



Enhancing CTC-based speech recognition with diverse modeling units

Shiyi Han, Zhihong Lei, Mingbin Xu, Xingyu Na, Zhen Huang

Apple Inc.

{shiyi.han, zlei, mingbinxu, zhen.huang, xingyu.na}@apple.com

Abstract

In recent years, the evolution of end-to-end (E2E) automatic speech recognition (ASR) models has been remarkable, largely due to advances in deep learning architectures like transformer. On top of E2E systems, researchers have achieved substantial accuracy improvement by rescoring E2E model's N-best hypotheses with a phoneme-based model. This raises an interesting question about where the improvements come from other than the system combination effect. We examine the underlying mechanisms driving these gains and propose an efficient joint training approach, where E2E models are trained jointly with diverse modeling units. This methodology does not only align the strengths of both phoneme and grapheme-based models but also reveals that using these diverse modeling units in a synergistic way can significantly enhance model accuracy. Our findings offer new insights into the optimal integration of heterogeneous modeling units in the development of more robust and accurate ASR systems.

Index Terms: Automatic Speech Recognition (ASR), Connectionist Temporal Classification (CTC), Joint training

1. Introduction

In recent years, both academia and industry have increasingly adopted end-to-end (E2E) ASR models [1, 2], which predict graphemic units such as wordpieces given audio. One of the production-friendly structures [3, 4] of the E2E system is a multi-pass system that uses a Connectionist Temporal Classification (CTC) [5] model in the first pass and an Attention-Encoder-Decoder (AED) [6, 7] in the second pass. While E2E ASR does not require a pronunciation lexicon, researchers have successfully enhanced the accuracy of end-to-end speech recognition with pronunciations. One approach is to incorporate pronunciations into the modeling units of the E2E system, such as [8, 9, 10, 11]. Without having to change the modeling units, another approach maps rare named entities to common wordpieces through pronunciation to boost the recognition of these named entities, such as [12, 13].

In addition to these works, a standalone phoneme-based acoustic model (AM) can be leveraged, for example, AM fusion [14]. The idea of AM fusion is to incorporate an external AM into the E2E system in a similar way to language model (LM) shallow fusion [15]. The authors adopt a strategy by interpolating scores across multiple recognition passes. In these passes, the N-best hypotheses are generated from the wordpiece E2E model, which are then re-ranked in the rescoring pass.

Despite considerable improvement from AM fusion, it is unclear yet whether AM fusion can generalize to other modeling units. Also, one can argue that the improvement is primarily from the system combination effect and additional model

parameters in the separate AM.

To answer the first question, we started by generalizing AM fusion from phoneme to character. We confirm that a character-based AM achieves similar improvement over a wordpiece-based E2E model. AM fusion with either phoneme or character achieved a relative word error rate (WER) reduction of 7.9% each on test-clean, and 5.6%, 6.0% respectively on test-other.

As an answer to the second question, we attempted to remove the need of a separate model by encoder sharing. More specifically, we jointly trained a wordpiece-phoneme/character model by adding an extra CTC output on an intermediate encoder layer in the baseline wordpiece model. This leads to a training criterion with two CTC losses and a wordpiece AED loss which are linearly combined. We then ran first pass wordpiece CTC decoding. Our experiments show that even without phoneme/character rescoring, the first pass wordpiece CTC WER was improved by up to 7.1% and 6.4% on test-clean and test-other, respectively. This joint training approach does not just provide a simple first pass alternative to AM fusion with similar improvement, but also confirms that the improvement from AM fusion can be achieved by simply providing representations of smaller units without an increase in model size or runtime computation.

To study the aforementioned improvement in details, we conducted an empirical study to understand the information each encoder layer learns. The result is in line with the findings of [16], which utilized attention visualization to demonstrate that initial layers of a model tend to learn pronunciation-related information, while subsequent layers focus on textual information. We also prove that intermediate CTC loss helps produce robust models which is aligned with [17] but our method uses different units and yields superior results. To generalize our finding on not only alphabetic language, we also experimented with logographic language and find similar trend as above.

Our main contributions are as follows: 1. We propose a new joint training strategy that leverages the naturally learned intermediate representations of the network. 2. We investigate the optimal layer matching between different units, resonating with previous studies on the interpretability of speech models. 3. We extend the AM fusion work by improving the Word Error Rate (WER) with models based on alternative units, such as characters.

