

Multi-Sentence Resampling: A Simple Approach to Alleviate Dataset Length Bias and Beam-Search Degradation

Ivan Provilkov

Yandex Research, Moscow
Moscow Institute of Physics and Technology
iv-provilkov@yandex-team.ru

Andrey Malinin

Yandex Research, Moscow
HSE University, Moscow
am969@yandex-team.ru

Abstract

Neural Machine Translation (NMT) is known to suffer from a beam-search problem: after a certain point, increasing beam size causes an overall drop in translation quality. This effect is especially pronounced for long sentences. While much work was done analyzing this phenomenon, primarily for autoregressive NMT models, there is still no consensus on its underlying cause. In this work, we analyze errors that cause major quality degradation with large beams in NMT and Automatic Speech Recognition (ASR). We show that a factor that strongly contributes to the quality degradation with large beams is *dataset length-bias - NMT datasets are strongly biased towards short sentences*. To mitigate this issue, we propose a new data augmentation technique – *Multi-Sentence Resampling (MSR)*. This technique extends the training examples by concatenating several sentences from the original dataset to make a long training example. We demonstrate that MSR significantly reduces degradation with growing beam size and improves final translation quality on the IWSTL15 En-Vi, IWSTL17 En-Fr, and WMT14 En-De datasets.

1 Introduction

In this work, we address the *beam-search problem* in Neural Machine Translation (Koehn and Knowles, 2017). Beam Search is the standard hypothesis search method for autoregressive sequence generation. Large beams provide more probable hypotheses than small beams; however, the overall translation quality drops with growing beam size after a certain point. This effect is especially strong for long sentences, connecting with the fact that NMT models are biased to giving high probabilities to short hypotheses. Stahlberg and Byrne (2019) showed that exact search by likelihood for neural machine translation finds empty string as the optimal hypothesis in more than 50% of cases.

One of the most famous methods to mitigate quality degradation with growing beam size is

length normalization (Bahdanau et al., 2016; Wu et al., 2016). This technique normalizes log-likelihoods of a hypothesis in beam search by its length, thus promoting long hypotheses. Other methods examine adding a reward for each token’s score during the decoding process (Yang et al., 2018; Murray and Chiang, 2018).

While the beam-search problem has been extensively studied (Sountsov and Sarawagi, 2016; Murray and Chiang, 2018; Kumar and Sarawagi, 2019; Meister et al., 2020; Eikema and Aziz, 2020; Yang et al., 2020; Wang and Sennrich, 2020) there is still no consensus on the underlying reason for such model behavior. Furthermore, prior work has investigated this problem primarily for NMT models, giving little attention to other domains that are also known to suffer from it, such as Automatic Speech Recognition (ASR) (Chorowski and Jaitly, 2017; Zhou et al., 2020). Murray and Chiang (2018) noticed that since in each step of beam-search generation, negative log-probability is added to the hypothesis’ score, if a model overestimates the probability of an already generated sequence of tokens, there is no way to downgrade this probability afterward. Consequently, models are biased towards finalizing a short hypothesis by generating an end-of-sequence token (EOS) rather than generating a long continuation. Their experiments show a connection between quality degradation and decreasing length of hypotheses with growing beam size. Our work is in agreement with their explanation. Moreover, we show that the main quality degradation with large beams in NMT and ASR comes from short translations obtained by early termination of long hypotheses from small beams.

Our work examines how the distribution of sentence lengths in a dataset affects the beam-search problem. We demonstrate that the *beam-search problem is strongly connected with the distribution of sentence lengths in training datasets*. Specifically, we show that common NMT datasets, such as

IWSLT and WMT, exhibit a strongly skewed distribution of sentence lengths, with the mode focused on short sentences. NMT models learn biased probability distributions and fail on long sentences during inference. In contrast, for ASR models trained on Librispeech (Panayotov et al., 2015), where the distribution of sentence lengths is more symmetrical and biased towards longer sentences, the beam-search degradation occurs at much larger lengths. Based on our findings, we propose a simple and effective dataset augmentation technique that makes training examples longer – *Multi-Sentence Resampling*. It creates a new dataset where each training sample can be a concatenation of multiple sentences. Our method alleviates quality degradation with growing beam size and increases the final quality of the model.

The key contributions of our work are as follows:

- We show that quality degradation with growing beam size comes mostly from short translations, which are early finalized prefixes of long hypotheses;
- We show that training datasets that are biased towards short sentences strongly contribute to the beam-search problem;
- We introduce *Multi-Sentence Resampling* – a simple and effective dataset augmentation technique that alleviates beam search problem and increases the final translation quality¹.

