

Correcting Length Bias in Neural Machine Translation

Kenton Murray and David Chiang

Department of Computer Science and Engineering

University of Notre Dame

{kmurray4, dchiang}@nd.edu

Abstract

We study two problems in neural machine translation (NMT). First, in beam search, whereas a wider beam should in principle help translation, it often hurts NMT. Second, NMT has a tendency to produce translations that are too short. Here, we argue that these problems are closely related and both rooted in label bias. We show that correcting the brevity problem almost eliminates the beam problem; we compare some commonly-used methods for doing this, finding that a simple per-word reward works well; and we introduce a simple and quick way to tune this reward using the perceptron algorithm.

1 Introduction

Although highly successful, neural machine translation (NMT) systems continue to be plagued by a number of problems. We focus on two here: the beam problem and the brevity problem.

First, machine translation systems rely on heuristics to search through the intractably large space of possible translations. Most commonly, beam search is used during the decoding process. Traditional statistical machine translation systems often rely on large beams to find good translations. However, in neural machine translation, increasing the beam size has been shown to degrade performance. This is the last of the six challenges identified by Koehn and Knowles (2017).

The second problem, noted by several authors, is that NMT tends to generate translations that are too short. Jean et al. (2015) and Koehn and Knowles address this by dividing translation scores by their length, inspired by work on audio chords (Boulanger-Lewandowski et al., 2013). A similar method is also used by Google’s production system (Wu et al., 2016). A third simple method used by various authors (Och and Ney, 2002; He et al., 2016; Neubig, 2016) is a tunable

reward added for each output word. Huang et al. (2017) and Yang et al. (2018) propose variations of this reward that enable better guarantees during search.

In this paper, we argue that these two problems are related (as hinted at by Koehn and Knowles) and that both stem from *label bias*, an undesirable property of models that generate sentences word by word instead of all at once.

The typical solution is to introduce a sentence-level correction to the model. We show that making such a correction almost completely eliminates the beam problem. We compare two commonly-used corrections, length normalization and a word reward, and show that the word reward is slightly better.

Finally, instead of tuning the word reward using grid search, we introduce a way to learn it using a perceptron-like tuning method. We show that the optimal value is sensitive both to task and beam size, implying that it is important to tune for every model trained. Fortunately, tuning is a quick post-training step.

2 Problem

Current neural machine translation models are examples of locally normalized models, which estimate the probability of generating an output sequence $e = e_{1:m}$ as

$$P(e_{1:m}) = \prod_{i=1}^m P(e_i | e_{1:i-1}).$$

For any partial output sequence $e_{1:i}$, let us call $P(e' | e_{1:i})$, where e' ranges over all possible completions of $e_{1:i}$, the *suffix distribution* of $e_{1:i}$. The suffix distribution must sum to one, so if the model overestimates $P(e_{1:i})$, there is no way for the suffix distribution to downgrade it. This is known as *label bias* (Bottou, 1991; Lafferty et al., 2001).




Figure 1: Label bias causes this toy word-by-word translation model to translate French *un hélicoptère* incorrectly to *an autogyro*.

2.1 Label bias in sequence labeling

Label bias was originally identified in the context of HMMs and MEMMs for sequence-labeling tasks, where the input sequence f and output sequence e have the same length, and $P(e_{1:i})$ is conditioned only on the partial input sequence $f_{1:i}$. In this case, since $P(e_{1:i})$ has no knowledge of future inputs, it’s much more likely to be incorrectly estimated. For example, suppose we had to translate, word-by-word, *un hélicoptère* to *a helicopter* (Figure 1). Given just the partial input *un*, there is no way to know whether to translate it as *a* or *an*. Therefore, the probability for the incorrect translation $P(an)$ will turn out to be an overestimate. As a result, the model will overweight translations beginning with *an*, regardless of the next input word.

This effect is most noticeable when the suffix distribution has low entropy, because even when new input (*hélicoptère*) is revealed, the model will tend to ignore it. For example, suppose that the available translations for *hélicoptère* are *helicopter*, *chopper*, *whirlybird*, and *autogyro*. The partial translation *a* must divide its probability mass among the three translations that start with a consonant, while *an* gives all its probability mass to *autogyro*, causing the incorrect translation *an autogyro* to end up with the highest probability.

