



Improved Factorized Neural Transducer Model For Text-only Domain Adaptation

Junzhe Liu¹, Jianwei Yu², Xie Chen^{1,*}

¹ MoE Key Lab of Artificial Intelligence, AI Institute, X-LANCE Lab, Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China

² Tencent AI Lab, Shenzhen, China

junzheliu@sjtu.edu.cn, tomasyu@tencent.com, chenxie@sjtu.edu.cn

Abstract

Adapting End-to-End ASR models to out-of-domain datasets with text data is challenging. Factorized neural Transducer (FNT) aims to address this issue by introducing a separate vocabulary decoder to predict the vocabulary. Nonetheless, this approach has limitations in fusing acoustic and language information seamlessly. Moreover, a degradation in word error rate (WER) on the general test sets was also observed, leading to doubts about its overall performance. In response to this challenge, we present the improved factorized neural Transducer (IFNT) model structure designed to comprehensively integrate acoustic and language information while enabling effective text adaptation. We assess the performance of our proposed method on English and Mandarin datasets. The results indicate that IFNT not only surpasses the neural Transducer and FNT in baseline performance in both scenarios but also exhibits superior adaptation ability compared to FNT. On source domains, IFNT demonstrated statistically significant accuracy improvements, achieving a relative enhancement of 1.2% to 2.8% in baseline accuracy compared to the neural Transducer. On out-of-domain datasets, IFNT shows relative WER(CER) improvements of up to 30.2% over the standard neural Transducer with shallow fusion, and relative WER(CER) reductions ranging from 1.1% to 2.8% on test sets compared to the FNT model.

Index Terms: neural Transducer, text-only domain adaptation, end-to-end speech recognition, language model

1. Introduction

In recent years, end-to-end (E2E) [1] based models have gained great interest in automatic speech recognition (ASR) systems. Compared to traditional hybrid systems, E2E systems such as connectionist temporal classification (CTC) [2], attention-based encoder-decoder (AED) [3], and neural Transducer (NT) [4] predict word sequences using a single neural network. When there is a mismatch between the trained domain and the test domain, a significant degradation in accuracy is observed. Conventional domain adaptation methods [5, 6] typically rely on speech-text pairs from the target domain. However, collecting a large amount of speech-text matching data from the target domain is difficult, while obtaining text-only data is relatively easier. As a result, text-only adaptive methods have been widely proposed and studied [7, 8]. Since E2E systems are jointly optimized, there is no separate component that solely performs as a language model (LM), making it challenging to directly apply common LM adaptation methods.

One feasible solution is to fine-tune the E2E model using the synthesized audio-transcript pairs generated by a text-to-

speech (TTS) model [9, 10, 11], but this approach is computationally expensive. Another common practice is LM fusion [12, 13, 14, 15], such as shallow fusion [16], deep fusion [17], and cold fusion [18]. Among them, the most widely used is shallow fusion, which combines the E2E model score and the external LM score in the log-linear domain during beam search. Methods like density ratio [19] also work similarly. Internal language model estimation [20, 21] was also proposed recently, it calculates the interpolated log-likelihood score based on the maximum scores from the internal LM and the external LM respectively during decoding. However, LM fusion involves interpolation weights that are task-dependent and require tuning, making the performance sensitive to the weight selection.

There have been increasing research efforts [22, 23, 24] to modify the structure of neural Transducers. Factorized neural Transducer (FNT) [25, 26, 14, 27, 28] addresses this issue by introducing a standalone LM for vocabulary prediction, enabling the direct application of conventional LM adaptation methods. While the results have shown promising adaptation abilities, minor accuracy degradation has also been observed on the general test set compared to the standard neural Transducer baseline, raising doubts about their overall performance.

Building upon this work, we propose the improved factorized neural Transducer (IFNT) model that can better combine acoustic information with language information, and improve its baseline accuracy both before and after adapting to text data. Our proposed model incorporates an internal LM alongside the standard neural Transducer model, with the LM posterior probability directly integrated into the vocabulary prediction. Utilizing a standalone LM in our model facilitates the application of various LM adaptation methods for fast text-only domain adaptation, similar to the hybrid system. We validate the proposed method in both English and Mandarin datasets, both in-domain and out-of-domain scenarios results demonstrate the superior performance of our proposed model over the FNT and standard neural Transducer shallow fusion methods.

