FAST AND ACCURATE FACTORIZED NEURAL TRANSDUCER FOR TEXT ADAPTION OF END-TO-END SPEECH RECOGNITION MODELS

Rui Zhao, Jian Xue, Partha Parthasarathy, Veljko Miljanic, Jinyu Li

Microsoft Speech Group, Redmond, WA, USA

ABSTRACT

Neural transducer is now the most popular end-to-end model for speech recognition, due to its naturally streaming ability. However, it is challenging to adapt it with text-only data. Factorized neural transducer (FNT) model was proposed to mitigate this problem. The improved adaptation ability of FNT on text-only adaptation data came at the cost of lowered accuracy compared to the standard neural transducer model. We propose several methods to improve the performance of the FNT model. They are: adding CTC criterion during training, adding KL divergence loss during adaptation, using a pre-trained language model to seed the vocabulary predictor, and an efficient adaptation approach by interpolating the vocabulary predictor with the n-gram language model. A combination of these approaches results in a relative word-error-rate reduction of 9.48% from the standard FNT model. Furthermore, n-gram interpolation with the vocabulary predictor improves the adaptation speed hugely with satisfactory adaptation performance.

Index Terms— neural transducer model, factorized transducer model, KL divergence, n-gram

1. INTRODUCTION

Recently, neural transducer based end-to-end (E2E) models [1–11], such as recurrent neural network transducer (RNN-T) [1], transformer-transducer (T-T) [4, 5] and conformer-transducer (C-T) [12], have become the dominant model for automatic speech recognition (ASR) in industry due to its natural streaming property, as well as competitive accuracy with traditional hybrid speech recognition systems [13–15].

However, one of the main challenges for neural transducer models is adaptation using only text data. This is because in neural transducer models, there are no separated acoustic or language model like in traditional hybrid models. Although the prediction network could be considered as an internal language model (LM) because the input to it is the previously predicted token, it is not a real LM since the prediction output needs to be combined with the acoustic encoder in a non-linear way to generate posteriors over the vocabulary augmented with a *blank* token. Adapting the prediction network using text-only data is not as straightforward or effective as adapting the LM in hybrid systems. Paired audio and text data is needed to adapt a neural transducer model, however, collecting labeled audio data is both time and money costly.

There are several types of methods proposed to address this issue. One is to generate artificial audio for adaptation text instead of collecting real audio. Audio generation method could be based on multi-speaker neural text to speech (TTS) model [15–19] or spliced data method [20]. The neural transducer model could then be finetuned with artificial paired audio and text data. A major drawback of these kinds of methods is the high computational cost. It takes much longer for the TTS-based methods to generate audio even with GPU machines, while the spliced-data method has very small cost for generating audio. However, the adaptation step for both methods involves updating part of the encoder, full prediction and the joint network with the RNN-T loss. The result is high computational cost for training, need for GPUs, and too much delay for scenarios in which rapid adaptation is necessary.

Another class of text-only adaptation methods is LM fusion [21–25], such as shallow fusion [21] where an external LM trained on target-domain text is incorporated during the neural transducer model decoding. However, there is already an internal LM in the neural transducer model. Directly adding an external LM is not mathematically grounded. To solve such an issue, density ratio [26], hybrid autoregressive transducer model [27], and internal LM estimation [28, 29] were proposed to remove the influence of the internal LM contained in the neural transducer model. However, the performance is often sensitive to the interpolation weight of external LM for different tasks, and it needs to be well tuned based on development data to get optimal results [30].

Different from aforementioned methods, factorized neural transducer model (FNT) [31] modifies the neural transducer model architecture by factorizing the blank and vocabulary prediction so that a standalone LM can be used for the vocabulary prediction. In this way, various language model adaptation [32, 33, 34] techniques could be applied to FNT. But based on results in [31], FNT degrades the accuracy on general testing sets compared with the standard neural transducer model, although it significantly improves the accuracy in the new domain after adaptation. Besides, it still needs significant GPU time to finetune FNT with text only data for the adaptation, which may not meet the fast adaptation requirement in some real applications.