In this example, $P(an)$, even though overestimated, is still lower than $P(a)$, and wins only because its suffixes have higher probability. Greedy search would prune the incorrect prefix *an* and yield the correct output. In general, then, we might expect greedy or beam search to alleviate some symptoms of label bias. Namely, a prefix with a low-entropy suffix distribution can be pruned if its probability is, even though overestimated, not among the highest probabilities. Such an observation was made by Zhang and Nivre (2012) in the

context of dependency parsing, and we will see next that precisely such a situation affects output length in NMT.

2.2 Length bias in NMT

In NMT, unlike the word-by-word translation example in the previous section, each output symbol is conditioned on the entire input sequence. Nevertheless, it’s still possible to overestimate or underestimate $p(e_{1:i})$, so the possibility of label bias still exists. We expect that it will be more visible with weaker models, that is, with less training data.

Moreover, in NMT, the output sequence is of variable length, and generation of the output sequence stops when $\langle /s \rangle$ is generated. In effect, for any prefix ending with $\langle /s \rangle$, the suffix distribution has zero entropy. This situation parallels example of the previous section closely: if the model overestimates the probability of outputting $\langle /s \rangle$, it may proceed to ignore the rest of the input and generate a truncated translation.

Figure 2 illustrates how this can happen. Although the model can learn not to prefer shorter translations by predicting a low probability for $\langle /s \rangle$ early on, at each time step, the score of $\langle /s \rangle$ puts a limit on the total remaining score a translation can have; in the figure, the empty translation has score -10.1 , so that no translation can have score lower than -10.1 . This lays a heavy burden on the model to correctly guess the total score of the whole translation at the outset.

As in our label-bias example, greedy search would prune the incorrect empty translation. More generally, consider beam search: at time step t , only the top k partial or complete translations are retained while the rest are pruned. (Implementations of beam search vary in the details, but this variant is simplest for the sake of argument.) Even if a translation ending at time t scores higher than a longer translation, as long as it does not fall within the top k when compared with partial translations of length t (or complete translations of length at most t), it will be pruned and unable to block the longer translation. But if we widen the beam (k), then translation accuracy will suffer. We call this problem (which is Koehn and Knowles’s sixth challenge) the *beam* problem. Our claim, hinted at by Koehn and Knowles (2017), is that the brevity problem and the beam problem are essentially the same, and that solving one will solve the other.

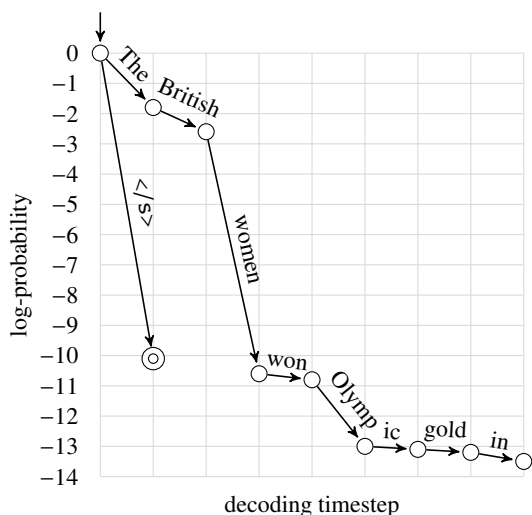


Figure 2: A locally normalized model must determine, at each time step, a “budget” for the total remaining log-probability. In this example sentence, “The British women won Olympic gold in pairs rowing,” the empty translation has initial position 622 in the beam. Already by the third step of decoding, the correct translation has a lower score than the empty translation. However, using greedy search, a nonempty translation would be returned.

3 Correcting Length

To address the brevity problem, many designers of NMT systems add corrections to the model. These corrections are often presented as modifications to the search procedure. But, in our view, the brevity problem is essentially a modeling problem, and these corrections should be seen as modifications to the model (Section 3.1). Furthermore, since the root of the problem is local normalization, our view is that these modifications should be trained as globally-normalized models (Section 3.2).

