

ATTENTION-BASED END-TO-END SPEECH RECOGNITION ON VOICE SEARCH

Changhao Shan^{1,2}, Junbo Zhang², Yujun Wang², Lei Xie¹

¹Shaanxi Provincial Key Laboratory of Speech and Image Information Processing,
School of Computer Science, Northwestern Polytechnical University, Xi'an, China

²Xiaomi Inc., Beijing, China

{shanchanghao, zhangjunbo, wangyujun}@xiaomi.com, lxie@nwpu.edu.cn

ABSTRACT

Recently, there has been a growing interest in end-to-end speech recognition that directly transcribes speech to text without any predefined alignments. In this paper, we explore the use of attention-based encoder-decoder model for Mandarin speech recognition on a voice search task. Previous attempts have shown that applying attention-based encoder-decoder to Mandarin speech recognition was quite difficult due to the logographic orthography of Mandarin, the large vocabulary and the conditional dependency of the attention model. In this paper, we use character embedding to deal with the large vocabulary. Several tricks are used for effective model training, including L2 regularization, Gaussian weight noise and frame skipping. We compare two attention mechanisms and use attention smoothing to cover long context in the attention model. Taken together, these tricks allow us to finally achieve a character error rate (CER) of 3.58% and a sentence error rate (SER) of 7.43% on the MiTV voice search dataset. While together with a trigram language model, CER and SER reach 2.81% and 5.77%, respectively.

Index Terms— automatic speech recognition, end-to-end speech recognition, attention model, voice search

1. INTRODUCTION

Voice search (VS) allows users to acquire information by a simple voice command. It has become a dominating function on various devices such as smart phones, speakers and TVs, etc. Automatic speech recognition (ASR) is the first step for a voice search task and thus its performance highly affects the user experience.

Deep Neural Networks (DNNs) have shown tremendous success and are widely used in ASR, usually in combination with Hidden Markov Models (HMMs) [12, 9, 21]. These systems are based on a complicated architecture with several separate components, including acoustic, phonetic and language models, which are usually trained separately, each with a different objective. Recently, some *end-to-end* neural network ASR approaches, such as connectionist temporal classification (CTC) [11, 1, 17] and attention-based encoder-decoder [6, 8, 3, 4, 5, 14, 23, 7], have emerged. These end-

to-end trained systems directly map the input acoustic speech to grapheme (or word) sequences and the acoustic, pronunciation, and language modeling components are trained jointly in a single system.

Attention-based models have become increasingly popular and with delightful performances on various sequence-to-sequence tasks, such as machine translation [2], text summarization [20], image captioning [22] and speech recognition. In speech recognition, the attention-based approaches usually consist of an encoder network, which maps the input acoustic speech into a higher-level representation, and an attention-based decoder that predicts the next output symbol conditioned on the sequence of previous predictions. A recent comparison of sequence-to-sequence models for speech recognition [19] has shown that Listen, Attend and Spell (LAS) [4], a typical attention-based approach, offered improvements over other sequence-to-sequence models.

Attention-based encoder-decoder performs considerably well in English speech recognition [7] and many attempts have been proposed to further optimize the model [8, 23]. However, applying attention-based encoder-decoder to Mandarin was found quite problematic. In [5], Chan *et al.* have pointed out that the attention model is difficult to converge with Mandarin data due to the logographic orthography of Mandarin, the large vocabulary and the conditional dependency of the attention model. They have proposed a joint Mandarin Character-Pinyin model but with limited success: the character error rate (CER) is as high as 59.3% on GALE broadcast news corpus. In this paper, we aim to improve the LAS approach for Mandarin speech recognition on a voice search task. Instead of using joint Character-Pinyin model, we directly use Chinese characters as network output. Specifically, we map the one-hot character representation to an embedding vector via a neural network layer. We also use several tricks for effective model training, including L2 regularization [13], Gaussian weight noise [15] and frame skipping [18]. We compare two attention mechanisms and use attention smoothing to cover long context in the attention model. Taken together, these tricks allow us to finally achieve a promising result on a Mandarin voice search task.

2. LISTEN, ATTEND AND SPELL

Listen, Attend and Spell (LAS) [4] is an attention-based encoder-decoder network which is often used to deal with variable-length input to variable-length output mapping problems. The encoder (the Listen module) extracts a higher-level feature representation (i.e., an embedding) from the input features. Then the attention mechanism (the Attend module) determines which encoder features should be attended in order to predict the next output symbol, resulting in a context vector. Finally, the decoder (the Spell module) takes the attention context vector and an embedding of the previous prediction to generate a prediction of the next output.

