



# Blank Collapse: Compressing CTC emission for the faster decoding

Minkyu Jung<sup>1</sup>, Ohhyeok Kwon<sup>2</sup>, Seunghyun Seo<sup>2</sup>, Sunshin Seo<sup>2</sup>

<sup>1</sup>Channel Corporation, Republic of Korea

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

lep@snu.ac.kr, {k.oh, real.seunghyun.seo, sunshin.seo}@navercorp.com

## Abstract

Connectionist Temporal Classification (CTC) model is a very efficient method for modeling sequences, especially for speech data. In order to use the CTC model as an Automatic Speech Recognition (ASR) task, the beam search decoding with an external language model like n-gram LM is necessary to obtain reasonable results.

In this paper, we analyze the blank label in CTC beam search decoding deeply and propose a very simple method to reduce the amount of calculation resulting in faster beam search decoding speed. With this method, we can get up to 78% faster decoding speed than ordinary beam search decoding with a very small loss of accuracy in LibriSpeech datasets.

We prove this method is effective not only practically by experiments but also theoretically by mathematical reasoning. We also observe that this reduction is more obvious if the accuracy of the model is higher.<sup>1</sup>

**Index Terms:** speech recognition, connectionist temporal classification, beam search.

## 1. Introduction

Recently, there has been a remarkable improvement in Automatic Speech Recognition (ASR) system thanks to the End-to-End (E2E) training approaches. Especially, most ASR models have one of three popular models as a base architecture, Connectionist Temporal Classification (CTC) based model [1], recurrent neural network (RNN) transducer-based model [2, 3] and Attention-based Encoder-Decoder (AED) model [4, 5].

Unlike other models, CTC encodes the waveform with sequential models like Recurrent Neural Network (RNN)[6] or Transformer[7, 8] and outputs the probability or logit for each sample frame. Since it does not have any decoder module itself and assumes conditional independence with respect to time, it can make the recognition result directly from its output. Moreover, this model tries to learn the frame-wise phoneme representation, it originally lacks temporal information. Therefore CTC-based model is often used with beam search decoder and external Language Model (LM) like n-gram LM[9, 7] or deep learning-based LM [10]. However, even though the beam search decoder shows a better result, it requires a longer decoding time. Both accuracy and decoding time is too important to give up so usually we have to take the saturation point between this trade-off.

There have been lots of efforts to reduce the beam search decoding time so far. In [11], they introduced beam threshold pruning which prunes some beam candidates with relatively low scores and showed significant improvements in beam search speed. [12] invented vectorized CTC which also shows big enhancements in CTC-Attention-based beam search decoding while [13] proposed an improved RNN-T beam search achieving decoding time

<sup>1</sup>Code and examples are available at <https://github.com/minkjung/blankcollapse>




Figure 1: Number of all consecutive frames for each size based on the frame type in LibriSpeech test-others. The blank type means the frame having the highest probability at the blank index while the non-blank type is for the other cases. Most of the non-blank frames last only one or two frames whereas blank frames seem to last longer.

reduction. Especially for CTC, there is a special label *blank* that has multiple functions in the decoding process and the calculation of CTC loss. [14, 15] proposed phone synchronous decoding (PSD) in WFST which can skip the unnecessary decoding steps whose blank probability is high enough. [16] also tried to reduce frames with high blank probability using CTC posterior in RNN-T model training. However, simply skipping blank frames should be more careful when we consider the end-to-end character-based or subword-based CTC models since the blank symbol is essential for some cases. Some heuristics might avoid relevant problems in PSD but they are not formulated well so far.

In this paper, we research the role of blank labels deeply in the CTC beam search decoding and find an efficient way to remove the redundant computations for the blank label. In conclusion, we propose a blank collapse method for reducing the calculations in decoding which results in the improvement of decoding speed with negligible loss of accuracy. This also can be regarded as a generalization from PSD for any decoding methods and the token units, accompanied by mathematical consideration and empirical results. This can be done without any further training and is available on every CTC emission.

## 2. Analysis on CTC blank

In this section, we study the characteristics of CTC blank. Most of the notations in this section follow those in [1].

