

# ON ADDRESSING PRACTICAL CHALLENGES FOR RNN-TRANSDUCER

Rui Zhao, Jian Xue, Jinyu Li, Wenning Wei, Lei He, Yifan Gong

Microsoft Speech and Language Group

## ABSTRACT

In this paper, several works are proposed to address practical challenges for deploying RNN Transducer (RNN-T) based speech recognition systems. These challenges are adapting a well-trained RNN-T model to a new domain without collecting the audio data, obtaining time stamps and confidence scores at word level. We solve the first challenge with a splicing data method which concatenates the speech segments extracted from the source domain data. To get time stamps, a phone prediction branch is added to the RNN-T model by sharing the encoder for the purpose of forced alignment. Finally, we obtain word level confidence scores by utilizing several types of features calculated during decoding and from a confusion network. Evaluated with Microsoft production data, the splicing data adaptation method improves the baseline and adaptation with the text to speech method by 58.03% and 15.25% relative word error rate reduction, respectively. The proposed time stamping method can get less than 50 millisecond word timing difference from the ground truth alignment on average while maintaining the recognition accuracy. We also obtain high confidence annotation performance with limited computation cost.

*Index Terms*— end-to-end, adaptation, word timing, RNN-T, confidence annotation

## 1. INTRODUCTION

Recently, all-neural end-to-end (E2E) models [1, 2, 3, 4, 5] have become the trend in automatic speech recognition (ASR) due to its simple training and decoding processes, as well as the similar or even better accuracy compared with traditional hybrid speech recognition systems [6, 7, 8]. Commonly used E2E ASR models include Connectionist Temporal Classification (CTC) model [1, 9, 10], Attention-based Encoder-Decoder (AED) model [3], and transducer models such as recurrent neural network transducer (RNN-T) [2] and transformer-transducer (T-T) [11]. The industry has more interest in transducer models including RNN-T and T-T because these models can be streaming in a natural way [5, 11, 12, 13, 14, 15, 16, 17]. However, there are lots of practical challenges that need to be addressed in order to deploy all these models. In this paper, we will address some of them by using RNN-T as a backbone model. All solutions can be applied to other

transducer models such as T-T.

The first challenge is to adapt a RNN-T model to a new domain without collecting the speech data of the new domain. Because there is no separate language model in RNN-T, we couldn't easily adapt it with text only data as what we do for a hybrid ASR system. Instead, we need to collect paired speech-text data of the new domain which contains enough speaker and environment variation to adapt the RNN-T model. Such collection usually costs a lot time and money, which forbids fast model adaption [18]. Therefore, how to get a large amount of paired speech-text data with a small cost is a practical problem for the domain adaptation of RNN-T.

Some methods have been proposed to solve this problem. The most popular method is to synthesize speech from the new domain texts using the text to speech (TTS) technology [8, 19, 20, 21, 22]. For example, [8] uses a multi-speaker neural TTS system to generate speech data using the text-only data of the new domain to adapt the RNN-T model. [20] improves a general ASR model's performance on the numeric data domain with TTS generated numeric speech data. [22] improves OOV detection by adding missing syllable sounds using a cross-lingual TTS system. Although no real speech data needs to be collected, the TTS-based method has its limitations: 1) the speaker variation in TTS generated data is limited compared with the real production data used for ASR model training, 2) The cost of training a multi-speaker neural TTS model and the generation of synthesized speech is large. There are also other methods proposed for the domain adaptation of E2E models [23, 24]. However, the speech data for the new domain is still needed in these methods.

Some applications also require that RNN-T can provide functions such as time stamps and confidence measure at word level. Getting word timings with RNN-T is challenging because it's not designed for this: no time alignment information is used during training, therefore it cannot generate reliable word start and end time for every recognized word. To solve this problem, [25] applies time constraints during training to improve the word timing accuracy. It uses an extra token "word boundary" to mark the word start time and the last word piece of the word to mark the word end time. This method increases the training cost because adding "word boundary" for each word increases the training target label length. Applying time constraints to RNN-T models also causes significant accuracy degradation [25]. Another

method ‘‘FastEmit’’ proposed in [26] reduces the latency of RNN-T by encouraging predicting labels instead of blank during training. This could help to get accurate word ending time. However, this may bring accuracy loss because it blindly promotes the prediction of labels at all the time-label positions. Confidence annotation is also a challenging task for RNN-T models. Most applications require word level or utterance level confidence scores, while output units of RNN-T models are usually subword units such as word pieces. In [27, 28], word level confidence scores are obtained through aggregating confidence scores of constituent word pieces and directly modeling last word piece using an end-of-word mask. Although these methods can effectively obtain word level confidence scores, the rich information included in all the word pieces instead of the last one for each word is lost to some extent.

