



Adaptive Contextual Biasing for Transducer Based Streaming Speech Recognition

Tianyi Xu*, Zhanheng Yang*, Kaixun Huang, Pengcheng Guo, Ao Zhang, Biao Li, Changru Chen, Chao Li, Lei Xie†

Audio, Speech and Language Processing Group (ASLP@NPU), School of Computer Science, Northwestern Polytechnical University, Xi'an, China

xutianyi@mail.nwpu.edu.cn, lxie@nwpu.edu.cn

Abstract

By incorporating additional contextual information, deep biasing methods have emerged as a promising solution for speech recognition of personalized words. However, for real-world voice assistants, always biasing on such personalized words with high prediction scores can significantly degrade the performance of recognizing common words. To address this issue, we propose an adaptive contextual biasing method based on Context-Aware Transducer (CATT) that utilizes the biased encoder and predictor embeddings to perform streaming prediction of contextual phrase occurrences. Such prediction is then used to dynamically switch the bias list on and off, enabling the model to adapt to both personalized and common scenarios. Experiments on LibriSpeech and internal voice assistant datasets show that our approach can achieve up to 6.7% and 20.7% relative reduction in WER and CER compared to the baseline respectively, mitigating up to 96.7% and 84.9% of the relative WER and CER increase for common cases. Furthermore, our approach has a minimal performance impact in personalized scenarios while maintaining a streaming inference pipeline with negligible RTF increase.

Index Terms: End-to-end Speech Recognition, RNN-T, Context-Aware Training, Contextual List Filtering

1. Introduction

Recent advancements in deep learning have propelled end-to-end automatic speech recognition (E2E ASR) systems to new heights, making them the backbone of many speech recognition applications. E2E ASR methods such as connectionist temporal classification (CTC) [1, 2, 3], recurrent neural network Transducer (RNN-T) [4, 5, 6], and attention-based encoder-decoder (AED) [7, 8, 9], have been widely adopted. However, since data coverage greatly impacts deep learning techniques, for sentences containing words that rarely occur in training data, ASR performance is often severely impaired. These rare words include but not limited to contact names, location names, internet buzzwords, or product names.

Generally, there are two types of methods that integrate context information into E2E ASR systems. The first method is to compile the biasing phrases into a finite state transducer (FST) and incorporate it into the decoding process [10]. However, the compiling process requires calibration to find the optimal fusion weight, so the adjustment of the context list can be non-trivial.

The second method is to use encoders to directly inject context information [11, 12, 13, 14, 15]. Previous works such as CLAS [16] and CATT [17] use an additional location-aware at-

tention mechanism to leverage contextual information into E2E ASR. The introduced entity encoder enables the entity list to be personalized for individual users. However, this personalization comes at a cost: the model has less prior knowledge of the customized words, which can result in false alarms. In other words, the model may mistakenly identify non-entity names as entity terms, leading to a decrease in overall recognition performance, particularly for words that are phonemically similar. For example, if we add “José” as a context phrase, the ASR system might falsely recognize “O say can you see” as “José can you see”. This issue is particularly acute for a general ASR system that is not restricted to a particular domain. As a result, this drawback makes biased models less competitive, as the benefits gained may be outweighed by the negative impact on overall performance.

In this paper, we focus on investigating how to mitigate the performance degradation seen in common scenarios caused by contextually biased Conformer-Transducer models. Previous research [18] has indicated that filtering out irrelevant context phrases can reduce the negative impact of biasing. However, previous filtering methods are based on phonemic similarity, thus can sometimes miss acoustically similar words, and compromise the model’s ability to prevent false alarms. Furthermore, the two-stage filtering approach proposed in [18] relies on the results of the first pass to filter context phrase candidates, which can hinder the model’s ability to perform streaming recognition. Additionally, the repetition of joint operations for Transducer models in different stages can complicate the inference pipeline and increase the computation costs. To address the filtering accuracy issue, we hypothesize that to differentiate those acoustically similar words, acoustic information is not enough, and semantic information is required as well. We introduce an Entity Detector (ED) that predicts the occurrence of context words with the multi-head attention matrix of the biased encoder and predictor output. To maintain a streaming transducer pipeline and avoid the additional joint cost, we use a single-staged approach – the filtering of the context list does not require the posterior matrix from previous stages and the biased encoder and predictor output we used instead can be calculated in a streaming way.

