

## Promising Accurate Prefix Boosting for Sequence-to-Sequence ASR

Baskar, M.K.; Burget, L.; Watanabe, S.; Karafiat, M.; Hori, T.; Cernocky, J.H.

TR2019-006 March 29, 2019

### Abstract

In this paper, we present promising accurate prefix boosting (PAPB), a discriminative training technique for attention based sequence-to-sequence (seq2seq) ASR. PAPB is devised to unify the training and testing scheme effectively. The training procedure involves maximizing the score of each partial correct sequence obtained during beam search compared to other hypotheses. The training objective also includes minimization of token (character) error rate. PAPB shows its efficacy by achieving 10.8% and 3.8% WER with and without external RNNLM respectively on Wall Street Journal dataset.

*IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.



# PROMISING ACCURATE PREFIX BOOSTING FOR SEQUENCE-TO-SEQUENCE ASR

Murali Karthick Baskar <sup>ϕ</sup>, Lukáš Burget <sup>ϕ</sup>, Shinji Watanabe <sup>π</sup>, Martin Karafiát <sup>ϕ</sup>,  
Takaaki Hori <sup>†</sup>, Jan “Honza” Černocký <sup>ϕ</sup>

<sup>ϕ</sup>Brno University of Technology, <sup>π</sup> Johns Hopkins University,  
<sup>†</sup>Mitsubishi Electric Research Laboratories (MERL)

{baskar,burget,karafiát,cernocký}@fit.vutbr.cz, shinjiw@jhu.edu, thori@merl.com

## ABSTRACT

In this paper, we present promising accurate prefix boosting (PAPB), a discriminative training technique for attention based sequence-to-sequence (seq2seq) ASR. PAPB is devised to unify the training and testing scheme effectively. The training procedure involves maximizing the score of each partial correct sequence obtained during beam search compared to other hypotheses. The training objective also includes minimization of token (character) error rate. PAPB shows its efficacy by achieving 10.8% and 3.8% WER with and without external RNNLM respectively on Wall Street Journal dataset.

**Index Terms**— Beam search training, sequence learning, discriminative training, Attention models, softmax-margin

## 1. INTRODUCTION

Sequence-to-sequence (seq2seq) modeling provides a simple framework to perform complex mapping between input and output sequence. In the original work, where seq2seq [1] model was applied to machine translation task, the model contains an encoder neural network with recurrent layers, to encode the entire input sequence (i.e the message in a source) into a internal fixed-length vector representation. This vector is an input to the decoder – another set of recurrent layers with final softmax layer, which, in each recurrent iterations, predicts probabilities for the next symbol of the output sequence (i.e. the message in the target). This work deals with the task of automatic speech recognition (ASR), where the seq2seq model is used to map a sequence of speech features into a sequence of characters. In particular, we use attention based seq2seq model [1], where the encoder encodes an input sequence into another internal sequence of the same length. The attention mechanism [1], then

---

The work reported here was carried out during the 2018 Jelinek Memorial Summer Workshop on Speech and Language Technologies, supported by Johns Hopkins University via gifts from Microsoft, Amazon, Google, Facebook, and MERL/Mitsubishi Electric. All the authors from Brno university of Technology was supported by Czech Ministry of Education, Youth and Sports from the National Programme of Sustainability (NPU II) project “IT4Innovations excellence in science - LQ1602” and by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA) MATERIAL program, via Air Force Research Laboratory (AFRL) contract # FA8650-17-C-9118. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, AFRL or the U.S. Government. Part of computing hardware was provided by Facebook within the FAIR GPU Partnership Program. We thank Ruizhi Li, for finding the hyper-parameters to obtain best baseline in WSJ. We also thank Hiroshi Seki, for providing the batch-wise beam search decoding implementation in ESPnet.

focuses on the relevant portion of the internal sequence in order to predict each next output symbol using the decoder. The seq2seq are typically trained to maximize the conditional likelihood (or minimize cross-entropy) of the correct output symbols. For predicting a current character, the previous character (e.g. its one-hot encoding) from ground truth sequence is typically fed as an auxiliary input to decoder during training. This so-called teacher-forcing [2] helped the decoder to learn an internal language model (LM) for the output sequences. During normal decoding, the last predicted character is fed back instead of the unavailable ground truth. Using such training strategy, attention [3] based seq2seq model has shown to absorb and jointly learn all the components of a traditional ASR (i.e. acoustic model, lexicon and language model). Two major drawbacks have been, however, identified with the training strategy described above:

- *Exposure bias*: The seq2seq training uses teacher forcing, where each output character is conditioned on the previous true character. However during testing, the model needs to rely on its own previous predictions. This mismatch between training and testing leads to poor generalization and is referred to as exposure bias [4, 5].
- *Error criterion mismatch*: Another potential issue is mismatch in error criterion between training and testing [6, 7]. ASR, uses character error rate (CER) or word error rate (WER) to validate the decoded output while the training objective is the conditional maximum likelihood (cross entropy) maximizing the probability of the correct sequence.

In this work, we first experiment with training objectives that better matches the CER or WER metric, namely minimum Bayes risk (MBR) [7] and softmax margin [8]. We show that such choice of training objective makes teacher-forcing strategy unnecessary and therefore effectively addresses both the aforementioned problems.

Both MBR and softmax margin objective needs to consider alternative sequences (hypotheses) besides the ground truth sequence. Unfortunately, seq2seq model does not make Markov assumptions and the alternative sequences cannot be efficiently represented with a lattice. Instead, we perform beam search to generate an (approximate) N-best list of alternative sequences. However, with the limited capacity of the N-best representation, some of the important hypotheses (i.e. sequences with a low error rate) can be easily pruned out by the beam search, which might result in less effective training. To address this problem, we propose a new training strategy, which we call *promising accurate prefix boosting (PAPB)*: The beam search keeps list of  $N$  promising prefixes (partial sequences) of the output sequence, which get extended by one character at each decoding iteration. In each iteration, we update parameters of the seq2seq model

to boost probabilities of such promising prefixes that are also accurate (i.e. partial sequences with low edit distance to partial ground truth). This is accomplished by using the softmax margin objective (and updates) not only for the whole sequences, but also for all the partial sequence obtained during the decoding.

There are existing works addressing the exposure bias or the error criterion mismatch problem with seq2seq models applied to natural language processing (NLP) problem. For example, scheduled sampling [5] and SEARN [4] handle the exposure bias by choosing either the model predictions or the true labels as the feedback to the decoder. The error criterion mismatch is handled using task loss minimization [9] using an edit-distance, RNN transducer [10] based expected error rate minimization, and minimum risk criterion based recurrent neural aligner [6]. Few works consider both the problems simultaneously: learning from character sampled from the model using reinforcement learning [11, 12] and actor-critic algorithm [13]. Our work is mostly inspired by beam search optimization (BSO) [14], where max-margin loss is used as sequence-level objective [15] for the machine translation task. All the mentioned works were applied to NLP problems, while the focus of this work is ASR. Also, none of the works considered the prefixes (partial sequences) during training. A recent work on seq2seq based ASR was trained with MBR objective [7] using N-best hypotheses obtained from a beam search. However, this work also did not consider the prefixes. Finally, optimal completion distillation [16] technique focuses on prefix learning, but it uses complex learning methods such as policy distillation and imitation learning.

## 2. ENCODER-DECODER

With the attention based Encoder-Decoder [3] architecture, the encoder  $H = \text{enc}(X)$  neural network provides an internal representations  $H = \{h_t\}_{t=1}^T$  of an input sequence  $X = \{x_t\}_{t=1}^T$ , where  $T$  is the number of frames in an utterance. In this work, the encoder is a recurrent network with bi-directional long short-term memory (BLSTM) layers [17, 18]. To predict the  $l$ -th output symbol, the attention component takes the sequence  $H$  and the previous hidden state of the decoder  $q_{l-1}$  as the input and produces per-frame attention weights

$$\{a_{lt}\}_{t=1}^T = \text{Attention}(q_{l-1}, H). \quad (1)$$

In this work, we use location aware attention [19]. The attention weights are expected to have high values for the frames that we need to pay attention to for predicting the current output and are typically normalized to sum-up to one over frames. Using such weights, the weighted average of the internal sequence  $H$  serves as an attention summary vector

