

# Semi-Supervised Speech Recognition via Graph-Based Temporal Classification

Moritz, Niko; Hori, Takaaki; Le Roux, Jonathan

TR2021-037 May 04, 2021

## Abstract

Semi-supervised learning has demonstrated promising results in automatic speech recognition (ASR) by self-training using a seed ASR model with pseudo-labels generated for unlabeled data. The effectiveness of this approach largely relies on the pseudo-label accuracy, for which typically only the 1-best ASR hypothesis is used. However, alternative ASR hypotheses of an N-best list can provide more accurate labels for an unlabeled speech utterance and also reflect uncertainties of the seed ASR model. In this paper, we propose a generalized form of the connectionist temporal classification (CTC) objective that accepts a graph representation of the training labels. The newly proposed graph-based temporal classification (GTC) objective is applied for self-training with WFST-based supervision, which is generated from an N-best list of pseudo-labels. In this setup, GTC is used to learn not only a temporal alignment, similarly to CTC, but also a label alignment to obtain the optimal pseudo-label sequence from the weighted graph. Results show that this approach can effectively exploit an N-best list of pseudo-labels with associated scores, considerably outperforming standard pseudo-labeling, with ASR results approaching an oracle experiment in which the best hypotheses of the N-best lists are selected manually.

*IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*



# SEMI-SUPERVISED SPEECH RECOGNITION VIA GRAPH-BASED TEMPORAL CLASSIFICATION

Niko Moritz, Takaaki Hori, Jonathan Le Roux

Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA

## ABSTRACT

Semi-supervised learning has demonstrated promising results in automatic speech recognition (ASR) by self-training using a seed ASR model with pseudo-labels generated for unlabeled data. The effectiveness of this approach largely relies on the pseudo-label accuracy, for which typically only the 1-best ASR hypothesis is used. However, alternative ASR hypotheses of an  $N$ -best list can provide more accurate labels for an unlabeled speech utterance and also reflect uncertainties of the seed ASR model. In this paper, we propose a generalized form of the connectionist temporal classification (CTC) objective that accepts a graph representation of the training labels. The newly proposed graph-based temporal classification (GTC) objective is applied for self-training with WFST-based supervision, which is generated from an  $N$ -best list of pseudo-labels. In this setup, GTC is used to learn not only a temporal alignment, similarly to CTC, but also a label alignment to obtain the optimal pseudo-label sequence from the weighted graph. Results show that this approach can effectively exploit an  $N$ -best list of pseudo-labels with associated scores, considerably outperforming standard pseudo-labeling, with ASR results approaching an oracle experiment in which the best hypotheses of the  $N$ -best lists are selected manually.

**Index Terms**— graph-based temporal classification, semi-supervised ASR, pseudo-labeling, self-training, WFST

## 1. INTRODUCTION

Over the last decade, automatic speech recognition (ASR) technologies have progressed to a point where ASR has become a central user interface for various electronic devices. This progress can largely be attributed to improved acoustic models using more sophisticated neural network architectures, as well as to improved sequence-to-sequence modeling approaches, known as end-to-end ASR, which have greatly contributed to optimizing the training and decoding process of ASR systems [1–3]. However, to achieve state-of-the-art performance, end-to-end ASR models are generally more data hungry compared to traditional hybrid HMM-DNN ASR solutions [4]. Although the available amount of manually transcribed training data has grown over the years, a lack of training data still exists, especially for low-resource domains and languages. To alleviate these problems, data augmentation as well as self- and semi-supervised learning can be applied in order to utilize unlabeled data for training. In self-supervised learning, typically a latent representation of speech is learned from unlabeled data to pre-train a neural network model for a following supervised or semi-supervised learning step [5–7]. Semi-supervised learning approaches exploit some amount of labeled data to support learning from unlabeled data [8,9]. Self-training uses a seed ASR model, trained from transcribed data, to generate pseudo-labels for unlabeled data for further training [10–

13]. This approach has recently become increasingly popular in end-to-end ASR as it has shown promising results [14–17]. However, as classical pseudo-labeling approaches only exploit the best ASR hypothesis for an unlabeled utterance, they ignore pseudo-label uncertainties of the system that may be useful [13] as well as alternative hypothesis with potentially fewer errors. In [18], an  $N$ -best list of ASR hypotheses is used by summing over the weighted losses of multiple pseudo-labels for a single speech utterance, where weights are estimated from scores of a strong language model (LM). In [19], multiple pseudo-labels are generated for each unlabeled speech utterance using different dropout settings, which are used for self-training with the purpose of capturing ASR uncertainties.

