Alignment Knowledge Distillation for Online Streaming Attention-Based Speech Recognition

Hirofumi Inaguma^(D), Member, IEEE, and Tatsuya Kawahara^(D), Fellow, IEEE

Abstract—This article describes an efficient training method for online streaming attention-based encoder-decoder (AED) automatic speech recognition (ASR) systems. AED models have achieved competitive performance in offline scenarios by jointly optimizing all components. They have recently been extended to an online streaming framework via models such as monotonie chunkwise attention (MoChA). However, the elaborate attention calculation process is not robust against long-form speech utterances. Moreover, the sequence-level training objective and time-restricted streaming encoder cause a nonnegligible delay in token emission during inference. To address these problems, we propose CTC synchronous training (CTC-ST), in which CTC alignments are leveraged as a reference for token boundaries to enable a MoChA model to learn optimal monotonie input-output alignments. We formulate a purely end-to-end training objective to synchronize the boundaries of MoChA to those of CTC. The CTC model shares an encoder with the MoChA model to enhance the encoder representation. Moreover, the proposed method provides alignment information learned in the CTC branch to the attention-based decoder. Therefore, CTC-ST can be regarded as self-distillation of alignment knowledge from CTC to MoChA. Experimental evaluations on a variety of benchmark datasets show that the proposed method significantly reduces recognition errors and emission latency simultaneously. The robustness to long-form and noisy speech is also demonstrated. We compare CTC-ST with several methods that distill alignment knowledge from a hybrid ASR system and show that the CTC-ST can achieve a comparable tradeoff of accuracy and latency without relying on external alignment information.

Index Terms—Attention-based encoder-decoder, connectionist temporal classification, knowledge distillation, monotonic chunkwise attention, streaming automatic speech recognition.

I. INTRODUCTION

O NLINE streaming automatic speech recognition (ASR) is a core technology for speech applications such as live captioning, simultaneous translation, voice search, and dialogue systems. The traditional but still dominant approach in production is a hybrid system that modularizes the entire system into an acoustic model, a pronunciation model, and a language model (LM). Recently, end-to-end (E2E) systems have achieved

The authors are with the Graduate School of Informatics, Kyoto University, Kyoto 606-8501, Japan (e-mail: inaguma@sap.ist.i.kyoto-u.ac.jp; kawahara@i.kyoto-u.ac.jp).

Digital Object Identifier 10.1109/TASLP.2021.3133217

comparable performance to that of hybrid systems by optimizing a direct mapping function from the input speech to the target transcription [1]–[3]. Representative approaches include the connectionist temporal classification (CTC) [4], recurrent neural network transducer (RNN-T) [5], recurrent neural aligner (RNA) [6], hybrid autoregressive transducer (HAT) [7], and attention-based encoder-decoder (AED) [8], [9] models. With the simplified architecture, the E2E approaches are advantageous for rapid system development, on-device applications with a small footprint, and fast inference when a large amount of training data is available.

The E2E models have been compared in offline scenarios [10]–[12], and AED models are typically the best choice because of the strong token dependency on the decoder side. In online streaming scenarios, however, AED models are not suitable because they require the entire input in order to generate the initial token. On the other hand, *frame-synchronous* models such as CTC and RNN-T can easily be extended to the streaming setting. RNN-T has been a practical choice because of its better performance than CTC with the help of token dependency modeling in the prediction network [3], [13], [14]. However, it is known that RNN-T consumes significant memory during training [15], [16] and requires a large search space during inference because of its frame-wise prediction, which significantly slows down the decoding speed.

To make AED models streamable, various methods have been proposed: neural transducer (NT) [17], triggered attention [18], adaptive computation steps (ACS) [19], continuous integrateand-fire (CIF) [20], streaming chunk-aware multi-head attention (SCAMA) [21], local windowing [22]–[25], Gaussian mixture model (GMM) attention [26], neural autoregressive transducer (NAT) [27], and monotonic chunkwise attention (MoChA) [28]. The major difference among these variants is the location where the input speech is segmented: at the encoder [17]–[21] or the decoder [22]–[28]). Among these methods, we focus on MoChA because it can generate tokens with linear-time complexity at test time and can be trained as efficiently as offline AED models [29]–[33]. Moreover, it can detect token boundaries by using contextual information captured in the decoder.

The main problem in making streaming AED models for practical systems is that they are not robust against long-form speech, which is not as problematic in frame-synchronous models [34]. Moreover, latency in the decision boundaries on token emission occurs in any E2E model [16], [32], [35]. This is because (1) the model is typically equipped with a time-restricted streaming

Manuscript received 18 February 2021; revised 21 August 2021; accepted 15 November 2021. Date of publication 7 December 2021; date of current version 7 April 2023. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Ozlem Kalinli. (*Corresponding author: Hirofumi Inaguma.*)

encoder having limited future contexts, and (2) the model is optimized with an end-to-end objective, which encourages the decoder to use as many future observations as possible.

Previous studies on frame-synchronous models tackled this problem by shifting output frames [36] and leveraging framewise alignment supervision [37]–[40]. As for label-synchronous models, Inaguma *et al.* leveraged alignment information extracted from a hybrid system to reduce the emission latency of MoChA [32]. However, this approach still depends on a hybrid model and is not a purely end-to-end solution. While the delayed token generation problem occurs similarly in framesynchronous models, CTC models are better than AED models in terms of latency, because they are optimized with the forwardbackward algorithm and assume conditional independence on a per-frame basis. Moreover, the peaky alignments learned in CTC are expected to be compatible with the token boundaries in MoChA.

In this article, we propose a novel purely end-to-end training method to enhance the alignment learning process of streaming MoChA models without external alignments.¹ We regard the peaks in CTC alignments as a reference for the tokens boundaries in MoChA; thus, we train a MoChA model to mimic a CTC model in order to detect token boundaries at similar positions. We refer to this method as CTC synchronous training (CTC-ST) [41]. This boundary supervision greatly eases the optimization of MoChA, especially with contaminated inputs as in SpecAugment [42]. This is because the accumulation of alignment errors in MoChA can be recovered with the help of the CTC alignments. The CTC model is jointly optimized with the MoChA model by having them share an encoder to encourage monotonic alignments in the MoChA decoder, similarly to the joint CTC/Attention framework [43]. In the proposed method, however, the CTC alignments are further provided to the MoChA decoder as supervision of token boundaries to restrict their positions. Because the alignment knowledge learned in CTC is transferred to improve MoChA's alignment in a unified architecture, we regard this framework as a form of self-distillation.²

Experimental evaluations on four benchmark datasets show that CTC-ST significantly improves the recognition accuracy, especially for long-form and noisy speech. We also demonstrate that CTC-ST can reduce the emission latency without external alignment information and achieve a tradeoff of the accuracy and emission latency comparable to that of alignment knowledge distillation from a hybrid system [32]. Finally, we compare MoChA with RNN-T in recognition accuracy and emission latency to demonstrate that CTC-ST can close the performance gap, so that the best MoChA system achieves recognition accuracy comparable to that of RNN-T and lower emission latency.³

²Unlike conventional knowledge distillation [44], we focus on token boundary positions instead of probability distributions.

II. RELATED WORK

A. Streaming Attention-Based Encoder-Decoder Model

We categorize streaming AED models into two groups in terms of how they segment speech frames for token generation.

1) Segmentation on Encoder Side: The first method in this category is NT [17], [45], [46], which performs labelsynchronous decoding on every fixed-size input block and moves to the next block if it detects no additional token boundaries. ACS [19] extends the idea to an adaptive segmentation policy based on a halting mechanism [47]. CIF [20] further enhances ACS by applying a fine-grained segmentation within the encoder output. SCAMA [21] learns to count the number of tokens to be generated in each fixed input chunk. Sterpu et al. followed a similar idea [48]. Triggered attention [18], [49], [50] is based on the joint CTC/Attention framework and truncates encoder outputs with CTC spikes to perform global attention over the past encoder outputs from each spike position, but the decoding complexity is quadratic. Our work is different in that we leverage CTC alignments only during training, and the decoding complexity is linear. A scout network [51] learns to detect word boundaries with alignment information from a hybrid system during training to reduce the emission latency, and it performs global attention similarly to triggered attention. However, it introduces a dependency on frame-wise alignment supervision. These methods can detect token boundaries regardless of the input length, but they are limited by not using contextual information for segmentation.

2) Segmentation on Decoder Side: The methods in this category leverage decoder states as a query to segment the input speech for every token. Local windowing methods were proposed first [22]-[25]. GMM attention [26], [34] forces the center of attention to move monotonically to the end of the encoder output. Kong et al. further incorporated source-side information [52]. NAT [27] trains stochastic variables with a policy gradient, and the optimization was further improved in [53]. Hard monotonic attention (HMA) [54] also introduces stochastic variables to detect token boundaries, but the model can be trained efficiently with a cross-entropy objective. MoChA [28] relieves the strong monotonic constraint in HMA by introducing additional soft attention over a small window. Miao et al. proposed stable MoChA (sMoChA) by simplifying the attention calculation in MoChA to ease the optimization [30]. They further proposed monotonic truncated attention (MTA) [33] by using soft attention scores in HMA at test time as well to reduce the gap between training and testing. HMA was further extended to the Transformer decoder [55], [56], while Li et al. extended the idea of ACS to the Transformer decoder [57]. Incremental decoding uses offline models for streaming applications, but the decoding complexity is quadratic [58]–[61].

B. Emission Latency in E2E ASR Model

Emission latency inevitably occurs in any E2E ASR model because sequence-level optimization allows the model to use as much future information as possible. This problem was tackled in CTC acoustic models for the first time by constraining CTC

¹Although we focus on MoChA in this work, the proposed method can be applied to any AED model that calculates attention scores.

³This work is an extension of our previous studies [41]. In this article, we add more evaluations regarding various datasets, the emission latency, and the inference speed. Also new are the comparisons with alignment knowledge distillation from a hybrid system [32] and RNN-T.

paths during marginalization with frame-level alignments from a hybrid system [62]. Zhang et al. trained a CTC model jointly with a frame-level cross-entropy [38]. Similar methods have also been investigated for RNN-T by pretraining the encoder with a frame-level cross-entropy [39] and applying joint training [40]. Inaguma et al. investigated the application of alignment information from a hybrid system to MoChA and reduced the recognition errors and emission latency simultaneously [32]. Sainath et al. constrained an attention head of the second-pass rescorer in a two-pass streaming E2E model by alignment information from a hybrid system to improve accuracy of word emission timing [63]. However, such alignment information is not necessarily available. Recently, FastEmit [64] was proposed by designing a new training objective for RNN-T to reduce the emission latency without any frame-level supervision, which was applied to two-pass E2E architectures in the voice search task successfully [65], [66]. The idea of FastEmit was extended to MoChA in [67], referred to as StableEmit. Self-alignment from RNN-T was also leveraged in [68]. Yu et al. trained a single RNN-T in both offline and streaming modes (dual-mode ASR) by sharing parameters [69].

