



# Contextualized End-to-End Speech Recognition with Contextual Phrase Prediction Network

Kaixun Huang, Ao Zhang, Zhanheng Yang, Pengcheng Guo, Bingshen Mu, Tianyi Xu, Lei Xie\*

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

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

## Abstract

Contextual information plays a crucial role in speech recognition technologies and incorporating it into the end-to-end speech recognition models has drawn immense interest recently. However, previous deep bias methods lacked explicit supervision for bias tasks. In this study, we introduce a contextual phrase prediction network for an attention-based deep bias method. This network predicts context phrases in utterances using contextual embeddings and calculates bias loss to assist in the training of the contextualized model. Our method achieved a significant word error rate (WER) reduction across various end-to-end speech recognition models. Experiments on the LibriSpeech corpus show that our proposed model obtains a 12.1% relative WER improvement over the baseline model, and the WER of the context phrases decreases relatively by 40.5%. Moreover, by applying a context phrase filtering strategy, we also effectively eliminate the WER degradation when using a larger biasing list.

**Index Terms:** End-to-end Speech Recognition, Deep Biasing, Contextual List Filter

## 1. Introduction

In recent years, end-to-end automatic speech recognition (ASR) systems have made remarkable progress, driven by the rapid development of deep neural networks, and have become the mainstream of speech recognition research. Various end-to-end ASR methods, such as connectionist temporal classification (CTC) [1, 2], recurrent neural network transducer (RNN-T) [3, 4], and attention-based encoder-decoder (AED) [5, 6, 7], have been widely used in scenarios such as voice assistants and online conferences. However, since deep learning techniques are highly data-dependent, the recognition accuracy of end-to-end ASR systems tends to be significantly reduced for rare phrases in the data, such as entity names, technical terms, etc. In practical applications, these rare phrases can often be obtained in advance as contextual information. For example, for voice assistants, *contact names in the address book* and *the names of installed applications* are likely to appear during the recognition process. Therefore, informing the model of contextual information is crucial to accurately predict the corresponding phrases. As a result, research on infusing contextual knowledge into end-to-end ASR models for biased decoding has become increasingly important.

In previous studies, a typical bias method is LM fusion [8, 9, 10, 11, 12], with shallow fusion being the most commonly used approach. In [9], an n-gram of contextual phrases is compiled into a weighted finite-state transducer (WFST) and the

state of the decoded path in the WFST is maintained during decoding. When a specific word is decoded, the score is added to the path, thereby biasing the recognition process toward the contextual phrase. While LM fusion can intuitively increase the posterior of tokens related to contextual phrases, it has limited performance improvement on the bias list and requires testing to determine the optimal fusion weight.

Another approach to address this problem is the deep bias method based on neural networks [13, 14, 15, 16, 17], which introduces a separate biasing component in the end-to-end model for modeling contextual phrases. Compared with LM fusion, the deep bias method is more flexible and allows quick adaptation to different contextual phrase lists for various scenarios. Several previous works have modified basic end-to-end speech recognition models to obtain contextualized models, such as CLAS [18] and CATT [19]. However, these methods lack explicit supervision for bias tasks and are often limited to specific model structures.

To solve this problem, we propose a novel approach called contextual phrase prediction network (CPP Network) for deep biasing, which combines a bidirectional long-short term memory network (BLSTM) based contextual encoder with an attention mechanism to extract embeddings of context phrases. Unlike CLAS [18] and CATT [19], our model predicts the biased phrases that appear in the utterance using the contextual phrase prediction network and calculates a CTC loss between the predicted phrases with labels that only contain context phrases as the bias loss. This loss function provides explicit supervision for the bias task, allowing the contextualized model to learn to focus on context phrases appearing in the audio while minimizing the impact on the original recognition results when there are no context phrases. The contextual phrase prediction network can be easily integrated into mainstream end-to-end speech recognition models, such as CTC, AED, and Transducer, to achieve contextualized recognition. Our proposed techniques achieve an average relative WER improvement of 12.1% on LibriSpeech, with a relative improvement of 40.5% in the WER of phrases included in the biasing list.