In this paper, we proposed several methods to advance FNT for effective and efficient adaptation with text only data. These methods include: 1) Adding Connectionist Temporal Classification (CTC) [32] criterion for the encoder network during training to make it work more like an acoustic model. Then, the combination of encoder output and vocabulary predictor output is similar to the combination of acoustic and language model in hybrid models. 2) Adding Kullback-Leibler (KL) divergence between the outputs of adapted model and baseline model to avoid over fitting to the adaptation data. 3) Initializing the vocabulary predictor with a neural LM trained with more text data. 4) Replacing the network fine-tuning with more efficient adaptation method using n-gram interpolation. Experimental results showed that on general testing sets, these methods help the modified FNT to get even a little better accuracy than the baseline neural transducer model. On adaptation sets, the word error rate (WER) after the adaptation of modified FNT is reduced by 9.48% relatively from the standard FNT model, and reduced by 29.21% relatively from the baseline C-T model. Besides, n-gram interpolation results in much faster adaptation speed.

The rest of this paper is organized as follows: Section 2 introduces the neural transducer model and FNT model. Section 3 presents the proposed methods for modified FNT in detail. Section 4 shows the experimental results and analysis. Section 5 gives the conclusions.

2. STANDARD NEURAL TRANSDUCER AND FNT MODEL

2.1. Standard neural transducer

A neural transducer model [1] consists of encoder, prediction, and joint networks. The encoder network is analogous to the acoustic model in hybrid models, which converts the acoustic feature x_t into a high-level representation f_t , where t is the time index. The prediction network works like a neural LM, which produces a high-level representation g_u 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 f_t and the prediction network output g_u to compute the output probability with

$$z_{t,u} = W * \operatorname{relu}(f_t + g_u) + b$$

$$P(\hat{y}_{t+1}|x_1^t, y_1^u) = \operatorname{softmax}(z_{t,u})$$
(1)

To address the length differences between the acoustic feature \mathbf{x}_1^T and label sequences \mathbf{y}_1^U , a special blank symbol, ϕ , is added to the output vocabulary. Therefore the output set is $\{\phi \cup \mathcal{V}\}$, where \mathcal{V} is the vocabulary set.

2.2. Factorized neural transducer

Two prediction networks are used in FNT [31], as shown in figure 1. One (*Predictor_b*) is for the prediction of the blank label ϕ , and the other (*Predictor_v*) is vocabulary prediction (rightmost orange part in figure 1). The vocabulary predictor could be considered as a standard LM. The combination methods with encoder output f_t for these two prediction outputs are different. For the blank prediction, it is the same as in standard neural transducer models.

$$z_{t,u}^b = W^b * \operatorname{relu}(f_t + g_u^b) + b^b \tag{2}$$

For the vocabulary prediction, it is firstly projected to the vocabulary size and converted to the log probability domain by the operation of log softmax. After this, it is added with the encoder output.

$$d_t^v = W_{enc}^v * \operatorname{relu}(f_t) + b_{enc}^v$$

$$d_u^v = W_{pred}^v * \operatorname{relu}(g_u^v) + b_{pred}^v$$

$$z_u^v = \log_\operatorname{softmax}(d_u^v)$$

$$z_{t,u}^v = d_t^v + z_u^v$$
(3)

Two combination outputs are concatenated and softmax is applied to get the final label probability

$$P(\hat{y}_{t+1}|x_1^t, y_1^u) = \text{softmax}([z_{t,u}^b; z_{t,u}^v])$$
(4)

The loss function of FNT is

$$\mathcal{J}_f = \mathcal{J}_t - \lambda \log P(\mathbf{y}_1^U) \tag{5}$$

where the first term is the standard neural transducer loss and the second term is the LM loss with cross entropy (CE). λ is a hyperparameter to tune the effect of LM loss.



Fig. 1. Flowchart of factorized neural transducer

3. IMPROVEMENT OF FNT

In this section, we will propose several methods to improve the accuracy and efficiency of FNT.

3.1. Adding CTC criterion in training

As showed in section 2.2, the encoder output and the vocabulary predictor output are combined by sum operation. The predictor output is log probability, but the encoder output is not. According to Bayes' theory, the acoustic and language model scores should be combined by weighted sum in log probability domain. Therefore we refine FNT by converting the encoder output to the log probability by adding log softmax

Furthermore, to force the encoder part act more like the acoustic model, CTC criterion is added for the encoder output as shown in the blue frame part in figure 2. The reason we choose CTC instead CE is that it's not easy to get the sentence piece level alignment for training data, while sentence piece unit is commonly used as the output unit for neural transducer E2E model.

With such changes, the combination of encoder output and vocabulary predictor output is shown in below equations.

$$z_t^v = \log_softmax(d_t^v)$$

$$z_{t,u}^v = z_t^v[:-1] + \gamma * z_u^v$$
(6)

where γ is a trainable parameter, which could be taken as LM weight. One thing needs to be mentioned is after adding CTC, the dimension of d_t^v and z_t^v become vocabulary_size+1 because CTC needs one extra output "blank". Here we put "blank" as the last dimension. and it's excluded when z_t^v is added with z_u^v .