In this paper, we detail our solutions to these practical challenges. To adapt RNN-T without collecting new speech data, a novel splicing data generation method is proposed by concatenating the sampled speech segments corresponding to underlying words of target texts into new utterances. The sampled speech segments are extracted from the existing training data randomly. It has the following advantages compared with TTS-based speech generation method : 1) the cost is almost zero since it is only based on the existing speech data, without the need of training any extra model and the cost-consuming TTS speech generation, 2) the constructed data could cover more speaker and acoustic environment variation, which makes the model more robust. The proposed method is evaluated by adapting a general RNN-T model to a new domain, achieving 15.25% relative word error rate reduction over the model adapted with TTS generated data.

To provide reliable word timings for RNN-T, we propose adding a context independent (CI) phone prediction branch on top of the encoder of the RNN-T model. The word timings are calculated by aligning the recognition results from the RNN-T model using the phone probability of each frame from the CI phone prediction model in the second pass. Since the CI phone prediction model shares the encoder with RNN-T model, the cost of phone level alignment is very cheap. The experiments proved that the proposed method could get less than 50 millisecond (ms) word timing difference on average compared with the ground truth while maintaining the recognition accuracy of the RNN-T model.

To provide high quality word level confidence scores for RNN-T models, we first obtain confidence features at word piece level directly from decoding, then aggregate them to word level features. Compared to aggregating confidence scores [27, 28], aggregating features could retain much richer information, thus has the potential to make the final word level confidence scores more reliable. We also generate extra features from confusion network constructed with the N-best list [29, 30], which provide complimentary information from other hypotheses. With all the confidence features, we train a

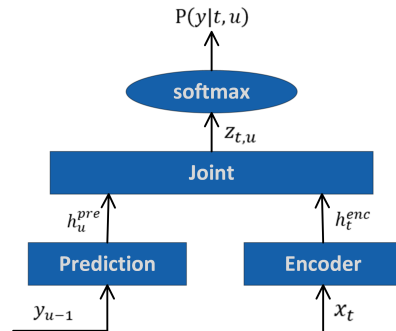


Fig. 1. RNN-T model structure.

two-layer feed forward neural network to classify each recognized word as ‘‘correct’’ or ‘‘incorrect’’. Experimental results showed that such a method could achieve high confidence annotation performance with limited computation cost.

## 2. RNN-T

A RNN-T model [2] consists of encoder, prediction, and joint networks as shown in Figure 1. The encoder network is analogous to the acoustic model in hybrid models, which converts the acoustic feature  $x_t$  into a high-level representation  $h_t^{enc}$ , where  $t$  is the time index. The prediction network works like a RNN language model, which produces a high-level representation  $h_u^{pre}$  by conditioning on the previous non-blank target  $y_{u-1}$  predicted by the RNN-T model, where  $u$  is output label index. The joint network combines the encoder network output  $h_t^{enc}$  and the prediction network output  $h_u^{pre}$  with a feed forward network as

$$z_{t,u} = f^{joint}(h_t^{enc}, h_u^{pre}). \quad (1)$$

Then  $z_{t,u}$  is connected to the output layer with a linear transform

$$h_{t,u} = W_y z_{t,u} + b_y. \quad (2)$$

The final posterior for each output token  $k$  is obtained after applying the softmax operation

$$P(k|t, u) = softmax(h_{t,u})_k. \quad (3)$$

The loss function of RNN-T is the negative log posterior of output label sequence  $\mathbf{y}$  given the input acoustic feature  $\mathbf{x}$ ,

$$L_{rnt} = -\ln P(\mathbf{y}|\mathbf{x}). \quad (4)$$

## 3. SPLICING DATA GENERATION FOR DOMAIN ADAPTATION

Domain adaptation is to adapt the well-trained source domain model to a target domain [18] which usually has mismatched

content from the source domain. Given the model was trained with a large amount of data, every word in the target domain usually has been observed while the word sequence is very different between the source and target domain.