### 2.1. Blanks on greedy decoding

Before introducing the method, we have to define the blank frame. Let  $\mathcal{N}_w(\mathbf{x}) = \mathbf{y} = (y_1, y_2, \dots, y_T) \in \mathbb{R}^{T \times |L'|}$  be the CTC emission probability of a certain waveform  $\mathbf{x}$  from CTC model  $\mathcal{N}_w$  where  $\mathbf{y}_t = (y_t^1, y_t^2, \dots, y_t^{|L'|}) \in \mathbb{R}^{|L'|}$  is the probability at a specified timestep  $t$  and  $L = \{w_1, w_2, \dots, w_{|L|}\}$  is a set of labels,  $L' = L \cup \{\epsilon\}$ ,  $\epsilon$  represents the blank label. CTC probability normally can be obtained by applying the softmax function to the logit from

the output of LSTM or Transformer trained with CTC loss. Each  $w_t$  can be a subword or character and we use a character dictionary.

If we want CTC greedy decoding (best path decoding), we calculate the CTC output sequence as  $(\arg \max_{k \in L'} y_t^k)_{t=1}^T$  for each timestep  $t$  and map the sequence with the mapping  $\mathcal{B}$  in [1].  $\mathcal{B}$  maps a sequence of CTC outputs to a label sequence by removing all blanks and repeated labels from the sequence. Thus by the definition of  $\mathcal{B}$ , consecutive blanks play the same role as a single blank. (e.g.  $\mathcal{B}(a\epsilon\epsilon\epsilon a\epsilon b) = \mathcal{B}(a\epsilon a\epsilon b) = aab$ ). This property motivates us to think that consecutive blank outputs could be replaced by a single blank output. Also, we note that we can't ignore a series of blanks entirely because the existence of a blank label plays an important role that used for representing a repetition of the non-blank label. That's why we have to leave at least one blank label on behalf of the following blanks.

However, consecutive blanks at the beginning and the ending of a sequence can be omitted entirely since they affect nothing to the result of greedy decoding. In other words, in greedy decoding, we are allowed to collapse the consecutive blanks to a single blank and drop all blanks before the first non-blank output and after the last non-blank output before decoding. More formally, we get the following.

**Definition 1**  $\mathcal{G} : \mathbb{R}^{T \times |L'|} \rightarrow \{1, \dots, |L'|\}^{\leq T}$  is the **CTC greedy decoding or best path decoding** given by

$$\mathcal{G}(\mathbf{y}) = \mathcal{B}((\arg \max_{k \in L'} y_t^k)_{t=1}^T)$$

for a CTC emission probability  $\mathbf{y}$ .  $\hat{E}_{\mathbf{y}}$  is called a set of **weak blank frames** of  $\mathbf{y}$  defined as

$$\hat{E}_{\mathbf{y}} := \{t \leq T : \arg \max_{k \in L'} y_t^k = \epsilon\}.$$

$\phi : \mathcal{P}(\{1, \dots, T\}) \rightarrow \mathcal{P}(\{1, \dots, T\})$  is a function called the **consecutive extension** defined by

$$\phi(E) := \{t \in E : (t-1 \in E) \vee (t=1) \vee (s \in E, \forall s \geq t)\}$$

By definition, it is straightforward that

$$\mathcal{G}(\mathbf{y}) = \mathcal{G}((\mathbf{y}_t)_{t \notin \phi(\hat{E}_{\mathbf{y}})}) \quad (1)$$

for a CTC emission probability  $\mathbf{y} \in \mathbb{R}^{T \times |L'|}$ . This means we may drop frames included in  $\phi(\hat{E}_{\mathbf{y}})$  for CTC greedy decoding, even though it is not useful in practice.

## 2.2. Blanks on CTC beam search

Equation 1 might not be valid for the CTC beam search decoding. Let  $C_t$  be a set of beam candidates for timestep  $t$ . During CTC beam search, we calculate  $P_b(c, t)$  and  $P_{nb}(c, t)$ , blank probability and non-blank probability respectively for each candidate  $c \in C_t$  and each timestep  $t \leq T$ . As  $t$  proceeds,  $P_b$  and  $P_{nb}$  have to be updated for the case of stay ( $c \in C_{t+1}$ ) or extend to another path ( $c' = c \cup \{k\} \in C_{t+1}$ ) according to the following rules:

$$\begin{aligned} P_b(c, t+1) &= P_{tot}(c, t) \cdot y_t^\epsilon \\ P_{nb}(c, t+1) &+= P_{nb}(c, t) \cdot y_t^k, \quad \text{for } k = c_{-1}, \\ P_{nb}(c', t+1) &+= \begin{cases} P_b(c, t) \cdot y_t^k, & \text{if } k = c_{-1}, \\ P_{tot}(c, t) \cdot y_t^k, & \text{otherwise,} \end{cases} \end{aligned} \quad (2)$$

where  $P_{tot}(c, t) = P_b(c, t) + P_{nb}(c, t)$  and  $c_{-1}$  represents the last label of  $c$ . We use  $+=$  instead of  $=$  for non-blank probability because they can be added from another beam extension. CTC

beam search with a language model (LM) sorts the candidates by their score defined by

$$Score(c, t) := P_{tot}(c, t) \cdot \lambda P_{LM}(c, t). \quad (3)$$

Here,  $P_{LM}$  is a probability of LM and  $\lambda$  is the LM weight (other hyper-parameters like a length penalty are not considered for this time). As we see Equation 2 even though  $y_t^\epsilon$  is the biggest probability,  $y_t^k$  is applied to the non-blank probability which can be accumulated as  $t$  goes. Thus collapsing *weak blank frames* is too risky to collapse carelessly.

Consequently, we consider a stronger condition for the blank frame which affects search little enough to collapse.

**Definition 2**  $E_{\mathbf{y}}^\theta$  is called a set of **strong blank frames** or **just blank frames of  $\theta$**  for  $\mathbf{y}$  defined by

$$E_{\mathbf{y}}^\theta := \{t \leq T : y_t^\epsilon > \theta\}$$

where  $\theta$  is called **blank threshold**.

Compared to *weak blank frames*, on *blank frames of  $\theta$* , it is more confident for CTC model to predict that these frames are for the blank label. Also, we can control the confidence by  $\theta$  for alleviating the difference from the result from original CTC emissions. We propose our method with this definition in the next chapter followed by its reasoning of it.

## 2.3. Comparison between blank and non-blank

Although other consecutive non-blank labels also might be collapsed into a single label in greedy decoding, we don't cover the case of non-blank labels because non-blank labels (a) don't occur often in a row and (b) are riskier than blank labels.

As we can see in Figure 1, non-blank usually occurs at most two frames consecutively whereas blank lasts longer once it occurs. Our ultimate purpose is to reduce the decoding time and collapsing consecutive non-blanks into a single non-blank will make little difference in decoding time.

## 3. Blank collapse method

Now we propose a new method called *blank collapse* which drops all collapsible frames before the CTC beam search in order to reduce the size of decoding frames. Additionally we study the theoretical reasoning behind this method, followed by its limitations and implementation.

### 3.1. Definition of the method

**Definition 3 (blank collapse)**  $\mathbf{F}_\theta : \mathbb{R}^{T \times |L'|} \rightarrow \mathbb{R}^{\leq T \times |L'|}$  is called **blank collapse method** with  $\theta$  defined by  $\mathbf{F}_\theta(\mathbf{y}) = (\mathbf{y}_t)_{t \in \mathbf{I}_\theta}$ , for a CTC emission probability  $\mathbf{y}$  where  $\mathbf{I}_\theta := \phi(E_{\mathbf{y}}^\theta)^C$ . We call  $\phi(E_{\mathbf{y}}^\theta)$  a set of **collapsible frames** of  $\mathbf{y}$  with  $\theta$ . The method using  $\hat{\mathbf{I}} = \phi(\hat{E}_{\mathbf{y}})^C$  as an index set is called **weak blank collapse**.

In other words, from the original CTC emissions, this method drops the blank frames if they occur in the front, the last, or following another blank frame. If we drop these collapsible frames, CTC emissions will be compressed on its length, resulting in a shorter beam search time.

PSD can be seen as a case with  $\mathbf{I} = (E_{\mathbf{y}}^\theta)^C$  which collapses every blank frame without the consecutive extension. This is generally fine when the model uses phoneme as a unit but can cause a significant difference otherwise according to Chapter 2.1.

This method certainly has to maintain the accuracy of beam search as much as possible and this is accomplished by the definition of blank frames. Since  $y_t^\epsilon > \theta$ , it automatically implies  $y_t^k < 1 - \theta, \forall k \in L$ . This means that we can assure that the change of the non-blank probability has to be limited on blank frames while the blank probability almost remains as same before for sufficiently large  $\theta$ . We can almost surely ignore these frames for this reason.

