

INTERMEDIATE LOSS REGULARIZATION FOR CTC-BASED SPEECH RECOGNITION

Jaesong Lee¹, Shinji Watanabe²

¹Naver Corporation ²Johns Hopkins University

ABSTRACT

We present a simple and efficient auxiliary loss function for automatic speech recognition (ASR) based on the connectionist temporal classification (CTC) objective. The proposed objective, an intermediate CTC loss, is attached to an intermediate layer in the CTC encoder network. This intermediate CTC loss well regularizes CTC training and improves the performance requiring only small modification of the code and small and no overhead during training and inference, respectively. In addition, we propose to combine this intermediate CTC loss with stochastic depth training, and apply this combination to a recently proposed Conformer network. We evaluate the proposed method on various corpora, reaching word error rate (WER) 9.9% on the WSJ corpus and character error rate (CER) 5.2% on the AISHELL-1 corpus respectively, based on CTC greedy search without a language model. Especially, the AISHELL-1 task is comparable to other state-of-the-art ASR systems based on autoregressive decoder with beam search.

Index Terms— end-to-end speech recognition, connectionist temporal classification, multitask learning, non-autoregressive

1. INTRODUCTION

End-to-end automatic speech recognition (ASR) has become a promising approach for the speech recognition community. It simplifies model design, training, and decoding procedure compared to conventional approaches like hybrid systems using hidden Markov model (HMM).

However, the improvement comes with a computational cost: many state-of-the-art ASR architectures employ attention-based deep encoder-decoder architecture [1–5], which requires heavy computational cost and large model size. Also, the decoder runs in an autoregressive fashion and requires sequential computation, i.e., the generation of an output token can be started only after the completion of the previous token.

Compared to the encoder-decoder modeling, the Connectionist Temporal Classification (CTC) [6] does not require a separate decoder, thus allows designing more compact and fast models. Also, CTC provides a greedy decoding algorithm for generating sentences in a fast and parallel way, especially compared to autoregressive decoder of encoder-decoder models.

Although recent advances on architectural design [7, 8] and pre-training method [9] have improved the performance with CTC, it is usually weaker than encoder-decoder models, often credited to its strong conditional independence assumption, and overcoming the performance often requires external language models (LMs) and beam search algorithm [10, 11], which demand extra computational costs and effectively makes the model an autoregressive one. Therefore, it is important to improve CTC modeling to reduce overall computational overhead, ideally without the help of LM and beam search.

There also has been a great interest on non-autoregressive speech recognition toward reaching the performance of autoregressive models [12–16], inspired by the success of non-autoregressive models in neural machine translation [17–20]. Non-autoregressive ASR would allow faster token generation compared to autoregressive ASR, as the generation of a token does not directly depend on the previous token. CTC itself can be viewed as an early instance of non-autoregressive ASR, and recently proposed methods, Mask CTC [13] and Imputer [14], use CTC as a part of non-autoregressive modeling: they first generate initial output from CTC, then refine it via the other network. Therefore, improving CTC is also important for improving non-autoregressive methods in general.

In this work, we show the performance of CTC can be improved with a proposed auxiliary task. The proposed task, named intermediate CTC loss, is constructed by first obtaining the intermediate representation of the model then computing its corresponding CTC loss. The model is trained with the original CTC loss in conjunction with the proposed loss, with a very small computational overhead. During inference, the usual CTC decoding algorithm is used, thus there is no overhead.

We show the proposed method can improve Transformer [21] with various depths, and also Conformer [5], recently proposed architecture combining self-attention and convolution layers. Also, we show the method can be combined with the other regularization method, stochastic depth [22, 23], for further enhancement.

The contributions of this paper are as follows:

- We present a simple yet efficient auxiliary loss, called intermediate CTC loss, for improving performance of CTC ASR network.
- We combine the intermediate CTC loss and stochastic depth regularization to achieve better performance than using only one of them.
- We show application to the Conformer encoder, recently proposed architecture. We show the proposed method is also effective for Conformer.
- We achieve comparable to state-of-the-art results, specifically word error rate (WER) 9.9% on Wall Street Journal (WSJ) and character error rate (CER) 5.2% on AISHELL-1, using CTC modeling and greedy decoding only.