The final loss function can be written as

$$\mathcal{J}_f = \mathcal{J}_t - \lambda \log P(\mathbf{y}_1^U) + \beta \mathcal{J}_{ctc}$$
(7)

where \mathcal{J}_{ctc} is CTC loss and β is a hyper-parameter to be tuned in the experiments.

3.2. Adding KL divergence in adaptation

To adapt FNT model with text data, the most straightforward way is to finetune the vocabulary predictor with the adaptation text based on cross-entropy loss. But this may degrade model performance on general domain. To avoid this, KL divergence between the vocabulary predictor outputs of adapted model and baseline model is added during the adaptation as shown in the orange frame part in figure 2.



Fig. 2. BLUE FRAME: adding CTC criterion for FNT training. OR-ANGE FRAME: adding KL divergence loss for FNT adaptation

The adaptation loss with KL divergence is

$$\mathcal{J}_{adapt} = \operatorname{CE}(Z_{u}^{v}, Y_{adapt}) + \alpha \operatorname{KL}(Z_{u}^{v}, Z_{u}^{'v})$$
(8)

where Y_{adapt} is the adaptation text data, Z_u^v is the log softmax of adaptation text from the adapted model and $Z_u^{'v}$ is the log softmax of adaptation text from the baseline model. α is the KL divergence weight to be tuned in the experiments.

3.3. External language model

Since the vocabulary predictor in FNT is designed to be an LM, we explore the possibility of training it independently on a much larger text corpus than the transcriptions in the FNT training data. The parameters of this pre-trained external LM could be further updated to potentially improve accuracy. In principle, we could choose a variety of architectures for the external LM. In this paper, we limit ourselves to an architecture that is very close to that of the standard prediction network for a fair comparison with the baseline system. The vocabulary of the external LM is the same as that of the FNT system, and is trained using the conventional cross-entropy loss. Experimental results in Section 4 show that external LM trained with more data improves model accuracy, and updating the external LM parameters during FNT model training further improves the results.

3.4. N-gram interpolation

As noted before, fine-tuning the vocabulary predictor is one of the straightforward adaptation method for FNT. Although updating vocabulary predictor is much faster than updating the whole FNT network, it could not meet the immediate adaptation requirement for some applications. In this paper, we propose to use n-gram integration for fast adaptation of FNT.

In this method, a n-gram based language model is firstly trained with the adaptation text data. Then it is interpolated with the probability output from the vocabulary predictor during the decoding. The vocabulary log probability after interpolation with n-gram probability $P(y_u|y_1^{u-1})_{ngram}$ is calculated as

$$z_u^v = \log((1-w) * P(y_u|y_1^{u-1})_{pred} + w * P(y_u|y_1^{u-1})_{ngram})$$
(9)

where $P(y_u|y_1^{u-1})_{pred} = \operatorname{softmax}(d_u^v)$ is the label probability output from vocabulary predictor. Then z_u^v is plugged into Equation (6) to calculate $z_{t,u}^v$ which is used to generate the final output of FNT.

testing set	adaptation data	testing data
task1	6,135	6,269
task2	193,047	21,960
Librispeech	18,740,565	210,246

 Table 1. Word count for testing sets.

In this method, no neural network training is involved, and ngram LM model training is super fast with the adaptation text data. Experimental results in section 4.3 show it has much lower computational cost compared to the fine-tuning based method.

4. EXPERIMENTS

In this section, the effectiveness of the proposed methods are evaluated based on conformer-transducer (C-T) model [12] for several adaptation tasks with different amount of adaptation text data.

In the baseline C-T model, the encoder network contains 18 conformer layers. The prediction network contains 2 LSTM [33] layers, and 1024 nodes per layer. The output label size is 4000 sentence pieces. We use the low-latency streaming implementation in [10] with 160 milliseconds (ms) encoder lookahead. The standard FNT model has the same encoder structure and output label inventory as the baseline C-T model. The blank and vocabulary predictor each consists of 2 LSTM layers, also 1024 nodes per layer. The acoustic feature is 80-dimension log Mel filter bank for every 10 ms speech.