2. Methodology

2.1. Backbone Model

Our backbone model follows the recipe of [3, 4, 18]. Consider T represents frame number and r is the subsampling factor, a conformer based encoder is used to map the audio feature $\mathbf{X} \in \mathbb{R}^{m \times T_1}$ to high-level representation $\mathbf{H}^{L+1} \in \mathbb{R}^{n \times T_2}$,

where $T_2 = \lfloor T_1/r \rfloor$. To combat the large vocabulary in real-life application, people tend to derive a wordpiece based sequence Y to make the output space tractable. The encoder is jointly trained by minimizing the CTC objective f_{CTC} and the AED objective f_{AED} , and they are parameterized by θ_{CTC} and θ_{AED} respectively:

$$\text{Loss}(X, Y) = f_{CTC}(H^{L+1}, Y; \theta_{CTC}) + \lambda_{AED} \cdot f_{AED}(H^{L+1}, Y; \theta_{AED}). \quad (1)$$

During inference, CTC prefix beam search [5] is used to find the most reasonable hypothesis \hat{Y} within the beam constraint and AED scoring is usually employed to further improve the accuracy.

2.2. Acoustic Model Fusion

On top of the above, AM-fusion firstly trains a frame-synchronous phoneme classifier with an arbitrary structure like conformer or transformer. We use H' and L' to emphasize that it's trained separately with its own parameters. In inference, the top hypothesis is chosen by the weighted summation of the first pass score (CTC in this backbone model) and the score from phoneme-based model:

$$\text{TopHyp}(X) = \arg \max_{\hat{Y}} [f_{CTC}(H^{L+1}, \hat{Y}; \theta_{CTC}) + \lambda_{ph} \cdot f_{ph}(H'^{L'+1}, \hat{Y}_{ph}; \theta_{ph})]. \quad (2)$$

where λ_{wp} is determined by some held-out data. The primary objective of AM fusion lies in integrating an external phoneme-based AM into an E2E ASR system. Notably, [14] has demonstrated a considerable enhancement in accuracy through this approach with a standalone AM fusion model.

While both wordpiece and phoneme-based AMs operate on a frame-synchronous basis, aligning the word boundary introduces challenges, making its implementation non-trivial. To overcome this, the authors adopt a different strategy by interpolating scores across multiple recognition passes. In these passes, the top-K hypotheses are generated using either CTC prefix beam search or Weighted Finite State Transducer [19] (WFST) search. Subsequently, these hypotheses are re-ranked in the rescoring pass using Equation 2.

In this paper, we proposed to further explore the possibility to use different modeling units such as phoneme, character and logographic decomposition when applicable.

2.3. Multiple modeling units joint training

One evident drawback associated with multi-pass recognition is the irrecoverability of early sub-optimal decisions. In particular, the challenge arises in the recognition of rare words, where the statistical prominence of wordpieces learned during training tends to overshadow the uncommonly spelled/pronounced words. To tackle this challenge, we propose an approach involving training the model with an additional CTC output layer followed by an objective focused on other modeling units like characters and phonemes. Given a set of decomposition of interest, we augment the overall objective function by appending the CTC loss corresponding to each representation which leads




Figure 1: An illustration of the proposed model with 4 encoder layers and 2 decoder layers. A phoneme CTC_{ph} is added on top of the 2nd encoder layer.

to Equation 3:

$$\text{Loss}(X, Y) = f_{CTC}(H^L, Y; \theta_{CTC}) + \lambda_{AED} \cdot f_{AED}(H^L, Y; \theta_{AED}) + \sum_{unit} \lambda_{unit} \cdot f_{CTC}(H^{l_{unit}}, Y_{unit}; \theta_{CTC_{unit}}). \quad (3)$$

In the described equation, the term *unit* denotes the specific unit chosen to enhance the base model, incorporating alternative transcription representations such as phonemes, characters, etc. It is also worth noting that the backbone CTC and AED are trained to condition on the last encoder layer, while the CTC_{unit} is trained on an intermediate layer, where $0 < l \leq L$. The notation $H^{l_{unit}}$ represents using different layers of encoder output as the input to the selected unit's CTC_{unit} layer. By using different layers of encoder output for different CTC_{unit} , we can achieve the joint training purpose. During decoding, CTC_{unit} is not involved. Both CTC_{unit} and the AED can be used in the rescoring stage or they can both be individually used in the first pass decoding. The benefits of our proposed approach are manifold. Firstly, leveraging external knowledge acquired during training enables its direct application in the initial pass, without introducing additional inference time complexity. Secondly, the synergy between the semantic knowledge embedded in wordpieces and the phonetic knowledge inherent in phonemes is seamlessly learned. Thirdly, the incorporation of multitask training streamlines the application of external knowledge, eliminating the necessity for multi-pass recognition. This does not only enhance the efficiency of the system but also reduces latency. Finally, the elimination of the standalone phoneme-based model confirms that the improvement from AM fusion comes from the pronunciation information rather than combination of two distinct systems. As a bonus, the space for an extra AM is saved for a deeper and likely more robust model, further contributing to overall system improvement.