Specifically, in Fig. 1, the encoder we used is a bidirectional long short term memory (BLSTM) recurrent neural network (RNN) that generates a high-level feature representation sequence $\mathbf{h} = (h_1, \dots, h_T)$ from the input time-frequency representation sequence of speech \mathbf{x} :

$$\mathbf{h} = \text{Listen}(\mathbf{x}) \quad (1)$$

In Fig. 2, the *AttendAndSpell* is an attention-based transducer:

$$p(\mathbf{y}|\mathbf{x}) = \text{AttendAndSpell}(\mathbf{y}, \mathbf{h}). \quad (2)$$

In practice, the process predicts the character y_i at a time according to the probability distribution:

$$p(y_i|\mathbf{x}, y_{i-1}, \dots, y_1) = \text{CharacterDist}(s_i, c_i), \quad (3)$$

where s_i is an LSTM hidden state for time i , computed by

$$s_i = \text{DecodeRNN}([y_{i-1}, c_{i-1}], s_{i-1}), \quad (4)$$

and the context vector

$$c_i = \text{AttentionContext}(s_i, \mathbf{h}). \quad (5)$$

The *DecodeRNN* is a unidirectional LSTM RNN which produces a transducer state s_i and the *AttentionContext* generates context c_i with a multi-layer perceptron (MLP) attention network. Finally, the probability distribution *CharacterDist* is computed by a softmax function.

3. METHODS

In this section, we detail the tricks we used in LAS-based speech recognition for the Mandarin voice search task.

3.1. Embedding and regularization

Chinese has a large set of characters and even the number of frequently-used characters can reach 3,500. Chan *et. al.* [5] have pointed out that the large vocabulary with limited training data made the model difficult to learn and generalize well. This means that, for an end-to-end Mandarin system that directly outputs characters, it is critical to use an appropriate embedding to ensure the converge of the model.

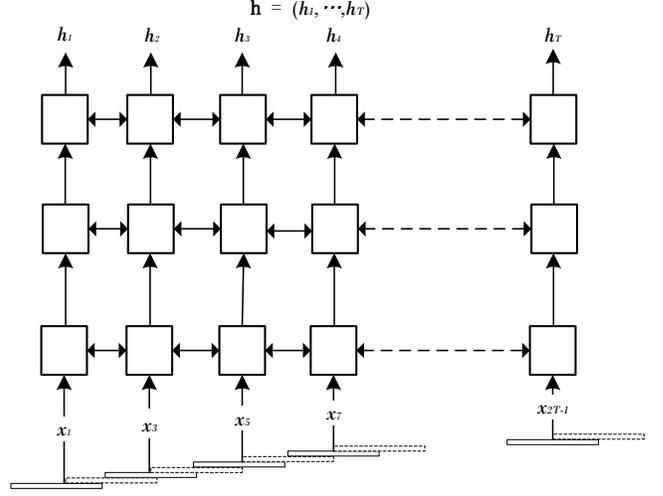


Fig. 1. The encoder model is a BLSTM that extracts \mathbf{h} from input \mathbf{x} . Frame skipping is employed during training.

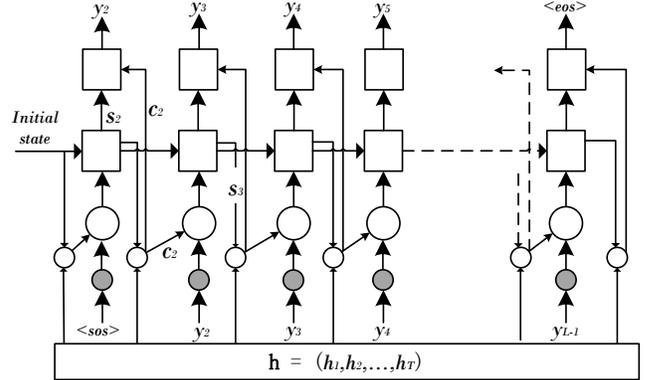


Fig. 2. The *AttendAndSpell* model composed by MLP (the Attention mechanism) and LSTM (the Decoder model).

In this paper, we first represent each character in a one-hot scheme and further embed it to a vector using neural network. Specifically, in Fig. 2, a fully-connected embedding layer (shaded circle) is used to connect the one-hot input and the subsequent BLSTM layer of the LAS encoder. The weight matrix \mathbf{W}_e of the character embedding layer is updated in the whole LAS model training procedure. The embedding layer works as follows. Assume the size of the vocabulary is n and the dimension of the embedding layer is m . Then the weight matrix \mathbf{W}_e is of size $n \times m$. When the character's index is i , the embedding layer will pass the i th row of \mathbf{W}_e to the subsequent encoder. That is, it acts as a look-up-table, making the training procedure more efficient. We find that this simple character embedding provides significant benefit to the model convergence and robustness.