When there is no budget for collecting and labeling speech data in the new domain, we could use TTS to synthesize speech for given text as in [8]. However, the speaker variation of the TTS generated audio is limited even with the state-of-the-art multi-speaker neural TTS model [31] when compared with that in the training data for ASR models of real productions. Besides, the cost of training a multi-speaker neural TTS model and generating TTS speech with it is high due to two reasons: 1) It needs the real speech and the transcription corresponding to it from multi speakers to train the model. 2) TTS model training and speech data generation require high performance computational resources, such as GPU.

In this paper, we propose a new method to generate speech data based on the existing ASR model training data. The proposed splicing data method generates new utterances by concatenating the audio segments of specific speech units (e.g., words) extracted from the existing training data. If there are multiple segments in the training data for the same word, we just randomly pick one segment. Figure 2 shows the implementation of the proposed methods. The new audio for the given text “Cortana open door” is generated by concatenating the audio segments of words “Cortana”, “open” and “door”. These audio segments are extracted randomly from the existing training utterances. In this way, we can generate audio for almost any texts. If there is an out of vocabulary (OOV) word in the new domain, speech units corresponding to phones are used to construct speech for this OOV word. It should be noted that there is no restriction when selecting audio segments, i.e., it is not required that the selected audio segments for one utterance should be from the same speaker, acoustic environment or corpus. They are all selected randomly.

The proposed method has clear advantages over the TTS-based audio generation method. Firstly, the cost is almost zero since it doesn’t need any extra model or data. Secondly, the speech data generated with the proposed method is “real” speech at each segment, hence it has the potential to cover all the speakers and acoustic environments in the existing training data. Therefore the speaker and acoustic variation in such data is much higher than that in the TTS-generated speech data. The model trained with the data generated by the proposed method should be more robust than that trained with TTS data.

A possible argument is that the generated audio in this way is not continuous at the transitions between words. As we know, one important feature of E2E models is that they make recognition decision after processing a segment of speech instead of frame by frame. For example, a RNN-T model will output “blank” if it isn’t confident about what content the speech until now should be. Hence this dis-fluency at the

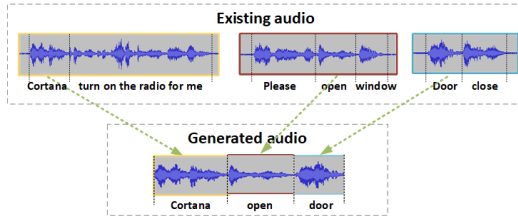


Fig. 2. Spliced data generation.

transitions between words won’t affect the RNN-T model too much. Besides, as shown in [19], when adapting a RNN-T model with TTS generated data, the lower layers of the encoder should be fixed and the natural speech data should be mixed to combat the over-fitting to TTS data. The same trick is also used in the proposed method to reduce the side effects of the spliced data in the experiments in section 6.1 and it is proved to be effective by the results.

#### 4. WORD TIMING WITH PHONE PREDICTION

To get accurate word timings of the recognition results with RNN-T model is challenging. Firstly, for a RNN-T model with word piece as the output token, it’s difficult to get both the word start and end time. For example, when a frequent word is tokenized into only one word piece, only the word start or end time could be estimated. Secondly, RNN-T model training is not guided by the ground truth alignment. Instead, all possible alignments are considered during the alignment, and the streaming model tends to delay its output to get better accuracy. To solve the above issues and get more accurate word timings, the authors in [25] add a “word boundary” token to get the word start time and use last word piece for word end time. Besides, the ground truth alignment is used to constraint the training. Although this method could help RNN-T model output better word timings, it results in recognition accuracy degradation due to the time constraint training.

We propose getting accurate word timings while maintaining the RNN-T model recognition accuracy by adding a CI phone prediction branch on top of the encoder. The CI phone prediction branch is used to obtain word timings with traditional forced alignment method, instead of forcing RNN-T to do the work which is not designed for. Specifically, the word sequence output from RNN-T model is aligned with time based on the CI phone probability output in the second pass. The CI phone prediction model is trained with ground truth alignment using the frame wise cross entropy (CE) criterion. Therefore, the time alignment based on such model is much more accurate than that obtained by a RNN-T model.