In this work, we propose a new objective function termed *graph-based temporal classification (GTC)*, which generalizes the popular connectionist temporal classification (CTC) loss function [20] to accept a graph representation as an input for modeling user-defined output symbol structures, including the possibility to assign probabilities to the generated output symbol sequences through transition weights. The proposed GTC loss function can be used to model all possible variations of CTC, including Gram-CTC [21], the automatic segmentation criterion (ASG) [22], and other still unexplored structures that can be modeled using weighted finite automata.

We apply the GTC loss to the semi-supervised learning problem, where we generate a WFST-based graph from an  $N$ -best list of pseudo-labels in order to leverage the information contained in alternative hypotheses as well as in their ASR scores. In such a setup, GTC is expected to find not only the best temporal alignments, similar to CTC, but also an optimal label sequence encoded in such a graph. Note that self-training with lattice-based supervision was also proposed in [23–25] using a hybrid ASR system and the LF-MMI objective in order to incorporate frame-level confidence scores and alternate pseudo-labels. However, we here consider the more general case of alternative alignment-free transcriptions using an end-to-end ASR system and the GTC objective function.

ASR results show that the proposed GTC-based semi-supervised learning approach can effectively utilize an  $N$ -best list of ASR hypotheses for self-training by achieving considerable improvements over the commonly used 1-best pseudo-labeling approach.

## 2. GRAPH-BASED TEMPORAL CLASSIFICATION

Let us consider a feature sequence  $X$  of length  $T'$  derived from a speech utterance, processed by a neural network to output a posterior probability distribution sequence  $Y = (\mathbf{y}^1, \dots, \mathbf{y}^T)$ , of length  $T$  potentially different from  $T'$  due to downsampling, where  $\mathbf{y}^t$  denotes the vector of posterior probabilities and  $y_k^t$  the posterior probability for output symbol  $k$  at time  $t$ . We assume that the labeling information for training is represented by a graph  $\mathcal{G}$ . The graph-based temporal classification (GTC) objective function marginalizes over all possible node sequences that can be output by graph  $\mathcal{G}$ , which in-

cludes all valid node patterns as well as all valid temporal alignment paths, the latter being known from CTC [20]. Thus, the conditional probability for a given graph  $\mathcal{G}$  is defined by the sum over all node sequences in  $\mathcal{G}$ , which can be written as:

$$p(\mathcal{G}|X) = \sum_{\pi \in \mathcal{S}(\mathcal{G}, T)} p(\pi|X), \quad (1)$$

where  $\mathcal{S}$  represents a search function that unfolds  $\mathcal{G}$  to all possible node sequences of length  $T$  (not counting non-emitting start and end nodes),  $\pi$  denotes a single node sequence and alignment path, and  $p(\pi|X)$  is the posterior probability for the path  $\pi$  given feature sequence  $X$ .

We introduce a few more notations that will be useful to derive  $p(\mathcal{G}|X)$ . We index the nodes using  $g = 0, \dots, G + 1$ , sorting them in a breadth-first search manner from 0 (non-emitting start node) to  $G + 1$  (non-emitting end node). We denote by  $l(g)$  the output symbol observed at node  $g$ , and by  $W_{(g, g')}$  the transition weight on edge  $(g, g')$ . Finally, we denote by  $\pi_{t:t'} = (\pi_t, \dots, \pi_{t'})$  the node sub-sequence of  $\pi$  from time index  $t$  to  $t'$ . Note that  $\pi_0$  and  $\pi_{T+1}$  correspond to the non-emitting start and end nodes 0 and  $G + 1$ .