C. Knowledge Distillation for Streaming ASR

Knowledge distillation [44] has also been investigated to improve the performance of streaming E2E ASR models. Previous studies focused on distilling knowledge from an offline or streaming teacher model to a weaker streaming student model within the same decoder topology [70]-[77]. In framesynchronous models, however, there exists a problem that the timing to emit tokens can differ between the teacher and student models, depending on the future context size in the encoder. Through an approach using bidirectional long short-term memory (BLSTM), Kurata et al. trained BLSTM-CTC to mimic LSTM-CTC in order to generate posterior spikes at similar positions and then distill the frame-level posterior probabilities to the LSTM-CTC model [78]. The idea was extended to RNN-T in [76]. Ding et al. jointly trained multiple teacher CTC models to synchronize their posterior spikes [79]. Distillation between different decoder topologies has also been investigated. Moriya et al. distilled knowledge from a teacher AED model to a student CTC model [80]. Self-distillation in a single E2E ASR model has also been proposed as an in-place operation, from an offline mode to a streaming mode [61], [69], and from a Transformer decoder to a CTC layer [81]. Unlike those previous methods, we focus on distilling the positions of token boundaries learned in a CTC model to an AED model, rather than distilling the posterior distributions. Moreover, the teacher and student models share the same encoder and are trained jointly from scratch.

III. BASICS

In this section, we review HMA and MoChA. Let $\boldsymbol{x} = (x_1, \ldots, x_T)$ be an input speech sequence, $\boldsymbol{h} = (h_1, \ldots, h_{T'})$ be encoder outputs $(T' \leq T)$, and $\boldsymbol{y} = (y_1, \ldots, y_U)$ be the corresponding output token sequence. The encoder performs downsampling to reduce the input sequence length from T to T'.

We use i and j as the time indices of the output and input sequences, respectively.

A. Hard Monotonic Attention (HMA)

Standard offline AED models are based on the global attention mechanism [8], in which relevant source information is selected according to the target context via attention scores calculated by normalizing energy activations over h. However, this prevents the model from performing online streaming recognition, because the decoder must see all the encoder outputs to generate the initial token. Moreover, the decoding complexity at each generation step is in proportion to the encoder output length T'. This results in a total decoding complexity $\mathcal{O}(T'U)$. To start generating tokens when given partial acoustic observations during inference, HMA introduces discrete binary decision processes. As a result, it can perform decoding with linear-time complexity $\mathcal{O}(T')$ during inference, but it behaves differently between the training and test times.

At test time, the decoder scans the encoder outputs $h_1, \ldots, h_{T'}$ from left to right. At every input frame index j, the decoder has the option to (1) stop at the current frame j to generate a token or (2) move forward to the next frame j + 1 according to a selection probability $p_{i,j} \in [0, 1]$. A discrete decision $z_{i,j} \in \{0, 1\}$ on whether to stop at the j-th frame is sampled from a Bernoulli random variable parameterized by $p_{i,j}$ as

$$e_{i,j} =$$
MonotonicEnergy $(h_j, s_i),$ (1)

$$p_{i,j} = \sigma(e_{i,j}),$$

$$z_{i,j} \sim \text{Bernoulli}(p_{i,j}),$$
(2)

where $e_{i,j}$ is a monotonic energy activation, s_i is the *i*-th decoder state⁴, and σ is a logistic sigmoid function. When $z_{i,j} = 1$, *i.e.*, $p_{i,j} \ge 0.5$, the decoder stops at an index $j = t_i$ (referred to as the *token boundary* of the *i*-th token). Then, only the corresponding single encoder output h_{t_i} is used for generating the next token and updating the decoder state. The next token boundary t_{i+1} is determined by resuming scanning from the previous token boundary $j = t_i$.

However, this hard assignment of $z_{i,j}$ is not differentiable. To perform the standard backpropagation training, the expected alignment scores $\alpha_{i,j}$ are calculated by marginalizing over all possible alignment paths as follows:

$$\alpha_{i,j} = p_{i,j} \sum_{k=1}^{j} \left(\alpha_{i-1,k} \prod_{l=k}^{j-1} (1-p_{i,l}) \right)$$
$$= p_{i,j} \left((1-p_{i,j-1}) \frac{\alpha_{i,j-1}}{p_{i,j-1}} + \alpha_{i-1,j} \right).$$
(3)

Because $\alpha_{i,j}$ in (3) introduces a recurrence relation, it is difficult to calculate it in parallel over the input indices. However, by substituting $q_{i,j} = \alpha_{i,j}/p_{i,j}$, it can be calculated efficiently with the cumulative sum and product operations, denoted respectively as cumsum and cumprod, as follows:

$$\boldsymbol{\alpha}_{i,:} = \boldsymbol{p}_i \cdot \texttt{cumprod}(1 - \boldsymbol{p}_{i,:}) \cdot \texttt{cumsum} \bigg(\frac{\boldsymbol{\alpha}_{i-1,:}}{\texttt{cumprod}(1 - \boldsymbol{p}_{i,:})} \bigg).$$

Finally, the monotonic energy activation $e_{i,j}$ in (1) is implemented as

MonotonicEnergy
$$(h_j, s_i) = g \frac{v^{\mathsf{I}}}{||v||} f(\mathbf{W}_{\mathrm{h}} h_j + \mathbf{W}_{\mathrm{s}} s_i + b) + r$$

where f is the nonlinear activation, g and v are parameters for weight normalization, and $\mathbf{W}_{\rm h}$, $\mathbf{W}_{\rm s}$, b, and r are trainable parameters. We use a rectified linear unit (ReLU) activation function as f. Following [54], the scalar offset parameter r is initialized as -4 in this work. To ensure the discreteness of $p_{i,j}$, a zero-mean, unit-variance Gaussian noise is added to the pre-sigmoid activation in (2) during training. The subsequent token generation processes are the same as in the global AED model.

B. Monotonic Chunkwise Attention (MoChA)

In HMA, source information is restricted to a single encoder output, and this strong constraint greatly sacrifices the accuracy in general. To overcome this problem by leveraging the surrounding contexts, MoChA introduces an additional soft attention mechanism over a fixed window of width w on top of HMA.

At test time, soft attention scores $\beta_{i,j}$ are calculated over w encoder outputs from every token boundary t_i as follows:

$$\beta_{i,j} = \exp(u_{i,j}) / \sum_{l=t_i - w + 1}^{t_i} \exp(u_{i,l}),$$

$$c_i = \sum_{j=t_i - w + 1}^{t_i} \beta_{i,j} h_j,$$
(4)

where $u_{i,j}$ is the chunk energy activation formulated as in (1) without weight normalization and the offset parameter r, and c_i is a context vector for the *i*-th token.

During training, $\beta_{i,j}$ can be calculated on top of the expected alignment score $\alpha_{i,j}$ as

$$\beta_{i,j} = \sum_{k=j}^{j+w-1} \left(\alpha_{i,k} \exp(u_{i,j}) / \sum_{l=k-w+1}^{k} \exp(u_{i,l}) \right).$$
(5)

The computation of $\beta_{i,j}$ is expensive because of the nested summation. Fortunately, it can be computed more efficiently with a moving sum operation, denoted as MovSum, as follows:

$$\begin{split} \beta_{i,j} &= \exp(u_{i,:}) \cdot \operatorname{MovSum}(\frac{\alpha_{i,:}}{\operatorname{MovSum}(\exp(u_{i,:},w,1))},1,w), \\ & \operatorname{MovSum}(\pmb{x},b,f)_n = \sum_{m=n-(b-1)}^{n+f-1} x_m. \end{split}$$

This computation can be implemented by a 1-dimensional convolution. c_i is calculated as $c_i = \sum_{j=1}^{T'} \beta_{i,j} h_j$. The objective function is formulated as the negative log-likelihood $\mathcal{L}_{\text{mocha}} = -\log P_{\text{mocha}}(\boldsymbol{y}|\boldsymbol{x})$, where P_{mocha} is the output probability distribution of MoChA.

To encourage MoChA to learn monotonic alignments, we train it jointly with an auxiliary CTC objective $\mathcal{L}_{ctc} = -\log P_{ctc}(\boldsymbol{y}|\boldsymbol{x})$, where P_{ctc} is the CTC output probability distribution, by sharing the encoder sub-network [43], [82]. Moreover, to avoid vanishing of $\alpha_{i,j}$, quantity loss \mathcal{L}_{qua} is introduced by making the expected total number of token boundaries closer to the reference output sequence length U as follows [20], [32]:

$$\mathcal{L}_{\text{qua}} = |U - \sum_{i=1}^{U} \sum_{j=1}^{T'} \alpha_{i,j}|.$$
 (6)

We refer to this technique as *quantity regularization* (QR).⁵ Note that any external alignment information is not used here. The total objective function \mathcal{L}_{total} is defined as a linear interpolation of \mathcal{L}_{mocha} , \mathcal{L}_{ctc} , and \mathcal{L}_{qua} as follows:

$$\mathcal{L}_{\text{total}} = (1 - \lambda_{\text{ctc}})\mathcal{L}_{\text{mocha}} + \lambda_{\text{ctc}}\mathcal{L}_{\text{ctc}} + \lambda_{\text{qua}}\mathcal{L}_{\text{qua}}, \quad (7)$$

where λ_{ctc} $(0 \le \lambda_{ctc} \le 1)$ and λ_{qua} (≥ 0) are tunable weights for the CTC loss and quantity loss, respectively.

IV. PROBLEMS IN MOCHA

In this section, we review two major problems in MoChA: vanishing alignment probabilities and delayed token generation.

A. Vanishing Alignment Probabilities

Conventional alignment models such as the hidden Markov model (HMM) [83] and CTC models use the forward-backward algorithm to calculate alignment probabilities. However, it is not straightforward to apply that algorithm to the HMA mechanism. This is because HMA does not normalize the monotonic energy $e_{i,j}$ across the entire set of encoder outputs h to obtain the expected alignment score $\alpha_{i,j}$, which means that $\alpha_{i,j}$ is not a valid probability. Moreover, the decoder is an autoregressive model, and an incremental left-to-right update of the decoder state is required for each token. Therefore, during the marginalization process at training time, $\alpha_{i,j}$ depends only on past alignments, as can be seen in (3), and it can quickly be attenuated as the number of decoding steps increases [33], [84], because $\sum_{j=1}^{T'} \alpha_{i,j} < 1$. This is especially problematic for long-form speech because the model is more likely to fail to learn the scale of $p_{i,j}$ properly in the latter steps. Accordingly, the gap in the HMA behaviors between training and testing is widened. This leads to premature endpointing, which increases deletion errors [85], [86].

B. Delayed Token Generation

To enable online inference, a streaming ASR model needs to be equipped with a time-restricted encoder, which does not use enough future information. However, because E2E models are optimized with sequence-level criteria, *i.e.*, the cross-entropy, future encoder outputs are used as much as possible [29], [31], [32], [37]. As a result, token boundaries are shifted several

⁵QR was investigated for DeCoT in [32] for the voice search task, and it did not bring any improvement for a naive MoChA model. However, we found that it gave large improvements for the naive model in our experiments. We attribute this to the longer input sequence length distributions and the smaller training data sizes of the corpora used in this work. Therefore, unless otherwise noted, we use QR for the baseline MoChA models in this work.



Fig. 1. Visualization of the token boundaries (yellow dots) of the baseline UniLSTM - MoChA model (top, *T5* in Table II) and the proposed model with CTC-ST (bottom, *T6*). Reference: "we might be putting lids and casting shadows on their power wouldn't we want to open doors for them instead."

frames ahead from the actual acoustic boundary positions. Moreover, MoChA allows emission of multiple tokens at the same input index. These problems inevitably cause a large perceived latency and make the model unusable in online ASR.