The output of the introduced ED module is used as a switch for the list of context phrases in the inference process so that the context embedding list will only be used if context phrases appear in the speech during the inference. We also investigate two different strategies of the proposed ED module. The first one is called the Predictor based Entity Detector (P-ED), which uses only the predictor output and has the advantage of less inference complexity. The other one is called the Encoder-Predictor based Entity Detector (EP-ED), which uses the encoder bias output and the predictor bias output to perform multi-headed attention.

*Equal contribution. †Lei Xie is the corresponding author.

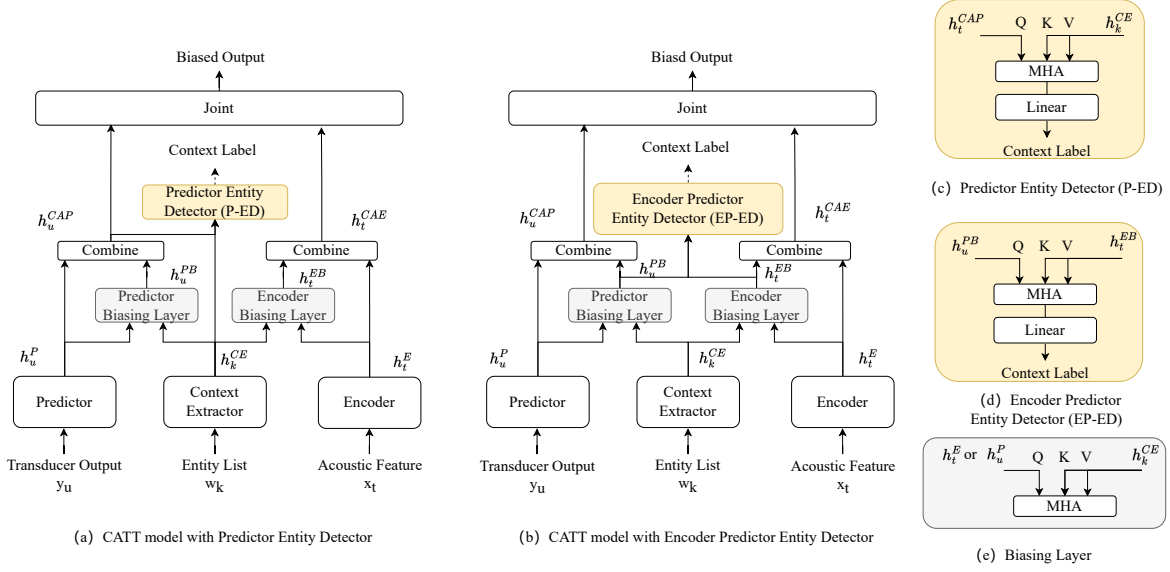


Figure 1: (a) CATT model with P-ED. (b) CATT model with EP-ED (c) P-ED module (d) EP-ED module (e) Biasing Layer

Because it uses biased context information as input, biasing operations are needed even for non-bias frames, thus leading to a slight increase in inference latency.

Experiments on Librispeech [19] and our internal voice assistant dataset show that our approach can achieve up to 6.7% and 20.7% relative reduction in word error rate (WER) compared to the baseline respectively, mitigating up to 96.7% and 84.9% of the relative WER increase for common cases. Furthermore, our approach has a minimal performance impact in personalized scenarios while maintaining a streaming inference pipeline with negligible real-time factor (RTF) increase.

2. Method

In this section, we introduce the overall framework of our base ASR system, as shown in Figure 1 (a) and (b). Generally, we add an additional Entity Detector module to the base CATT model.

2.1. Predictor Entity Detector (P-ED)

Because context words may cause false alarms, if we know that the current utterance contains no context phrases from the entity list, then it is better to inference with no entity list at all (an empty entity list). Therefore, we aim to design a module that can predict the occurrence of the context words, and use the prediction to assist biased inference.

