A study of reinforcement learning for neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Lijun Wu, Yingce Xia, Li Zhao, Fei Tian, Tao Qin, Jianhuang Lai, and Tie-Yan Liu. 2017. Adversarial neural machine translation. *arXiv preprint arXiv:1704.06933*.
- Yingce Xia, Xu Tan, Fei Tian, Tao Qin, Nenghai Yu, and Tie-Yan Liu. 2018. Model-level dual learning. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, pages 5379–5388.
- Xiangwen Zhang, Jinsong Su, Yue Qin, Yang Liu, Ron-grong Ji, and Hongji Wang. 2018. Asynchronous bidirectional decoding for neural machine translation. *CoRR*, abs/1801.05122.