The training data contains 30 thousand (K) hours of transcribed Microsoft data, mixed with 8K and 16K HZ sampled data [34]. All the data are anonymized with personally identifiable information removed. The general testing set covers different application scenarios including dictation, conversation, far-field speech and call center etc., consisting of a total of 6.4 million (M) words. For the adaptation testing sets, we selected 2 real application tasks with different size of adaptation text data, as well as Librispeech sets for better reference. The data size of the testing sets are listed in table 1. The model training never observes the data from these adaptation tasks.

The external LM has the same model structure as the vocabulary predictor in the FNT and was trained with text data containing about 1.5 billion words, which includes the transcription of 30k training data mentioned above for C-T and FNT model training.

We first evaluate the FNT model's accuracy on general testing set with the proposed methods, including adding CTC criterion, initializing the vocabulary predictor from a well trained external LM. Then we examine the performance of above FNT models on adaptation sets by fine-tuning the vocabulary predictor with text adaptation data. Finally, n-gram interpolation adaptation method is evaluated based on the best FNT model from above experiments.

4.1. Results on general testing set

All results on general testing set are given in table 2. The baseline model is a standard C-T model. The standard FNT model is with the same structure as in [31], and it is trained with 30k training data from scratch. Compared with the baseline C-T model, the standard FNT model got 1.29% relative WER increase. To reduce the accuracy degradation, FNT model is refined by adding the CTC criterion based on equation 6 and 7. The CTC loss weight β is 0.1. We can see that adding CTC decreases WER from 11.01 to 10.97, but is still worse than the baseline C-T model. The degradation is reduced further by initializing vocabulary predictor from an external language model. Two recipes are examined: one is the external LM is fixed

during the FNT model training, and the other is the external LM is updated together with other parts of FNT model. Updating the external LM got the best result, which is even better than the baseline C-T model.

Model	General set
Baseline C-T (B0)	10.87
Standard FNT (F0)	11.01
+CTC (F1)	10.97
+ext. LM fix (F2)	10.89
+ext. LM update (F3)	10.70

Table 2. WER(%) on the general testing set.

4.2. Results on adaptation testing sets

In this section, We first evaluated the impact of KL divergence using Librispeech set based on standard FNT model. The results are given in table 3. The results showed the adapted model without KL divergence degraded the accuracy on general testing set largely. FNT adaptation with KL divergence helped to recovered the loss on general testing set obviously with very small WER increase on adaptation set compared to the standard FNT adaptation. In the following adaptation experiments, KL divergence weight α is always set to 0.1.

Table 4 shows the adaptation results for different FNT models on all adaptation sets. For each task, the results in "base" column are WER before adaptation, and the results in "adapt" column are WER after adaptation. Simple average WERs are also reported by averaging the WERs from all three tasks. Comparing "base" results for B0 and F0, we could find the accuracy gap between the baseline C-T model and the standard FNT model on these adaptation set is much larger than that on general set, especially for task1 and task2. The possible reason is that the domains in the training data may have some coverage for the general testing set, but they are totally irrelevant to these adaptation sets. With the proposed refinements for FNT, this gap is decreased step by step, the best FNT model F3, which is with full combination of the proposed methods, could get the similar accuracy as the baseline C-T model. The same trend could be also observed for the "adapt" results. Each method contributes accuracy improvement to the adapted FNT model. Compared with the adapted standard FNT model (F0), the adapted model F3 reduces WER for three adaptation sets by relatively 9.48% in average (from 12.80% to 11.59%). And compared with the baseline C-T model, the adapted model F3 gets 29.21% relative WER reduction (from 16.37% to 11.59%).

	Standard	KLD weight				
	FNT	0.0	0.1	0.2	0.3	
General set	10.87	12.78	11.69	11.52	11.4	
Librispeech	8.32	7.17	7.17	7.25	7.33	

Table 3.	WER(%) for Lit	orispeech	with	different	KLD	weights.
		/ -					

4.3. Results of n-gram interpolation

In this section, we evaluated n-gram interpolation performance based on the best FNT model (F3) for three adaptation tasks. For each task, a 5-gram LM is trained with the adaptation text. To make the interpolation simple and efficient, sentence piece instead of word is used as the basic unit for the 5-gram LM. The interpolation weight

Model	Task1		Task2		Librispeech		Simple average	
	Base	Adapt	Base	Adapt	Base	Adapt	Base	Adapt
B0	16.59	-	24.19	-	8.32	-	16.37	-
F0	17.45	11.8	25.21	19.43	8.44	7.17	17.03	12.80
F1	17.21	10.89	25.05	19.04	8.43	7.07	16.90	12.33
F2	17.63	10.8	24.96	18.55	8.41	7.16	17.00	12.17
F3	16.38	10.19	24.44	17.7	8.33	6.87	16.38	11.59
F3+n-gram	16.38	11.49	24.44	19.89	8.33	7.39	16.38	12.92

