



An investigation of phone-based subword units for end-to-end speech recognition

Weiran Wang, Guangsen Wang, Aadyot Bhatnagar, Yingbo Zhou, Caiming Xiong, Richard Socher

Salesforce Research
575 High St, Palo Alto, CA 94301, USA

{weiran.wang, guangsen.wang, abhatnagar, yingbo.zhou, cxiong, rsocher}@salesforce.com

Abstract

Phones and their context-dependent variants have been the standard modeling units for conventional speech recognition systems, while characters and subwords have demonstrated their effectiveness for end-to-end recognition systems. We investigate the use of phone-based subwords, in particular, byte pair encoder (BPE), as modeling units for end-to-end speech recognition. In addition, we also developed multi-level language model-based decoding algorithms based on a pronunciation dictionary. Besides the use of the lexicon, which is easily available, our system avoids the need of additional expert knowledge or processing steps from conventional systems. Experimental results show that phone-based BPEs tend to yield more accurate recognition systems than the character-based counterpart. In addition, further improvement can be obtained with a novel one-pass joint beam search decoder, which efficiently combines phone- and character-based BPE systems. For Switchboard, our phone-based BPE system achieves 6.8%/14.4% word error rate (WER) on the Switchboard/CallHome portion of the test set while joint decoding achieves 6.3%/13.3% WER. On Fisher + Switchboard, joint decoding leads to 4.9%/9.5% WER, setting new milestones for telephony speech recognition.

Index Terms: end-to-end speech recognition, byte pair encoding, multi-level language model, one-pass decoding

1. Introduction

For English speech recognition, phones and the context-dependent variants have long been the standard modeling units for conventional automatic speech recognition (ASR) systems [1, 2]. However, there has been a surge of interest in using character and character-based subwords, such as byte pair encoding (BPE, [3]) and word-pieces [4], in modern end-to-end systems [5, 6, 7, 8, 9, 10, 11, 12, 13, 14]. The advantages of character and subwords mainly lie in the simplicity. They not only enable straightforward open-vocabulary recognition, but also avoid the need of domain knowledge such as pronunciation dictionary and phonetic questions in state tying [15] for context-dependent phones, and additional processing steps such as hidden Markov models and Gaussian mixture models training for phone state modeling. On the other hand, phones have tighter correspondence to audio than characters, and often leads to higher recognition accuracy. As a concrete example, for the Switchboard corpus, it takes more modeling effort for the char CTC system to match performance of the phone CTC system, as noted by [16] and [17]. This motivates the question of whether phone-based subwords would be more effective for ASR.

In this work, we investigate the use of phone-based BPEs in the context of end-to-end speech recognition. We use a pronunciation dictionary to convert transcription into phone sequences while maintaining the word boundaries, and extract

BPEs by gradually merging frequent pairs of phones or phone sequences, as is done for character BPEs.¹ Intuitively, similar to the context-dependent phones, phone BPEs shall capture correlations between contiguous phones. On the other hand, it allows a trade-off between the size of modeling units and output length, which influences both the performance and efficiency of end-to-end systems. We train the acoustic model using the phone BPE targets as usual, with the hybrid CTC/attention loss [19].

At decoding time, we use the pronunciation dictionary again to convert decoded phone BPE sequence back into words, with a novel multi-level RNN language model (LM). To take advantage of the complementarity of both phone and character BPE systems, we develop a one-pass joint beam search decoder that efficiently combines the two. Our experimental results on both the Wall Street Journal and Switchboard show that the phone BPE systems tend to outperform the character-based counterpart in accuracy, and joint decoding with the two leads to significant improvement. Specifically, our phone BPE system achieves 6.8%/14.4% word error rate (WER) on the SWBD/CALLHM portion of the test set while joint decoding achieves 6.3%/13.3% WER. This, to the best of our knowledge, sets the new state-of-the-art on Switchboard 300 hours. By adding the Fisher corpora into training, we obtain best single systems and joint decoding yields 4.9% WER on the SWBD portion, setting a new milestone for this setup.

