



# Guiding Frame-Level CTC Alignments Using Self-knowledge Distillation

Eungbeom Kim<sup>1\*</sup>, Hantae Kim<sup>2†</sup>, Kyogu Lee<sup>1,3</sup>

<sup>1</sup>IPAI, Seoul National University, Republic of Korea

<sup>2</sup>NAVER Cloud, Republic of Korea

<sup>3</sup>Dept. of Intelligence and Information, AIIS, Seoul National University, Republic of Korea

eb.kim@snu.ac.kr, hantae.kim@navercorp.com, kglee@snu.ac.kr

## Abstract

Transformer encoder with connectionist temporal classification (CTC) framework is widely used for automatic speech recognition (ASR). However, knowledge distillation (KD) for ASR displays a problem of disagreement between teacher-student models in frame-level alignment which ultimately hinders it from improving the student model's performance. In order to resolve this problem, this paper introduces a self-knowledge distillation (SKD) method that guides the frame-level alignment during the training time. In contrast to the conventional method using separate teacher and student models, this study introduces a simple and effective method sharing encoder layers and applying the sub-model as the student model. Overall, our approach is effective in improving both the resource efficiency as well as performance. We also conducted an experimental analysis of the spike timings to illustrate that the proposed method improves performance by reducing the alignment disagreement.

**Index Terms:** speech recognition, knowledge distillation

## 1. Introduction

Automatic speech recognition (ASR) is a sequence labeling task that aims to translate a given speech utterance into a transcript [1, 2, 3, 4, 5]. To train the ASR model, Connectionist Temporal Classification (CTC) [6] is a widely used method when the speech-transcript alignment is unknown, which is a common setting for real-world data. By combining CTC and scalable self-supervised learning methods, recent studies [7, 8, 9, 10, 11, 12] lead to a substantial performance gain with large-scale ASR architectures.

Moreover, the large-scale models are computationally expensive and memory-inefficient. For this reason, various methods including knowledge distillation (KD) [13] are applied to attain competitive results under limited environment. Conventional KD leverages a large teacher model to train a small student model, in which the student and the teacher models have independent models. In addition, [14, 15, 16] explain the self-knowledge distillation framework (SKD) which distills the knowledge of the model within itself.

Although CTC enables us to train the ASR models successfully without the frame-level input-label alignment, the alignment estimated by CTC can be unstable and arbitrary [17, 18, 19]. Thus, directly applying the knowledge distillation of the teacher model's frame-level alignment to the student model may possess two major risks in performance degradation [17].

The first issue is originated from instability of the trained alignment. Since a single transcript can be mapped over various

alignments, the teacher and student models' alignments can also be diverse. This alignment disagreement distracts the knowledge distillation process which disturbs the convergence of the student model [20, 21]. To mitigate this issue, several methods are proposed. [20, 22] only utilize sequence-level information extracted from the frame-level alignment. However, the removal of useful frame-level information results in a limited performance increase. Although Guide-CTC [18] distills the useful frame-level information from the teacher model alignment, teacher-student alignment disagreement remains problematic [23].

Likewise, most of the previous methods are limited to post-processing of the teacher model's alignment because the previous methods adhere to the teacher-student knowledge distillation framework from which the alignment disagreement problem stems. This study aims to alleviate this alignment disagreement problem in advance by proposing an in-processing solution. In order to achieve our goal, we explore a new knowledge distillation strategy for CTC based on a self-knowledge distillation upon the intermediate CTC [24] method. Since self-knowledge distillation method such as [14] includes teacher and student in the same model, with the teacher inheriting the student's output sequences, the disagreement in the teacher-student alignments is intrinsically mitigated compared to the independent teacher and student framework. We experimentally verify this argument by quantitatively comparing the alignments of the teacher and student models for the self-knowledge distillation and the conventional knowledge distillation.