2. Standard and Factorized Neural Transducer Models

2.1. Standard neural Transducer

ASR systems predicts a conditional distribution over blank-augmented token sequences $\hat{\mathbf{Y}} = \{\hat{y}_1, \dots, \hat{y}_{T+U}\}$, where T, U are acoustic and label sequence lengths, $\hat{y}_i \in \mathcal{V} \cup \phi$, and \mathcal{V}, ϕ denotes vocabulary and blank respectively. The standard neural Transducer model can be divided into three parts: an acoustic encoder, which takes the acoustic feature \mathbf{x}_1^t and generates the acoustic representation \mathbf{f}_t ; a label decoder, which consumes the history of the previously predicted label sequence \mathbf{y}_1^u and com-

* Corresponding author

puts the label representation \mathbf{g}_u ; and a joint network, which takes both representations and combines them to compute the probability distribution over $\mathcal{V} \cup \phi$:

$$P(\hat{y}_{t+1} | \mathbf{x}_1^t, \mathbf{y}_1^u) = \text{softmax}(\sigma(\mathbf{f}_t + \mathbf{g}_u)) \quad (1)$$

σ denotes some non-linear activation function, e.g. relu. In Figure 1(a), the part enclosed by the gray dashed box is the joint network. It first projects the outputs of the encoder and decoder to a joint dimension (represented as D in the figure), then adds the two together, and finally maps the output to a $V + 1$ dimension corresponding to the vocabulary plus the blank token. The objective function of the neural Transducer is to minimize the negative log probability over all possible alignments, which could be written as:

$$\mathcal{J}_t = -\log P(\mathbf{Y} | \mathbf{x}) = -\log \sum_{\alpha \in \beta^{-1}(\mathbf{y})} P(\alpha | \mathbf{x}) \quad (2)$$

where β is the function to convert the alignment α to label sequence \mathbf{Y} by removing the blank ϕ .

2.2. LM shallow fusion

In shallow fusion, an LM trained on target domain training text is integrated with the E2E model during inference to optimize a log-linear interpolation between the E2E and LM probabilities. The optimal token sequence \mathbf{Y} is obtained via beam search:

$$\mathbf{Y} = \arg \max_{\mathbf{Y}} \left[\log P(\mathbf{Y} | \mathbf{X}; \theta_{\text{E2E}}^S) + \lambda_T \log P(\mathbf{Y}; \theta_{\text{LM}}^T) \right] \quad (3)$$

where $P(\mathbf{Y}; \theta_{\text{LM}}^T)$ is the posterior probability given by the external LM, and λ_T is a hyper-parameter for tuning.

2.3. Factorized Neural Transducer

Although the neural Transducer’s structure includes a label decoder, it is not entirely equivalent to a language model. Because the output of the decoder is a high dimensional representation of the tokens, rather than a posterior distribution; Moreover, typical LM only predicts the vocabulary tokens \mathcal{V} , while the joint network also has to predict the blank token ϕ .

Recognizing this distinction, the factorized neural Transducer (FNT) adopts a structure consisting of two separate decoders. As shown in Fig 1(b), the original joint network portion (enclosed in the gray box) remains the same as the standard neural Transducer structure, with the exception that the projection layer maps the output to a dimension of 1, predicting only the blank token. Consequently, this decoder is referred to as the blank decoder. In the vocabulary section, FNT introduces a separate language model component into the model, generating a probability distribution over \mathcal{V} . The encoder output is projected to a dimension of V (indicated by the yellow projection layer). Acoustic and label information are then combined at the logit level to predict vocabulary tokens. The logit for the blank token and the vocabulary logits are concatenated to compute the Transducer loss. The total training loss of FNT is expressed as:

$$\mathcal{J}_f = \mathcal{J}_t - \lambda_f \log P_{LM}(\mathbf{Y}) \quad (4)$$

where the first term is the Transducer loss, and the second term is the language model loss with cross-entropy. λ_f is a hyper-parameter for LM tuning. Since the vocabulary decoder works as a standalone language model, we could use the target domain’s text data to fine-tune this part directly.

