



Robust Beam Search for Encoder-Decoder Attention Based Speech Recognition without Length Bias

Wei Zhou, Ralf Schlüter, Hermann Ney

Human Language Technology and Pattern Recognition, Computer Science Department,
RWTH Aachen University, 52074 Aachen, Germany
AppTek GmbH, 52062 Aachen, Germany
{zhou, schlueter, ney}@cs.rwth-aachen.de

Abstract

As one popular modeling approach for end-to-end speech recognition, attention-based encoder-decoder models are known to suffer the length bias and corresponding beam problem. Different approaches have been applied in simple beam search to ease the problem, most of which are heuristic-based and require considerable tuning. We show that heuristics are not proper modeling refinement, which results in severe performance degradation with largely increased beam sizes. We propose a novel beam search derived from reinterpreting the sequence posterior with an explicit length modeling. By applying the reinterpreted probability together with beam pruning, the obtained final probability leads to a robust model modification, which allows reliable comparison among output sequences of different lengths. Experimental verification on the LibriSpeech corpus shows that the proposed approach solves the length bias problem without heuristics or additional tuning effort. It provides robust decision making and consistently good performance under both small and very large beam sizes. Compared with the best results of the heuristic baseline, the proposed approach achieves the same WER on the ‘clean’ sets and 4% relative improvement on the ‘other’ sets. We also show that it is more efficient with the additional derived early stopping criterion.

Index Terms: speech recognition, encoder-decoder, beam search, length bias

1. Introduction & Related Work

So called “end-to-end” speech recognition enables the direct mapping of acoustic feature sequences to sub-word or word sequences. One of the most successful end-to-end approaches is the attention-based encoder-decoder model [1], which has achieved promising results in speech recognition [2, 3, 4, 5]. For attention-based encoder-decoder systems without monotonic constraints, there is generally no explicit time or positional information in the output sequences w.r.t the input sequences. Such systems usually apply label-synchronous search for decoding, where mostly a sequence end label is used for termination. Simple beam search is used for most end-to-end systems, where only an absolute beam size limit controls the complete search procedure.

Encoder-decoder models are known to suffer the length bias problem due to the locally normalized training objective [6]. In short, models can produce higher sequence posterior for much shorter output sequences than for the correct ones. This behavior becomes more obvious with larger beam sizes in beam search, which leads to the beam problem. Reasonable performance is only achieved with very small beam sizes, where search errors are adopted to avoid model errors [7]. Such issues are observed in many applications such as speech recogni-

tion [8, 9] and neural machine translation (NMT) [10, 11, 12]. While an additional length model may solve the problem, it is often ignored or considered to be implicitly learned with existing models.

Instead, many different approaches have been applied in simple beam search to ease the problem. The general goal is to prevent too short output sequences and to allow reliable comparison among output sequences of different lengths. In decoding, scores are commonly used by taking logarithm of probabilities. One straightforward and widely-used approach is the length normalization [13, 11], which divides the score of a sequence by its length. Another common approach is the end-of-sentence (EOS) threshold [9, 8], which allows a sequence end label to appear only if its score is better than the current best non-end one multiplied by a predefined factor. [4] combined these two approaches and obtained good results with beam size 64, which we will use as our baseline. Another approach is a length reward term added to the score of a sequence based on its length [14, 15, 16], where the scaling value requires careful tuning on each data set. A more sophisticated but less heuristic-based approach is the coverage term [8, 17], which is added to the score based on the output sequence’s coverage over the input. It requires more complicated coverage computation based on all accumulated attention weights for each hypothesis up to the current search step and involves several threshold values to be tuned. [18] combined length reward and coverage term, and obtained stable results up to beam size 240. However, the resulting system has seven individual hyper-parameters to be optimized for decoding, which is a huge tuning effort.

While these approaches largely eliminate the length bias problem, they are either pure heuristics or difficult to optimize. Their usage and results are mostly reported with beam sizes below 240. The potential side-effect due to the additional bias introduced towards longer sequences is often disregarded. We show that with a much larger beam size, such bias leads to more wrong decisions towards too long transcriptions, which results in severe performance degradation. This suggests that the heuristic approaches are not proper modeling refinement and make the decisions less robust w.r.t. length variation and search beam size.

