

# Right Label Context in End-to-End Training of Time-Synchronous ASR Models

Tina Raissi<sup>1</sup>, Ralf Schlüter<sup>1,2</sup>, Hermann Ney<sup>1,2</sup>

<sup>1</sup>Machine Learning and Human Language Technology Group, RWTH Aachen University

<sup>2</sup>AppTek GmbH  
Germany

**Abstract**—Current time-synchronous sequence-to-sequence automatic speech recognition (ASR) models are trained by using sequence level cross-entropy that sums over all alignments. Due to the discriminative formulation, incorporating the right label context into the training criterion’s gradient causes normalization problems and is not mathematically well-defined. The classic hybrid neural network hidden Markov model (NN-HMM) with its inherent generative formulation enables conditioning on the right label context. However, due to the HMM state-tying the identity of the right label context is never modeled explicitly. In this work, we propose a factored loss with auxiliary left and right label contexts that sums over all alignments. We show that the inclusion of the right label context is particularly beneficial when training data resources are limited. Moreover, we also show that it is possible to build a factored hybrid HMM system by relying exclusively on the full-sum criterion. Experiments were conducted on Switchboard 300h and LibriSpeech 960h.

**Index Terms**—CTC, HMM, factored hybrid HMM, full-sum, end-to-end,

## I. INTRODUCTION

The training of time-synchronous automatic speech recognition (ASR) models requires a set of possible alignment paths between the input feature and output label sequences, following a certain label topology [1]–[4]. The common sequence level cross-entropy training criterion that uses the sum over all alignments can be replaced by a frame-wise cross-entropy on the single best alignment path without significant loss of performance. This approach is the prevalent technique for training Gaussian mixture hidden Markov model (GM-HMM) and later the hybrid neural network (NN)-HMMs [5]. In these frameworks, after the initial linear segmentation, the alignment path is iteratively refined through multiple realignment steps. This type of training is designed to address convergence issues that occur when using the sum over all alignment paths [4]. Recent advances in encoder architecture design leveraging global dependencies through the self-attention mechanism and recurrent layers on the feature input side, along with the use of downsampling and larger label output units such as word-pieces and byte-pair encoding units [6] have mitigated part of this convergence issue. The mentioned modeling choices reduce the number of alignment paths, thereby imposing fewer constraints on the alignment quality.

Another difference between the classic and the sequence-to-sequence (seq2seq) frameworks is the use of the output label context in the frame-wise probability distribution. The

autoregressive factorization in seq2seq models that use label context does not allow for the inclusion of right (future) context. Whereas, in principle, there is no such constraint on right label context modeling in the classic generative approach. Nevertheless, even the right label context modeled via generalized triphone states in classic approaches necessitates a parameter tying with several inconsistencies in the resulting statistical model [7], [8]. Moreover, it has been shown that models such as recurrent neural network transducers (RNN-T) [2] do not require the full label history to achieve full performance. The label context history in these models can be restricted to one or two previous output labels with no word error rate (WER) degradation [9]–[12].

The issues investigated in this work are two-fold: first, whether the inclusion of the right label context in the model definition leads to ASR performance gain with full-sum training, and second, whether it is possible to build an end-to-end hybrid HMM system. For the first question we expanded on an initial investigation [13] with three key extensions and differences: (1) we evaluate an ASR model trained from scratch and with no intermediate forced alignment step, (2) we omit the HMM state prior during training, and (3) we use the Conformer encoder architecture [14]. We show that the inclusion of the right label context is particularly beneficial when data resources are limited. Moreover, we build a factored hybrid HMM [8], [15] using only the full-sum criterion, and achieve comparable performance to a multi-stage pipeline system that relies on an external alignment. To the best of our knowledge, no prior approach in the literature has explicitly modeled the right label context for full-sum training.

## II. TRAINING CRITERIA AND INFERENCE RULES

### A. Modeling Approach

The standard Bayes decision rule for automatic speech recognition for an input feature sequence  $X$  and an output word sequence  $W$  is defined in Eq. (1).

$$X \rightarrow \tilde{W}(X) = \operatorname{argmax}_W \{P_\theta(W|X)\} \quad (1)$$

$$= \operatorname{argmax}_W \{P_{\theta_{\text{am}}}(X|W) \cdot P_{\theta_{\text{lm}}}(W)\} \quad (2)$$

In classic generative approaches the acoustic model (AM) and the language model (LM) are treated as distinct components, as indicated in Eq. (2). In contrast, modern seq2seq models, commonly referred to as end-to-end approaches, eliminate this

separation by integrating the acoustic and language modeling into a single, unified framework, as in Eq. (1).

