

UniEnc-CASSNAT: An Encoder-only Non-autoregressive ASR for Speech SSL Models

Ruchao Fan, *Student Member, IEEE*, Natarajan Balaji Shankar, and Abeer Alwan, *Fellow, IEEE*

Abstract—Non-autoregressive automatic speech recognition (NASR) models have gained attention due to their parallelism and fast inference. The encoder-based NASR, e.g. connectionist temporal classification (CTC), can be initialized from the speech foundation models (SFM) but does not account for any dependencies among intermediate tokens. The encoder-decoder-based NASR, like CTC alignment-based single-step non-autoregressive transformer (CASS-NAT), can mitigate the dependency problem but is not able to efficiently integrate SFM. Inspired by the success of recent work of speech-text joint pre-training with a shared transformer encoder, we propose a new encoder-based NASR, UniEnc-CASSNAT, to combine the advantages of CTC and CASS-NAT. UniEnc-CASSNAT consists of only an encoder as the major module, which can be the SFM. The encoder plays the role of both the CASS-NAT encoder and decoder by two forward passes. The first pass of the encoder accepts the speech signal as input, while the concatenation of the speech signal and the token-level acoustic embedding is used as the input for the second pass. Examined on the Librispeech 100h, MyST, and Aishell1 datasets, the proposed UniEnc-CASSNAT achieves state-of-the-art NASR results and is better or comparable to CASS-NAT with only an encoder and hence, fewer model parameters. Our codes¹ are publicly available.

Index Terms—Non-autoregressive ASR, E2E ASR, Self-supervised Learning, Speech Foundation Model

I. INTRODUCTION

IN recent years, self-supervised learning (SSL) has become popular in speech [1], [2], [3] and natural language [4], [5] processing. The SSL models learn prior knowledge from a large amount of unannotated data and are called pre-trained or foundation models. Widely-used speech foundation models include APC [6] that predicts future frames from their histories, and Wav2vec2.0 [7], HuBERT [8], and WavLM [9] that reconstruct or predict pseudo labels via the masked portions of the speech signal. The speech foundation models are proven effective in improving low-resource tasks by fine-tuning [10].

Concurrently, non-autoregressive automatic speech recognition (NASR) has attracted considerable interest due to its fast inference [11], [12], [13]. Although it is not naturally designed for streaming ASR, NASR can greatly improve the inference efficiency for offline applications. As the earliest end-to-end ASR framework, connectionist temporal classification (CTC) [14], [15] can be regarded as an encoder-based NASR model when using greedy decoding. However, the performance of

CTC is always constrained by the output independence assumption. On the other hand, most NASR models are proposed based on the encoder-decoder framework where the decoder can mitigate the output independence problem. For example, the decoder of Mask-CTC [16] is a masked language model to correct the low confidence tokens in CTC output. Align-Refine [17] uses the decoder to refine the CTC alignment iteratively. LASO [18], CASS-NAT [19], and Paraformer [20] extract acoustic embedding as the decoder input for token-level contextual representation learning. However, the encoder-decoder framework does not perfectly fit the current foundation models, which are pre-trained with the transformer encoder structure. Although previous work developed pre-trained models [21], [22] for the encoder-decoder framework, it is specifically designed for autoregressive transformers. Additionally, [23] trains the transformer decoder from scratch with the encoder initialized from the speech foundation model. The work in [24] and [25] introduce BERT to the NASR model for better output dependency modeling. However, these methods may contain unnecessary model parameters.

