



Analyzing the Importance of Blank for CTC-Based Knowledge Distillation

Benedikt Hilmes^{1,2}, Nick Rossenbach^{1,2}, Ralf Schlüter^{1,2}

¹RWTH Aachen University, Aachen, Germany

²Apptek GmbH, Aachen, Germany

<lastname>@ml.rwth-aachen.de

Abstract

With the rise of large pre-trained foundation models for automatic speech recognition new challenges appear. While the performance of these models is good, runtime and cost of inference increases. One approach to make use of their strength while retaining efficiency is to distill their knowledge to smaller models during training. In this work, we explore different CTC-based distillation variants, focusing on blank token handling. We show that common approaches like blank elimination do not always work off the shelf. We explore new blank selection patterns as a potential sweet spot between standard knowledge distillation and blank elimination mechanisms. Through the introduction of a symmetric selection method, we are able to remove the CTC loss during knowledge distillation with minimal to no performance degradation. With this, we make the training independent from target labels, potentially allowing for distillation on untranscribed audio data.

Index Terms: speech recognition, knowledge distillation, blank distribution, foundation models

1. Introduction

Recently, foundation models emerged as the new state-of-the-art in automatic speech recognition (ASR). These models are trained on large amounts of data either in a supervised or unsupervised fashion. Examples for these large pre-trained models are Whisper [1], Hubert [2] or Wav2Vec [3]. While the training methods differ, they perform well on a number of downstream tasks. When deploying models to mobile or edge devices, resources are limited and inference speed as well as memory footprint are a major limitation of foundation models. Still, it is desirable to use the capabilities provided by these models, as their performance is often superior to a smaller and directly trained model. One way to do that is through knowledge distillation (KD), where a larger teacher model serves to improve a smaller student model through training the model both on the given data and the output distribution of the teacher model.

There has been quite some work done in order to tackle the challenges of distilling the knowledge of connectionist temporal classification (CTC)-based ASR models. KD was initially proposed in [4] as a general method to transfer information from one neural network to another. One of the problems of applying KD to CTC is the amount of blanks in the predictions of the models. To deal with this, [5] and [6] propose sequence level KD, as an alternative to frame level KD. In [7], the student is trained in a way that the CTC spike timings align with those of the teacher. This alignment problem is taken up in [8], which proposed to counteract the mismatch in alignment between teacher and student distributions through distilling over frame areas. [9] introduces a two-stage training, where first the

features of the student are trained to match those of the teacher. In a second step, they calculate the distillation loss by combining a Kullback-Leibler divergence (KLD) loss on the non-blank positions of the teacher with a KLD loss tying the student and teacher features together. In [10], blank and non-blank distributions are factorized separately and combined in a three stage process: 1. Binary distinction between blank and non-blank; 2. Balanced distillation on blank and non-blank; 3. Learning a posterior matrix over a factorized KD loss.

For architectures like RNN-T, distilling the model on the teachers outputs only is more common [11, 12]. Most CTC-based approaches combine the KD objective with the CTC objective [9, 10]. This is why as a baseline for this work we use a setup that includes both the CTC and KD objective, and confirm that for base KD the CTC loss is helpful. A potentially undesired side-effect of this combination is the dependence on labeled data for the CTC loss. While one solution would be to have the teacher generate pseudo labels as targets, we refrain from this approach, as this has been shown not to be helpful when doing distillation [12]. Usually, work on CTC-based KD starts with a medium sized teacher to distill the information to a small student, resulting in maximum size reduction for the final model [5, 6, 9, 10]. In this work we focus on keeping the student medium sized and instead try to deploy a larger foundation model as teacher, pre-trained on unlabeled data, in order to provide maximum information to the student. We do this with the goal of having an efficiently deployable model with performance enhanced by a large foundation model teacher.

