



Rethinking Speech Recognition with A Multimodal Perspective via Acoustic and Semantic Cooperative Decoding

Tian-Hao Zhang^{1,2}, Hai-Bo Qin¹, Zhi-Hao Lai^{1,2}, Song-Lu Chen^{1,2}, Qi Liu^{1,2},
Feng Chen^{2,3}, Xinyuan Qian^{1,*}, Xu-Cheng Yin^{1,2}

¹School of Computer and Communication Engineering, University of Science and Technology Beijing, Beijing 100083, China

²USTB-EEasyTech Joint Lab of Artificial Intelligence, University of Science and Technology Beijing, Beijing 100083, China

³EEasy Technology Company Ltd., Zhuhai 519000, China

{tianhaozhang, zhihaolai}@xs.ustb.edu.cn, {songluchen, xinyuanqian}@ustb.edu.cn

Abstract

Attention-based encoder-decoder (AED) models have shown impressive performance in ASR. However, most existing AED methods neglect to simultaneously leverage both acoustic and semantic features in decoder, which is crucial for generating more accurate and informative semantic states. In this paper, we propose an Acoustic and Semantic Cooperative Decoder (ASCD) for ASR. In particular, unlike vanilla decoders that process acoustic and semantic features in two separate stages, ASCD integrates them cooperatively. To prevent information leakage during training, we design a Causal Multimodal Mask. Moreover, a variant Semi-ASCD is proposed to balance accuracy and computational cost. Our proposal is evaluated on the publicly available AISHELL-1 and a1datatang_200zh datasets using Transformer, Conformer, and Branchformer as encoders, respectively. The experimental results show that ASCD significantly improves the performance by leveraging both the acoustic and semantic information cooperatively.

Index Terms: speech recognition, transformer decoder, acoustic and semantic, cooperative decoding

1. Introduction

Automatic speech recognition (ASR) has gained significant advances in recent years, largely driven by the development of deep learning models [1, 2, 3, 4, 5]. Among these models, the attention-based encoder-decoder (AED) architecture has emerged as a powerful methods for improving ASR performance. Specifically, AED combines an encoder network that processes acoustic features with a decoder network that produces text transcriptions. The attention mechanism [6] allows the decoder to selectively focus on different parts of the acoustic features, enabling the model to better capture long-range dependencies, which eventually improves the ASR accuracy.

In the transformer-based AED models [5, 7, 8, 9, 10], the decoder is significant for text transcription from the encoded acoustic features and the previously generated tokens. Therefore, recently, there has been considerable attention on improving the decoder side. Some works [11, 12, 13, 14] focus on improving decoding speed, with the introduction of non-autoregressive decoders to generate predictions of the entire sentence. A bidirectional decoder is proposed in [15] to capture the contextual information in forward and backward directions. To exploit complementary information from different layers, Blockformer [16] is proposed with promising results. There are also works focus on integrating an external language model (LM) into the decoder for better domain adaptation [17, 18, 19]. Despite numerous efforts have been made

* Corresponding author.

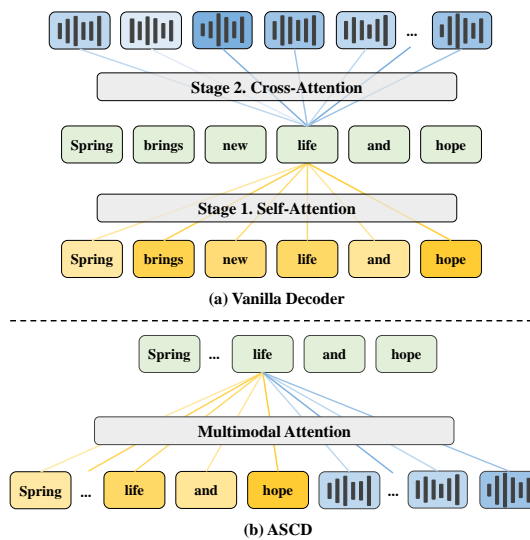


Figure 1: The distinct interaction processes of audio and text features in (a) the Vanilla Decoder and (b) our proposed ASCD. The shading of the color blocks and lines represents the attention weight assigned to each feature.

on decoder design, they overlooked the fundamental issue: the decoder takes acoustic features and semantic features as inputs, how to leverage the complementary information of the two modalities simultaneously is crucial for obtaining high ASR performance.