In this work, based on previous method (CASS-NAT) [26], we present a new encoder-only NASR (UniEnc-CASSNAT) that can function in a way that is similar to CASS-NAT encoder and decoder. Like CTC, UniEnc-CASSNAT can be initialized from speech foundation models (HuBERT base model [8] is used). To behave as both the CASS-NAT encoder and decoder, UniEnc-CASSNAT has two forward passes and accepts two types of input for each. In the first pass, speech features (output of HuBERT Conv. encoder) are fed into the contextual encoder to generate token-level acoustic embeddings (TAEs). In the second pass, the concatenation of speech features and the TAEs (along the time dimension) are used as the contextual encoder inputs. The TAE corresponding outputs are selected for ASR loss computation. The outputs in the second pass can generate better quality TAEs than those in the first pass and hence lead to better ASR performance. We, therefore, further propose a multi-pass CTC (MP-CTC) training and iterative decoding method to improve the WER performance. Experiments on Librispeech 100-hour [27], MyST [28], and Aishell1 [29] datasets show that the proposed methods can achieve better or comparable WERs to CASS-NAT, and contain fewer parameters. The framework can be applied to other encoder-decoder-based NASR.

The remainder of this paper is organized as follows. Section II introduces the framework of UniEnc-CASSNAT and iterative decoding process. Experimental setups are described in Section III. Results are shown and discussed in Section IV. We conclude the paper in Section V.

This work was supported in part by the NSF.

R. Fan and N. B. Shankar are students with the UCLA ECE Department, California, 90095, USA (e-mail: {fanruchao,balaji1312}@ucla.edu).

A. Alwan is a Professor with the UCLA ECE Department, Los Angeles, California, 90095, USA (e-mail: alwan@ee.ucla.edu).

¹https://github.com/Diamondfan/cassnat_asr

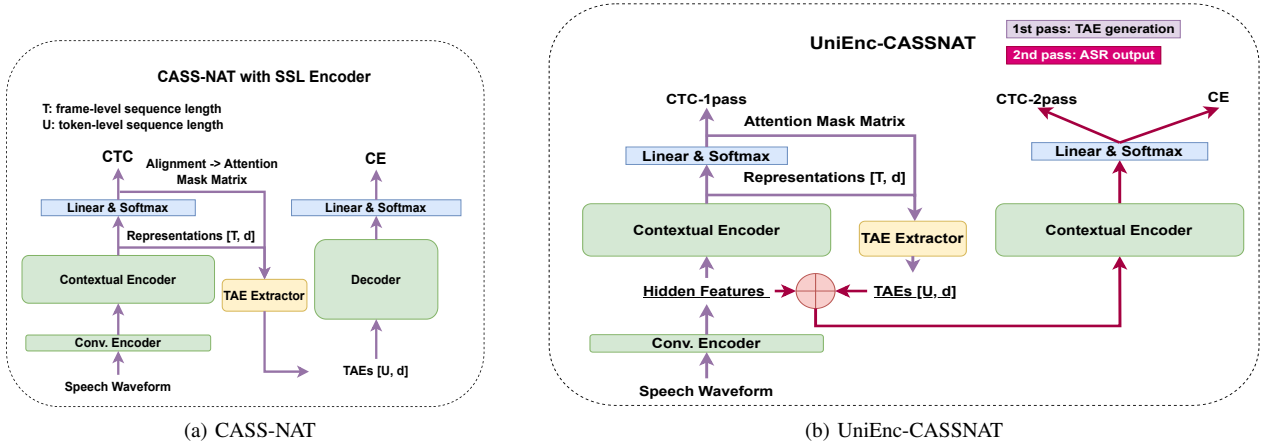


Fig. 1: (a): the diagram of CASS-NAT. (b): the proposed UniEnc-CASSNAT. HuBERT conv. and contextual encoders are used. The TAE extractor is a self-attention module that transforms the acoustic representations with length T to TAEs with length U . The generation of TAEs and second pass forward computation are repeated during iterative decoding.