3.1 Models

Without any length correction, the standard model score (higher is better) is:

$$s(e) = \sum_{i=1}^m \log P(e_i | e_{1:i}).$$

To our knowledge, there are three methods in common use for adjusting the model to favor longer sentences.

Length normalization divides the score by m (Koehn and Knowles, 2017; Jean et al., 2015; Boulanger-Lewandowski et al., 2013):

$$s'(e) = s(e) / m.$$

Google’s NMT system (Wu et al., 2016) relies on a more complicated correction:

$$s'(e) = s(e) \left/ \frac{(5 + m)^\alpha}{(5 + 1)^\alpha} \right.$$

Finally, some systems add a constant *word reward* (He et al., 2016):

$$s'(e) = s(e) + \gamma m.$$

If $\gamma = 0$, this reduces to the baseline model. The advantage of this simple reward is that it can be computed on partial translations, making it easier to integrate into beam search.

3.2 Training

All of the above modifications can be viewed as modifications to the base model so that it is no longer a locally-normalized probability model.

To train this model, in principle, we should use something like the globally-normalized negative log-likelihood:

$$L = -\log \frac{\exp s'(e^*)}{\sum_e \exp s'(e)}$$

where e^* is the reference translation. However, optimizing this is expensive, as it requires performing inference on every training example or heuristic approximations (Andor et al., 2016; Shen et al., 2016).

Alternatively, we can adopt a two-tiered model, familiar from phrase-based translation (Och and Ney, 2002), first training s and then training s' while keeping the parameters of s fixed, possibly on a smaller dataset. A variety of methods, like minimum error rate training (Och, 2003; He et al., 2016), are possible, but keeping with the globally-normalized negative log-likelihood, we obtain, for the constant word reward, the gradient:

$$\frac{\partial L}{\partial \gamma} = -|e^*| + E[|e|].$$

If we approximate the expectation using the mode of the distribution, we get

$$\frac{\partial L}{\partial \gamma} \approx -|e^*| + |\hat{e}|$$

where \hat{e} is the 1-best translation. Then the stochastic gradient descent update is just the familiar perceptron rule:

$$\gamma \leftarrow \gamma + \eta (|e^*| - |\hat{e}|),$$

although below, we update on a batch of sentences rather than a single sentence. Since there is only one parameter to train, we can train it on a relatively small dataset.

Length normalization does not have any additional parameters, with the result (in our opinion, strange) that a change is made to the model without any corresponding change to training. We could use gradient-based methods to tune the α in the GNMT correction, but the perceptron approximation turns out to drive α to ∞ , so a different method would be needed.

4 Experiments

We compare the above methods in four settings, a high-resource German–English system, a medium-resource Russian–English system, and two low-resource French–English and English–French systems. For all settings, we show that larger beams lead to large BLEU and METEOR drops if not corrected. We also show that the optimal parameters can depend on the task, language pair, training data size, as well as the beam size. These values can affect performance strongly.

4.1 Data and settings

Most of the experimental settings below follow the recommendations of [Denkowski and Neubig \(2017\)](#). Our high-resource, German–English data is from the 2016 WMT shared task ([Bojar et al., 2016](#)). We use a bidirectional encoder-decoder model with attention ([Bahdanau et al., 2015](#)).¹ Our word representation layer has 512 hidden units, while other hidden layers have 1024 nodes. Our model is trained using Adam with a learning rate of 0.0002. We use 32k byte-pair encoding (BPE) operations learned on the combined source and target training data ([Sennrich et al., 2016](#)). We train on minibatches of size 2012 words and validate every 100k sentences, selecting the final model based on development perplexity.

Our medium-resource, Russian–English system uses data from the 2017 WMT translation task, which consists of roughly 1 million training sentences ([Bojar et al., 2017](#)). We use the same architecture as our German–English system, but only have 512 nodes in all layers. We use 16k BPE operations and dropout of 0.2. We train on mini-

¹We use Lamtram ([Neubig, 2015](#)) for all experiments and our modifications have been added to the project.

batches of 512 words and validate every 50k sentences.