3. Improved factorized neural Transducer

Despite the good adaptability of FNT, an accuracy degradation on general test sets was also observed. To narrow this gap, we propose the IFNT model with enhanced performance and adaptability. We noted that FNT employs a different approach in combining acoustic and label information compared to the standard neural Transducer, in which the encoder and decoder output are fused in the joint dimension space (D -dimension) before mapping to the vocabulary size. In contrast, FNT directly adds both information together in the vocabulary space (V -dimension) on a logit basis. Inspired by this difference, we revert to the standard neural Transducer style of integrating these two types of information in our proposed IFNT model. An illustration of our IFNT model is presented in Fig 1(c).

We introduced several major modifications in IFNT: Firstly, we apply a sigmoid layer to the output of the vocabulary decoder. Our goal is to constrain the distribution of text vectors in the feature space, thereby facilitating faster convergence during adaptation. Secondly, LM output is projected to the joint dimension D and then fused with the encoder output, following the standard neural Transducer approach. Furthermore, we directly incorporate the LM posterior probability over \mathcal{V} into the final output of the vocabulary component. Including the probability distribution in the final layer helps ensure the preservation and effective integration of essential linguistic information with acoustic features, resulting in improved recognition accuracy. The total training loss is the same as FNT.

Compared to deep fusion [17] and cold fusion [18], our proposed IFNT model stands out in that the E2E model and LM component are jointly trained from scratch, while deep fusion and cold fusion rely on pre-trained LMs for integration. Moreover, the most significant distinction lies in the direct incorporation of the LM posterior probability into the final output, without which the IFNT model would lose its adaptation ability. We have also explored the use of a pre-trained language model for LM initialization in IFNT. However, experiments revealed that training the whole model from scratch yields more stable results. This may be because joint training allows the LM to synchronize more effectively with the acoustic encoder.

4. Experiments

4.1. Datasets

We evaluate our model on both English and Mandarin datasets. In the English context, we utilize the GigaSpeech-M [29] subset as our training set, comprising 1,000 hours of speech data. For text-only domain adaptation, we have chosen three target domain datasets: EuroParl [30], TED-LIUM [31], and a Medical [32] dataset. In the Mandarin scenario, we employ the Wenetspeech [33] train-M subset, which encompasses 1,000 hours of speech data, as our training set. Additionally, for out-of-domain adaptation, we incorporate four datasets: Thchs-30 [34], Aishell-1 [35], Aishell-2 (iOS) [36], and Aishell-4 (L subset) [37]. Detailed dataset statistics can be found in table 1.

4.2. Experiment setups

We build 80-dimensional mel filterbank features with global-level cepstral mean and variance normalization for acoustic feature extraction. Regarding the model structure, the encoder consists of 12 Conformer [38] layers. The inner size of the feed-forward layer is 2,048, and the attention dimension is 512 with 8 heads. To maintain a roughly equal number of param-

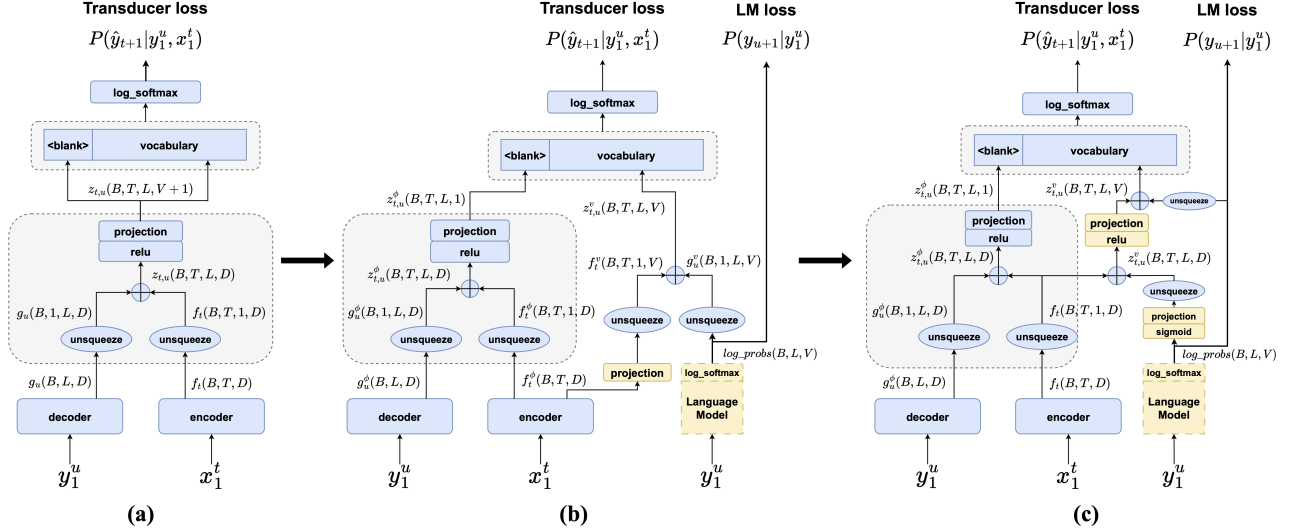


Figure 1: The illustration of three model structures: (a) standard neural Transducer (NT). (b) factorized neural Transducer (FNT). (c) proposed improved factorized neural Transducer (IFNT).