II. PROPOSED FRAMEWORK: UNIENC-CASSNAT

A. Encoder-Decoder CASS-NAT

CASS-NAT [19] consists of an encoder, a token-level embedding extractor (TAE), and a decoder as plotted in Figure 1(a). The connectionist temporal classification (CTC) [14] loss is added to learn the alignment between the acoustic and token sequences. The alignment can provide segmentation information for each token. TAE extracts an embedding for each token from encoder outputs (with a shape of $[T, d]$, where T is the frame length and d is the hidden dimension) using the segmentation information. The extracted token-level acoustic embeddings (TAEs) (with a shape of $[U, d]$, where U is the token sequence length) are fed into the decoder, which models the relationship between tokens. Suppose the input sequence is $X = \{x_1, \dots, x_t, \dots, x_T\}$, the ground truth is $Y = \{y_1, \dots, y_u, \dots, y_U\}$ and the CTC alignment is Z , then the objective function on the decoder side can be written as:

$$\begin{aligned} L_{\text{dec}} &= \log P(Y|X) \\ &\geq \mathbb{E}_{Z|X} [\log P(Y|Z, X)] \\ &\approx \max_Z \log \prod_{u=1}^U P(y_u | z_{t_{u-1}+1:t_u}, x_{1:T}) \end{aligned} \quad (1)$$

where $t_{u-1} + 1 : t_u$ represents the acoustical boundary for token u provided by the alignment Z . We use a maximum approximation for the expectation (Viterbi-alignment during training). CASS-NAT is then trained by jointly maximizing the decoder loss in Eq. 1 and the CTC loss on the encoder side with a task ratio λ .

$$L_{\text{joint}} = L_{\text{dec}} + \lambda \cdot \log \sum_{Z \in \mathcal{q}} \prod_{i=1}^T P(z_i | X) \quad (2)$$

where \mathcal{q} is all the alignments that can be mapped to the label Y by removing blank tokens in CTC and repetitions.

During decoding, Viterbi-alignment is not available. We therefore use error-based sampled alignments (ESA) (see details in [19]), where the multiple alignments Z are sampled based on the CTC greedy search output with low confidence

scores. The TAEs computed from the sampled alignments are fed into the decoder to obtain multiple ASR outputs. The autoregressive transformer provides a ranking score for each ASR output (one ASR output corresponds to one alignment).

B. Encoder-only CASS-NAT: UniEnc-CASSNAT

Speech foundation models are proven to be useful in downstream ASR tasks. The encoder of CASS-NAT can be inherited from a speech foundation model and extracts better acoustic representations [23]. However, the CASS-NAT decoder has to be trained from scratch. Inspired by the success of recent work of speech-text joint pre-training [30], [31] with a shared encoder, we rethought the necessity of CASS-NAT decoder and propose an encoder-only CASS-NAT, denoted as UniEnc-CASSNAT to fit the size of the speech foundation models.

The UniEnc-CASSNAT is shown in Figure 1(b) with two forward passes. In the first pass, the hidden features extracted from the conv. encoder are fed into the contextual encoder for CTC modeling and the token-level acoustic embeddings (TAEs) are extracted using the alignment information from CTC outputs. In the second pass, the extracted TAEs ($[U, d]$) are concatenated with the hidden features ($[T, d]$) (along the time dimension) to be the input to the contextual encoder. The self-attention layer in the contextual encoder enables frame-frame, frame-token, and token-token interactions between hidden features and TAEs. Note that the goal of the first pass is to obtain TAEs, whose quality is highly related to the ASR performance. The better the speech foundation model, the better the quality of the TAEs extracted by UniEnc-CASSNAT. The second pass is similar to the role of the CASS-NAT decoder for modeling the relationships between TAEs and frame-level hidden features. We investigate whether the encoder is capable of both frame-level acoustic representation learning and contextual modeling between tokens.

C. MP-CTC Training and Iterative Decoding

The output of the second pass is a sequence of $T + U$ vectors, where the first T vectors correspond to hidden fea-

tures, and the U vectors correspond to TAEs. Since the quality of TAEs is essential to the performance of the CASS-NAT decoder, we propose to add another CTC loss to the first T outputs of the second pass and formulate a multi-pass CTC (MP-CTC) training. With the CE loss used on the U outputs, the final objective function of UniEnc-CASSNAT can be written as:

$$L_{\text{unienc-cassnat}} = L_{\text{dec}} + \lambda_1 \cdot L_{\text{CTC-1pass}} + \lambda_2 \cdot L_{\text{CTC-2pass}} \quad (3)$$