The training criterion that determines the optimal parameters  $\theta$  and  $\theta_{AM}$  for the acoustic models in the two approaches makes use of a frame-level probability distribution that either generates an input feature at time frame  $t$  or predicts a label conditioned on the input and, optionally, a label history. The input and output at each time frame are matched via an alignment or path through the lattice of all possible alignments between the two feature and label sequences. Given a phoneme sequence  $\phi$  corresponding to  $W$ , the marginalization over all alignments is done by using the hidden Markov sequence  $S$  and the blank augmented sequence  $Y$  both of length  $T$ . Denote  $h_1^T = Enc(X)$  as the neural representation of the input acoustic feature sequence  $X$  with downsampling by a factor of four for all our models. We focus only on *time synchronous* modeling approaches where independent of the intermediate factorization or use of auxiliary factors, the final learned probability distribution used in Bayes rule for decoding is only conditioned on the feature input. Namely, a first order label context model proposed as (diphone) factored hybrid [8], [16], and three zero order label context models: a connectionist temporal classification [1], a posterior HMM [17], and a hybrid NN-HMM trained with left and right auxiliary factors [13]. The label set across all topologies consists in single state phonemes from the pronunciation lexicon, with the end-of-word (EOW) distinction. We consider a single pronunciation for each word. For the CTC model, an additional blank token is used, while only the HMM model explicitly utilizes a silence token.

### B. Right Label Context in Full-Sum training

To discuss the integration of the right label context, we first perform the frame-level factorization under model specific assumptions for the full-sum training criterion. The frame-level factorization of CTC model can be seen in Eq. (3).

$$\sum_Y P(\phi, Y | h_1^T) = \sum_{y_1^T: \phi} \prod_{t=1}^T P(y_t | h_t) \quad (3)$$

The frame-wise posterior of CTC does not allow for the integration of any type of label context due to the independence assumption. Discriminative approaches such as RNN-T that do not drop the dependency on the label context rely on a factorization that yields merely a left label history, with arbitrary length, while excluding any right context. Access to the right context differs in the generative approach. For models following the HMM assumption, we first carry out the decomposition into the emission and transition probabilities as in Eq. (4a). In neural based HMMs, Eq. (4b) is derived by applying Bayes identity.

$$\sum_S P(h_1^T, S | \phi) = \sum_{s_1^T: \phi} \prod_{t=1}^T P(h_t | s_t, \phi) P(s_t | s_{t-1}) \quad (4a)$$

$$\sim \sum_{s_1^T: \phi} \prod_{t=1}^T \frac{P(s_t, \phi | h_t)}{P_{\text{prior}}(s_t, \phi)} P(s_t | s_{t-1}) \quad (4b)$$

The additional term  $P(h_t)$  is constant with respect to both the training criterion and the Bayes decision rule, and is therefore

omitted here.

In contrast to CTC, the emission probability in Eq. (4a) generates the input conditioned on the tuple  $\{s_t, \phi\}$ . The main emphasis of the analysis is the extent to which the HMM state  $s$  at time frame  $t$  has access to the sequence  $\phi$ . During training, the phoneme sequence  $\phi$  is fully available, since we have access to the transcript and the pronunciation lexicon. For decoding, when using a finite state acceptor, the structure of the graph can allow the distinction between alignment states of a certain center phoneme in different contexts following the context composition  $C$  in the HCLG composition [18]. We consider the augmentation of the phoneme aligned to the state  $s_t$  at time frame  $t$  with its right and left phonemes from the original phoneme sequence  $\phi$ . Let  $a_{s_t}$  be a position aware labeling function that given the hidden Markov alignment state  $s$  at time frame  $t$  provides the label of the aligned phoneme within  $\phi$ . For simplicity the access to the adjacent phoneme labels is done via  $a_{s_{t-1}}$  and  $a_{s_{t+1}}$ , for left and right phonemes, respectively. Explicit modeling of the left, center, and right phonemes follows the substitution of  $\{s_t, \phi\}$  with  $\{a_{s_{t-1}}, a_{s_t}, a_{s_{t+1}}\}$ , yielding Eq. (5).