Secondly, it is believed that the sparse spike timing of CTC [6, 25, 26] hinders knowledge distillation because it means that the other frames are filled with dummy blank tokens [25, 20, 18, 23]. Therefore, selection of informative frames has been regarded as the key to a successful knowledge distillation, and [18, 23] propose blank frame masking for knowledge distillation by removing blank frames in which the blank token has the maximum probability. However, as shown in Figure 1, although blank frame masking is able to alleviate undesired knowledge distillation, it also mask useful frames, that results in a limited performance improvement. Moreover, blank frame is also an important component of the alignment and contributes to the output transcription [27, 28]. For this reason, we claim that the limited performance improvement of knowledge distillation of CTC is not originated from sparse spike timing but the alignment disagreement problem. Thus, we explore the effect of blank frame masking under self-knowledge distillation environment in which the alignment disagreement problem has been mitigated. Unlike the previous teacher-student knowledge distillation method, it is observed that self-knowledge distillation without blank frame masking surpasses the other with blank frame masking, which highlights the importance of dis-

\* Work done during internship at NAVER Cloud.

† Corresponding author.




Figure 1: Illustrations of knowledge distillation for ASR, comparing frame-level teacher-student alignments when the teacher model outputs the word "CAT". Red arrow denotes undesired distillation and blue arrow denotes desired distillation. (a) and (b) include undesired distillation and miss desired distillation while (c) only includes desired distillation.

tilling blank frames. In other words, we reveal that sparse spike timing is not a major factor of limited performance enhancement under the self-knowledge distillation method.

In summary, we introduce a simple and effective frame-level self-knowledge distillation framework for CTC-based ASR models. We address the alignment disagreement problem of the teacher-student knowledge distillation framework by utilizing the self-knowledge distillation framework. At the same time, we lessen the information loss which stems from blank frame masking strategy of the conventional knowledge distillation methods. To verify this, we examine the belief that blank frames deteriorate the knowledge distillation performance and confirm that it is false when the alignment disagreement problem is insignificant such as in self-knowledge distillation. Based on this result, we experimentally prove the usefulness of the proposed method.

## 2. Background

### 2.1. Connectionist Temporal Classification (CTC)

Given the input sequence  $x = [x_1, \dots, x_F]$  where  $x_f \in \mathbb{R}^D$  denotes  $D$ -dimensional acoustic speech feature of frame  $f$  and label sequence  $y = [y_1, \dots, y_N]$  of  $y_n \in \mathcal{Y}$  where  $\mathcal{Y}$  denotes a label set, automatic speech recognition aims to translate  $x$  into  $y$ . In many cases, however, the sequence length of the input and the label are different (i.e.,  $F \gg N$ ), and the optimal alignment between them is unknown. CTC [6] is a useful framework to address this issue. For input  $x$ , CTC computes the possible alignment set  $y^a = [y_1^a, \dots, y_F^a] \in \beta^{-1}(y)$  where  $y_f^a \in \mathcal{Y}'$  is the element of the blank-augmented label set  $\mathcal{Y}' = \mathcal{Y} \cup \phi$  and  $\phi$  is a blank token. An automatic speech recognition model  $h$  with  $L$ -layer is trained to minimize  $\mathcal{L}_{\text{CTC}}$  under the CTC framework as follows:

$$\mathcal{L}_{\text{CTC}} = -\log \sum_{y^a \in \beta^{-1}(y)} p(y^a | x^L), \quad (1)$$

where  $x^L$  is a  $L$ -th layer output (i.e.,  $x^L = [x_1^L, \dots, x_F^L] = h(x) = h^L \circ h^{L-1} \dots \circ h^1(x)$ ).

### 2.2. Frame-Level Knowledge Distillation

[20] explores the frame-level knowledge distillation method for CTC following the conventional knowledge distillation framework. Frame-level knowledge distillation uses the frame-level output of the teacher model to teach the student model. Formally, the loss of frame-level knowledge distillation with cross-

entropy is formulated as follows:

$$\mathcal{L}_{\text{KD}}^{\text{frame}} = -\sum_{f=1}^F \sum_{a \in \mathcal{Y}'} p(a|x_f^T) \log p(a|x_f^S), \quad (2)$$

where  $x^T = [x_1^T, \dots, x_F^T]$  is the output of the teacher model and  $x^S = [x_1^S, \dots, x_F^S]$  is the output of the student model.