The model structure of the proposed method is shown in Figure 3. To save the computational cost, the CI phone prediction model shares the encoder with the RNN-T model. In order to maintain a good recognition accuracy, it is better to

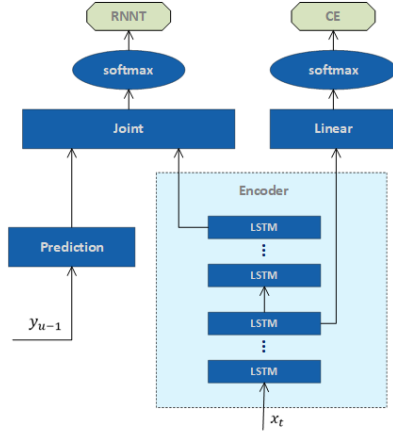


Fig. 3. Word timing model structure

not apply any time constraint to the RNN-T model during training. In contrast, the CI phone prediction model should be trained with ground truth alignment. If the whole encoder network is shared between these two models, they may have negative impacts on each other. Therefore, the CI phone prediction model shares only the lower layers of the encoder network of RNN-T model.

The CI phone prediction and RNN-T models could be trained jointly with the multitask learning (MTL) method from scratch. In this case, the total loss is represented as follows by combining both RNN-T loss in Eq. 4 and CE loss.

$$L_{MTL} = \alpha L_{ce} + (1 - \alpha) L_{rnnnt}. \quad (5)$$

Where  $\alpha$  is the MTL weight.

If we already have a well trained RNN-T model, we could also directly add the CI phone prediction branch and train it by fixing the encoder part. This ensures the RNN-T model won't be affected and saves training time.

Because the CI phone prediction model shares the encoder with RNN-T model and the phone level alignment is fast, the total cost of the second pass alignment is quite small.

## 5. CONFIDENCE ANNOTATION

We calculate word confidence scores using a binary classifier trained with features obtained from decoding and confusion network. The features calculated directly from decoding are at word piece level, and then we aggregate them to word level.

### 5.1. Feature generation during decoding

We generate three features on word pieces during decoding:

- wp-prob: Log posterior probability of the word piece.
- hyp-prob: Log posterior probability of the partial hypothesis ended with the current word piece.

- neg-entropy: Negative entropy calculated over all the nodes on the output layer when output the word piece, which is  $\sum_i p_i \log(p_i)$ .

Then we aggregate word piece level features into five word level features:

- avg-hyp-prob: Log posterior probability of the partial hypothesis, averaged by the number of tokens, including word pieces and “blank” token.
- min-wp-prob: Minimum of the log posterior probabilities for all word pieces in the word.
- avg-wp-prob: Average of the log posterior probabilities for all word pieces in the word.
- min-neg-entropy: Minimum of the negative entropy for all word pieces in the word.
- avg-neg-entropy: Average of the negative entropy for all word pieces in the word.

### 5.2. Feature generation from confusion network

We construct a word level confusion network with the N-best list generated from decoding. We use two different scores for each hypothesis in the N-best list, one is the posterior probability of the hypothesis, and another one is the length normalized posterior probability. Hence we obtain two confusion networks, and use the word posterior probabilities in each confusion network as the extra confidence features, noted as cn-prob and cn-norm-prob.

With the 7 word level features described above, we train a 2-layer feed forward neural network to classify each word into two classes (correct and incorrect). The classifier is trained to minimize the binary entropy between the estimated confidence score  $p$  and target class  $c$ , where  $c$  is 1 if the word is correct, or 0 if the word is an insertion or substitution error.

$$L(c, p) = \sum_i (c_i \log(p_i) + (1 - c_i) \log(1 - p_i)). \quad (6)$$

## 6. EXPERIMENTS

The baseline RNN-T model was trained with 65 thousand (K) hours transcribed Microsoft data. All the data are anonymized with personally identifiable information removed. The encoder network of the RNN-T model consists of 6 LSTM layers, with 768 nodes per layer. The prediction network consists of 2 LSTM layers, with 768 nodes per layer. The joint network is a single feed forward layer with 640 nodes. Singular value decomposition (SVD) [32] compression method is applied to further reduce the model size. The feature is 80-dimension log Mel filter bank for every 10 milliseconds (ms) speech. Three of them are stacked together to form a frame of 240-dimension input acoustic feature to the encoder network. The output targets are 4000 word piece units.