The vanilla transformer decoder has two attention computational stages as shown in Fig. 1(a). Specifically, self-attention models the semantic features in text to get semantic states for current time. Differently, cross-attention identifies the most relevant encoded acoustic features to subsequently obtain the acoustic states. The residual connection combines the current acoustic and semantic states which are then fed into the feedforward network for further feature extraction. It is worth noting that the self- and cross-attention operate independently when producing the semantic and acoustic states. Considering the inherent correlation between acoustic and semantic features, their separate analysis may lead to inaccurate ASR results, as there is no acoustic information leveraged to generate semantic states.

A recent study [20] in brain science shows that human brain does not convert speech signals into specific tokens in a serial manner. Instead, the primary auditory cortex simply processes acoustic features where the adjacent superior temporal gyrus (STG) extracts semantic features from audio for better speech intelligibility. Experiments verify that the primary auditory cortex and STG are activated concurrently once the cerebral cortex receives an acoustic stimuli. This indicates the parallel

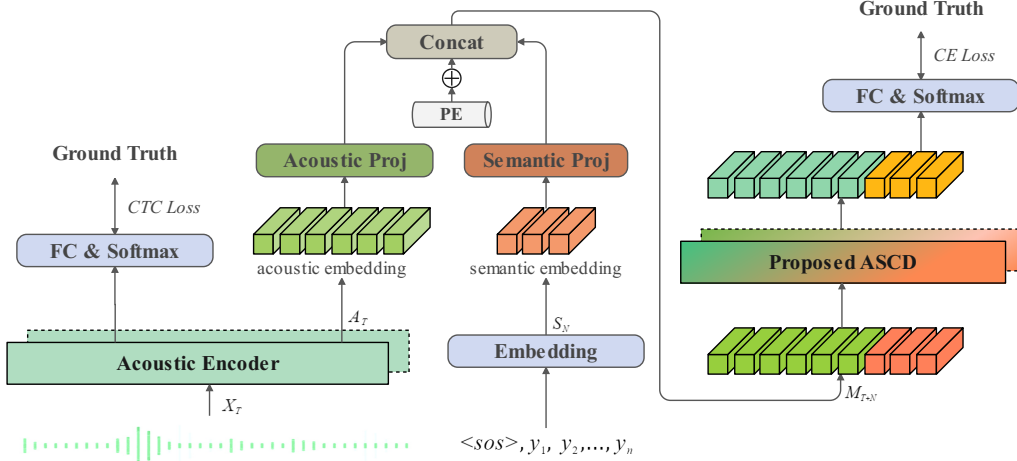


Figure 2: The pipeline of our proposed ASCD based ASR model. The color change of the input and output of modules represents whether the feature is updated or not. Since the encoder side is not the focus of our study, the down-sampling module is omitted for brevity. (PE: positional encoding; FC: fully connected layer; $\langle \text{sos} \rangle$: the start of a sentence).

processing of acoustic and semantic information by the auditory cortex. In agreement with [20], we believe the interaction between acoustic and semantic features in the decoder stage matches their inherent correlation and complementary characteristics where their collaborative processing strategy can promote the ASR results.

Inspired by the remarkable success of multimodal fusion [21, 22], in this paper, we integrate the two fundamental modalities in ASR: the acoustic and semantic information, with a novel proposal of **Acoustic and Semantic Cooperative Decoder (ASCD)**. As shown in Fig. 1(b), ASCD retains only one attention module, where acoustic and semantic features are considered simultaneously and fused dynamically. Different with the general multimodal fusion methods, only the previous semantic information is visible due to the autoregressive decoding, thus a Causal Multimodal Mask is proposed to prevent the information leakage of the training process. Moreover, considering the computational overhead of the models, a variant of ASCD, Semi-ASCD (S-ASCD) is proposed with less computational cost. We conduct experiments on the public available AISHELL-1 and aidaatang_200zh datasets. Transformer [23], Conformer [24], and Branchformer [25] are used as the encoder, respectively. In all cases, our proposed ASCD leads to improved ASR results. In particular, when using the Branchformer as the encoder, we achieve state-of-the-art results where CER equals 4.08% and 4.33% on the AISHELL-1 dev set and test set, respectively. The efficacy of integrating acoustic and semantic modalities through ASCD is validated by visualizing the attention maps. Our contributions are listed as follows:

1. We propose a novel multimodal collaboration strategy for ASR, namely ASCD, which cooperatively integrates the acoustic and semantic modalities in AED decoder with improved performance.
2. Considering the computational overhead, we propose a variant of ASCD, namely S-ASCD, that achieves comparable performance with lower computational resources.
3. We conduct extensive experiments on two publicly available datasets to evaluate our proposed ASCD model. Through visualization of the generated attention maps, we verify that collaboration of the acoustic and semantic modalities in ASCD significantly improves the ASR accuracy.

2. Methods

2.1. Overviews

We build our model upon the AED framework with a multi-task training strategy [26]. The overall architecture includes a convolutional neural network based down-sampling module, an acoustic encoder and the proposed ASCD. The flow chart of the proposed model is shown in Fig. 2. Given the down-sampled speech features \mathbf{X}_T , the acoustic encoder considers the contextual information of the speech to obtain a high-level acoustic embedding \mathbf{A}_T , formulated as:

$$\mathbf{A}_T = \text{Encoder}(\mathbf{X}_T) \quad (1)$$

where $\mathbf{A}_T \in \mathbf{R}^{T \times D_m}$, T is the length of acoustic embedding and D_m is the model dimension. Encoder is a general notation which represents the employed encoder structure. We choose Transformer, Conformer, and Branchformer due to their widespread usage in ASR community.

We use the teacher-forcing method during training. We denote $\mathbf{Y}_N = \{y_1, y_2, \dots, y_N\}$ as the ground truth of the input speech where N indicates the target length. In addition, a special token $\langle \text{sos} \rangle$, which indicates the beginning of a sentence, is added in \mathbf{Y}_N to generate the first prediction. For the expression brevity, we still use N to denote the length of \mathbf{Y} . Then, \mathbf{A}_T and \mathbf{Y}_N are subsequently processed together and passed to ASCD for joint decoding.

2.2. Acoustic and Semantic Cooperative Decoder

Given the text inputs \mathbf{Y}_N , each token is mapped to a high-dimensional representation through the Embedding layer:

$$\mathbf{S}_N = \text{Embedding}(\mathbf{Y}_N) \quad (2)$$

where $\mathbf{S}_N \in \mathbf{R}^{T \times D_e}$ is a semantic embedding with D_e the output dimension. After that, acoustic embedding \mathbf{A}_T and semantic embedding \mathbf{S}_N are projected from the space of their individual modalities to the same subspace through two fully connected acoustic projection layer and semantic projection layer, respectively. Finally, the two projected embeddings are concatenated along time to obtain the multimodal embedding and the positional encoding (PE) is added to provide the model with

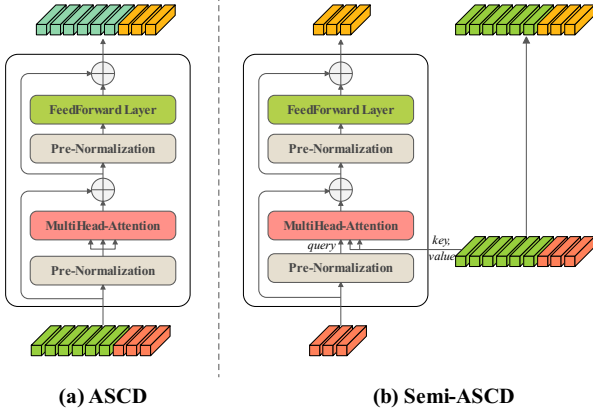


Figure 3: The structure of ASCD and Semi-ASCD layer.

information about the relative position of each token within the input sequence. The process is formulated as:

$$\mathbf{M}_L = [\mathbf{A}_T \mathbf{W}^A; \mathbf{S}_N \mathbf{W}^S] \oplus \text{PE} \quad (3)$$

where L is the sum of T and N . $\mathbf{W}^A \in \mathbf{R}^{D_m \times D_m}$ and $\mathbf{W}^S \in \mathbf{R}^{D_e \times D_m}$ are the weights of the acoustic and semantic projection layers, respectively. \oplus denotes addition operations.