Fig. 1 (top) shows an example of token boundaries emitted from a baseline MoChA model trained with an auxiliary CTC loss. The yellow dots and the spikes below represent the token boundaries from MoChA and the CTC posterior probabilities, respectively. It is known that the posterior probabilities in a well-trained CTC model tend to peak in sharp spikes [4]. In the figure, we can observe that MoChA's token boundaries are indeed shifted to the right (future side) from the actual acoustic boundaries and are poorly aligned to the CTC spikes.

V. ALIGNMENT KNOWLEDGE DISTILLATION FROM HYBRID ASR System

This section describes previous approaches to tackle the delayed token generation problem in MoChA by leveraging word alignments extracted from a hybrid ASR system [32]. We review two methods in [32]: **De**lay **Constrained Training** (DeCoT) and **Min**imum Latency Training (MinLT), which can reduce both the emission latency and recognition errors. Because the alignment knowledge bootstrapped in a hybrid system is transferred to an AED model, we refer to the procedure as *alignment knowledge distillation* from the hybrid system.

Let $A = (a_1, \ldots, a_T)$ $(a_j: \mathcal{V}$ -dimensional one-hot vector; \mathcal{V} : vocabulary size of the hybrid system) be a frame-level word alignment corresponding to the input sequence x, and let $\mathbf{b}^{\text{ref}} = (\mathbf{b}_1^{\text{ref}}, \ldots, \mathbf{b}_U^{\text{ref}})$ be a sequence of endpoints of token boundaries for a reference transcription $\mathbf{y} = (y_1, \ldots, y_U)$. To convert the word alignment to a subword alignment compatible with MoChA, we divide the total time duration of each word in A by the ratio of the character length of each subword. Finally, we select the end timestamp as the token boundary.

A. Delay Constrained Training (DeCoT)

In DeCoT, inappropriate alignment paths that poorly match the reference alignment are removed by masking out their scores $\alpha_{i,j}$ during marginalization. Then, $\alpha_{i,j}$ in (3) is reformulated to $\alpha_{i,j}^{\text{decot}}$ as follows:

$$\alpha_{i,j}^{\text{decot}} = \begin{cases} p_{i,j} \left((1 - p_{i,j-1}) \frac{\alpha_{i,j-1}}{p_{i,j-1}} + \alpha_{i-1,j} \right) & (j \le \mathbf{b}_i^{\text{ref}} + \delta_{\text{decot}}) \\ 0 & (\text{otherwise}), \end{cases}$$

where δ_{decot} is a hyperparameter to control the acceptable delay, and b_i^{ref} is a reference boundary of the *i*-th token transferred from the hybrid system.

As explained in Section IV-A, MoChA has a problem of an exponential decay of $\alpha_{i,j}$, and the masking for $\alpha_{i,j}^{\text{decot}}$ accelerates it further. To recover the proper scale of $\alpha_{i,j}^{\text{decot}}$, \mathcal{L}_{qua} in (6) is calculated with $\alpha_{i,j}^{\text{decot}}$ instead of $\alpha_{i,j}$ as

$$\mathcal{L}_{\text{qua}}^{\text{decot}} = |U - \sum_{i=1}^{U} \sum_{j=1}^{T'} \alpha_{i,j}^{\text{decot}}|.$$
(8)

Intuitively, QR in DeCoT emphasizes the valid alignment paths during marginalization. Moreover, this also leads to better estimation of $\beta_{i,j}$ in (5). Accordingly, the total objective function in (7) is modified as follows:

$$\mathcal{L}_{\text{total}} = (1 - \lambda_{\text{ctc}})\mathcal{L}_{\text{mocha}} + \lambda_{\text{ctc}}\mathcal{L}_{\text{ctc}} + \lambda_{\text{qua}}\mathcal{L}_{\text{qua}}^{\text{decot}}.$$
 (9)

Unlike in [32], we also add a CTC loss as an auxiliary objective for DeCoT.

B. Minimum Latency Training (MinLT)

While DeCoT can effectively reduce the emission latency, a fixed buffer size δ_{decot} must be predefined for every token. However, the emission latency can vary depending on the speaking rate, subwords, and so on. To reduce the latency of each token more flexibly, MinLT directly minimizes the expected latency so that the expected token boundaries in MoChA get closer to the corresponding reference boundaries. A differentiable expected latency loss \mathcal{L}_{minlt} is specified as

$$\mathbf{b}_{i}^{\text{mocha}} = \sum_{j=1}^{T'} j \cdot \alpha_{i,j},\tag{10}$$

$$\mathcal{L}_{\text{minlt}} = \frac{1}{U} \sum_{i=1}^{U} |\mathbf{b}_i^{\text{ref}} - \mathbf{b}_i^{\text{mocha}}|, \qquad (11)$$

where b_i^{mocha} is the expected boundary position in MoChA for the *i*-th token during training. The total objective function in (7) is modified as follows:

$$\mathcal{L}_{\text{total}} = (1 - \lambda_{\text{ctc}})\mathcal{L}_{\text{mocha}} + \lambda_{\text{ctc}}\mathcal{L}_{\text{ctc}} + \lambda_{\text{minlt}}\mathcal{L}_{\text{minlt}}, \quad (12)$$

where $\lambda_{\text{minlt}} (\geq 0)$ is a hyperparameter to control the latency.

VI. PROPOSED METHOD: CTC SYNCHRONOUS TRAINING

In this section, we propose a novel training method to alleviate the problems of vanishing alignment probabilities and delayed token generation in MoChA by leveraging CTC alignments as a reference for token boundaries.



Fig. 2. Overview of CTC synchronous training (CTC-ST).

A. Overview

As observed in Fig. 1, there exists a gap in the timing to emit tokens between the CTC and MoChA models even when they share the same encoder. However, we can see that the CTC spikes are closer to the actual acoustic boundaries. The reason is that CTC does not suffer from alignment error propagation because of the optimization with the forward-backward algorithm and the assumption of conditional independence on a per-frame basis. Therefore, we expect that CTC can generate more reliable alignments than MoChA and serve as an effective guide for MoChA to learn to detect token boundaries more accurately.

Motivated by this reasoning, we propose CTC synchronous training (CTC-ST), in which a MoChA model is trained to mimic a CTC model in order to generate token boundaries at similar positions. Fig. 2 shows a system overview. Both the MoChA and CTC branches are jointly optimized by sharing the encoder, and reference token boundaries are obtained from the CTC alignments generated from the CTC branch. Therefore, the CTC model serves not only as a regularizer to enhance the encoder representations but also as a teacher alignment model to estimate accurate token boundary positions. In this sense, we regard CTC-ST as a form of self-distillation from CTC to MoChA. Specifically, the synchronization of token boundaries in CTC-ST can be viewed as explicit interaction between MoChA and CTC models on the *decoder* side, unlike in the conventional joint CTC/Attention framework [43]. However, we only leverage the discrete token boundary positions rather than the probability distributions as in the conventional knowledge distillation method [44].

Because CTC is not allowed to emit multiple symbols at the same input index, the reference token boundaries from the CTC alignments can also enforce the monotonicity of $\alpha_{i,j}$ in MoChA, which leads to the emission latency reduction as well. Moreover, unlike the methods described in Section V, the entire model can be trained in an end-to-end manner without relying on external alignments extracted from a hybrid system or manual annotation.

B. Extraction of CTC Alignments

We use the most plausible CTC path $\bar{\pi} = \operatorname{argmax}_{\pi} p(\pi|x)$ $(|\bar{\pi}| = T')$, given by Viterbi alignment, via forced alignment with the forward-backward algorithm in a manner similar to triggered attention [18]. The time indices of non-blank tokens in $\bar{\pi}$ are used as the reference token boundaries $\mathbf{b}^{\text{ctc}} =$ $(\mathbf{b}_1^{\text{ctc}}, \ldots, \mathbf{b}_U^{\text{ctc}})$. When repeated non-blank labels exist, the leftmost index corresponding to the same non-blank token is used as a reference token boundary. The last time index T' is used for the end-of-sentence (EOS) mark, (eos). For instance, given a CTC path $\bar{\pi} = [\phi, c, c, \phi, \phi, a, a, a, \phi, t, t, \phi]$ (ϕ : blank) corresponding to a reference transcription "c a t $\langle eos \rangle$," we convert it to $[\phi, c, \phi, \phi, \phi, a, \phi, \phi, \phi, t, \phi, \langle eos \rangle]$ and then extract the time indices of the non-blank tokens $\mathbf{b}^{ctc} = (2, 5, 9, 11)$ (1-indexed). Unless otherwise specified, \mathbf{b}^{ctc} is generated with the model parameters at each training step *on the fly* and is expected to get more accurate as the training continues. We can also pre-compute \mathbf{b}^{ctc} and use the fixed boundaries throughout training when adopting curriculum learning; this approach is discussed in Section XI-A.

C. Optimization

We define the objective function of CTC-ST $\mathcal{L}_{\rm sync}$ (hereafter, the CTC-ST loss) as

$$\mathcal{L}_{\text{sync}} = \frac{1}{U} \sum_{i=1}^{U} |\mathbf{b}_i^{\text{ctc}} - \mathbf{b}_i^{\text{mocha}}|.$$
(13)

The total objective function in (7) is reformulated as

$$\mathcal{L}_{\text{total}} = (1 - \lambda_{\text{ctc}})\mathcal{L}_{\text{mocha}} + \lambda_{\text{ctc}}\mathcal{L}_{\text{ctc}} + \lambda_{\text{sync}}\mathcal{L}_{\text{sync}}, \quad (14)$$

where λ_{sync} (≥ 0) is a tunable hyperparameter. When using CTC-ST, we do not use the quantity loss and CTC-ST loss simultaneously. This is because we found that the combination was not effective in our experiments, as described in Section VIII. Instead, we propose an effective curriculum learning strategy.

D. Curriculum Learning Strategy

To calculate the effective gradient via (13) in CTC-ST, it is necessary to estimate reasonable expected token boundary positions in MoChA during training. However, $\alpha_{i,j}$ tends to diffuse over several frames in the early training stage. Again, we note that $\alpha_{i,j}$ is not explicitly normalized to sum up to one. Moreover, CTC alignments would not be very accurate in the early training stage either. Therefore, applying CTC-ST from a random parameter initialization leads to unstable, slower convergence. To tackle this problem, we propose a simple but effective curriculum learning strategy composed of the following two stages.

1) Stage 1: We first train a MoChA model equipped with a bidirectional encoder (e.g., BLSTM) together with QR by applying (7) from scratch until convergence. As the bidirectional encoder can see the entire context, we refer to this model as "offline." In this stage, we expect the model to learn a proper scale of $\alpha_{i,j}$.

2) Stage 2: Next, we optimize a MoChA model equipped with a latency-controlled bidirectional encoder (*e.g.*, LC-BLSTM [87]) with CTC-ST by applying CTC-ST with (14). We initialize the parameters with values optimized in stage 1. Because the BLSTM and LC-BLSTM encoders have the same model structure, we can reuse all of the parameters; the only difference between them is the lookahead context size. The optimizer's parameters and the learning rate are reset at the beginning of stage 2. In this stage, we expect the model to learn accurate token boundary location. When using a unidirectional LSTM encoder, the same encoder is used in both stages. We also

Corpus	Hours	Domain	Language	BPE size
TEDLIUM2 [87]	210	Lecture	English	10k
CSJ [89]	980 586	Lecture	Japanese	10k
AMI (SDM) [90]	100	Meeting	English	500

TABLE I Corpus Statistics

apply this curriculum learning strategy to DeCoT and MinLT to stabilize the training [32].