We share the final feed-forward layer for the two CTC losses. Theoretically, the second-pass CTC loss would have better performance than the first pass because it accepts additional input information (TAEs). An intuitive idea is to iteratively improve the quality of TAEs by repeating the second pass with newly extracted TAEs. Hence, we propose an iterative decoding method for UniEnc-CASSNAT. Specifically, we define the hidden features as H , and the first pass of UniEnc-CASSNAT encoder as Iter_0 . Iter_0 would generate TAE_0 . The second pass uses $H + \text{TAE}_0$ as input and generates TAE_1 , which we define as Iter_1 . Generally, for iteration n , the contextual encoder accepts H and TAE_{n-1} as input and generates TAE_n for the iteration $n + 1$. In each iteration, ESA generates multiple TAEs for the next iteration. We define the number of sampled alignments in each iteration as S_n . The total number of the sampled alignments would be $\prod_{n=0}^{N-1} S_n$, where N is the number of iterations used in the decoding. We empirically found that two iterations are sufficient for a desirable word error rate (WER).

III. EXPERIMENTAL SETUP

A. Data Settings

The experiments were conducted on three datasets: the 100-hour subset of LibriSpeech English corpus [27], the 240-hour (annotated section) My Science Tutor (MyST) children’s speech corpus [28], and 170-hour Aishell1 Mandarin corpus [29]. We chose the 100-hour subset of Librispeech to enable comparisons with previous work on NASR. We conducted pre-processing on MyST dataset to get a better baseline compared to [23]. For example, we mapped filling pauses, non-speech events, and truncated words to $\langle \text{UNK} \rangle$. The $\langle \text{UNK} \rangle$ is not considered when computing WER.

The sets of output labels consist of 1024 word-pieces for Librispeech 100h and 500 word-pieces for the MyST, obtained by the SentencePiece method [32]. For Aishell, 4230 characters are used as the vocabulary.

B. Model Settings

A CTC/Attention autoregressive transformer (AT) baseline was first trained with an architecture of a 12-block encoder and a 6-block decoder. Suppose the tuple of a transformer setting is represented by (model dimension, feed-forward layer dimension, number of heads in self-attention), we define three settings: d_{512} for (512, 2048, 8), d_{768} for (768, 3072, 12), and d_{256} for (256, 2048, 4). d_{512} is used for the two English datasets, and d_{256} is used for the Aishell1 dataset. Later on, we follow the same setting as in [23] for CASS-NAT training. For a fair comparison, we also include a CTC baseline as

an encoder-only NASR architecture. When training with the speech foundation models, the 12-block encoder was replaced with a HuBERT-base model, either the English² version for Librispeech and MyST, or the Chinese version³ for Aishell1. We also conducted experiments on the TAE extractor in UniEnc-CASSNAT to examine the trade-off between performance and model size.

All models are optimized using a noam scheduler [33] with warmup steps of 15k (10k for Librispeech 100h), a peak learning rate of $5e-5$ for the encoder, and $1e-3$ ($5e-4$ for MyST) for uninitialized modules. The models were trained using a batch size of 80s audio samples (40s for MyST because it contains longer utterances). The training either stops when the WER of the valid set doesn’t improve for 10 epochs or is terminated at 30 epochs. For MP-CTC training, the task ratio of CTC loss in the second pass is set to one.

All results are decoded without the usage of the external language model. For the AT baseline, the beam search decoding is applied with a beam size of 20 for Librispeech and MyST, and 10 for Aishell1. For CASS-NAT, the number of sampled alignments is 50 and the threshold is 0.9. We explore the effects of the number of sampled alignments in two iterations, and the threshold for each iteration is set to 0.9 as well.