$$r_l = \sum_t a_{lt} h_t. \quad (2)$$

The decoder is a recurrent network with LSTM layers, which receives  $r_l$  along with the previous predicted output character  $y_{l-1}$  (e.g. its one-hot encoding) as the input and estimates the hidden state vector

$$q_l = \text{dec}(r_l, q_{l-1}, y_{l-1}). \quad (3)$$

This vector is further subject to an affine transformation (LinB) and Softmax non-linearity to obtain the probabilities for the current output symbol  $y_l$ :

$$s_l = \text{LinB}(q_l) \quad (4)$$

$$p(y_l|y_{1:l-1}, X) = \text{Softmax}(s_l) \quad (5)$$

The probability of a whole sequence  $y = \{y_l\}_{l=1}^L$  is

$$p(y|X) = \prod_l p(y_l|y_{1:l-1}, X) \quad (6)$$

To decode the output sequence, simple greedy search can be performed, where the most likely symbol is chosen according to (5) in each decoding iteration until the dedicated end-of-sentence symbol is decoded. This procedure, however, does not guarantee in finding the most likely sequence according to (6). To find the optimal sequence, multiple hypotheses explored by beam search usually provides better results. Note, however, that each partial hypothesis in the beam search has its own hidden state (3) as it depends on the previously predicted symbols  $y_l$  in that hypothesis.

During training, model parameters are typically updated to minimize cross-entropy (CE) loss for correct output  $y^*$ :

$$\mathcal{L}_{CE} = -\log p(y^*|X) = -\sum_{l=1} \log p(y_l^*|y_{1:l-1}^*, X). \quad (7)$$

This is particularly easy with the teacher forcing, when the symbol from the ground truth sequence is always used in (7) as the previously predicted symbol and therefore no alternative hypotheses needs to be considered.

## 3. TRAINING CRITERION

We compare our proposed PAPB approach with two other objective functions that serves as our baseline. Namely, we use minimum Bayes risk criterion and softmax margin loss, which both perform sequence level discriminative training. Both objectives need to estimate character error rate (CER) for alternative hypotheses, which are explored using beam search. In the following equations, the symbol  $\text{cer}(y^*, y)$  denotes the edit distance between the ground truth sequence  $y^*$  and hypothesized sequence  $y$ .

### 3.1. Minimum Bayes risk (MBR)

In minimum Bayes risk [20, 21, 22, 23], the expectation of character error rates  $\text{cer}(y^*, y)$  is taken w.r.t model distribution  $p(y|X)$ :

$$\mathcal{L}_{MBR} = E_{p(y|X)} [\text{cer}(y^*, y)] = \sum_{y \in Y} p(y|X) \text{cer}(y^*, y), \quad (8)$$

In practice, the total set of hypotheses  $Y$  generated is reduced to  $N$ -best hypotheses  $Y_N$  for computational efficiency. MBR training objective effectively performs sequence level discriminative training in ASR [22] and provides substantial gains when used as secondary objective after performing cross-entropy loss based optimization [20].

### 3.2. Softmax margin loss (SM)

Softmax margin loss [8] falls under the category of maximum-margin [24] classification technique. It is a generalization of boosted maximum mutual information (b-MMI) [21].

$$\mathcal{L}_{SM} = -s(y^*, X) + \log \left( \sum_{y \in Y} \exp(s(y, X) + \alpha \text{cer}(y^*, y)) \right), \quad (9)$$

where  $\alpha$  is a tunable margin factor ( $\alpha = 1$ ) and the un-normalized score of a chosen sequence,

$$s(y, X) = \sum_l^L s_l \quad (10)$$

is the sum of the pre-softmax outputs  $s_l$  from (4). Note the dependence of the scores  $s_l$  on the chosen hypothesis  $y$  through the predictions fed back to decoder in (3), which is not explicitly denoted in our notation. The function aims to boost the score of the true sequence,  $s(y^*, X)$ , to stay above the other hypotheses  $s(y, X)$  with a margin defined by CER of the alternative hypotheses.