**Table 1.** WER for RNN-T models on a new domain

| RNN-T models              | WER(%) |
|---------------------------|--------|
| baseline                  | 9.27   |
| adapted with TTS data     | 4.59   |
| adapted with spliced data | 3.89   |

### 6.1. Domain adaptation

In this experiment, the baseline RNN-T model is adapted to a new command and control domain. The testing data contains about 800 utterances collected in the real application environment with personally identifiable information removed. To adapt the RNN-T model into the new domain, first, the texts of the new domain are obtained by randomly parsing the grammar in the new domain and also using the crowd sourcing method, with about 200 thousand(K) sentences in total. Then, the speech of these texts are generated with either TTS or the proposed splicing data method.

The multi speaker neural TTS model used to generate adaptation speech data is the best TTS model we have. It was trained by following the approach in [31]. Firstly, a multi speaker neural TTS model is built with in-house TTS corpus, then it is fine-tuned by adding LibriSpeech [34] data with thousands of speakers onboard. To evaluate the speech quality generated with this model, 10 speakers in LibriSpeech corpus are randomly selected, and 60 TTS samples are generated for each speaker. The subjective listening test (Mean Opinion Score, MOS) is carried out with crowd-sourcing judges. The results show the neural TTS quality is close to human recording of corresponding speakers. For the adaptation speech data generation, about 1000 speakers are selected, and for each speaker, 2.5K utterances are generated based on the texts randomly chosen from those 200K sentences. Then, the TTS speech data is doubled by adding simulated data, which is obtained by adding different noises and room impulse responses (RIR) to the original data to improve the robustness. In total, we got 5 million(M) utterances of TTS speech data for the new domain.

In the splicing data method, for each text sentence, the audio segment for each word (or phone) in the text sentence is randomly extracted from the 65k hours general data. Similar to TTS data, a total of about 5M utterances are generated.

When updating the baseline RNN-T model with the TTS or spliced data, we use two methods to prevent model from over fitting to the TTS or spliced data: 1) The lower 4 layers of the encoder network are frozen. 2) Similar amount of (about 5M utterances) normal speech data are randomly selected from the 65k hours data and mixed with the TTS or spliced data.

Table 1 shows the word error rate (WER) of the baseline RNN-T model, as well as the RNN-T models adapted with the TTS and spliced data. Adapting with spliced data gets 58.03% and 15.25% relative word error rate reduction over

the baseline and the one adapted with TTS data, respectively.

### 6.2. Word timing

In this experiment, the recognition accuracy evaluation set is a general test set which covers 13 application scenarios such as Cortana and far-field speech, containing a total of 1.8M words. The personally identifiable information is also removed for the testing data. 3500 utterances are randomly selected from the above 1.8M word set to evaluate the word timing accuracy. The reference word boundary is obtained by the forced alignment with a traditional hybrid model.

As in [25], we measure the accuracy of word timings with below 4 metrics:

- Average start time delta (Ave. ST  $\Delta$ ): the average start time difference between the ground truth word start time and the estimated word start time.
- Average end time delta (Ave. ET  $\Delta$ ): the average end time difference between the ground truth word end time and the estimated word end time.
- Percentage of word start time less than 200ms (% WS < 200ms): the percentage of word start time difference that is less than 200ms.
- Percentage of word end time less than 200ms (% WE < 200ms): the percentage of word end time difference that is less than 200ms.

Table 2 gives the WER and word timing accuracy for different methods. For the baseline model, the word end time is the emitting time of last word piece in this word. No word start time is estimated as explained in Section 4. For the proposed methods, we examined the performance of two different training recipes: one is adding CI phone prediction model to a well trained RNN-T model and only updating the CI phone model by fixing the shared encoder part. Another is training RNN-T and CI phone models with the MTL method from scratch. In this recipe,  $\alpha$  is set to 0.1. We also build a RNN-T model based on the time constraint method in [25].

From the results, we can see all models trained with the proposed methods get smaller average word start/end time delta than the model trained with the time constraint method in [25]. The WER of RNN-T models trained with our proposed methods is not increased compared to the baseline RNN-T model, while the RNN-T model trained with time constraint method [25] has a higher WER. MTL Training from scratch gives a little lower WER than training only the CI phone prediction branch with CE. These two different training recipes could get similar word boundary accuracy.