2 Quality degradation analysis

This section analyzes quality degradation with the growing beam size of two systems: Neural Machine Translation (NMT) and Automatic Speech Recognition (ASR). ASR is also known to suffer from the beam-search problem (Zhou et al., 2020; Chorowski and Jaitly, 2017). The models, training, and evaluation processes of NMT and ASR models in our work are almost identical. However, the ASR dataset (Librispeech) has some properties that are differ from the Machine Translation setting: the average length of target sentences in the training dataset is much larger than in the test, which is why we chose this task for comparison.

2.1 Experimental setup

In order to make an informative comparative analysis of beam-search quality degradation between

¹Our code is available at <https://github.com/yandex-research/msr>

NMT and ASR, we aimed to make the model and experimental setup for both tasks as similar as possible to minimize mismatch. Specifically, the vocabulary, pre-processing, and models (except the first two layers of the ASR encoder) are identical between the two tasks.

Datasets and preprocessing

We use IWSLT2017 Fr-En, IWSLT2015 En-Vi, WMT2014 En-De, and Librispeech (Panayotov et al., 2015) datasets. The information about them is summarized in Table 1. We use standard validation and test splits for WMT, for the IWSLT En-Vi pair we used test 2012 for validation and test 2013 as a test set, for the IWSLT En-Fr pair we used development set 2010 for validation and test 2015 as a test set. The bulk of the analysis is done on IWSLT17 Fr-En and LibriSpeech-clean, as they are similar in the number of target-side tokens. As there is no information about the case in Librispeech, we converted Librispeech and IWSLT2017 Fr-En to lowercase to have similar conditions for these datasets. For all other datasets the casing is unchanged. We preprocess all datasets with the Moses toolkit³, and use BPE (Sennrich et al., 2016) with vocabulary size 32k for WMT and 5k for other datasets, as small vocabularies are beneficial for small datasets (Ding et al., 2019).

Model and Optimization

For NMT, we use Transformer base (Vaswani et al., 2017) model from fairseq (Ott et al., 2019). For IWSLT, we use the batch size of 8k tokens and dropout coefficient 0.2; all other parameters are kept as in (Vaswani et al., 2017). Models are trained until convergence on a validation dataset.

For ASR, we used Transformer-base with two additional convolutional layers in the encoder, as suggested in (Wang et al., 2020a), all parameters for ASR are kept as in the original paper.

Inference and Evaluation

To produce length-normalized hypotheses, we use standard beam-search from fairseq (Ott et al., 2019). For evaluation, we averaged the last 5 checkpoints and use BLEU (Papineni et al., 2002) computed via Sacrebleu (Post, 2018).

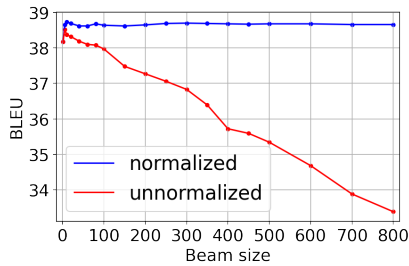
For evaluating the ASR system, we used word-error-rate (WER) (Marzal and Vidal, 1993) – a standard metric that shows edit distance from generated sequence to reference.

³<https://github.com/amosm/amos-smt-mosesdecoder>

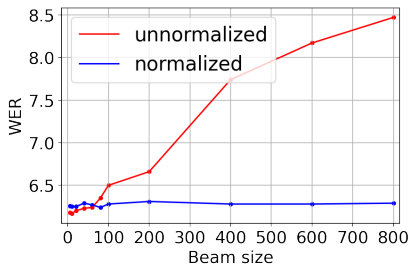
		Number of sentences (train/dev/test)	Number of tokens in train (en)	Avg. sentence length in tokens
Librispeech	clean	133k / 2703 / 2620	4.6M	34.9
IWSLT17	En ↔ Fr	232k / 890 / 1210	4.8M	20.5
IWSLT15	En ↔ Vi	133k / 1553 / 1268	2.7M	20.3
WMT14	En ↔ De	4.5M / 3000 / 3003	11.4M	25.2

Table 1: Overview of the datasets. The number of tokens is calculated after Moses preprocessing².

2.2 Degradation analysis



(a) IWSLT17 Fr-En



(b) Librispeech

Figure 1: Comparison of BLEU and WER scores of hypotheses normalized and not normalized by length for different beam sizes. The Librispeech model was trained and tested on the 'clean' subset.

In this section, we analyse which errors contribute to quality degradation with growing beam size in ASR and NMT and provide additional evidence to connect the beam-search quality degradation with the hypotheses shortening on large beams.