The LAS model often gives poor generalization to new data without regularizations. Thus two popular regularization tricks are used in this paper: L2 regularization and Gaussian weight noise [13, 15].

3.2. Attention mechanism

The attention mechanism selects (or weights) the input frames to generate the next output element. In this study, we compared the content-based attention and the location-based attention.

Content-based attention: Borrowed from neural machine translation [2], content-based attention can be directly used in speech recognition. Here, the context vector c_i is computed as a weighted sum of h_j :

$$c_i = \sum_{j=1}^T \alpha_{i,j} h_j. \quad (6)$$

The weight $\alpha_{i,j}$ of each h_j is computed by

$$\alpha_{i,j} = \exp(e_{i,j}) / \sum_{j=1}^T \exp(e_{i,j}), \quad (7)$$

where

$$e_{i,j} = \text{Score}(s_{i-1}, h_j). \quad (8)$$

Here the *Score* is an MLP network which measures how well the inputs around position j and the output at position i match. It is based on the LSTM hidden state s_{i-1} and h_j of the input sentence. Specifically, it can be further described by

$$e_{i,j} = \mathbf{w}^\top \tanh(\mathbf{W}s_{i-1} + \mathbf{V}h_j + \mathbf{b}), \quad (9)$$

where \mathbf{w} and \mathbf{b} are vectors, and \mathbf{W} and \mathbf{V} are matrices.

Location-based attention: In [8], location-awareness was added to the attention mechanism to better fit the speech recognition task. Specifically, the content-based attention mechanism is extended by making it take into account the alignment at the previous step. k vectors $\mathbf{f}_{i,j}$ are extracted for every position j of the previous alignment α_{i-1} by convolving it with a matrix \mathbf{F} :

$$\mathbf{f}_i = \mathbf{F} * \alpha_{i-1}. \quad (10)$$

By adding $\mathbf{f}_{i,j}$, the scoring mechanism is changed to

$$e_{i,j} = \mathbf{w}^\top \tanh(\mathbf{W}s_{i-1} + \mathbf{V}h_j + \mathbf{U}\mathbf{f}_{i,j} + \mathbf{b}). \quad (11)$$

3.3. Attention smoothing

We found that long context information is important for the voice search task. Hence we explore attention smoothing to get longer context in the attention mechanism. When the input sequence \mathbf{h} is long, the α_i distribution is typically very sharp on convergence, and thus it focuses on only a few frames of \mathbf{h} . To keep the diversity of the model, similar to [8], we replace the softmax function in Eq. (7) with the logistic sigmoid σ :

$$\alpha_{i,j} = \sigma(e_{i,j}). \quad (12)$$

3.4. Frame skipping

Frame skipping is a simple-but-effective trick that has been previously used for fast model training and decoding [18]. As training BLSTM is notoriously time-consuming, we borrow this idea in the training of LAS encoder which is BLSTM. As our task does not consider online decoding, we use all frames to generate the context \mathbf{h} during decoding.

3.5. Language model

At each time step, the decoder generates a character depending on the previous ones, similar to the mechanism of a language model (LM). Therefore, the attention model works pretty good without using any explicit language model. However, the model itself is insufficient to learn a complex language model [3]. Hence we build a character-level language model T from a word-level language model G that is trained using the training transcripts and a lexicon L that simply spells out the characters of each word. In other words, the input of L is characters and output is words. More specifically, we build a finite state transducer (FST) $T = \min(\det(L \circ G))$ to calculate the log-probability for the character sequences. We add T to the cost of decoder's output:

$$C = - \sum_i [\log p(y_i | \mathbf{x}, y_{i-1}, \dots, y_1) + \gamma T] \quad (13)$$

During decoding, we minimize the cost C which combines the attention-based model and the external language model with a tunable parameter γ .

4. EXPERIMENTS

4.1. Data

We used a 3000-hour dataset for LAS model training, which contains approximately 4M voice search utterances, collected from the microphone on the MiTV remote controller. The dataset was composed of diverse search entries on popular TV programs, movies, songs and personal names (e.g. movie stars). The test set and held-out validation set were also from the MiTV voice search and each was composed of 3,000 utterances. As input features, we used 80 Mel-scale filterbank coefficients computed every 10ms with delta and delta-delta acceleration coefficients. Mean and variance normalization was conducted for each speaker. For the decoder model, we used 6,925 labels: 6,922 common Chinese characters, unknown token and sentence start and end tokens (< sos > / < eos >).