#### 4. PROMISING ACCURATE PREFIX BOOSTING (PAPB)

In PAPB, we perform training at prefix level in similar fashion to decoding, by incorporating an appropriate training objective with beam search. The primary motivation to carry out prefix level training is because, seq2seq models predict a sequence, character by character. MBR aims to improve the score of the completed hypothesis with less error, but it might get pruned out during the beam search. However, in our approach, the model will be exposed to all prefixes  $y_{1:l}$  obtained from N-best as generated by beam search and optimized to maximize the scores of true hypothesis  $y_{1:l}^*$ . In brief, we consider not only the fully completed hypotheses, but also prefixes so that the promising prefixes with low error keep scoring high and therefore are likely to survive the pruning. The  $\mathcal{L}_{PAPB}$  loss is computed for each prefix by modifying the softmax margin loss  $\mathcal{L}_{SM}$  as:

$$\mathcal{L}_{PAPB} = \sum_{l=1}^L \left( -s(y_{1:l}^*, X) + \log \sum_{y \in Y_N} \{\exp [s(y_{1:l}, X) + B]\} \right) \quad (11)$$

where  $B = \text{cer}(y_{1:l}^*, y_{1:l})$  and  $Y_N$  denotes the  $N$ -best set hypothesis obtained using beam search. In equation (11), the prefix scores of predicted hypotheses  $s(y_{1:l}, X)$  and true hypothesis  $s(y_{1:l}^*, X)$  are computed by summing the scores  $s_l$  given by (4) from 1 to  $l$ , while, in the standard sequence objective  $\mathcal{L}_{SM}$ , the summation is performed only across a whole sequence as noted in (10). The contributions of our proposed approach are as follows:

- The output scores  $s(y_l, X)$  (as in (4)) are computed for each character conditioned on the previous character from the corresponding explored hypothesis (i.e. no teacher-forcing used).
- In our experiments, we select the hypothesis from N-best that obtains the lowest CER as the pseudo-true hypothesis  $y_{1:l}^*$  to compute the score  $s(y_{1:l}^*, X)$ , instead of using the true hypothesis. This is to avoid harmful effects during model training by abruptly including the true sequence  $y_{1:l}$  into the beam, which might have very small score. Defining the true hypothesis with a pseudo-true hypothesis brings our objective analogous to MBR criterion where very unlikely hypotheses do not affect model parameter updates.
- The  $\text{cer}(y_{1:l}^*, y_{1:l})$  is calculated using edit-distance between the prefixes  $y_{1:l}^*$  and  $y_{1:l}$ . Here, the number of characters are kept equal between true prefix  $y_{1:l}^*$  and prefix hypothesis  $y_{1:l}$ , which, according to our assumption, should contribute to reduction of insertion and deletion errors.

#### 5. EXPERIMENTAL SETUP

*Database details:* Voxforge-Italian [25] and Wall Street Journal (WSJ) [26] corpora were used for our experimental analysis. Voxforge-Italian is a broadband speech corpus (16 hours) and is

split into 80%, 10% and 10% to training, development, and evaluation sets by ensuring that no sentence was repeated in any of the sets. WSJ with 284 speakers comprising 37,416 utterances (82 hours of speech) is used for training, and eval92 test set is used for decoding.

*Training:* Filter-bank features containing 83 dimensional (80 Mel-filter bank coefficients plus 3 pitch features) coefficients are used as input. In this work, the encoder-decoder model is aligned and trained using attention based approach. Location aware attention [19] is used in our experiments. For WSJ experiments, the encoder comprises 3 bi-directional LSTM layers [18, 17] each with 1024 units and the decoder comprises 2 (uni-directional) LSTM layers with 1024 units. For VoxForge experiments, the encoder comprises 3 bi-directional LSTM layers with 320 units and the decoder contains one LSTM layer with 320 units. The CE training is optimized using AdaDelta [27] optimizer with an initial learning rate set to 1.0. The training batch size is 30 and the number of training epochs is 20. The learning rate decay is based on the validation performance computed using the character error rate (min. edit distance). ESPnet [28] is used to implement and execute all our experiments. The MBR, softmax margin and prefix training configuration has initial learning rate 0.01, the number of training epochs is set to 10 and the batch-size to 10. The beam-size for training and testing is set to 10. The model weights are initialized with pre-trained CE model. The rest of configuration is kept the same as for CE training. In our experiments, we use a modified MBR objective:

$$\mathcal{L}'_{MBR} = \mathcal{L}_{MBR} + \lambda \mathcal{L}_{CE}, \quad (12)$$

which is a weighted combination of the original MBR objective (8) and CE objective (12). Adding the CE objective is analogous to f-smoothing [20] and provides gains when applied for seq2seq models [7]. Similarly, we also use a modified prefix boosting objective

$$\mathcal{L}'_{PAPB} = \mathcal{L}_{PAPB} + \lambda \mathcal{L}_{CE} \quad (13)$$

where  $\lambda$  is the CE objective weight empirically set to 0.001 for both MBR (also noted in [7]) and prefix training experiments. Altering the  $\lambda$  did not show much difference in performance.