We would like to emphasize that, besides the use of pronunciation dictionary for subword extraction and decoding², our method avoids the extra processing steps from conventional systems. On the other hand, large collections of pronunciations are readily accessible [20], and pronunciation of out-of-collection words can be constructed with grapheme-to-phoneme methods [21, 22, 23, 24], and therefore our approach maintains the simplicity of end-to-end methods.

2. Multi-level LM for decoding with BPEs

Compared to decoding with character BPEs which just needs to output the highest scoring sequence of modeling units, possibly without additional language models (subword-level or word-level), there are a few challenges for decoding with phone BPEs. First, it is necessary to use a pronunciation dictionary to convert the decoded phone sequence into a word sequence. Second, unlike in the character case where the spelling uniquely determines a word, different words can have the same pronunciations for the case of phones (i.e., homophones) and therefore a word LM is helpful for distinguishing them.

We develop a multi-level LM that combines scores from both a subword LM and a word LM, and use it in beam search

¹We use implementation of [18] for training and inference of BPEs.

²We pick one pronunciation for each word if the word has multiple pronunciations in the dictionary.

by shallow fusion [25]. Intuitively, the method is similar to that proposed in [26], which combines character LM and word LM. We build a prefix tree storing the pronunciation of words in the dictionary, and as we move down the tree from the root according to the hypothesized subwords and accumulating subword LM scores from each step, we may come across tree nodes that contain words, whose pronunciations match the sequences of subwords on the paths stemming from root. At that time, we may decide to output the words and replace accumulated subword LM scores with word LM scores, and subsequently move back to the root. We highlight challenges not present in [26]:

- To build the prefix tree, we need to first decompose the pronunciation of each word (i.e., a phone sequence) into a phone BPE sequence, using the BPEs extracted from transcription. This decomposition is greedy (utilizing large subwords as much as possible) and deterministic.
- In [26], the modeling units are characters, and one can determine the completion of a word when the word boundary '␣' is proposed. In the BPE case, however, many subwords contain '␣' as the first symbol. Therefore, all these subwords indicate both the word boundary and a new word is started, *simultaneously*.
- Due to the issue of homophones, a node in the prefix tree may contain multiple words. If the word boundary is met at such a node, we have to output multiple word hypotheses. As a result, one decoding beam branches into multiple beams, which have the same subword LM state but different word LM states.

In Algorithm 1, we detail the forwarding function of our multi-level LM M , constructed from a subword LM S and a word LM W . This function $\text{forward}(\text{state}, s)$ takes as input the current state state and a subword s , and returns the updated state after accepting s , the vector of look-ahead scores la_scores for all subwords, and word outputs if the word boundary is met (output is set to special token `<incomplete>` otherwise). la_scores are computed based on scores from S and W with a relative weight α (S is not activated if $\alpha = 0$), and they are combined with the acoustic model scores for evaluating partial hypotheses. The state for multi-level LM is a 6-tuple

$(S\text{state}, S\text{logp}, W\text{state}, W\text{logp}, \text{node}, \text{accum})$

containing the state and associated log-probabilities (for subwords) from S , the state and associated log-probabilities (for words) from W , the position in the prefix tree T , accumulated subword score since the last output word. To start decoding, we initialize states and log-probabilities by accepting the start of sentence token `<sos>`, and set node to the root of T .

3. Joint one-pass decoding of BPE systems

The multi-level LM in Section 2 applies to both phone and character BPE systems. Since the two types of units can be complementary to each other (capturing different aspects of the language), to take advantage of this, we develop a one-pass beam search algorithm utilizing both systems. Note that, while ideas of (hierarchically) combining phone and character labels have been explored for acoustic model training in end-to-end systems (e.g., [27, 28, 29, 30]), our approach of combining different units at decoding time is orthogonal to them. In [31], the authors used the pronunciation dictionary to assist the extraction of character-based subwords, sharing the intuition that phones and characters are complementary. However, in the end they train only one subword system, and the subword extraction requires a non-trivial step of aligning phones and characters.

Algorithm 1 The forwarding function of multi-level RNNLM.