The obtained multimodal embedding \mathbf{M}_L is fed into the ASCD with each layer structured as shown in Fig. 3(a), consisting of only a multi-head attention module and FFN module, both with pre-normalization and shortcut connection. In ASCD, each element of the semantic features interacts with the both acoustic features and the previous predicted semantic features through multi-head attention, such that the relevant features of both modalities are simultaneously involved in reasoning about the current semantic state. At the same time, the acoustic features also interacts internally, and in the same attention computation process as the interaction between the two modalities, so that the acoustic features are updated guided by the semantic information. The process of ASCD can be written as follows:

$$\begin{aligned} \hat{\mathbf{M}}_L &= \text{LayerNorm}(\mathbf{M}_L) \\ \mathbf{M}'_L &= \text{MHA}(\mathbf{q}, \mathbf{k}, \mathbf{v} = \hat{\mathbf{M}}_L) + \mathbf{M}_L \\ \hat{\mathbf{M}}'_L &= \text{LayerNorm}(\mathbf{M}'_L) \\ \mathbf{M}''_L &= \text{FFN}(\hat{\mathbf{M}}'_L) + \mathbf{M}'_L \end{aligned} \quad (4)$$

where MHA represents MultiHead Attention and FFN represents FeedForward Layer.

However, compared to the semantic-only input of vanilla decoder, the multimodal embedding \mathbf{M}_L has an additional length T of acoustic embedding, which is involved in the computation of all components in each ASCD layer and therefore brings additional computation overhead. Therefore, we propose an alternative cooperative method, S-ASCD. As shown in Fig. 3(b), only the semantic part \mathbf{S}_N is computed in all components in S-ASCD, while the multimodal embedding \mathbf{M}_L is only participated in the multi-head attention as key and value. In this way, the reasoning of semantic states still considers the information of both modalities, but the acoustic features is no longer updated. With the high computational overhead removed, the time complexity of the algorithm is decreased from $O((T+N)^2)$ to $O(N(T+N))$, which is consistent with vanilla decoder. The formula process follows:




Figure 4: Details of the mask strategy for ASCD and S-ASCD. The red block represents the masked position which will be assigned the $-inf$ value.

$$\begin{aligned} \hat{\mathbf{S}}_N &= \text{LayerNorm}(\mathbf{S}_N) \\ \mathbf{S}'_N &= \text{MHA}(\mathbf{q} = \hat{\mathbf{S}}_N; \mathbf{k}, \mathbf{v} = \mathbf{M}_L) + \mathbf{S}_N \\ \hat{\mathbf{S}}'_N &= \text{LayerNorm}(\mathbf{S}'_N) \\ \mathbf{S}''_N &= \text{FFN}(\hat{\mathbf{S}}'_N) + \mathbf{S}'_N \\ \mathbf{M}'_L &= [\mathbf{A}_T; \mathbf{S}''_N] \end{aligned} \quad (5)$$

where \mathbf{M}'_L is the updated multimodal embedding used for the next layer. For both ASCD and S-ASCD, we only constrain the semantic features of the output with a cross-entropy loss.

2.3. Masking Strategies

In the attention computation of ASCD, there are four types of interaction between modalities: acoustic-to-acoustic (A-A), acoustic-to-semantic (A-S), semantic-to-acoustic (S-A), and semantic-to-semantic (S-S). To ensure that each element in the S-S interaction is only influenced by the decoded tokens that precede it, an upper triangular mask is employed to shield information from future time steps, effectively avoiding the possibility of any future information leakage. However, the precautionary measures do not end there. In order to prevent the acoustic features from carrying semantic information of future moments, which can cause an unintended breach of information in the subsequent layer, a mask matrix is also necessary for the interaction of S-A. As illustrated in Fig. 4(a), the integration of these two masks produces a Causal Mask that effectively forestalls any potential leakage of semantic information from future time steps.