We also compared the performance of sharing different number of encoder network layers: lower 4 layers or full encoder (6 layers). Sharing lower 4 layers yields much better word boundary accuracy than sharing the full encoder: the

**Table 2.** WER and word timing accuracy for different methods with RNN-T model

|                       | RNN-T Baseline | RNN-T + CE<br>update CE only | RNN-T + CE<br>update CE only | RNN-T+CE<br>MTL | RNN-T + CE<br>MTL | RNN-T with<br>time constraint |
|-----------------------|----------------|------------------------------|------------------------------|-----------------|-------------------|-------------------------------|
| Shared encoder layers | -              | lower 4                      | full                         | lower 4         | full              | -                             |
| WER(%)                | 12.02          | 12.02                        | 12.02                        | 11.89           | 11.93             | 12.84                         |
| Ave. ST $\Delta$      | -              | 45ms                         | 89ms                         | 46ms            | 86ms              | 120ms                         |
| Ave. ET $\Delta$      | 202ms          | 42ms                         | 85ms                         | 44ms            | 82ms              | 97ms                          |
| % WS < 200ms          | -              | 97.86                        | 94.42                        | 97.69           | 94.19             | 98.63                         |
| % WE < 200ms          | 49.88          | 97.80                        | 94.28                        | 97.75           | 94.39             | 98.62                         |

**Table 3.** Confidence annotation results using each single feature and classifier

|                   | AUPR-incorrect | AUPR-correct |
|-------------------|----------------|--------------|
| avg-hyp-prob      | 14.56          | 94.60        |
| min-wp-prob       | 25.70          | 96.68        |
| avg-wp-prob       | 21.71          | 96.45        |
| min-neg-entropy   | 38.58          | 97.35        |
| avg-neg-entropy   | 33.68          | 96.91        |
| cn-prob           | 35.99          | 96.40        |
| cn-norm-prob      | 31.27          | 96.36        |
| classifier        | 45.86          | 98.69        |
| hybrid model [33] | 23.86          | 97.28        |

average word timing difference is decreased from more than 80ms to less than 50ms.

The average decoding time is increased by about 1% with the second pass alignment for word timings.

### 6.3. Confidence annotation

We conducted confidence experiments on the general test set described in section 6.2. We use the area under precision-recall curve (AUPR) as the evaluation metric. We calculate AUPR for both correctly and incorrectly recognized words. The larger AUPR is, the better the confidence classifier performs. Table 3 summarizes the performance. From the results we can see the min-neg-entropy feature gives the best performance among all features. We further improve the performance by using the classifier to combine all features together. In the last row, we also listed the confidence performance obtained with the method in [35] of the hybrid model [33] as a reference.

We evaluated the computation cost caused by the confidence annotation function during online decoding. For each utterance in the test set, we calculated the percentage of increased computation time after adding confidence annotation, compared to decoding the utterance without generating confidence features and score. The average percentage of increased computation time is 2.9% for the whole test.

## 7. CONCLUSIONS

In this paper, three challenges for deploying RNN-T models to real applications have been addressed. To adapt a well-trained RNN-T model to a new domain without collecting new speech data, for all the words in the text sentences of the new domain, we randomly extracted the corresponding audio segments from the source training data, and then concatenated them to form new speech utterances. Experimental result showed the model adapted with such data obtained 58.03% and 15.25% relative WER reduction over the baseline model and the model adapted with TTS generated data, respectively. To get accurate word timings of the recognition results, we added a context independent phone prediction branch by sharing the lower layers of the RNN-T encoder network to perform forced alignment based on the phone probability of the phone prediction model. Experimental results showed that we could get less than 50ms average word boundary differences compared with the ground truth alignment without degrading the recognition accuracy of the RNN-T model. With the encoder sharing, the additional cost of the alignment is about 1%. To obtain reliable word level confidence scores for the RNN-T model with word piece units, we designed 7 word level confidence features and a binary classifier. Experimental results showed the effectiveness of all the features and the binary classifier can further boost performance. The additional computational cost increase of confidence measure function is about 2.9%.