Table 1: Statistics of the source domain dataset (GigaSpeech, Wenetspeech) and target domain datasets including their durations in hours (h) and number of training text sentences.

Dataset	#sentences	Train	Dev	Test
GigaSpeech-M	909,401	1000.0	12.0	40.0
EuroParl	32,584	77.2	3.1	2.9
TED-LIUM	268,263	346.2	3.7	3.8
Medical	20,107	35.1	4.9	5.2
Wenetspeech-M	1,514,500	1000.0	20.0	38.0
Thchs-30	10,001	27.1	-	6.2
Aishell-1	120,099	150.0	10.0	5.0
Aishell-2 (iOS)	1,009,224	1000.0	2.5	5.1
Aishell-4 (L)	12,494	16.2	-	12.7

eters across three models, we keep the encoder part consistent while modifying the decoder’s structural parameters, and the final number of parameters of the models remains around 106M. The hyperparameters are tuned for the best accuracy, and the final choices are $\lambda_T = 0.1$ and $\lambda_f = 0.1$. Our experiments are implemented using the fairseq [39] framework. During adaptation, factorized Transducer models are fine-tuned using the target domain’s training text and then evaluated on the test sets. Our training utilizes fp32 precision and the Adam optimizer. All models are trained under the same training configuration. We apply a model averaging over 5 checkpoints during inference, and beam search with a beam size of 5 is used.

4.3. In-domain evaluation

Table 2 illustrates the accuracy of the three models on the source domain test sets. IFNT exhibits superior baseline accuracy over both Transducer and FNT across English and Mandarin datasets, achieving noteworthy relative WER(CER) reductions ranging from 2.5% to 4.6% compared to FNT, and 1.2% to 2.8% compared to the neural Transducer.

Table 2: Comparison of ASR accuracy on GigaSpeech (WER %) and Wenetspeech (CER %) between standard neural Transducer, FNT, and IFNT. Model parameters are all 106M.

Model	GigaSpeech		Wenetspeech	
	dev	test	dev	test_net
① Transducer	13.96	13.99	10.80	11.25
② FNT	14.11	14.19	11.29	11.76
③ IFNT (proposed)	13.67	13.84	10.77	11.23

4.4. Out-of-domain evaluation

In the out-of-domain scenarios, we evaluate the accuracy of the models both before and after adaptation to training texts. For the standard neural Transducer, we incorporate an external language model trained on the training text with a parameter of $4M$ for shallow fusion. Results are presented in Table 3.

In the English scenario, IFNT exhibits significant improvements across three datasets. Prior to text-only adaptation, IFNT outperforms FNT in baseline results and surpasses the neural Transducer, achieving up to 2.6% relative WER reductions. After fine-tuning with target domain training text, IFNT demonstrates substantial relative WER reductions of 9.5%, 30.2%, and 20.7% compared to the neural Transducer with shallow fusion on the respective test sets. In comparison to FNT, IFNT displays enhanced adaptation capabilities, resulting in relative WER reductions of 2.8%, 2.5%, and 1.1% respectively.

For the Mandarin datasets, IFNT has surpassed the baseline accuracy of FNT entirely and demonstrated better results over the conventional neural Transducer on three datasets. Following text adaptation, IFNT achieves the lowest CER results, exhibiting a relative CER reduction of 1.1% to 9.8% compared to the neural Transducer with shallow fusion. Furthermore, in comparison to FNT, IFNT demonstrates a relative CER improvement of up to 2.5% after completing text adaptation.

Table 3: Comparison of Perplexity (PPL) on the dev set and ASR accuracy among the standard neural Transducer, FNT, and IFNT model before and after adaptation. Upper table: Results for English datasets (WER %). Lower table: Results for Mandarin datasets (CER %). The number of parameters for all three models is equally set to 106M, with the external language model used for shallow fusion comprising 4M parameters.