To maintain the length consistency in each batch, we apply zero-padding to each speech signal and target label during the pre-processing phase. To prevent the model from using the padded region during training, a Padding Mask, as shown in Fig. 4(b), is applied to conceal it which not only ensures consistency but also accelerates the model's convergence. The union of the Causal Mask and Padding Mask results in the ASCD mask as delineated in Fig. 4(c). Moreover, the S-ASCD mask is composed of the A-A and A-S components of the ASCD mask as depicted in the Fig. 4(d).

3. Experiments

3.1. Experimental Setup

We conduct experiments on AISHELL-1 [28] and ai-datatang_200zh¹. To represent the input vectors, we employ a sequence of 80-dimensional log-Mel filter bank with 3-dimensional pitch features. Each frame is computed with a window length of 25 ms and a shift of 10 ms and normalized using Global CMVN [29]. To enhance the robustness of the model, we adopt the speed perturbation [30] and SpecAugment [31] during the training process. Our experiments are conducted using the NVIDIA Tesla V100 GPUs with the same configuration

¹<https://www.datatang.com>

Table 1: Results on the AISHELL-1 dataset. In our Implementations, CTC loss is used on the Conformer and Branchformer, but not on the Transformer.

Models	CER (%)		Params (M)
	Dev	Test	
<i>Published on the ESPnet [27]</i>			
Transformer	5.9	6.4	30.4
Conformer	4.3	4.6	46.2
Branchformer	4.2	4.4	45.4
<i>Ours (Implemented on the ESPnet)</i>			
Transformer	6.07	6.53	29.26
+ ASCD	5.37	5.82	
+ S-ASCD	5.36	5.80	27.81
Conformer	4.28	4.69	46.25
+ ASCD	4.23	4.52	
+ S-ASCD	4.22	4.54	44.80
Branchformer	4.12	4.47	45.43
+ ASCD	4.08	4.33	
+ S-ASCD	4.07	4.38	43.98

as the official ESPnet². We employ the widely used character error rate (CER) metric for evaluation.

Table 2: Results on the aidatang_200zh dataset.

Models	CER (%)		Params (M)
	Dev	Test	
Conformer	3.65	4.31	45.98
+ ASCD	3.58	4.18	
+ S-ASCD	3.54	4.19	44.53

3.2. Results

Table 1 shows the results on the AISHELL-1 dataset. We first report the public results on the ESPnet, where the Conformer and Branchformer results represent the state-of-the-art public level. Subsequently, we present the baseline results obtained through our independent implementation of the three models, followed by the results obtained with the incorporation of ASCD. It is noteworthy that in the case of the Transformer model, we intentionally abstain from incorporating an additional CTC loss [26] for encoder supervision. This deliberate omission is aimed at generating an acoustic embedding of insufficient quality, thereby facilitating a more accurate evaluation of the effectiveness of ASCD in enhancing the performance of the model. As a result, ASCD achieves 5.37% and 5.82% CER results on the dev and test sets, respectively, and S-ASCD achieves 5.36% and 5.80% CER results, with a relative CER reduction (CERR) of 11.18% on the test set compared to the transformer baseline. For Conformer, ASCD and S-ASCD achieve 3.62% and 3.20% CERR on the test set, respectively. When using Branchformer as encoder, ASCD achieves a state-of-the-art results of 4.08%\4.33% CER on the dev\test set, and S-ASCD achieves 4.07%\4.38% CER on the dev\test set. It is observed that the results of ASCD is marginally superior to S-ASCD. This advantage may be attributed to the synchronized update of acoustic features. The outstanding results demonstrate the importance of simultaneously considering acoustic and semantic features for generating accurate semantic states. Additionally,

²<https://github.com/espnet/espnet/tree/master/egs2>

Ground truth

即一些财大气粗的科技公司自己不创新

Branchformer inference

即一些财大**切磋**的科技公司自己不创新

Branchformer with S-ASCD inference

即一些财大**气粗**的科技公司自己不创新

Attention map of top layer in S-ASCD




Figure 5: A decoding example of the AISHELL-1 test set. Red denotes incorrectly predicted tokens, while green represents tokens correctly predicted by considering both acoustic and semantic features (red box in the attention map).

as each layer of ASCD comprises one less attention layer than the conventional decoder layer, the former inherently possesses a lower number of parameters.

The results on the aidatang_200zh dataset are shown in Table 2. ASCD and S-ASCD achieve 3.58%\4.18% and 3.54%\4.19% CER on the dev\test sets, respectively. Compared to the previous best result on the test set, a relative CERR of 3.02% is achieved.