Table 4. WER(%) on adaptation testing sets.

w is always set as 0.3. The results are shown in the last row of table 4. Compared with the adaptation of fine-tuning vocabulary predictor, n-gram interpolation based adaptation got a little higher WER, but the relative WER reduction over the baseline C-T model is still satisfying, which is 21.04% (from 16.37% to 12.92%). More importantly, the adaptation speed is improved hugely. Experiments show that for fine-tuning method, the adaptation process cost about 10 seconds per 1,000 words on GPU and the cost is formidable when adapting with CPU. In contrast, it only needs about 0.002 seconds per 1,000 words on CPU for n-gram LM training. This is very useful for those application scenarios which need immediate adaptation.

5. CONCLUSIONS

In this paper, several methods are proposed to improve the accuracy and efficiency for FNT adaptation with text-only data. These methods include: 1) during the FNT model training, adding CTC criterion to make the encoder act more like an acoustic model and initializing vocabulary predictor with a well trained external LM to use more text data. 2) during the FNT model adaptation, adding KL divergence to avoid over fitting to the adaptation data. 3) using n-gram interpolation with the vocabulary-prediction LM module inside FNT instead of fine-tuning the vocabulary-prediction LM module to improve the adaptation speed. The experimental results proved that, compared with standard FNT, the proposed methods could get better accuracy on general testing set, and decrease the adaptation WER by 9.48% percent relatively. In total, compared with the baseline C-T model, the adaptation WER is decreased by 29.21% relatively. Besides, n-gram interpolation could get much faster adaptation than the fine-tuning method, enabling the scenarios which require immediate adaptation.

6. REFERENCES

- [1] Alex Graves, "Sequence transduction with recurrent neural networks," in *International Conference on Machine Learning: Representation Learning Workshop*, 2012.
- [2] Yanzhang He, Tara N Sainath, Rohit Prabhavalkar, Ian Mc-Graw, Raziel Alvarez, Ding Zhao, David Rybach, et al., "Streaming end-to-end speech recognition for mobile devices," in *Proc. ICASSP*, 2019, pp. 6381–6385.
- [3] 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.
- [4] Ching-Feng Yeh, Jay Mahadeokar, et al., "Transformertransducer: End-to-end speech recognition with self-attention," arXiv preprint arXiv:1910.12977, 2019.
- [5] Qian Zhang, Han Lu, Hasim Sak, Anshuman Tripathi, Erik McDermott, Stephen Koo, and Shankar Kumar, "Transformer

transducer: A streamable speech recognition model with transformer encoders and RNN-T loss," *Proc. ICASSP*, 2020.

- [6] Jinyu Li, Rui Zhao, Hu Hu, and Yifan Gong, "Improving RNN transducer modeling for end-to-end speech recognition," in *Proc. ASRU*, 2019.
- [7] Eric Battenberg, Jitong Chen, Rewon Child, Adam Coates, Yashesh Gaur, Yi Li, Hairong Liu, and et al., "Exploring neural transducers for end-to-end speech recognition," in *Proc. ASRU*, 2017.
- [8] Chung-Cheng Chiu, Tara N. Sainath, Yonghui Wu, Rohit Prabhavalkar, Patrick Nguyen, Zhifeng Chen, Anjuli Kannan, et al., "State-of-the-art speech recognition with sequence-tosequence models," in *Proc. ICASSP*, 2018, pp. 4774–4778.
- [9] 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.
- [10] 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 *Proc. ICASSP*, 2021.
- [11] Jinyu Li, "Recent advances in end-to-end automatic speech recognition," APSIPA Transactions on Signal and Information Processing, vol. 11, no. 1, 2022.
- [12] Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, and et al., "Conformer: Convolution-augmented transformer for speech recognition," in *Proc. Interspeech*, 2020, pp. 5036–5040.
- [13] 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.
- [14] Tara N. Sainath, Yanzhang He, Bo Li, Arun Narayanan, Ruoming Pang, Antoine Bruguier, Shuo yiin Chang, and 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.
- [15] Jinyu Li, Rui Zhao, Zhong Meng, Yanqing Liu, Wenning Wei, Sarangarajan Parthasarathy, Vadim Mazalov, and et. al., "Developing RNN-T models surpassing high-performance hybrid models with customization capability," in *Proc. Interspeech*, 2020.
- [16] Khe Chai Sim, Francoise Beaufays, et al., "Personalization of end-to-end speech recognition on mobile devices for named entities," in *Proc. ASRU*, 2019.
- [17] Yan Deng, Rui Zhao, Zhong Meng, Xie Chen, Bin Liu, Jinyu Li, Yifan Gong, and Lei He, "Improving RNN-T for domain scaling using semi-supervised training with neural TTS," in *Proc. Interspeech*, 2021.
- [18] Xianrui Zheng, Yulan Liu, Deniz Gunceler, and Daniel Willett, "Using synthetic audio to improve the recognition of outof-vocabulary words in end-to-end ASR systems," in *Proc. ICASSP*, 2021.
- [19] Amin Fazel, Wei Yang, Yulan Liu, Roberto Barra-Chicote, Yixiong Meng, Roland Maas, and Jasha Droppo, "SynthASR: Unlocking Synthetic Data for Speech Recognition," in *Proc. Interspeech*, 2021.