In this work, we propose a novel beam search derived from reinterpreting the sequence posterior with an explicit length modeling. By applying the reinterpreted probability together with beam pruning, the resulting final probability is obtained from pure estimations based on models’ output without heuristics. This leads to a robust model modification which allows reliable comparison among output sequences of different lengths. Experimental verification on the LibriSpeech corpus [19] shows that the proposed approach eliminates the length bias problem without heuristics or additional tuning effort. Compared with

the heuristic baseline, it achieves better performance and shows better efficiency with the additional derived early stopping criterion. More importantly, without introducing additional side-effects, the proposed approach provides robust decision making and consistently good performance under both small and very large beams. It is also applicable to streaming usage as well as other tasks such as NMT.

2. Proposed Beam Search

2.1. Probability reinterpretation

Let x_1^T denote an input sequence of length T and a_1^N denote partial output sequence hypotheses at the N -th step of beam search, where N also represents output label position. The original sequence posterior probability of a_1^N is quantified by:

$$\begin{aligned} q(a_1^N | x_1^T) &= p(a_1^N | x_1^T) \cdot p^\alpha(a_1^N) \\ &= \prod_{n=1}^N p(a_n | a_0^{n-1}, x_1^T) \cdot p^\alpha(a_n | a_0^{n-1}) \end{aligned} \quad (1)$$

The optional language model (LM) shallow fusion [20] with scale α can be omitted without influencing the derivation.

Let $V \cup \{\$\}$ define the output label vocabulary, where $\$$ is the sequence end label. If $a_N = \$$, then a_1^N represents ending sequences at position N . Ending sequences are terminated without further expansion and are stored separately. Therefore, $a_N = \$$ also implies $a_1^{N-1} \in V^{N-1}$, which we omit in all equations for simplicity. By considering $\$$ as the last output label, the output sequence length of ending sequences a_1^N is obtained as $len = N$, which reversely implies $a_N = \$$. For ending sequences at position N , we rewrite their final probability with an explicit length modeling:

$$\begin{aligned} p(a_1^N, len = N | x_1^T) &= p(a_1^N | len = N, x_1^T) \cdot p(len = N | x_1^T) \\ p(a_1^N | len = N, x_1^T) &= \frac{q(a_1^N | x_1^T)}{\sum_{\{\hat{a}_1^N : \hat{a}_N = \$\} \in B_N} q(\hat{a}_1^N | x_1^T)} \\ p(len = N | x_1^T) &= p_N(\$ | x_1^T) \prod_{n=1}^{N-1} (1 - p_n(\$ | x_1^T)) \end{aligned} \quad (2)$$

$$p_N(\$ | x_1^T) = \frac{\sum_{\{\hat{a}_1^N : \hat{a}_N = \$\} \in B_N} q(\hat{a}_1^N | x_1^T)}{\sum_{\hat{a}_1^N \in B_N} q(\hat{a}_1^N | x_1^T)} \quad (3)$$

Here $B_N = \{a_1^N | a_1^{N-1} \in V^{N-1}, a_N \in V \cup \{\$\}\}$ is an unlimited beam of all label sequence hypotheses reaching position N , which can end at positions larger or equal to N . Additionally, we define $p_N(\$ | x_1^T)$ as the ending probability at position N . It is obtained by re-normalizing the probability mass of all label sequences ending at position N over the probability mass of all label sequences reaching position N . Accordingly, $1 - p_N(\$ | x_1^T)$ accounts for the non-ending probability at position N . Therefore, the probability of finishing with output sequence length $len = N$, i.e. Eq. (2), can be obtained by multiplying the accumulated non-ending probabilities from positions 1 to $N - 1$ with the ending probability at position N . By merging all terms, we obtain the final probability of ending sequences at position N as:

$$p(a_1^N, len = N | x_1^T) = \underbrace{\frac{q(a_1^N | x_1^T)}{\sum_{\hat{a}_1^N \in B_N} q(\hat{a}_1^N | x_1^T)}}_{p_B} \cdot \underbrace{\prod_{n=1}^{N-1} (1 - p_n(\$ | x_1^T))}_{p_{1\$}} \quad (4)$$

Note that with an unlimited beam B at each step, $\prod_{n=1}^{N-1} (1 - p_n(\$ | x_1^T))$ is equal to $\sum_{\hat{a}_1^N \in B_N} q(\hat{a}_1^N | x_1^T)$. Both represent the probability mass of all label sequence hypotheses

reaching position N . This verifies the derivation of the reinterpreted final probability which leads to the same sequence posterior as in Eq. (1). Note that no additional parameters or model training are introduced here.