Input: subword s , previous state state . $W\text{logp}(w)$ gives the score at position w of vector $W\text{logp}$. Function $\text{node.getWords}()$ returns the list of complete words associated with node of prefix tree T , $\text{node.getTokens}()$ returns the list of subwords branching out from node , and $\text{node.branch}(s)$ returns the new node after accepting s at node . α is used for weighing scores of S versus M .

```

(Sstate, Slogp, Wstate, Wlogp, node, accum) ← state
if  $s.\text{startswith}('␣')$  and (not node == root) then
  # Word boundary is met, inter-word transition
  if  $\text{node.getWords}()$  is not empty then
    wordlist ←  $\text{node.getWords}()$ 
  else
    wordlist ← [<unk>]
  end if
  output ← [] (empty list)
  for  $w$  in wordlist do
    if  $w == \text{<unk>}$  then
      adjust ←  $W\text{logp}(\text{<unk>}) + \text{oov\_penalty}$ 
    else
      adjust ←  $W\text{logp}(w) - \text{accum}$ 
    end if
    # Update word LM state
    ( $W\text{state\_new}, W\text{logp\_new}$ ) ←  $W.\text{forward}(W\text{state}, w)$ 
     $\text{accum\_new} \leftarrow \alpha \cdot S\text{logp}(s)$ ,  $\text{node\_new} \leftarrow \text{node.branch}(s)$ 
    # Update subword LM state
    ( $S\text{state\_new}, S\text{logp\_new}$ ) ←  $S.\text{forward}(S\text{state}, s)$ 
     $\text{la\_scores\_new} \leftarrow \text{adjust} + \alpha \cdot S\text{logp\_new}$ 
    ( $\text{state\_new}, W\text{logp\_new}, \text{node\_new}, \text{accum\_new}$ )
    output.append(( $\text{state\_new}, \text{la\_scores\_new}, w$ ))
  end for
  return output
else
  # Intra-word transition, no word output
   $w \leftarrow \text{<incomplete>}$ 
  if  $s$  in  $\text{node.getTokens}()$  then
     $\text{node} \leftarrow \text{node.branch}(s)$ 
     $\text{accum} \leftarrow \text{accum} + \alpha \cdot S\text{logp}(s)$ 
    ( $S\text{state}, S\text{logp}$ ) ←  $S.\text{forward}(S\text{state}, s)$ 
     $\text{la\_scores} \leftarrow \alpha \cdot S\text{logp}$ 
  else
     $\text{la\_scores} \leftarrow -\infty$  (vector of all  $-\infty$ 's)
  end if
   $\text{state} \leftarrow (S\text{state}, S\text{logp}, W\text{state}, W\text{logp}, \text{node}, \text{accum})$ 
  return [( $\text{state}, \text{la\_scores}, w$ )]
end if

```

Our main idea is to use the phone BPE system, denoted as Model1 , to propose subwords. After seeing a word boundary, we decompose the word into a character BPE sequence, and run the character BPE system, denoted as Model2 , to accept the sequence. Scores from both systems are then linearly combined up to the word boundary. In other words, the phone BPE system leads the decoding process and the character BPE system verifies its hypotheses; the two systems synchronize at each word boundary. In such a way, we incorporate the evidence from Model2 as early as possible to adjust scores of word hypotheses. Compared to second-pass rescoring, our one-pass approach avoids generating large amount of hypotheses by Model1 .

To simplify presentation, we divide each BPE system into acoustic model (AM) and language model (LM). The AM refers to the model trained with end-to-end objectives, e.g., the hybrid CTC/attention model [19] in our case (although the decoder implicitly models the language of subwords). The AM also provides a scoring function which computes the score of the next subword given acoustic inputs and previously decoded subwords (in our case, the score is a linear combination of log-probabilities from the attention decoder and CTC prefix score). The LM refers to the one described in previous section with its components (subword and word LMs) trained separately on text data, and provides the `forward()` function which computes the score of next subword given previously decoded results.