1.1. Contribution

Our contributions are as follows: We start from investigating the behavior of blank elimination [9] for CTC-based KD. We then introduce different extensions with the goal of using the information contained in the blank predictions of our teacher. We analyze these approaches in training settings with differently scaled distillation losses. One of these approaches, symmetric blank selection, includes only the neighboring positions of non-blank positions. We found that with this particular subset of blanks we can remove the CTC objective from the distillation process without losing performance. This is in contrast to the standard blank elimination and other newly introduced approaches, where the performance can not be maintained without the CTC component. With a distillation process using symmetric blank selection and dropping the CTC loss, it becomes possible to include unsupervised data in the distillation process. We run experiments on two English corpora, with baseline models close to state-of-the-art by using HuBERT as a large foundation model as teacher. Our work and software [13, 14] is public.¹

¹<https://github.com/rwth-i6/returnn-experiments/tree/master/2025-ctc-blank-kd>

2. Connectionist Temporal Classification

In this work we consider ASR models using connectionist temporal classification (CTC) [15]. CTC is used to align an acoustic signal x_1^T and a label sequence a_1^S through the introduction of a blank label, which extends a_1^S to length T . An alignment y_1^T is a valid alignment for x_1^T and a_1^S iff $B(y_1^T) = a_1^S$, where B is a function that removes all blanks and consecutive label repetitions from the sequence. A model can now be trained with the CTC criterion by marginalizing over all possible alignments and minimizing the following objective:

$$\mathcal{L}_{\text{CTC}} = -\log p(a_1^S | x_1^T) = -\log \sum_{y_1^T: B(y_1^T) = a_1^S} p(y_1^T | x_1^T) \quad (1)$$

It is a well known observation that CTC produces alignments which are dominated by the blank label in what is called a “peaky behavior” [16, 17, 18, 19]. This means that for the frame-wise predictions of a model trained with CTC criterion, a majority of the probability mass lies in the position of the blank symbol. While this is not hurtful during recognition, when distilling knowledge from a teacher to a student, some problems may occur, as will be discussed in the following section.

3. Knowledge Distillation

Knowledge distillation (KD) describes the process of transferring the information embedded in a so called teacher model into a (usually smaller) student model [4]. There are multiple variations of KD, which range from pseudo-labeling unsupervised data via the teacher, to training the student to have a similar output distribution as the teacher. Depending on the architecture of the student and teacher model different approaches have shown to be the most effective.

A major distinction is if the model output simultaneously serves as label-to-acoustic alignment, as is the case for CTC and RNN-T-based models. For such models, minimizing the frame-level cross entropy error between the posterior distributions of the teacher and the student model is difficult [5]. Instead, in this work we focus on distilling the knowledge of the teacher to the student by minimizing the KLD of the posterior distributions. [4, 9, 10].

The criterion for a given sequence can be formalized as:

$$\mathcal{L}_{\text{KD}} = \sum_{t=1}^T \sum_{c=1}^C p_{\text{tea}}(c | x_1^T, t) \cdot \frac{p_{\text{tea}}(c | x_1^T, t)}{p_{\text{stu}}(c | x_1^T, t)} \quad (2)$$

with given acoustic signal x_1^T of length T , C classes including blank, teacher distribution p_{tea} and student distribution p_{stu} .

3.1. Blank Elimination

The calculation of Equation (2) includes both blank and non-blank positions of the teacher distribution. As stated in Section 2, CTC tends to produce peaky alignments where a majority of the positions are blank. Consequently, the influence of blank positions seen during training proportionally surpasses that of non-blank positions. When viewed on a per-class comparison, this effect increases by another magnitude. This means that making use of the KD criterion without further modifications hinders the convergence of the distilled model and hurts performance, which was already observed in previous works [5, 20, 21]. To deal with this problem [9] proposes to calculate the KLD only on the non-blank positions of the sequence which we are going to call *blank elimination* in the following. These positions are calculated in a greedy fashion by looking at

the posterior distribution of the teacher and selecting the positions where a non-blank token has the highest probability. The modified KD criterion can be described as:

$$\mathcal{L}_{\text{KD-BE}} = \sum_{t=1}^T \sum_{c=1}^C p_{\text{tea}}(c | x_1^T, t) \cdot \frac{p_{\text{tea}}(c | x_1^T, t)}{p_{\text{stu}}(c | x_1^T, t)} \cdot (1 - \delta(\langle \mathbf{b} \rangle, \arg \max_{c'} \{p_{\text{tea}}(c' | x_1^T, t)\})) \quad (3)$$

where δ describes the Kronecker delta and $\langle \mathbf{b} \rangle$ is the index of the blank token.

