
WHO NEEDS WORDS? LEXICON-FREE SPEECH RECOGNITION

A PREPRINT

Tatiana Likhomanenko
Facebook AI Research
Menlo Park, USA
antares@fb.com

Gabriel Synnaeve
Facebook AI Research
New York, USA
gab@fb.com

Ronan Collobert
Facebook AI Research
Menlo Park, USA
locronan@fb.com

September 16, 2019

ABSTRACT

Lexicon-free speech recognition naturally deals with the problem of out-of-vocabulary (OOV) words. In this paper, we show that character-based language models (LM) can perform as well as word-based LMs for speech recognition, in word error rates (WER), even without restricting the decoding to a lexicon. We study character-based LMs and show that convolutional LMs can effectively leverage large (character) contexts, which is key for good speech recognition performance downstream. We specifically show that the lexicon-free decoding performance (WER) on utterances with OOV words using character-based LMs is better than lexicon-based decoding, both with character or word-based LMs.

Keywords speech recognition, beam-search decoder, out-of-vocabulary words, lexicon-free

1 Introduction

Character-based models permeated text classification [1], language modeling [2, 3, 4], machine translation [5, 6, 7, 8, 9, 10], and automatic speech recognition (ASR) [11, 12, 13]. However, most competitive ASR systems, character based or not, use a beam search decoder constrained on a word-level language model and lexicon [14, 15, 16, 17]. In recent works [18, 19], authors achieved competitive results with acoustic models (AM) and LMs operating on word pieces, and a lexicon-free decoder. To the best of our knowledge, the first ASR system to achieve competitive results with a character-based LM and without a lexicon (on Switchboard and WSJ) was [20], that our lexicon-free character-based ConvLM surpasses on WSJ (see Table 4).

The main advantage of a lexicon-free approach is that it allows the decoder to handle out-of-vocabulary (OOV) words: the decoder and the language model are responsible not only for scoring words but usually also restrict the vocabulary. Drawbacks sometimes include system complexity and most often poorer performance than in the lexicon based case. The first lexicon-free beam-search decoder aiming at dealing with OOV was benchmarked on Switchboard [21], although with a significantly worse word error rate (WER) than lexicon-based systems. Other recent works in this direction include [22, 23] on the English and [24, 25] on the Arabic and Finnish languages.

Here, we study a simple end-to-end ASR system combining a character level acoustic model with a character level language model through beam search. We show that it can yield competitive word error rates on the WSJ and Librispeech corporas, even without a lexicon. Finally, our model shows significant word error rates improvement on utterances that include out-of-vocabulary words.

2 Setup

Acoustic model (AM) We consider in this paper 1D gated convolutional neural networks [26, 15], trained to map speech features (log-mel filterbanks) to their corresponding letter transcription. The training criterion is the auto segmentation criterion (ASG) [27]. The token set contains 31 graphemes: the standard English alphabet, the apostrophe

and period, two repetition characters (e.g. the word `ann` is transcribed as `an1`), and a silence token (`l`) used as word boundary.

Language model (LM) Our language models are character-based. We evaluated n -gram language models, as well as gated convolutional language models (ConvLMs) [26], and show that with enough context these language models can match (in perplexity) their word-based counterparts. The LM training data was pre-processed to be consistent with the AM training data: the silence character (`l`) defines word boundaries, and repetition symbols are used when letter repetitions occur.

Beam-search decoder We extended the beam-search decoder from [15] to support character-level language models. Given a word transcription y , we denote $\text{AM}(y)$ the corresponding acoustic score and $P_{\text{LM}}(y)$ the corresponding LM likelihood. The beam-search decoder generates transcriptions by finding the argmax of the following score [27]:

$$\text{AM}(y) + \alpha \log P_{\text{LM}}(y) + \beta |y| + \gamma \sum_{i=1}^T [\pi_i = 'l'], \quad (1)$$

where $\pi = \pi_1, \dots, \pi_T$ is the sequence of letters corresponding to the transcription y . The hyper-parameters α , β and γ weight the language model, word penalty and silence penalty, respectively. The decoder has two additional parameters: (i) the beam size and (ii) a beam threshold, controlling which hypothesis can make it to the beam.

Experiments We experiment with the Wall Street Journal (WSJ) dataset [28, 29, 30] (about 81 hours of transcribed audio data) and the Librispeech dataset [31] (1000 hours with clean and noisy speech).