$$\sum_S P(h_1^T, S | \phi) = \sum_{s_1^T: \phi} \prod_{t=1}^T \frac{P(a_{s_{t-1}}, a_{s_t}, a_{s_{t+1}} | h_t)}{P_{\text{prior}}(a_{s_{t-1}}, a_{s_t}, a_{s_{t+1}})} P(s_t | s_{t-1}) \quad (5)$$

### C. Decoding

The best word sequence for all approaches is obtained via Viterbi decoding and by using a scaled log linear formulation, using prior, transition, and language model (LM) scales  $\beta$ ,  $\eta$ , and  $\lambda$ , respectively. The decoding of all hidden Markov models mentioned in Sec. II-A relies on Eq. (6), except for the diphone model. This model uses  $P(a_{s_{t-1}}, a_{s_t} | h_t)$  and  $P(a_{s_{t-1}}, a_{s_t})$  for the label posterior and priors, respectively [8].

$$\arg \max_W [P_{\text{LM}}^\lambda(W) \max_{s_1^T: \phi} \prod_{t=1}^T \frac{P(a_{s_t} | h_t)}{P_{\text{prior}}^\beta(a_{s_t})} P^\eta(s_t, s_{t-1})] \quad (6)$$

A similar formulation is used also for CTC by omitting the transition probability [4].

### D. Right Label Context Contribution to the Gradient

Following the model definition of Eq. (5), the gradient for the optimization of the log likelihood criterion with respect to the acoustic model parameters  $\theta_{AM}$  is

$$\frac{\partial}{\partial \theta_{AM}} LL(\theta_{AM}) = \quad (7a)$$

$$\sum_{t,s} \left[ \gamma_t(a_{s-1}, a_s, a_{s+1} | h_1^T, \theta_{AM}) \frac{\partial}{\partial \theta_{AM}} \log \frac{P(a_{s-1}, a_s, a_{s+1} | h_t)}{P_{\text{prior}}(a_{s-1}, a_s, a_{s+1})} \right]$$

$$\approx \sum_t \sum_{a \in \{a_{s-1}, a_s, a_{s+1}\}} \left[ \gamma_t(a | h_1^T, \theta_{AM}) \frac{\partial}{\partial \theta_{AM}} \log \frac{P^{\alpha_a}(a | h_t)}{P_{\text{prior}}^{\beta_a}(a)} \right] \quad (7b)$$

Where  $\gamma$  defined in Eq. (8) guarantees the local normalization over the set of acoustic model labels at each time step and is calculated efficiently via forward-backward algorithm.

$$\gamma_t(a_{s_{t-1}}, a_{s_t}, a_{s_{t+1}} | h_1^T, \theta_{AM}) = \frac{\sum_{S: \phi, s_t=s} P(h_1^T, S | \phi, \theta_{AM})}{\sum_{S: \phi} P(h_1^T, S | \phi, \theta_{AM})} \quad (8)$$

For efficient implementation of Eq. (7a), we apply two additional simplifying assumptions motivated later in Sec. II-F: (1) elimination of the dependency to the label context for each factor, and (2) a state prior scale  $\beta$  equal to zero. The resulting

gradient after the application of the mentioned assumptions can be seen in Eq. (7b).

### E. Role of Scales in Training

The standard solution for training high performance HMM based systems is to use an external alignment, while the novel seq2seq approaches such as CTC use the full-sum criterion. The optimization based on full-sum criterion is performed as a form of expectation-maximization algorithm. The training can begin with randomly initialized parameters, eliminating the need for an alignment between the speech input and the output label sequences. There are several aspects in the choice of common CTC approaches that play an important role for the convergence of the model. This includes the choice of output label unit, the acoustic input downsampling factor, and the label topology [4], [19]. For the standard CTC label topology, the blank label assumes both the role of continuation of the duration of the last emitted label, and the role of silence. Due to this ambiguity, the use of blank in the CTC label topology allows the model to have more freedom in the choice of the output label start and end boundaries. Given a segment, the label can be emitted at any arbitrary position. This is not the case for the HMM topology, where we have a partition that starts with the emission of the label and ends after the last label loop. This distinction results also in different convergence behaviors in the two models. When training a model from scratch using full-sum training, it is desirable at the beginning that the model is not too confident with respect to the segmentation. Too much confidence leads to the concentration of the sequence-level probability mass into fewer alignment paths. Early in training, these paths tend to be of low quality due to the random initialization of parameters. The uncertainty introduced by the blank label plays a crucial role in ensuring convergence. The scaling of label posterior and transition probabilities within the HMM topology plays a comparable role. The scales regulate the smoothness of the distribution, influencing the model’s confidence in deciding whether to remain in the current label or transition to a different one [20]. Previous work has shown also the relation between these scales and the input frame shift [21]. Here, we argue that the contribution of the additional right and left factors in the gradient of Eq. (7b) have also a role on the convergence, due to the inclusion of the output labels information on a wider range.