We use a Context Encoder to convert a list of context words to context embedding. Each contextual word or phrase w_k is first represented as a sub-word for English and a character for Mandarin, and then fed into the context encoder to produce fixed dimensional vector representations h_k^{CE} . Then, the Encoder and Predictor Biasing Layers (Figure 1(e)) are introduced to allow the model to learn the relationship between contextual phrases and speech. This relationship is learned with a multi-head attention (MHA) layer. The audio embedding h_t^E serves as the query vector, while the contextual phrase embedding h_k^{CE} refers to the keys and values vectors. The Combine module consists of a LayerNorm layer, a concatenation operation, and a feed-forward projection layer. This process can be described

as:

$$h_u^{concat} = [\text{LayerNorm}(h_u^{PB}), \text{LayerNorm}(h_u^P)], \quad (1)$$

$$h_u^{CAP} = \text{FeedForward}(h_t^{concat}) \quad (2)$$

To predict the occurrence of the context words, we add a Predictor based Entity Detector (P-ED) as shown in Figure (a) and (c). Inspired by the biasing layers, P-ED is an MHA layer designed to predict the occurrence of context phrases according to Predictor output h_u^P . Specifically, the predictor output h_u^P is used as the query vector, and the contextual phrase embedding h_k^{CE} is used as the key and value. More specifically, the query, key, and value are calculated as:

$$\begin{aligned} Q^{ed} &= \sigma(XW^{ed,q} + 1(b_i^{ed,q})^\top), \\ K^{ed} &= \sigma(CW^{ed,k} + 1(b_i^{ed,k})^\top), \\ V^{ed} &= \sigma(CW^{ed,v} + 1(b_i^{ed,v})^\top), \end{aligned} \quad (3)$$

where $X = [h_1^{CAP}, \dots, h_U^{CAP}]^\top$ is the concatenated predictor bias output of an utterance, and U is the number of tokens. $C = [h_1^{CE}, \dots, h_k^{CE}]^\top$ is the context embeddings, where K is the number of context phrases. The entity detection embedding h_t^{ed} is calculated via multihead attention:

$$H^{ed} = \text{Softmax}\left(\frac{Q^{ed}(K^{ed})^\top}{\sqrt{d}}\right)V^{ed}, \quad (4)$$

where the scaling factor \sqrt{d} is for numerical stability. The entity detection embedding h_t^{ed} is then used to perform a 2-class classification task against the ground truth label of the text via a linear layer. A cross-entropy loss, L_{bias} is calculated. This process can be described as:

$$h_t^{predict} = \text{FeedForward}(h_t^{ed}). \quad (5)$$

The final joint loss function for our model is:

$$L = L_{transducer} + \lambda_1 L_{bias}, \quad (6)$$

where λ_1 is a weight hyperparameter we set as 0.4 in all our experiments.

Table 1: Performance (WER%) of different biasing list size N on the LibriSpeech benchmark.

Model	Test Set	N=0		N=20		N=50		N=100	
		test-clean	test-other	test-clean	test-other	test-clean	test-other	test-clean	test-other
C-T	Personalized/Common	4.39	9.05	4.39	9.05	4.39	9.05	4.39	9.05
CATT	Personalized	8.62	9.07	3.92	8.42	3.94	8.53	4.04	8.62
+P-ED		4.41	9.02	4.04	7.90	3.97	8.10	3.87	8.15
+EP-ED		4.41	9.03	3.91	7.89	3.94	8.06	3.90	8.13
CATT	Common	4.40	9.07	4.62	9.37	4.81	9.41	4.85	9.56
+P-ED		4.41	9.02	4.41	9.02	4.45	9.12	4.72	9.42
+EP-ED		4.41	9.03	4.41	9.03	4.43	9.09	4.71	9.41

2.2. Encoder-Predictor based Entity Detector (EP-ED)

Because the predictor embedding may lack the relevant acoustic information for the ED to deduce the presence of a context phrase, we design another framework to incorporate acoustic information as well. Different from Section 2.1, the EP-ED do not take context embedding h_t^{CE} and predictor embedding h_t^{PE} as input, instead, it uses biased predictor embedding h_t^{PB} as query and biased encoder embedding h_t^{EB} as key and value. Figure 1 (b) and (d) illustrates how we use both audio embeddings and label embeddings to detect the occurrence of context word. Specifically, use biasing layers, one taking the biased encoder embeddings h_t^{EB} as keys and values, and the biased predictor embeddings h_u^{PB} as queries to obtain the entity prediction.

Similar to that in Section 2.1, to predict the context phrase, the context-aware predictor-encoder embedding is then used for a 2-class classification task. This enables the ED to predict whether the current token is contained within a context word or not.