3.2. Blank Selection Mechanisms

Learning a proper blank distribution is a crucial part of training and decoding a model with the CTC criterion. Since the gradient of the KD objective for the case of blank elimination is only computed on non-blank positions, an additional CTC loss is used in prior work to update positions where the teacher does predict blank. This means that without further modification training requires supervised data and a larger portion of the teacher outputs is not considered for KD. As shown in [9], distilling knowledge over all blanks is hurtful for the convergence of the model. Nevertheless, completely removing them from the distillation means that the teacher’s distributions at blank positions, and consequently parts of the alignment, are not propagated to the student. Thus, in the following we explain novel approaches to consider only specific blank positions.

3.2.1. Non-blank dependent selection

For our first approach, we propose to extend the blank elimination symmetrically around the non-blank positions of the teacher output. This means that for every frame where the argmax of the teacher is not the blank symbol, the surrounding n frames will also be considered for the calculation of the KLD. This way, the blank information at the boundary is distilled, while at positions with larger amounts of sequential blank positions the predictions are left out.

As we intentionally only select positions around non-blank predictions of the teacher, we do not check if the selected position would already be included due to being non-blank. While in theory the amount of blanks selected is upper bounded by $2n$, the amount of actually selected blank positions is less, which we are going to analyze in Section 5.2.

The amount of blanks included into the distillation is a tunable parameter, which can differ for different corpora and models. For our experiments we test a value of n ranging from 1 to 5. Since we hypothesize that within-sequence blanks might be more important than blank predictions outside of the sequence, we add a chain of experiments where we only trim the blanks before the first and after the last non blank prediction of the teacher from the KLD calculation.

3.2.2. Probability-based selection

Instead of relying on our non-blank positions to select which blank positions we include during training, we do a selection based on the probability distributions of the blank positions. For this, we consider a threshold α and include all positions into our KD process for which the probability of blank is below α . This way, we aim to select positions where the teacher is less sure about its blank prediction. Thus, a potential transition into non-blank occurs, which might be an important information for the student to learn. We try the values 0.95, 0.9 and 0.8 for α .

As additional control experiment, we include a random sampling approach in our experiments, where the blank posi-

tions seen during KD are selected randomly. Here we use a scale β which decides how many positions proportionally to the non-blank positions are selected, where $\beta = 1$ means that we randomly select as many positions as there are non-blank positions in the teacher output. For β we test 0.5, 1.0 and 2.0.

4. Experimental Setup

4.1. Data

In this work we make use of the two English datasets TEDLIUMv2 (TEDv2) [22] and LibriSpeech (LBS) [23]. TEDv2 consists of 207 hours of TED talks, while LBS offers 960 hours of audio book recordings for training. For LibriSpeech we report results on *dev-other* and *test-other*. We split the data into sub-epochs of 5 for TEDv2 and 10 for LBS. We use phonemes as targets, with the phoneme set consisting of ARPABET phoneme symbols without stress marker. In order to predict phoneme sequences missing for words in the lexicon during training, we use Sequitur [24]. For recognition, we use a 4-gram language model (LM) trained on the corresponding text-only data provided with each of the corpora.

4.2. Teacher

As a teacher for this work, we use the publicly available *large* variant of the HuBERT model² (317M parameters) [2], only trained on Libri-Light [25]. In order to adjust the model to the task, we add a linear layer at the end and fine-tune the whole model to each corpus for 10 full epochs. Using the publicly available LBS fine-tuned version of the HuBERT model did not yield a performance difference, so we stick to the unsupervised model. We evaluate each full epoch and take the checkpoint with the best dev word error rate (WER) as teacher checkpoint. We down-sample the outputs by another factor of two in order to match the time resolution of the student. We found that this does not result in a worse WER for the teacher. For knowledge distillation, we disable dropout and other data augmentation techniques, as this improves the KD [26].