2. ARCHITECTURE

We consider a multi-layer architecture with the CTC loss function. For given input $x_0 \in \mathbb{R}^{T \times D}$ of length T and dimension D , the encoder consists of L layers as follows:

$$x_l = \text{EncoderLayer}_l(x_{l-1}), \quad (1)$$

where EncoderLayer_l is the l -th layer of the network explained at Section 2.2.

2.1. Connectionist Temporal Classification

CTC [6] computes the likelihood of target sequence y by considering all possible alignments for the label and the input length T . For the encoder output x_L and target sequence y , the likelihood is defined as:

$$P_{\text{CTC}}(y|x_L) := \sum_{a \in \beta^{-1}(y)} P(a|x_L) \quad (2)$$

where $\beta^{-1}(y)$ is the set of alignment a of length T compatible to y including the special blank token. The alignment probability $P(a|x_L)$ is factorized with the following conditional independence assumption:

$$P(a|x_L) = \prod_t P(a[t]|x_L[t]) \quad (3)$$

where $a[t]$ and $x_L[t]$ denote the t -th symbol of a and the t -th representation vector of x_L , respectively.

At training time, we minimize the negative log-likelihood induced by CTC by using $P_{\text{CTC}}(y|x_L)$ in Eq. (2):

$$\mathcal{L}_{\text{CTC}} := -\log P_{\text{CTC}}(y|x_L). \quad (4)$$

At test time, we use greedy search to find the most probable alignment for fast inference.

2.2. Encoder

We use two encoder architectures: Transformer [21] and Conformer [5]. Transformer uses self-attention ($\text{SelfAttention}(\cdot)$) shown in Eq. (5) for learning global representation, and layer normalization [24] and residual connection [25] for stabilizing learning.

With Transformer, **EncoderLayer**(\cdot) in Eq. (1) consists of:

$$x_l^{\text{MHA}} = \text{SelfAttention}(x_{l-1}) + x_{l-1}, \quad (5)$$

$$x_l = \text{FFN}(x_l^{\text{MHA}}) + x_l^{\text{MHA}}, \quad (6)$$

where $\text{FFN}(\cdot)$ denotes the feed forward layers.

Conformer combines Transformer and convolution neural layers for efficient learning of both global and local representations.

With Conformer, **EncoderLayer**(\cdot) in Eq. (1) consists of:

$$x_l^{\text{FFN}} = \frac{1}{2}\text{FFN}(x_{l-1}) + x_{l-1} \quad (7)$$

$$x_l^{\text{MHA}} = \text{SelfAttention}(x_l^{\text{FFN}}) + x_l^{\text{FFN}} \quad (8)$$

$$x_l^{\text{Conv}} = \text{Convolution}(x_l^{\text{MHA}}) + x_l^{\text{MHA}} \quad (9)$$

$$x_l = \text{LayerNorm}(\frac{1}{2}\text{FFN}(x_l^{\text{Conv}}) + x_l^{\text{Conv}}). \quad (10)$$

2.3. Stochastic Depth

Stochastic depth [22, 23] is a regularization technique for residual network. It helps training of very deep networks by randomly skipping some layers. It can be viewed as training an ensemble of 2^L sub-models, induced by removing some layers of the model.

Consider **EncoderLayer**(\cdot) in Eq. (1) with residual connection:

$$x_l = x_{l-1} + f_l(x_{l-1}) \quad (11)$$

for some layer $f_l(\cdot)$.

Let b_l be a Bernoulli random variable which takes value 1 with probability p_l . During training, the layer is computed as:

$$x_l = \begin{cases} x_{l-1} & \text{if } b_l = 0, \\ x_{l-1} + \frac{1}{p_l} \cdot f_l(x_{l-1}) & \text{otherwise.} \end{cases} \quad (12)$$

Thus, with probability $1 - p_l$, the layer skips the $f_l(x_{l-1})$ part. The denominator $\frac{1}{p_l}$ ensures the expectation matches the Eq. (1). During testing, we do not skip the layers and use Eq. (11).

The per-layer survival probability is given as $p_l = 1 - \frac{l}{L}(1 - p_L)$ with hyper-parameter p_L . This assigns higher skipping probability to higher layers, as skipping lower layers may harm the overall performance [22]. We use $p_L = 0.7$ for all experiments.