E. Combination With SpecAugment

Recently, SpecAugment [42] has shown the capability to improve the performance of E2E ASR models greatly. SpecAugment is an on-the-fly data augmentation method that introduces stochastic time and frequency masks into input speech. M_T time masks, whose size is sampled from a uniform distribution $\mathcal{U}(0,T)$ are applied to the input log-mel spectrogram. Similarly, frequency masks are also applied with mask parameters M_F and F. However, such input masks easily collapse the recurrence of $\alpha_{i,j}$ in (3) right after the masked region. Although such a problem does not exist in offline global AED models or frame-synchronous models, it is a severe problem in MoChA. In our experiments, in fact, the performance of the baseline MoChA model was degraded by applying SpecAugment.

In contrast, CTC-ST can help MoChA recover the collapsed $\alpha_{i,j}$ in the masked region by leveraging the CTC spikes because CTC assumes conditional independence on a per-frame basis. Therefore, we expect that CTC-ST is beneficial for MoChA to learn monotonic alignments that are robust against noisy inputs.

VII. EXPERIMENTAL EVALUATION

A. Datasets

We used the TEDLIUM release 2 (TEDLIUM2) [88] and Librispeech [89], the Corpus of Spontaneous Japanese (CSJ) [90], and the single distant microphone (SDM) portion of the AMI Meeting Corpus [91] for our experimental evaluations. The corpus statistics and the utterance length distributions are presented in Table I and Fig. 3, respectively. We used 10 k vocabularies based on the byte pair encoding (BPE) algorithm [92] except for the AMI corpus, for which a vocabulary of 500 BPE units was used.

B. Experimental Configuration

Using the Kaldi toolkit [93], we extracted 80-channel logmel filterbank coefficients computed with a 25-ms window that was shifted every 10 ms. Input features were normalized by the global mean and variance calculated on each training set. We removed utterances longer than 16 seconds from the training data to conserve the GPU memory capacity. The training utterances were sorted by their input lengths in ascending order during the entire training stage. We applied 3-fold speed perturbation [94] to the TEDLIUM2 and AMI corpora with factors of 0.9, 1.0, and 1.1.



Fig. 3. Utterance length distributions in TEDLIUM2, Librispeech, CSJ, and AMI (SDM) corpora.

The encoders consisted of two CNN blocks followed by five layers of (LC-)BLSTM or unidirectional LSTM. Each CNN block consisted of two CNN layers, each of which had a 3×3 filter followed by a max-pooling layer with a stride of 2×2 . This resulted in a 4-fold frame rate reduction in total and introduced a 60 ms lookahead latency for every output of the CNN blocks. We set the number of units in each (LC-)BLSTM layer to 512 per direction. To reduce the input dimension of the subsequent (LC-)BLSTM layer, we summed the LSTM outputs in both directions at every layer [95]. For a unidirectional LSTM encoder, the unit size was increased to 1024. In this article, we denote an LC-BLSTM encoder with a hop size of $N_{\rm c}$ frames and a future context of N_r frames as "LC-BLSTM- N_c+N_r ." The decoder was a single layer of unidirectional LSTM with 1024-dimensional units. We set the window size w of chunkwise attention in MoChA to 4. Offline global AED models used the location-based attention [8].

We also trained RNN-T models with the same encoder for comparison. The RNN-T models had a two-layer LSTM prediction network with 1024 memory units and a joint network with 512 units. The 1 k BPE was used for the vocabulary except for 500 units on the AMI corpus. These vocabulary sizes were selected to achieve the best performance for RNN-T. We also used an auxiliary CTC loss for RNN-T as $\mathcal{L}_{total} = (1 - \lambda_{ctc})\mathcal{L}_{rnnt} + \lambda_{ctc}\mathcal{L}_{ctc}$, where \mathcal{L}_{rnnt} is a RNN-T loss. λ_{ctc} was set to 0.3.

The Adam optimizer [96] was used with an initial learning rate of 1e - 3, which was then decayed exponentially. For the TEDLIUM2 and AMI corpora, 4 k warmup steps were used. We applied dropout and label smoothing [97] with probabilities of 0.4 and 0.1, respectively. The weight of the quantity loss λ_{qua} in (7) was set to 2.0, 0.1, 1.0, and 1.0 for the TEDLIUM2, Librispeech, CSJ, and AMI corpora, respectively. For DeCoT, λ_{qua} in (9) was set to 2.0. The CTC-ST weight λ_{sync} in (14) was set to 1.0 for all corpora, unless otherwise noted. λ_{ctc} was set to 0.3 in all models. All training was performed with a single GPU.

TABLE II RESULTS WITHOUT SPECAUGMENT ON TEDLIUM2

Model Seed		WER [WER [%] (\downarrow)	
		dev	test	
Offline				
UniLSTM - Global AED		12.9	11.9	
BLSTM - Global AED (T1)		9.7	9.5	
BLSTM - RNN-T (T2)		9.4	9.2	
BLSTM - MoChA	-	13.8	12.6	
+ QR (I3)		10.8	9.8	
+ CTC-ST		11.0	10.2	
Streaming (UniLSTM encoder)				
RNN-T (T4)	_	12.8	13.0	
MoChA + QR (T5)	-	16.7	15.0	
MoChA + CTC-ST $(T6)$	Τ5	14.8	13.2	
Streaming (LC-BLSTM-40+20 encoder)				
RNN-T	T2	10.5	10.4	
MoChA + QR	ΤЗ	13.8	12.2	
MoChA + CTC-ST	ТЗ	12.0	10.5	
Streaming (LC-BLSTM-40+40 encoder)				
RNN-T	T2	10.0	10.0	
MoChA + QR (T7)	ΤЗ	13.4	11.3	
MoChA + CTC-ST (T8)	Τ3	11.5	9.9	

QR: Quantity Regularization. CTC-ST: CTC-Synchronous Training.

For inference, we used a 4-layer LSTM LM with 1024 units per layer. For AED models, we used a beam width of 10 and scores normalized by the output sequence length at every output timestep, except for global AED models on the AMI corpus.⁶ Joint CTC decoding was performed for global AED models [82]. For RNN-T models, we used the breadth-first time-synchronous decoding (TSD) algorithm with a monotonic constraint (hereafter, *mono-TSD*) [98] to speed up decoding,⁷ and we merged paths corresponding to the same label history, except for the AMI corpus.⁸ We also reduced the beam width of streaming RNN-T models to 5 because of the inference speed constraint. Our implementation is publicly available.⁹

VIII. RESULTS

A. TEDLIUM2

1) Effectiveness of CTC-ST: Table II summarizes the results on the TEDLIUM2 corpus. In the offline scenario, the naive implementation of MoChA showed very poor performance. Quantity regularization (QR) drastically improved the performance (T3). Although CTC-ST also improved the performance, it was less effective than QR when it was applied from scratch. The BLSTM RNN-T outperformed the global AED model.

In the streaming scenario, however, our proposed CTC-ST significantly improved the baseline performance as compared to QR regardless of the encoder type. We obtained relative word error rate (WER) improvements of 12.0, 13.9, and 12.3% for the UniLSTM, LC-BLSTM-40+20, and LC-BLSTM-40+40 MoChA models, respectively, on the *test* set. Although using a larger number of future lookahead frames improved the WER



Fig. 4. WER bucketed by different input length on TEDLIUM2 test set.

TABLE III Results of Regularization Methods and Curriculum Learning Strategy on TEDLIUM2 With LC-BLSTM-40+40 Encoder Configuration

Initialization	Seed	\mathcal{L}_{aua}	\mathcal{L}_{sync}	WER [[%] (↓)
		quu	59 110	dev	test
Enome constals		1	X	13.8	12.3
FIOIII SCIAICII	_	×	\checkmark	12.6	10.9
		X	X	18.1	16.9
Curriculum	m D	1	X	13.4	11.3
learning	13	X	1	11.5	9.9
-		1	\checkmark	11.5	10.1

as expected, the effectiveness of CTC-ST was consistent. The token boundaries from the MoChA and CTC branches in T5 and T6 are visualized in Fig. 1. We can see that the MoChA boundaries moved to the left, and that the gap in the timing to emit tokens in both branches was reduced. This is desirable for reducing the perceived latency [32], which will be evaluated in Section X-B. Moreover, we found that the CTC spikes slightly shifted to the left. Although RNN-T was more robust against long-form utterances on the *dev* set, MoChA optimized with CTC-ST matched the performance of RNN-T on the *test* set.

2) Bucketing by Input Length: Fig. 4 shows a plot of the WER bucketed by the input length on the *test* set. The plot confirms that the largest gains by CTC-ST were for utterances longer than 20 seconds. The offline global AED model with joint CTC decoding (T1) did not have difficulty in recognizing long utterances, whereas the baseline streaming MoChA models did, regardless of the encoder type (T5, T7). As we had expected, the proposed CTC-ST successfully mitigated this problem (T6, T8).

3) Effectiveness of Curriculum Learning: Next, we investigate the effectiveness of the regularization methods (CTC-ST and QR) and the curriculum learning strategy by using LC-BLSTM-40+40 MoChA. Table III summarizes the results. Initialization with the offline model T3 was very helpful for both regularization methods, which is consistent with a previous study [31]. However, regularization with the CTC-ST loss or quantity loss in stage 2 was essential for achieving a performance gain even with curriculum learning. When using the LC-BLSTM encoder, CTC-ST was more effective than QR regardless of the use of curriculum learning, unlike the results listed in Table II for the BLSTM encoder. Combining both losses did not lead to any further improvement, although it was more effective than applying the quantity loss alone. Therefore, CTC-ST has an overlapping effect of encouraging MoChA to learn the scale of $\alpha_{i,i}$ properly.

⁶This was because the utterance lengths in the AMI corpus are relatively short. ⁷Although we did not explicitly use a monotonic RNN-T loss, the joint CTC training enforced a similar effect.

⁸Again, this was because the utterance lengths are relatively short.

⁹[Online]. Available: https://github.com/hirofumi0810/neural_sp

TABLE IV Results With SpecAugment on TEDLIUM2

Model	Seed	WER [%] (↓)
		dev	test
Offline			
Transformer [2]	_	9.3	8.1
Transformer + semantic mask [98]	_	-	7.7
BLSTM - Global AED [99]	_	10.3	8.8
BLSTM - Global AED	-	8.1	7.5
BLSTM - RNN-T	-	8.5	7.8
Streaming (UniLSTM)			
RNN-T	Τ4	11.6	11.7
MoChA + QR	Т5	17.7	16.0
MoChA + CTC-ST	Т5	12.9	11.5
Streaming (LC-BLSTM-40+40)			
RNN-T	Т2	8.9	8.5
MoChA + QR	ΤЗ	12.9	11.2
MoChA + CTC-ST	ТЗ	10.6	8.6

4) Combination With SpecAugment: We next investigate the combination of CTC-ST and SpecAugment, whose results are summarized in Table IV. We set (F_{sp}, T_{sp}) to (13, 50) for UniLSTM MoChA, (27, 50) for LC-BLSTM MoChA, and (27, 100) for the other models. We used the same configurations on other corpora as well. Because of the convergence issue, we applied SpecAugment only in stage 2 only. Moreover, we increased $\lambda_{\rm sync}$ to 4.0 for the UniLSTM MoChA model.¹⁰ The naive streaming MoChA models only with QR did not obtain any improvement with SpecAugment, whereas the performance of the RNN-T models improved. Therefore, this was a problem on the decoder side, rather than the encoder side. However, CTC-ST mitigated this problem and showed additional 12.8% and 13.1% relative improvements for the UniSLTM and LC-BLSTM models, respectively, on the test set. Finally, MoChA optimized with CTC-ST matched the performance of RNN-T on the test set when SpecAugment was applied as well.