Model	text adaptation	EuroParl			TED-LIUM			Medical			Avg. WER
		PPL	dev	test	PPL	dev	test	PPL	dev	test	
① neural Transducer	✗	-	20.22	20.81	-	11.17	11.90	-	29.12	28.47	20.28
② FNT	✗	107.31	20.53	21.21	120.13	11.12	11.91	173.12	29.39	28.64	20.47
③ IFNT (proposed)	✗	102.91	20.21	20.69	123.30	11.11	11.88	158.33	28.63	28.20	20.12
① + shallow fusion	✓	55.89	19.86	20.19	67.83	10.86	11.71	18.46	28.09	27.10	19.64
② + text-only finetune	✓	37.65	18.15	18.80	61.03	8.55	8.38	12.06	22.16	21.51	16.26
③ + text-only finetune	✓	39.84	18.09	18.27	60.44	8.21	8.17	11.61	21.91	21.50	16.02

Model	text adaptation	Thchs-30		Aishell-1			Aishell-2			Aishell-4		Avg. CER
		PPL	test	PPL	dev	test	PPL	dev	test	PPL	test	
① neural Transducer	✗	-	10.19	-	5.57	6.60	-	6.37	6.39	-	32.38	11.25
② FNT	✗	289.31	10.51	73.25	5.97	6.67	68.93	6.77	6.63	81.24	31.59	11.36
③ IFNT (proposed)	✗	275.61	10.07	73.11	5.73	6.49	69.90	6.64	6.49	77.91	31.29	11.12
① + shallow fusion	✓	183.11	10.01	57.81	5.27	6.14	58.14	6.27	6.24	67.01	30.97	10.82
② + text-only finetune	✓	214.16	9.99	43.80	4.97	5.57	47.28	6.43	6.24	39.09	29.91	10.52
③ + text-only finetune	✓	208.11	9.84	46.89	4.90	5.54	50.44	6.27	6.17	38.64	29.88	10.43

4.5. Ablation study and analysis

We tried to re-model the decoder network as a language model thereby eliminating the need for a separate third component. However, this alteration did not result in any enhancement in the overall accuracy, nor did it succeed in adapting to various domains. The findings clearly show that a distinct language model component is essential for the desired adaptability.

We then conducted ablation experiments to validate the efficacy of the modifications introduced in IFNT by removing the `log_probs` directly injected into the final projection layer, which is denoted as *IFNT w/ lprobs* and index 4. The most significant distinction between this model and FNT lies in the structure of vocabulary prediction. The baseline results are presented in Table 4, indicating that the structural modifications indeed led to improvements in baseline accuracy compared to FNT.

To assess its adaptation ability, we chose EuroParl and Aishell-1 as out-of-domain test sets. Experimental results reveal a substantial decline in IFNT’s adaptation capability upon the removal of `log_probs`, emphasizing the crucial role of incorporating `log_probs` for maintaining adaptation capability.

In summary, our modifications in IFNT prove beneficial to the model’s baseline and adaptation results, aligning with the findings of [40] that earlier fusion of acoustic and linguistic information leads to better performance in Transducer models.

4.6. Limitations

Despite our proposed IFNT exhibiting impressive baseline and adaptation performances, it still presents certain limitations: 1) It does not leverage audio data for adaptation, and 2) It still remains confined within the Transducer framework structure. Future exploration could consider re-designing the blank token ϕ for further improvements.

Table 4: Ablation study: ④ *IFNT w/ lprobs* denotes the IFNT model after removing the `log_probs`. Upper table: Baseline results on source domain datasets. Lower table: Adaptation results on target domain datasets.

Model	GigaSpeech		Wenetspeech	
	dev	test	dev	test_net
② FNT	14.11	14.19	11.29	11.76
④ IFNT w/ lprobs	13.48	13.49	11.26	11.75
③ IFNT (proposed)	13.67	13.84	10.77	11.23

Model	EuroParl			Aishell-1		
	PPL	dev	test	PPL	dev	test
④	97.89	20.16	20.25	87.23	6.01	6.93
④ + adapt	47.66	20.03	20.21	56.84	6.01	6.90

5. Conclusion

In this work, we proposed the IFNT model. By redesigning the model structure of the FNT, we addressed the drawback of FNT experiencing a drop in baseline accuracy compared to the neural Transducer, while also enhancing its capability for text-only domain adaptation. On both Mandarin and English datasets, IFNT demonstrated statistically significant accuracy improvements, achieving a relative enhancement of 1.2% to 2.8% in baseline accuracy compared to the neural Transducer. On out-of-domain datasets, its adaptation capability demonstrated a relative reduction of up to 2.8% compared to FNT and up to 30.2% compared to the neural Transducer.