Our low-resource systems use French and English data from the 2010 IWSLT TALK shared task ([Paul et al., 2010](#)). We build both French–English and English–French systems. These networks are the same as for the medium Russian–English task, but use only 6k BPE operations. We train on minibatches of 512 words and validate every 30k sentences, restarting Adam when the development perplexity goes up.

To tune our correction parameters, we use 1000 sentences from the German–English development dataset, 1000 sentences from the Russian–English development dataset, and the entire development dataset for French–English (892 sentences)². We initialize the parameter, $\gamma = 0.2$. We use batch gradient descent, which we found to be much more stable than stochastic gradient descent, and use a learning rate of $\eta = 0.2$, clipping gradients for γ to 0.5. Training stops if all parameters have an update of less than 0.03 or a max of 25 epochs was reached.

4.2 Solving the length problem solves the beam problem

Here, we first show that the beam problem is indeed the brevity problem. We then demonstrate that solving the length problem does solve the beam problem. Tables 1, 2, and 3 show the results of our German–English, Russian–English, and French–English systems respectively. Each table looks at the impact on BLEU, METEOR, and the ratio of the lengths of generated sentences compared to the gold lengths ([Papineni et al., 2002](#); [Denkowski and Lavie, 2014](#)). The baseline method is a standard model without any length correction. The reward method is the tuned constant word reward discussed in the previous section. Norm refers to the normalization method, where a hypothesis’ score is divided by its length.

4.2.1 Baseline

The top sections of Tables 1, 2, 3 illustrate the brevity and beam problems in the baseline models. As beam size increases, the BLEU and METEOR scores drop significantly. This is due to the brevity problem, which is illustrated by the length ratio numbers that also drop with increased

²We found through preliminary experiments that this size of dev subset was an adequate trade-off between tuning speed and performance.

Russian–English (medium)		Beam Size					
		10	50	75	100	150	1000
baseline	BLEU	24.9	23.8	23.6	23.3	22.5	3.7
	METEOR	30.9	30.0	29.7	29.4	28.8	12.8
	length	0.90	0.86	0.85	0.84	0.81	0.31
reward	BLEU	26.5	26.6	26.5	26.5	26.5	25.7
	METEOR	32.0	32.0	31.9	31.9	31.9	31.2
	length	0.98	0.98	0.98	0.98	0.98	1.02
	γ	0.716	0.643	0.640	0.633	0.617	0.562
norm	BLEU	26.2	26.3	26.3	26.3	26.3	25.3
	METEOR	31.8	31.8	31.8	31.7	31.7	31.2
	length	0.96	0.96	0.96	0.96	0.97	1.02

Table 1: Results of the Russian–English translation system. We report BLEU and METEOR scores, as well as the ratio of the length of generated sentences compared to the correct translations (length). γ is the word reward score discovered during training. Here, we examine a much larger beam (1000). The beam problem is more pronounced at this scale, with the baseline system losing over 20 BLEU points when increasing the beam from size 10 to 1000. However, both our tuned length reward score and length normalization recover most of this loss.

German–English (large)		Beam Size		
		10	50	75
baseline	BLEU	29.6	28.6	28.2
	METEOR	34.0	33.1	32.8
	length	0.95	0.90	0.89
reward	BLEU	30.3	30.6	30.6
	METEOR	34.9	34.8	34.9
	length	1.02	1.00	1.00
	γ	0.67	0.57	0.58
norm	BLEU	30.7	31.0	30.9
	METEOR	34.9	35.0	35.0
	length	1.00	1.00	1.00

Table 2: Results of the high-resource German–English system. Rows: BLEU, METEOR, length = ratio of output to reference length; γ = learned parameter value. While baseline performance decreases with beam size due to the brevity problem, other methods perform more consistently across beam sizes. Length normalization (norm) gets the best BLEU scores, but similar METEOR scores to the word reward.