3. INTERMEDIATE CTC LOSS

The stochastic depth aims to improve training of multi-layer network using a stochastic ensemble approach, but the experiments show the improvement only comes with sufficiently deep networks, e.g. with 24 or more layers [23].

We hypothesize that while the stochastic depth is effective for regularizing higher layers, it is not effective for regularizing lower layers, due to the its ensemble strategy. As each layer has own random variable for skipping, the probability of skipping all high layers is very low. Therefore, for most cases, the lower layers may rely on the remaining higher layers rather than learn regularized representation by themselves.

In this context, we propose to skip the higher layers as a whole. We choose a layer, called ‘‘intermediate layer’’, and induce a sub-model by skipping all layers after the intermediate layer. The sub-model relies on the lower layers rather than higher layers, thus training the sub-model would regularize the lower part of the full model.

For the position of the intermediate layer, this paper mainly uses $\lfloor L/2 \rfloor$, as it seems a safe choice between lower and higher layers. We later discuss other choices at Section 3.1.

As the sub-model and the full model share lower structure, it is possible to denote the output of the sub-model as $x_{\lfloor L/2 \rfloor}$, the intermediate representation of the full-model. Like the full-model, we use a CTC loss for the sub-model:

$$\mathcal{L}_{\text{InterCTC}} := -\log P_{\text{CTC}}(y|x_{\lfloor L/2 \rfloor}). \quad (13)$$

Then we note that the sub-model representation $x_{\lfloor L/2 \rfloor}$ is naturally obtained when we compute the full model. Thus, after computing the CTC loss of the full model, we can compute the CTC loss of the sub-model with a very small overhead. The proposed training objective is the weighted sum of the two losses:

$$\mathcal{L} := (1 - w)\mathcal{L}_{\text{CTC}} + w\mathcal{L}_{\text{InterCTC}}, \quad (14)$$

where we use $w = 0.3$ for all experiments.

During testing, we do not use the intermediate prediction and only use the final representation x_L for decoding.

The intermediate loss can also be used jointly with stochastic depth. We expect the intermediate loss regularizes the lower layers, and the stochastic depth regularizes the higher layers, thus combining them further improves the whole model. We show the empirical result at Section 4.

3.1. Position variants

We also consider different sub-model configurations and investigate their effects. We consider the following variants:

- **Lower than the middle.** Depending on the number of layers L , the optimal ratio of lower layers to higher layers may differ. To find the effect of position of the intermediate loss, we consider lower position than middle, e.g., $\lfloor L/4 \rfloor$, for the sub-model.

- **Multiple sub-models.** We consider multiple sub-models rather than only one. For the number of sub-models K , we compute the following loss:

$$-\frac{1}{K} \sum_{k=1}^K P_{\text{CTC}}(y|x_{\lfloor \frac{kL}{K+1} \rfloor}). \quad (15)$$

For $K = 1$, the loss corresponds to Eq. (13).

- **Random position.** We also consider randomly choosing sub-model among multiple models. We introduce a uniform random variable u with range from $\lfloor L/2 \rfloor$ to $L - 1$, and choose u -th layer for the intermediate representation.

We show the experimental results at Section 4.2.

3.2. Stochastic variant of Intermediate Loss

In Eq. (14), we compute the weighted sum of the two sub-models. Instead, we may compute the stochastic variant of the loss, like stochastic depth, as follows. Let b a Bernoulli random variable which takes value 1 with probability w . the stochastic intermediate CTC objective is:

$$\mathcal{L}' := \begin{cases} \mathcal{L}_{\text{CTC}} & \text{if } b = 0, \\ \mathcal{L}_{\text{InterCTC}} & \text{otherwise.} \end{cases} \quad (16)$$

The loss coincides with Eq. (14) in expectation.

We argue the deterministic version is better than stochastic one for gradient-based learning even if they have same expected value. For the stochastic variant, the loss and its gradient only have access to $\mathcal{L}_{\text{InterCTC}}$ if $b = 1$, and the model may forget features useful for \mathcal{L}_{CTC} but not for $\mathcal{L}_{\text{InterCTC}}$. On the other hand, the deterministic variant always computes two losses at the same time, therefore, the risk of forgetting features is low.