2.3. Adaptive Contextual Inference

For the baseline CATT method, every frame of the encoder unit need to be biased using encoder bias during inference, and with every none blank token, a predictor out h_t^{EP} is generated and used to compute predictor biased output h_t^{PB} . With the introduction of the entity detector, if it predicts that a context word is present, then the context encoder will use the context embedding of w_i , otherwise, it will use the context embedding of an empty list.

For the P-ED approach, biased encoder embedding h_t^{EB} is not required when the Entity Detector predicts that context phrases are not present in the current token. Thus only the computation is needed only for tokens where an entity is detected.

For the EP-ED approach, the inference process is a bit more complex. Because it integrates the context-aware acoustic information from the encoder, it comes Because EP-ED requires both biased encoder embedding and biased predictor embedding. For every non-blank token, a prediction is made. If no context phrase is detected, the predictor bias and encoder bias are recomputed using the context embedding of an empty context word list.

3. Experiments

3.1. Data

LibriSpeech: We used the LibriSpeech corpus [19] for some of our experiments. It contains 960 hours of English audiobook recordings. The training was performed using the full training set of 960h, with SpecAugment [20] applied. We evaluate the models on the dev and test sets provided with the LibriSpeech corpus: *dev-clean*, *dev-other*, *test-clean* and *test-other*. For

Librispeech, *test-clean* and *test-other* are used to create both *common* and *personalized* test sets. The only difference between *common* and *personalized* is the entity list. For the *personalized* set, entity lists for the LibriSpeech corpus from [14] were used. The entity lists are created by selecting words from each utterance that appear in a rare-word list. and then adding a certain number of distractor words. For *common* set, the entity list contains only distractor words.

Voice assistant dataset: To validate the model’s performance, a larger in-house dataset containing 15,000 hours anonymized Mandarin speech collected from a speech assistant product was used. The *personalized* and *common* sets are created individually. The *personalized* test set contains approximately 500 utterances with 50 contextual words. Most of the contextual words are personal names. For the *common* set, there are about 10000 utterances, the same entity list for *common* set was used, and the utterances contain no context phrases from the provided entity list.

3.2. Experimental Setups

Our proposed method and the baseline Conformer Transducer model (referred to as C-T) shares the following parameters: The encoder is a 12-layer conformer. Each Conformer block consists of 4-head 256-dim multi-head attention block. The predictor is a 2-layer LSTM. Both the encoder and predictors have an embedding size of 256. The joint network is a fully-connected feed-forward component with one hidden layer followed by a Tanh activation function.

For models other than the the C-T baseline, we used the LSTM context encoder as mentioned in Section 2.1. It consists of a 1-layer BLSTM with a dimension of 256. The bias layer is a transformer block with an embedding size of 256 and 4 attention heads. The linear layer in the entity detector keeps a 256-dimensional input and output.

In addition to the word error rate (WER) for the English corpus and character error rate(CER) for the mandarin corpus, results were further evaluated via the biased label character error rate (L-CER) which is defined by:

$$L-CER = \frac{\text{EditDistance}(ref, hyp)}{\max(\text{len}(ref), \text{len}(hyp))} \quad (7)$$

where *ref* is the ground truth of the context phrase label, *hyp* is the prediction of ED. L-CER is a direct assessment of the ED prediction, compared with WER metrics, we use it as a more direct indicator of the performance of the ED module.

3.3. Contextual ASR Accuracy

In this section, we tested both entity detection methods and non-contextual baselines for both *personalized* and *common* sets. As shown in Table 1, Experiments for the *common* set show a sig-

Table 2: Ablation study on Librispeech benchmark

Model	Personalized				Common			
	N=20		N=50		N=20		N=50	
	test-clean	test-other	test-clean	test-other	test-clean	test-other	test-clean	test-other
CATT	3.92	8.42	3.94	8.53	4.65	9.37	4.81	9.41
+ P-ED	4.04	7.90	3.97	8.10	4.41	9.03	4.45	9.12
- detector	3.77	7.83	3.87	7.99	4.62	9.32	4.73	9.38
- 50% detector	4.12	8.31	4.19	8.40	4.53	9.15	4.57	9.19
+ EP-ED	3.91	7.89	3.93	8.06	4.41	9.04	4.43	9.09
- detector	3.79	7.87	3.89	7.96	4.59	9.30	4.69	9.34
- 50% detector	4.13	8.29	4.21	8.39	4.52	9.14	4.57	9.17

Table 3: Performance(L-CER%) on Librispeech, TO stands for test other, TC stands for test clean, Pers. stands for Personalized, Comm. stands for Common.