IV. RESULTS AND DISCUSSION

A. Main Results

The main WER results of UniEnc-CASSNAT on the Librispeech 100h, MyST, and Aishell1 datasets are shown in Table I. We first train two autoregressive transformer baselines with or without the usage of self-supervised learning. The results in the table again show the effectiveness of the speech foundation models. CASS-NAT achieves close performance to their AT counterpart, which is consistent with previous work. We also present the results of CTC on the three datasets. Due to the output-independent assumption, CTC is worse than the AT baseline and CASS-NAT although it requires fewer parameters. Note that the motivation of UniEnc-CASSNAT is to investigate whether the encoder can jointly model the frame-level and token-level acoustic embedding without the use of the decoder and thus has fewer model parameters. We expect to obtain a model with similar model parameters compared to CTC but close performance to the CASS-NAT. As shown in Table I, the proposed UniEnc-CASSNAT achieves comparable or better results than CASS-NAT, for example, a WER of 11.0% for UniEnc-CASSNAT vs. 11.2% for CASS-NAT on the Librispeech test-other set, but is superior to CASS-NAT in terms of model size (99.3M vs. 130.5M). A smaller model size can be helpful for on-device deployment. Compared to CTC, the UniEnc-CASSNAT achieves much better performance than CTC with a similar model size. The additional 3M parameters compared to CTC (95.7M) are from the TAE extractor. The limitation of UniEnc-CASSNAT could be its slower inference than CTC and CASSNAT because of the multiple forward computations of the encoder with a longer input sequence (concatenation of frames and tokens). The RTF values in

²https://dl.fbaipublicfiles.com/hubert/hubert_base_ls960.pt

³<https://huggingface.co/TencentGameMate/chinese-hubert-base>

TABLE I: WER performance of UniEnc-CASSNAT and comparisons to previous methods on Librispeech-100h, MyST, and Aishell1 datasets. State-of-the-art (SOTA) results with the usage of speech foundation models that are pre-trained with the same amount of unannotated data to ours are reported. The real-time factor (RTF) of each method on Librispeech test-other data is presented for speed comparison. All bold-faced improvements are statistically significant.

Model Type	Model Size	Librispeech-100h				RTF	MyST		Model Size	Aishell1	
		dev-clean	dev-other	test-clean	test-other		dev	test		dev	test
AT-w/o SSL	85.1M	6.6	18.2	6.9	18.2	0.325	13.5	14.9	33.6M	4.6	5.0
AT-w/ SSL	121.6M	4.8	11.0	4.8	10.8	0.486	11.4	13.1	107.3M	4.0	4.3
Non-autoregressive ASR											
Previous SOTA	w/ SSL	4.6	11.3	4.8	11.3 ^[34]	-	16.0	15.6 ^[23]	w/ SSL	3.6	3.8 ^[35]
BERT-CTC [25]	-	7.0	16.3	7.2	16.6	-	-	-	143M	3.9	3.9
CTC	95.7M	6.1	13.8	6.2	13.8	0.005	12.9	14.5	95.7M	4.5	4.9
CASS-NAT	130.5M	4.7	11.4	4.9	11.2	0.014	11.9	13.5	109.7M	4.0	4.3
UniEnc-CASSNAT	99.3M	4.9	11.0	4.8	11.0	0.093	11.8	13.5	102.7M	4.2	4.5

TABLE II: Ablation study of MP-CTC training, the size of the TAE module, and the iterative decoding. d_{256} , d_{512} , and d_{768} are defined in Section III-B and their model sizes (including encoder) are 96.1M, 99.3M, 104.2M, respectively. S_n is the number of sampled alignments in the iteration n .