Here, and later in this work, we abbreviate models with length-normalized beam search as *normalized*, while models without length-normalization as *unnormalized*. Figure 1 shows quality of IWSLT Fr-En and Librispeech models with beam size growing from 1 to 800 on test sets. Normalized models do not show quality degradation with growing beam size. However, without length normalization, quality drops significantly with increasing beam-size.

To show which test samples cause a drop in quality, we divide the test sets into several categories,

based on hypotheses from beam size 5 and beam size 400. These categories are following:

- $b_{400} \geq b_5$ – sentences on which sentence-level BLEU of a top hypothesis from beam size 400 is greater or equal than sentence-level BLEU of a correspondent hypothesis from beam size 5. In other words, all cases where quality improved or didn't change with the large beam size.
- $b_{400} \lesssim b_5$ – sentences where best hypothesis from beam 400 is a prefix of a corresponding best hypothesis from beam 5 (except EOS token and "." before EOS). An example of this category is a pair of hypotheses: "I can" from beam size 400, and "I can do this tomorrow if you wait." from beam size 5, the first is a prefix of the second;
- $b_{400} < b_5$ – all other cases that are not in the first 2 categories. In other words, examples where quality drops, however top hypothesis from beam 400 is not a prefix of a correspondent top hypothesis from beam 5.

Dataset	IWSLT17		Librispeech	
Subset	unnorm.	norm.	unnorm.	norm.
$b_{400} \geq b_5$	90	95.6	97.6	98.7
$b_{400} \lesssim b_5$	3	0	0.9	0
$b_{400} < b_5$	7	4.4	1.5	1.3

Table 2: Distribution of cases (%) according to defined categories between beam 5 and beam 400. Norm. means length-normalized versions of translations.

Table 2 shows how hypotheses are distributed among categories. The smallest is the category with prefixes – " $b_{400} \lesssim b_5$ ". Such examples are only 3% of cases in IWSLT and nearly 1% in Librispeech in unnormalized versions. This is significantly less than the category " $b_{400} < b_5$ " which represents all other cases where quality drops.

Subset	beam 5	beam 400	Contribution to degradation
IWSLT Fr → En, BLEU			
b400 ≥ b5	38.83	39.29	+0.41
b400 ≤ b5	38.71	0.13	-1.16
b400 < b5	36.70	30.75	-0.42
Librispeech, WER			
b400 ≥ b5	5.5	5.4	-0.1
b400 ≤ b5	22.9	86.5	+0.57
b400 < b5	17.6	26.6	+0.13

Table 3: BLEU scores for IWSLT17 and WER scores for Librispeech corresponding to the different categories. Column "Contribution to degradation" represents the difference between beam 400 and beam 5 scores, weighted by the percentage of the dataset in the corresponding category from Table 2.

Although examples where the EOS token appeared too early during the generation of a reasonable, long hypothesis, are *smallest* category, they have the *greatest* contribution to the overall quality degradation with growing beam size. Consider Table 3, which shows quality in terms of BLEU/WER for different categories and beam sizes. The biggest drop in quality is in the prefixes category. It drops from 38.71 BLEU to almost 0 for IWSLT. For Librispeech, WER increases from 22.9 to 86.5 in the same category. Performance within the category "b400 < b5" degrades more modestly, losing 5 BLEU and gaining 3.52 WER, respectively. Weighed by the fraction of each category within the datasets, the prefixes category contributes nearly 3 times more than the non-prefix category to the *overall* BLEU on IWSLT (1.16 vs. 0.42), and nearly 4 times more to the *overall* WER on Librispeech (0.57 vs. 0.13)⁴.

Table 4 shows that the prefix category (early EOS) is also the most significant in terms of length reduction with growing beam size. Length for beam 400 in this category is nearly 84% lower compared to beam 5. Interestingly, early EOS occurs mainly in examples where the top hypothesis from beam 5 is long, on average 53.8 tokens in IWSLT and 67.17 in Librispeech, which is much longer than average lengths over the whole test datasets, 24.53 and 26.73 respectively. This observation adds further evidence to the connection between hypotheses shortening and the quality degradation with growing beam size.

⁴This is an approximate analysis, as BLEU is a non-linear corpus-level statistic.

Subset	beam 5	beam 400	Contribution to shortening
IWSLT Fr → En			
b400 ≥ b5	22.02	21.69	-0.3
b400 ≤ b5	53.8	8.72	-1.35
b400 < b5	44.2	39.3	-0.34
whole test	24.53	22.54	-1.99
Librispeech			
b400 ≥ b5	26.12	26.11	-0.01
b400 ≤ b5	67.17	10.65	-0.51
b400 < b5	42.67	39.15	-0.05
whole test	26.73	26.17	-0.56

Table 4: Average token lengths of best hypotheses from different beam sizes and categories. Column "Contribution to shortening" represents the difference between columns "beam 400" and "beam 5" weighted by the fraction of the corresponding category in the dataset.