B. Librispeech

Table V summarizes the results on the LibriSpeech corpus. For the UniLSTM MoChA model, CTC-ST achieved relative improvements of 14.8% and 5.8% on the *test-clean* and *testother* sets, respectively. For the LC-BLSTM MoChA model, we obtained gains of 27.0% and 2.1% with CTC-ST on the *test-clean* and *test-other* sets, respectively. Therefore, we conclude that CTC-ST is effective for large-scale data as well. Furthermore, the best MoChA models outperformed their RNN-T counterparts except for the *dev-clean* set.

We also compared our models with streaming RNN-based E2E systems reported in the literature. Our enhanced MoChA models optimized with CTC-ST showed the best performance. To compare our model with sMoChA [30] and MTA [33], we deactivated SpecAugment in the LC-BLSTM MoChA model. The results were 3.9% and 11.2% on the *test-clean* and *test-other* sets, and we still confirmed that our method outperformed them. The average lookahead latency of our LC-BLSTM MoChA was

¹⁰We found that this was beneficial when using SpecAugment for a weak encoder like UniLSTM. The effect of λ_{sync} will be discussed in Section X-B.

TABLE V Results on Librispeech

		WER [%] (↓)	
Model	de	ev	tes	st
	clean	other	clean	other
Offline				
BLSTM - Global AED	2.5	7.2	2.6	7.5
BLSTM - RNN-T	2.8	8.1	3.1	8.7
Streaming (UniLSTM)				
PTDLSTM - Triggered attention [49]	5.6	16.2	5.7	16.9
MoChA + MWER [29]	-	-	5.6	15.5
MoChA + {char, BPE}-CTC [100]	-	-	4.4	15.2
RNN-T	3.7	11.6	4.0	11.6
MoChA + QR	4.5	11.4	4.7	11.9
MoChA + CTC-ST	3.8	10.9	4.0	11.2
Streaming (LC-BLSTM-64+32)				
sMoChA [30]	-	-	6.0	16.7
MTA [33]	-	-	4.2	13.3
Streaming (LC-BLSTM-40+40)				
RNN-T	3.3	9.7	3.5	10.1
MoChA + QR	3.9	9.0	4.8	9.3
MoChA + CTC-ST	3.3	8.8	3.5	9.1

SpecAugment was Used for All Models. QR: Quantity Regularization. CTC-ST: CTC-Synchronous Training.

TABLE VI RESULTS ON CSJ

Model	WER / CER [%] (\downarrow)			
	eval1	eval2	eval3	
Offline CNN-TDNN-LSTM +				
LF-MMI→sMBR [103] BLSTM - Global AED	7.5 / 5.7 6.9 / 5.5	6.2 / 4.9 5.3 / 4.2	6.2 / 4.5 5.9 / 4.7	
Streaming (UniLSTM) MoChA + QR MoChA + CTC-ST	9.1 / 7.1 8.7 / 6.4	7.0 / 5.1 6.4 / 5.3	7.6 / 5.8 7.2 / 5.6	
Streaming (LC-BLSTM-40+40) MoChA + QR MoChA + CTC-ST	7.6 / 6.0 7.4 / 5.8	5.9 / 4.8 5.6 / 4.5	6.5 / 5.0 6.4 / 4.9	

SpecAugment was Used for All Models.

660 ms (= 400 ms $(N_c)/2$ + 400 ms (N_r) + 60 ms (CNN) while that was 640 ms (= 640 ms/2 + 320 ms) in [30], [33]. Therefore, we consider that the difference in the lookahead latency is negligible. Because CTC-ST is a method to enhance attention-based decoders, any type of encoder can be used. Further improvement is expected by adapting Transformer [102] and Conformer [103] encoders, which we leave for a future work.

C. Csj

Table VI summarizes the results on the CSJ. For both UniL-STM and LC-BLSTM MoChA models, we observed clear improvements with CTC-ST, which was consistent with the previous experiments. However, the relative gains for LC-BLSTM MoChA were smaller than those in other corpora. We reason that this was because the utterance lengths in the CSJ are relatively shorter, as shown in Fig. 3. We will analyze this behavior by simulating long-form speech utterances in Section IX.

Model	WER	[%] (↓)	
	dev	eval	
Offline			
TDNN + LF-MMI [104]	42.8	46.6	
BLSTM - Global AED	37.4	40.1	
BLSTM - RNN-T	38.1	41.2	
Streaming (LC-BLSTM-40+4	0)		
RNN-T	41.1	44.1	
MoChA + QR	46.5	49.7	
MoChA + CTC-ST	40.6	43.0	

TABLE VII Results on AMI (SDM)

SpecAugment was Used for All Models.



Fig. 5. WER of simulated long-form speech on CSJ *eval1* set, with different maximum input length threshold $T_{\rm cat}$. LC-BLSTM-40+40 encoder was used.

D. AMI

Table VII summarizes the results on the AMI corpus. We did not use any external LMs on this corpus because we did not observe any improvement. We observed significant improvement with CTC-ST in the far-field ASR task as well. CTC-ST gave 12.6% and 13.4% relative improvements over the baseline LC-BLSTM MoChA model on the *dev* and *eval* sets, respectively. LC-BLSTM MoChA also achieved better performance than that of LC-BLSTM RNN-T. Note that our baseline was very strong, as demonstrated by its superior performance compared to the TDNN+LF-MMI system [105]. As shown in Fig. 3, utterance lengths in the AMI corpus are relatively short compared to other corpora. Therefore, we can conclude that CTC-ST is also very effective for noisy speech, for which AED models have trouble learning alignments.

IX. EVALUATION OF ROBUSTNESS TO LONG-FORM SPEECH

In Section VIII-A, we have observed that CTC-ST was effective for reducing WER of long-form utterances on TEDLIUM2. In this section, we further analyzed this behavior by simulating long-form evaluation sets on other domains. We used CSJ and Librispeech for this purpose because input lengths of the original utterances in the evaluation sets were seen during training, as shown in Fig. 3. We simulated long-form utterances by merging adjacent utterances according to timestamps. Specifically, given a maximum input length threshold $T_{\rm cat}$ [sec.], we concatenated adjacent utterances from the first utterance to the last utterance of the same speaker in a greedy way until the accumulated utterance length surpassed $T_{\rm cat}$.

Figs. 5 and 6 show the results on CSJ and Librispeech, respectively. On CSJ, the baseline MoChA model without SpecAugment (blue bars) performed well with manual audio segmentation. However, as $T_{\rm cat}$ increased, the performance was



Fig. 6. WER of simulated long-form speech on Librispeech *test-other* set, with different maximum input length threshold T_{cat} . LC-BLSTM-40+40 encoder was used.

gradually degraded, whereas MoChA models trained with CTC-ST only (red bars) were robust in recognizing the long-form utterances. On the other hand, the WER of the naive MoChA trained with SpecAugment (green bars) was increased quickly for longer utterances. This indicates that SpecAugment affected the training of the naive MoChA model, which could not be observed with the original test sets because they did not include unseen input lengths. However, we confirmed that CTC-ST mitigated this problem, showing the lowest WER on all lengths (purple bars). We also confirmed the effectiveness of CTC-ST on Librispeech in all length bins as well although SpecAugment without CTC-ST did not degrade WER as severely as on CSJ. Still, the absolute gains by CTC-ST were larger in long-form speech. Therefore, we can conclude that CTC-ST is effective for recognizing long-form speech, which is generally challenging for AED models [34], [85].

X. EVALUATION OF EMISSION LATENCY

In this section, we evaluate the emission latency and compare CTC-ST with alignment knowledge distillation from a hybrid system [32]. Moreover, we also compare MoChA with RNN-T.

A. Emission Latency Metric

1) Token Emission Latency (TEL): Unlike the algorithmic latency introduced by lookahead frames in the encoder, the token emission latency (TEL) represents the user-perceived latency in a real application [106]. Although some previous works investigated the endpoint latency corresponding the last token in voice search and assistance tasks [3], [106], [107], we mainly focus here on the *per-token* emission latency because we are interested in long-form speech applications, as in lectures and meetings.

Following [32], we define the TEL as the difference in timing between the reference and predicted boundaries. To obtain the reference token boundaries, we perform forced alignment with Kaldi. The predicted boundaries in an utterance are obtained from the input timesteps at which monotonic attention in MoChA is activated, *i.e.*, $\{j | z_{i,j} = 1\}_{i=1,...,U}$. The TEL of the *i*-th token in the *n*-th utterance, $\Delta_{n,i}$, is calculated as

$$\Delta_{n,i}[ms] = \mathbf{b}_i^n - \mathbf{b}_i^n, \tag{15}$$

where \hat{b}_i^n and b_i^n are the *i*-th predicted and reference boundaries, respectively, in the *n*-th utterance. We do not include the $\langle eos \rangle$ token for TEL calculation. A negative latency can be observed for some tokens because of premature boundary detection. To match the lengths of a hypothesis and the corresponding reference when calculating the TEL, we apply teacher forcing by conditioning the decoder on the ground-truth transcript. However, the WER is reported with beam search decoding using the same model. Therefore, the TEL is a *corpus-level* latency metric. In the following, we report the median (PT@50) and 90th (PT@90) percentile values of the corpus-level TEL distributions.

We also evaluate the TEL of CTC. In this case, we perform forced alignment and use the most plausible alignment path to calculate the TEL.

2) Word Emission Latency (WEL): To enable a comparison of the emission latency between MoChA and RNN-T, which have different output units in our experiments, we also evaluate the word emission latency (WEL). The WEL only considers the time differences of the last subword in each word. When calculating the WEL of RNN-T, we conduct forced alignment similarly to CTC and use the most plausible alignment path. Moreover, we also evaluate the three kinds of WEL; (1) per-word WEL (*average WEL*), (2) WEL corresponding to the first word in an utterance (*first WEL*) [106], and (3) WEL corresponding to the last word (*last WEL*) [64], [69].

B. Latency Evaluation of CTC-ST

As shown in Fig. 1, the naive MoChA tended to emit tokens later than the corresponding CTC spikes, and CTC-ST reduced the gap in the example. To evaluate this quantitatively, we calculated the TEL. We used UniLSTM encoders for this purpose, because LC-BLSTM encoders introduce the algorithmic latency on the encoder side, whereas we are interested in the emission latency on the decoder side. Tables VIII and IX summarize the results on the TEDLIUM2 and Librispeech corpora, respectively. The TEL on Librispeech was averaged over the test-clean and test-other sets. We first observed that SpecAugment significantly increased the TEL of the baseline MoChA model. We also evaluated the TEL from the CTC branch,¹¹ and it also increased slightly by applying SpecAugment. On the other hand, we confirmed that CTC-ST significantly reduced both TEL and WER on both corpora. Increasing λ_{sync} up to 4.0 showed improvements of both metrics on TEDLIUM2. On the other hand, on Librispeech, WER was best at $\lambda_{sync} = 1.0$ while the TEL was continuously reduced by increasing λ_{sync} up to 3.0. PT@50 of the baseline MoChA with SpecAugment was reduced by 240 ms and 240 ms on TEDLIUM2 and Librispeech, respectively. PT@90 was reduced by 600 ms and 280 ms on TEDLIUM2 and Librispeech, respectively. The TEL of MoChA matched that of CTC in most conditions, confirming the function of CTC-ST. CTC-ST traded WER and TEL effectively by changing λ_{sync} . Interestingly, we found that CTC-ST also reduced the TEL of the CTC branch by increasing λ_{sync} . This indicates that joint training reduced the TEL of the other branch interactively via the shared encoder. Although CTC itself had a delay from the reference acoustic boundaries, it provided better timing to emit tokens for MoChA.