Our algorithm maintains a set of decoding beams (hypotheses), each of which is an 8-tuple

$(\text{score}, \text{ws}, \text{sc1}, \text{ys1}, \text{st1}, \text{sc2}, \text{ys2}, \text{st2})$

containing the final score of the beam (for pruning), the word hypothesis, followed by the score (sc), output subword sequence (ys), and multi-level LM state (st) from `Model1` and `Model2` respectively. The detailed procedure is given in Algorithm 2. We use the parameter β for combining the LM score with the AM score within each system, and $\gamma \in [0, 1]$ for combining scores from both systems, as shown in (*). We use the end detection method of [19] for terminating the algorithm, when longer hypotheses consistently yield much lower scores than the best finished hypothesis (ending with `<eos>`).

In our algorithm, for each beam we run `Model1` once through the phone BPE sequence, run `Model2` once through the corresponding character BPE sequence; `Model2` does not propose additional hypotheses but simply follows `Model1`. Therefore, the time complexity of Algorithm 2 is roughly the sum of that of individual systems for the same beam size.

4. Experiments

In the experiments, we largely adopt the acoustic modeling recipe based on transformers [32] from Espnet [11], as detailed in [33], for training the hybrid attention + CTC model [19].³

We explore with two model architectures of different size. For the default, small architecture, the encoder shared by both attention and CTC consists of 2 convolutional layers that reduce the time and frequency dimension by a factor of 4, and 12 transformer layers, while the decoder consists of 6 transformer layers. Every encoder layer employs self-attention, and every decoder layer employs self-attention (to the previously decoded labels) followed by source-attention (to encoder outputs). All attention operations use 4 heads of 64 dimensions each, and the output of multi-head attention goes through a one-hidden-layer position-wise feed-forward network of 2048 ReLU units, before it is fed to the next layer. For the large architecture, we increase the number of transformer layers to 24 and 12 for the encoder and decoder respectively. The number of attention heads is increased to 6, yielding an attention dimension of 384. To improve generalization of the large architecture, we apply the technique of stochastic residual connections from [34]. During training, we randomly skip the attention and feed-forward operations with a probability for each layer so the layer reduces to the identity mapping, and the layer dropout probability linearly increases with depth up to p . The large architecture uses mostly the same hyperparameters tuned with the default architecture, and its results are denoted with “stochastic layers (p)”.

We extract 80D fbank features plus 3D pitch features from audio (resampled to 16KHz) as inputs to acoustic model.

³We mainly adapt Espnet’s character BPE recipe for Switchboard.

Algorithm 2 Beam search algorithm for joint BPE system.

Input: Input x , trained models, and parameters (β, γ) . `top(sc, bs)` return the list of (score, subword)-tuples of the bs highest values in vector sc . `prune(\mathcal{H} , bs)` returns the bs highest scoring beams (in the score field) from \mathcal{H} . `finish(beam)` forwards both systems to accept `<eos>` and the final output words), and compute final score as in (*).

```

 $\mathcal{H} \leftarrow [(0.0, [\text{<sos>}], 0.0, [\text{<sos>}], \text{init\_st1},$ 
               $0.0, [\text{<sos>}], \text{init\_st2})]$ 
 $\mathcal{C} \leftarrow []$  (set of completed beams)
while end_detection( $\mathcal{C}$ ) == false do
   $\mathcal{T} \leftarrow []$ 
  for beam in  $\mathcal{H}$  do
    (score, ws, sc1, ys1, st1, sc2, ys2, st2)  $\leftarrow$  beam
    lm1_output  $\leftarrow$  LM1.forward(st1, ys1[-1])
    for (st1_n, la1, w) in lm1_output do
      yscores1  $\leftarrow$  AM1.score(x, ys1) +  $\beta \cdot$  la1
      sc2_n  $\leftarrow$  sc2, ys2_n  $\leftarrow$  ys2, st2_n  $\leftarrow$  st2
      if not w == <incomplete> then
        # Word boundary, forward Model2
         $\Delta \leftarrow$  spm.encode(w) (decomp. w into char BPEs)
        for y in  $\Delta$  do
          (st2_n, la2, _)  $\leftarrow$  LM2.forward(st2_n, ys2_n[-1])
          sc2_n  $\leftarrow$  sc2_n + AM2.score(x, ys2_n)(y) +  $\beta \cdot$  la2(y)
          ys2_n.append(y)
        end for
      end if
      for (c, y) in top(yscores1, beamsize) do
        ws_n  $\leftarrow$  ws, sc1_n  $\leftarrow$  sc1 + c, ys1_n.append(y)
        if not w == <incomplete> then
          # Incorporate Model2 score at word boundary
          score_n  $\leftarrow$  (1 -  $\gamma$ )  $\cdot$  sc1 +  $\gamma \cdot$  sc2_n + c (*)
          ws_n.append(w)
        else
          # Otherwise update score with Model1 only
          score_n  $\leftarrow$  score + c
        end if
         $\mathcal{T}$ .append((score_n, ws_n, sc1_n, ys1_n,
                  st1_n, sc2_n, ys2_n, st2_n))
      end for
    end for
  end for
   $\mathcal{H} \leftarrow$  prune( $\mathcal{T}$ , beamsize)
  for beam in  $\mathcal{H}$  do
    if beam.ys1[-1] == <eos> then
       $\mathcal{C}$ .append(finish(beam))
    end if
  end for
end while