3.3. Analysis

In order to have a comprehensive understanding of our proposed ASCD, we conduct detailed visualizations of specific samples. As shown in Fig. 5, the topmost row represents the ground truth sequence of characters, whereas the second and third rows correspond to the output hypotheses generated by the baseline model and our S-ASCD proposal, respectively. The bottom row is the attention map generated by the top layer of S-ASCD. In this instance, Branchformer erroneously recognizes a result with a similar pronunciation. However, S-ASCD produces a correct prediction by simultaneously attending to both acoustic and semantic features as shown in the attention map. This becomes evident that cooperatively decoding with acoustic and semantic features leads to more accurate results.

4. Conclusion

This study takes a fresh perspective on ASR task by adopting a multimodal fusion approach to effectively leverage both acoustic and semantic features for generating semantic states. Considering the computational cost, we propose a variant of ASCD, namely S-ASCD, which consumes less computation. The proposed methods achieve state-of-the-art results on two public datasets. Furthermore, the effectiveness of considering both acoustic and semantic features simultaneously is demonstrated through visualization analysis. In future work, we intend to explore the combination of ASCD and non-autoregressive methods, as well as investigate more efficient interaction patterns between the acoustic and semantic features.

Acknowledgements: The research is supported by National Key Research and Development Program of China (2020AAA0109700), National Natural Science Foundation of China (62076024, 62006018, U22B2055).