*External language model for WSJ:* Beside the internal language train by the decoder, an external RNN language model (RNNLM) [29] is also experimented with by training on the text used for seq2seq training along with additional text data accounting to 1.6 million utterances from WSJ corpus. Both character and word-level language models are used in our experiments. The vocabulary size is 50 for character LM and 65k for word LM. The word level RNNLM is trained using 1 recurrent layer of 1000 units, 300 batch size and SGD optimizer. The character level RNNLM configuration contains 2 recurrent layers of 650 units, 1024 batch size and uses Adam [30].

## 6. RESULTS AND DISCUSSION


Initial investigation on the proposed approach is performed with Voxforge-Italian dataset and later tested on WSJ.

### 6.1. Comparison with scheduled sampling (SS)

Our best performing CE baseline model is with 50% SS (which denotes 50 % true labels) as mentioned in the first row in Table 1. The SS-50% model is compared with SS-0% (0% true labels) to investigate the impact of only feeding model predictions. The second row shows that WER of SS-0% degrades by 8.2 % on test set and by 5% on dev set compared to SS-50% model. The SS-0% re-trained using weights initialized from SS-50% model (acts as prior), still resulted

**Table 1.** Comparison of recognition performance between scheduled sampling (SS), MPE and PAPB on Voxforge-Italian

	%WER	Test	Dev
SS-50% (from random init.)		50.9	52.3
SS-0% (from random init.)		59.1	57.3
SS-0% (fine-tuned from SS-50%)		52.9	52.5
MBR		49.9	50.8
Softmargin		50.1	50.4
PAPB		47.4	47.7



**Fig. 1.** Changes in character error rate (CER) during training with different criterion’s on dev set of Voxforge-Italian dataset. The plot shows, PAPB improves over both softmax margin loss and MBR objectives

in performance degradation but the gap got reduced to 2.0% on test set and 0.2% on dev set. These results highlight the limitation of using scheduled sampling with 0% true labels (or 100% model predictions) as it lead to loss in recognition performance. Thus, a need to use a specific objective which can train only with model predictions is necessary and justifies the focus of this paper.

## 6.2. Comparison of PAPB with MBR

Table 1 shows that the performance of both MBR and softmax margin loss objectives are comparable to each other. While MBR and softmax margin loss provide considerable gains over scheduled sampling, they do not consider the prefixes (partial sequences generated during beam search) for training. In the following experiment, we show that the performance of PAPB justify our intuition to use prefix information by providing improvement from 49.9 % to 47.4 % WER on test set and from 50.8 % to 47.7 % WER on dev set compared to MBR objective. PAPB shows an improvement of 2.7 % WER for both test and dev sets compared to softmax margin loss. Figure 1 also shows that PAPB allows smoother convergence during training, and gaining better CER over MBR and softmax margin objectives.

## 6.3. Effect of varying beam-size during training and testing

Further analysis on prefix training method is performed to understand the impact of beam-size used during training and testing. The beam-size decides the number of hypotheses to retain during beam search and is denoted as N-best. The results obtained by varying this hyper-parameter showcases the importance of using multiple hypotheses in the loss objective. Table 2 introduces the effect of retain-

**Table 2.** Comparison of recognition performance between different beam sizes obtained during training  $N_{tr}$  and decoding  $N_{de}$  for prefix training on Voxforge-Italian.