To stabilize the frame-level knowledge distillation, softmax-level knowledge distillation (Sftmx-KD) [21] is proposed. Softmax-level knowledge distillation substitutes the loss function of the frame-level knowledge distillation with  $l_2$  loss function, which is stable due to the  $l_2$  loss function being bounded. The loss is formulated as follows:

$$\mathcal{L}_{\text{KD}}^{\text{sftmx}} = \sum_{f=1}^F \sum_{a \in \mathcal{Y}'} (p(a|x_f^T) - p(a|x_f^S))^2. \quad (3)$$

### 2.3. Frame-Level Knowledge Distillation with Masking

Previous studies [20, 22] confirm that frame-level knowledge distillation has limited performance improvement. One of the possible reasons is that spike timings are sparse, meaning that considerable frames are filled with a dummy blank token. Guide-CTC (G-CTC) [18] challenges this problem by masking the blank frames of the teacher model in knowledge distillation as follows:

$$\mathcal{L}_{\text{KD}}^{\text{Guide-CTC}} = -\sum_{f=1}^F \sum_{a \in \mathcal{Y}'} M_{f,a} \log p(a|x_f^S), \quad (4)$$

$$M_{f,a} = \mathbb{1}(\arg \max_{\bar{a} \in \mathcal{Y}'} p(\bar{a}|x_f^T) = a \neq \phi), \quad (5)$$

where  $M_{f,a}$  denotes a mask function. In short, Guide-CTC utilizes the frame-level greedy prediction of the teacher as a label only when the teacher's prediction is not a blank token.

### 2.4. Intermediate CTC and Layer Pruning

Based on the conventional CTC loss, Intermediate CTC [24] also utilizes the intermediate output of the  $l$ -th layer  $x^l = [x_1^l, \dots, x_F^l] = h^l \circ h^{l-1} \dots \circ h^1(x)$  to regularize the model using the intermediate CTC loss as follows:

$$\mathcal{L}_{\text{iCTC}} = -\log \sum_{y^a \in \beta^{-1}(y)} p(y^a | x^l), \quad (6)$$

where  $0 < l < L$ . The total loss function  $\mathcal{L}$  with intermediate CTC is formulated as  $\mathcal{L} = (1 - \alpha)\mathcal{L}_{\text{CTC}} + \alpha\mathcal{L}_{\text{iCTC}}$ . Upon this, it is possible to directly distill the teacher's output to the student's

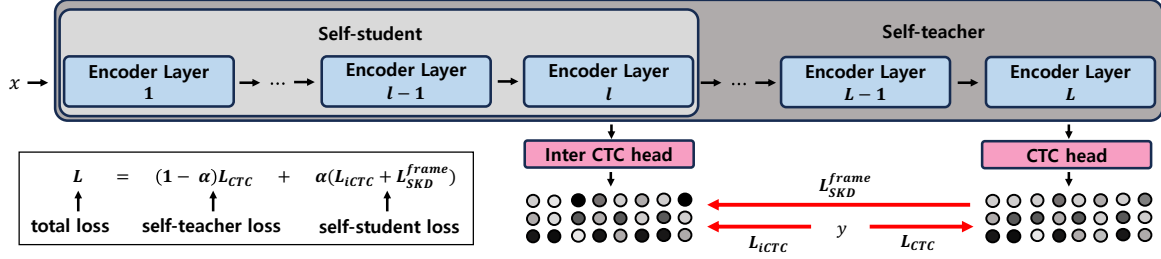


Figure 2: A framework for self-knowledge distillation of  $L$ -layer CTC-based ASR model with  $l$ -layer student model. The total training loss contains self-teacher loss and self-student loss.

intermediate outputs as presented in [29]. Also, [30] extends intermediate CTC to a layer pruning method by removing the layers after the intermediate CTC layer. For the experiment, we utilize this method for the layer pruning baseline.