3 Language Model Experiments

We consider word-level n -gram and ConvLM-based language models as baseline, and compare them in word perplexity with their character-level counterpart.

Data preparation Language models for both WSJ and Librispeech are trained with the corresponding language model data available for these datasets. For word-level model training, we keep all words (162K) for WSJ and use only the most frequent 200K (out of 900K) words for Librispeech (words appearing less than 10 times are dropped). Words outside this scope are replaced by `unknown`.

n -gram LMs All models were trained with KenLM [32]. For both Librispeech and WSJ, we trained 4-gram word-level language models as a baseline. For character-level language models, we study how the context width impacts perplexity, training n -grams ranging from 5 to 20. For large values of n , we pruned the models by thresholding rarely-occurring n -grams: 6,7,8-grams appearing once, 9-grams appearing once or twice, and all n -grams for $n \geq 10$ appearing ≤ 3 times were dropped.

ConvLMs As a baseline for ConvLMs, we use the ‘GCNN-14B’ word-level LM architecture from [26], which achieved competitive results on several language model benchmarks. This network contains 14 convolutional-residual blocks with a growing number of channels and gated linear units as activation functions, resulting in 318M parameters and an effective receptive field of 57 tokens. An *adaptive softmax* [33] over words follows the convolutional layers.

For character-level LMs, we consider both the ‘GCNN-14B’ architecture and a deeper variant (20 convolutional layers) dubbed ‘GCNN-20B’, with a larger receptive field of 81 tokens. For both configurations a *softmax* (over letters) follows the last convolutional layer. The resulting number of parameters was 163M for ‘GCNN-14B’ and 224M for ‘GCNN-20B’. Dropout is used at each convolutional and linear layer: with probabilities 0.2 for WSJ and 0.1 for Librispeech.

ConvLMs were trained with the *fairseq* toolkit¹ [34], using Nesterov accelerated gradient descent [35] with fixed learning rate. Gradient clipping and weight normalization are used following [26].

¹<https://github.com/pytorch/fairseq>

Table 1: Word perplexity on the validation set of WSJ. For character-level LMs, we display lower and upper perplexity bounds. For models marked with *, pruning is applied during training. The receptive field is in characters.

Language Model	Size	Rcp. field (char.)	nov93dev
word 4-gram	878 M	32	156
char 5-gram	3.3 M	5	(714, 1285)
char 10-gram	447 M	10	(211, 243)
char 15-gram*	546 M	15	(185, 205)
char 15-gram	3.5 G	15	(186, 203)
char 20-gram*	836 M	20	(177, 196)
char 20-gram	9.7 G	20	(180, 196)
word GCNN-14B	1.1 G	450	80
char GCNN-14B	936 M	57	(79, 95)
char GCNN-20B	1.3 G	81	(76, 90)

Table 2: Word perplexity on the validation sets of Librispeech, shown as lower and upper bounds for character-level LMs. For models marked with *, pruning is applied during training. The average receptive fields are 31 characters and 439 characters for the word 4-gram and word GCNN-14B LMs, respectively.

Language Model	Size	dev-clean	dev-other
word 4-gram*	13 G	148	137
char 5-gram	7.7 M	(597, 1000)	(515, 869)
char 10-gram	2.5 G	(198, 230)	(180, 210)
char 15-gram*	6.5 G	(165, 180)	(151, 165)
char 17-gram*	9.5 G	(163, 178)	(148, 162)
char 20-gram*	13 G	(162, 177)	(147, 161)
word GCNN-14B*	1.8 G	57	58
char GCNN-14B	936 M	(72, 88)	(70, 84)
char GCNN-20B	1.3 G	(65, 76)	(64, 75)

Word-level perplexity for character-level LMs To compare word and character-level LMs, we estimate a word-level perplexity² for character-level LMs. The word probability can be estimated with:

$$P(\text{word}|\mathcal{C}) = P(l_1|\mathcal{C}) \prod_{i=2}^n P(l_i|\mathcal{C} l_1 \dots l_{i-1}), \tag{2}$$

where l_1, \dots, l_n are letters in a word and the last letter l_n is a silence symbol ‘|’ with which the word finishes, \mathcal{C} — the previous context. However, this approach does not take into account that word-level LMs are constrained to a fixed-size lexicon, while character-based LMs have virtually an infinite vocabulary. We thus re-normalize (2), taking into account only words from the word-level LM vocabulary \mathcal{V} :

$$P_{\mathcal{V}}(\text{word}|\mathcal{C}) = \frac{P(\text{word}|\mathcal{C})}{\sum_{\text{word}_i \in \mathcal{V}} P(\text{word}_i|\mathcal{C})}. \tag{3}$$