### 3.2. Limitations of the method

Even though we can almost surely omit the consecutive blank frames, this method always as a nonzero opportunity that could harm the original result. No matter how  $\theta$  is large, sometimes consecutive blank frames can change the score quite a lot and reverse the order. As Equation 2 shows, the non-blank accumulates the score not only from its beam but also from another beam as they merge their scores. Namely,

$$P_{nb}(c,t+1) += P_{tot}(c_{-1},t) \cdot y_t^{c-1}$$

where  $c_{-1}$  is the last label of the beam  $c$  and  $c_{-1} := c - c_{-1}$ . In this equation, we can limit the scale of  $y_t^{c-1}$  but we don't know how  $P_{tot}(c_{-1},t)$  be larger than the other beam scores. Again, every beam always has a chance to add a relatively large score from another beam having a much higher score. LM score can cause deepen this kind of side effect. Additionally, this tells us why we collapse only blanks not non-blanks, mentioned in section 2.3 (b). Collapsing frames with high non-blank emission probability can't limit  $P_{nb}$  which can cause serious distortion.

Nevertheless, this method is still effective making little difference from the original. This is because the non-blank probability still has an upper limit  $\theta$  and proper *beam threshold pruning* [11] prevents the beams from overspreading. Trivially, the higher  $\theta$ , the less distortion we could expect.

### 3.3. Implementation

To implement this method, we take advantage of a special utility function *unique\_consecutive* from PyTorch[17]. First, we get a vector  $\Phi$  where each value  $\phi_t$  represents whether the frame at timestep  $t$  belongs to the blank frames or not. This function returns each unique value of a vector in order and how many times a such value occurs consecutively. Using the returned value from this function with a vector  $\Phi$ , we can get  $\mathbf{u}, \mathbf{c}$  representing how many blank/non-blank frames occur in a row. By these values, we can leave only non-collapsible frames,  $\mathbf{I}_\theta$ .

This method is also compatible with timestep alignment [18] since we can reorder the timestep alignment result with  $\mathbf{I}_\theta$ . The detail method is described in algorithm 1. Figure 2 shows the resulting image of collapsed CTC emission before and after the blank collapse. Yellow represents high log probability and the top row is for the blank label. As we can see, consecutive blank frames disappear in collapsed emission. The total length of the collapsed emission reduces from 169 to 102. The ground truth transcript is *I had that curiosity beside me at this moment.*




Figure 2: Log probability of original/collapsed CTC emission for a sample waveform in LibriSpeech.

## 4. Experiments

### 4.1. ASR model and datasets

We use the *wav2vec 2.0* BASE / LARGE [7] model pre-trained on the unlabeled audio data of LibriVox dataset [19] and fine-tuned

---

### Algorithm 1: blank collapse

---

**Data:** CTC probability  $\mathbf{y} \in \mathbb{R}^{T \times |L'|}$ , blank threshold  $\theta$   
**Result:** collapsed probability  $\mathbf{F}_\theta(\mathbf{y}) \in \mathbb{R}^{\leq T \times |L'|}$ , indices  $\mathbf{I}_\theta$   
 $\Psi \leftarrow (\psi_1, \dots, \psi_T)$  where  $\psi_t = [y_t^e \geq \theta], \forall t \leq T$ ;  
 $N_0 \leftarrow 0$  // number of first blanks;  
 $\mathbf{u}, \mathbf{c} \leftarrow \text{unique\_consecutive}(\Phi)$ ;  
 $\mathbf{C} \leftarrow []$ ;  
 $k \leftarrow 0$ ;  
**for**  $i \leftarrow 1$  **to**  $\text{length}(\mathbf{c})$  **do**  
  **if**  $u_i$  **is** *True* **then**  
    **if**  $i = 1$  **then**  
       $N_0 \leftarrow c_i$   
    **else**  
      **if**  $i = \text{length}(\mathbf{c})$  **then**  
        **break**  
      **else**  
         $\mathbf{C} \leftarrow [\mathbf{C}; c_i]$   
        // append number of blanks  
      **end**  
    **end**  
  **else**  
     $\mathbf{C} \leftarrow [\mathbf{C}; [1, \dots, 1]]$  // append  
    1's with number of non-blanks  
  **end**  
   $k \leftarrow k + c_i$   
**end**  
 $\mathbf{I}_\theta \leftarrow \text{cumsum}(\mathbf{C}) - 1 + N_0$ ;  
 $\mathbf{F}_\theta(\mathbf{y}) \leftarrow (\mathbf{y}_t)_{t \in \mathbf{I}_\theta}$ ;