In addition, using a biasing list containing thousands of contextual phrases as input for a contextualized ASR model poses challenges in retrieving meaningful contextual encodings from the biasing layer. Moreover, the biasing layer requires computations on thousands of contextual phrase encodings, leading to a noticeable increase in inference time. Inspired by [20], we employ a two-stage contextual phrase filtering method. Firstly, the model performs inference using an empty biasing list to calculate posterior score confidence and sequence order confidence for each context phrase in the actual biasing list. Then, we use the filtered biasing list to generate the final bias output by combining it with the audio embed-

\* Corresponding author.

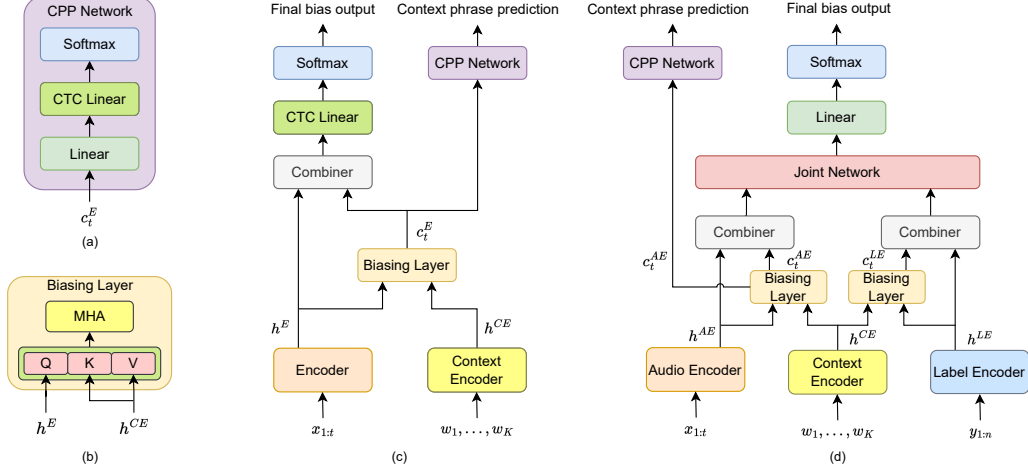


Figure 1: (a) The proposed contextual phrase prediction network, where the CTC linear shares parameters with contextualized models. (b) Biasing layer via multi-head attention (MHA) (c) Contextualized CTC model (d) Contextualized Transducer model

dings retained from the first stage through a small amount of computation. This approach reduces the computational burden and inference time, while ensuring that only relevant contextual phrases are used for biasing. Compared to [20], we extended this contextual phrase filtering method to be applicable to multiple structures, rather than being limited to Transducer models with cascaded encoders.

## 2. Contextual Bias with CPP Network

In this section, we introduce the modification on contextual bias module with our proposed CPP Network. Different modifications made on different ASR framework including CTC, AED and Transducer will be introduced respectively.

### 2.1. Contextualized CTC/AED

For both contextualized CTC and AED frameworks, we made modifications only to the encoder. It should be noted that we also deploy CTC as auxiliary training criterion for AED encoder. Thus these modifications for both frameworks were similar. Generally, we add a context encoder, a multi-head attention-based biasing layer, and a contextual phrase prediction network (CPP Network) for the framework to integrate contextual information, as shown in Figure 1(c).

The context encoder is used to encode the custom contextual phrase list into embeddings with the same length. Each phrase in the list is tokenized with the same tokenizer as the base CTC model, and the token sequence is used as input for the BLSTM. The hidden state and cell state of the last time step from both the forward and backward LSTM are concatenated and then passed through a linear layer to obtain the embedding  $h_i^{CE}$  for the contextual phrase. In addition, we added a phrase to the contextual phrase list that contains only a single blank token, referred to as *no-bias*, to enable the model to focus on it when there is no contextual phrase in the audio and ignore other contextual phrases.