### 3. Method

#### 3.1. Frame-Level Self-knowledge Distillation

Despite the success of the teacher-student knowledge distillation framework, knowledge distillation for CTC model suffers from the alignment disagreement problem. The alignments of teacher and student models can be diverse since CTC takes into account every possible alignment between an input and a label for training, meaning that the alignment is not fixed. We tackle this issue by introducing frame-level self-knowledge distillation (SKD) for CTC. The overall architecture of SKD is illustrated in Figure 2. The proposed method utilizes the intermediate CTC framework into the self-knowledge distillation framework. Specifically, the self-teacher model consists of CTC head and layers  $1, \dots, L$ , and the self-student model consists of intermediate CTC head and layers  $1, \dots, l$  so that  $1 \leq l < L$  which are shared with the self-teacher. In other words, SKD contains the teacher and student models in a single architecture, and the teacher inherits frame-level information of the student. Thus, the alignment disagreement problem can be intrinsically decreased. Also, SKD is efficient in memory and training time compared to KD because SKD only needs the single architecture. The loss function of the proposed method is formulated as follows:

$$\mathcal{L}_{SKD}^{\text{frame}} = - \sum_{f=1}^F \sum_{a \in \mathcal{Y}'} p(a|x_f^L) \log p(a|x_f^l), \quad (7)$$

where a stop gradient operation is applied to the output of the teacher model. Blank frame masking is not applied to the proposed method because we observe that blank frame masking is inappropriate for self-knowledge distillation in which the alignment disagreement is not significant. Detailed experimental results are introduced in section 4.2. Finally, the total loss function is formulated as follows:

$$\mathcal{L} = (1 - \alpha)\mathcal{L}_{CTC} + \alpha(L_{iCTC} + \mathcal{L}_{SKD}^{\text{frame}}), \quad (8)$$

where  $\alpha \in [0, 1]$  is a hyperparameter that decides the training strength of the student.

#### 3.2. Frame-Level Self-knowledge Distillation with Scheduling

The teacher and the student are independent and trained in a row at the conventional teacher-student KD methods. SKD

can also directly utilize this strategy by sequentially setting  $\alpha = 0$  and  $\alpha = 1$  for the total loss,  $\mathcal{L}$ . However, setting  $\alpha = 1$  for self-knowledge distillation leads to a catastrophic forgetting of the self-teacher model since the shared model is trained only with the self-student model's loss without considering the self-teacher model's loss. So, we introduce a scheduling function  $\alpha = \tau(e, E)$  such that  $\tau(e, E) < 1$  for self-knowledge distillation where  $e$  is training epoch and  $E$  is the total training epoch. Also, we consider  $\tau(e, E)$  which satisfies  $\frac{1}{E} \sum_{e=1}^E \tau(e, E) = 0.5$ , meaning that the average of hyperparameter  $\alpha$  is fixed to 0.5 to preserve the total training volume of the student model. Based on the linear scheduling  $\tau(e, E) = \frac{e-1}{E-1}$  from  $\alpha = 1$  to  $\alpha = 0$ , we introduce the clipped linear scheduling for  $\alpha$  to enforce  $\tau(e, E) < 1$  by  $\tau(e, E) = \min\{\max\{\frac{e-1}{E-1}, t\}, 1-t\}$ . Following the original intermediate CTC [24] which uses  $a = 0.3$ , we empirically set  $t = 0.3$  for all experiments.

## 4. Experiments

#### 4.1. Experimental Setup

For the experiment, pre-trained 12-layer transformer encoder-based models are considered: HuBERT Base [9] and WavLM Base+ [12]. 6, 8 and 10-layer models are used as the student models with the 12-layer teacher models. All models smaller than the 12 layers is simply initialized using the first  $l$  layers from the pre-trained models for a fair comparison of the baseline, KD, and SKD methods. We also experiment with a larger model setting using an 18-layer teacher model and a 12-layer student model. To initialize the 18-layer teacher model, the last 6 layers from the 12-layer pre-trained models are duplicated and repeated after the 12 layers.