French–English (small)		Beam Size				
		10	50	100	150	200
baseline	BLEU	30.0	28.9	25.4	21.9	19.4
	METEOR	32.4	31.3	28.6	25.9	24.1
	length	0.94	0.89	0.80	0.71	0.64
reward	BLEU	29.4	29.7	29.7	29.8	29.8
	METEOR	32.8	32.9	32.9	32.9	32.9
	length	1.03	1.03	1.03	1.03	1.03
	γ	1.20	1.05	1.01	0.99	0.97
norm	BLEU	30.7	30.8	30.7	30.7	30.7
	METEOR	32.8	32.8	32.8	32.7	32.7
	length	0.97	0.97	0.97	0.96	0.96

English–French (small)		Beam Size				
		10	50	100	150	200
baseline	BLEU	25.8	26.1	26.1	25.5	24.3
	METEOR	47.8	47.5	47.2	46.3	44.2
	length	1.03	1.01	1.00	0.97	0.92
reward	BLEU	25.5	25.5	25.5	25.5	25.5
	METEOR	48.3	48.5	48.5	48.5	48.4
	length	1.05	1.05	1.05	1.05	1.05
	γ	0.353	0.444	0.465	0.474	0.475
norm	BLEU	25.4	25.5	25.5	25.5	25.5
	METEOR	48.4	48.4	48.4	48.4	48.4
	length	1.06	1.05	1.05	1.05	1.05

Table 3: Results of low-resource French–English and English–French systems. Rows: BLEU, METEOR, length = ratio of output to reference length; γ = learned parameter value. While baseline performance decreases with beam size due to the brevity problem, other methods perform more consistently across beam sizes. Word reward gets the best scores in both directions on METEOR. Length normalization (norm) gets the best BLEU scores in Fra-Eng due to the slight bias of BLEU towards shorter translations.

beam	10	50	75	100	150	200
French–English (small)	6.9	27.2	52.4	71.1	105.9	176.6
English–French (small)	12.6	44.2	67.3	88.1	107.5	111.2
German–English (large)	6.8	132.6	1066			

Table 4: Tuning time on top of baseline training time. Times are in minutes on 1000 dev examples (German–English) or 892 dev examples (French–English). Due to the much larger model size, we only looked at beam sizes up to 75 for German–English.

beam size. For larger beam sizes, the length of the generated output sentences are a fraction of the lengths of the correct translations. For the lower-resource French–English task, the drop is more than 8 BLEU when increasing the beam size from 10 to 150. The issue is even more evident in our Russian–English system where we increase the beam to 1000 and BLEU scores drop by more than 20 points.

4.2.2 Word reward

The results of tuning the word reward, γ , as described in Section 3.2, is shown in the second section of Tables 1, 2, and 3. In contrast to our baseline systems, our tuned word reward always fixes the brevity problem (length ratios are approximately 1.0), and generally fixes the beam problem. An optimized word reward score always leads to improvements in METEOR scores over any of the best baselines. Across all language pairs, reward and norm have close METEOR scores, though the reward method wins out slightly. BLEU scores for reward and norm also increase over the baseline in most cases, despite BLEU’s inherent bias towards shorter sentences. Most notably, whereas the baseline Russian–English system lost more than 20 BLEU points when the beam was increased to 1000, our tuned reward score resulted in a BLEU gain over any baseline beam size. Whereas in our baseline systems, the length ratio decreases with larger beam sizes, our tuned word reward results in length ratios of nearly 1.0 across all language pairs, mitigating many of the issues of the brevity problem.

4.2.3 Wider beam

We note that the beam problem in NMT exists for relatively small beam sizes – especially when compared to traditional beam sizes in SMT systems. On our medium-resource Russian–English system, we investigate the full impact of this problem using a much larger beam size of 1000. In Table 1, we can see that the beam problem is particularly pronounced. The first row of the table shows the uncorrected, baseline score. From a beam of 10 to a beam of 1000, the drop in BLEU scores is over 20 points. This is largely due to the brevity problem discussed earlier. The second row of the table shows the length of the translated outputs compared to the lengths of the correct translations. Though the problem persists even at a beam size of 10, at a beam size of 1000, our baseline system

generates less than one third the number of words that are in the correct translations. Furthermore, 37.3% of our translated outputs have sentences of length 0. In other words, the most likely translation is to immediately generate the stop symbol. This is the problem visualized in Figure 2.