2.2. Beam search with pruning

We then apply this reinterpreted final probability Eq. (4) into normal beam search, where B_N becomes a limited beam after pruning. At each search step N , we use the sequence posterior in Eq. (1) to directly prune all partial label sequence hypotheses a_1^N . Since all of them have the same length up to this position, they are directly comparable. We first apply score-based pruning to prune away hypotheses whose score difference to the current best is more than a predefined threshold. A predefined beam size is then applied if the remaining number of hypotheses still exceeds this upper bound.

Ending sequences are then detected from the remaining hypotheses in the beam B_N , which are used to compute the ending probability of position N according to Eq. (3). We apply the accumulated non-ending probability from all previous positions 1 to $N - 1$ into Eq. (4) to compute the final probability for each ending sequence within B_N . Since all computation only has a dependency on the past, no additional delay is introduced here. All ended sequences up to this search step are stored separately and we only keep the best k of them based on their final probability. Note that we explicitly do not use ended sequences to prune away ongoing sequences in further steps, since they may not be directly comparable.

2.3. Final probability

With such limited beam at each search step, the obtained final probability no longer equals to the original sequence posterior in Eq. (1). It essentially leads to a beam-dependent model modification. The fraction term (denoted as p_B in Eq. (4)) can be interpreted as renormalization within B_N , which estimates the relative quality of the ending sequence within the beam of the current step. The non-ending probability of each previous position is effectively also renormalization within the corresponding beam, which indicates how probable it is to not end at that position. The accumulated non-ending probability from all previous positions (denoted as $p_{1\$}$ in Eq. (4)) then estimates the probability of not finishing before the current position N . Both p_B and $p_{1\$}$ are pure estimations based on the models' output without heuristics, which jointly decide the final probability of ending sequences.

Note that this final probability depends on the beam pruning. For extremely large beams with little pruning, it approaches the original sequence posterior which may still suffer the length bias problem. For extremely small beams with very strong pruning, it can have overestimation problem and search become less reliable, which however does not contradict the concept of beam search. Both cases are very unlikely by simply applying a reasonable threshold for the score-based pruning. This leads to an optimal beam at each step based on scores, which prunes away bad hypotheses while keeping a proper probability mass for renormalization. We observe that even without score-based pruning, the approach works consistently well with both small and very large beams.

In terms of reliable comparison among ending sequences of different lengths based on this final probability, some intuitive interpretation can be given as following. Let M_{opt} denote the correct output sequence length for a given input sequence. At positions much smaller than M_{opt} , ending sequences should have rather small sequence posterior based on a reasonable model. Even if they survive pruning, their final probability

should have a high $p_{1\text{s}}$ but suffer a very low p_B . At positions around M_{opt} , sequence posterior of ending sequences close to the correct transcription become more dominant in the beam. This leads to an increasing p_B and a one-step-delayed decreasing $p_{1\text{s}}$. These ending sequences should have a rather high final probability. Finally at positions much larger than M_{opt} , ending sequences might have a good p_B , but suffer a very low $p_{1\text{s}}$.

2.4. Decision and early stopping

The final best output sequence can be decided using the maximum a posteriori (MAP) decision rule:

$$x_1^T \rightarrow a_1^M = \arg \max_{a_1^M, M} p(a_1^M, \text{len} = M | x_1^T)$$

Since we do not apply pruning between ended sequences and ongoing sequences in further steps, we need to derive a stopping criterion to avoid unnecessary search steps. This can be easily obtained from Eq. (4). Let \tilde{a}_1^M denote the current best ended sequence, where $1 \leq M \leq N$ and N is the current step. All future hypotheses from further steps after N can not be better than \tilde{a}_1^M , if the following holds:

$$\prod_{n=1}^N (1 - p_n(\$ | x_1^T)) \leq p(\tilde{a}_1^M, \text{len} = M | x_1^T)$$

An additional maximum length constraint with respect to the input sequence length can also be added to stop decoding, which is generally valid for ASR. The pseudo code of the proposed beam search is given in Algorithm 1, where the choice of a_0 and the initial computation with or without a_0 can be model-specific.