Model	Test Set	N=20		N=50		N=100	
		TC	TO	TC	TO	TC	TO
CATT + P-ED	Pers.	10.98	9.72	19.34	18.36	43.24	35.42
		7.04	8.53	15.84	13.57	39.84	37.36
CATT + P-ED	Comm.	4.97	3.40	17.82	13.86	91.52	85.87
		4.38	2.02	13.62	10.72	90.93	83.25

nificant decrease in WER for the common dataset, with up to 6.7 % relative WERR for the test other with an entity list of 20. Significantly reducing the WER increase incurred by bias word by 96.7%, as the entity detector has a low false alarm rate, and can predict the occurrence of context phrases with an L-CER under 5% as shown in Table 3.

The result of the *personalized* set showed a WER reduction for all entity list sizes compared with CATT baseline. The difference is especially significant in the test-other set. This is unlikely caused by the additional parameters, since the increase in model size is negligible. We hypothesize that this is caused by the introduced loss function helping to supervise the gradient descent process and thus leading to better performance. This hypothesis is further validated in the ablation study.

However, when the size of the biasing list grows over 100, both P-ED and EP-ED showed a significant increase in L-CER, because when the biasing list grows too large, the ED tends to predict always on, thus reducing the performance for common scenarios. This suggests that the model is sensitive to the size of the biasing list, which is related to the number of context words that are used during training.

3.4. Ablation Study

As mentioned in Section 2.1, ED is introduced to predict the occurrence of context phrases. To verify it, we removed features from removing features from P-ED and EP-ED (- detection means always attending to context phrase regardless of TD prediction, -50% detection means attending to context phrase with 50% chance, regardless of TD prediction). Table 2 presents the result. Both P-ED and EP-ED show improvement over the 50% bias, with less WER in both personalized and common scenarios. This implies that through training both entity detectors have learned the ability to predict context words. When compared to a always-on bias (-detector), it shows that even without using the entity detector during inference, both entity detector methods show a significant improvement over the base CATT method, probably due to the introduced supervision helping the model to be better optimized during training via the introduced supervision. Both P-ED and EP-ED, improve *common* set WERs with EP-ED have better performance over P-ED overall, compared with always-on bias.

Table 4: Performance of Entity Detector on in-house dataset

Model	CER(Personalized)		CER(Common)	RTF
	N=0	N=50	N=50	
C-T	15.75	15.75	6.64	0.108
CATT	15.89	6.93	8.82	0.128
+P-ED	16.03	7.33	7.19	0.130
+EP-ED	14.43	7.02	6.97	0.135

3.5. RTF Analysis

To evaluate the performance of our proposed filter module, we compared the runtime RTF of the ASR system with and without the biasing module. We conducted the experiment on a 3.40GHz AMD Ryzen 5950X CPU using a single thread. As shown in Table 4, the results indicate that the addition of the proposed module leads to a reduction in RTF when compared to the baseline system. This is mainly due to the fact that the encoder bias h_t^{EB} is unnecessary when the proposed module predicts the absence of a contextual phrase. Additionally, the computational cost of the prediction process is negligible. However, when comparing the RTF of the EP-ED system with the baseline, we observed a slight increase in RTF. This is because the EP-ED system requires the computation of every context-aware acoustic embedding h_t^{EB} and h_t^{PB} , as in CATT, to make a prediction.

3.6. Test on Voice Assistant Dataset

We further validate our result on an in-house voice assistance dataset. For the common set, When utilizing a biased list, the P-ED achieved a 16.4% CERR and the EP-ED achieved a 20.7% CERR compared to the baseline CATT. Mitigating the CER increase introduced by biasing modules by 84.9% and 74.7% respectively. When not using a bias list, the performance has no significant difference compared to the baseline. Both P-ED and EP-ED did not incur significant performance degradation, having 5.4% and 1.2% relative CER increase respectively.