The biasing layer (Figure 1(b)) is a multi-head attention (MHA) layer designed to help the model learn the relationship between contextual phrases and speech. While computing the MHA score, the audio embedding  $h_t^E$  serves as the query, and the contextual phrase embeddings  $h_i^{CE}$  serve as the keys and values to obtain the context representations  $c_t^E$  as:

$$\alpha_{it} = \text{Softmax} \left( \left( \mathbf{W}_q \mathbf{h}_t^E \right) \left( \mathbf{W}_k \mathbf{h}_i^{CE} \right) / \text{sqrt}(d) \right), \quad (1)$$

$$\mathbf{c}_t^E = \sum_{i=0}^K \alpha_{it} \mathbf{h}_i^{CE}, \quad (2)$$

where the scaling factor  $\text{sqrt}(d)$  is for numerical stability, and  $K$  represents the number of contextual phrases. The context embeddings and audio embeddings are then combined to obtain context-aware audio embeddings  $h_t^{CA}$ . The combiner consists of a LayerNorm layer, a concatenation operation, and a feed-forward projection layer. This process can be described as:

$$\mathbf{h}_t^{\text{concat}} = [\text{LayerNorm}(\mathbf{h}_t^E), \text{LayerNorm}(\mathbf{c}_t^E)], \quad (3)$$

$$\mathbf{h}_t^{CA} = \text{FeedForward}(\mathbf{h}_t^{\text{concat}}). \quad (4)$$

The CPP network consists of two feed-forward projection layers and a Softmax function (Figure 1(a)), where the second feed-forward projection layer shares parameters with the linear layer in the CTC model that maps audio embeddings to a posterior probability distribution over vocabularies, denoted as CTC linear in Figure 1. Receiving the context embedding  $c_t^E$  as input, CPP Network predicts the contextual phrases that appear in the dialogue utterance. We use the posterior generated by CPP Network as reference for CTC loss calculation. Such loss can supervise the biasing layer to extract relevant contextual phrases information from the audio embeddings  $h^E$  when the contextual phrases appear and minimizes the influence of introducing additional information when no contextual phrases appear. Note that this part of the model is only used for auxiliary training and does not participate in model inference.

The final joint loss function for the CTC model is:

$$L = L_{ctc} + \lambda_2 L_{CPP}, \quad (5)$$

where  $\lambda_2$  is a weight hyperparameter we set as 0.1 in experiment.

For the AED model, with  $\lambda_1$  as the CTC weight, the final joint loss function is:

$$L = \lambda_1 L_{ctc} + (1 - \lambda_1) L_{att} + \lambda_1 \lambda_2 L_{CPP}. \quad (6)$$

### 2.2. Contextualized Transducer

For the Transducer, we augment the architecture with a contextual phrase prediction network to implement the contextualized Transducer, as illustrated in Figure 1(d). Specifically, compared to the standard Transducer, we add a context encoder, two

Table 1: LibriSpeech results with different biasing list size  $N$ . Reported metrics are in the following format: WER (U-WER/B-WER).

Model	N=0		N=100		N=500		N=1000	
	test-clean	test-other	test-clean	test-other	test-clean	test-other	test-clean	test-other
CTC Baseline	5.59 (3.5 / 23.4)	12.17 (8.1 / 48.3)	5.59 (3.5 / 23.4)	12.17 (8.1 / 48.3)	5.59 (3.5 / 23.4)	12.17 (8.1 / 48.3)	5.59 (3.5 / 23.4)	12.17 (8.1 / 48.3)
CTC + bias	5.67 (3.6 / 23.4)	12.14 (8.2 / 47.3)	<b>5.13</b> <b>(4.2 / 12.7)</b>	<b>11.76</b> <b>(9.9 / 29.0)</b>	6.11 (5.2 / 13.7)	13.77 (11.8 / 31.5)	6.89 (6.0 / 14.7)	14.85 (12.9 / 32.6)
AED Baseline	4.22 (2.6 / 18.1)	8.88 (5.6 / 37.6)	4.22 (2.6 / 18.1)	8.88 (5.6 / 37.6)	4.22 (2.6 / 18.1)	8.88 (5.6 / 37.6)	4.22 (2.6 / 18.1)	8.88 (5.6 / 37.6)
AED + bias	4.29 (2.6 / 18.3)	9.16 (5.9 / 37.5)	<b>3.40</b> <b>(2.6 / 10.4)</b>	<b>7.77</b> <b>(6.0 / 23.0)</b>	3.68 (2.8 / 10.9)	8.31 (6.5 / 24.3)	3.81 (2.9 / 11.4)	8.75 (6.9 / 25.3)
Transducer Baseline	4.30 (2.8 / 17.2)	8.90 (5.9 / 35.5)	4.30 (2.8 / 17.2)	8.90 (5.9 / 35.5)	4.30 (2.8 / 17.2)	8.90 (5.9 / 35.5)	4.30 (2.8 / 17.2)	8.90 (5.9 / 35.5)
Transducer + bias	4.38 (2.7 / 18.3)	9.12 (5.9 / 37.6)	<b>3.66</b> <b>(2.8 / 11.2)</b>	<b>7.63</b> <b>(6.0 / 22.1)</b>	3.78 (2.9 / 11.5)	7.99 (6.2 / 23.4)	3.88 (2.9 / 11.9)	8.28 (6.4 / 24.5)