While Equation 3 outlines a plausible approach that incorporates secondary knowledge, our experiment indicates that the resulting improvement was less than that achieved by applying the standalone model when the extra unit is trained on the top encoder layer. Although counter-intuitive, this observation aligns with findings from [16]. In their study, two metrics, cumulative attention diagonality (CAD) and phoneme attention relationship (PAR), are introduced to analyze and quantify the linguistic and phonetic modeling capabilities of each layer. Both metrics converge to the same conclusion: lower layers excel at extracting phonetic features, while upper layers are adept at aggregating linguistic information. Furthermore, adopting our notation, [16] constructs a series of classifiers, denoted as

$[C_0, C_1, \dots, C_{L-1}]$, utilizing outputs from different hidden layers $[H^1, H^2, \dots, H^L]$ respectively. The accuracy curve concerning the layer index l exhibits a bell-shaped pattern, with $H^{(L+1)/2}$ ($L = 17$) yielding the most precise phoneme classification.

Such observations suggest that the output of the last hidden layer may not necessarily be the optimal choice for serving the additional objective. Instead, the multitask objective should be extended to encompass the appropriate representation of audio features, denoted as H^l . Furthermore, it is crucial that the objective function aligns the right representation of the audio with the correct representation of the transcripts to maximize its effectiveness. This necessitates a careful pairing of the audio and text representations within the objective function.

It is worth noting that there exist alternative methods for representing the transcripts, utilizing a different set of symbols beyond phonemes and wordpieces. We categorize them into 3 groups:

Orthographic The orthographies of words are decomposed into smaller subword units, such as characters or wordpieces. Byte Pair Encoding (BPE) [20] stands out as one of the most widely used representatives to obtain wordpieces.

Phonetic The phonetic decomposition is derived by linguists, though in practice, often obtained from a grapheme-to-phoneme (G2P) model. The pronunciation dictionary in hybrid ASR [19], the character spelling of phonograms, and Pinyin (a romanization system representing pronunciation) in Chinese fall under this category.

Logographic This form of decomposition finds applicability primarily in logographic languages, where the written system often reflects semantic composition. An illustrative example is the Wubi input method, which decomposes Chinese characters based on their structural components and maps each character to a unique sequence of keystrokes.

Moreover, our observations indicate that certain units (*unit*) require mapping to an appropriately matched encoder layer to optimize outcomes, a finding that echoes the insights of the aforementioned paper [16]. This alignment between specific units and encoder layers underscores the nuanced relationship between model architecture and the efficacy of joint training, highlighting the strategic importance of layer selection in achieving our stated aims.

3. Experiments

3.1. Overview

In this work, we use a conformer CTC-AED model as our baseline modeling structure. This model employs a 12-layer conformer encoder and a 6-layer transformer decoder, each layer using a 4-head self-attention with 256 units. In addition, the decoder layers attend to the encoder outputs using a 4-head, 256-unit cross-attention. All models feature a 4x subsampling convolution block which consumes 80-dimensional log-Mel filterbank features obtained from a 25ms window with a 10ms stride. This model has a total of 47M parameters.

We conduct experiments on English and Chinese Mandarin. We use LibriSpeech [21] to train our English models, while AISHELL-2 [22] for Chinese Mandarin. All baseline English models use a vocabulary of 5000 wordpieces, constructed with the SentencePiece library [23], meanwhile Chinese transcripts are tokenized into characters.

3.2. LibriSpeech

We start by training a baseline wordpiece CTC-AED model using the aforementioned architecture, a phoneme CTC (without AED) model and a character CTC (without AED) model. The phoneme/character model use the same encoder architecture as the baseline model. The phoneme representation is sourced from [24, 25], encompassing 70 distinct symbols. While the character representation is inherently self-contained, we follow the SentencePiece library, treating initial characters and character within a word as distinct units, even though they are originally the same character. The quality of the system is measured by word error rate (WER). For the rest of the evaluation, if the rescoring model is not explicitly stated, a wordpiece CTC prefix beam search is assumed. It is worth noting that the phoneme CTC model is evaluated in conjunction with an n-gram LM provided by LibriSpeech to produce transcription.