In CTC, the conditional probabilities  $p(l|X)$  for a given label sequence  $l$  are computed efficiently by a dynamic programming algorithm, which is based on computing the forward and backward variables and stitching both together to compute  $p(l|X)$  at any given time index  $t$  [20]. In a similar fashion, the GTC forward probability can be computed for  $g = 1, \dots, G$  using

$$\alpha_t(g) = \sum_{\substack{\pi \in \mathcal{S}(\mathcal{G}, T): \\ \pi_{0:t} \in \mathcal{S}(\mathcal{G}_{0:g}, t)}} \prod_{\tau=1}^t W_{(\pi_{\tau-1}, \pi_{\tau})} y_{l(\pi_{\tau})}^{\tau}, \quad (2)$$

where  $\mathcal{G}_{0:g}$  denotes the sub-graph of  $\mathcal{G}$  starting at node 0 and terminating at node  $g$ . The sum is taken over all possible  $\pi$  whose sub-sequence up to time index  $t$  can be generated in  $t$  steps from the sub-graph  $\mathcal{G}_{0:g}$ . The backward variable  $\beta$  is computed similarly for  $g = 1, \dots, G$  using

$$\beta_t(g) = \sum_{\substack{\pi \in \mathcal{S}(\mathcal{G}, T): \\ \pi_{t:T+1} \in \mathcal{S}(\mathcal{G}_{g:T+1}, T-t+1)}} \prod_{\tau=t}^T W_{(\pi_{\tau}, \pi_{\tau+1})} y_{l(\pi_{\tau})}^{\tau}, \quad (3)$$

where  $\mathcal{G}_{g:T+1}$  denotes the sub-graph of  $\mathcal{G}$  starting at node  $g$  and terminating at node  $G + 1$ . By using the forward and backward variables, the probability function  $p(\mathcal{G}|X)$  can be computed for any  $t$  by summing over all  $g$ :

$$p(\mathcal{G}|X) = \sum_{g \in \mathcal{G}} \frac{\alpha_t(g) \beta_t(g)}{y_{l(g)}^t}. \quad (4)$$

For gradient descent training, the loss function

$$\mathcal{L} = -\ln p(\mathcal{G}|X) \quad (5)$$

must be differentiated with respect to the network outputs, which can be written as:

$$-\frac{\partial \ln p(\mathcal{G}|X)}{\partial y_k^t} = -\frac{1}{p(\mathcal{G}|X)} \frac{\partial p(\mathcal{G}|X)}{\partial y_k^t}, \quad (6)$$

for any symbol  $k \in \mathcal{U}$ , where  $\mathcal{U}$  denotes a set of all possible output symbols.

Because  $\alpha_t(g) \beta_t(g) / y_{l(g)}^t$  is proportional to  $y_{l(g)}^t$ ,

$$\frac{\partial(\alpha_t(g) \beta_t(g) / y_{l(g)}^t)}{\partial y_{l(g)}^t} = \frac{\alpha_t(g) \beta_t(g)}{y_{l(g)}^{t+2}}, \quad (7)$$

and from (4), we can derive

$$\frac{\partial p(\mathcal{G}|X)}{\partial y_k^t} = \frac{1}{y_k^{t+2}} \sum_{g \in \Psi(\mathcal{G}, k)} \alpha_t(g) \beta_t(g), \quad (8)$$

where  $\Psi(\mathcal{G}, k) = \{g \in \mathcal{G} : l(g) = k\}$  denotes the set of nodes in  $\mathcal{G}$  at which symbol  $k$  is observed.

To backpropagate the gradients through the softmax function, we need the derivative with respect to the unnormalized network outputs  $u_k^t$  before the softmax is applied, which is

$$-\frac{\partial \ln p(\mathcal{G}|X)}{\partial u_k^t} = -\sum_{k' \in \mathcal{U}} \frac{\partial \ln p(\mathcal{G}|X)}{\partial y_{k'}^t} \frac{\partial y_{k'}^t}{\partial u_k^t}. \quad (9)$$