MHA biasing layers, and a contextual phrase prediction network. Similar to CTC/AED framework, both the audio encoder and the label encoder are followed by MHA biasing layers to integrate contextual information from  $h_i^{CE}$ . Then the context-aware audio embeddings  $c^{AE}$  and the context-aware label embeddings  $c^{LE}$  are fed into the joint network.

For CPP Network, We use the context embeddings  $c_t^{AE}$  obtained from the audio encoder as input to predict the contextual phrases contained in the speech and compute the bias loss.

Finally, we combine the bias loss with the standard Transducer loss as the optimization objective:

$$L = \lambda_1 L_{ctc} + (1 - \lambda_1) L_{transducer} + \lambda_1 \lambda_2 L_{CPP}, \quad (7)$$

where  $\lambda_1$  and  $\lambda_2$  are the weight hyperparameters for the CTC loss and the CPP loss.

### 3. Contextual Phrase Filtering

Inspired by [20], we apply a two-stage filtering algorithm, using the posterior of the CTC to filter the contextual phrases. Differently, we simplified the structure of the Cascade Encoders, performing contextual phrase filtering without altering the original model structure.

Firstly, the calculation is performed using a biasing list that contains only *no-bias* to obtain the CTC posterior without bias. The audio embedding  $h^E$  is retained. The decoding results from the CTC during the first pass inference can serve as streaming results.

The contextual phrases are then filtered in two stages using the CTC posterior. The first stage calculates the posterior and the phrase score confidence (PSC), ignoring the order of the contextual phrases. It computes the sum of the maximum posterior values of each token contained in the contextual phrases within a sliding window and normalizes it with the sequence length. This stage eliminates most of the contextual phrases with low correlation with the audio, retaining only the remaining ones for the second stage of filtering. In the second stage, the sequence order confidence (SOC) is calculated by using a dynamic programming algorithm to compute the maximum posterior sum of tokens contained in contextual phrases within the sliding window, considering token order.

After obtaining the filtered contextual phrases, subsequent inference is performed with the previously retained audio embedding  $c_t^E$ , resulting in a more accurate biased recognition result. The complexity of context phrase filtering is low for inference. Compared to the original recognition process, only the biasing layer, combiner, and CTC linear are repeated in the calculation, and the increase in computation is relatively small.

## 4. Experiments

### 4.1. Data

A majority of our experiments were conducted on the LibriSpeech corpus [21], which comprises 960 hours of English audiobook readings. The model was trained on the complete 960 hours of training data to demonstrate its performance on relatively low-resource tasks. The *dev-clean* and *dev-other* sets were used for validation, while the *test-clean* and *test-other* sets were used for evaluation. SpecAugment [22] was applied for data augmentation.

During the training of contextualized models, we extract three random phrases of word length between 1 to 3 from each utterance and add them to the biasing list in the batch. Additionally, we incorporate a certain number of distractors into the biasing list to maintain the list size at 60.

During testing, we utilized the biasing lists for the LibriSpeech corpus provided in [16] as the contextual phrase lists. As outlined in [16], the biasing lists were first formed for each utterance by identifying words belonging to the full rare-word list from the reference of each utterance, and adding a certain number of distractor words.