4.2. Training

We trained LAS models, in which the encoder was a 3-layer BLSTM with 256 LSTM units per-direction (or 512 in total) and the decoder was a 1-layer LSTM with 256 LSTM units. All the weight matrices were initialized with the normalized

Table 1. Results of our attention-based models with a beam size of 30, $\tau = 2$ and $\gamma = 0.1$.

model	CER/%	SER/%
<i>CTC</i>	5.29	14.57
<i>Content based attention</i>	4.05	9.10
+ trigram LM	3.60	7.20
<i>Location based attention</i>	3.82	8.17
+ trigram LM	3.26	6.33
<i>Attention smoothing</i>	3.58	7.43
+ trigram LM	2.81	5.77

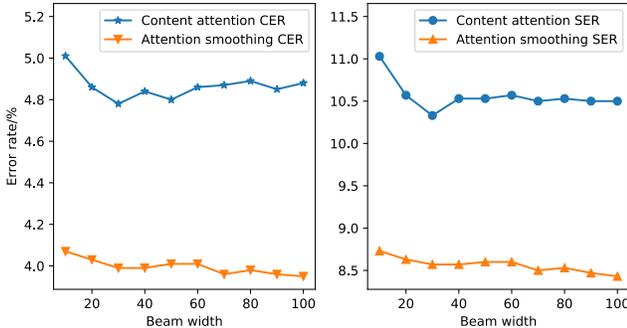


Fig. 3. The effect of the decoding beam width for the content-based attention and attention smoothing ($\tau = 1$).

initialization [10] and the bias vectors were initialized to 0. Gradient norm clipping to 1 was applied, together with Gaussian weight noise and L2 weight decay $1e-5$. We used ADAM as the optimization method [16] while we decayed the learning rate from $1e-3$ to $1e-4$ after it converged. The softmax output and the cross entropy cost were combined as the model cost. Frame skipping was used in the encoder during training. For comparison, we also constructed a CTC model that has the same structure with the the LAS encoder.

4.3. Decoding

We used a simple left-to-right beam search algorithm during decoding [8]. We investigated the importance of the beam-search width on decoding accuracy [8] and the impact of the temperature of the softmax function [5]. The temperature can smooth the distribution of characters. We changed the character probability distribution by a temperature hyperparameter τ :

$$y_t = \exp(o_t/\tau) / \sum_j \exp(o_j/\tau). \quad (14)$$

where o_t is the input of the softmax function.

4.4. Results

Table 1 shows that our models performed extremely well in the Mandarin voice search task. The content-based attention model achieved a CER of 4.05% and a SER of 9.1%. The location-based attention model achieved a CER of 3.82% and

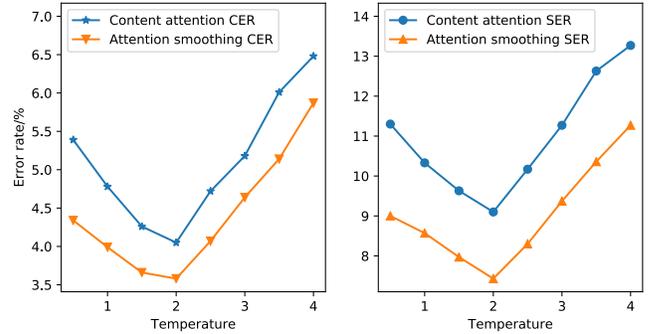


Fig. 4. The impact of the temperature for content-based attention and attention smoothing (beam-size=30).

a SER of 8.17%, which outperformed the content-based attention model. By using attention smoothing on the content-based attention model, the CER was reduced to 3.58% (or 11.6% relative gain over the content-based attention). We believe that the improvement is mainly because the sigmoid function keeps the diversity of the model and smooths the focus found by the attention mechanism.

Fig. 3 shows the effect of the decoding beam width on the WER/SER for the test set. The CER reached the lowest (4.78%) at a beam width of 30. We cannot observe extra benefit when further increasing the beam width. In Fig. 4, we can see that attention smoothing achieves the best performance when $\tau = 2$ and there is no additional benefits when we further increase the temperature. We see the same observation on the validation set as well. Meanwhile, we investigated the effect of adding language model. During decoding, with the help of a trigram LM that was trained using 4M voice search entries, further gains can be observed. Finally, *attention smoothing + trigram LM* achieved the lowest CER of 2.81%. This was obtained when $\gamma = 0.1$.