We then follow [14] to rescore the N-best hypotheses from the wordpiece model using either the phoneme or the character CTC. Results in Table 1 show a significant relative WER reduction, 7.9% ($3.65 \rightarrow 3.36$) each on test-clean, and 5.6% ($9.26 \rightarrow 8.76$), 6.0% ($9.26 \rightarrow 8.70$) on test-other, using phoneme/character CTC respectively, despite the fact that the phoneme/character CTC is much weaker than the wordpiece baseline.

Loss	Rescoring model	WER	
		clean	other
Separately trained AM			
CTC_{char}^s	-	3.93	10.53
CTC_{ph}^s	-	4.46	10.26
	-	3.65	9.26
$CTC + AED^j$	AED^j	3.21	8.45
	CTC_{mono}^s	3.36	8.74
	CTC_{char}^s	3.36	8.70
Jointly-trained AM			
+ CTC_{ph}^j	-	3.53	8.76
	AED^j	3.13	8.00
+ CTC_{char}^j	-	3.39	8.67
	AED^j	3.06	8.07
	CTC_{char}^j	3.31	8.57

Table 1: We use superscripts s and j to refer to models trained separately as a standalone and models trained jointly along with the backbone, respectively. When the phoneme model is used in first pass, we combine the lexicon and LM from LibriSpeech in order to properly search.

Next, we jointly train wordpiece-phoneme models with an additional phoneme CTC output affixed to the 6th layer with an interpolation weight of 0.1 applied to the loss. Without introducing extra modeling parameters or rescoring pass, the model achieved 3.3% ($3.65 \rightarrow 3.53$) WER reduction on test-clean, and 5.4% ($9.26 \rightarrow 8.76$) WER reduction on test-other.

We further exhaustively affix the CTC loss of phoneme to various layers l . As l increases from encoder layer 1, the WER shows gradual improvement, reaching its peak around the 4th layer before declining. When the loss is attached to the last conformer encoder layer, it starts to hurt word error rate. This trend aligns with the findings of [16], wherein the examination of phoneme modeling capability across different layers similarly demonstrates optimal accuracy in the middle layer. With such compelling evidence, we conclude that the middle layer serves as the most appropriate phonetic representation of the

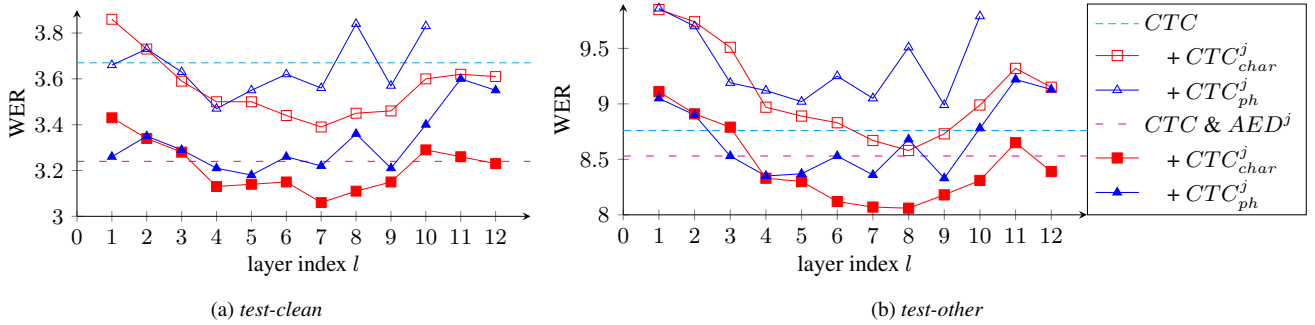


Figure 2: WER on LibriSpeech when different CTC_{units} are attached.

deep Conformer model and, consequently, is the optimal location for integrating pronunciation knowledge.