To validate the model’s performance with larger amounts of data, we also evaluate our method on GigaSpeech corpus [23] and tested on *Earnings21* dataset [24]. *Earnings21* set consists of earnings call transcripts from various publicly traded companies and provides a range of proprietary terms, including names of companies, products, and executives.

### 4.2. Experimental Setups

The baseline CTC model consists of a 12-layer conformer with 256-dimensional input and 4 self-attention heads and the baseline AED model uses 12 conformer layers for the encoder and 6 transformer layers for the decoder. The baseline Transducer model uses a 12-layer conformer as the audio encoder and 2-layer LSTM as the label encoder, referred to as C-T.

In the contextualized version, the context encoder consists of a 1-layer BLSTM with a dimensionality of 256 and a linear layer, while the biasing layer is an MHA layer with an embedding size of 256 and 4 attention heads. The first linear layer in the context prediction network maintains a 256-dimensional input and output, followed by a tanh activation function, while the second linear layer projects the input onto the vocabulary size and shares parameters with the CTC linear.

In addition to the word error rate (WER), we also evaluated the results using the biased word error rate (B-WER) and unbiased word error rate (U-WER). U-WER is measured on words that are NOT in the biasing list, while B-WER is measured on

Table 2: Results of contextual phrase filtering on contextual AED model with biasing list size  $N = 1000$

Model	without filter		with filter	
	test-clean	test-other	test-clean	test-other
CTC + bias	6.89 (6.0 / 14.7)	14.85 (12.9 / 32.6)	4.95 (3.72 / 15.3)	10.81 (8.3 / 32.7)
AED + bias	3.81 (2.9 / 11.4)	8.75 (6.9 / 25.3)	3.61 (2.6 / 12.1)	7.90 (5.9 / 25.8)
C-T + bias	3.88 (2.9 / 11.9)	8.28 (6.4 / 24.5)	3.76 (2.7 / 12.8)	7.96 (5.8 / 26.8)

Table 3: Comparison of proposed method and shallow fusion

Model	test-clean	test-other
AED baseline	4.29 (2.6 / 18.3)	9.16 (5.9 / 37.5)
+ SF	4.19 (2.7 / 16.8)	8.79 (6.0 / 33.1)
+ bias	3.88 (2.8 / 13.0)	8.47 (6.4 / 26.6)
++ SF	3.86 (2.9 / 12.4)	8.45 (6.5 / 25.5)

words that are IN the biasing list. For insertion errors, if the inserted phrase is in the biasing list, it is counted towards B-WER, otherwise it is counted towards U-WER.

### 4.3. Contextual ASR Accuracy

In this section, we tested the contextual recognition models and their non-contextual baseline on all these types of models. As shown in Table 1, the results of experiments containing a non-empty biasing list demonstrate a decrease in B-WER. The best results were achieved when using a biasing list of 100, yielding an average 40.5% relative improvement of B-WER compared to the baseline. When the size of the biasing list increased to 1000, B-WER still showed a noticeable decrease, but the reduction was less significant, and U-WER was also impacted to some degree. This suggests that the model is sensitive to the size of the biasing list. We found that using more distractors during the training process, or not sampling contextual phrases for some utterances can reduce U-WER when using a large biasing list, but will also decrease the improvement of B-WER.

### 4.4. Contextual Phrase Filtering Performance

As shown in Table 2, we tested the effect of contextual phrase filtering on all these types of models with a biasing list of size 1000. After applying the filtering method, the impact of the biasing list on U-WER was significantly reduced and largely regressed to the level when an empty biasing list was used, while B-WER showed only a slight increase compared to the results without contextual phrase filtering.