### F. Simplifying Assumptions

#### 1) Context-independent factors

The frame-wise label posteriors in the modern ASR approaches are modeled using neural networks. In locally normalized models, the final softmax layer defines a probability distribution over the set of acoustic model labels. Historically, due to data sparsity issues, the acoustic label units associated with triphones were defined through parameter tying, clustering similar triphones according to a predefined criterion [22]. Even under the assumption that all triphones are well represented in the data, training a model with a softmax layer of size equal to the phoneme set raised to the power

of three, as the posterior used in Eq. (7a), is impractical. Recently, a factored hybrid HMM with no state-tying has been implemented and extensively evaluated [8], [15], [17], [21]. Factorization of the joint probability  $P(a_{s_t-1}, a_{s_t}, a_{s_t+1}|h_t)$  enables the original softmax layer parameters to be split into three distinct softmax layers, each normalized over the set of phonemes [8]. The triphone factored hybrid model is generally trained with frame-wise cross-entropy criterion with a fixed given alignment, known as Viterbi training. In order to calculate the original joint distribution, each of the network branches with context-dependent output needs to be forwarded with the given context embeddings and the output of all three softmax layers should be combined. The score computation for this model is generally more expensive in full-sum than in Viterbi. At a certain time frame the center phoneme of an alignment state could appear in different contexts on different alignment paths. Therefore, we propose a simplifying assumption for training that after the factorization of the joint probability in Eq. (7a) drops the dependency to the phoneme contexts for each factor. The resulting gradient is the weighted sum of each  $\partial \log$  probabilities of left, center, and right phonemes, where the weight for each  $\partial \log$  is obtained by marginalizing the original joint  $\gamma$  over the remaining phonemes, as shown in Eq. (7b). In our experiments, we observed that despite this simplifying assumption, given the input features, the three separate distributions for left, center, and right phonemes, learn the triphone structure implicitly.

#### 2) Omitting the Prior During Training

The scaled prior of Eq. (7b) can be estimated on-the-fly during the training with exponential decaying average [20] or as a fixed prior estimated on the transcriptions [4]. We observed convergence problems during training when using a prior and therefore considered the special case of prior scale  $\beta = 0$ . Note that for the monophone HMM model that uses only the center phoneme state, setting the prior scale to zero effectively results in a posterior HMM [4].

## III. EXPERIMENTAL RESULTS AND SETTING

### A. Setting

The experiments are carried out on 300h Switchboard-1 (SWB) Release 2 (LDC97S62) [23] and 960h LibriSpeech (LBS) [24]. We evaluate our models on SWB and Call-Home subsets of Hub5‘00 (LDC2002S09), the three subsets of Hub5‘01 (LDC2002S13), as well as LBS dev and test sets. The loss augmented with the right and left context is an extension of the CUDA based forward-backward within RETURNN [25]. Decoding of HMM based models use RASR for the core algorithms, and its recent extension for CTC decoding [26], [27]. Our experimental workflow is managed by Sisyphus [28]. The speech signal is extracted using a 25ms window with a 10ms shift, yielding Gammatone filterbank features with dimensions of 40 for SWB and 50 for LBS [29]. SpecAugment is applied across all models [30]. All encoder architectures consist of a 12-layer Conformer encoder with 75 million parameters [31], [32]. The experiments of Table I are trained for 50 and 25 epochs for SWB and LBS, respectively.

TABLE I: ASR performance models trained for 50 epochs on SWB300h and 25 epochs on LBS960h, using a 4-gram LM. We distinguish between the loss outputs used during training and the output used for decoding. The notation stands for a conditional probability of left, center, and right labels  $\ell$ ,  $c$ , and  $r$  given the encoder output  $h$ , respectively.