4. Conclusions

In this work, we studied how to reduce the WER increase introduced by biasing modules in common scenarios. We proposed a novel adaptive method based on entity detector. The entity detector predicts context phrases contained in the utterances and calculates cross-entropy loss on the prediction results to train the ED to filter out the irrelevant entity list during inference. Compared with previous methods, our approach has better performance overall, achieving significant word error rate reduction for common scenarios.

5. References

- [1] A. Graves, S. Fernández, F. J. Gomez, and J. Schmidhuber, “Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks,” in *Proc. ICML*. ACM, 2006, pp. 369–376.
- [2] A. Graves and N. Jaitly, “Towards End-To-End Speech Recognition with Recurrent Neural Networks,” in *Proc. ICML*. JMLR.org, 2014, pp. 1764–1772.
- [3] K. Li, J. Li, G. Ye, R. Zhao, and Y. Gong, “Towards Code-Switching ASR for End-to-End CTC Models,” in *Proc. ICASSP*. IEEE, 2019, pp. 6076–6080.
- [4] Y. He, T. N. Sainath, R. Prabhavalkar, I. McGraw, R. Alvarez, D. Zhao, D. Rybach, A. Kannan, Y. Wu, R. Pang, Q. Liang, D. Bhatia, Y. Shangguan, B. Li, G. Pundak, K. C. Sim, T. Bagby, S. Chang, K. Rao, and A. Gruenstein, “Streaming End-to-end Speech Recognition for Mobile Devices,” in *Proc. ICASSP*. IEEE, 2019, pp. 6381–6385.
- [5] Y. Zhang, S. Sun, and L. Ma, “Tiny Transducer: A Highly-Efficient Speech Recognition Model on Edge Devices,” in *Proc. ICASSP*. IEEE, 2021, pp. 6024–6028.
- [6] X. Wang, S. Sun, L. Xie, and L. Ma, “Efficient Conformer with Prob-Sparse Attention Mechanism for End-to-End Speech Recognition,” in *Proc. INTERSPEECH*. ISCA, 2021, pp. 4578–4582.
- [7] J. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, “Attention-Based Models for Speech Recognition,” in *Proc. NeurIPS*, 2015, pp. 577–585.
- [8] K. Cho, B. van Merriënboer, Ç. Gülçehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, “Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation,” in *Proc. EMNLP*, A. Moschitti, B. Pang, and W. Daelemans, Eds. ACL, 2014, pp. 1724–1734.
- [9] W. Chan, N. Jaitly, Q. V. Le, and O. Vinyals, “Listen, Attend and Spell: A Neural Network for Large Vocabulary Conversational Speech Recognition,” in *Proc. ICASSP*. IEEE, 2016, pp. 4960–4964.
- [10] M. Mohri, F. Pereira, and M. Riley, “Weighted finite-state transducers in speech recognition,” *Computer Speech & Language*, vol. 16, no. 1, pp. 69–88, 2002.
- [11] Z. Chen, M. Jain, Y. Wang, M. L. Seltzer, and C. Fuegen, “Joint Grapheme and Phoneme Embeddings for Contextual End-to-End ASR,” in *Proc. INTERSPEECH*, G. Kubin and Z. Kacic, Eds. ISCA, 2019, pp. 3490–3494.
- [12] M. Han, L. Dong, S. Zhou, and B. Xu, “CIF-Based Collaborative Decoding for End-to-End Contextual Speech Recognition,” in *Proc. ICASSP*. IEEE, 2021, pp. 6528–6532.
- [13] G. Sun, C. Zhang, and P. C. Woodland, “Tree-Constrained Pointer Generator for End-to-End Contextual Speech Recognition,” in *Proc. ASRU*. IEEE, 2021, pp. 780–787.
- [14] D. Le, M. Jain, G. Keren, S. Kim, Y. Shi, J. Mahadeokar, J. Chan, Y. Shangguan, C. Fuegen, O. Kalinli, Y. Saraf, and M. L. Seltzer, “Contextualized Streaming End-to-End Speech Recognition with Trie-Based Deep Biasing and Shallow Fusion,” in *Proc. INTERSPEECH*. ISCA, 2021, pp. 1772–1776.