For fine-tuning, we use 100-hour train-clean subset from LibriSpeech [31] dataset. We use the learning rate of  $3e-5$  with AdamW [32] optimizer using 10% warmup steps and the cosine learning rate decaying for  $E = 200$ . The models are optimized with a batch size of 128. The feature encoder is fixed during training and the model is frozen for the first 12.5% of the steps except the linear CTC head. We use a character-level tokenizer, and the language model is not utilized for decoding. Other details follow the original HuBERT and WavLM models. Note that the baseline results are reproduced.

#### 4.2. Results

Table 1 shows the word error rate (WER, %) of the automatic speech recognition task. We compare the two KD methods including Guide-CTC (G-CTC) and softmax-based KD

Table 1: Overall performances in terms of WER (%) of the LibriSpeech dataset. Four different student layers are utilized: 12, 10, 8, 6.

Model	Method	18 → 12		12 → 10		12 → 8		12 → 6	
		test-clean	test-other	test-clean	test-other	test-clean	test-other	test-clean	test-other
HuBERT	Baseline (T)	5.8	13.0	6.3	14.4	6.3	14.4	6.3	14.4
	Baseline (S)	6.3	14.4	10.3	22.8	18.4	36.2	34.9	55.8
	G-CTC (S)	6.0	13.7	8.7	19.7	16.8	33.8	31.7	53.1
	Sftmx-KD (S)	6.4	14.6	9.8	22.0	21.8	41.0	34.7	55.7
	Layer Prun (S)	<b>5.6</b>	<b>13.3</b>	8.4	20.2	16.3	34.2	30.9	53.1
	SKD (S)	5.8	13.5	<b>8.1</b>	<b>19.4</b>	<b>15.5</b>	<b>33.4</b>	<b>28.3</b>	<b>51.1</b>
WavLM	Baseline (T)	4.7	9.8	5.2	11.4	5.2	11.4	5.2	11.4
	Baseline (S)	5.2	11.4	9.6	21.5	20.7	38.8	28.9	51.1
	G-CTC (S)	8.5	12.5	9.3	18.0	18.5	34.1	26.9	48.3
	Sftmx-KD (S)	5.2	11.6	9.6	21.5	18.5	37.1	28.2	50.7
	Layer Prun (S)	7.2	11.5	9.2	17.1	14.3	30.8	25.8	48.6
	SKD (S)	<b>5.0</b>	<b>11.0</b>	<b>7.1</b>	<b>16.8</b>	<b>13.4</b>	<b>30.0</b>	<b>22.8</b>	<b>46.0</b>

Table 2: WER (%) of LibriSpeech test-other dataset. Student models are distilled from the 12-layer HuBERT Base model.

# student layers	Method	test-other WER (%)	
		w/ mask	w/o mask
10 layers	G-CTC	19.7	19.7
	SKD	19.9	<b>19.4</b>
8 layers	G-CTC	33.8	34.3
	SKD	33.3	<b>33.0</b>
6 layers	G-CTC	53.1	53.5
	SKD	51.7	<b>51.1</b>

(Sftmx-KD), one layer pruning method using an intermediate CTC (Layer Prun), and the proposed self-knowledge distillation method (SKD). It is observed that the KD-based methods and the layer pruning method improve performance over the baselines for most of the cases. However, the proposed SKD method outperforms the other methods except for a single case: the 18-layer HuBERT teacher with the 12-layer student. For instance, the WavLM 8-layer model achieves 20.7% WER for test-clean dataset, and the layer pruning achieves 14.3%, surpassing the 18.5% of G-CTC and Sftmx-KD. SKD achieves 13.4%, which is the lowest WER.