TABLE VIII TOKEN EMISSION LATENCY (TEL) ON TEDLIUM2, WITH COLUMNS **A** INDICATING ALIGNMENT

Δ	Model	WER [%] (↓)		TEL [ms] (\downarrow)	
21	Model	dev	test	PT@50	PT@90
_	MoChA + QR^{\diamond} + from CTC branch MoChA + QR + from CTC branch	16.7 15.1 17.7 13.3	15.0 14.7 16.0 12.8	280 200 320 240	680 280 840 360
	MoChA + CTC-ST - $\lambda_{sync} = 0.5$ + from CTC branch	13.0 13.4	11.6	240 200	400 320
	- $\lambda_{\text{sync}} = 1.0$ + from CTC branch	13.6 13.6	11.6 13.1	200 200	360 280
CTC	$-\lambda_{\text{sync}} = 2.0 + \text{from CTC branch}$	13.3 13.7	11.5 13.2	120 120	280 240
	- $\lambda_{\text{sync}} = 3.0$ + from CTC branch	13.0 14.0	11.4 13.4	120 120	240 240
	$-\lambda_{\text{sync}} = 4.0 + \text{from CTC branch}$	12.9 14.1	11.5 13.8	80 120	240 200
	- $\lambda_{\text{sync}} = 5.0$ + from CTC branch	13.4 14.6	11.9 14.5	80 80	200 200
Hybrid ASR		13.4 12.9 12.5 13.0 13.5	11.5 11.2 11.0 11.3 11.7	160 200 280 240 280	240 320 440 400 480
	$-\lambda_{\text{minlt}} = 1.0$ $-\lambda_{\text{minlt}} = 2.0$ $-\lambda_{\text{minlt}} = 3.0$	14.4 13.5 13.9	12.7 11.7 11.6	240 240 200	520 360 320

UniLSTM Encoder Was Used for All Models. TEL Was Calculated on *Test* Set. PT@X Represents X-Th Percentile. δ_{decot} Corresponds to Number of Encoder Outputs [40ms]. \diamond SpecAugment Was NOT Used.

TABLE IX Token Emission Latency (TEL) on Librispeech

Δ		WER	WER [%] (↓)		ns] (↓)
A	Model	te clean	est other	PT@50	PT@90
_	MoChA + QR [♦] + from CTC branch MoChA + QR + from CTC branch	4.8 6.2 4.7 5.3	14.2 17.5 11.9 14.4	320 280 360 320	520 400 560 440
CTC		4.0 5.3 4.0 5.7 - 4.1	11.2 14.6 11.5 15.2 11.7	240 240 160 160 120	$ \begin{array}{r} 400 \\ 360 \\ 320 \\ 320 \\ 280 \\ 280 \\ \end{array} $
Hybrid	+ from CTC branch MoChA + DeCoT - $\delta_{decot} = 12$ - $\delta_{decot} = 16$ - $\delta_{decot} = 20$	6.0 3.9 3.9 4.0	15.7 11.6 11.4 11.2	120 240 320 360	280 360 440 480
ASR	$\begin{aligned} & \text{MoChA + MinLT} \\ & -\lambda_{\text{minlt}} = 1.0 \\ & -\lambda_{\text{minlt}} = 2.0 \end{aligned}$	4.5 5.6	11.7 12.4	320 280	400 400

UniLSTM Encoder Was Used for All Models. TEL Was Averaged Over *Test-Clean* and *Test-Other* Sets. \diamond SpecAugment Was NOT Used.

¹¹The corresponding WER was calculated with a beam width of 10 and shallow fusion.

TABLE X Word Emission Latency (WEL) of MoChA and RNN-T on TEDLIUM2, With All Models Using SpecAugment

Model	Avg WEI	Ĺ [ms] (↓)	First W	EL [ms]	Last WE	L [ms]
model	PT@50	PT@90	PT@50	PT@90	PT@50	PT@90
RNN-T	240	320	240	400	0	160
MoChA	320	840	320	640	80	440
+ CTC-ST	80	240	80	200	-80	40

TABLE XI Word Emission Latency (WEL) of MoChA and RNN-T on Librispeech, With All Models Using SpecAugment

Model	Avg WEL [ms] (\downarrow)		First W	First WEL [ms]		Last WEL [ms]	
moder	PT@50	PT@90	PT@50	PT@90	PT@50	PT@90	
RNN-T	280	400	320	400	160	280	
MoChA	360	480	400	520	240	360	
+ CTC-ST	160	280	200	320	0	120	

C. Comparison With Alignment Distillation From Hybrid System

Next, we compared CTC-ST with methods of alignment knowledge distillation methods from a hybrid system. We trained MoChA models with MinLT and DeCoT on the same model configuration. Tables VIII and IX summarize the results on the TEDLIUM2 and Librispeech corpora, respectively. We observed that both DeCoT and MinLT also reduced the WER and TEL from those of the baseline model.¹² DeCoT with the optimal δ_{decot} outperformed MinLT in both metrics on both corpora. CTC-ST also outperformed MinLT in both metrics. On Librispeech, increasing λ_{sync} brought a large TEL reduction without hurting the WER so much while MinLT sacrificed the WER a lot with a small TEL reduction. Compared to DeCoT, CTC-ST achieved a lower TEL, especially for PT@50, with a comparable WER. This was because DeCoT focused on tokens whose emission latency surpassed δ_{decot} , and thus the TEL reduction was large in tail parts (PT@90). On the other hand, CTC-ST reduced the emission latency of all tokens. Therefore, we conclude that CTC-ST can achieve a similar or better tradeoff compared to alignment knowledge distillation from a hybrid system without relying on external alignment information.

D. Comparison With RNN-T

Finally, we compared the emission latency between MoChA and RNN-T. In addition to the average per-word statistics, we also calculated the WEL corresponding to the first and last tokens. The results in Tables X and XI present the WEL on TEDLIUM2 and Librispeech, respectively. We confirmed that MoChA trained with CTC-ST achieved lower WELs than those of RNN-T in all the conditions on both corpora. Note that RNN-T was also jointly trained with the CTC objective, and thus can be regarded as a strong baseline. Comparing the first WEL and the last WEL, we found that the latter had a lower latency. We reason that more acoustic contexts were necessary to emit the first word because there was no linguistic context on the decoder side. On

TABLE XII Comparison of CTC Alignment Generation in CTC-ST on TEDLIUM2, With All Models Using SpecAugment

Model	WER [%] (↓)		
	dev	test	
UniLSTM - MoChA			
on the fly	13.6	11.6	
precomputing	13.9	12.0	
LC-BLSTM-40+40 - MoChA			
on the fly	10.6	8.6	
precomputing	11.2	9.2	

TABLE XIII COMPARISON OF CTC ALIGNMENT GENERATION IN CTC-ST ON LIBRISPEECH, WITH ALL MODELS USING SPECAUGMENT

Model	WER [%] (↓)			
	dev		test	
	clean	other	clean	other
UniLSTM - MoChA				
on the fly	3.8	10.9	4.0	11.2
precomputing	3.8	10.7	4.0	11.1
LC-BLSTM-40+40 - MoChA				
on the fly	3.3	8.8	3.5	9.1
precomputing	3.4	8.8	3.5	9.1

the other hand, the last WEL of MoChA with CTC-ST was close to zero.

XI. ANALYSIS

In this section, we perform an ablation study of alignment generation in CTC-ST. Finally, we compare MoChA and RNN-T in terms of the inference speed.

A. Effect of Incremental Alignment Update

In the above experiments, we generated the reference boundaries from CTC alignments with the model parameters at each training step on the fly. Here, we investigated the effect of using *fixed* reference boundaries throughout stage 2 by using parameters optimized in stage 1. We refer to this strategy as precomputing. When generating the CTC alignments for precomputing, we deactivated SpecAugment and other regularization methods such as dropout. We used $\lambda_{sync} = 1.0$ in this experiment. Tables XII and XIII summarize the results on the TEDLIUM2 and Librispeech corpora, respectively. For TEDLIUM2, the on-the-fly CTC alignment generation consistently outperformed the precomputing strategy, regardless of the encoder type. Note, however, that the precomputing strategy also significantly outperformed the baseline listed in Table IV. For Librispeech, precomputing showed similar performances to those of on-the-fly computing. This was because the parameters learned in stage 1 had already provided good CTC alignments by leveraging more training data. This also confirms the observation that CTC-ST achieved similar WERs to those of DeCoT on Librispeech in Section X-C.

¹²Unlike in [32], we applied SpecAugment to DeCoT and MinLT. We found that those methods can also tolerate noisy inputs to some extent.



Fig. 7. Inference speed as measured by real time factor (RTF). TSD: timesynchronous decoding.

B. Inference Speed

While we have shown that MoChA can match RNN-T in terms of accuracy, we also evaluated its efficiency as compared to RNN-T in terms of the inference speed. For both models, we precomputed token embeddings before decoding. For RNN-T, we cached prediction network states corresponding to the same hypothesis [13] and batched all hypotheses in the beam for updating the prediction network and joint network [108]. We applied both TSD and mono-TSD as the search algorithm [98]. The maximum expansion number was set to 3 per frame in the TSD algorithm. We used the best MoChA (optimized with CTC-ST) and RNN-T models trained with SpecAugment, with the beam width set to 10 and $\{5, 10\}$, respectively. The inference speed was measured with a 6-core Intel(R) Xeon(R) Gold 6128 CPU @ 3.4 GHz. We investigated {1, 2, 4} threads, and we report the real-time factor (RTF) obtained by averaging five trials. Fig. 7 shows the results on the TEDLIUM2 test set. We observed that all MoChA models achieved an RTF of less than 1.0 with a single thread in a Python implementation. Using more threads led to faster decoding. The UniLSTM encoder was slower than the LC-BLSTM encoder because of the incremental state update on a per-frame basis. On the other hand, RNN-T required using the mono-TSD algorithm with the half beam width to achieve a similar inference speed. Moreover, RNN-T with the TSD algorithm was much slower because of the multiple symbol expansions per frame.

XII. CONCLUSION

In this article, we have proposed CTC synchronous training (CTC-ST), a self-distillation method for knowledge in inputoutput alignment to improve the performance of streaming AED models. Specifically, we distill knowledge of token boundary positions from a CTC model to a MoChA model, both of which share an encoder and are trained jointly. The proposed method forces MoChA to generate tokens in positions similar to those predicted by CTC, by synchronizing both sets of token boundaries during training. Experimental evaluations on four benchmark datasets demonstrated that the proposed method significantly improved MoChA in terms of both the recognition accuracy and the emission latency, especially for long-form and noisy utterances. We also compared the proposed method with methods of alignment knowledge distillation from an external hybrid ASR system and achieved a similar tradeoff of the accuracy and latency without any external alignments. Finally,

we showed that MoChA can achieve comparable recognition accuracy, lower emission latency, and faster inference speed compared to RNN-T.

In future work, we would like to further reduce the gap in recognition accuracy between RNN-T and MoChA in very long utterances. Reducing flicker of MoChA by selecting stable partial hypotheses that do not change in the subsequent prefix expansion, which were studied for incremental systems [59], [60], [109]–[113] and RNN-T [114], is also an interesting research direction.