Model Type	(S_1, S_2)	dev-clean	dev-other	test-clean	test-other
CASS-NAT	(50, NA)	4.7	11.4	4.9	11.2
UniEnc-CASSNAT					
SP-CTC	(50, NA)	4.9	11.9	5.0	11.6
MP-CTC- d_{512}	(50, NA)	5.0	11.7	5.2	11.8
	(50, 1)	5.0	11.1	4.9	11.1
	(25, 2)	4.9	11.0	4.8	11.0
	(10, 5)	4.9	11.1	4.9	11.1
	(5, 10)	4.9	11.1	4.9	11.2
	(2, 25)	5.0	11.4	5.1	11.4
	(1, 50)	5.2	11.5	5.3	11.6
MP-CTC- d_{256}	(25, 2)	4.9	11.4	4.9	11.2
MP-CTC- d_{768}	(25, 2)	4.7	11.2	4.8	11.0

Table I show that the UniEnc-CASSNAT is still 3-5x faster than the AT models although it is 6x slower than CASS-NAT.

The proposed UniEnc-CASSNAT achieves the best-performing NASR results so far in the literature [34], [23] on Librispeech 100h and MyST. One can find better WER performance on the Librispeech 100h data, for example, in [21], [36]. However, in that work, the authors either use a larger model trained with Libri60k hours of data or extra text data. We compare the UniEnc-CASSNAT results to a similar work BERT-CTC [25], which also uses an encoder-only structure. Differently, UniEnc-CASSNAT generates ASR outputs with CE loss instead of CTC loss in BERT-CTC and does not require a pre-trained BERT module (smaller in size than BERT-CTC). In addition, UniEnc-CASSNAT achieves the best performance with two iterations only instead of more than 10 iterations in BERT-CTC (faster inference). Based on the results in Table I, UniEnc-CASSNAT is better on Librispeech-100h but worse on Aishell1 than BERT-CTC. The reason could be that Aishell1 contains simple sentences where a pre-trained BERT model is more beneficial [25] and the BERT-CTC has 143M parameters versus 102M in UniEnc-CASSNAT.

B. Ablation Study of UniEnc-CASSNAT

We present more results on the Librispeech 100h data to show the importance of the proposed MP-CTC training and iterative decoding. First, we set λ_2 in Eq. 3 to zero and train

a UniEnc-CASSNAT with only first-pass CTC. The results in Table II show that the single-pass CTC (SP-CTC) training has a performance gap compared to the CASS-NAT. Additionally, SP-CTC training is not able to perform iterative decoding because $TAE_{n>1}$ is not constrained by CTC outputs. MP-CTC training is also worse than the CASS-NAT without iterative decoding (e.g. (50, NA)). When applying iterative decoding, we explore different combinations of the number of sampled alignments S_n in each iteration. The total number of sampled alignments is set to the same as that used in CASS-NAT for a fair comparison. As shown in Table II, iterative decoding with a setting of (25, 2) achieves the best WER performance and is better than the WER of CASS-NAT. Most of the combinations of S_n achieve comparable WERs to CASS-NAT. It is also noted that the diversity of sampled alignments in the first iteration is more important than that in the second iteration.

Finally, since the TAE extractor introduces extra model parameters besides the foundation model, we conduct experiments of UniEnc-CASSNAT with different transformer settings (d_{256} , d_{512} , d_{768}) described in Section III-B). The results are also shown in Table II. We can see from the table that with a bigger TAE module, the performance tends to be better. However, we select MP-CTC- d_{512} as the final results to show in Table I because MP-CTC- d_{768} did not achieve significant improvements with additional 5M parameters.

V. CONCLUSION

We present a novel encoder-only non-autoregressive ASR (NASR) model, UniEnc-CASSNAT, which integrates the advantage of CTC and CASS-NAT. The encoder of UniEnc-CASSNAT acts as both the encoder and decoder in CASS-NAT to reduce the model parameters and can be well initialized from the speech foundation models. Furthermore, MP-CTC training and iterative decoding are proposed for UniEnc-CASSNAT to further improve the performance to be better or comparable to CASS-NAT. We examined the effectiveness of the proposed methods on the Librispeech 100h, MyST, and Aishell1 datasets. To the best of our knowledge, we have achieved the best-performing WER results for NASR on the first two datasets with the same settings as those in the literature. Future work includes model compression and distillation to further reduce the parameters.