#	Model	Outputs		WER [%]			
		Loss	Decode	Hub5'00	Hub5'01	dev-other	test-other
1	CTC	$(c h)$	$(c h)$	12.8	11.8	6.9	7.4
2				12.8	12.2	6.9	7.3
3	HMM	$(\ell h)(c h)(r h)$	$(\ell, c h)$	<b>12.0</b>	<b>11.6</b>	<b>6.6</b>	7.1
4		$(\ell, c h)$		12.4	11.9	<b>6.6</b>	<b>7.0</b>

TABLE II: Similar experiments as in Table I with double amount of epochs. Two diphone Exps. 9 and 10 are initialized with the models from Exps. 2 and 3. The total number of epochs is the same as all other models in the table.

#	Model	Outputs		Init	WER [%]			
		Loss	Decode		Hub5'00	Hub5'01	dev-other	test-other
5	CTC	$(c h)$	$(c h)$	-	12.3	11.3	6.6	6.9
6					12.2	11.7	6.3	6.6
7	HMM	$(\ell h)(c h)(r h)$	$(\ell, c h)$	#2	<b>11.6</b>	11.0	6.1	6.5
8		11.8			11.0	<b>6.0</b>	<b>6.4</b>	
9	HMM	$(\ell, c h)$	$(\ell, c h)$	#3	11.8	10.9	6.3	6.7
10					<b>11.6</b>	<b>10.8</b>	6.2	6.8

The total number of epochs is doubled for the experiments in Table II, also in case of experiments that used the seed models of Table I for initialization. We use one cycle learning rate schedule (OCLR) up to peak LR of  $6e-4$  over 90% of the training epochs, followed by a linear decrease to  $1e-6$  [31], [33]. The Exp. 9 and 10 from Table II are initialized with the seed model and further trained with a constant learning rate of  $5e-5$  on 90% of the training epochs followed by a linear decrease to  $1e-6$ . An Adam optimizer with Nesterov momentum, together with optimizer epsilon of  $1e-8$  are used [34]. Decoding is performed using time-synchronous beam search based on dynamic programming principles with a lexical prefix tree. We use the official 4-gram LM for both tasks. For the real time factor measurement experiment, we used an AMD CPU (released 2021), with 2 logical cores. Example setups with more details are available.<sup>1</sup>

## B. Results

We start with the evaluation of a first set of models trained from scratch in Table I. All models are trained for the same number of epochs, i.e., 50 epochs for SWB 300h and 25 epochs for LBS 960h. It is possible to see that among the zero order label context models, the inclusion of the right context in Exp. 3 yields the best results for SWB task. The performance gap is smaller when using LBS. However, model of Exp. 3 still outperforms all other zero-order label context models. We then evaluate the effect of longer training and the model initialization in Table II. Here, we use the two models from Exps. 2 and 3 to initialize the diphone factored hybrid that is trained for additional 50 and 25 epochs for SWB and LBS, respectively. All models in Table II are trained for the

TABLE III: Comparisons of training time and real time factor (RTF) for two experiments from Table II on dev-other.

#	Outputs		Train[h]	RTF	WER [%] dev-other
	Loss	Decode			
7	$(\ell h)(c h)(r h)$	$(c h)$	<b>300</b>	<b>0.099</b>	6.1
8	$(\ell, c h)$	$(\ell, c h)$	330	0.118	<b>6.0</b>

TABLE IV: Diphone factored and monophone hybrid models on LBS 960h decoded with 4-gram LM. The external alignment (ext. align) is a from-scratch posterior HMM with BLSTM encoder trained for 25 epochs. The Viterbi training for the remaining 20 epochs relies on factored outputs without the simplifying assumption. The other two models use only full-sum and do not use any external alignment.

Ext. Align	Train			WER [%]	
	Loss	Decode	Epochs	dev-other	test-other
yes	$(\ell h)(c \ell, h)$	$(\ell, c h)$	45	6.1	6.8
no	$(\ell, c h)$	$(c h)$	50	<b>6.0</b>	<b>6.4</b>
	$(\ell h)(c h)(r h)$			6.1	6.5

same number of epochs, also when they are trained without initialization. The diphone model can take advantage of the pretrained model of Exp. 3 on SWB 300. This does not seem to be the case for LBS task where the best performance is obtained by the from-scratch diphone model. However, the model of Exp. 7 with the right context is also performing in the same range as the diphone of Exp. 8. In Table III, we take a closer look at the training time and real time factor of the monophone HMM with left and right auxiliary losses (Exp. 7) and diphone factored hybrid (Exp. 8). The model in Exp. 7 shows a 1.6% relative increase of WER, which is negligible considering the 10% relative speedup in training time and 16% relative improvement in RTF. Furthermore, comparison between the Exps. 1-4 of Table I and the Exps. 5-8 of Table II highlights the advantage of the inclusion of the right label context even when training for fewer epochs. Moreover, in all experiments we observed that HMM based models outperformed the CTC under similar conditions.