REFERENCES

- C.-C. Chiu *et al.*, "State-of-the-art speech recognition with sequenceto-sequence models," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2018, pp. 4774–4778.
- [2] S. Karita *et al.*, "A comparative study on Transformer vs RNN in speech applications," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2019, pp. 499–456.
- [3] T. N. Sainath *et al.*, "A streaming on-device end-to-end model surpassing server-side conventional model quality and latency," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2020, pp. 6059–6063.
- [4] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks," in *Proc. 23rd Int. Conf. Mach. Learn.*, 2006, pp. 369–376.
- [5] A. Graves, "Sequence transduction with recurrent neural networks," 2012, arXiv:1211.3711.
- [6] H. Sak, M. Shannon, K. Rao, and F. Beaufays, "Recurrent Neural Aligner: An encoder-decoder neural network model for sequence to sequence mapping," in *Interspeech*, 2017, pp. 1298–1302.
- [7] E. Variani, D. Rybach, C. Allauzen, and M. Riley, "Hybrid autoregressive transducer (HAT)," in *Proc. Int. Conf. Acoust. Speech Signal Process.*, 2020, pp. 6139–6143.
- [8] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, "Attention-based models for speech recognition," in *Proc. 28th Int. Conf. Neural Inf. Process. Syst.*, 2015, pp. 577–585.
- [9] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2016, pp. 4960– 4964.
- [10] R. Prabhavalkar, K. Rao, T. N. Sainath, B. Li, L. Johnson, and N. Jaitly, "A comparison of sequence-to-sequence models for speech recognition," in *Proc. Interspeech*, 2017, pp. 939–943.
- [11] E. Battenberg et al., "Exploring neural transducers for end-to-end speech recognition," in Proc. IEEE Autom. Speech Recognit. Understanding Workshop, 2017, pp. 206–213.
- [12] M. Huang, Y. Lu, L. Wang, Y. Qian, and K. Yu, "Exploring model units and training strategies for end-to-end speech recognition," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2019, pp. 524–531.
- [13] Y. He et al., "Streaming end-to-end speech recognition for mobile devices," in Proc. IEEE Int. Conf. Acoust. Speech Signal Process., 2019, pp. 6381–6385.
- [14] J. Li, Y. Wu, Y. Gaur, C. Wang, R. Zhao, and S. Liu, "On the comparison of popular end-to-end models for large scale speech recognition," in *Proc. Interspeech*, 2020, pp. 1–5.
- [15] T. Bagby, K. Rao, and K. C. Sim, "Efficient implementation of recurrent neural network transducer in Tensorflow," in *Proc. IEEE Spoken Lang. Technol. Workshop*, 2018, pp. 506–512.
- [16] J. Li, R. Zhao, H. Hu, and Y. Gong, "Improving RNN transducer modeling for end-to-end speech recognition," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2019, pp. 114–121.
- [17] N. Jaitly, D. Sussillo, Q. V. Le, O. Vinyals, I. Sutskever, and S. Bengio, "A neural transducer," 2015, arXiv:1511.04868.
- [18] N. Moritz, T. Hori, and J. Le Roux, "Triggered attention for end-toend speech recognition," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2019, pp. 5666–5670.
- [19] M. Li, M. Liu, and H. Masanori, "End-to-end speech recognition with adaptive computation steps," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2019, pp. 6246–6250.

- [20] L. Dong and B. Xu, "CIF: Continuous integrate-and-fire for end-to-end speech recognition," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2020, pp. 6079–6083.
- [21] S. Zhang *et al.*, "Streaming chunk-aware multihead attention for online end-to-end speech recognition," in *Proc. Interspeech*, 2020, pp. 2142–2146.
- [22] M.-T. Luong, H. Pham, and C. D. Manning, "Effective approaches to attention-based neural machine translation," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2015, pp. 1412–1421.
- [23] W. Chan and I. Lane, "On online attention-based speech recognition and joint mandarin character-pinyin training," in *Proc. Interspeech*, 2016, pp. 3404–3408.
- [24] A. Tjandra, S. Sakti, and S. Nakamura, "Local monotonic attention mechanism for end-to-end speech and language processing," in *Proc. Int. Joint Conf. Natural Lang. Process.*, 2017, pp. 431–440.
- [25] J. Hou, S. Zhang, and L.-R. Dai, "Gaussian prediction based attention for online end-to-end speech recognition," in *Proc. Interspeech*, 2017, pp. 3692–3696.
- [26] A. Graves, "Generating sequences with recurrent neural networks," 2013, arXiv:1308.0850.
- [27] Y. Luo, C.-C. Chiu, N. Jaitly, and I. Sutskever, "Learning online alignments with continuous rewards policy gradient," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2017, pp. 2801–2805.
- [28] C.-C. Chiu and C. Raffel, "Monotonic chunkwise attention," in Proc. Int. Conf. Learn. Representations, 2018.
- [29] K. Kim et al., "Attention based on-device streaming speech recognition with large speech corpus," in Proc. IEEE Autom. Speech Recognit. Understanding Workshop, 2019, pp. 956–963.
- [30] H. Miao, G. Cheng, P. Zhang, T. Li, and Y. Yan, "Online hybrid CTC/attention architecture for end-to-end speech recognition," in *Proc. Interspeech*, 2019, pp. 2623–2627.
- [31] R. Fan, P. Zhou, W. Chen, J. Jia, and G. Liu, "An online attentionbased model for speech recognition," in *Proc. Interspeech*, 2019, pp. 4390–4394.
- [32] H. Inaguma, Y. Gaur, L. Lu, J. Li, and Y. Gong, "Minimum latency training strategies for streaming sequence-to-sequence ASR," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2020, pp. 6064–6068.
- [33] H. Miao, G. Cheng, P. Zhang, T. Li, and Y. Yan, "Online hybrid CTC/attention end-to-end automatic speech recognition architecture," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 28, pp. 1452–1465, Apr. 2020, doi: 10.1109/TASLP.2020.2987752.
- [34] C.-C. Chiu et al., "A comparison of end-to-end models for long-form speech recognition," in Proc. IEEE Autom. Speech Recognit. Understanding Workshop, 2019, pp. 889–896.
- [35] H. Sak, A. Senior, K. Rao, and F. Beaufays, "Fast and accurate recurrent neural network acoustic models for speech recognition," in *Proc. Interspeech*, 2015, pp. 1468–1472.
- [36] G. Pundak and T. N. Sainath, "Lower frame rate neural network acoustic models," in *Proc. Interspeech*, 2016, pp. 22–26.
- [37] H. Sak et al., "Acoustic modelling with CD-CTC-sMBR LSTM RNNs," in Proc. IEEE Autom. Speech Recognit. Understanding Workshop, 2015, pp. 604–609.
- [38] S. Zhang and M. Lei, "Acoustic modeling with DFSMN-CTC and joint CTC-CE learning," in *Proc. Interspeech*, 2018, pp. 771–775.
- [39] H. Hu, R. Zhao, J. Li, L. Lu, and Y. Gong, "Exploring pre-training with alignments for RNN transducer based end-to-end speech recognition," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2020, pp. 7079–7083.
- [40] J. Mahadeokar *et al.*, "Alignment restricted streaming recurrent neural network transducer," in *Proc. IEEE Spoken Lang. Technol. Workshop*, 2021, pp. 52–59.
- [41] H. Inaguma, M. Mimura, and T. Kawahara, "CTC-synchronous training for monotonic attention model," in *Proc. Interspeech*, 2020, pp. 571–575.
- [42] D. S. Park *et al.*, "SpecAugment: A simple data augmentation method for automatic speech recognition," in *Proc. Interspeech*, 2019, pp. 2613– 2617.
- [43] S. Kim, T. Hori, and S. Watanabe, "Joint CTC-attention based end-to-end speech recognition using multi-task learning," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2017, pp. 4835–4839.
- [44] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," 2015, arXiv:1503.02531.
- [45] N. Jaitly, Q. V. Le, O. Vinyals, I. Sutskever, D. Sussillo, and S. Bengio, "An online sequence-to-sequence model using partial conditioning," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2016, pp. 5067–5075.
- [46] T. N. Sainath *et al.*, "Improving the performance of online neural transducer models," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2018, pp. 5864–5868.

- [47] A. Graves, "Adaptive computation time for recurrent neural networks," 2016, arXiv:1603.08983.
- [48] G. Sterpu, C. Saam, and N. Harte, "Learning to count words in fluent speech enables online speech recognition," in *Proc. IEEE Spoken Lang. Technol. Workshop*, 2021, pp. 38–45.
- [49] N. Moritz, T. Hori, and J. Le Roux, "Streaming end-to-end speech recognition with joint CTC-attention based models," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2019, pp. 936–943.
- [50] N. Moritz, T. Hori, and J. L. Roux, "Streaming automatic speech recognition with the Transformer model," 2020, arXiv:2001.02674.
- [51] C. Wang *et al.*, "Low latency end-to-end streaming speech recognition with a scout network," in *Proc. Interspeech*, 2020, pp. 2112–2116.
- [52] T. G. Kang et al., "Learning monotonic alignments with source-aware GMM attention," 2021. [Online]. Available: https://openreview.net/ forum?id=y13JLBiNMsf
- [53] C.-C. Chiu et al., "An online sequence-to-sequence model for noisy speech recognition," 2017, arXiv:1706.06428.
- [54] C. Raffel, M.-T. Luong, P. J. Liu, R. J. Weiss, and D. Eck, "Online and linear-time attention by enforcing monotonic alignments," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 2837–2846.
- [55] E. Tsunoo, Y. Kashiwagi, T. Kumakura, and S. Watanabe, "Towards online end-to-end Transformer automatic speech recognition," 2019, arXiv:1910.11871.
- [56] H. Inaguma, M. Mimura, and T. Kawahara, "Enhancing monotonic multihead attention for streaming ASR," in *Proc. Interspeech*, 2020, pp. 2137–2141.
- [57] M. Li, C. Zorila, and R. Doddipatla, "Transformer-based online speech recognition with decoder-end adaptive computation steps," in *Proc. IEEE Spoken Lang. Technol. Workshop*, 2021, pp. 1–7.
- [58] S. Novitasari, A. Tjandra, S. Sakti, and S. Nakamura, "Sequence-tosequence learning via attention transfer for incremental speech recognition," in *Proc. Interspeech*, 2019, pp. 3835–3839.
- [59] D. Liu, G. Spanakis, and J. Niehues, "Low-latency sequence-to-sequence speech recognition and translation by partial hypothesis selection," in *Proc. Interspeech*, 2020, pp. 3620–3624.
- [60] T.-S. Nguyen, S. Stueker, and A. Waibel, "Super-human performance in online low-latency recognition of conversational speech," in *Proc. Interspeech*, 2021, pp. 1762–1766.
- [61] E. Tsunoo, Y. Kashiwagi, and S. Watanabe, "Streaming Transformer ASR with blockwise synchronous beam search," in *Proc. IEEE Spoken Lang. Technol. Workshop*, 2020, pp. 22–29.
- [62] H. Sak et al., "Learning acoustic frame labeling for speech recognition with recurrent neural networks," in Proc. IEEE Int. Conf. Acoust. Speech Signal Process., 2015, pp. 4280–4284.
- [63] T. N. Sainath, R. Pang, D. Rybach, B. García, and T. Strohman, "Emitting word timings with end-to-end models," in *Proc. Interspeech*, 2020, pp. 3615–3619.
- [64] J. Yu *et al.*, "FastEmit: Low-latency streaming ASR with sequence-level emission regularization," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2021, pp. 6004–6008.
- [65] B. Li et al., "A better and faster end-to-end model for streaming ASR," in Proc. IEEE Int. Conf. Acoust. Speech Signal Process., 2021, pp. 5634–5638.
- [66] A. Narayanan et al., "Cascaded encoders for unifying streaming and non-streaming ASR," in Proc. IEEE Int. Conf. Acoust. Speech Signal Process., 2021, pp. 5629–5633.
- [67] H. Inaguma and T. Kawahara, "StableEmit: Selection probability discount for reducing emission latency of streaming monotonic attention ASR," in *Proc. Interspeech*, 2021, pp. 1817–1821.
- [68] J. Kim, H. Lu, A. Tripathi, Q. Zhang, and H. Sak, "Reducing streaming ASR model delay with self alignment," in *Proc. Interspeech*, 2021, pp. 3440–3444.
- [69] J. Yu et al., "Dual-mode ASR: Unify and improve streaming ASR with full-context modeling," in Proc. Int. Conf. Learn. Representations, 2021.
- [70] R. Takashima, S. Li, and H. Kawai, "An investigation of a knowledge distillation method for CTC acoustic models," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2018, pp. 5809–5813.
- [71] S. Kim, M. Seltzer, J. Li, and R. Zhao, "Improved training for online end-to-end speech recognition systems," in *Proc. Interspeech*, 2018, pp. 2913–2917.
- [72] M. Huang, Y. You, Z. Chen, Y. Qian, and K. Yu, "Knowledge distillation for sequence model," in *Proc. Interspeech*, 2018, pp. 3703–3707.
- [73] R. Pang et al., "Compression of end-to-end models," in Proc. Interspeech, 2018, pp. 27–31.
- [74] G. Kurata and K. Audhkhasi, "Improved knowledge distillation from bidirectional to uni-directional LSTM CTC for end-to-end speech recognition," in *Proc. IEEE Spoken Lang. Technol. Workshop*, 2018, pp. 411–417.