REFERENCES

- [1] Yu Zhang, Daniel S Park, Wei Han, James Qin, Anmol Gulati, Joel Shor, Aren Jansen, Yuanzhong Xu, Yanping Huang, Shibo Wang, et al., “Bigssl: Exploring the frontier of large-scale semi-supervised learning for automatic speech recognition,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1519–1532, 2022.
- [2] Abdelrahman Mohamed, Hung-yi Lee, Lasse Borgholt, Jakob D. Havtorn, Joakim Edin, Christian Igel, Katrin Kirchhoff, Shang-Wen Li, Karen Livescu, Lars Maaløe, et al., “Self-supervised speech representation learning: A review,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1179–1210, 2022.
- [3] Xuankai Chang, Takashi Maekaku, Pengcheng Guo, Jing Shi, Yen-Ju Lu, Aswin Shanmugam Subramanian, Tianzi Wang, Shu-wen Yang, et al., “An exploration of self-supervised pretrained representations for end-to-end speech recognition,” in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 228–235.
- [4] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, et al., “Exploring the limits of transfer learning with a unified text-to-text transformer,” *The Journal of Machine Learning Research*, vol. 21, no. 1, pp. 5485–5551, 2020.
- [5] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of NAACL-HLT*, 2019, pp. 4171–4186.
- [6] Yu-An Chung and James R. Glass, “Generative pre-training for speech with autoregressive predictive coding,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 3497–3501.
- [7] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in neural information processing systems*, vol. 33, pp. 12449–12460, 2020.
- [8] Wei-Ning Hsu et al., “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 29, pp. 3451–3460, 2021.
- [9] Sanyuan Chen, Chengyi Wang, Zhengyang Chen, et al., “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, 2022.
- [10] Ruchao Fan, Yunzheng Zhu, Jinhan Wang, and Abeer Alwan, “Towards better domain adaptation for self-supervised models: A case study of child asr,” *IEEE Journal of Selected Topics in Signal Processing*, 2022.
- [11] Nanxin Chen, Shinji Watanabe, et al., “Non-autoregressive transformer for speech recognition,” *IEEE Signal Process. Letter*, vol. 28, pp. 121–125, 2021.
- [12] Jumon Nozaki and Tatsuya Komatsu, “Relaxing the conditional independence assumption of ctc-based ASR by conditioning on intermediate predictions,” in *Interspeech*, 2021, pp. 3735–3739.
- [13] Yosuke Higuchi, Nanxin Chen, Yuya Fujita, Hirofumi Inaguma, Tatsuya Komatsu, Jaesong Lee, Jumon Nozaki, Tianzi Wang, and Shinji Watanabe, “A comparative study on non-autoregressive modelings for speech-to-text generation,” in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 47–54.
- [14] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen 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.
- [15] Edwin G. Ng, Chung-Cheng Chiu, Yu Zhang, and William Chan, “Pushing the Limits of Non-Autoregressive Speech Recognition,” in *Proc. Interspeech 2021*, 2021, pp. 3725–3729.
- [16] Yosuke Higuchi, Shinji Watanabe, Nanxin Chen, Tetsuji Ogawa, and Tetsunori Kobayashi, “Mask CTC: non-autoregressive end-to-end ASR with CTC and mask predict,” in *Interspeech*, 2020, pp. 3655–3659.
- [17] Ethan A. Chi, Julian Salazar, et al., “Align-refine: Non-autoregressive speech recognition via iterative realignment,” in *Proceedings of NAACL-HLT*, 2021, pp. 1920–1927.
- [18] Ye Bai, Jiangyan Yi, Jianhua Tao, Zhengkun Tian, Zhengqi Wen, and Shuai Zhang, “Fast end-to-end speech recognition via non-autoregressive models and cross-modal knowledge transferring from bert,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1897–1911, 2021.
- [19] Ruchao Fan, Wei Chu, et al., “CASS-NAT: CTC alignment-based single step non-autoregressive transformer for speech recognition,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 5889–5893.