At Section 4.2, we experimentally show while the stochastic variant also improves the model, it is not so effective as the deterministic one.

3.3. Application to other non-autoregressive ASR: Mask CTC

Mask CTC [13] consists of an encoder, a CTC layer on top of the encoder, and a conditional masked language model (CMLM) [18].

During decoding, the model first generates initial hypotheses from the CTC layer, and replaces any token of low probability (below a given threshold) with special token $\langle \text{MASK} \rangle$. The CMLM predicts the token of masked position given the masked hypothesis.

During training, the target y is randomly masked and fed to CMLM. The CMLM predicts the token of masked position for the masked input. Let y_{obs} the masked input and y_{mask} the prediction for the mask. The training objective is:

$$-w_{\text{CTC}} \log P_{\text{CTC}}(y|x_L) - (1-w_{\text{CTC}}) \log P_{\text{CMLM}}(y_{\text{mask}}|y_{\text{obs}}, x_L) \quad (17)$$

with hyper-parameter w_{CTC} .

As the initial hypothesis is predicted from the CTC layer, its performance is crucial for the overall performance. We aim to improve the CTC layer using the proposed intermediate loss. We take the intermediate output $x_{\lfloor L/2 \rfloor}$ from the encoder and compute the intermediate CTC probability $P_{\text{CTC}}(y|x_{\lfloor L/2 \rfloor})$. The extended training objective is:

$$-w_{\text{CTC}} \log P_{\text{CTC}}(y|x_L) - w_{\text{InterCTC}} \log P_{\text{CTC}}(y|x_{\lfloor L/2 \rfloor}) - (1 - w_{\text{CTC}} - w_{\text{InterCTC}}) \log P_{\text{CMLM}}(y_{\text{mask}}|y_{\text{obs}}, x_L). \quad (18)$$

We present the experimental result for Mask CTC at Section 4.3.

3.4. Related work

Hierarchical CTC [26–28] (HCTC) introduced an auxiliary CTC task based on the assumption that different layers learn different level of abstraction. While HCTC looks similar to intermediate loss, it requires additional labeling effort (e.g., phoneme) or various tokenization (e.g., sub-word for high-level and character for low-level), which may not be applicable for certain cases, e.g., when the character-based tokenization is the best effort for the data like Mandarin and Japanese [29]. In contrast, intermediate CTC is based on the sub-model regularization, therefore it does not require additional low-level labels, and it is natural to combine intermediate CTC with stochastic depth.

[30] and [31] introduced additional networks to train the intermediate layer of the encoder for CTC and RNN-Transducer [32] respectively. We note that intermediate CTC does not require additional network, and has very little overhead at the training time, in contrast to [31], due to the structure of CTC architecture.

4. EXPERIMENTS

We evaluate the performance of intermediate CTC loss on the three corpora: Wall Street Journal (WSJ) [33] (English, 81 hours), TED-LIUM2 [34] (English, 207 hours), and AISHELL-1 [35] (Chinese, 170 hours). We use ESPnet [29] for all experiments. We use 80-dimensional log-mel feature and 3-dimensional pitch feature for the input, and apply SpecAugment [36] during training. For WSJ and AISHELL-1, we tokenize label sentences as characters. For TED-LIUM2, we tokenize label sentences as sub-words with sentence-piece [37].

For WSJ, the model is trained for 100 epochs. For TED-LIUM2 and AISHELL-1, the model is trained for 50 epochs. After training, the model parameter is obtained by averaging models from last 10 epochs. Note that we do not use any external language models (LMs) or beam search, and only use greedy decoding for CTC. Thus, all experiments are based on the *non-autoregressive* setup in order to keep the benefit of fast and parallel inference of CTC.

4.1. Results

We show the experimental results for Transformer and Conformer architectures. For each architecture, we compare four regularization configurations:

- Baseline (no regularization)
- Intermediate CTC (“InterCTC”)
- Stochastic depth (“StochDepth”)
- Intermediate CTC + Stochastic depth (“both”)

For Transformer, we use 12-layer, 24-layer and 48-layer models. Table 1 shows the word error rates (WERs) for WSJ and TED-LIUM2, and character error rates (CERs) for AISHELL-1.