6. Acknowledgements

This work was supported by the National Natural Science Foundation of China (No.62206171 and No.U23B2018), Shanghai Municipal Science and Technology Major Project under Grant 2021SHZDZX0102, and the International Cooperation Project of PCL.

7. References

- [1] J. Li, "Recent advances in end-to-end automatic speech recognition," *ArXiv*, 2022.
- [2] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks," in *Proc. ICML*, 2006.
- [3] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell," *arXiv preprint:1508.01211*, 2015.
- [4] A. Graves, "Sequence Transduction with recurrent neural networks," *ArXiv*, vol. abs/1211.3711, 2012.
- [5] P. Bell, J. Fainberg, O. Klejch, J. Li, S. Renals, and P. Swietojanski, "Adaptation algorithms for neural network-based speech recognition: An overview," *IEEE Open Journal of Signal Processing*, 2021.
- [6] K. Deng and P. C. Woodland, "Adaptable end-to-end asr models using replaceable internal lms and residual softmax," *ArXiv*, vol. abs/2302.08579, 2023.
- [7] J. Pytköinen, A. Ukkonen, J. Kilpikoski, S. Tamminen, and H. Heikinheimo, "Fast text-only domain adaptation of rnn-transducer prediction network," *preprint:2104.11127*, 2021.
- [8] C. Choudhury, A. Gandhe, X. Ding, and I. Bulyko, "A likelihood ratio-based domain adaptation method for end-to-end models," in *Proc. ICASSP*, 2022.
- [9] K. C. Sim, F. Beaufays, A. Benard, D. Guliani *et al.*, "Personalization of end-to-end speech recognition on mobile devices for named entities," in *IEEE Proc. ASRU*, 2019.
- [10] X. Zheng, Y. Liu, D. Gunceler, and D. Willett, "Using synthetic audio to improve the recognition of out-of-vocabulary words in end-to-end asr systems," in *IEEE Proc. ICASSP*, 2021.
- [11] Y. Deng, R. Zhao, Z. Meng, X. Chen, B. Liu, J. Li, Y. Gong, and L. He, "Improving RNN-T for domain scaling using semi-supervised training with neural TTS," in *Proc. Interspeech*, 2021.
- [12] R. Cabrera, X. Liu, M. Ghodsi, Z. Matteson, E. Weinstein, and A. Kannan, "Language model fusion for streaming end to end speech recognition," *arXiv preprint:2104.04487*, 2021.
- [13] S. Toshniwal, A. Kannan, C.-C. Chiu, Y. Wu, T. N. Sainath, and K. Livescu, "A comparison of techniques for language model integration in encoder-decoder speech recognition," in *Proc. SLT*, 2018.
- [14] M. Levit, S. Parthasarathy, C. Aksoylar, M. S. Rasooli, and S. Chang, "External language model integration for factorized neural transducers," *preprint:2305.17304*, 2023.
- [15] Y. Li, Y. Wu, J. Li, and S. Liu, "Prompting large language models for zero-shot domain adaptation in speech recognition," in *2023 IEEE ASRU*. IEEE, 2023, pp. 1–8.
- [16] A. Kannan, Y. Wu, P. Nguyen, T. N. Sainath, Z. Chen, and R. Prabhavalkar, "An analysis of incorporating an external language model into a sequence-to-sequence model," *Proc. ICASSP*, 2018.
- [17] C. Gulcehre, O. Firat, K. Xu, K. Cho *et al.*, "On using monolingual corpora in neural machine translation," *arXiv preprint:1503.03535*, 2015.
- [18] A. Sriram, H. Jun, S. Sathesh, and A. Coates, "Cold fusion: Training seq2seq models together with language models," in *Proc. Interspeech*, 2018.
- [19] E. McDermott, H. Sak, and E. Variani, "A density ratio approach to language model fusion in end-to-end automatic speech recognition," *Proc. ASRU*, 2019.
- [20] W. Zhou, Z. Zheng, R. Schlüter, and H. Ney, "On language model integration for RNN Transducer based speech recognition," in *Proc. ICASSP*, 2022.