We replicate the previous experiment with a character-CTC. We observed a more significant improvement, 7.1% (3.65 \rightarrow 3.39) on test-clean and 6.4% (9.26 \rightarrow 8.67) on test-other. Following a similar trend, where the peak WER reduction was most pronounced in the 7th layer, which is 3 layers deeper than phoneme’s optimal layer. Furthermore, we notice that even if we use last layer of encoder for character-based CTC, we still get an on-par accuracy compared to the base model. Our rationale is that, English, as a alphabetic system, incorporates numerous clues to pronunciation within its written form, allowing characters to carry phonetic information. In comparison to phonemes, characters also contain linguistic information to some extent, as certain combinations of spellings convey specific meanings. When we concatenate the character CTC layer into the middle layers, it helps the model fit the pronunciation representations, thereby improving the word error rate. Additionally, as another representation of textual information, characters hold similar information to wordpieces. When concatenated with the last encoder layer, the accuracy is almost on par. Nevertheless, phonemes primarily serve as a phonetic representation without specific meanings. We also try to cooperate both phoneme and character loss. It turns out that there is no further improvement, as the secondary knowledge they bring duplicates.

3.3. AISHELL-2




Figure 3: CER trend on test-ios partition of AISHELL2

To evaluate the effectiveness of logographic representation in Mandarin, we conduct experiments on the AISHELL-2 dataset [22] using Pinyin [26] and Wubi, which respectively represent the phonetic decomposition and structural decomposition of Chinese characters. We use a similar model setting as English experiments’, a 12-layered conformer with CTC-AED loss. Character error rate (CER) serves as the accuracy metric in the corresponding experiments. Results in Table 2 reveals CER

Loss	Pure CTC			Rescored by AED		
	ios	and	mic	ios	and	mic
Baseline	6.84	7.39	7.47	5.97	6.52	6.55
+ CTC_{pinyin}	6.29	6.93	7.07	5.71	6.28	6.43
+ CTC_{wubi}	6.48	7.00	7.10	5.95	6.51	6.57

Table 2: CER on ios / android / mic partitions of AISHELL2

improvements ranging from 5.4% to 8.0% using Pinyin without further AED rescoring, higher than improvements from Wubi keystrokes. This is likely because Pinyin introduces more information that is otherwise less well represented by Chinese characters. Notably, when AED rescoring is used, Pinyin still preserves the majority of the gains, while Wubi’s effect is marginal. We conduct an exhaustive search for suitable intermediate representations. As illustrated in Figure 3, we find that layer 5 provides the most improvement when jointing training with Pinyin CTC, while layer 9 is the most suitable for Wubi CTC joint training. This observation aligns with the premise that the effectiveness of phonetic modeling reaches its peak in the middle layers, whereas the effectiveness of linguistic modeling tends to increase with depth. Wubi’s effectiveness starts to decline after reaching its peak at layer 9, possibly because the word structure does not always fully embed word semantics, and as the model goes deeper, the surrounding context plays a more significant role in representing semantics [27, 28].

3.4. Highlights

These experiments validate the improvements from additional pronunciation information without any form of system combination. More importantly, our method offers additional advantages. It eliminates the need for additional training-time compute resources by avoiding the inclusion of an additional AM. This reduction in computational demand leads to a decrease in the overall system footprint. Our approach obviates the necessity for a second pass rescoring to benefit from the accuracy improvements achieved through AM fusion. Consequently, this reduces both engineering efforts and inference compute resources required during deployment.

4. Conclusion

In this work, we achieve training-time knowledge injection by attaching objective functions corresponding to various modeling units. We show that these objectives should be linked to the hidden layers that excel at representing the information carried by such units in order to properly inject the knowledge.