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 *Machine Learning, Proceedings of the Twenty-Third International Conference ICML*, vol. 148, 2006, pp. 369–376.
- [2] A. Graves, "Sequence transduction with recurrent neural networks," *arXiv preprint arXiv:1211.3711*, 2012.
- [3] T. N. Sainath, Y. He, B. Li, A. Narayanan, R. Pang *et al.*, "A streaming on-device end-to-end model surpassing server-side conventional model quality and latency," in *International Conference on Acoustics, Speech and Signal Processing, ICASSP*, 2020, pp. 6059–6063.
- [4] J. Li, R. Zhao, Z. Meng, Y. Liu, W. Wei *et al.*, "Developing RNN-T models surpassing high-performance hybrid models with customization capability," in *21st Annual Conference of the International Speech Communication Association, INTERSPEECH*, 2020, pp. 3590–3594.
- [5] C. Chiu, T. N. Sainath, Y. Wu, R. Prabhavalkar, P. Nguyen *et al.*, "State-of-the-art speech recognition with sequence-to-sequence models," in *International Conference on Acoustics, Speech and Signal Processing, ICASSP*, 2018, pp. 4774–4778.
- [6] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [7] J. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, "Attention-based models for speech recognition," in *Annual Conference on Neural Information Processing Systems, NIPS*, 2015, pp. 577–585.
- [8] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, "End-to-end attention-based large vocabulary speech recognition," in *International Conference on Acoustics, Speech and Signal Processing, ICASSP*, 2016, pp. 4945–4949.
- [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 *International Conference on Acoustics, Speech and Signal Processing, ICASSP*, 2016, pp. 4960–4964.
- [10] S. Karita, X. Wang, S. Watanabe, T. Yoshimura, W. Zhang *et al.*, "A comparative study on transformer vs RNN in speech applications," in *Automatic Speech Recognition and Understanding Workshop, ASRU*, 2019, pp. 449–456.
- [11] Y. Bai, J. Yi, J. Tao, Z. Tian, Z. Wen *et al.*, "Listen attentively, and spell once: Whole sentence generation via a non-autoregressive architecture for low-latency speech recognition," in *21st Annual Conference of the International Speech Communication Association, Virtual Event, INTERSPEECH*, 2020, pp. 3381–3385.
- [12] Y. Higuchi, S. Watanabe, N. Chen, T. Ogawa, and T. Kobayashi, "Mask CTC: non-autoregressive end-to-end ASR with CTC and mask predict," in *21st Annual Conference of the International Speech Communication Association, INTERSPEECH*, 2020, pp. 3655–3659.
- [13] C. Zhang, Y. Liu, T. Zhang, S. Chen, F. Chen, and X. Yin, "Non-autoregressive transformer with unified bidirectional decoder for automatic speech recognition," in *International Conference on Acoustics, Speech and Signal Processing, ICASSP*, 2022, pp. 6527–6531.
- [14] Z. Gao, S. Zhang, I. McLoughlin, and Z. Yan, "Paraformer: Fast and accurate parallel transformer for non-autoregressive end-to-end speech recognition," in *23rd Annual Conference of the International Speech Communication Association, INTERSPEECH*, 2022, pp. 2063–2067.
- [15] X. Chen, S. Zhang, D. Song, P. Ouyang, and S. Yin, "Transformer with bidirectional decoder for speech recognition," in *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, INTERSPEECH*, 2020, pp. 1773–1777.
- [16] X. Ren, H. Zhu, L. Wei, M. Wu, and J. Hao, "Improving mandarin speech recognition with block-augmented transformer," *arXiv preprint arXiv:2207.11697*, 2022.
- [17] Z. Meng, S. Parthasarathy, E. Sun, Y. Gaur, N. Kanda *et al.*, "Internal language model estimation for domain-adaptive end-to-end speech recognition," in *Spoken Language Technology Workshop, SLT*, 2021, pp. 243–250.
- [18] E. Tsunoo, Y. Kashiwagi, C. P. Narisetty, and S. Watanabe, "Residual language model for end-to-end speech recognition," in *23rd Annual Conference of the International Speech Communication Association, INTERSPEECH*, 2022, pp. 3899–3903.
- [19] Z. Meng, N. Kanda, Y. Gaur, S. Parthasarathy, E. Sun *et al.*, "Internal language model training for domain-adaptive end-to-end speech recognition," in *International Conference on Acoustics, Speech and Signal Processing, ICASSP*, 2021, pp. 7338–7342.
- [20] L. S. Hamilton, Y. Oganian, J. Hall, and E. F. Chang, "Parallel and distributed encoding of speech across human auditory cortex," *Cell*, vol. 184, no. 18, pp. 4626–4639, 2021.
- [21] P. K. Atrey, M. A. Hossain, A. El-Saddik, and M. S. Kankanhalli, "Multimodal fusion for multimedia analysis: a survey," *Multim. Syst.*, vol. 16, no. 6, pp. 345–379, 2010.
- [22] L. B. Smith and M. Gasser, "The development of embodied cognition: Six lessons from babies," *Artif. Life*, vol. 11, no. 1-2, pp. 13–29, 2005.
- [23] L. Dong, S. Xu, and B. Xu, "Speech-transformer: A no-recurrence sequence-to-sequence model for speech recognition," in *International Conference on Acoustics, Speech and Signal Processing, ICASSP*, 2018, pp. 5884–5888.
- [24] 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," *ISCA*, vol. abs/2005.08100, 2020.
- [25] Y. Peng, S. Dalmia, I. Lane, and S. Watanabe, "Branchformer: Parallel mlp-attention architectures to capture local and global context for speech recognition and understanding," in *ICML*. PMLR, 2022, pp. 17 627–17 643.
- [26] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, "Hybrid ctc/attention architecture for end-to-end speech recognition," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [27] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba *et al.*, "Espnet: End-to-end speech processing toolkit," in *19th Annual Conference of the International Speech Communication Association, Interspeech*, B. Yegnanarayana, Ed., pp. 2207–2211.
- [28] 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 *20th Conference of the Oriental Chapter of the International Coordinating Committee on Speech Databases and Speech I/O Systems and Assessment, O-COCOSDA*, 2017, pp. 1–5.
- [29] O. Viikki and K. Laurila, "Cepstral domain segmental feature vector normalization for noise robust speech recognition," *Speech Commun.*, vol. 25, no. 1-3, pp. 133–147, 1998.
- [30] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, "Audio augmentation for speech recognition," in *16th Annual Conference of the International Speech Communication Association, INTERSPEECH*, 2015, pp. 3586–3589.
- [31] D. S. Park, W. Chan, Y. Zhang, C. Chiu, B. Zoph *et al.*, "SpecAugment: A simple data augmentation method for automatic speech recognition," in *20th Annual Conference of the International Speech Communication Association, INTERSPEECH*, 2019, pp. 2613–2617.