For large vocabulary \mathcal{V} the denominator in (3) is computationally expensive. The probability (2) can be used as an upper bound of (3) (see, for example, [36]) while the lower bound can be obtained by taking in the denominator of (3) the sum over most probable (by word-level LM) words which cover 95% of word-level LM distribution.

We then exclude from the perplexity computation words which are not presented in the word-level LM vocabulary \mathcal{V} (n -gram and ConvLM models have the same \mathcal{V}).³

Results The comparison of different language models is presented in Tables 1 and 2 for WSJ and Librispeech, respectively. It can be seen that (as expected) increasing the context decreases perplexity for both n -gram models and ConvLMs. With n -grams, pruning is critical to avoid overfitting. On both benchmarks character-level language models

²Defined as $\left(\prod_{i=1}^N p(\text{word}_i)\right)^{-\frac{1}{N}}$, N is a number of words in data.

³Only about 20 word occurrences for *nov93dev* in WSJ, and around 200 (300) for *clean (other)* in Librispeech.

already have similar performance for n -grams with $n \geq 15$, and are clearly outperformed by ConvLMs. With enough context, character-level LMs appear to be in the same ballpark as word-level LMs.

4 ASR Experiments

In this section, we decode the output of a single acoustic model trained on WSJ or Librispeech, through a beam-search procedure constrained by the LMs trained in Section 3. Both AM training and decoding were performed with the *wav2letter++* open source library⁴ [37]. The decoder was adapted to support character-level LMs, as well as lexicon-free decoding, alleviating the need for a word lexicon while decoding.

Table 3: Word and character error rates (%) on Librispeech data.

Language Model	Lexicon	dev-clean		dev-other		test-clean		test-other	
		WER	CER	WER	CER	WER	CER	WER	CER
CAPIO (Ensemble) [38]	yes	2.7		7.6		3.2		7.6	
CAPIO (Single) ⁵ [38]	yes	3.0		8.3		3.5		8.6	
Learnable front-end [17]	yes	3.2		10.1		3.4		11.2	
DeepSpeech2 ⁶ [14]	yes					5.2		12.7	
Sequence-to-sequence [18]	yes	3.5		11.5		3.8		12.8	
word 4-gram*	yes	4.6	1.8	13.3	6.8	5.0	1.8	14.4	7.2
char 15-gram*	yes	4.8	1.9	14.1	7.1	5.2	1.9	14.9	7.4
char 20-gram*	yes	4.8	1.9	13.9	7.0	5.1	1.9	14.8	7.3
char 15-gram*	no	4.9	1.9	14.1	7.1	5.1	1.9	14.9	7.4
char 20-gram*	no	4.8	1.9	13.9	6.9	4.9	1.8	14.7	7.2
word GCNN-14B*	yes	3.1	1.3	10.3	5.4	3.4	1.3	11.1	5.7
char GCNN-20B	yes	3.4	1.4	10.7	5.7	3.7	1.4	11.7	6.1
char GCNN-20B	no	3.4	1.4	10.8	5.7	3.5	1.3	11.7	6.1

Table 4: Word and character error rates (%) on WSJ data.

Language Model	Lexicon	nov93dev		nov92	
		WER	CER	WER	CER
Regular LF-MMI [20]	yes			2.8	
DeepSpeech2 ⁶ [14]	yes			3.1	
CNN-BLSTM-HMM ⁷ [39]	yes	6.6		3.5	
Learnable front-end [17]	yes	6.8		3.5	
EE-LF-MMI (word LM) ⁸ [20]	yes			4.1	
EE-LF-MMI (char LM) ⁸ [20]	no			5.4	
word 4-gram	yes	8.5	3.4	5.5	2.2
char 15-gram*	yes	9.4	3.7	6.2	2.3
char 20-gram*	yes	9.3	3.6	5.9	2.2
char 15-gram*	no	9.5	3.8	6.2	2.3
char 20-gram*	no	9.4	3.7	6.1	2.3
word GCNN-14B	yes	6.5	2.7	4.3	1.8
char GCNN-20B	yes	6.4	2.8	3.6	1.5
char GCNN-20B	no	6.4	2.7	3.6	1.5