- [15] S. Dingliwal, M. Sunkara, S. Ronanki, J. Farris, K. Kirchhoff, and S. Bodapati, “Personalization of CTC Speech Recognition Models,” in *Proc. SLT*. IEEE, 2022, pp. 302–309.
- [16] G. Pundak, T. N. Sainath, R. Prabhavalkar, A. Kannan, and D. Zhao, “Deep Context: End-to-end Contextual Speech Recognition,” in *Proc. SLT*. IEEE, 2018, pp. 418–425.
- [17] F. Chang, J. Liu, M. Radfar, A. Mouchtaris, M. Omologo, A. Rastrow, and S. Kunzmann, “Context-aware transformer transducer for speech recognition,” in *IEEE Automatic Speech Recognition and Understanding Workshop, ASRU 2021*. IEEE, 2021, pp. 503–510.
- [18] Z. Yang, S. Sun, X. Wang, Y. Zhang, L. Ma, and L. Xie, “Two Stage Contextual Word Filtering for Context bias in Unified Streaming and Non-streaming Transducer,” *CoRR*, vol. abs/2301.06735, 2023.
- [19] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An ASR Corpus Based on Public Domain Audio Books,” in *Proc. ICASSP*. IEEE, 2015, pp. 5206–5210.
- [20] D. S. Park, W. Chan, Y. Zhang, C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition,” in *Proc. INTERSPEECH*. ISCA, 2019, pp. 2613–2617.
- [21] P. Aleksic, M. Ghodsi, A. Michaely, C. Allauzen, K. Hall, B. Roark, D. Rybach, and P. Moreno, “Bringing contextual information to google speech recognition,” in *Proc. INTERSPEECH*, 2015, pp. 468–472.
- [22] I. Williams, A. Kannan, P. S. Aleksic, D. Rybach, and T. N. Sainath, “Contextual Speech Recognition in End-to-end Neural Network Systems Using Beam Search,” in *Proc. INTERSPEECH*. ISCA, 2018, pp. 2227–2231.
- [23] Z. Chen, M. Jain, Y. Wang, M. L. Seltzer, and C. Fuegen, “End-to-end Contextual Speech Recognition Using Class Language Models and a Token Passing Decoder,” in *Proc. ICASSP*. IEEE, 2019, pp. 6186–6190.
- [24] D. Zhao, T. N. Sainath, D. Rybach, P. Rondon, D. Bhatia, B. Li, and R. Pang, “Shallow-Fusion End-to-End Contextual Biasing,” in *Proc. INTERSPEECH*. ISCA, 2019, pp. 1418–1422.
- [25] R. Huang, O. Abdel-Hamid, X. Li, and G. Evermann, “Class LM and Word Mapping for Contextual Biasing in End-to-End ASR,” in *Proc. INTERSPEECH*. ISCA, 2020, pp. 4348–4351.
- [26] S. Kim, Y. Shangguan, J. Mahadeokar, A. Bruguier, C. Fuegen, M. L. Seltzer, and D. Le, “Improved Neural Language Model Fusion for Streaming Recurrent Neural Network Transducer,” in *Proc. ICASSP*. IEEE, 2021, pp. 7333–7337.
- [27] G. Chen, S. Chai, G. Wang, J. Du, W. Zhang, C. Weng, D. Su, D. Povey, J. Trmal, J. Zhang, M. Jin, S. Khudanpur, S. Watanabe, S. Zhao, W. Zou, X. Li, X. Yao, Y. Wang, Z. You, and Z. Yan, “GigaSpeech: An Evolving, Multi-Domain ASR Corpus with 10,000 Hours of Transcribed Audio,” in *Proc. INTERSPEECH*, H. Hermansky, H. Cernocký, L. Burget, L. Lamel, O. Scharenborg, and P. Motlíček, Eds. ISCA, 2021, pp. 3670–3674.
- [28] J. D. Fox and N. Delworth, “Improving Contextual Recognition of Rare Words with an Alternate Spelling Prediction Model,” in *Proc. INTERSPEECH*. ISCA, 2022, pp. 3914–3918.
- [29] B. Zhang, D. Wu, Z. Peng, X. Song, Z. Yao, H. Lv, L. Xie, C. Yang, F. Pan, and J. Niu, “WeNet 2.0: More Productive End-to-End Speech Recognition Toolkit,” in *Proc. INTERSPEECH*. ISCA, 2022, pp. 1661–1665.