Our findings relate to work studying calibration⁵ problems of NMT, which show that NMT architectures are poorly calibrated, especially the EOS token (Kumar and Sarawagi, 2019; Wang et al., 2020b).

2.3 Dataset length-bias

Having found further evidence to link length bias with the beam search problem, we examine and compare the distribution of sentence lengths in typical NMT and ASR datasets.

Consider Figure 2, which shows the distribution of sentence lengths in IWSLT17 Fr-En, IWSLT15 En-Vi, WMT14 De-En, and Librispeech datasets. The NMT datasets have an average sentence length between 20 and 25 tokens and exhibit a strong, asymmetric skew towards short sentences. In contrast, the Librispeech *training* dataset exhibits a more symmetric distribution of sentence length, with an average length of more than 40. At the same time, the distribution of lengths in the test and validation sets is similar to NMT. As a result, *during training on Librispeech, the model sees a far larger and diverse set of long sentences than is encountered in evaluation.*

Let's investigate how quality relates to the length of the target sentences. Figure 3 shows BLEU/WER scores on the test sets for buckets based on the reference length. If we compare Figure 2 with Figure 3 we will note an interesting feature: quality degradation on unnormalized beam 400 in machine translation tasks starts after length around 30 and mostly monotonically in-

⁵Calibration measures how well the model's probability distribution matches the actual data distribution.

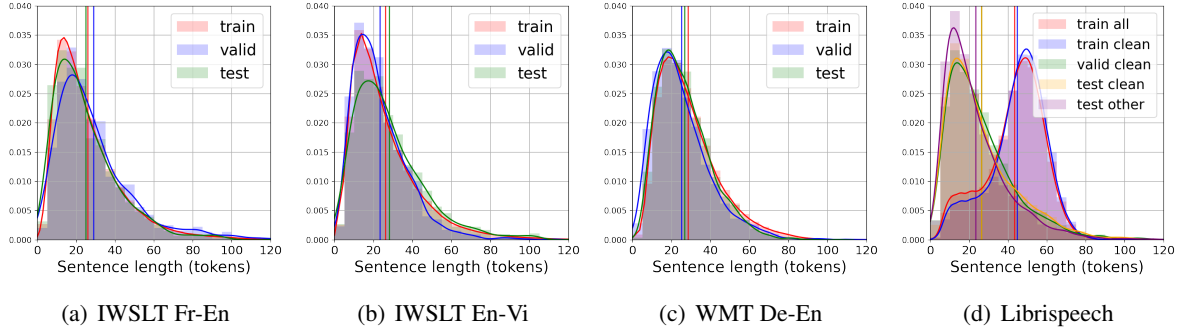


Figure 2: Distribution of target side sentence lengths after Moses and BPE for different datasets. Vertical lines represent mean of the corresponding distribution.

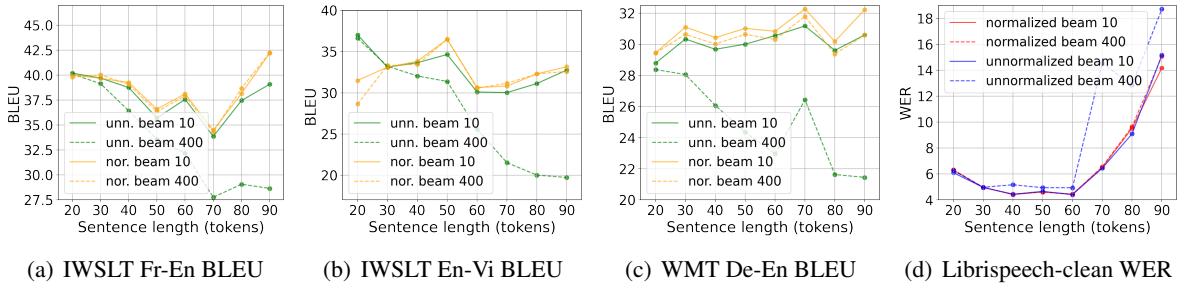


Figure 3: BLEU scores for buckets with growing sentence length. Each point represents quality calculated on sentences with lengths from the previous point to the current position, starting from 0.

creases as we go to longer sentences. In contrast, on Librispeech, the quality starts to drop only after the length 60, which correlates with the distribution of lengths of train examples. Specifically, in IWSLT17 Fr-En dataset, only 30% of training sentences have lengths greater than 30 and less than 10% are longer than 50 tokens. In contrast, Librispeech has many training sentences with a length of 60 tokens or less, and their amount drops rapidly only after this value, with about 5% of sentences having a length greater than 65 tokens.