$N_{tr}$	% WER		
	$N_{de}$		
	2	5	10
2	50.4	51.1	51.5
5	50.8	49.3	48.8
10	51.3	47.9	<b>47.4</b>

**Table 3.** % CER and %WER on WSJ corpus for test set with and without LM.

LM weight	LM type	CE		MBR		PAPB	
		%CER	%WER	%CER	%WER	%CER	%WER
0	-	4.6	12.9	4.3	11.5	<b>4.0</b>	<b>10.8</b>
0.1	char.	4.6	11.2	4.3	10.1	4.0	9.9
0.2	char.	4.5	10.9	4.1	9.9	3.9	9.1
1.0	char.	2.5	5.8	2.5	5.4	2.1	4.5
1.0	word	2.0	4.8	2.1	4.3	<b>2.0</b>	<b>3.8</b>

ing best paths (2,5, and 10) during training  $N_{tr}$  and testing  $N_{de}$ . A noticeable pattern observed in our experiments is that increasing the beam-size led to significant improvement in performance. Further, increase in beam-size did not provide considerable gains.

## 6.4. Results on WSJ

The results in Table 3 showcase the importance of using character level, word level RNNLM over no RNNLM. For decoding with RNNLM, we use look-ahead word LM decoding procedure recently introduced in [31] to integrate the word based RNNLM and the character RNNLM is decoded by following the procedure in [32]. The LM weight is optimized to show the impact of language model across CE, MBR and our proposed PAPB models. The Table 3 also show that results of PAPB and MBR shows complementary effect with both word and character LM. The PAPB and MBR shows improvement in %WER but not in %CER is due to the effect of including word-level LM and requires further investigation.

*State-of-the-art results on WSJ:* Without using external LM deep CNN [33] achieves 10.5% WER and 9.6% WER using OCD [16]. OCD’s nice performance is the resultant effect of their strong baseline with 10.6% WER. 4.1% WER are obtained with word RNNLM using end-to-end LF-MMI [34].

## 7. CONCLUSION

In this paper, we proposed PAPB, a strategy to train on N-best partial sequences generated using beam search. This method suggests that improving the hypothesis at prefix level can attain better model predictions for refining the feed back to predict next character. The softmax margin loss function is inherited in our approach to serve this purpose. The experimental results shows the efficacy of the proposed approach compared to CE and MBR objectives with consistent gains across two datasets. The PAPB also has its drawbacks in-terms of time complexity, as it consumes 20% more training time compared to CE training. This work can be further extended to use complete set of lattices instead of N-best list by exploiting the capabilities of GPU for improving time complexity. Also, modified MBR training in-place of softmax margin can be used to learn prefixes.