## 8. ACKNOWLEDGEMENT

We thank Kshitiz Kumar at Microsoft for providing the confidence measure evaluation of the hybrid model. We also thank Yuhui Wang and Min Hu at Microsoft for the runtime support of the time alignment and confidence measure of the RNN-T model.

## 9. REFERENCES

- [1] Alex Graves, Santiago Fernández, Faustino Gomez, and Jurgen Schmidhuber, “Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks,” in *ICML*, 2006, pp. 369–376.
- [2] Alex Graves, “Sequence transduction with recurrent neural networks,” in *International Conference on Machine Learning: Representation Learning Workshop*, 2012.
- [3] William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *Proc. ICASSP*, 2016, pp. 4960–4964.
- [4] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” in *Advances in neural information processing systems*, 2017.
- [5] Yanzhang He, Tara N Sainath, Rohit Prabhavalkar, Ian McGraw, Raziq Alvarez, Ding Zhao, David Rybach, Anjali Kannan, Yonghui Wu, Ruoming Pang, et al., “Streaming end-to-end speech recognition for mobile devices,” in *Proc. ICASSP*, 2019, pp. 6381–6385.
- [6] Shinji Watanabe, Takaaki Hori, Suyoun Kim, John R Hershey, and Tomoki 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.
- [7] Tara N Sainath, Yanzhang He, Bo Li, et al., “A streaming on-device end-to-end model surpassing server-side conventional model quality and latency,” in *Proc. ICASSP*, 2020, pp. 6059–6063.
- [8] Jinyu Li, Rui Zhao, Zhong Meng, et al., “Developing RNN-T models surpassing high-performance hybrid models with customization capability,” in *Proc. Interspeech*, 2020.
- [9] Yajie Miao, Mohammad Gowayyed, and Florian Metze, “EESSEN: end-to-end speech recognition using deep RNN models and WFST-based decoding,” in *Proc. ASRU*, 2015, pp. 167–174.
- [10] Jinyu Li, Guoli Ye, Amit Das, Rui Zhao, and Yifan Gong, “Advancing acoustic-to-word CTC model,” in *Proc. ICASSP*, 2018, pp. 5794–5798.
- [11] Qian Zhang, Han Lu, Hasim Sak, Anshuman Tripathi, Stephen Koo Erik McDermott, and Shankar Kumar, “Transformer transducer: A streamable speech recognition model with transformer encoders and RNN-T loss,” *ICASSP*, 2020.
- [12] Jinyu Li, Rui Zhao, Hu Hu, and Yifan Gong, “Improving RNN transducer modeling for end-to-end speech recognition,” in *Proc. ASRU*, 2019.
- [13] Eric Battenberg, Jitong Chen, Rewon Child, Adam Coates, Yashesh Gaur, Yi Li, Hairong Liu, Sanjeev Satheesh, David Seetapun, Anuroop Sriram, and Zhenyao Zhu, “Exploring neural transducers for end-to-end speech recognition,” in *Proc. ASRU*, 2017.
- [14] Chung-Cheng Chiu, Tara N Sainath, Yonghui Wu, Rohit Prabhavalkar, Patrick Nguyen, Zhifeng Chen, Anjali Kannan, Ron J Weiss, Kanishka Rao, Katya Gonnina, et al., “State-of-the-art speech recognition with sequence-to-sequence models,” in *Proc. ICASSP*, 2018.
- [15] Jinyu Li, Yu Wu, Yashesh Gaur, Chengyi Wang, Rui Zhao, and Shujie Liu, “On the comparison of popular end-to-end models for large scale speech recognition,” in *Proc. Interspeech*, 2020.
- [16] Ching-Feng Yeh, Jay Mahadeokar, Kaustubh Kalgankar, Yongqiang Wang, Duc Le, Mahaveer Jain, Kjell Schubert, Christian Fuegen, and Michael L Seltzer, “Transformer-transducer: End-to-end speech recognition with self-attention,” *arXiv preprint arXiv:1910.12977*, 2019.
- [17] Xie Chen, Yu Wu, Zhenghao Wang, Shujie Liu, and Jinyu Li, “Developing real-time streaming transformer transducer for speech recognition on large-scale dataset,” in *ICASSP*, 2021.
- [18] Peter Bell, Joachim Fainberg, Ondrej Klejch, Jinyu Li, Steve Renals, and Pawel Swietojanski, “Adaptation algorithms for neural network-based speech recognition: An overview,” *IEEE Open Journal of signal processing*, vol. 2, 2021.
- [19] Eva Sharma, Guoli Ye, Wenning Wei, Rui Zhao, Yao Tian, Jian Wu, Lei He, Edward Lin, and Yifan Gong, “Adaptation of RNN transducer with text-to-speech technology for keyword spotting,” in *Proc. ICASSP*, 2020.
- [20] Cal Peysers, Hao Zhang, Tara N. Sainath, and Zelin Wu, “Improving performance of end-to-end ASR on numeric sequences,” in *Proc. Interspeech*, 2019.
- [21] Murali Karthick Baskar, Shinji Watanabe, Ramon Astudillo, Takaaki Hori, Lukáš Burget, and Jan Černocký, “Self-supervised sequence-to-sequence ASR using unpaired speech and text,” in *Conference of the International Speech Communication Association*, 2019.
- [22] Savitha Murthy, Dinkar Sitaram, and Sunayana Sitaram, “Effect of TTS generated audio on OOV detection and word error rate in ASR for low-resource languages,” in *Proc. Interspeech*, 2018, pp. 1026–1030.