- [20] Zhifu Gao, Shiliang Zhang, et al., “Paraformer: Fast and accurate parallel transformer for non-autoregressive end-to-end speech recognition,” in *Interspeech*, 2022.
- [21] Ziqiang Zhang, Long Zhou, Junyi Ao, Shujie Liu, Lirong Dai, Jinyu Li, and Furu Wei, “SpeechUT: Bridging speech and text with hidden-unit for encoder-decoder based speech-text pre-training,” in *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Abu Dhabi, United Arab Emirates, Dec. 2022, pp. 1663–1676, Association for Computational Linguistics.
- [22] Junyi Ao, Rui Wang, Long Zhou, Chengyi Wang, Shuo Ren, Yu Wu, Shujie Liu, et al., “SpeechT5: Unified-modal encoder-decoder pre-training for spoken language processing,” in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2022, pp. 5723–5738.
- [23] Ruchao Fan, Wei Chu, Peng Chang, and Abeer Alwan, “A ctc alignment-based non-autoregressive transformer for end-to-end automatic speech recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 1436–1448, 2023.
- [24] Keqi Deng, Zehui Yang, et al., “Improving non-autoregressive end-to-end speech recognition with pre-trained acoustic and language models,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 8522–8526.
- [25] Yosuke Higuchi, Brian Yan, Siddhant Arora, Tetsuji Ogawa, Tetsunori Kobayashi, and Shinji Watanabe, “BERT meets CTC: New formulation of end-to-end speech recognition with pre-trained masked language model,” in *Findings of the Association for Computational Linguistics: EMNLP 2022*, Abu Dhabi, United Arab Emirates, Dec. 2022, pp. 5486–5503, Association for Computational Linguistics.
- [26] Ruchao Fan, Wei Chu, Peng Chang, Jing Xiao, and Abeer Alwan, “An improved single step non-autoregressive transformer for automatic speech recognition,” in *Interspeech*, 2021, pp. 3715–3719.
- [27] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [28] Wayne Ward, Ronald Cole, Daniel Bolanos, Cindy Buchenroth-Martin, Edward Svirsky, Sarel Van Vuuren, Timothy Weston, Jing Zheng, and Lee Becker, “My science tutor: A conversational multimedia virtual tutor for elementary school science,” *ACM Transactions on Speech and Language Processing (TSLP)*, vol. 7, no. 4, pp. 1–29, 2011.
- [29] Hui Bu, Jiayu Du, Xingyu Na, Bengu Wu, and Hao 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)*. IEEE, 2017, pp. 1–5.
- [30] Yun Tang, Hongyu Gong, Ning Dong, Changhan Wang, Wei-Ning Hsu, Jiatao Gu, Alexei Baevski, Xian Li, Abdelrahman Mohamed, Michael Auli, and Juan Pino, “Unified speech-text pre-training for speech translation and recognition,” in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, May 2022.
- [31] Gary Wang, Kyle Kastner, Ankur Bapna, Zhehuai Chen, Andrew Rosenberg, Bhuvana Ramabhadran, and Yu Zhang, “Understanding shared speech-text representations,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [32] Taku Kudo and John 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: System Demonstrations*, Brussels, Belgium, Nov. 2018, pp. 66–71, Association for Computational Linguistics.
- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [34] Ruchao Fan, Yiming Wang, Yashesh Gaur, and Jinyu Li, “Ctebert: Advancing hidden-unit bert with ctc objectives,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [35] Guolin Zheng, Yubei Xiao, Ke Gong, Pan Zhou, Xiaodan Liang, and Liang Lin, “Wav-bert: Cooperative acoustic and linguistic representation learning for low-resource speech recognition,” in *Findings of the Association for Computational Linguistics: EMNLP 2021*, 2021, pp. 2765–2777.
- [36] Felix Wu, Kwangyoung Kim, Shinji Watanabe, Kyu J Han, Ryan McDonald, Kilian Q Weinberger, and Yoav Artzi, “Wav2seq: Pre-training speech-to-text encoder-decoder models using pseudo languages,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.