For all of the experiment, intermediate CTC gives an improvement over the baseline model. Stochastic depth improves 24-layer and 48-layer models, but does not improve 12-layer models well for WSJ and AISHELL-1. Using both the intermediate loss and the stochastic depth gives better result than using only one of them. Thus, we conclude the two methods have complimentary effects.

Additionally, we apply intermediate CTC to 6-layer Transformer for WSJ, and get WER improvement from 21.1% to 18.3%. This suggests the intermediate CTC is still beneficial for smaller networks.

For Conformer, we use 12-layer model. The results are at Table 2. Again, intermediate CTC gives consistent improvement

Table 1. Word error rates (WERs) and character error rates (CERs) for Transformer. See section 4.1 for details.

	WSJ (WER)		TED-LIUM2 (WER)		AISHELL-1 (CER)	
	dev93	eval92	dev	test	dev	test
12-layer	20.1	16.5	14.8	14.0	5.8	6.3
+ InterCTC	17.5	13.6	13.3	12.3	5.7	6.2
+ StochDepth	19.8	16.2	13.8	13.1	5.9	6.4
+ both	16.8	13.7	13.2	12.1	5.7	6.1
24-layer	17.8	13.9	12.6	12.2	5.4	5.9
+ InterCTC	15.3	12.4	11.5	10.6	5.1	5.6
+ StochDepth	16.3	12.7	11.9	11.2	5.2	5.7
+ both	14.9	11.8	10.9	10.2	5.2	5.5
48-layer	16.6	13.8	11.6	10.9	5.1	5.7
+ InterCTC	14.9	12.6	10.7	10.3	5.1	5.5
+ StochDepth	15.6	12.9	11.0	10.2	5.0	5.4
+ both	14.2	11.8	10.3	9.9	4.9	5.3

Table 2. Word error rates (WERs) and character error rates (CERs) for Conformer. See section 4.1 for details.

	WSJ (WER)		TED-LIUM2 (WER)		AISHELL-1 (CER)	
	dev93	eval92	dev	test	dev	test
12-layer	15.2	12.4	10.5	9.8	5.4	6.0
+ InterCTC	13.4	10.8	9.7	9.1	5.1	5.6
+ StochDepth	13.1	10.8	11.1	10.7	5.2	5.8
+ both	12.0	9.9	10.8	9.9	4.7	5.2

over baseline. Stochastic depth gives improvement for WSJ and AISHELL-1, but does not give improvement for TED-LIUM2.

The combination of intermediate loss and stochastic depth achieves WER **9.9%** for WSJ and CER **5.2%** for AISHELL-1. For WSJ, it outperforms the previously published non-autoregressive results [13, 14, 38], and is close to the state-of-the-art autoregressive result (9.3%) [39]. Also, for AISHELL-1, it outperforms Transformer-based encoder-decoder models [40, 41], and is close to the state-of-the-art autoregressive result (5.1%) [42]. Note that the referred state-of-the-art results use an autoregressive decoder and [42] also uses an external LM. On the other hand, our result is solely based on CTC with greedy decoding, without LM or beam search.

4.2. Study on Intermediate Loss design

To compare the proposed intermediate loss to the position variants (Section 3.1) and the stochastic variant (Section 3.2), we conduct additional experiments for WSJ corpus. The result is at Table 3. We conduct the following experiments:

- **Lower position.** We conduct this variant for the 24-layer model, which is sufficiently deep to consider a lower position. We used 6th layer for the experiment. Despite the deep network, the variant performs slightly worse than the default.
- **Multiple positions.** We conduct this variant for the 24-layer and the 48-layers, which are very deep and more sub-models may help. We use $K = 3$ for 24-layer and $K = 7$ for 48-layer, to select all layer positions of power of 6. It gives a small improvement for the 24-layer, but gives a mixed result for the 48-layer.
- **Random position.** We conduct this variant for all models. The result is mixed: it gives no improvement for the 12-layer and the 24-layer, although a small improvement for 48-layer.