3. Experiments

3.1. Setups

The proposed beam search is implemented based on the RWTH ASR toolkit¹ [21] with an extension described in [22]. Experiments are conducted on the LibriSpeech corpus [19]. Both the long short term memory [23] (LSTM)-based encoder-decoder attention model and the LSTM LM are the same as described in [4]. They are trained on the LibriSpeech acoustic and LM training data respectively using the RETURNN toolkit [24, 25]. Both models share the same set of about 10k byte-pair encoding (BPE) units. We refer the readers to [4] for more model and training details. Different beam sizes from {32, 64, 128, 5000} are evaluated. Decoding parameters are optimized on each development set and applied to the corresponding test set. All results are obtained with the MAP decision rule.

3.2. Simple beam search with heuristics

We follow [4] to apply simple beam search with heuristics using length normalization and EOS threshold. Here the scale for LM shallow fusion and the EOS threshold factor need to be optimized. [4] reported to achieve the best result with beam size 64. We apply the optimal parameter settings for beam size 64 to all other beam sizes. For a complete comparison, we also include the results of simple beam search without heuristics under beam sizes 64 and 5000. The word error rate (WER) results are shown in Table 1. Additionally for the dev-other set, we show insertion, deletion and substitution errors as well as the average length of recognized transcriptions under beam sizes 64 and 5000. The trend remains the same also for other subsets.

Without heuristics, simple beam search suffers a huge increment of deletion errors from beam size 64 to 5000. The length bias problem and corresponding beam problem are directly visible from the largely degraded performance and much

¹Source code will be published in the next release of RASR.

Algorithm 1: Proposed Beam Search

```

Initialize:  $N = 0, B_0 = \{a_0\}, k_{\text{best}} = \{\},$ 
 $p_{1\text{s}} = 1.0, \text{Stop} = \text{false}$ 
while not Stop do
   $N += 1;$ 
  for  $a_1^{N-1}$  in  $B_{N-1}$  do
    extend to all  $a_1^N$  and compute  $q(a_1^N | x_1^T);$ 
    add all  $a_1^N$  to  $B_N;$ 
  end
  remove  $B_{N-1};$ 
  apply beam pruning in  $B_N;$ 
   $p_{\Sigma} = 0, p_{\Sigma\text{s}} = 0, B_{\text{s}} = \{\};$ 
  for  $a_1^N$  in  $B_N$  do
     $p_{\Sigma} += q(a_1^N | x_1^T);$ 
    if  $a_N == \$$  then
       $p_{\Sigma\text{s}} += q(a_1^N | x_1^T);$ 
      move  $a_1^N$  from  $B_N$  to  $B_{\text{s}};$ 
    end
  end
  for  $a_1^N$  in  $B_{\text{s}}$  do
     $p_{\text{final}}(a_1^N, \text{len} = N | x_1^T) = q(a_1^N | x_1^T) / p_{\Sigma} \cdot p_{1\text{s}};$ 
    insert  $a_1^N$  to  $k_{\text{best}}$  based on  $p_{\text{final}}(a_1^N, \text{len} = N | x_1^T);$ 
  end
   $p_{1\text{s}} := (1 - p_{\Sigma\text{s}} / p_{\Sigma});$ 
  if  $p_{1\text{s}} \leq \text{best } p_{\text{final}}$  in  $k_{\text{best}}$  or  $N \geq T$  then
     $\text{Stop} = \text{true};$ 
  end
end
return  $k_{\text{best}}$ 

```

Table 1: WER comparison of different beam search with different beam sizes on the LibriSpeech corpus. Additional analysis on the dev-other set including insertion, deletion and substitution errors, and average transcription length (BPE units are merged to words already). Reference transcriptions of dev-other set have an average length of 17.8 words.