Thus, we can clearly see that beam search quality degrades when operating on sentences, which are underrepresented in the training datasets in terms of *reference* length. This brings us to one of the main ideas of our work: *training datasets biased towards short sentences strongly contribute to the quality degradation with growing beam size*. In typical training datasets in Machine Translation (IWSLT and WMT) long sentences are significantly underrepresented, causing models to overfit to shorter sentences and overestimate probabilities of short prefixes. In the next section, we propose a data-augmentation strategy to alleviate this issue.

Algorithm 1: Multi-Sentence Resampling

```

 $D \leftarrow$  old train dataset with pairs of sentences;
 $N \leftarrow$  maximum number of sentences in one example in a new train dataset;
 $S = |D| \cdot M \leftarrow$  number of examples in the new train dataset;
 $R \leftarrow \emptyset$  – new dataset;
for  $i$  in  $1..S$  do
     $n \leftarrow$  random integer from 1 to  $N$ 
    new_source = ""
    new_target = ""
    for  $k$  in  $1..n$  do
        (cur_source, cur_target) = sample random example from  $D$ 
        new_source += cur_source
        new_target += cur_target
    end
     $R.append((new\_source, new\_target))$ 
end
return  $R$ ;

```

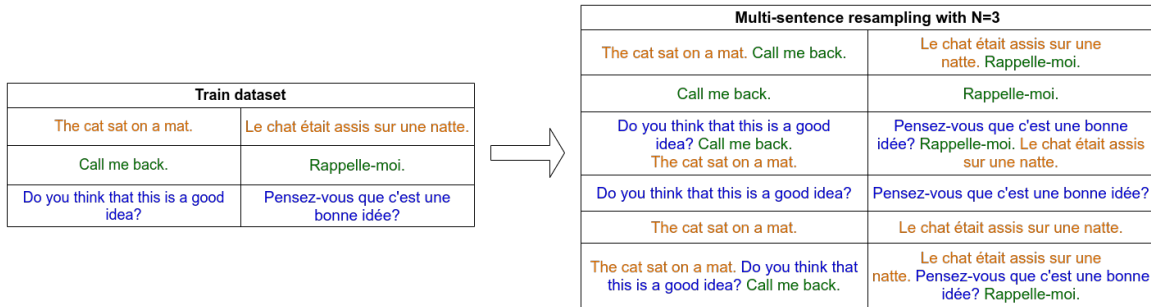


Figure 4: Example of Multi-Sentence Resampling with N=3. In the new train dataset with equal probability there are examples with 1,2 and 3 sentences from the original dataset.

3 Multi-Sentence Resampling

In this section, we introduce *Multi-Sentence Resampling* (MSR) – a simple data augmentation method that alleviates the beam search problem by addressing dataset length bias and which increases the overall quality of translation models. Specifically, MSR augments a dataset such that instead of one sentence, each training example consists of 1 to N sentences, randomly chosen from a dataset and concatenated one after another. It preserves the order of sentences: the source side is concatenated to the source side, and the target side is concatenated to the target side of a new train example. For each new training example, MSR randomly chooses from 1 to N sentences, so that the model does not overfit to the particular number of sentences.

Figure 4 illustrates the algorithm. The original training dataset, with 3 examples, is on the left, and the dataset created by *Multi-Sentence Resampling* is shown on the right. The new dataset contains 6 train examples: 2 examples with one sentence, 2 examples with a concatenation of two original sentences, and 2 examples with a concatenation of 3 original sentences. The full procedure is described in Algorithm 1. The algorithm takes an original train dataset, desired number of examples in the augmented dataset S , and the maximum number of sentences in an example N . *Multi-Sentence Resampling* increases the average length of the dataset and provides a diverse set of long training examples. In contrast to other methods that use rescoring of hypotheses and per-token rewards (Yang et al., 2018) or predict target length separately (Yang et al., 2020), our method does not change the search procedure.

Figure 5 illustrates how train examples length distribution changes in IWSLT17 Fr-En dataset for N from 2 to 5. With growing N distributions be-

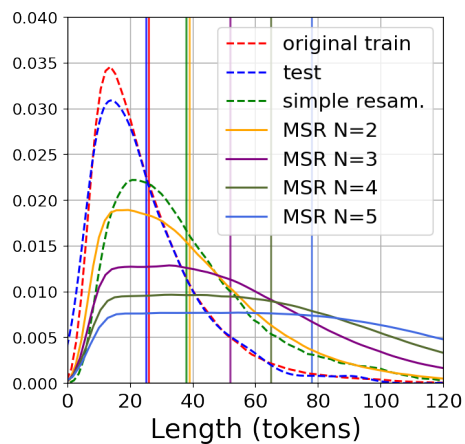


Figure 5: Distribution of sentence lengths for different N in MSR for IWSLT17 Fr-En dataset (en part). Vertical lines represent mean of the corresponding distribution.

come more flatten for lengths presented in the test set. The average length of examples in a new train dataset can be approximately calculated as

$$new_length \simeq \sum_{n=1}^N \frac{L \cdot n}{N} = L \cdot \frac{N+1}{2},$$

where L is the average length of the original dataset.