4.3. Training

Our student model is a 12-layer Conformer [27] with relative positional encodings [28]. The hidden dimension of our student is 384, going up to 1536 for the feed-forward module. In total the model consists of around 42M parameters. We use 80-dimensional log-mel features with a frame shift of 10ms as input to our network, augmented by SpecAugment [29]. Our frontend consisting of a stack of convolutions downsamples the features by a factor of 4, so that each output label corresponds to 40ms. We use a dropout [30] probability of 0.2 for TEDv2 and 0.1 for LBS. We apply the AdamW [31] optimizer with a weight decay scale of 0.01. We use a linear one cycle learning rate scheduling starting from $7e-6$ peaking at $5e-4$ after 235 sub-epochs, going down to $5e-5$ and finishing with a cooldown of 30 sub-epochs down to $1e-7$. We train our models for 500 sub-epochs, so 50 full epochs for TEDv2 and 100 epochs LBS. For TEDv2 we use a batch size of 180 seconds of audio, while for LBS we use 300 seconds. We train our baseline with CTC loss and use Flashlight [32] for decoding, using the last checkpoint for recognition. For training our models with knowledge distillation, the loss consists of a KLD loss with the teacher’s posterior distribution as target, and an optional (for $\lambda < 1$) interpolation with a CTC loss with the phoneme labels as targets:

$$\mathcal{L} = \lambda \mathcal{L}_{KD} + (1 - \lambda) \mathcal{L}_{CTC} \quad (4)$$

²<https://huggingface.co/facebook/hubert-large-ll60k>

Table 1: *Baseline performance of TEDv2 and LBS with 4-gram LM. For LibriSpeech we evaluate on the other portions of the evaluation data. A distillation scale of 1.0 means no CTC loss is used.*

| Model | Distillation Scale | WER [%] | | | |
|-------------|--------------------|------------|------------|------------|------------|
| | | TEDv2 | | LBS | |
| | | dev | test | dev | test |
| Teacher | - | 4.8 | 5.4 | 4.0 | 4.0 |
| Baseline | - | 6.4 | 7.0 | 6.0 | 6.6 |
| KD | 0.25 | 6.6 | 7.1 | 5.6 | 6.1 |
| | 0.9 | 6.0 | 6.7 | 5.8 | 6.0 |
| | 1.0 | 6.3 | 6.6 | 5.9 | 6.4 |
| Blank Elim. | 0.25 | 6.1 | 6.9 | 5.7 | 6.1 |
| | 0.9 | 5.8 | 6.5 | 6.1 | 6.5 |
| | 1.0 | 6.4 | 6.8 | 6.9 | 7.3 |

In initial experiments, we tried a number of different scales between 0 and 1. We found that 0.25, 0.9 and 1.0 give the best impression on the behavior of different KD conditions, exhibiting the best results. This is why we put our focus on these three settings: 1. Emphasis on CTC loss (0.25); 2. Emphasis on KD loss (0.9); 3. No CTC loss (1.0). All experiments can be run on a single consumer GPU with 24gb VRAM (e.g. Nvidia RTX 3090), resulting in a low barrier for reproduction.

5. Experiments

5.1. Baselines

The baseline performance of our models can be seen in Table 1. Comparing the teacher with our baseline trained without any KD, we can see that the teacher outperforms the baseline by around 30% WER relative across multiple corpora and test sets. This is to be expected, as the teacher not only was trained on a lot more data, but also contains around 8 times the number of parameters compared to the baseline.


When adding KD to train the student first differences become visible. In general, both models benefit from the additional information given by the teacher, but different distillation scales seem to be optimal. For TEDv2 a higher scale (0.9 or 1.0) produces better results than the lower scale (0.25) which is only on par with the baseline. For LBS a lower scale seems to perform better, where the best result of 5.6% on dev-other is achieved with a scale of 0.25. Nevertheless, the two LBS models with higher scales also outperform the baseline.

When removing blank positions from the KD objective, a major difference is visible. Blank elimination helps in the case of TEDv2 where the WER goes down to 5.8% on the dev set, but for the models trained on LBS performance degrades. Especially for the case where the distillation scale is 1.0 this difference is strongly visible for LBS dev-other, going up from 5.9% to 6.9%. Since a distillation scale of 1.0 effectively disables the CTC objective, this hints that the CTC objective for the case of blank elimination is able to counteract losing teacher information about the blank positions. Looking at TEDv2, this difference does not seem to be important as much, meaning that the blank positions for TEDv2 are not as important as for LBS. We also experimented with prior correction, where the blank prior is removed from the teacher distribution, but results were consistently worse.

5.2. Blank Aware Approaches

As a first study we present the results of symmetric selection evaluated on the TEDv2 dev set over the three distillation scales

Figure 1: WER [%] of different symmetric selection hyperparameters for TEDv2 dev with 4-gram LM.