- [20] Rui Zhao, Jian Xue, Jinyu Li, Wenning Wei, Lei He, and Yifan Gong, "On addressing practical challenges for RNN-Transducer," in *Proc. ASRU*, 2021.
- [21] Anjuli Kannan, Yonghui Wu, Patrick Nguyen, Tara N. Sainath, Zhifeng Chen, and Rohit Prabhavalkar, "An analysis of incorporating an external language model into a sequence-tosequence model," in *Proc. ICASSP*, 2018.
- [22] Vijay Ravi, Yile Gu, Ankur Gandhe, Ariya Rastrow, Linda Liu, Denis Filimonov, Scott Novotney, and Ivan Bulyko, "Improving accuracy of rare words for rnn-transducer through unigram shallow fusion," arXiv preprint arXiv:2012.00133, 2020.
- [23] Rodrigo Cabrera, Xiaofeng Liu, Mohammadreza Ghodsi, Zebulun Matteson, Eugene Weinstein, and Anjuli Kannan, "Language model fusion for streaming end to end speech recognition," arXiv preprint arXiv:2104.04487, 2021.
- [24] Chhavi Choudhury, Ankur Gandhe, Xiaohan Ding, and Ivan Bulyko, "A likelihood ratio-based domain adaptation method for end-to-end models," in *Proc. ICASSP*, 2022.
- [25] Duc Le, Mahaveer Jain, Suyoun Kim Gil Keren, Yangyang Shi, Julian Chan Jay Mahadeokar, and et. al., "Contextualized streaming end-to-end speech recognition with trie-based deep biasing and shallow fusion," in *Proc. Interspeech*, 2020.
- [26] Erik McDermott, Hasim Sak, and Ehsan Variani, "A density ratio approach to language model fusion in end-to-end automatic speech recognition," in *Proc. ASRU*, 2019.
- [27] Ehsan Variani, David Rybach, Cyril Allauzen, and Michael Riley, "Hybrid autoregressive transducer (HAT)," in *Proc. ICASSP*, 2020, pp. 6139–6143.
- [28] Zhong Meng, Sarangarajan Parthasarathy, Eric Sun, Yashesh Gaur, Naoyuki Kanda, Liang Lu, Xie Chen, and et. al., "Internal language model estimation for domain-adaptive end-to-end speech recognition," in *Proc. SLT*, 2021.
- [29] Vinit Unni, Shreya Khare, Ashish Mittal, Preethi Jyothi, Sunita Sarawagi, and Samarth Bharadwaj, "Adaptive discounting of implicit language models in rnn-transducers," in *Proc. ICASSP*, 2022.
- [30] Zhong Meng, Naoyuki Kanda, Yashesh Gaur, Sarangarajan Parthasarathy, Eric Sun, Liang Lu, Xie Chen, and et. al., "Internal language model training for domain-adaptive end-to-end speech recognition," in *Proc. ICASSP*, 2021.
- [31] Xie Chen, Zhong Meng, Sarangarajan Parthasarathy, and Jinyu Li, "Factorized neural transducer for efficient language model adaptation," in *Proc. ICASSP*, 2022.
- [32] 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.
- [33] Sepp Hochreiter and Jurgen Schmidhuber, "Long Short-Term Memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [34] Jinyu Li, Dong Yu, Jui-Ting Huang, and Yifan Gong, "Improving wideband speech recognition using mixed-bandwidth training data in CD-DNN-HMM," in *Proc. SLT*, 2012, pp. 131– 136.