By substituting (8) and the derivative of the softmax function  $\partial y_{k'}^t / \partial u_k^t = y_{k'}^t \delta_{kk'} - y_{k'}^t y_k^t$  into (9), we finally derive

$$-\frac{\partial \ln p(\mathcal{G}|X)}{\partial u_k^t} = y_k^t - \frac{1}{y_k^t p(\mathcal{G}|X)} \sum_{g \in \Psi(\mathcal{G}, k)} \alpha_t(g) \beta_t(g), \quad (10)$$

where we used the fact that

$$\begin{aligned} \sum_{k' \in \mathcal{U}} \frac{1}{y_{k'}^t} \sum_{g \in \Psi(\mathcal{G}, k')} \alpha_t(g) \beta_t(g) &= \sum_{k' \in \mathcal{U}} \sum_{g \in \Psi(\mathcal{G}, k')} \frac{\alpha_t(g) \beta_t(g)}{y_{l(g)}^t} \\ &= \sum_{g \in \mathcal{G}} \frac{\alpha_t(g) \beta_t(g)}{y_{l(g)}^t} \\ &= p(\mathcal{G}|X), \end{aligned} \quad (11)$$

and

$$\sum_{k' \in \mathcal{U}} \frac{\partial \ln p(\mathcal{G}|X)}{\partial y_{k'}^t} y_{k'}^t y_k^t = \frac{1}{p(\mathcal{G}|X)} p(\mathcal{G}|X) y_k^t = y_k^t. \quad (12)$$

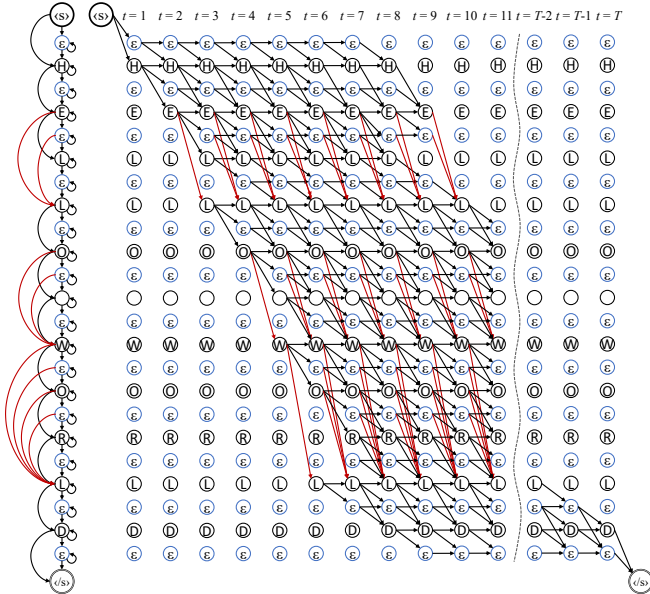
For efficiency reason, we implemented the GTC objective in CUDA as an extension for PyTorch.

### 3. GRAPH GENERATION FOR SELF-TRAINING

In the supervised case, GTC can be used to train an ASR system similarly to CTC by setting the graph to insert blank symbols between the labels of the ground truth transcription and include CTC-like transition rules [20]. However, we here wish to consider more general cases in which there may be multiple transcriptions obtained from a seed ASR system for an unlabeled utterance. The proposed GTC loss makes it possible to learn the model parameters with label information in a graph format. In a semi-supervised ASR scenario, we can use  $N$ -best hypotheses or word/token lattices as pseudo labels, which are typically given by a baseline system trained with a small amount of labeled data. As an example, a graph representation of an  $N$ -best hypotheses list is shown in Fig. 1. For computational reasons, it is preferable to make the graph compact while retaining correct predictions as much as possible.

In this work, we generate a compact CTC-like graph from  $N$ -best hypotheses according to the following steps:

1. Convert the  $N$ -best hypotheses to a sausage-form confusion network (CN) using minimum Bayes risk decoding [26].



**Fig. 1.** CTC-style confusion network  $\mathcal{G}$  on the left side with the unfolded graph  $\mathcal{S}(\mathcal{G}, T)$  on the right side, which is generated from an  $N$ -best list using character-based label symbols: "HELO WORLD", "HELO WOLD", "HELO WOLD", "HELLOWLD". Circles denote nodes with a corresponding label symbol observation inside, where  $\varepsilon$  is the blank symbol. Black arrows correspond to transitions of correct label paths and red arrows to transitions of alternative label paths. Although the correct transcription is not present in the given  $N$ -best list, it can be recovered from the presented graph. Each transition is also associated with a transition weight, which are not shown in this figure.