## 8. REFERENCES

- [1] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *arXiv preprint arXiv:1409.0473*, 2014.
- [2] R. J. Williams and D. Zipser, "A learning algorithm for continually running fully recurrent neural networks," *Neural computation*, vol. 1, no. 2, pp. 270–280, 1989.
- [3] J. Chorowski, D. Bahdanau, K. Cho, and Y. Bengio, "End-to-end continuous speech recognition using attention-based recurrent NN: first results," *arXiv preprint arXiv:1412.1602*, 2014.
- [4] H. D. III, J. Langford, and D. Marcu, "Search-based structured prediction," *CoRR*, vol. abs/0907.0786, 2009.
- [5] S. Bengio, O. Vinyals, N. Jaitly, and N. Shazeer, "Scheduled sampling for sequence prediction with recurrent neural networks," in *Advances in Neural Information Processing Systems*, pp. 1171–1179, 2015.
- [6] H. Sak, M. Shannon, K. Rao, and F. Beaufays, "Recurrent neural aligner: An encoder-decoder neural network model for sequence to sequence mapping," in *Proc. Interspeech*, pp. 1298–1302, 2017.
- [7] R. Prabhavalkar, T. N. Sainath, Y. Wu, P. Nguyen, Z. Chen, C.-C. Chiu, and A. Kannan, "Minimum word error rate training for attention-based sequence-to-sequence models," in *ICASSP, 2018*, pp. 4839–4843, IEEE, 2018.
- [8] K. Gimpel and N. A. Smith, "Softmax-margin training for structured log-linear models," 2010.
- [9] D. Bahdanau, D. Serdyuk, P. Brakel, N. R. Ke, J. Chorowski, A. Courville, and Y. Bengio, "Task loss estimation for sequence prediction," *arXiv preprint arXiv:1511.06456*, 2015.
- [10] A. Graves and N. Jaitly, "Towards end-to-end speech recognition with recurrent neural networks.," in *ICML*, vol. 14, pp. 1764–1772, 2014.
- [11] M. Ranzato, S. Chopra, M. Auli, and W. Zaremba, "Sequence level training with recurrent neural networks," *arXiv preprint arXiv:1511.06732*, 2015.
- [12] S. Karita, A. Ogawa, M. Delcroix, and T. Nakatani, "Sequence training of encoder-decoder model using policy gradient for end-to-end speech recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5839–5843, IEEE, 2018.
- [13] D. Bahdanau, P. Brakel, K. Xu, A. Goyal, R. Lowe, J. Pineau, A. Courville, and Y. Bengio, "An actor-critic algorithm for sequence prediction," *arXiv preprint arXiv:1607.07086*, 2016.
- [14] S. Wiseman and A. M. Rush, "Sequence-to-sequence learning as beam-search optimization," *arXiv preprint arXiv:1606.02960*, 2016.
- [15] I. Tsochantaridis, T. Joachims, T. Hofmann, and Y. Altun, "Large margin methods for structured and interdependent output variables," *Journal of machine learning research*, vol. 6, no. Sep, pp. 1453–1484, 2005.
- [16] S. Sabour, W. Chan, and M. Norouzi, "Optimal completion distillation for sequence learning," *arXiv preprint arXiv:1810.01398*, 2018.
- [17] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [18] M. Schuster and K. K. Paliwal, "Bidirectional recurrent neural networks," *IEEE Transactions on Signal Processing*, vol. 45, no. 11, pp. 2673–2681, 1997.
- [19] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, "End-to-end attention-based large vocabulary speech recognition," in *ICASSP, 2016*, pp. 4945–4949, IEEE, 2016.
- [20] D. Povey and P. C. Woodland, "Minimum phone error and i-smoothing for improved discriminative training," in *ICASSP, 2002*, vol. 1, pp. I–105, IEEE, 2002.
- [21] D. Povey, D. Kanevsky, B. Kingsbury, B. Ramabhadran, G. Saon, and K. Visweswariah, "Boosted MMI for model and feature-space discriminative training," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4057–4060, IEEE, 2008.
- [22] K. Vesely, A. Ghoshal, L. Burget, and D. Povey, "Sequence-discriminative training of deep neural networks.," in *INTER-SPEECH*, pp. 2345–2349, 2013.
- [23] H. Su, G. Li, D. Yu, and F. Seide, "Error back propagation for sequence training of context-dependent deep networks for conversational speech transcription," in *ICASSP, 2013*, pp. 6664–6668, IEEE, 2013.
- [24] C. H. Lampert, "Maximum margin multi-label structured prediction," in *Advances in Neural Information Processing Systems 24 (J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, eds.)*, pp. 289–297, Curran Associates, Inc., 2011.
- [25] "Voxforge.org, "Free speech recognition"." <http://www.voxforge.org/>. Accessed: 2014-06-25.
- [26] D. B. Paul and J. M. Baker, "The design for the Wall Street Journal-based CSR corpus," in *Proc. of the workshop on Speech and Natural Language*, pp. 357–362, Association for Computational Linguistics, 1992.
- [27] M. D. Zeiler, "ADADELTA: an adaptive learning rate method," *CoRR*, vol. abs/1212.5701, 2012.
- [28] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen, et al., "Espnet: End-to-end speech processing toolkit," *arXiv preprint arXiv:1804.00015*, 2018.
- [29] T. Mikolov, M. Karafiát, L. Burget, J. Černocký, and S. Khudanpur, "Recurrent neural network based language model," in *Eleventh Annual Conference of the International Speech Communication Association*, 2010.
- [30] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014.
- [31] T. Hori, J. Cho, and S. Watanabe, "End-to-end speech recognition with word-based RNN language models," *arXiv preprint arXiv:1808.02608*, 2018.
- [32] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, "Hybrid CTC/attention architecture for end-to-end speech recognition," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [33] Y. Zhang, W. Chan, and N. Jaitly, "Very deep convolutional networks for end-to-end speech recognition," in *ICASSP, 2017*, pp. 4845–4849, IEEE, 2017.
- [34] H. Hadian, H. Sameti, D. Povey, and S. Khudanpur, "End-to-end speech recognition using lattice-free MMI," *Proc. Interspeech 2018*, pp. 12–16, 2018.