⁴<https://github.com/facebookresearch/wav2letter>

⁵Speaker adaptation; pronunciation lexicon

⁶12k hours AM train set and common crawl LM

⁷Speaker adaptation; 3k acoustic states

⁸Data augmentation; n -gram LM

Data preparation For WSJ, we consider the standard subsets *si284*, *nov93dev* and *nov92* for training, validation and test, respectively. For Librispeech, all the available training data was used for training. Validation and test were achieved according to the available two configurations (*clean* for clean speech and *other* for "noisy" speech). All hyper-parameter tuning was performed on validation sets, and only final performance was evaluated on the test sets. We kept the original 16kHz sampling rate and computed log-mel filterbanks with 40 (for Librispeech) or 80 (for WSJ) coefficients for a 25ms sliding window, strided by 10ms. All features are normalized to have zero mean and unit variance per input sequence before feeding into the neural network. No data augmentation or speaker adaptation was performed.

Acoustic model training Models are trained with stochastic gradient descent (SGD), gradient clipping [40] and weight normalization [41]. We followed [15] for the architecture choices, picking the "high dropout" model with 19 convolutional layers for Librispeech, and the lighter version with 17 layers for WSJ. Batch size was set to 4 and 16, for Librispeech, and WSJ respectively.

Tuning the beam-search decoder Hyper-parameters of the decoder were selected via a random search. A large fixed beam size and beam threshold were set before running a random search. The LM weight α was randomly sampled from the interval $(0, 5)$, the word β and silence γ penalties were sampled from the interval $(-5, 5)$ for both language model types. For each configuration and dataset up to 100 attempts of random search were run. Hyper-parameters that lead to the best WER were chosen for the final evaluation on the test sets.

Results Models are evaluated in Word Error Rate (WER) and Character Error Rate (CER), as reported in Table 4 for WSJ and in Table 3 for Librispeech. The ASR system based on either n -gram or ConvLM character-level language model achieves performance similar to its word-level language model configuration both on WSJ and Librispeech. Furthermore, the lexicon-free ASR systems, where the beam-search is not conditioned by a word-level lexicon vocabulary, are very close to their lexicon-based counterparts for both types of language models, n -gram and ConvLM. On the WSJ test set (see Table 4), decoding with the character-level ConvLM (even in a lexicon-free setup) leads to better performance than with word-level and achieves state-of-the-art results.

Beam size analysis For each utterance we can define an effective beam size as the maximum position index in the sorted beam over all frames for the final transcription. In other words, selecting a beam size larger than the effective beam size should not affect decoding. In Figure 1, we show that the effective beam size for the character-level models is significantly smaller than one for the word-level models. When decoding with word-level LMs, a large (2000-8000) beam size is often needed, while for the character-level LMs, a beam size of 2000 is always enough. This property looks promising from the computational point of view, when switching to character-level LMs.

Lexicon-free decoding analysis We investigated how out-of-vocabulary (OOV) words are transcribed by our lexicon-free decoder, as those words cannot be output by a standard beam-search decoder conditioned with a word-level lexicon. The OOV words we consider are not present in the lexicon vocabulary, meaning these words are beyond the top- k most frequent words chosen in our lexicon, or did not even appear in the acoustic and language models training/validation

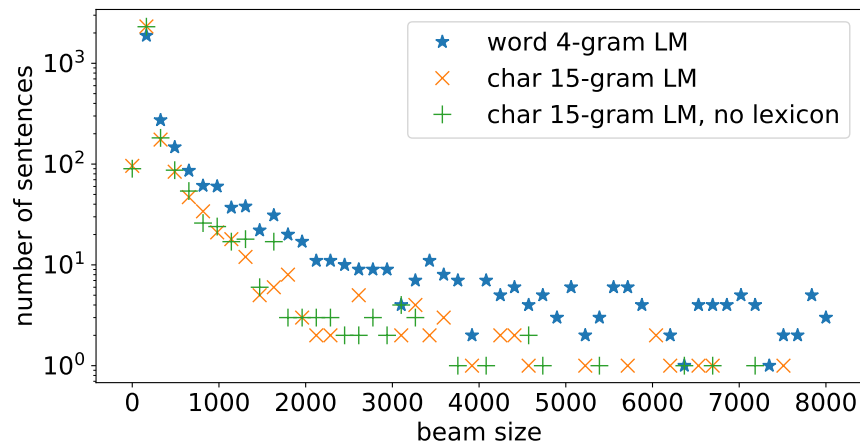


Figure 1: Effective beam size for the beam-search decoder (8000 beam size is set) with different LMs on ‘other’ part of Librispeech validation set.