0.25, 0.9 and 1.0, and the symmetric selection parameter range of 1 to 5. As visible in Figure 1, for a distillation scale of 0.25 there is a marginal improvement for the case of $n = 2$, but otherwise the performance degrades up to the baseline KD performance. For a distillation scale of 0.9 the performance stagnates for lower keep values and reaches the baseline KD performance at $n = 3$. Interesting is the case of a scale of 1.0, where both the blank elimination and the baseline knowledge distillation perform worse than the best value of 5.6 % for the case of $n = 2$. This means that the model is able to benefit from the additional but limited amount of blanks included during distillation, producing a result that outperforms the other scales. Not only does the introduction of symmetric blank selection improve the WER, but additionally the dependence on any labels is dropped, as the model is only optimized on the teacher outputs.

When analyzing the distributions, we find that on average 54% of a sequence is determined to be non-blank by the teacher. Applying symmetric selection these values increase to 66%, 69%, 72%, 74% and 76% ($1 \leq n \leq 5$). This means that in the largest case a quarter of the distilled positions is discarded, while overall three quarters of the sequence are considered for distillation. Consequently, a quarter of the sequence consists of larger areas of blank positions, and they were not selected.

We combine these results with the other previously introduced approaches and extend the analysis to LBS in Table 2. For all cases where the modification of KD has a hyperparameter, we present the best value for the corresponding set in the table, even though in all cases the ranges as stated in Section 3.2 where tested.

Symmetric selection on LBS is not able to outperform the simple KD baseline, but overall stabilizes blank elimination for the distillation scales 0.9 and 1.0. The WER of a distillation scale of 1.0 is even slightly better than for KD, meaning that without labels keeping some blanks in a symmetric fashion is helpful. Trimming blanks before and after the sequence does not help for TEDv2, but produces a marginally better result than the KD baseline for LBS. Selecting blanks based on a threshold or in a random fashion does not improve the distillation for TEDv2. For LBS we see a similar trend, with the exception of a distillation scale of 0.25. Here the models perform similar to the KD baseline. From this we conclude that for LBS and a distillation scale of 0.25 the position of the blank information is not important for the distillation, but rather a certain amount of blanks need to be seen during training for a good performance.

Overall we can conclude that when enabling symmetric selection during KD, the dependence on the labeled data can be dropped and the model can be trained only on the teachers output, as visible in the results with a distillation scale of 1.0. For TEDv2 this symmetric selection without CTC loss produced the

Table 2: Performance of blank selection mechanisms TEDv2 and LBS with 4-gram LM. For LBS we evaluate on the other portions of the evaluation data. In the first two rows \dagger represents the distillation scale for TEDv2 and $*$ the scale for LBS.

| Blank Select. | KD Scale | Hyper-param. | TEDv2 | | Hyper-param. | LBS | |
|---------------|---------------|----------------|------------|------------|----------------|------------|------------|
| | | | WER [%] | | | WER [%] | |
| | | | dev | test | | dev | test |
| - | 0.9 \dagger | - | 6.0 | 6.7 | - | 5.6 | 6.1 |
| Bl. Elim. | / 0.25* | - | 5.8 | 6.5 | - | 5.7 | 6.1 |
| Sym. | 0.25 | $n = 2$ | 6.0 | 6.8 | $n = 4$ | 5.7 | 6.0 |
| | 0.9 | $n = 1$ | 5.8 | 6.3 | $n = 3$ | 5.8 | 6.2 |
| | 1.0 | $n = 2$ | 5.6 | 6.3 | - | 5.8 | 6.2 |
| Trim | 0.25 | - | 6.5 | 6.8 | - | 5.6 | 5.9 |
| | 0.9 | - | 6.1 | 6.5 | - | 5.9 | 6.3 |
| | 1.0 | - | 6.2 | 6.6 | - | 6.0 | 6.4 |
| Thresh. | 0.25 | $\alpha = 0.9$ | 6.1 | 7.0 | $\alpha = 0.8$ | 5.6 | 6.1 |
| | 0.9 | $\alpha = 0.8$ | 5.9 | 6.4 | $\alpha = 0.9$ | 5.9 | 6.4 |
| | 1.0 | $\alpha = 0.9$ | 6.0 | 6.5 | - | 6.2 | 6.6 |
| Random | 0.25 | $\beta = 1.0$ | 6.3 | 7.0 | $\beta = 0.5$ | 5.6 | 6.1 |
| | 0.9 | $\beta = 1.0$ | 6.0 | 6.6 | $\beta = 2.0$ | 6.0 | 6.2 |
| | 1.0 | $\beta = 0.5$ | 6.1 | 6.4 | - | 6.0 | 6.4 |

best results, while for LBS it helped reduce the gap to the results of training with a distillation scale of 0.25.