- [75] R. M. Mun'im, N. Inoue, and K. Shinoda, "Sequence-level knowledge distillation for model compression of attention-based sequence-tosequence speech recognition," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2019, pp. 6151–6155.
- [76] G. Kurata and G. Saon, "Knowledge distillation from offline to streaming RNN transducer for end-to-end speech recognition," in *Proc. Interspeech*, 2020, pp. 2117–2121.
- [77] S. Panchapagesan, D. S. Park, C.-C. Chiu, Y. Shangguan, Q. Liang, and A. Gruenstein, "Efficient knowledge distillation for RNN-transducer models," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2021, pp. 5639–5643.
- [78] G. Kurata and K. Audhkhasi, "Guiding CTC posterior spike timings for improved posterior fusion and knowledge distillation," in *Proc. Inter*speech, 2019, pp. 1616–1620.
- [79] H. Ding, K. Chen, and Q. Huo, "Improving knowledge distillation of CTC-trained acoustic models with alignment-consistent ensemble and target delay," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 28, pp. 2561–2571, 2020.
- [80] T. Moriya, H. Sato, T. Tanaka, T. Ashihara, R. Masumura, and Y. Shinohara, "Distilling attention weights for CTC-based ASR systems," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2020, pp. 6894–6898.
- [81] T. Moriya et al., "Self-distillation for improving CTC-Transformer-based ASR systems," in Proc. Interspeech, 2020, pp. 546–550.
- [82] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, "Hybrid CTC/attention architecture for end-to-end speech recognition," *IEEE J. Sel. Top. Signal Process.*, vol. 11, no. 8, pp. 1240–1253, Dec. 2017.
- [83] M. Gales and S. Young, *The Application of Hidden Markov Models in Speech Recognit.*, Boston, MA, USA: Now, 2008.
- [84] M. He, Y. Deng, and L. He, "Robust sequence-to-sequence acoustic modeling with stepwise monotonic attention for neural TTS," in *Proc. Interspeech*, 2019, pp. 1293–1297.
- [85] R. Hsiao, D. Can, T. Ng, R. Travadi, and A. Ghoshal, "Online automatic speech recognition with listen, attend and spell model," *IEEE Signal Process. Lett.*, vol. 27, pp. 1889–1893, Oct. 2020, doi: 10.1109/LSP.2020.3031480.
- [86] Y. Zhao, C. Ni, C.-C. Leung, S. Joty, E. S. Chng, and B. Ma, "Preventing early endpointing for online automatic speech recognition," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2021, pp. 6813–6817.
- [87] Y. Zhang, G. Chen, D. Yu, K. Yaco, S. Khudanpur, and J. Glass, "Highway long short-term memory RNNs for distant speech recognition," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2016, pp. 5755–5759.
- [88] A. Rousseau, P. Deléglise, and Y. Estève, "TED-LIUM: An automatic speech recognition dedicated corpus," in *Proc. 8th Int. Conf. Lang. Resour. Eval.*, 2012, pp. 125–129.
- [89] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "LibriSpeech: An ASR corpus based on public domain audio books," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2015, pp. 5206–5210.
- [90] K. Maekawa, "Corpus of Spontaneous Japanese: Its design and evaluation," in Proc. ISCA IEEE Workshop Spontaneous Speech Process. Recognit., 2003.
- [91] J. Carletta et al., "The AMI meeting corpus: A pre-announcement" in Proc. Int. Workshop Mach. Learn. Multimodal Interact., 2005, pp. 28–39.
- [92] R. Sennrich, B. Haddow, and A. Birch, "Neural machine translation of rare words with subword units," in *Proc. 54th Annu. Meeting Assoc. Comput. Linguistics*, 2016, pp. 1715–1725.
- [93] D. Povey et al., "The Kaldi speech recognition toolkit," in Proc. IEEE Autom. Speech Recognit. Understanding Workshop, 2011.
- [94] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, "Audio augmentation for speech recognition," in *Proc. Interspeech*, 2015, pp. 3586–3589.
- [95] Z. Tüske, K. Audhkhasi, and G. Saon, "Advancing sequence-to-sequence based speech recognition," in *Proc. Interspeech*, 2019, pp. 3780–3784.
- [96] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, arXiv:1412.6980.
- [97] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 2818–2826.
- [98] A. Tripathi, H. Lu, H. Sak, and H. Soltau, "Monotonic recurrent neural network transducer and decoding strategies," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2019, pp. 944–948.
- [99] C. Wang *et al.*, "Semantic mask for Transformer based end-to-end speech recognition," in *Proc. Interspeech*, 2020, pp. 971–975.
- [100] A. Zeyer, P. Bahar, K. Irie, R. Schlüter, and H. Ney, "A comparison of Transformer and LSTM encoder decoder models for ASR," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2019, pp. 8–15.

- [101] A. Garg, D. Gowda, A. Kumar, K. Kim, M. Kumar, and C. Kim, "Improved multi-stage training of online attention-based encoder-decoder models," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2019, pp. 70–77.
- [102] A. Vaswani et al., "Attention is all you need," in Proc. Int. Conf. Neural Inf. Process. Syst., 2017, pp. 5998–6008.
- [103] A. Gulati *et al.*, "Conformer: convolution-augmented transformer for speech recognition," in *Proc. Interspeech*, 2020, pp. 5036–5040.
- [104] N. Kanda, Y. Fujita, and K. Nagamatsu, "Lattice-free state-level minimum Bayes risk training of acoustic models," in *Proc. Interspeech*, 2018, pp. 2923–2927.
- [105] V. Peddinti, V. Manohar, Y. Wang, D. Povey, and S. Khudanpur, "Far-field ASR without parallel data," in *Proc. Interspeech*, 2016, pp. 1996–2000.
- [106] Y. Shangguan *et al.*, "Dissecting user-perceived latency of on-device E2E speech recognition," in *Proc. Interspeech*, 2021, pp. 4553–4557.
- [107] B. Li et al., "Towards fast and accurate streaming end-to-end ASR," in Proc. IEEE Int. Conf. Acoust. Speech Signal Process., 2020, pp. 6069–6073.
- [108] G. Saon, Z. Tüske, and K. Audhkhasi, "Alignment-length synchronous decoding for RNN transducer," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2020, pp. 7804–7808.
- [109] C. Fügen, "A system for simultaneous translation of lectures and speeches," Ph.D. dissertation, Karlsruhe Inst. Technol., Karlsruhe, Germany, 2009.
- [110] T. Baumann, M. Atterer, and D. Schlangen, "Assessing and improving the performance of speech recognition for incremental systems," in *Proc. Hum. Lang. Technol. Annu. Conf. North Amer. Chapter Assoc. Comput. Linguistics*, 2009, pp. 380–388.
- [111] I. McGraw and A. Gruenstein, "Estimating word-stability during incremental speech recognition," in *Proc. Interspeech*, 2012, pp. 1019–1022.
- [112] E. Selfridge, I. Arizmendi, P. A. Heeman, and J. D. Williams, "Stability and accuracy in incremental speech recognition," in *Proc. 12th Annu. Meeting Special Interest Group Discourse Dialogue*, 2011, pp. 110–119.
- [113] T. S. Nguyen *et al.*, "Low latency ASR for simultaneous speech translation," 2020, arXiv:2003.09891.
- [114] Y. Shangguan, K. Knister, Y. He, I. McGraw, and F. Beaufays, "Analyzing the quality and stability of a streaming end-to-end on-device speech recognizer," in *Proc. Interspeech*, 2020, pp. 591–595.



Hirofumi Inaguma (Member, IEEE) received the B.E. degree in engineering, and the M.S. and Ph.D. degrees in informatics from Kyoto University, Kyoto, Japan, in 2016, 2018, and 2021, respectively. His research interests include automatic speech recognition and speech translation. He is a Member of ISCA, ASJ, and IPSJ.



Tatsuya Kawahara (Fellow, IEEE) received the B.E., M.E., and Ph.D. degrees in information science from Kyoto University, Kyoto, Japan, in 1987, 1989, and 1995, respectively. From 1995 to 1996, he was a Visiting Researcher with Bell Laboratories, Murray Hill, NJ, USA. He is currently a Professor and the Dean of the School of Informatics, Kyoto University. He was also an Invited Researcher with ATR and NICT. He has authored or coauthored more than 400 academic papers on speech recognition, spoken language processing, and spoken dialogue systems. He

is conducting several projects including speech recognition software Julius, the automatic transcription system deployed in the Japanese Parliament (Diet), and the autonomous android ERICA. Dr. Kawahara received the Commendation for Science and Technology by the Minister of Education, Culture, Sports, Science and Technology (MEXT) in 2012. From 2003 to 2006, he was a Member of IEEE SPS Speech Technical Committee. He was the General Chair of IEEE ASRU 2007. He was the Tutorial Chair of INTERSPEECH 2010, the Local Arrangement Chair of ICASSP 2012, and the General Chair of APSIPA ASC 2020. He was an Editorial Board Member of the Elsevier Journal of Computer Speech and Language and IEEE/ACM TRANSACTIONS ON AUDIO, SPEECH, AND LANGUAGE PROCESSING. He is the Editor-in-Chief of the APSIPA Transactions on Signal and Information Processing. Dr. Kawahara is a Board Member of the APSIPA and ISCA.