- [23] Zhong Meng, Jinyu Li, Yashesh Gaur, and Yifan Gong, "Domain adaptation via teacher-student learning for end-to-end speech recognition," in *ASRU*, 2019.
- [24] Lahiru Samarakoon, Brian Mak, and Albert Y.S. Lam, "Domain adaptation of end-to-end speech recognition in low-resource settings," in *Proc. SLT*, 2018.
- [25] Tara N Sainath, Ruoming Pang, David Rybach, Basi García, and Trevor Strohman, "Emitting word timings with end-to-end models," in *Proc. Interspeech*, 2020.
- [26] Jiahui Yu, ChungCheng Chiu, Bo Li, Shuoyiin Chang, Tara N Sainath, Yanzhang He, Arun Narayanan, Wei Han, Anmol Gulati, Yonghui Wu, and Ruoming Pang, "FastEmit: Low-latency Streaming ASR with Sequence-level Emission Regularization," in *ICASSP*, 2021.
- [27] Dan Oneata, Alexandru Caranica, Adriana Stan, and Horia Cucu, "An evaluation of word-level confidence estimation for end-to-end automatic speech recognition," in *Proc. SLT*, 2021, pp. 258–265.
- [28] David Qiu, Qiujia Li, Yanzhang He, Yu Zhang, Bo Li, Liangliang Cao, Roohit Prabhavalkar, Deepti Bhatia, Wei Li, Ke Hu, T N. Sainath, and Ian McGraw, "Learning word-level confidence for subword end-to-end ASR," *ICASSP*, 2021.
- [29] Lidia Mangu, Eric Brill, and Stolcke Andreas, "Finding consensus in speech recognition: Word error minimization and other applications of confusion network," *Computer Speech and Language*, 2000.
- [30] Jian Xue and Yunxin Zhao, "Random forests-based confidence annotation using novel features from confusion network," in *Proc. ICASSP*, 2006, pp. I-1149–I-1152.
- [31] Yan Deng, Lei He, and Frank Soong, "Modeling multi-speaker latent space to improve neural TTS: Quick enrolling new speaker and enhancing premium voice," in <https://arxiv.org/abs/1812.05253>, 2018.
- [32] Jian Xue, Jinyu Li, and Yifan Gong, "Restructuring of deep neural network acoustic models with singular value decomposition.," in *Interspeech*, 2013, pp. 2365–2369.
- [33] Jinyu Li, Rui Zhao, Eric Sun, Jeremy HM Wong, Amit Das, Zhong Meng, and Yifan Gong, "High-accuracy and low-latency speech recognition with two-head contextual layer trajectory LSTM model," in *Proc. ICASSP*, 2020, pp. 7699–7703.
- [34] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, "Librispeech: an ASR corpus based on public domain audio books," in *Proc. ICASSP*, 2015, pp. 5206–5210.
- [35] Po-Sen Huang, Kshitiz Kumar, Chaojun Liu, Yifan Gong, and Li Deng, "Predicting speech recognition confidence using deep learning with word identity and score features," in *Proc. ICASSP*, 2013, pp. 7413–7417.