In our previous works, we proposed a Viterbi trained diphone factored hybrid using the alignment from a monophone posterior HMM, with comparable or superior performance to a strictly monotonic RNN-T [15], [21]. As a further step, in Table IV we show that our proposed end-to-end training can lead to similar ASR accuracy as the standard hybrid HMM pipeline, bringing our HMM based systems closer to the end-to-end approaches, in terms of simplicity.

## IV. CONCLUSIONS

We showed that the inclusion of right label context in training of time-synchronous ASR models lead to performance gain particularly when data resources are limited. We also proposed an end-to-end training solution for hybrid HMM systems eliminating the need for the standard complex multi-stage pipeline and the external alignments while still achieving comparable performance.

## V. ACKNOWLEDGMENTS

This work was partially supported by NeuroSys, which as part of the initiative ‘‘Clusters4Future’’ is funded by the Federal Ministry of Education and Research BMBF (03ZU1106DA). The authors thank Albert Zeyer for the invaluable discussions and Eugen Beck for their unwavering support of the factored hybrid concept.

<sup>1</sup> <https://github.com/rwth-ifa/returnn-experiments/tree/master/2025-factored-fullsum-rightcontext>

## REFERENCES

- [1] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proc. ICML*, 2006, pp. 369–376.
- [2] A. Graves, "Sequence transduction with recurrent neural networks," in *Proc. ICML*, Edinburgh, Scotland, Jun. 2012, Workshop on Representation Learning, arXiv:1211.3711.
- [3] A. Tripathi, H. Lu, H. Sak, and H. Soltau, "Monotonic recurrent neural network transducer and decoding strategies," in *Proc. IEEE ASRU*, 2019, pp. 944–948.
- [4] T. Raissi, W. Zhou, S. Berger, R. Schlüter, and H. Ney, "HMM vs. CTC for Automatic Speech Recognition: Comparison Based on Full-Sum Training from Scratch," in *Proc. IEEE SLT*, 2023.
- [5] H. A. Bourlard and N. Morgan, *Connectionist Speech Recognition: a Hybrid Approach*. Norwell, MA: Kluwer Academic Publishers, 1993.
- [6] R. Sennrich, B. Haddow, and A. Birch, "Neural machine translation of rare words with subword units," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, K. Erk and N. A. Smith, Eds. Berlin, Germany: Association for Computational Linguistics, Aug. 2016, pp. 1715–1725. [Online]. Available: <https://aclanthology.org/P16-1162>
- [7] P. Bell, P. Swietojanski, and S. Renals, "Multitask learning of context-dependent targets in deep neural network acoustic models," *IEEE/ACM Transactions on Audio, Speech, and Language Process.*, vol. 25, no. 2, pp. 238–247, 2016.
- [8] T. Raissi, E. Beck, R. Schlüter, and H. Ney, "Context-dependent acoustic modeling without explicit phone clustering," in *Proc. IEEE ICASSP*, 2020.
- [9] M. Ghodsi, X. Liu, J. Apfel, R. Cabrera, and E. Weinstein, "RNN-transducer with stateless prediction network," in *Proc. IEEE ICASSP*, 2020, pp. 7049–7053.
- [10] Z. Yao, L. Guo, X. Yang, W. Kang, F. Kuang, Y. Yang, Z. Jin, L. Lin, and D. Povey, "Zipformer: A faster and better encoder for automatic speech recognition," in *The Twelfth International Conference on Learning Representations*, 2023.
- [11] W. Zhou, S. Berger, R. Schlüter, and H. Ney, "Phoneme based neural transducer for large vocabulary speech recognition," in *Proc. IEEE ICASSP*, Jun. 2021, pp. 5644–5648.
- [12] E. Variani, K. Wu, M. D. Riley, D. Rybach, M. Shannon, and C. Al-lauzen, "Global normalization for streaming speech recognition in a modular framework," *Proc. NeurIPS*, vol. 35, pp. 4257–4269, 2022.
- [13] T. Raissi, E. Beck, R. Schlüter, and H. Ney, "Towards Consistent Hybrid HMM Acoustic Modeling," Apr. 2021, arXiv:2104.02387.