We analyze the effect of blank frame masking which have been regarded as an important component of KD. Table 2 shows that the conventional KD (Guide-CTC) with blank frame masking contributes to a performance improvement compared to the Guide-CTC without the masking function. However, we observe that distilling the entire frame without masking for SKD leads to an even lower WER. This supports our argument that distilling blank frames also contributes to an increased performance when the alignment disagreement of the teacher and student is not significant, as depicted in Figure 1 (c) for which the alignment disagreement is not as severe as in Figure 1 (a). Furthermore, we examine the frame-level teacher-student output accuracy with and without the masking, as shown in Table 3. Overall, SKD achieves a similar or higher frame-level output accuracy between the teacher model and the student model than Guide-CTC, which accounts for the improved performance of

Table 3: Accuracy (%) of frame-level outputs between the teacher model and the student model, measured under two settings: total frames and active frames. The former considers all frames, while the latter excludes blank frames. The experiments are conducted based on HuBERT models.

# layers	Method	test-clean (%)		test-other (%)	
		active	total	active	total
18 → 12	G-CTC w/ m	96.62	76.58	96.13	77.53
	G-CTC w/o m	95.21	94.66	93.13	93.49
	SKD w/ m	<b>99.13</b>	88.89	<b>98.58</b>	86.15
	SKD w/o m	98.83	<b>98.31</b>	98.30	<b>97.77</b>
12 → 8	G-CTC w/ m	92.15	78.43	86.27	76.92
	G-CTC w/o m	88.33	90.35	86.59	<b>90.35</b>
	SKD w/ m	<b>93.72</b>	88.43	<b>89.14</b>	85.94
	SKD w/o m	91.48	<b>92.69</b>	86.47	89.62

SKD. Guide-CTC with masking shows the worst total frame accuracy (total) for in all experiments. Guide-CTC without masking increases total accuracy but decreases the non-blank frames’ teacher-student output accuracy (active). Still, SKD without masking achieves higher active accuracy than Guide-CTC with masking while achieving the highest total accuracy except in a single setting.

## 5. Conclusion

In this paper, we introduce a self-knowledge distillation method for CTC-based ASR model. Experimental results confirm that the proposed method outperforms the conventional knowledge distillation methods and the layer pruning method across the various model settings. The proposed method is effective in that the proposed method only requires extra linear CTC head computation to the teacher model since the teacher model and the student model share the layers. Furthermore, we investigate the necessity of the blank frame masking for distillation, and observe that the blank frame masking diminishes the performance improvement of self-knowledge distillation. For future work, we consider adapting different teacher-student architectures to the proposed method.

## 6. Acknowledgements

This work was partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No. RS-2022-II220641, XVoice: Multi-Modal Voice Meta Learning, 1/2) and (No. RS-2022-II220320, Artificial intelligence research about cross-modal dialogue modeling for one-on-one multi-modal interactions, 1/2).