Table 3. Word error rates (WERs) of the intermediate loss variants for WSJ. See Section 4.2 for details.

		dev93	eval92
12-layer	Default	17.5	13.6
	Random	17.4	14.3
	Stochastic	19.0	15.0
24-layer	Default	15.3	12.4
	Lower	15.8	12.9
	Multiple	15.1	12.0
	Random	15.4	12.4
48-layer	Default	14.9	12.6
	Multiple	15.4	12.1
	Random	14.7	12.0

Table 4. Word error rates (WERs) of Mask CTC-based non-autoregressive ASR for WSJ. See Section 4.3 for details.

	threshold	dev93	eval92
12enc-6dec	0.0	16.5	13.5
	0.999	15.7	12.9
+ InterCTC	0.0	14.4	11.6
	0.999	14.1	11.3
Mask CTC [13]	0.999	15.4	12.1
Align-Refine [38]	-	13.7	11.4

- **Stochastic variant.** We conduct this variant for 12-layer model. As discussed in Section 3.2, the stochastic variant is worse than the deterministic one, although it is still better than no regularization.

From the experimental results, we conclude that the proposed design is a simple yet reasonable choice among the variants.

4.3. Application to other non-autoregressive ASR

We present an experimental result of Mask CTC-based non autoregressive ASR and intermediate loss, as described at Section 3.3. The WSJ corpus is used for the experiment. We use $w_{CTC} = 0.3$, and for intermediate CTC variant, we also use $w_{InterCTC} = 0.3$. We use a Transformer model with 12-layer encoder and 6-layer decoder. The model is trained for 500 epochs and parameters of last 60 epochs are averaged.

Table 4 shows the WERs for Mask CTC. The second column indicates the threshold of probability for CTC prediction; Mask CTC uses 0.999 by default. If threshold is 0.0, the model does not use the decoder and just treats the CTC result as the final prediction. We see the intermediate CTC improves the performance of CTC prediction, from 13.5% to 11.6%. We also see the improvement of CTC leads the overall improvement of Mask CTC, as the WER reduced from 12.9% to 11.3%. It is also lower than Align-Refine [38] (11.4%) which improves Mask CTC by modifying the role of CMLM. This shows the intermediate loss helps the training of Mask CTC.

5. CONCLUSION

We present intermediate CTC loss, an auxiliary task for improving CTC-based speech recognition. The proposed loss is easy to implement, has small overhead at training time and no overhead at test time. We empirically show the intermediate CTC loss improves Transformer and Conformer architectures, and combining the loss with stochastic depth further improves training, reaching word error rate (WER) 9.9% on WSJ and character error rate (CER) 5.2% on AISHELL-1, without an autoregressive decoder or external language model.