However, when we tune our word reward score with a beam of 1000, the problem mostly goes away. Over the uncorrected baseline, we see a 22.0 BLEU point difference for a beam of 1000. Over the uncorrected baseline with a beam of 10, the corrected beam of 1000 gets a BLEU gain of 0.8 BLEU. However, the beam of 1000 still sees a drop of less than 1.0 BLEU over the best corrected version. The word reward method beats the uncorrected baseline and the length normalization correction in almost all cases.

4.2.4 Short sentences

Another way to demonstrate that the beam problem is the same as the brevity problem is to look at the translations generated by baseline systems on shorter sentences. Figure 3 shows the BLEU scores of the Russian–English system for beams of size 10 and 1000 on sentences of varying lengths, with and without correcting lengths. The x-axes of the figure are cumulative: length 20 includes sentences of length 0–20, while length 10 includes 0–10. It is worth noting that BLEU is a word-level metric, but the systems were built using BPE; so the sequences actually generated are longer than the x-axes would suggest.

The baseline system on sentences with 10 words or less still has relatively high BLEU scores—even for a beam of 1000. Though there is a slight drop in BLEU (less than 2), it is not nearly as severe as when looking at the entire test set (more than 20). When correcting for length with normalization or word reward, the problem nearly disappears when considering the entire test set, with reward doing slightly better. For comparison, the rightmost points in each of the subplots correspond to the BLEU scores in columns 10 and 1000 of Table 1. This suggests that the beam problem is strongly related to the brevity problem.

4.2.5 Length ratio

The interaction between the length problem and the beam problem can be visualized in the histograms of Figure 4 on the Russian–English system. In the upper left plot, the uncorrected model with beam 10 has the majority of the generated




Figure 3: Impact of beam size on BLEU score when varying reference sentence lengths (in words) for Russian–English. The x-axis is cumulative moving right; length 20 includes sentences of length 0-20, while length 10 includes 0-10. As reference length increases, the BLEU scores of a baseline system with beam size of 10 remain nearly constant. However, a baseline system with beam 1000 has a high BLEU score for shorter sentences, but a very low score when the entire test set is used. Our tuned reward and normalized models do not suffer from this problem on the entire test set, but take a slight performance hit on the shortest sentences.




Figure 4: Histogram of length ratio between generated sentences and gold varied across methods and beam size for Russian–English. Note that the baseline method skews closer 0 as the beam size increases, while our other methods remain peaked around 1.0. There are a few outliers to the right that have been cut off, as well as the peaks at 0.0 and 1.0.

sentences with a length ratio close to 1.0, the gold lengths. Going down the column, as the beam size increases, the distribution of length ratios skews closer to 0. By a beam size of 1000, 37% of the sentences have a length of 0. However, both the word reward and the normalized models remain very peaked around a length ratio of 1.0 even as the beam size increases.

4.3 Tuning word reward

Above, we have shown that fixing the length problem with a word reward score fixes the beam problem. However these results are contingent upon choosing an adequate word reward score, which we have done in our experiments by optimization using a perceptron loss. Here, we show the sensitivity of systems to the value of this penalty, as well as the fact that there is not one correct penalty for all tasks. It is dependent on a myriad of factors including, beam size, dataset, and language pair.

4.3.1 Sensitivity to γ

In order to investigate how sensitive a system is to the reward score, we varied values of γ from 0 to 1.2 on both our German–English and Russian–English systems with a beam size of 50. BLEU scores and length ratios on 1000 heldout development sentences are shown in Figure 5. The length ratio is correlated with the word reward as expected, and the BLEU score varies by more than 5 points for German–English and over 4.5 points for Russian–English. On German–English, our method found a value of $\gamma = 0.57$, which is slightly higher than optimal; this is because the heldout sentences have a slightly shorter length ratio than the training sentences. Conversely, on Russian–English, our found value of $\gamma = 0.64$ is slightly lower than optimal as these heldout sentences have a slightly higher length ratio than the sentences used in training.