2. Convert the CN into an optimized WFST by applying the epsilon removal, determinization, and minimization operations [27] to the CN, where the arc weights are operated in log-semiring to ensure the WFST remains probabilistic.
3. Convert the WFST to a CTC-like graph by replacing each state with a blank node and each arc with a non-blank node to which the arc label is assigned, where edges between the nodes are made based on the optimized WFST structure and the CTC rule, i.e., each blank node can be skipped if the adjacent non-blank nodes have different labels.

In Step 1, a scaling factor  $\mu$  can be applied to the ASR score (log probability) of each hypothesis, where the scaling factor indicates the degree to which the ASR scores are reflected in the label probabilities in the CN:  $\mu = 1$  means the ASR scores are used without alteration, and  $\mu = 0$  means the hypotheses are treated equally without considering the ASR scores. We optionally add pruning steps after steps 1 and 2 to reduce the size of the CTC graph, which eliminate arcs if the assigned probabilities are less than a threshold  $\eta$ .

## 4. EXPERIMENTS

### 4.1. Dataset

We use as the ASR benchmark the LibriSpeech corpus of read English audio books [28], which provides about 960 hours of training data, 10.7 hours of development data, and 10.5 hours of test data. The development and test data sets are both split into approximately two halves named "clean" and "other" based on the quality of the



**Fig. 2.** ASR system architecture.

recorded utterances [28]. The training data is also split into three subsets: "clean" 100 hours, "clean" 360 hours, and "other" 500 hours. We use the "clean" 100 hours subset for supervised training and consider the remaining 860 hours as unlabeled data.

### 4.2. ASR System

Figure 2 illustrates the ASR system used in this work, a transformer-based neural network architecture that employs the proposed GTC loss function of Section 2 for training. We use 80-dimensional log-mel spectral energies plus 3 extra features for pitch information as acoustic features input to the neural network [29]. The neural network architecture consists of a two-layer convolutional neural network (CNN) module followed by a stack of  $E = 12$  transformer-based encoder layers with a linear layer plus softmax function at the end to project the neural network outputs to posterior probabilities. Each layer of the 2-layer CNN module is using a stride of 2, a kernel size of  $3 \times 3$ , 320 channels, and a rectified linear unit (ReLU) activation function. In addition, a linear neural network layer is applied to the output of the last CNN layer. Sinusoidal positional encodings [30] are added to the 2-layer CNN module output before feeding it to the transformer-based encoder. Each transformer layer employs 320-dimensional self-attention layers with 4 attention heads, layer normalization, and a feed-forward neural network module of inner dimension 1540. Residual connections are applied to the self-attention and feed-forward module outputs. Dropout with a probability of 0.1 is used after self-attention and after the feed-forward module as well as for the inner dimension of the feed-forward module. In addition, SpecAugment-based data augmentation [31] is utilized for training. ASR output symbols consist of a blank symbol plus 5000 subwords obtained by the SentencePiece method [32], which we generated from the transcripts of the "clean" 100h LibriSpeech training data subset only. The ASR model is trained for 100 epochs using the Adam optimizer with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.98$ ,  $\epsilon = 10^{-9}$ , and learning rate scheduling similar to [30] with 25000 warmup steps and a learning rate factor of 5.0.

A language model (LM) is employed via shallow fusion at inference time, which consists of 2 long short-term memory (LSTM) neural network layers with 1024 units each trained using stochastic gradient descent and the official LM training text data of LibriSpeech [28], where we excluded sentences that occur in the 860h training data subsets. ASR decoding is based on a time-synchronous prefix beam search algorithm similar to [33, 34]. We use a decoding

**Table 1.** Oracle label error rates (LER) [%] for the 860 hours “unlabeled” training data subsets “clean” and “other” using different pseudo-label representations.  $CN^{20}$  denotes a confusion network generated from the 20-best ASR hypotheses for each utterance, where  $low$  and  $high$  indicate low and high pruning settings.

data set	1-best	10-best	20-best	$CN^{20}$	$CN^{20}_{low}$	$CN^{20}_{high}$
clean 360h	11.3	9.0	8.7	8.3	8.6	8.9
other 500h	18.3	15.7	15.3	14.5	14.8	15.3

beam size of 30, a score-based pruning threshold of 14.0, an LM-weight of 0.8, and an insertion bonus factor of 2.0.