```

We apply SpecAugment [35], with “max_time_warp” set to 5 (frames), two frequency masks of widths up to 30 frequency bins, and two time masks of widths up to 40 frames during training. A warmup schedule is applied to ADAM [36] learning rate. We average weight parameters of last 10 epochs to obtain the final model. A beam size of 20 is used for decoding.

4.1. Results on Wall Street Journal (WSJ)

We first verify the effectiveness of phone BPE systems on the WSJ corpora (LDC93S6B and LDC94S13B). Partitions `si284/dev93/eval92` are used as the training/development/test set respectively. Mini-batch size is set to 16 for 100 training

Table 1: Dev WERs (%) of BPE systems with different number of units k for WSJ. LM weights (α , β) are given in parenthesis.

Systems	$k=75$	100	150	250	500	1000
Char BPE Subword ($\beta=0.8$)	9.6	9.8	10.3	10.9	11.2	12.0
Char BPE Multi-level (0.6, 1.0)	7.4	7.5	8.3	9.0	9.1	10.1
Phone BPE Multi-level (0.6, 1.0)	6.2	6.5	7.0	7.6	8.3	9.1

Table 2: Test WERs (%) on WSJ. $k=75$ for BPE systems.

Unit/ LM	Char Word LM	Char BPE Subword	Char BPE Multi-level	Phone BPE Multi-level	Joint $\gamma=0.2$
WER	4.9	7.1	5.1	3.6	3.4
Stoc. layers (0.5)		6.0	4.1	3.1	3.0

epochs. Both subword and word LMs are trained on the WSJ language model training data, and the word LM has a vocabulary size of 65K. To build the lexicon, we use *cmudict* [20] and apply the g2p model from [37] to words not in *cmudict*.

We vary the number of subword units, denoted by k , for both character and phone BPE systems. Dev set WERs of these systems are given in Table 1. For character BPE systems, we use both subword RNNLM and multi-level RNNLM for decoding. We observe that small k is clearly preferred by both systems, indicating that the WSJ training set (80 hours) probably has poor coverage for large set of BPE units; this observation is in line with that of [31].⁴ Furthermore, phone BPE systems consistently outperforms character BPE systems for all k with multi-level LM. By investigating dev set learning curves for BPE systems with different k (not shown due to space limit), we observe too large k (e.g., 1000) yields significantly worse losses, agreeing with the trend in WER. And for the same k , the phone BPE systems have consistently lower losses than character BPE systems, implying less confusion for the acoustic model.

We obtain test WERs of systems with $k=75$ in Table 2. As a reference, we provide the WER of Espnet’s character recipe trained with our setup (specAugment, same mini-batch size). With word RNNLM for decoding, it is not surprising that the character system, using 52 character units, performs very similarly to the character BPE system with $k=75$ (4.9% vs. 5.1% test WER). The phone BPE system significantly outperforms character based systems, and joint decoding which emphasizes the phone BPE system (using $\gamma=0.2$) yields small further improvement. We note that the improvement from using phone-based subwords is on par with the one obtained from discriminative training for the character system achieved by [38]. We then train the large architecture and give resulting WERs in Table 2 (last row). Stochastic transformer layers [34] leads to improved generalization for large models.