6. Limitation and Future Work

One of the caveats of most of our blank aware approaches to improving CTC-based KD is the introduction of an additional hyperparameter to tune, as our experiments show different settings are optimal for different corpora and distillation scales. This might be due to the fact that blank distributions induced by CTC training might depend a lot on the data, as blank serves both as a wait and a silence token. Consequently, corpora with different amounts of silence will have largely different blank distributions. Future work on this topic should aim to remove additional hyperparameters by introducing a mechanism which decides the blanks to be kept in a more automatic fashion. Thresholding is a step in that direction, but still has a different (although easier to tune) parameter. Furthermore, it should be investigated how to deal with the different blank amounts by potentially taking the induced blank distribution of the corpus into account before blank selection. This way different corpora with different blank amounts could behave in a more similar fashion reducing tuning efforts.

7. Conclusion

In this work we analyzed the role of blank symbols during knowledge distillation for CTC-based ASR. For this we used a large pre-trained teacher to distill knowledge to a medium sized student for two different English corpora. We show that baseline extensions as blank elimination do not always improve the model performance and a dependence on the target labels is important to produce good results. We analyzed different new ways of selecting blanks for distillation, namely trimming, thresholding and random selection. These methods showed no improvements over the baseline on LibriSpeech, while degrading on TED-LIUMv2. Lastly, we introduced symmetric selection, selecting blank positions close to the non blank positions during distillation, while ignoring further away blanks. With a distillation scale of 1.0 and symmetric blank selection, we outperform the blank elimination method on TED-LIUMv2 and close the gap to knowledge distillation with a distillation scale of 0.25 on LibriSpeech. Through this, symmetric selection enables training on unsupervised data, as a CTC loss is no longer required for good performance.

8. 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 (funding IDs 03ZU2106DA and 03ZU2106DD), and by the project RESCALE within the program *AI Lighthouse Projects for the Environment, Climate, Nature and Resources* funded by the Federal Ministry for the Environment, Nature Conservation, Nuclear Safety and Consumer Protection (BMUV), funding ID: 67KI32006A.