5. References

- [1] M. Xu, A. Jin, S. Wang, M. Su, T. Ng, H. Mason, S. Han, Y. Deng, Z. Huang, and M. Krishnamoorthy, "Conformer-based speech recognition on extreme edge-computing devices," *arXiv preprint arXiv:2312.10359*, 2023.
- [2] M. Xu, C. Song, Y. Tian, N. Agrawal, F. Granqvist, R. van Dalen, X. Zhang, A. Argueta, S. Han, Y. Deng *et al.*, "Training large-vocabulary neural language models by private federated learning for resource-constrained devices," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [3] Z. Yao, D. Wu, X. Wang, B. Zhang, F. Yu, C. Yang, Z. Peng, X. Chen, L. Xie, and X. Lei, "Wenet: Production oriented streaming and non-streaming end-to-end speech recognition toolkit," in *Interspeech 2021, 22nd Annual Conference of the International Speech Communication Association, Brno, Czechia, 30 August - 3 September 2021*, H. Hermansky, H. Cernocký, L. Burget, L. Lamel, O. Scharenborg, and P. Motlíček, Eds. ISCA, 2021, pp. 4054–4058.
- [4] 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, Incheon, Korea, 18-22 September 2022*, H. Ko and J. H. L. Hansen, Eds. ISCA, 2022, pp. 1661–1665. [Online]. Available: <https://doi.org/10.21437/Interspeech.2022-483>
- [5] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 369–376.
- [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, October 25-29, 2014, Doha, Qatar: A meeting of SIGDAT, a Special Interest Group of the ACL*, A. Moschitti, B. Pang, and W. Daelemans, Eds. ACL, 2014, pp. 1724–1734. [Online]. Available: <https://doi.org/10.3115/v1/d14-1179>
- [7] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015. Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015.
- [8] H. Xu, S. Ding, and S. Watanabe, "Improving end-to-end speech recognition with pronunciation-assisted sub-word modeling," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 7110–7114.
- [9] L. Li, D. Xu, H. Wei, and Y. Long, "Phonetic-assisted multi-target units modeling for improving conformer-transducer asr system," *arXiv preprint arXiv:2211.01571*, 2022.
- [10] Y. Higuchi, K. Karube, T. Ogawa, and T. Kobayashi, "Hierarchical conditional end-to-end asr with ctc and multi-granular subword units," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7797–7801.
- [11] G. Ma, P. Hu, N. Yolwas, S. Huang, and H. Huang, "PM-MMUT: boosted phone-mask data augmentation using multi-modeling unit training for phonetic-reduction-robust E2E speech recognition," in *Interspeech 2022*, H. Ko and J. H. L. Hansen, Eds. ISCA, 2022, pp. 1021–1025.
- [12] Z. Lei, E. Pusateri, S. Han, L. Liu, M. Xu, T. Ng, R. Travadi, Y. Zhang, M. Hannemann, M.-H. Siu *et al.*, "Personalization of ctc-based end-to-end speech recognition using pronunciation-driven subword tokenization," in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 10 096–10 100.
- [13] 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*, H. Meng, B. Xu, and T. F. Zheng, Eds. ISCA, 2020, pp. 4348–4351.
- [14] Z. Lei, M. Xu, S. Han, L. Liu, Z. Huang, T. Ng, Y. Zhang, E. Pusateri, M. Hannemann, Y. Deng, and M. Siu, "Acoustic model fusion for end-to-end speech recognition," in *IEEE Automatic Speech Recognition and Understanding Workshop, ASRU 2023, Taipei, Taiwan, December 16-20, 2023*. IEEE, 2023, pp. 1–7.
- [15] 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," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 1–5828.
- [16] K. Shim, J. Choi, and W. Sung, "Understanding the role of self attention for efficient speech recognition," in *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [17] J. Lee and S. Watanabe, "Intermediate loss regularization for ctc-based speech recognition," in *ICASSP 2021, 2021*, pp. 6224–6228.
- [18] A. Gulati, J. Qin, 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 *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, H. Meng, B. Xu, and T. F. Zheng, Eds. ISCA, 2020, pp. 5036–5040.
- [19] 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.
- [20] R. Sennrich, B. Haddow, and A. Birch, "Neural machine translation of rare words with subword units," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics, 2016.
- [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)*. IEEE, 2015, pp. 5206–5210.
- [22] J. Du, X. Na, X. Liu, and H. Bu, "Aishell-2: Transforming mandarin asr research into industrial scale," *arXiv preprint arXiv:1808.10583*, 2018.
- [23] T. Kudo and J. Richardson, "Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018: System Demonstrations, Brussels, Belgium, October 31 - November 4, 2018*, E. Blanco and W. Lu, Eds. Association for Computational Linguistics, 2018, pp. 66–71.
- [24] L. Lugosch, M. Ravanelli, P. Ignoto, V. S. Tomar, and Y. Bengio, "Speech model pre-training for end-to-end spoken language understanding," in *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*, G. Kubin and Z. Kacic, Eds. ISCA, 2019, pp. 814–818.
- [25] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, "Montreal forced aligner: Trainable text-speech alignment using kaldii," in *Interspeech 2017*, F. Lacerda, Ed. ISCA, 2017, pp. 498–502.
- [26] S. Duanmu, *The phonology of standard Chinese*. OUP Oxford, 2007.
- [27] Z. S. Harris, "Distributional structure," *Word*, vol. 10, no. 2-3, pp. 146–162, 1954.
- [28] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," *Advances in neural information processing systems*, vol. 26, 2013.