Table 5: Word and character error rates (%) for Librispeech: ‘in vocabulary’ (IV) utterances and ‘out-of-vocabulary’ (OOV) utterances. Proportions for test-clean data: IV — 93.3% utterances, OOV — 6.7% utterances (176 among 2620). Proportions for test-other data: IV — 91.6% utterances, OOV — 8.4% utterances (247 among 2939).

Language Model	Lexicon	test-clean IV		test-other IV		test-clean OOV		test-other OOV	
		WER	CER	WER	CER	WER	CER	WER	CER
word 4-gram*	yes	4.4	1.6	13.5	6.8	11.2	3.8	22.2	10.8
char 15-gram*	yes	4.6	1.7	14.0	7.0	11.4	3.8	22.8	10.9
char 20-gram*	yes	4.5	1.7	13.9	6.9	11.8	3.7	22.4	10.8
char 15-gram*	no	4.7	1.7	14.1	7.0	9.4	3.2	21.4	10.4
char 20-gram*	no	4.5	1.7	13.9	6.8	9.4	3.2	21.7	10.3
word ConvLM*	yes	2.7	1.0	10.1	5.2	9.9	3.5	19.4	9.4
char ConvLM-20B	yes	3.1	1.2	10.9	5.6	10.4	3.4	18.8	9.6
char ConvLM-20B	no	3.1	1.2	11.0	5.7	7.7	2.7	18.3	9.4

Table 6: Librispeech transcriptions for decoder with n -gram LMs: target (T), word-level LM and lexicon (W), character-level LM and lexicon (C), lexicon-free decoder with character-level LM (F). Underlined words are not presented in the lexicon.

T fauchelevant limped along behind the hearse in a very contented frame of mind
W lochleven limped along behind the hearse in a very contented frame of mind
C lochleven limped along behind the hearse in a very contented frame of mind
F fauchelevant limped along behind the hearse in a very contented frame of mind
T ... he did not want to join his own friends that is sergey ivanovitch stepan arkadyevitch sviazhsy and ...
W ... he did not come to join his own friends that a soldier ivanovitch step on markovitch the sky and ...
C ... he did not own to join his own friends that a sojer ivanovitch step on radovitch totski and ...
F ... he did not own to join his own friends that a sojer ivanovitch stepan arkadyevitch tievski and ...
T menahem king of israel had died and was succeeded by his son pekahiah
W many a king of israel had died and was succeeded by his son pekah
C many a king of israel had died and was succeeded by his son pekah
F many a king of israel had died and was succeeded by his son pekaiah

sets. In Table 5, we evaluated WER and CER on the isolated “out-of-vocabulary” utterances, which contain at least one OOV word. For comparison, we also report performance on “in vocabulary” utterances (IV), which contain only words present in the lexicon. The lexicon-free decoder performs significantly better on the utterances that include OOV words, while holding competitive performance on the IV utterances. The lexicon-free decoder with an n -gram language model recognizes up to 25% (28.5%) and 13.5% (13.5%) OOV words (occurrences) for the *clean* and *other* test parts of Librispeech, respectively. With a ConvLM language model, this performance raises up to 31% (33%) and 10% (8%) OOV words (occurrences) for the *clean* and *other* test parts of Librispeech, respectively. A few examples of decoded transcriptions are reported in Table 6.

5 Conclusion

We built an ASR system with a beam-search decoder based on a character-level language model that achieves performance close to the state-of-the-art among end-to-end models. We also showed that transcribing with a lexicon-free beam-search decoder achieves a performance similar to transcribing with a beam-search decoder conditioned by a word-level lexicon, itself already close to the state-of-the-art. Moreover, lexicon-free ASR naturally handles out-of-vocabulary words: it has significantly better performance on the utterances with out-of-vocabulary words than when the system is conditioned by a word-level lexicon.