4 Experiments

In this section we provide an empirical evaluation of the *Multi-Sentence Resampling* on the IWSLT and WMT datasets.

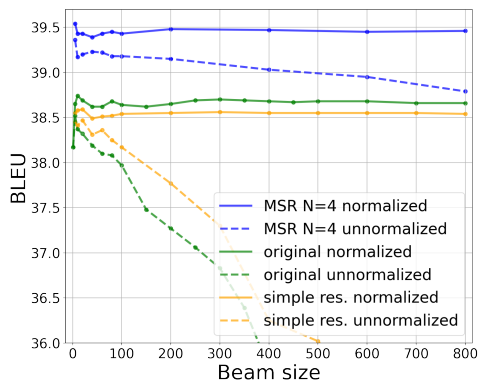
4.1 Experimental setup

To compare with the original Transformer paper (Vaswani et al., 2017), where authors used another beam-search parameters and BLEU computation, we added additional part to Table 5. For

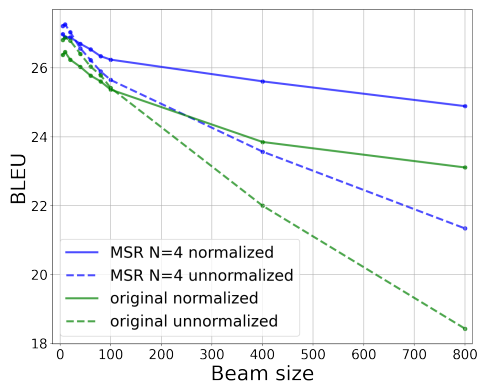
this part we changed length penalty to 0.6 in beam-search and compute BLEU as in (Vaswani et al., 2017)⁶. As baselines we use standard models trained without data augmentation.

In our experiments, for IWSLT datasets $M = 10$ – the new train dataset is 10 times larger than the original one, for WMT2014 En-De $M = 5$, as this dataset is much bigger.

4.2 Quality with growing beam size



(a) IWSLT17 Fr-En



(b) WMT14 En-De

Figure 6: Quality with growing beam size of the baseline, MSR $N = 4$ and simple resampling.

Figure 6 compares quality degradation with growing beam size for the baseline and MSR with $N = 4$ for IWSLT17 Fr-En datasets and WMT14 En-De. As an additional baseline, we compare MSR with a simpler strategy on IWSLT – resampling the dataset multiple times so that the probability of a sentence is proportional to its length. This way, long sentences occur more frequently during training. There are several interesting points

⁶https://github.com/pytorch/fairseq/blob/master/scripts/compound_split_bleu.sh

in this comparison. Firstly, Multi-Sentence Resampling achieves significantly better quality than the baseline on both datasets. Secondly, while the baseline’s quality rapidly drops with the growing beam size, the quality of the *Multi-Sentence Resampling* drops much more slowly. In particular, MSR with $N = 4$ with beam size 800 has quality better than the baseline with *any* beam size on IWSLT. On WMT, improvements for large beam sizes are more modest, which is expected, as data augmentation has less effect on larger datasets. However, MSR works on par with the length-normalized baseline up to the beam size 400. Third, simple resampling in IWSLT works slightly better than the baseline in the unnormalized setting; however, it drops quality in the normalized case. The benefits of simple resampling are limited because, unlike with MSR, the set of long sentences severely lacks diversity, and the model overfits to it during training.

We analyze how the value of the hyperparameter N in MSR affects beam search quality in Figure 8, which shows the quality of trained models for different values of N across a range of beam sizes. On small beams, all N behave without significant difference. On large beams, quality grows with N up to 4 and decreases for larger N .

An additional analysis of how effects of the *Multi-Sentence Resampling* vary with reference length and number of concatenated sentences is provided in Figure 7. Baseline unnormalized beam-search with beam-width 400 works badly on long sentences: quality degradation increases starting from length 30. On the other hand, MSR with $N = 4$ has almost no degradation for long sentences. Additionally, experiments with $N = 2$ and $N = 3$ show that quality degradation on long sentences decreases as we increase N , likely because we start fitting the model to far longer sequence lengths than encountered in the test set.