- [14] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-Augmented Transformer for Speech Recognition," in *Proc. Interspeech*, Shanghai, China, Oct. 2020, pp. 5036–5040.
- [15] T. Raissi, C. Lüscher, M. Gunz, R. Schlüter, and H. Ney, "Competitive and resource efficient factored hybrid HMM systems are simpler than you think," *Proc. Interspeech*, 2023.
- [16] H. Bourlard, N. Morgan, C. Wooters, and S. Renals, "CDNN: A context dependent neural network for continuous speech recognition," in *Proc. IEEE ICASSP*, 1992.
- [17] T. Raissi, W. Zhou, S. Berger, R. Schlüter, and H. Ney, "HMM vs. CTC for automatic speech recognition: Comparison based on full-sum training from scratch," 2022.
- [18] M. Mohri, F. Pereira, and M. Riley, "Weighted finite-state transducers in speech recognition," *Computer Speech & Language*, vol. 16, no. 1, pp. 69–88, 2002.
- [19] Z. Zhao and P. Bell, "Regarding topology and variant frame rates for differentiable WFST-based end-to-end asr," in *Proc. Interspeech*, 2023.
- [20] A. Zeyer, E. Beck, R. Schlüter, and H. Ney, "CTC in the context of generalized full-sum HMM training," in *Proc. Interspeech*, 2017, pp. 944–948.
- [21] T. Raissi, C. Lüscher, S. Berger, R. Schlüter, and H. Ney, "Investigating the effect of label topology and training criterion on asr performance and alignment quality," *Proc. Interspeech*, 2024.
- [22] S. J. Young, J. J. Odell, and P. C. Woodland, "Tree-based state tying for high accuracy acoustic modelling," in *Proc. ACL*, 1994, pp. 307–312.
- [23] J. Godfrey, E. Holliman, and J. McDaniel, "Switchboard: Telephone Speech Corpus for Research and Development," in *Proc. IEEE ICASSP*, 1992.
- [24] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "LibriSpeech: An ASR Corpus Based on Public Domain Audio Books," in *Proc. IEEE ICASSP*, 2015.
- [25] P. Doetsch, A. Zeyer, P. Voigtlaender, I. Kulikov, R. Schlüter, and H. Ney, "RETURNN: The RWTH extensible training framework for universal recurrent neural networks," in *Proc. IEEE ICASSP*, 2017, pp. 5345–5349.
- [26] D. Rybach, S. Hahn, P. Lehnen, D. Nolden, M. Sundermeyer, Z. Tüske, S. Wiesler, R. Schlüter, and H. Ney, "RASR-the RWTH Aachen university open source speech recognition toolkit," in *Proc. IEEE automatic speech recognition and understanding workshop*, 2011.
- [27] W. Zhou, E. Beck, S. Berger, R. Schlüter, and H. Ney, "RASR2: The RWTH ASR toolkit for generic sequence-to-sequence speech recognition," 2023.
- [28] J.-T. Peter, E. Beck, and H. Ney, "Sisyphus, a workflow manager designed for machine translation and automatic speech recognition," in *Proc. EMNLP*, 2018, pp. 84–89.
- [29] R. Schlüter, I. Bezrukov, H. Wagner, and H. Ney, "Gammatone features and feature combination for large vocabulary speech recognition," in *Proc. IEEE ICASSP*, 2007.
- [30] D. S. Park, Y. Zhang, C.-C. Chiu, Y. Chen, B. Li, W. Chan, Q. V. Le, and Y. Wu, "SpecAugment on Large Scale Datasets," in *Proc. IEEE ICASSP*, Brighton, UK, May 2019, pp. 6879–6883.
- [31] W. Zhou, W. Michel, R. Schlüter, and H. Ney, "Efficient Training of Neural Transducer for Speech Recognition," in *Proc. Interspeech*, Sep. 2022, arXiv:2204.10586.
- [32] A. Gulati, J. Qin, C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented Transformer for Speech Recognition," in *Proc. Interspeech*, 2020.
- [33] L. N. Smith and T. Nicholay, "Super-convergence: Very fast training of neural networks using large learning rates," in *Artificial intelligence and machine learning for multi-domain operations applications*, 2019.
- [34] T. Dozat, "Incorporating Nesterov Momentum into Adam," in *Proc. ICLR*, 2016.