4.3.2 Optimized γ values

Tuning the reward penalty using the method described in Section 3.2 resulted in consistent improvements in METEOR scores and length ratios across all of our systems and language pairs. Tables 1, 2, and 3 show the optimized value of γ for each beam size. Within a language pair, the optimal value of γ is different for every beam size. Likewise, for a given beam size, the optimal value is different for every system. Our French–English and English–French systems in Table 3 have the




Figure 5: Effect of word penalty on BLEU and hypothesis length for Russian–English (top) and German–English (bottom) on 1000 unseen dev examples with beams of 50. Note that the vertical bars represent the word reward that was found during tuning.

exact same architecture, data, and training criteria. Yet, even for the same beam size, the tuned word reward scores are very different.

Training dataset size Low-resource neural machine translation performs significantly worse than high-resource machine translation (Koehn and Knowles, 2017). Table 5 looks at the impact of training data size on BLEU scores and the beam problem by using 10% and 50% of the available Russian–English data. Once again, the optimal value of γ is different across all systems and beam sizes. Interestingly, as the amount of training data decreases, the gains in BLEU using a tuned reward penalty increase with larger beam sizes. This suggests that the beam problem is more prevalent in lower-resource settings, likely due to the fact that less training data can increase the effects of label bias.

Russian–English (medium) Dataset Size		Beam Size				
		10	50	75	100	150
100%	baseline	24.9	23.8	23.6	23.3	22.5
	reward	26.5	26.6	26.5	26.5	26.5
	γ	0.716	0.643	0.640	0.633	0.617
50%	baseline	22.8	21.4	20.8	20.4	19.2
	reward	24.7	25.0	24.9	24.9	25.0
	γ	0.697	0.645	0.638	0.636	0.646
10%	baseline	17.0	16.2	15.8	15.6	15.1
	reward	17.6	18.0	18.0	18.0	18.1
	γ	0.892	0.835	0.773	0.750	0.800

Table 5: Varying the size of the Russian–English training dataset results in different optimal word reward scores (γ). In all settings, the tuned score alleviates the beam problem. As the datasets get smaller, using a tuned larger beam improves the BLEU score over a smaller tuned beam. This suggests that lower-resource systems are more susceptible to the beam problem.

4.3.3 Tuning time

Fortunately, the tuning process is very inexpensive. Although it requires decoding on a development dataset multiple times, we only need a small dataset. The time required for tuning our French–English and German–English systems is shown in Table 4. These experiments were run on an Nvidia GeForce GTX 1080Ti. The tuning usually takes a few minutes to hours, which is just a fraction of the overall training time. We note that there are numerous optimizations that could be taken to speed this up even more, such as storing the decoding lattice for partial reuse. However, we leave this for future work.

4.4 Word reward vs. length normalization

Tuning the word reward score generally had higher METEOR scores than length normalization across all of our settings. With BLEU, length normalization beat the word reward on German–English and French–English, but tied on English–French and lost on Russian–English. For the largest beam of 1000, the tuned word reward had a higher BLEU than length normalization. Overall, the two methods have relatively similar performance, but the tuned word reward has the more theoretically justified, globally-normalized derivation – especially in the context of label bias’ influence on the brevity problem.

5 Conclusion

We have explored simple and effective ways to alleviate or eliminate the beam problem. We showed

that the beam problem can largely be explained by the brevity problem, which results from the locally-normalized structure of the model. We compared two corrections to the model and introduced a method to learn the parameters of these corrections. Because this method is helpful and easy, we hope to see it included to make stronger baseline NMT systems.