4.2. Results on Switchboard (SWBD)

We now experiment with the Switchboard corpus (LDC97S62), with 300 hours of training data. In the first set of experiments, to explore the effect of k , we hold out 4K utterances from the full training set as development set, and set the mini-batch size to 256 for training the small architecture on the remaining 193K utterances for 100 epochs. For decoding, the word LM has a vocabulary size of 31K and is trained on the SWBD training set transcription. When $k=50$, the BPEs include only individ-

⁴We have trained a phone BPE system with $k=46$, where the BPEs include only the individual phones; this simulates a phone-based baseline and gives 6.3% WER on the dev set.

Table 3: RT-03 WERs (%) of phone-based BPE systems with different k trained on SWBD (193K utterances). $\alpha=0$, $\beta=0.4$.

	$k=50$	250	350	500	1000	2000
Phone BPE	15.2	14.7	14.7	14.5	14.4	15.5

Table 4: WERs (%) of BPE systems on eval2000 and RT-03.

Modeling units	eval2000		RT03
	SWBD	CALLHM	
Char + sMBR [39]	12.0	23.1	
Char BPE [40]	11.0	23.1	
SentencePiece [14]	9.2	19.1	
WordPiece [35]	6.8	14.1	
Ours: Stoc. layers (0.5), $\alpha=0$			
Char BPE (Subword, $\beta=0.2$)	7.0	14.5	12.8
Phone BPE (Multi, $\beta=0.4$)	6.8	14.4	12.3
Joint ($\beta=0.4$, $\gamma=0.4$)	6.3	13.3	11.4
+ Fisher: Stoc. layers (0.2)			
Char BPE (Subword, $\beta=0.1$)	5.1	9.5	8.5
Phone BPE (Multi, $\beta=0.2$)	5.5	10.0	9.1
Joint ($\beta=0.2$, $\gamma=0.4$)	4.9	9.5	

ual phones, simulating a phone-based baseline. We show the performance of phone BPE systems on RT-03 (LDC2007S10) in Table 3. Observe that SWBD prefers a much larger k than WSJ, and the performance is stable for a large range of k .

We then add the 4K utterances back for training the large architecture, with minibatch size of 512 for 150 epochs. With PyTorch’s DistributedDataParallel scheme, training takes 2 days with 8 Tesla V100 GPUs. We use a word RNNLM trained also on Fisher text for decoding, with a vocabulary size of 67K. Test WERs on *eval2000* (LDC2002S09 and LDC2002T43) are provided in Table 4, together with recent results obtained by attention-based models.⁵ The character BPE system no longer benefits from the word LM. With stochastic layers, the larger architecture significantly outperforms the small one. Our individual systems are on par with the large model from [35], and joint decoding improves over known results in this setup.

Finally, we add the Fisher corpora (LDC2004T19 and LDC2005T19) to training, without tuning the architecture and most hyperparameters. Models converge after 60 epochs over the joint training set. We tune LM weight for decoding, and with more acoustic training data, systems now prefer smaller β ’s. Our phone BPE system performs similarly to previous best single system [41], while the char BPE system, to our knowledge, sets the new state-of-the-art for single system with 5.1%/9.5% test WER. Joint decoding further improves the WER on the SWBD portion to 4.9%, outperforming known SWBD results even by large ensemble of systems [41].

5. Conclusions

We have investigated the use of phone-based subwords in end-to-end ASR, proposed a novel multi-level language model for subword-based decoding, and a new beam search algorithm for joint decoding with phone- and character-based subword systems. Experiments on two benchmark datasets show that phone-based BPE systems tend to achieve higher accuracy while maintaining the simplicity of end-to-end methods. In the future, we can explore other types of subwords than BPE, and incorporate subword regularization [42] which is shown to improve character-based subword systems.