### 4.3. Graph Analysis

A seed ASR model is trained by using the 100h “clean” LibriSpeech training data set and an  $N$ -best list of ASR hypotheses is generated for each utterance of the remaining 860h of training data. Each  $N$ -best list of pseudo-labels is used to generate a CTC-style confusion network (CN) as discussed in Section 3 with different settings for pruning. We compare “no”, “low”, and “high” pruning settings in our experiments, where we use a scaling factor of  $\mu = 0.6$  and a threshold of  $\eta = 0.02$  for “low” and  $\eta = 0.05$  for “high” pruning. The pruning settings are determined based on experiments using the development data sets of LibriSpeech, where selected parameters resulted in a good trade-off between the oracle label error rates (LERs) and the graph densities, which ideally should both be small. Pruning reduces the size and density of a graph, as can be identified by the ratio of the number of non-blank nodes in the graph to the number of labels in a reference sequence, the ground-truth transcription: the average graph densities for “no”, “low”, and “high” pruning are 1.510, 1.233, and 1.119 for the “clean” 360h training data set, and 1.545, 1.275, and 1.147 for the “other” 500h training data set. Table 1 shows the oracle LERs for  $N$ -best lists of different sizes  $N$  as well as for CNs that are generated from the 20-best ASR hypotheses for each utterance. Oracle LERs of an  $N$ -best list are obtained by selecting the best pseudo-label sequence from that list, i.e., the sequence that has the smallest edit distance compared to the ground-truth transcription. Oracle LERs of a graph correspond to the minimum edit distance between an FST and a reference sequence, the ground-truth transcription, which is computed by a composition operation between an acyclic FST and an edit-distance FST, followed by a single-source shortest-path algorithm [35]. We can see from Table 1 that an  $N$ -best list contains ASR hypotheses with much lower error rates compared to the 1-best hypothesis. For example, selecting the oracle hypothesis from the 20-best list reduces the average LER by 2.5% (clean 360h) and 3.0% (other 500h) on an absolute scale. Using an oracle pseudo-label computed from an  $N$ -best list in a graph format reduces the LER even further, since a graph representation of an  $N$ -best list allows for more flexible label combinations, as illustrated in Fig. 1.

### 4.4. ASR Results

GTC-based ASR results for different pseudo-label representations are shown in Table 2. “N/A” denotes results of the seed ASR model, which is trained using the 100h of labeled clean LibriSpeech training data only.  $CN^{20}$  indicates the use of CTC-style confusion networks that are generated from the 20-best ASR hypotheses obtained from the seed model for each utterance of the 860h of unlabeled training data. Four different CN setups are compared: 1) without pruning and with all transition weights set to 1.0, 2) with probabilistic

**Table 2.** WERs [%] of models trained with different pseudo-label representations. N/A denotes the seed ASR model, which is used to generate pseudo-labels. 1-best denotes self-training results using the best ASR hypothesis only.  $W = 1.0$  indicates that all transition weights of the graph  $CN^{20}$  are set to 1, and  $W = p$  that probabilistic transition weights are used according to Section 3.

pseudo-labels	$W$	pruning	dev		test	
			clean	other	clean	other
N/A			7.9	20.0	8.2	21.2
1-best			5.8	14.9	6.0	15.6
$CN^{20}$	1.0	no	5.8	14.7	6.0	15.5
$CN^{20}$	$p$	low	5.7	14.7	5.9	15.1
$CN^{20}$	1.0	high	5.6	14.4	6.0	15.2
$CN^{20}$	$p$	high	5.6	14.2	5.9	15.0
oracle 20-best			5.1	13.6	5.4	14.2
ground-truth			3.3	8.2	3.5	8.4

transition weights and low pruning (see Section 4.3), 3) with transition weights set to 1.0 and high pruning, and 4) with probabilistic transition weights and high pruning. In addition, ASR results obtained when training on the best pseudo-label sequence manually selected in an oracle fashion from the 20-best list are shown as “oracle 20-best” and results for supervised ASR using the ground-truth transcriptions as “ground-truth”. Table 2 shows that 1-best pseudo-labeling improves the word error rates (WERs) of the LibriSpeech test data sets by a large margin, e.g., from 21.2% to 15.6% for test-other. Training on multiple pseudo-label sequences encoded in graph format further improves the WERs, whereby the best results are obtained when using a graph with high pruning settings and probabilistic transition weights. Although an unpruned graph is more likely to contain the correct transcription, we suppose the large variance in such a graph makes it harder to learn the best label sequence due to more label noise. Therefore, pruning and the use of transition weights can guide self-training to find pseudo-label sequences of lower error rates. In the best setting, the proposed GTC-based self-training approach achieves 0.7% (dev-other) and 0.6% (test-other) lower WERs compared to 1-best pseudo-labeling and 0.6% and 0.8% higher WERs compared to the “oracle 20-best” ASR results, which approximately correspond to a lower bound for training on multiple pseudo-labels obtained from an  $N$ -best list.