9. References

- [1] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” in *International Conference on Machine Learning*, 2022.
- [2] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. rahman Mohamed, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [3] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: a framework for self-supervised learning of speech representations,” in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, ser. NIPS ’20. Red Hook, NY, USA: Curran Associates Inc., 2020.
- [4] G. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” in *NIPS Deep Learning and Representation Learning Workshop 2014*, 2014.
- [5] R. Takashima, S. Li, and H. Kawai, “An investigation of a knowledge distillation method for CTC acoustic models,” *ICASSP 2018*, pp. 5809–5813, 2018.
- [6] M. Huang, Y. You, Z. Chen, Y. Qian, and K. Yu, “Knowledge distillation for sequence model,” in *Interspeech 2018*, 2018, pp. 3703–3707.
- [7] G. Kurata and K. Audhkhasi, “Guiding CTC posterior spike timings for improved posterior fusion and knowledge distillation,” in *Interspeech*, 2019.
- [8] —, “Improved knowledge distillation from bi-directional to uni-directional LSTM CTC for end-to-end speech recognition,” in *2018 IEEE Spoken Language Technology Workshop (SLT)*, 2018, pp. 411–417.
- [9] S. Tian, K. Deng, Z. Li, L. Ye, G. Cheng, T. Li, and Y. Yan, “Knowledge distillation for CTC-based speech recognition via consistent acoustic representation learning,” in *Interspeech*, 2022.
- [10] S. Tian, Z. Li, Z. Lyv, G. Cheng, Q. Xiao, T. Li, and Q. Zhao, “Factorized and progressive knowledge distillation for CTC-based asr models,” *Speech Communication*, vol. 160, p. 103071, 2024.
- [11] S. Panchapagesan, D. S. Park, C.-C. Chiu, Y. Shangguan, Q. Liang, and A. Gruenstein, “Efficient knowledge distillation for RNN-Transducer models,” *ICASSP 2021*, pp. 5639–5643, 2021.
- [12] M. Zeineldeen, K. Audhkhasi, M. K. Baskar, and B. Ramabhadran, “Robust knowledge distillation from RNN-T models with noisy training labels using full-sum loss,” in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Rhodes, Greece, Jun. 2023.
- [13] J. Peter, E. Beck, and H. Ney, “Sisyphus, a workflow manager designed for machine translation and automatic speech recognition,” in *EMNLP 2018: System Demonstrations, Brussels, Belgium, October 31 - November 4, 2018*, pp. 84–89.
- [14] 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 *ICASSP 2017, New Orleans, LA, USA, March 5-9, 2017*, pp. 5345–5349.
- [15] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *Proceedings of the 23rd International Conference on Machine Learning*, ser. ICML ’06. New York, NY, USA: Association for Computing Machinery, 2006, p. 369–376.
- [16] A. Zeyer, R. Schlüter, and H. Ney, “Why does CTC result in peaky behavior?” Preprint arXiv:2105.14849, May 2021.
- [17] H. Liu, S. Jin, and C. Zhang, “Connectionist temporal classification with maximum entropy regularization,” in *Neural Information Processing Systems*, 2018.
- [18] H. Li and W. Wang, “Reinterpreting CTC training as iterative fitting,” *Pattern Recognition*, vol. 105, p. 107392, 2020.
- [19] R. Huang, X. Zhang, Z. Ni, L. Sun, M. Hira, J. Hwang, V. Manohar, V. Pratap, M. Wiesner, S. Watanabe, D. Povey, and S. Khudanpur, “Less peaky and more accurate CTC forced alignment by label priors,” *ICASSP 2024*, pp. 11 831–11 835, 2024.
- [20] R. Takashima, S. Li, and H. Kawai, “Investigation of sequence-level knowledge distillation methods for CTC acoustic models,” *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6156–6160, 2019.
- [21] A. W. Senior, H. Sak, F. de Chaumont Quiry, T. N. Sainath, and K. Rao, “Acoustic modelling with CD-CTC-SMBR LSTM RNNs,” *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, pp. 604–609, 2015.
- [22] A. Rousseau, P. Deléglise, and Y. Estève, “Enhancing the tedlium corpus with selected data for language modeling and more ted talks,” in *International Conference on Language Resources and Evaluation*, 2014.
- [23] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An ASR corpus based on public domain audio books,” in *ICASSP 2015*, 2015, pp. 5206–5210.
- [24] M. Bisani and H. Ney, “Joint-sequence models for grapheme-to-phoneme conversion,” *Speech Communication*, vol. 50, no. 5, pp. 434–451, May 2008.
- [25] J. Kahn, M. Rivière, W. Zheng, E. Kharitonov, Q. Xu, P.-E. Mazar’e, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen, T. Likhomanenko, G. Synnaeve, A. Joulin, A. rahman Mohamed, and E. Dupoux, “Libri-light: A benchmark for ASR with limited or no supervision,” *ICASSP 2020*, pp. 7669–7673, 2019.
- [26] D. Hwang, K. Chai Sim, Y. Zhang, and T. Strohmaier, “Comparison of soft and hard target RNN-T distillation for large-scale ASR,” in *ICASSP 2023*, 2023, pp. 1–5.
- [27] 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 *Interspeech 2020, Virtual Event, Shanghai, China, October 25-29, 2020*. ISCA, 2020, pp. 5036–5040.
- [28] P. Shaw, J. Uszkoreit, and A. Vaswani, “Self-attention with relative position representations,” in *North American Chapter of the Association for Computational Linguistics*, 2018.
- [29] 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,” in *Interspeech*, 2019.
- [30] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: A simple way to prevent neural networks from overfitting,” *Journal of Machine Learning Research*, vol. 15, no. 56, pp. 1929–1958, 2014.
- [31] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” in *International Conference on Learning Representations*, 2019.
- [32] J. Kahn, V. Pratap, T. Likhomanenko, Q. Xu, A. Hannun, J. Cai, P. Tomasello, A. Lee, E. Grave, G. Avidov, B. Steiner, V. Liptchinsky, G. Synnaeve, and R. Collobert, “Flashlight: Enabling innovation in tools for machine learning,” 2022.