We have argued that the brevity problem is an example of label bias, and that the solution is a very limited form of globally-normalized model. These can be seen as the simplest case of the more general problem of label bias and the more general solution of globally-normalized models for NMT (Wiseman and Rush, 2016; Venkatraman et al., 2015; Ranzato et al., 2015; Shen et al., 2016). Some questions for future research are:

- Solving the brevity problem leads to significant BLEU gains; how much, if any, improvement remains to be gained by solving label bias in general?
- Our solution to the brevity problem requires globally-normalized training on only a small dataset; can more general globally-normalized models be trained in a similarly inexpensive way?

Acknowledgements

This research was supported in part by University of Southern California, subcontract 67108176 under DARPA contract HR0011-15-C-0115, and an Amazon Research Award to Chiang.

References

- Daniel Andor, Chris Alberti, David Weiss, Aliaksei Severyn, Alessandro Presta, Kuzman Ganchev, Slav Petrov, and Michael Collins. 2016. Globally normalized transition-based neural networks. In *Proc. ACL*, pages 2442–2452.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proc. ICLR*.
- 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 *Proc. Conference on Machine Translation*, pages 169–214.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, et al. 2016. Findings of the 2016 Conference on Machine Translation. In *Proc. Conference on Machine Translation*, volume 2, pages 131–198.
- Léon Bottou. 1991. *Une Approche théorique de l'Apprentissage Connexionniste; Applications à la reconnaissance de la Parole*. Ph.D. thesis, Université de Paris Sud.
- Nicolas Boulanger-Lewandowski, Yoshua Bengio, and Pascal Vincent. 2013. Audio chord recognition with recurrent neural networks. In *ISMIR*, pages 335–340. Citeseer.
- Michael Denkowski and Alon Lavie. 2014. Meteor universal: Language specific translation evaluation for any target language. In *Proc. Workshop on Statistical Machine Translation*.
- Michael Denkowski and Graham Neubig. 2017. Stronger baselines for trustable results in neural machine translation. In *Proceedings of the First Workshop on Neural Machine Translation*, pages 18–27.
- Wei He, Zhongjun He, Hua Wu, and Haifeng Wang. 2016. Improved neural machine translation with SMT features. In *Proc. AAAI*.
- Liang Huang, Kai Zhao, and Mingbo Ma. 2017. When to finish? optimal beam search for neural text generation (modulo beam size). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2134–2139.
- Sébastien Jean, Orhan Firat, Kyunghyun Cho, Roland Memisevic, and Yoshua Bengio. 2015. Montreal neural machine translation systems for wmt'15. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 134–140.
- 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.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proc. ICML*, pages 282–289.
- Graham Neubig. 2015. lamtram: A toolkit for language and translation modeling using neural networks. <http://www.github.com/neubig/lamtram>.
- Graham Neubig. 2016. Lexicons and minimum risk training for neural machine translation: NAIST-CMU at WAT2016. In *Proceedings of the 3rd Workshop on Asian Translation*, pages 119–125.
- Franz Josef Och. 2003. Minimum error rate training in statistical machine translation. In *Proc. ACL*, pages 160–167.
- Franz Josef Och and Hermann Ney. 2002. Discriminative training and maximum entropy models for statistical machine translation. In *Proc. ACL*, pages 295–302. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proc. ACL*, pages 311–318.
- Michael Paul, Marcello Federico, and Sebastian Stüker. 2010. Overview of the IWSLT 2010 evaluation campaign. In *International Workshop on Spoken Language Translation (IWSLT) 2010*.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. Sequence level training with recurrent neural networks. In *Proceedings of ICLR*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proc. ACL*, pages 1715–1725.
- Shiqi Shen, Yong Cheng, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2016. Minimum risk training for neural machine translation. In *Proc. of ACL*.
- Arun Venkatraman, Martial Hebert, and J Andrew Bagnell. 2015. Improving multi-step prediction of learned time series models. In *AAAI*, pages 3024–3030.
- Sam Wiseman and Alexander M Rush. 2016. Sequence-to-sequence learning as beam-search optimization. In *Proceedings of EMNLP*.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. arXiv: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*.
- Yue Zhang and Joakim Nivre. 2012. Analyzing the effect of global learning and beam-search on transition-based dependency parsing. In *Proceedings of COLING 2012: Posters*, pages 1391–1400.