## 5. CONCLUSIONS

We proposed a new objective function that generalizes the popular CTC loss function to accept weighted finite automata in order to train with label information in a graph format with user-defined transition rules and transition weights. The proposed graph-based temporal classification (GTC) loss is applied to a semi-supervised ASR problem in order to leverage a graph with a CTC-like structure generated from an  $N$ -best list of pseudo-labels for self-training. We demonstrate that GTC-based self-training improves ASR results compared to 1-best pseudo-labeling. The use of pruned graphs and probabilistic transition weights further helps GTC to better exploit pseudo-label information from such a graph with improved ASR results. In our experiments, GTC-based self-training achieved up to 0.7% better WERs compared to the commonly used 1-best pseudo-labeling approach, reducing the gap to an oracle experiment where the best pseudo-label sequence was selected manually from the 20-best list of ASR hypotheses.

## 6. REFERENCES

- [1] B. Li, S.-Y. Chang, T. Sainath, R. Pang, Y. R. He, T. Strohmaier, and Y. Wu, "Towards fast and accurate streaming end-to-end ASR," in *Proc. ICASSP*, May 2020.
- [2] N. Moritz, T. Hori, and J. Le Roux, "Streaming automatic speech recognition with the transformer model," in *Proc. ICASSP*, May 2020, pp. 6074–6078.
- [3] C. Wang, Y. Wu, S. Liu, J. Li, L. Lu, G. Ye, and M. Zhou, "Low latency end-to-end streaming speech recognition with a scout network," in *Proc. Interspeech*, Oct. 2020.
- [4] G. Pundak and T. Sainath, "Lower frame rate neural network acoustic models," in *Proc. Interspeech*, Sep. 2016.
- [5] A. Baeviski, H. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *arXiv preprint arXiv:2006.11477*, 2020.
- [6] A. T. Liu, S.-w. Yang, P.-H. Chi, P.-c. Hsu, and H.-y. Lee, "Mockingjay: Unsupervised speech representation learning with deep bidirectional transformer encoders," in *Proc. ICASSP*, May 2020.
- [7] S. Khurana, A. Laurent, and J. Glass, "Cstnet: Contrastive speech translation network for self-supervised speech representation learning," *arXiv preprint arXiv:2006.02814*, 2020.
- [8] T. Hori, R. Astudillo, T. Hayashi, Y. Zhang, S. Watanabe, and J. Le Roux, "Cycle-consistency training for end-to-end speech recognition," in *Proc. ICASSP*, May 2019, pp. 6271–6275.
- [9] A. H. Liu, H. Lee, and L. Lee, "Adversarial training of end-to-end speech recognition using a criticizing language model," in *Proc. ICASSP*, May 2019, pp. 6176–6180.
- [10] L. Lamel, J.-L. Gauvain, and G. Adda, "Lightly supervised and unsupervised acoustic model training," *Comput. Speech Lang.*, vol. 16, no. 1, pp. 115–129, 2002.
- [11] S. Novotney, R. Schwartz, and J. Ma, "Unsupervised acoustic and language model training with small amounts of labelled data," in *Proc. ICASSP*, Apr. 2009, pp. 4297–4300.
- [12] Y. Huang, D. Yu, Y. Gong, and C. Liu, "Semi-supervised GMM and DNN acoustic model training with multi-system combination and confidence re-calibration," in *Proc. Interspeech*, Aug. 2013.
- [13] K. Veselý, M. Hannemann, and L. Burget, "Semi-supervised training of deep neural networks," in *Proc. ASRU*, Dec. 2013, pp. 267–272.
- [14] B. Li, T. N. Sainath, R. Pang, and Z. Wu, "Semi-supervised training for end-to-end models via weak distillation," in *Proc. ICASSP*, May 2019, pp. 2837–2841.