## 7. References

- [1] A. Graves, "Sequence transduction with recurrent neural networks," *arXiv preprint arXiv:1211.3711*, 2012.
- [2] A. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Sathesh, S. Sengupta, A. Coates *et al.*, "Deep speech: Scaling up end-to-end speech recognition," *arXiv preprint arXiv:1412.5567*, 2014.
- [3] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *2016 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2016, pp. 4960–4964.
- [4] K. Rao, H. Sak, and R. Prabhavalkar, "Exploring architectures, data and units for streaming end-to-end speech recognition with rnn-transducer," in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 193–199.
- [5] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [6] 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.
- [7] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Unsupervised pre-training for speech recognition," *arXiv preprint arXiv:1904.05862*, 2019.
- [8] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [9] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [10] A. Babu, C. Wang, A. Tjandra, K. Lakhota, Q. Xu, N. Goyal, K. Singh, P. von Platen, Y. Saraf, J. Pino, A. Baevski, A. Conneau, and M. Auli, "XLS-R: Self-supervised Cross-lingual Speech Representation Learning at Scale," in *Proc. Interspeech 2022*, 2022, pp. 2278–2282.
- [11] A. Baevski, W.-N. Hsu, Q. Xu, A. Babu, J. Gu, and M. Auli, "Data2vec: A general framework for self-supervised learning in speech, vision and language," in *International Conference on Machine Learning*. PMLR, 2022, pp. 1298–1312.
- [12] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [13] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," in *NeurIPS Deep Learning and Representation Learning Workshop*, 2015. [Online]. Available: <http://arxiv.org/abs/1503.02531>
- [14] L. Zhang, J. Song, A. Gao, J. Chen, C. Bao, and K. Ma, "Be your own teacher: Improve the performance of convolutional neural networks via self distillation," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 3713–3722.
- [15] T.-B. Xu and C.-L. Liu, "Data-distortion guided self-distillation for deep neural networks," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, 2019, pp. 5565–5572.
- [16] M. Ji, S. Shin, S. Hwang, G. Park, and I.-C. Moon, "Refine myself by teaching myself: Feature refinement via self-knowledge distillation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 10 664–10 673.
- [17] H. Sak, A. Senior, K. Rao, O. Irsoy, A. Graves, F. Beaufays, and J. Schalkwyk, "Learning acoustic frame labeling for speech recognition with recurrent neural networks," in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 4280–4284.
- [18] G. Kurata and K. Audhkhasi, "Guiding CTC Posterior Spike Timings for Improved Posterior Fusion and Knowledge Distillation," in *Proc. INTERSPEECH 2019*, 2019, pp. 1616–1620.
- [19] J. Tian, B. Yan, J. Yu, C. WENG, D. Yu, and S. Watanabe, "BAYES RISK CTC: CONTROLLABLE CTC ALIGNMENT IN SEQUENCE-TO-SEQUENCE TASKS," in *The Eleventh International Conference on Learning Representations*, 2023.
- [20] R. Takashima, S. Li, and H. Kawai, "An investigation of a knowledge distillation method for ctc acoustic models," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5809–5813.
- [21] J. W. Yoon, H. Lee, H. Y. Kim, W. I. Cho, and N. S. Kim, "Tutornet: Towards flexible knowledge distillation for end-to-end speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1626–1638, 2021.
- [22] R. Takashima, L. Sheng, and H. Kawai, "Investigation of sequence-level knowledge distillation methods for ctc acoustic models," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6156–6160.
- [23] S. Tian, K. Deng, Z. Li, L. Ye, G. Cheng, T. Li, and Y. Yan, "Knowledge distillation for ctc-based speech recognition via consistent acoustic representation learning," in *Proc. INTERSPEECH 2022*, 2022, pp. 2633–2637.
- [24] J. Lee and S. Watanabe, "Intermediate loss regularization for ctc-based speech recognition," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6224–6228.
- [25] A. Senior, H. Sak, F. de Chaumont Quiry, T. Sainath, and K. Rao, "Acoustic modelling with cd-ctc-smbc lstm rnns," in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. IEEE, 2015, pp. 604–609.
- [26] Z. Chen, Y. Zhuang, Y. Qian, and K. Yu, "Phone synchronous speech recognition with ctc lattices," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 1, pp. 90–101, 2016.
- [27] M. Jung, O. Kwon, S. Seo, and S. Seo, "Blank Collapse: Compressing CTC Emission for the Faster Decoding," in *Proc. INTERSPEECH 2023*, 2023, pp. 1349–1353.
- [28] P. Fan, C. Shan, S. Sun, Q. Yang, and J. Zhang, "Key frame mechanism for efficient conformer based end-to-end speech recognition," *IEEE Signal Processing Letters*, 2023.
- [29] J. W. Yoon, B. J. Woo, S. Ahn, H. Lee, and N. S. Kim, "Interkd: Intermediate knowledge distillation for ctc-based automatic speech recognition," in *2022 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2023, pp. 280–286.
- [30] J. Lee, J. Kang, and S. Watanabe, "Layer pruning on demand with intermediate ctc," in *Proc. INTERSPEECH 2021*, 2021, pp. 1381–1385.
- [31] 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.
- [32] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *International Conference on Learning Representations*, 2019.