### 4.5. Comparison with Shallow Fusion

In order to compare the proposed deep bias method and shallow fusion, we separately tested the effect of using each method alone and using them together on the contextualized AED model. To ensure fairness, we extracted the contextual phrases in each utterance in the test set into a list, with a list size of 4250 for the *test-clean* set and 3838 for the *test-other* set. Two methods use the same biasing list. When evaluating the results of the shallow fusion, we used the shallow fusion method based on OTF rescoring described in [25]. As shown in Table 3, the shallow fusion method reduced the average U-WER by 6.9% relative to the baseline and the proposed deep bias method reduced the average U-WER by 29.2%. Compared to shallow fusion, the improvement in B-WER from deep bias was appar-

Table 4: Results of ablation study

Model	N=100		N=1000	
	test-clean	test-other	test-clean	test-other
AED + bias	<b>3.40</b> ( <b>2.6 / 10.4</b> )	7.77 (6.0 / 23.0)	<b>3.81</b> ( <b>2.9 / 11.4</b> )	<b>8.75</b> ( <b>6.9 / 25.3</b> )
- share	3.50 (2.7 / 10.2)	<b>7.68</b> ( <b>6.2 / 21.0</b> )	4.13 (3.3 / 11.3)	9.43 (7.7 / 24.4)
- loss	4.26 (2.6 / 18.0)	9.31 (5.9 / 39.0)	4.26 (2.6 / 18.0)	9.31 (5.9 / 39.0)
C-T + bias	<b>3.66</b> ( <b>2.8 / 11.2</b> )	<b>7.63</b> ( <b>6.0 / 22.1</b> )	<b>3.88</b> ( <b>2.9 / 11.9</b> )	<b>8.28</b> ( <b>6.4 / 24.5</b> )
- loss	4.05 (2.8 / 14.5)	8.65 (6.2 / 30.5)	4.24 (2.9 / 15.4)	8.99 (6.3 / 32.8)

Table 5: Results of contextualized AED model on Earnings21

Model	WER	Recall	Precision	F1
AED Baseline	14.96	58.66	91.52	71.50
AED + bias	15.12	79.11	80.67	79.88

ently greater, and the use of both methods together can further reduce the B-WER, although the impact on U-WER will also accumulate.

### 4.6. Ablation Study

Table 4 presents the results of ablation study. We tested the performance of the contextualized AED model without sharing the CTC linear parameters in the CPP Network and without the addition of the bias loss. The results of not sharing the CTC linear parameters showed better B-WER performance, but the U-WER was more sensitive to the size of the biasing list. Without the bias loss, the contextualized AED model completely lost its ability to perform bias recognition. When a larger biasing list was used, there was a noticeable increase in the WER loss. In the contextualized Transducer, when the bias loss is removed, the resulting model is essentially CATT, which had the ability to perform bias recognition, but the improvement in the B-WER was not as significant as before.

### 4.7. Test on Earnings21

We trained a contextualized AED model on the GigaSpeech dataset [23] and evaluated it on the Earnings21 dataset [24], with results presented in Table 5. Similar to [24], we demonstrate the WER and phrase-level recall, precision, and F1 scores of two models on the Earnings21 dataset. When using the contextualized AED model and applying contextual phrase filtering, the recall of contextual phrases is significantly improved, resulting in a relative improvement of 33.2% compared to the baseline. Although the precision was impacted to a certain extent, the overall F1 score still improves by 11.7% relative. Due to the small proportion of biased phrases in the test set, contextual biasing has a minimal impact on overall WER.

## 5. Conclusions

In this paper, we propose a deep biasing method based on the contextual phrase prediction network. The context phrase prediction network predicts context phrases contained in the utterances and calculates CTC loss on the prediction results as an auxiliary loss for the training of the contextualized model. Compared with previous works, our approach achieves significant B-WER improvement on all CTC, AED, and Transducer models. In addition, we apply a context phrase filtering method to the final context-aware ASR model, which effectively reduces the U-WER loss when using a larger biasing list.