- [15] F. Weninger, F. Mana, R. Gemello, J. Andrés-Ferrer, and P. Zhan, "Semi-supervised learning with data augmentation for end-to-end ASR," in *Proc. Interspeech*, Oct. 2020.
- [16] Q. Xu, T. Likhomanenko, J. Kahn, A. Hannun, G. Synnaeve, and R. Collobert, "Iterative pseudo-labeling for speech recognition," *arXiv preprint arXiv:2005.09267*, 2020.
- [17] S. Khurana, N. Moritz, T. Hori, and J. L. Roux, "Unsupervised domain adaptation for speech recognition via uncertainty driven self-training," *arXiv preprint arXiv:2011.13439*, 2020.
- [18] W.-N. Hsu, A. Lee, G. Synnaeve, and A. Hannun, "Semi-supervised speech recognition via local prior matching," *arXiv preprint arXiv:2002.10336*, 2020.
- [19] S. Dey, P. Motlicek, T. Bui, and F. Dernoncourt, "Exploiting semi-supervised training through a dropout regularization in end-to-end speech recognition," *arXiv preprint arXiv:1908.05227*, 2019.
- [20] A. Graves, S. Fernández, F. J. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proc. ICML*, vol. 148, Jun. 2006, pp. 369–376.
- [21] H. Liu, Z. Zhu, X. Li, and S. Satheesh, "Gram-CTC: Automatic unit selection and target decomposition for sequence labelling," in *Proc. ICML*, Aug. 2017, p. 2188–2197.
- [22] R. Collobert, C. Puhersch, and G. Synnaeve, "Wav2Letter: an end-to-end ConvNet-based speech recognition system," *arXiv preprint arXiv:1609.03193*, 2016.
- [23] V. Manohar, H. Hadian, D. Povey, and S. Khudanpur, "Semi-supervised training of acoustic models using lattice-free MMI," in *Proc. ICASSP*, Apr. 2018, pp. 4844–4848.
- [24] S. Tong, A. Vyas, P. N. Garner, and H. Bourlard, "Unbiased semi-supervised LF-MMI training using dropout," in *Proc. Interspeech*, Sep. 2019, pp. 1576–1580.
- [25] I. Sheikh, E. Vincent, and I. Illina, "On semi-supervised LF-MMI training of acoustic models with limited data," in *Proc. Interspeech*, Oct. 2020.
- [26] H. Xu, D. Povey, L. Mangu, and J. Zhu, "Minimum Bayes risk decoding and system combination based on a recursion for edit distance," *Comput. Speech Lang.*, vol. 25, no. 4, pp. 802–828, 2011.
- [27] M. Mohri, "Weighted automata algorithms," in *Handbook of Weighted Automata*. Springer, 2009, pp. 213–254.
- [28] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "LibriSpeech: An ASR corpus based on public domain audio books," in *Proc. ICASSP*, Apr. 2015.
- [29] T. Hori, S. Watanabe, Y. Zhang, and W. Chan, "Advances in joint CTC-attention based end-to-end speech recognition with a deep CNN encoder and RNN-LM," in *Proc. Interspeech*, Aug. 2017, pp. 949–953.
- [30] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. NIPS*, Dec. 2017, pp. 6000–6010.
- [31] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," *arXiv preprint arXiv:1904.08779*, 2019.
- [32] T. Kudo and J. Richardson, "SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing," *arXiv preprint arXiv:1808.06226*, 2018.
- [33] A. L. Maas, A. Y. Hannun, D. Jurafsky, and A. Y. Ng, "First-pass large vocabulary continuous speech recognition using bidirectional recurrent DNNs," *arXiv preprint arXiv:1408.2873*, 2014.
- [34] N. Moritz, T. Hori, and J. Le Roux, "Streaming end-to-end speech recognition with joint CTC-attention based models," in *Proc. ASRU*, Dec. 2019, pp. 936–943.
- [35] M. Mohri, "Edit-distance of weighted automata: General definitions and algorithms," in *Int. J. Found. Comput. Sci.*, 2003, pp. 957–982.