⁵Trained on 193K utterances with the small architecture, the original Espnet recipe obtains 8.5%/17.3% WER on SWBD/CALLHM.

6. References

- [1] S. J. Young, D. Kernshaw, J. Odell, D. Ollason, V. Valtchev, and P. Woodland, "The HTK book version 2.2," Tech. Rep., 1999.
- [2] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, J. Silovsky, G. Stemmer, and K. Vesely, "The Kaldi speech recognition toolkit," in *ASRU*, 2011.
- [3] R. Sennrich, B. Haddow, and A. Birch, "Neural machine translation of rare words with subword units," in *ACL*, 2016.
- [4] M. Schuster and K. Nakajima, "Japanese and korean voice search," in *ICASSP*, 2012.
- [5] A. Graves and N. Jaitly, "Towards end-to-end speech recognition with recurrent neural networks," in *ICML*, 2014.
- [6] Y. Miao, M. Gowayed, and F. Metze, "EESSEN: End-to-end speech recognition using deep RNN models and WFST-based decoding," in *ASRU*, 2015.
- [7] H. Sak, A. Senior, K. Rao, and F. Beaufays, "Fast and accurate recurrent neural network acoustic models for speech recognition," in *Interspeech*, 2015.
- [8] D. Amodei, R. Anubhai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, J. Chen, M. Chrzanowski, A. Coates, G. Diamos, E. Elsen, J. Engel, L. Fan, C. Fougner, A. Hannun, B. Jun, T. Han, P. LeGresley, X. Li, L. Lin, S. Narang, A. Ng, S. Ozair, R. Prenger, S. Qian, J. Raiman, S. Satheesh, D. Seetapun, S. Sengupta, C. Wang, Y. Wang, Z. Wang, B. Xiao, Y. Xie, D. Yogatama, J. Zhan, and Z. Zhu, "Deep speech 2: End-to-end speech recognition in english and mandarin," in *ICML*, 2016.
- [9] R. Collobert, C. Puhersch, and G. Synnaeve, "Wav2letter: An end-to-end convnet-based speech recognition system," 2016, arXiv:1609.03193 [cs.LG].
- [10] W. Chan, N. Jaitly, Q. V. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *ICASSP*, 2016.
- [11] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. Enrique Yalta Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, "ESPnet: End-to-end speech processing toolkit," in *Interspeech*, 2018.
- [12] A. Zeyer, K. Irie, R. Schluter, and H. Ney, "Improved training of end-to-end attention models for speech recognition," in *Interspeech*, 2018.
- [13] Y. He, T. N. Sainath, R. Prabhavalkar, I. McGraw, R. Alvarez, D. Zhao, D. Rybach, A. Kannan, Y. Wu, R. Pang, Q. Liang, D. Bhatia, Y. Shangguan, B. Li, G. Pundak, K. C. Sim, T. Bagby, S. Chang, K. Rao, and A. Gruenstein, "Streaming end-to-end speech recognition for mobile devices," in *ICASSP*, 2019.
- [14] Y. Wang, T. Chen, H. Xu, S. Ding, H. Lv, Y. Shao, N. Peng, L. Xie, S. Watanabe, and S. Khudanpur, "Espresso: A fast end-to-end neural speech recognition toolkit," in *ASRU*, 2019.
- [15] S. J. Young, J. J. Odell, and P. C. Woodland, "Tree-based state tying for high accuracy acoustic modelling," in *Proceedings of the Workshop on Human Language Technology*, 1994.
- [16] Y. Miao, M. Gowayed, X. Na, T. Ko, F. Metze, and A. Waibel, "An empirical exploration of CTC acoustic models," in *ICASSP*, 2016.
- [17] G. Zweig, C. Yu, J. Droppo, and A. Stolcke, "Advances in all-neural speech recognition," in *ICASSP*, 2017.
- [18] T. Kudo, "Subword regularization: Improving neural network translation models with multiple subword candidates," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, 2018.