## 6. References

- [1] A. Graves, S. Fernández, F. J. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006)*, W. W. Cohen and A. W. Moore, Eds. ACM, 2006, pp. 369–376.
- [2] A. Graves and N. Jaitly, "Towards end-to-end speech recognition with recurrent neural networks," in *Proceedings of the 31th International Conference on Machine Learning, ICML 2014*. JMLR.org, 2014, pp. 1764–1772.
- [3] 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 *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2019*. IEEE, 2019, pp. 6381–6385.
- [4] Y. Zhang, S. Sun, and L. Ma, "Tiny transducer: A highly-efficient speech recognition model on edge devices," in *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2021*. IEEE, 2021, pp. 6024–6028.
- [5] J. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, "Attention-based models for speech recognition," in *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015*, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds., 2015, pp. 577–585.
- [6] 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 *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014*, A. Moschitti, B. Pang, and W. Daelemans, Eds. ACL, 2014, pp. 1724–1734.
- [7] W. Chan, N. Jaitly, Q. V. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2016*. IEEE, 2016, pp. 4960–4964.
- [8] 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 *Interspeech 2018, 19th Annual Conference of the International Speech Communication Association*, B. Yegnanarayana, Ed. ISCA, 2018, pp. 2227–2231.
- [9] 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 *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2019, Brighton, United Kingdom, May 12-17, 2019*. IEEE, 2019, pp. 6186–6190.
- [10] D. Zhao, T. N. Sainath, D. Rybach, P. Rondon, D. Bhatia, B. Li, and R. Pang, "Shallow-fusion end-to-end contextual biasing," in *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association*, G. Kubin and Z. Kacic, Eds. ISCA, 2019, pp. 1418–1422.
- [11] R. Huang, O. Abdel-Hamid, X. Li, and G. Evermann, "Class LM and word mapping for contextual biasing in end-to-end ASR," in *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association*, H. Meng, B. Xu, and T. F. Zheng, Eds. ISCA, 2020, pp. 4348–4351.
- [12] 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 *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2021*. IEEE, 2021, pp. 7333–7337.
- [13] Z. Chen, M. Jain, Y. Wang, M. L. Seltzer, and C. Fuegen, "Joint grapheme and phoneme embeddings for contextual end-to-end ASR," in *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association*, G. Kubin and Z. Kacic, Eds. ISCA, 2019, pp. 3490–3494.
- [14] M. Han, L. Dong, S. Zhou, and B. Xu, "Cif-based collaborative decoding for end-to-end contextual speech recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2021*. IEEE, 2021, pp. 6528–6532.
- [15] G. Sun, C. Zhang, and P. C. Woodland, "Tree-constrained pointer generator for end-to-end contextual speech recognition," in *IEEE Automatic Speech Recognition and Understanding Workshop, ASRU 2021*. IEEE, 2021, pp. 780–787.
- [16] 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 *Interspeech 2021, 22nd Annual Conference of the International Speech Communication Association*, H. Hermansky, H. Cernocký, L. Burget, L. Lamel, O. Scharenborg, and P. Motlíček, Eds. ISCA, 2021, pp. 1772–1776.
- [17] S. Dingliwal, M. Sunkara, S. Ronanki, J. Farris, K. Kirchhoff, and S. Bodapati, "Personalization of CTC speech recognition models," in *IEEE Spoken Language Technology Workshop, SLT 2022*. IEEE, 2022, pp. 302–309.
- [18] G. Pundak, T. N. Sainath, R. Prabhavalkar, A. Kannan, and D. Zhao, "Deep context: End-to-end contextual speech recognition," in *2018 IEEE Spoken Language Technology Workshop, SLT 2018*. IEEE, 2018, pp. 418–425.
- [19] 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.
- [20] 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.
- [21] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2015*. IEEE, 2015, pp. 5206–5210.
- [22] 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 *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association*, G. Kubin and Z. Kacic, Eds. ISCA, 2019, pp. 2613–2617.
- [23] 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 *Interspeech 2021, 22nd Annual Conference of the International Speech Communication Association*, H. Hermansky, H. Cernocký, L. Burget, L. Lamel, O. Scharenborg, and P. Motlíček, Eds. ISCA, 2021, pp. 3670–3674.
- [24] J. D. Fox and N. Delworth, "Improving contextual recognition of rare words with an alternate spelling prediction model," in *Interspeech 2022, 23rd Annual Conference of the International Speech Communication Association*, H. Ko and J. H. L. Hansen, Eds. ISCA, 2022, pp. 3914–3918.
- [25] 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 *Interspeech 2022, 23rd Annual Conference of the International Speech Communication Association*, H. Ko and J. H. L. Hansen, Eds. ISCA, 2022, pp. 1661–1665.