5. CONCLUSION

In this paper, we reported our preliminary results on attention-based encoder-decoder for Mandarin speech recognition. With several tricks, our model finally achieves a CER of 3.58% and a SER of 7.43% on a Mandarin voice search task without a language model. Note that the voice search content on MiTV is kind of limited with closed domains. In the future, we will further investigate our approach on general ASR tasks through public datasets.

6. ACKNOWLEDGMENTS

The authors would like to thank the Xiaomi Deep Learning Team and MiAI SRE Team for Xiaomi Cloud-ML and GPU cluster support. We also thank Yu Zhang and Jian Li for helpful comments and suggestions.

7. REFERENCES

- [1] D. Amodei, S. Ananthanarayanan, R. Anubhai, J. Bai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, Q. Cheng, G. Chen *et al.*, “Deep speech 2: End-to-end speech recognition in English and Mandarin,” in *ICML*, 2016, pp. 173–182.
- [2] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” *arXiv preprint arXiv:1409.0473*, 2014.
- [3] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, “End-to-end attention-based large vocabulary speech recognition,” in *ICASSP*. IEEE, 2016, pp. 4945–4949.
- [4] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *ICASSP*. IEEE, 2016, pp. 4960–4964.
- [5] W. Chan and I. Lane, “On online attention-based speech recognition and joint Mandarin character-pinyin training,” in *INTERSPEECH*, 2016, pp. 3404–3408.
- [6] 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.
- [7] J. Chorowski and N. Jaitly, “Towards better decoding and language model integration in sequence to sequence models,” *arXiv preprint arXiv:1612.02695*, 2016.
- [8] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, “Attention-based models for speech recognition,” in *Advances in Neural Information Processing Systems*, 2015, pp. 577–585.
- [9] L. Deng, J. Li, J. Huang, K. Yao, D. Yu, F. Seide, M. Seltzer, G. Zweig, X. He, J. Williams *et al.*, “Recent advances in deep learning for speech research at Microsoft,” in *ICASSP*. IEEE, 2013, pp. 8604–8608.
- [10] X. Glorot and Y. Bengio, “Understanding the difficulty of training deep feedforward neural networks,” in *Proceedings of the 13rd International Conference on Artificial Intelligence and Statistics*, 2010, pp. 249–256.
- [11] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *ICML*. ACM, 2006, pp. 369–376.
- [12] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath *et al.*, “Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups,” *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 82–97, 2012.
- [13] G. E. Hinton and D. Van Camp, “Keeping the neural networks simple by minimizing the description length of the weights,” in *Proceedings of the sixth annual conference on Computational learning theory*. ACM, 1993, pp. 5–13.
- [14] T. Hori, S. Watanabe, Y. Zhang, and W. Chan, “Advances in joint CTC-attention based end-to-end speech recognition with a deep CNN encoder and RNN-LM,” *arXiv preprint arXiv:1706.02737*, 2017.
- [15] K. Jim, C. L. Giles, and B. G. Horne, “An analysis of noise in recurrent neural networks: convergence and generalization,” *IEEE Transactions on neural networks*, vol. 7, no. 6, pp. 1424–1438, 1996.
- [16] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [17] Y. Miao, M. Gowayyed, and F. Metze, “Eesen: End-to-end speech recognition using deep RNN models and WFST-based decoding,” in *ASRU*. IEEE, 2015, pp. 167–174.
- [18] Y. Miao, J. Li, Y. Wang, S.-X. Zhang, and Y. Gong, “Simplifying long short-term memory acoustic models for fast training and decoding,” in *ICASSP*. IEEE, 2016, pp. 2284–2288.
- [19] R. Prabhavalkar, K. Rao, T. N. Sainath, B. Li, L. Johnson, and N. Jaitly, “A comparison of sequence-to-sequence models for speech recognition,” in *Interspeech*, 2017.
- [20] A. M. Rush, S. Chopra, and J. Weston, “A neural attention model for abstractive sentence summarization,” *arXiv preprint arXiv:1509.00685*, 2015.
- [21] J. Schmidhuber, “Deep learning in neural networks: An overview,” *Neural networks*, vol. 61, pp. 85–117, 2015.
- [22] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio, “Show, attend and tell: Neural image caption generation with visual attention,” in *ICML*, 2015, pp. 2048–2057.
- [23] Y. Zhang, W. Chan, and N. Jaitly, “Very deep convolutional networks for end-to-end speech recognition,” in *ICASSP*. IEEE, 2017, pp. 4845–4849.