Beam Search	Beam Size	dev		test		dev-other			
		clean	other	clean	other	ins	del	sub	len
simple	64	5.9	11.1	6.5	12.2	0.6k	1.3k	3.8k	17.5
	5000	19.7	32.0	20.3	35.3	0.3k	13.5k	2.5k	13.2
+ heuristics	32	2.8	7.6	3.1	8.3	n.a.			
	64	2.8	7.5	3.1	8.2	0.6k	0.2k	3.0k	17.9
	128	2.8	7.7	3.1	8.7	n.a.			
	5000	5.2	15.7	5.7	17.8	4.6k	0.2k	3.1k	19.3
proposed	32	2.8	7.4	3.1	8.0	n.a.			
	64	2.8	7.2	3.1	7.9	0.5k	0.2k	2.9k	17.8
	128	2.8	7.2	3.1	7.9	n.a.			
	5000	2.8	7.2	3.1	8.0	0.5k	0.3k	2.8k	17.8
	optimal	2.8	7.1	3.1	7.8	n.a.			

shorter transcription lengths. This indicates a major flaw in modeling, which clearly requires modeling refinement. Heuristics using length normalization and EOS threshold reduce the deletion errors and improve the results dramatically. Rather stable performance is obtained with beam sizes 32, 64 and 128 except a small degradation on the test-other set. For beam size 64, the average length of recognized transcriptions closely approaches the one of the reference transcriptions (17.8 words), which shows a good effectiveness against the length bias problem. However, a considerable performance degradation is observed with beam size 5000. The major impact comes from a large increment of insertion errors, which is also visible from the longer transcription length. This raises a new beam problem suggesting that heuristics are not proper modeling refinement.

We also conduct informal experiments to apply separate pruning between ended and ongoing sequences, and use the input length constraint to stop decoding. This gives worse results for both simple beam search with and without heuristics.

Table 2: Example transcription with scores and average number of search steps for heuristic-based and proposed beam search with different beam sizes on the LibriSpeech dev sets (BPE units are merged to words already).

Beam Search	Beam Size	Example of Recognized Transcription (utterance 1585-157660-0003)	Original Score	Final Score	Search Steps	
					dev-clean	dev-other
simple+ heuristics	64	“GLORIOUS LONDON”	-10.55	-3.52	27.0	25.6
	5000	“ZARATHUSTRA DE L’OISEAU DE L’OISEAU DE L’OISEAU DE L’OISEAU”	-111.39	-3.28	48.0	49.1
proposed	64	“GLORIOUS LONDON”	-10.55	-0.32	24.2	21.7
	5000			-0.45	24.2	21.8

3.3. Proposed beam search

For the proposed beam search, only one scale for LM shallow fusion needs to be optimized. For a fair comparison under the same beam sizes, we explicitly deactivate the score-based pruning. We optimize the LM scale for beam size 64 and apply it to all other beam sizes. The results are also shown in Table 1.

Compared with simple beam search without heuristics, the proposed approach clearly eliminates the length bias problem based on the largely reduced deletion errors and improved accuracy. Unlike the heuristic approach, this effectiveness is maintained when the beam size is increased from 64 to 5000. For both beam sizes, it produces the same average transcription length as the reference, which strongly supports the intuitive interpretation given in Section 2.3 about reliable comparison among output sequences of different lengths. In fact, consistent and good performance is obtained under all beam sizes. This suggests a more robust capability for modeling refinement, even though the approach does not provide a theoretical final solution to the length bias problem. Compared with the best WER of the heuristic baseline using beam size 64, the proposed approach achieves the same performance on the ‘clean’ sets and 4% relative improvement on the ‘other’ sets. To show the performance of the complete approach, we also include results using score-based pruning with a threshold of 8 and beam size 5000 as upper bound. Further improvement is obtained on the ‘other’ sets by using such optimal beam at each step.

3.4. Analysis

For more insights into the new beam problem of the heuristic approach, we further check those utterances of degraded performance from beam size 64 to 5000. We find out that they actually point out a robustness issue of the heuristic-based score for decision making. For better illustration, we show one example of such utterances in Table 2. We denote the score of Eq. (1) as original score and the approach-specific score as final score.

Based on the length-normalized final score, the heuristic-based beam search decides for the correct transcription with beam size 64 and a much longer transcription with beam size 5000. However, this wrong transcription actually has a much worse original score based on the models’ output. This indicates a strong bias introduced by the heuristics towards longer output sequences, which can over-correct the length bias and cause new modeling problems. With much larger beam and more hypotheses considered, this leads to more wrong decisions towards too long transcriptions. Therefore, good performance is still only achievable with rather small beam sizes and careful tuning, where search errors are adopted to cover the new modeling errors. This completely contradicts the concept of beam search. Similar effect is also possible with other heuristics such as length reward. [15] applied length reward in decoding and reported issues about looping transcriptions. This is very similar as the example shown here, which is actually resulted from the same reason.