---

on either 10 minutes, 100 hours, and 960 hours of transcribed LibriSpeech dataset[20]. We do not fine-tune further and use a CTC beam search decoder and 4-gram word LM provided by torchaudio [21] which uses [22] for the decoder. Since this decoder uses the logit of each beam for sorting beam candidates we get the probability vector explicitly by the softmax function to apply our method. Also, we try the beam threshold for  $\gamma = 10, 30, 50$ .

Beam search decoding in an end-to-end manner uses 32 batch size, 1,500 beams, LM weight 1.57, and length penalty -0.64 for every experiment in order to simulate the experiments done by [7]. For the analysis of their effectiveness, we use various beam thresholds ( $\gamma$ ) and blank thresholds ( $\theta$ ). Every experiment is done on Intel(R) Xeon(R) Gold 5120 CPU @ 2.20GHz.

We evaluate our method on LibriSpeech dev and test sets. Firstly we check the proportion of collapsible frames to all frames from a CTC emission with  $\theta = 0.999, 0.99, 0.9$ , and weak collapse (Definition 2). As we can see at Table 1, about half of CTC frames seem to be collapsible frames. Additionally, there are more collapsible frames with the lower  $\theta$  than the higher one, though the gap between them is not that big. With this observation, we can estimate that the CTC model predicts a certain frame to a blank frame with very high confidence so that there are not many frames with ambiguous certainty.

Table 1: The percentage of collapsible frames to all frames from CTC emissions with wav2vec 2.0 LARGE model fine-tuned on 960 hours.

$\theta$	0.999	0.99	0.9	weak
dev-clean	42.87	43.24	43.60	43.97
dev-other	42.98	43.77	44.51	45.20
test-clean	43.88	44.27	44.65	45.05
test-other	43.83	44.67	45.44	46.15

Table 2: Word Error Rate (WER) (%) and Real Time Factor (RTF) with its reduction ratio compared to the original on LibriSpeech dev/test sets with  $\gamma = 50$ . RTF includes the time spent executing blank collapse which is less than a second.

$\theta$		dev-clean	dev-other	test-clean	test-other
original	WER	1.781	3.509	2.031	3.681
	RTF	0.278	0.290	0.279	0.291
0.999	WER	1.781	3.511	2.031	3.681
	RTF	<b>0.160</b> (0.42)	<b>0.167</b> (0.42)	<b>0.157</b> (0.44)	<b>0.165</b> (0.43)
0.99	WER	1.783	3.513	2.029	3.683
	RTF	<b>0.160</b> (0.42)	<b>0.165</b> (0.43)	<b>0.156</b> (0.44)	<b>0.163</b> (0.44)
weak	WER	1.866	3.749	2.109	3.834
	RTF	0.158 (0.43)	0.161 (0.45)	0.154 (0.45)	0.159 (0.45)

## 4.2. Experimental results

Table 2 shows the accuracy and decoding time of each experiment with  $\gamma = 50$ . As we can see, blank collapse with sufficiently high  $\theta$  shows a significant improvement in its decoding speed with little difference in accuracy. For test-clean,  $\theta = 0.99$  shows about 44% time reduction at most, which equals to about 78% speed enhancement without any accuracy loss. Weak blank collapse shows the best speed among all settings with not a small distortion on the accuracy. Thus we may know that sufficient  $\theta$  is safe for this method.

Interestingly, the reduction ratio of the decoding time seems to be almost as much as the ratio of their frame sizes, 43.2% and 43.8% respectively for test-other. However, this is not always true. As we see in Figure 3, with  $\gamma = 50$ , the reduction ratio of decoding time is very similar to that of the size of frames. On the contrary,  $\gamma = 10, 30$  shows a little bit less effect than  $\gamma = 50$ . With this phenomenon, we can estimate that the ratio of the time consuming on blank frames to that on non-blank frames is relatively lower on the smaller  $\gamma$ . This is because the beam threshold pruning with small  $\gamma$  reduces the time more on the blank frames than the others. In other words, a higher  $\gamma$  beam search takes longer time on blank frames than lower  $\gamma$  because it has more candidates which must be pruned with the lower  $\gamma$ .