6 Acknowledgements

We thank Qiantong Xu for helpful discussions.

7 Addendum

In the Interspeech2019 submission paper lower bounds on the word perplexity for character-level language models were computed excluding end of sentence token while estimating the denominator from (3). In Tables 1 and 2 updated results are presented where the end of sentence is correctly accounted. This does not affect any outcomes and conclusions in the paper.

The code for experiments is open-sourced and available at https://github.com/facebookresearch/wav2letter/tree/master/recipes/models/lexicon_free.

References

- [1] Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In *Advances in neural information processing systems*, pages 649–657, 2015.
- [2] Gerasimos Potamianos and Frederick Jelinek. A study of n-gram and decision tree letter language modeling methods. *Speech Communication*, 24(3):171–192, 1998.
- [3] Bob Carpenter. Scaling high-order character language models to gigabytes. In *Proceedings of the Workshop on Software*, pages 86–99. Association for Computational Linguistics, 2005.
- [4] Yoon Kim, Yacine Jernite, David Sontag, and Alexander M Rush. Character-aware neural language models. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [5] David Vilar, Jan-T. Peter, and Hermann Ney. Can We Translate Letters? In *Proceedings of the Second Workshop on Statistical Machine Translation*, StatMT ’07, pages 33–39, Stroudsburg, PA, USA, 2007. Association for Computational Linguistics.
- [6] Preslav Nakov and Jörg Tiedemann. Combining word-level and character-level models for machine translation between closely-related languages. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers-Volume 2*, pages 301–305. Association for Computational Linguistics, 2012.
- [7] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*, 2015.
- [8] Wang Ling, Isabel Trancoso, Chris Dyer, and Alan W Black. Character-based neural machine translation. *arXiv preprint arXiv:1511.04586*, 2015.
- [9] Marta R Costa-Jussa and José AR Fonollosa. Character-based neural machine translation. *arXiv preprint arXiv:1603.00810*, 2016.
- [10] Junyoung Chung, Kyunghyun Cho, and Yoshua Bengio. A character-level decoder without explicit segmentation for neural machine translation. *arXiv preprint arXiv:1603.06147*, 2016.
- [11] Alex Graves and Navdeep Jaitly. Towards end-to-end speech recognition with recurrent neural networks. In *International Conference on Machine Learning*, pages 1764–1772, 2014.
- [12] William Chan, Navdeep Jaitly, Quoc Le, and Oriol 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)*, pages 4960–4964. IEEE, 2016.
- [13] Dzmitry Bahdanau, Jan Chorowski, Dmitriy Serdyuk, Philemon Brakel, and Yoshua Bengio. End-to-end attention-based large vocabulary speech recognition. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4945–4949. IEEE, 2016.
- [14] Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, et al. Deep speech 2: End-to-end speech recognition in english and mandarin. In *International Conference on Machine Learning*, pages 173–182, 2016.
- [15] V Liptchinsky, G Synnaeve, and R Collobert. Letterbased speech recognition with gated convnets. *CoRR*, vol. abs/1712.09444, 1, 2017.
- [16] Ying Zhang, Mohammad Pezeshki, Philémon Brakel, Saizheng Zhang, Cesar Laurent Yoshua Bengio, and Aaron Courville. Towards end-to-end speech recognition with deep convolutional neural networks. *arXiv preprint arXiv:1701.02720*, 2017.
- [17] Neil Zeghidour, Qiantong Xu, Vitaliy Liptchinsky, Nicolas Usunier, Gabriel Synnaeve, and Ronan Collobert. Fully Convolutional Speech Recognition. *arXiv preprint arXiv:1812.06864*, 2018.