In contrast, without introducing any artificial terms, the proposed approach gives the best final score for the correct tran-

scription under both beam sizes. This reflects the robustness of the proposed final probability for decision making, which serves as a robust model modification as described in Section 2.3. For both small and very large beam sizes, the approach solves the length bias problem without introducing additional side-effects, which also explains the performance difference in Table 1.

3.5. Efficiency

In terms of computation at each search step under the same beam size, there is not much difference between the baseline and proposed approach. However, since we do not use ended sequences to prune away ongoing sequences, we need to check if our derived stopping criterion really avoids unnecessary search steps. We verify this by comparing the average number of search steps needed to finish recognition for each dev set under beam sizes 64 and 5000. As shown in the last two columns of Table 2, the numbers needed for the simple beam search with heuristics largely increase with beam sizes. Therefore, its efficiency decreases with increasing beam sizes. On the other hand, the numbers needed for the proposed approach are consistently small under both beam sizes. This approves the derived early stopping criterion and the better efficiency of the approach.

4. Conclusion

In this work, we presented a novel beam search derived from reinterpreting the sequence posterior with an explicit length modeling. By applying the reinterpreted probability together with beam pruning, the obtained final probability leads to a robust model modification without heuristics, which allows reliable comparison among output sequences of different lengths. Experiments on the LibriSpeech corpus show that the proposed approach solves the length bias problem without heuristics or additional tuning effort. We showed that simple heuristics are not proper modeling refinement and introduce strong bias for decision making, which results in severe performance degradation with largely increased beam size. In contrast, the proposed approach provides robust decision making and consistently good performance under both small and very large beam sizes. Compared with the best WER of the heuristic baseline using small beam size as in practice, the approach achieves the same performance on the ‘clean’ sets and 4% relative improvement on the ‘other’ sets. It is also more efficient with the additional derived early stopping criterion.

Future work includes verifying the proposed approach with different data and more complicated decision rules, and extension to a more general label-synchronous search framework. It is also worthy to further research into better modeling approaches that in principle would work even without pruning and thus fully retain a proper beam search behavior.

5. Acknowledgements

This work has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No 694537, project “SEQCLAS”) and from a Google Focused Award. The work reflects only the authors’ views and none of the funding parties is responsible for any use that may be made of the information it contains.

We thank Wilfried Michel for useful discussions.