Table 3 shows the correlation between the model type and the improvement of decoding time by blank collapse. The larger the size of the model and the larger the dataset has been fine-tuned, the improvement of the speed gets more evident. Actually, these two factors directly affect the accuracy of the model, which is the key factor in deciding how many frames can be collapsed out. This is because the model with better accuracy tends to provide the CTC emission probability with higher confidence and it makes the blank probability high enough to be dropped. In other words, a good model can tell whether a certain frame can be collapsed out or not.

Table 3: The percentage of decoding time improvement on various model types, depending on the size of the model and the fine-tuning dataset.

$\theta$	0.999	0.99	0.9	weak
BASE / 10 min	23.56	28.22	31.54	35.81
LARGE / 10 min	35.95	39.05	40.34	40.59
LARGE / 100h	41.08	42.50	43.54	43.96
LARGE / 960h	<b>43.30</b>	<b>43.99</b>	<b>45.36</b>	<b>45.70</b>

## 4.3. Results on other decoder settings

In this section, we discuss proposed methods in two different decoder settings on LibriSpeech test-other subset. Firstly, we decode both vanilla frames and collapsed ones with a CTC beam search decoder fused with a Transformer based LM. We use word-level




Figure 3: Reduction ratio of the decoding time by blank collapse in LibriSpeech test-other according to  $\theta$  for each  $\gamma$ , compared to the reduction ratio of the number of frames.

20 decoder layers Transformer for LM[9, 7] and the same decoder in section 4.1. However, unlike the n-gram fusion which shows significant improvement in inference speed with a negligible WER degradation, there is only 1% gain in decoding time when  $\theta$  is 0.999. This must be because most of the time is spent on huge LM inference when decoding with a neural LM.

Nextly, we experiment with a WFST-based beam search decoder [23]. We compile the same 4-gram LM in section 4.1 for a language graph and a CTC graph using k2 framework [24]. And we implement a WFST decoder using Kaldi [25]. We also observe the improvement of decoding time by blank collapse with WFST-based decoding. The reduction ratio of the decoding time is about 33% with a very small loss in accuracy similar to the case for the end-to-end decoder. Since the hyper-parameters of the WFST decoder are different from the end-to-end decoder, its improvement seems to be relatively low. If we find more optimized hyper-parameters for the WFST decoder, we expect a similar speed improvement to that of the end-to-end decoder.

## 4.4. Analysis of side-effects

It seems a bit strange that the reduction ratio on  $\theta = 0.9$  is lower than  $\theta = 0.999, 0.99$  at  $\gamma = 50$  as shown in Figure 3. It's rare but it also does happen on our internal dataset as well. We interpret that over-collapsing might drop some good candidates unintentionally which might lead to faster decoding with appropriate pruning for the rest of the time.

Additionally, there are some cases in which the accuracy turns out to be better when collapsed than the original. For example, WER of test-clean with  $\theta = 0.99$  (2.029) is lower than that of the original (2.031) though it is a small amount. We guess that blank collapse may drop some frames having harmful information potentially. However, such cases do not occur very often and the difference is usually small throughout our experiments.

## 5. Conclusions

In this paper, we analyze the characteristics of the blank label which is used as a special role in the CTC model. We define the blank frame as a frame with a high blank probability  $y_t^e$  and find out that usually CTC emission has a bunch of blank frames. It is also discussed that the blank frames can be omitted in CTC beam search decoding in almost every case.

By this, we finally propose a new method called blank collapse which intends to reduce the collapsible frames in order to improve the decoding speed with minimal loss of accuracy. It is shown by several experiments that this method actually can improve the decoding speed. We also find that we can collapse more frames when the model is well-trained on a dataset and with a higher beam threshold.

As future work, we expect that our method can be plugged into any E2E ASR models using CTC loss as a regularization.