Table 5 examines the effects of MSR on a range of translation tasks. All scores are computed with Sacrebleu and default beam search from fairseq (Ott et al., 2019), except "WMT14, eval as in (Vaswani et al., 2017)". In this table, all MSR experiments are conducted with $N=4$, which is a simple baseline to choose N . We can make the following observations. Firstly, on all datasets and translation directions, models trained with *Multi-Sentence Resampling* statistically significantly outperform baselines: from 0.42 to 0.76 BLEU on Fr-En, En-Fr, Vi-En and En-Vi, and nearly 0.3 BLEU

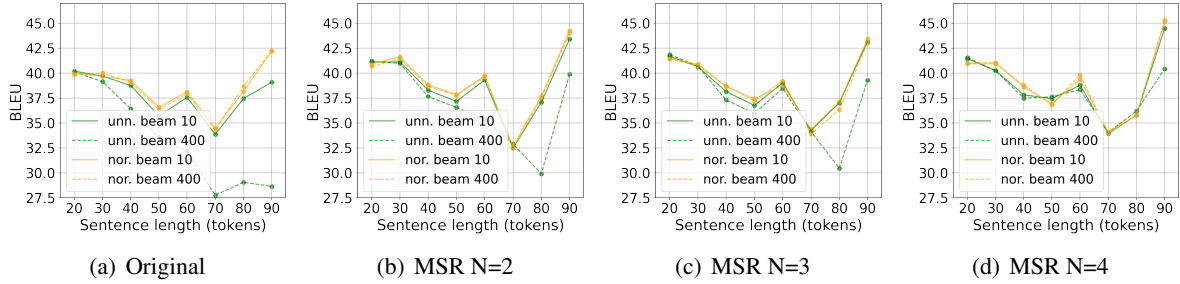
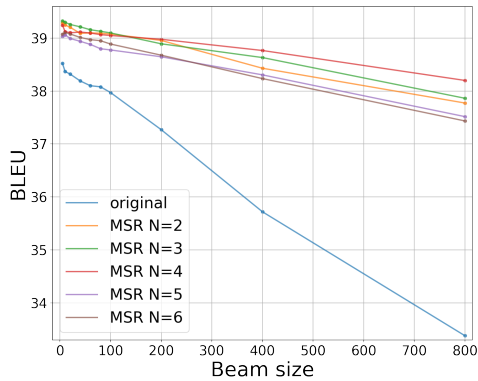


Figure 7: Changes of quality with growing sentence length for IWSLT17 Fr-En. Each point represents quality calculated on sentences with lengths from previous point to the current position, starting from 0.



(a) IWSLT unnormalized

Figure 8: BLEU with growing beam size for Multi-Sentence resampling with different N on IWSLT2017 Fr-En dataset. Each point is an average between 3 models trained with different random seeds.

	original		MSR	
	norm.	unn.	norm.	unn.
IWSLT17				
Fr-En	38.74	38.37	39.4	39.23
En-Fr	40.06	40.53	40.64	40.95
IWSLT15				
Vi-En	29.57	29.56	30.01	30.08
En-Vi	31.78	30.99	32.54	31.17
WMT14, sacrebleu				
En-De	26.61	27.01	26.92	27.29
De-En	30.96	30.15	31.25	30.59
WMT14, eval as in (Vaswani et al., 2017)				
En-De	27.37	27.48	27.71	27.67

Table 5: BLEU scores. Bold indicates the best score and all scores whose difference from the best is not statistically significant (with p -value of 0.05). (Statistical significance is computed via bootstrap (Koehn, 2004).)

for En-De and De-En. Secondly, length-normalized models work significantly better than models without length-normalization only in 2 directions out of 6. We did not tune length-normalization hyperparameters in our experiments; however, our results suggest that length-normalization may be unnecessary in some cases.

4.3 Training time

As with any regularization, Multi-Sentence Resampling increases the training time of models. Although we expanded the dataset by 10x and 5x times for IWSLT and WMT, respectively, the training time in both cases did not increase with the size of the dataset. Table 5 shows that MSR with $N = 4$ increases training time before convergence on average 80% among used datasets, compared to the default training. This suggests that it is possible to make a more memory efficient MSR implementation as part of the data processing pipeline which does MSR on-the-fly, without the need to pre-process a training dataset which is 5-10x larger than the original. However, we leave this to future work.

	original	MSR
IWSLT17		
Fr-En	60	118
En-Fr	77	195
IWSLT15		
Vi-En	55	75
En-Vi	53	85
WMT14		
En-De	150	320
De-En	174	250

Table 6: Number of thousands of training steps before convergence for experiments from Table 5.