6. References

- [1] D. Bahdanau, K. Cho, and Y. Bengio, "Neural Machine Translation by Jointly Learning to Align and Translate," in *International Conference on Learning Representations (ICLR), San Diego, CA, USA, May 7-9, 2015*.
- [2] W. Chan, N. Jaitly, Q. V. Le, and O. Vinyals, "Listen, Attend and Spell: A Neural Network for Large Vocabulary Conversational Speech Recognition," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP), Shanghai, China, March 20-25, 2016*, pp. 4960–4964.
- [3] D. S. Park, W. Chan, Y. Zhang, C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition," in *Interspeech 2019, Graz, Austria, 15-19 September 2019*, pp. 2613–2617.
- [4] A. Zeyer, P. Bahar, K. Irie, R. Schlüter, and H. Ney, "A Comparison of Transformer and LSTM Encoder Decoder Models for ASR," in *Automatic Speech Recognition and Understanding Workshop (ASRU), Sentosa, Singapore, Dec. 2019*, pp. 8–15.
- [5] Z. Tüske, G. Saon, K. Audhkhasi, and B. Kingsbury, "Single Headed Attention based Sequence-to-sequence Model for state-of-the-art Results on Switchboard-300," 2020. [Online]. Available: <https://arxiv.org/abs/2001.07263>
- [6] P. Sountsov and S. Sarawagi, "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 (EMNLP), Austin, Texas, USA, November 1-4, 2016*, pp. 1516–1525.
- [7] F. Stahlberg and B. Byrne, "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), Hong Kong, China, November 3-7, 2019*. Association for Computational Linguistics, 2019, pp. 3354–3360.
- [8] J. Chorowski and N. Jaitly, "Towards Better Decoding and Language Model Integration in Sequence to Sequence Models," in *Interspeech 2017, Stockholm, Sweden, August 20-24, 2017*, pp. 523–527.
- [9] A. Hannun, A. Lee, Q. Xu, and R. Collobert, "Sequence-to-Sequence Speech Recognition with Time-Depth Separable Convolutions," 2019. [Online]. Available: <http://arxiv.org/abs/1904.02619>
- [10] J. Pouget-Abadie, D. Bahdanau, B. van Merriënboer, K. Cho, and Y. Bengio, "Overcoming the Curse of Sentence Length for Neural Machine Translation using Automatic Segmentation," in *Proceedings of SSST@EMNLP 2014, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation, Doha, Qatar, 25 October 2014*, pp. 78–85.
- [11] K. Murray and D. Chiang, "Correcting Length Bias in Neural Machine Translation," in *Proceedings of the Third Conference on Machine Translation: Research Papers, WMT 2018, Belgium, Brussels, October 31 - November 1, 2018*, pp. 212–223.
- [12] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey, J. Klingner, A. Shah, M. Johnson, X. Liu, ukasz Kaiser, S. Gouws, Y. Kato, T. Kudo, H. Kazawa, K. Stevens, G. Kurian, N. Patil, W. Wang, C. Young, J. Smith, J. Riesa, A. Rudnick, O. Vinyals, G. Corrado, M. Hughes, and J. Dean, "Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation," 2016. [Online]. Available: <http://arxiv.org/abs/1609.08144>
- [13] P. Koehn and R. Knowles, "Six Challenges for Neural Machine Translation," in *Proceedings of the First Workshop on Neural Machine Translation, NMT@ACL, Vancouver, Canada, August 4, 2017*, pp. 28–39.
- [14] A. Y. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Sathesh, S. Sengupta, A. Coates, and A. Y. Ng, "Deep Speech: Scaling up end-to-end speech recognition," 2014. [Online]. Available: <http://arxiv.org/abs/1412.5567>
- [15] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, "End-to-end Attention-based Large vocabulary speech recognition," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP), Shanghai, China, March 20-25, 2016*, pp. 4945–4949.
- [16] W. He, Z. He, H. Wu, and H. Wang, "Improved neural machine translation with SMT features," in *AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pp. 151–157.
- [17] Z. Tu, Z. Lu, Y. Liu, X. Liu, and H. Li, "Modeling Coverage for Neural Machine Translation," in *Annual Meeting of the Association for Computational Linguistics (ACL) 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*.
- [18] Z. Tüske, K. Audhkhasi, and G. Saon, "Advancing Sequence-to-Sequence Based Speech Recognition," in *Interspeech 2019, Graz Austria, September 15-19, 2019*.
- [19] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," *International Conference on Acoustics, Speech and Signal Processing (ICASSP) 2015*, pp. 5206–5210.
- [20] Ç. Gülçehre, O. Firat, K. Xu, K. Cho, L. Barrault, H. Lin, F. Bougares, H. Schwenk, and Y. Bengio, "On Using Monolingual Corpora in Neural Machine Translation," 2015. [Online]. Available: <http://arxiv.org/abs/1503.03535>
- [21] S. Wiesler, A. Richard, P. Golik, R. Schlüter, and H. Ney, "RASR/NN: The RWTH Neural Network Toolkit for Speech Recognition," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP) 2014*, pp. 3281–3285.
- [22] E. Beck, W. Zhou, R. Schlüter, and H. Ney, "LSTM Language Models for LVCSR in First-Pass Decoding and Lattice-Rescoring," 2019. [Online]. Available: <https://arxiv.org/abs/1907.01030>
- [23] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [24] P. Doetsch, A. Zeyer, P. Voigtlaender, I. Kulikov, R. Schlüter, and H. Ney, "RETURNN: the RWTH extensible training framework for universal recurrent neural networks," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP) 2017*, pp. 5345–5349.
- [25] A. Zeyer, T. Alkhoul, and H. Ney, "RETURNN as a Generic Flexible Neural Toolkit with Application to Translation and Speech Recognition," in *Annual Meeting of the Association for Computational Linguistics (ACL), System Demonstrations, Melbourne, Australia, Jul. 2018*, pp. 128–133.