## 6. References

- [1] A. Graves, S. Fernández, F. Gomez *et al.*, “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 369–376.
- [2] A. Graves, “Sequence transduction with recurrent neural networks,” *ICML workshop on representation learning*, 2012.
- [3] Q. Zhang, H. Lu, H. Sak *et al.*, “Transformer transducer: A streamable speech recognition model with transformer encoders and rnn-t loss,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7829–7833.
- [4] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *2016 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2016, pp. 4960–4964.
- [5] L. Dong, S. Xu, and B. Xu, “Speech-transformer: a no-recurrence sequence-to-sequence model for speech recognition,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5884–5888.
- [6] A. Graves and N. Jaitly, “Towards end-to-end speech recognition with recurrent neural networks,” in *Proceedings of the 31st International Conference on International Conference on Machine Learning-Volume 32*, 2014, pp. II–1764.
- [7] A. Baevski, Y. Zhou, A. Mohamed *et al.*, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 12 449–12 460, 2020.
- [8] A. Gulati, J. Qin, C.-C. Chiu *et al.*, “Conformer: Convolution-augmented transformer for speech recognition,” *Proc. Interspeech 2020*, pp. 5036–5040, 2020.
- [9] G. Synnaeve, Q. Xu, J. Kahn, T. Likhomanenko, E. Grave, V. Pratap, A. Sriram, V. Liptchinsky, and R. Collobert, “End-to-end asr: from supervised to semi-supervised learning with modern architectures,” in *ICML 2020 Workshop on Self-supervision in Audio and Speech*.
- [10] A. Y. Hannun, A. L. Maas, D. Jurafsky, and A. Y. Ng, “First-pass large vocabulary continuous speech recognition using bi-directional recurrent dnns,” *arXiv preprint arXiv:1408.2873*, 2014.
- [11] M. Freitag and Y. Al-Onaizan, “Beam search strategies for neural machine translation,” in *Proceedings of the First Workshop on Neural Machine Translation*, 2017, pp. 56–60.
- [12] H. Seki, T. Hori, S. Watanabe *et al.*, “Vectorized beam search for ctc-attention-based speech recognition,” in *INTERSPEECH*, 2019, pp. 3825–3829.
- [13] M. Jain, K. Schubert, J. Mahadeokar *et al.*, “Rnn-t for latency controlled asr with improved beam search,” *arXiv preprint arXiv:1911.01629*, 2019.
- [14] C. *et al.*, “Phone synchronous decoding with ctc lattice,” in *Interspeech*, 2016, pp. 1923–1927.
- [15] Z. *et al.*, “Tiny transducer: A highly-efficient speech recognition model on edge devices,” in *ICASSP 2021*. IEEE, 2021, pp. 6024–6028.
- [16] Y. Wang, Z. Chen, C. Zheng, Y. Zhang, W. Han, and P. Haghani, “Accelerating rnn-t training and inference using ctc guidance,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [17] A. Paszke, S. Gross, F. Massa *et al.*, “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, Eds. Curran Associates, Inc., 2019, pp. 8024–8035.
- [18] L. Kürzinger, D. Winkelbauer, L. Li, T. Watzel, and G. Rigoll, “Ctc-segmentation of large corpora for german end-to-end speech recognition,” in *Speech and Computer*, A. Karpov and R. Potapova, Eds. Cham: Springer International Publishing, 2020, pp. 267–278.
- [19] “[https://librivox.org/!](https://librivox.org/)”
- [20] V. Panayotov, G. Chen, D. Povey *et al.*, “Librispeech: an asr corpus based on public domain audio books,” IEEE, 2015, pp. 5206–5210.
- [21] Y.-Y. Yang, M. Hira, Z. Ni *et al.*, “Torchaudio: Building blocks for audio and speech processing,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6982–6986.
- [22] J. D. Kahn, V. Pratap, T. Likhomanenko, Q. Xu, A. Hannun, J. Cai, P. Tomasello, A. Lee, E. Grave, G. Avidov *et al.*, “Flashlight: Enabling innovation in tools for machine learning,” in *International Conference on Machine Learning*. PMLR, 2022, pp. 10 557–10 574.
- [23] 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.
- [24] D. Povey, P. Zelasko, and S. Khudanpur, “Speech recognition with next-generation kaldi (k2, hotse, icefall),” *Interspeech: tutorials*, 2021.
- [25] D. Povey, A. Ghoshal, G. Boulianne *et al.*, “The kaldi speech recognition toolkit,” in *IEEE 2011 workshop on automatic speech recognition and understanding*, no. CONF. IEEE Signal Processing Society, 2011.