- [19] S. Watanabe, T. Hori, S. Kim, J. Hershey, and T. Hayashi, "Hybrid CTC/attention architecture for end-to-end speech recognition," *IEEE Journal of Selected Topics in Signal Processing*, 2017.
- [20] "The CMU pronouncing dictionary," <http://www.speech.cs.cmu.edu/cgi-bin/cmudict>.
- [21] M. Bisani and H. Ney, "Joint-sequence models for grapheme-to-phoneme conversion," *Speech Communication*, vol. 50, no. 5, pp. 434–451, 2008.
- [22] K. Rao, F. Peng, H. Sak, and F. Beaufays, "Grapheme-to-phoneme conversion using long short-term memory recurrent neural networks," in *ICASSP*, 2015.
- [23] S. Toshniwal and K. Livescu, "Jointly learning to align and convert graphemes to phonemes with neural attention models," in *SLT*, 2016.
- [24] S. Yolchuyeva, G. Németh, and B. Gyires-Tóth, "Transformer based grapheme-to-phoneme conversion," in *Interspeech*, 2019.
- [25] C. Gulcehre, O. Firat, K. Xu, K. Cho, L. Barrault, H.-C. Lin, F. Bougares, H. Schwenk, and Y. Bengio, "On using monolingual corpora in neural machine translation," 2015, arXiv:1503.03535.
- [26] T. Hori, S. Watanabe, and J. R. Hershey, "Multilevel language modeling and decoding for open vocabulary end-to-end speech recognition," in *ASRU*, 2017.
- [27] H. Sak and K. Rao, "Multi-accent speech recognition with hierarchical grapheme based models," in *ICASSP*, 2017.
- [28] S. Toshniwal, H. Tang, L. Lu, and K. Livescu, "Multitask learning with low-level auxiliary tasks for encoder-decoder based speech recognition," in *Interspeech*, 2017.
- [29] K. Rao, H. Sak, and R. Prabhavalkar, "Exploring architectures, data and units for streaming end-to-end speech recognition with RNN-transducer," in *ASRU*, 2017.
- [30] C. Yu, C. Zhang, C. Weng, J. Cui, and D. Yu, "A multistage training framework for acoustic-to-word model," in *Interspeech*, 2018.
- [31] H. Xu, S. Ding, and S. Watanabe, "Improving end-to-end speech recognition with pronunciation-assisted sub-word modeling," in *ICASSP*, 2019.
- [32] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *NeurIPS*, 2017.
- [33] S. Karita, N. Chen, T. Hayashi, T. Hori, H. Inaguma, Z. Jiang, M. Someki, N. E. Y. Soplin, R. Yamamoto, X. Wang, S. Watanabe, T. Yoshimura, and W. Zhang, "A comparative study on transformer vs rnn in speech applications," in *ASRU*, 2019.
- [34] N.-Q. Pham, T.-S. Nguyen, J. Niehues, M. Müller, S. Stüker, and A. Waibel, "Very deep self-attention networks for end-to-end speech recognition," in *Interspeech*, 2019.
- [35] 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.
- [36] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *ICLR*, 2015.
- [37] "g2p-seq2seq," <https://github.com/cmuspinx/g2p-seq2seq>.
- [38] M. K. Baskar, L. Burget, S. Watanabe, M. Karafiat, T. Hori, and J. H. Cernocky, "Promising accurate prefix boosting for sequence-to-sequence ASR," in *ICASSP*, 2019.
- [39] J. Cui, C. Weng, G. Wang, J. Wang, P. Wang, C. Yu, D. Su, and D. Yu, "Improving attention-based end-to-end ASR systems with sequence-based loss functions," in *SLT*, 2018.
- [40] A. Zeyer, K. Irie, R. Schluter, and H. Ney, "A comprehensive analysis on attention models," in *Proc. IRASL Workshop, NeurIPS*, 2018.
- [41] K. J. Han, A. Chandrasekaran, J. Kim, and I. Lane, "The CAPIO 2017 conversational speech recognition system," arXiv:1801.00059 [cs.CL].
- [42] J. Drexler and J. Glass, "Subword regularization and beam search decoding for end-to-end automatic speech recognition," in *ICASSP*, 2019.