6. REFERENCES

- [1] Jan K Chorowski et al., “Attention-based models for speech recognition,” in *Proc. NeurIPS*, 2015.
- [2] W. Chan et al., “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *Proc. ICASSP*, 2016.
- [3] L. Dong, S. Xu, and B. Xu, “Speech-transformer: A no-recurrence sequence-to-sequence model for speech recognition,” in *Proc. ICASSP*, 2018.
- [4] Shigeki Karita et al., “Improving Transformer-Based End-to-End Speech Recognition with Connectionist Temporal Classification and Language Model Integration,” in *Proc. Interspeech*, 2019.
- [5] Anmol Gulati et al., “Conformer: Convolution-augmented transformer for speech recognition,” in *Proc. Interspeech*, 2020.
- [6] Alex Graves et al., “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *Proc. ICML*, 2006.
- [7] Vineel Pratap et al., “Scaling up online speech recognition using convnets,” in *Proc. Interspeech*, 2020.
- [8] S. Krivan et al., “Quartznet: Deep automatic speech recognition with 1d time-channel separable convolutions,” in *Proc. ICASSP*, 2020.
- [9] Alexei Baevski et al., “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Proc. NeurIPS*, 2020.
- [10] Yajie Miao, Mohammad Gowayyed, and Florian Metze, “EESSEN: End-to-end speech recognition using deep rnn models and wfst-based decoding,” in *Proc. ASRU*, 2015.
- [11] Sei Ueno et al., “Acoustic-to-word attention-based model complemented with character-level ctc-based model,” in *Proc. ICASSP*, 2018.
- [12] Nanxin Chen et al., “Listen and fill in the missing letters: Non-autoregressive transformer for speech recognition,” 2020.
- [13] Yosuke Higuchi et al., “Mask ctc: Non-autoregressive end-to-end asr with ctc and mask predict,” in *Proc. Interspeech*, 2020.
- [14] William Chan et al., “Imputer: Sequence modelling via imputation and dynamic programming,” in *Proc. ICML*, 2020.
- [15] Yuya Fujita et al., “Insertion-based modeling for end-to-end automatic speech recognition,” in *Proc. Interspeech*, 2020.
- [16] Zhengkun Tian et al., “Spike-triggered non-autoregressive transformer for end-to-end speech recognition,” in *Proc. Interspeech*, 2020.
- [17] Jiatao Gu et al., “Non-autoregressive neural machine translation,” in *Proc. ICLR*, 2018.
- [18] Marjan Ghazvininejad et al., “Mask-predict: Parallel decoding of conditional masked language models,” in *Proc. EMNLP-IJCNLP*, 2019.
- [19] Xuezhe Ma et al., “Flowseq: Non-autoregressive conditional sequence generation with generative flow,” in *Proc. EMNLP-IJCNLP*, 2019.
- [20] Raphael Shu et al., “Latent-variable non-autoregressive neural machine translation with deterministic inference using a delta posterior,” in *Proc. AAAI*, 2020.
- [21] Ashish Vaswani et al., “Attention is all you need,” in *Proc. NeurIPS*, 2017.
- [22] Gao Huang et al., “Deep networks with stochastic depth,” in *Proc. ECCV*, 2016.
- [23] Ngoc-Quan Pham et al., “Very Deep Self-Attention Networks for End-to-End Speech Recognition,” in *Proc. Interspeech*, 2019.
- [24] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton, “Layer normalization,” 2016.
- [25] K. He et al., “Deep residual learning for image recognition,” in *Proc. CVPR*, 2016.
- [26] Santiago Fernández, Alex Graves, and Jürgen Schmidhuber, “Sequence labelling in structured domains with hierarchical recurrent neural networks,” in *Proc. IJCAI*, 2007.
- [27] Kalpesh Krishna, Shubham Toshniwal, and Karen Livescu, “Hierarchical multitask learning for ctc-based speech recognition,” 2019.
- [28] Shubham Toshniwal et al., “Multitask learning with low-level auxiliary tasks for encoder-decoder based speech recognition,” in *Proc. Interspeech*, 2017.
- [29] Shinji Watanabe et al., “ESPnet: End-to-end speech processing toolkit,” in *Proc. Interspeech*, 2018.
- [30] A. Tjandra et al., “Deja-vu: Double feature presentation and iterated loss in deep transformer networks,” in *Proc. ICASSP*, 2020.
- [31] Chunxi Liu et al., “Improving rnn transducer based asr with auxiliary tasks,” in *Proc. SLT*, 2020.
- [32] Alex Graves, “Sequence transduction with recurrent neural networks,” 2012.
- [33] Douglas B Paul and Janet M Baker, “The design for the wall street journal-based CSR corpus,” in *Proc. Workshop on Speech and Natural Language*, 1992.
- [34] Anthony Rousseau, Paul Deléglise, and Yannick Estève, “Enhancing the TED-LIUM corpus with selected data for language modeling and more TED talks,” in *Proc. LREC*, May 2014.
- [35] Hui Bu et al., “Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline,” in *Proc. O-COCOSDA*, 2017.
- [36] Daniel S Park et al., “SpecAugment: A simple data augmentation method for automatic speech recognition,” in *Proc. Interspeech*, 2019.
- [37] Taku Kudo and John Richardson, “SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing,” in *Proc. EMNLP: System Demonstrations*, Nov. 2018.
- [38] Ethan A. Chi, Julian Salazar, and Katrin Kirchhoff, “Align-refine: Non-autoregressive speech recognition via iterative realignment,” 2020.
- [39] Sara Sabour, William Chan, and Mohammad Norouzi, “Optimal completion distillation for sequence learning,” in *Proc. ICLR*, 2019.
- [40] S. Karita et al., “A comparative study on transformer vs rnn in speech applications,” in *Proc. ASRU*, 2019.
- [41] Zhifu Gao et al., “San-m: Memory equipped self-attention for end-to-end speech recognition,” in *Proc. Interspeech*, 2020.
- [42] Xinyuan Zhou et al., “Self-and-mixed attention decoder with deep acoustic structure for transformer-based lvcsr,” in *Proc. Interspeech*, 2020.