- [21] Z. Meng, N. Kanda, Y. Gaur, S. Parthasarathy, E. Sun, L. Lu, X. Chen, J. Li, and Y. Gong, "Internal language model training for domain-adaptive end-to-end speech recognition," *Proc. ICASSP*, 2021.
- [22] E. Variani, D. Rybach, C. Allauzen, and M. Riley, "Hybrid autoregressive Transducer (HAT)," *Proc. ICASSP*, 2020.
- [23] X. Gong, W. Wang, H. Shao, X. Chen, and Y. Qian, "Factorized AED: Factorized attention-based encoder-decoder for text-only domain adaptive ASR," in *Proc. ICASSP*, 2023.
- [24] Z. Meng, T. Chen, R. Prabhavalkar, Y. Zhang *et al.*, "Modular hybrid autoregressive Transducer," in *Proc. SLT*, 2023.
- [25] R. Zhao, J. Xue, P. Parthasarathy, V. Miljanic, and J. Li, "Fast and accurate factorized neural Transducer for text adaption of end-to-end speech recognition models," in *Proc. ICASSP*, 2023.
- [26] X. Chen, Z. Meng, S. Parthasarathy, and J. Li, "Factorized neural Transducer for efficient language model adaptation," *Proc. ICASSP*, 2021.
- [27] X. Gong, Y. Wu, J. Li, S. Liu, R. Zhao, X. Chen, and Y. Qian, "Longfnt: Long-form speech recognition with factorized neural transducer," in *IEEE ICASSP 2023*. IEEE, 2023, pp. 1–5.
- [28] D. Le, F. Seide, Y. Wang, Y. Li, K. Schubert, O. Kalinli, and M. L. Seltzer, "Factorized blank thresholding for improved runtime efficiency of neural transducers," in *IEEE ICASSP 2023*, pp. 1–5.
- [29] G. Chen, S. Chai, G. Wang, J. Du *et al.*, "Gigaspeech: An evolving, multi-domain asr corpus with 10,000 hours of transcribed audio," *arXiv preprint:2106.06909*, 2021.
- [30] P. Koehn, "Europarl: A parallel corpus for statistical machine translation," in *Proc. Machine Translation Summit X: Papers*, 2005.
- [31] A. Rousseau, P. Deléglise, and Y. Estève, "TED-LIUM: an automatic speech recognition dedicated corpus," in *Proc. LREC*, 2012.
- [32] F. Fareez, T. Parikh, C. Wavell, S. Shahab *et al.*, "A dataset of simulated patient-physician medical interviews with a focus on respiratory cases," *Scientific Data*, 2022.
- [33] B. Zhang, H. Lv, P. Guo, Q. Shao, C. Yang, L. Xie, X. Xu, H. Bu, X. Chen, C. Zeng *et al.*, "Wenetspeech: A 10000+ hours multi-domain mandarin corpus for speech recognition," in *IEEE ICASSP 2022*. IEEE, 2022, pp. 6182–6186.
- [34] D. Wang and X. Zhang, "Thchs-30: A free chinese speech corpus," *arXiv preprint:1512.01882*, 2015.
- [35] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, "Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline," in *2017 20th conference of the oriental chapter of the international coordinating committee on speech databases and speech I/O systems and assessment*. IEEE, 2017, pp. 1–5.
- [36] J. Du, X. Na, X. Liu, and H. Bu, "Aishell-2: Transforming mandarin asr research into industrial scale," *arXiv preprint:1808.10583*, 2018.
- [37] Y. Fu, L. Cheng, S. Lv, Y. Jv, Y. Kong, Z. Chen, Y. Hu, L. Xie, J. Wu, H. Bu *et al.*, "Aishell-4: An open source dataset for speech enhancement, separation, recognition and speaker diarization in conference scenario," *arXiv preprint:2104.03603*, 2021.
- [38] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented transformer for speech recognition," in *Proc. Interspeech*, 2020.
- [39] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, "fairseq: A fast, extensible toolkit for sequence modeling," in *Proc. NAACL-HLT*, 2019.
- [40] A. Graves, r. Mohamed, Abdel, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *2013 IEEE ICASSP*, pp. 6645–6649.