- [18] Albert Zeyer, Kazuki Irie, Ralf Schlüter, and Hermann Ney. Improved training of end-to-end attention models for speech recognition. *arXiv preprint arXiv:1805.03294*, 2018.
- [19] Chung-Cheng Chiu, Tara N Sainath, Yonghui Wu, Rohit Prabhavalkar, Patrick Nguyen, Zhifeng Chen, Anjali Kannan, Ron J Weiss, Kanishka Rao, Ekaterina Gonina, et al. State-of-the-art speech recognition with sequence-to-sequence models. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4774–4778. IEEE, 2018.
- [20] Hossein Hadian, Hossein Sameti, Daniel Povey, and Sanjeev Khudanpur. End-to-end speech recognition using lattice-free MMI. In *Proc. Interspeech*, pages 12–16, 2018.
- [21] Andrew Maas, Ziang Xie, Dan Jurafsky, and Andrew Ng. Lexicon-free conversational speech recognition with neural networks. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 345–354, 2015.
- [22] Kyu-yeon Hwang and Wonyong Sung. Character-level incremental speech recognition with recurrent neural networks. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5335–5339. IEEE, 2016.
- [23] Thomas Zenkel, Ramon Sanabria, Florian Metz, Jan Niehues, Matthias Sperber, Sebastian Stüker, and Alex Waibel. Comparison of decoding strategies for ctc acoustic models. *arXiv preprint arXiv:1708.04469*, 2017.
- [24] Abdelrahman Ahmed, Yasser Hifny, Khaled Shaalan, and Sergio Toral. Lexicon free Arabic speech recognition recipe. In *International Conference on Advanced Intelligent Systems and Informatics*, pages 147–159. Springer, 2016.
- [25] Peter Smit, Siva Reddy Gangireddy, Seppo Enarvi, Sami Virpioja, Mikko Kurimo, et al. Character-based units for unlimited vocabulary continuous speech recognition. In *ASRU 2017–IEEE Workshop on Automatic Speech Recognition & Understanding*, 2017.
- [26] Yann N Dauphin, Angela Fan, Michael Auli, and David Grangier. Language modeling with gated convolutional networks. *arXiv preprint arXiv:1612.08083*, 2016.
- [27] Ronan Collobert, Christian Puhres, and Gabriel Synnaeve. Wav2letter: an end-to-end convnet-based speech recognition system. *arXiv preprint arXiv:1609.03193*, 2016.
- [28] John Garofolo, David Graff, Doug Paul, and David Pallett. CSR-I (wsj0) complete. Web Download. Philadelphia: Linguistic Data Consortium, vol. LDC93S6A, 1993.
- [29] Linguistic Data Consortium and others. CSR-II (wsj1) complete. Linguistic Data Consortium, Philadelphia, vol. LDC94S13A, 1994.
- [30] Philip C Woodland, Julian J Odell, Valtcho Valtchev, and Steve J Young. Large vocabulary continuous speech recognition using HTK. In *Acoustics, Speech, and Signal Processing, 1994. ICASSP-94., 1994 IEEE International Conference on*, volume 2, pages II–125. Ieee, 1994.
- [31] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: an ASR corpus based on public domain audio books. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pages 5206–5210. IEEE, 2015.
- [32] Kenneth Heafield, Ivan Pouzyrevsky, Jonathan H Clark, and Philipp Koehn. Scalable modified Kneser-Ney language model estimation. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 690–696, 2013.
- [33] Edouard Grave, Armand Joulin, Moustapha Cissé, Hervé Jégou, et al. Efficient softmax approximation for GPUs. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1302–1310. JMLR.org, 2017.
- [34] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. Convolutional Sequence to Sequence Learning. *ArXiv e-prints*, May 2017.
- [35] Yurii Nesterov. A method for unconstrained convex minimization problem with the rate of convergence $O(1/k^2)$. In *Doklady AN USSR*, volume 269, pages 543–547, 1983.
- [36] Ben Krause, Liang Lu, Iain Murray, and Steve Renals. Multiplicative LSTM for sequence modelling. *arXiv preprint arXiv:1609.07959*, 2016.
- [37] Vineel Pratap, Awni Hannun, Qiantong Xu, Jeff Cai, Jacob Kahn, Gabriel Synnaeve, Vitaliy Liptchinsky, and Ronan Collobert. wav2letter++: The Fastest Open-source Speech Recognition System. *arXiv preprint arXiv:1812.07625*, 2018.
- [38] Kyu J Han, Akshay Chandrashekar, Jungsuk Kim, and Ian Lane. The CAPIO 2017 conversational speech recognition system. *arXiv preprint arXiv:1801.00059*, 2017.

- [39] William Chan and Ian Lane. Deep recurrent neural networks for acoustic modelling. *arXiv preprint arXiv:1504.01482*, 2015.
- [40] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *International Conference on Machine Learning*, pages 1310–1318, 2013.
- [41] Tim Salimans and Durk P Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In *Advances in Neural Information Processing Systems*, pages 901–909, 2016.