5 Conclusions

In this work, we analyzed errors that cause quality degradation with growing beam size in NMT and ASR. We demonstrated that the major contribution to quality degradation on large beams comes from short translations, which are early terminated prefixes of hypotheses which are long when decoding with a small beams. In contrast to ASR, we showed that the reference length on which beam search degradation begins to grow is connected with the low number of sentences longer than this length during training. Thus, usual NMT datasets, that are biased towards short sentences, strongly contribute to the degradation with large beams. Based on this finding, we introduced *Multi-Sentence Resampling* – a simple data augmentation technique. It concatenates several sentences from a dataset, increasing the length of training examples. Models trained with *Multi-Sentence Resampling* were shown to consistently outperform baseline models on IWSLT15 En-Vi, IWSLT17 En-Fr, and WMT14 En-De datasets. Thus, we demonstrate that it is possible to mitigate beam search degradation with data augmentation. Future research directions include adapting *Multi-Sentence Resampling* to other domains like ASR and studying beam search problems for document-level Machine Translation, where adjacent sentences are naturally connected.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2016. [Neural machine translation by jointly learning to align and translate](#).
- Jan Chorowski and Navdeep Jaitly. 2017. [Towards better decoding and language model integration in sequence to sequence models](#).
- Shuoyang Ding, Adithya Renduchintala, and Kevin Duh. 2019. [A call for prudent choice of subword merge operations in neural machine translation](#). In *Proceedings of Machine Translation Summit XVII Volume 1: Research Track*, pages 204–213, Dublin, Ireland. European Association for Machine Translation.
- Bryan Eikema and Wilker Aziz. 2020. [Is map decoding all you need? the inadequacy of the mode in neural machine translation](#).
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*.
- Philipp Koehn and Rebecca Knowles. 2017. [Six challenges for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.
- Aviral Kumar and Sunita Sarawagi. 2019. [Calibration of encoder decoder models for neural machine translation](#). *CoRR*, abs/1903.00802.
- A. Marzal and E. Vidal. 1993. [Computation of normalized edit distance and applications](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(9):926–932.
- Clara Meister, Tim Vieira, and Ryan Cotterell. 2020. [If beam search is the answer, what was the question?](#)
- Kenton Murray and David Chiang. 2018. [Correcting length bias in neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 212–223, Brussels, Belgium. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of NAACL-HLT 2019: Demonstrations*.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. [Librispeech: An asr corpus based on public domain audio books](#). In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Pavel Sountsov and Sunita Sarawagi. 2016. [Length bias in encoder decoder models and a case for global conditioning](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1516–1525, Austin, Texas. Association for Computational Linguistics.

- Felix Stahlberg and Bill Byrne. 2019. [On NMT search errors and model errors: Cat got your tongue?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3356–3362, Hong Kong, China. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need.](#) In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Changhan Wang, Yun Tang, Xutai Ma, Anne Wu, Dmytro Okhonko, and Juan Pino. 2020a. [Fairseq S2T: Fast speech-to-text modeling with fairseq.](#) In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 33–39, Suzhou, China. Association for Computational Linguistics.
- Chaojun Wang and Rico Sennrich. 2020. [On exposure bias, hallucination and domain shift in neural machine translation.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3544–3552, Online. Association for Computational Linguistics.
- Shuo Wang, Zhaopeng Tu, Shuming Shi, and Yang Liu. 2020b. [On the inference calibration of neural machine translation.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3070–3079, Online. Association for Computational Linguistics.
- Y. Wu, Mike Schuster, Z. Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, M. Krikun, Yuan Cao, Q. Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, M. Johnson, X. Liu, L. Kaiser, S. Gouws, Y. Kato, Taku Kudo, H. Kazawa, K. Stevens, G. Kurian, Nishant Patil, W. Wang, C. Young, J. Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, G. S. Corrado, Macduff Hughes, and J. Dean. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *ArXiv*, abs/1609.08144.
- Yilin Yang, Liang Huang, and Mingbo Ma. 2018. [Breaking the beam search curse: A study of \(re-\)scoring methods and stopping criteria for neural machine translation.](#) In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3054–3059, Brussels, Belgium. Association for Computational Linguistics.
- Zijian Yang, Yingbo Gao, Weiyue Wang, and Hermann Ney. 2020. [Predicting and using target length in neural machine translation.](#) In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 389–395, Suzhou, China. Association for Computational Linguistics.
- Wei Zhou, Ralf Schlüter, and Hermann Ney. 2020. Robust beam search for encoder-decoder attention based speech recognition without length bias. In *Interspeech*, pages 1768–1772, Shanghai, China.