

# Who Needs Words? Lexicon-Free Speech Recognition

Tatiana Likhomanenko, Gabriel Synnaeve, Ronan Collobert

Facebook AI Research, Menlo Park & New York, USA

antares@fb.com, gab@fb.com, locronan@fb.com

## 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, with character or word-based LMs.

**Index Terms:** 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, 18]. To the best of our knowledge, the first ASR system to achieve competitive results with a character-based model and without a lexicon (on Switchboard and WSJ) was [19], 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 [20], although with a significantly worse word error rate (WER) than lexicon-based systems. Other recent works in this direction on the Arabic and Finnish languages include [21, 22].

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 [23, 15], trained to map speech features (log-mel filterbanks) to their corresponding letter transcription. The training criterion is the auto segmentation criterion (ASG) [24]. 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 (`()`) 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) [23], 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 (`()`) 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  $AM(y)$  the corresponding acoustic score and  $P_{LM}(y)$  the corresponding LM likelihood. The beam-search decoder generates transcriptions by finding the argmax of the following score [24]:

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

where  $\pi = \pi_1, \dots, \pi_T$  is the sequence of letters corresponding to the transcription  $y$ . The hyper-parameters  $\alpha$ ,  $\beta$  and  $\gamma$  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 [25] (about 81 hours of transcribed audio data) and the Librispeech dataset [26] (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 [27]. For both Librispeech and WSJ, we trained 4-gram word-level language models as a baseline. For character-level language mod-

Table 1: Word perplexity on validation data of WSJ corpus. For character-level LMs, lower and upper perplexity bounds are shown. For models marked with star (\*) pruning is applied during training. The receptive field of all those language models is in characters.

Language Model	Size	Rcp. field (char.)	nov93dev
word 4-gram	878 M	32	156
char 5-gram	3.3 M	5	(927, 1285)
char 10-gram	447 M	10	(221, 243)
char 15-gram*	546 M	15	(186, 205)
char 15-gram	3.5 G	15	(186, 203)
char 20-gram*	836 M	20	(178, 196)
char 20-gram	9.7 G	20	(180, 196)
word GCNN-14B	1.1 G	450	80
char GCNN-14B	936 M	57	(76, 95)
char GCNN-20B	1.3 G	81	(74, 90)

Table 2: Word perplexity on validation sets of Librispeech. For character-level LMs lower and upper bounds are presented. For models marked with star (\*) 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	(748, 1000)	(649, 869)
char 10-gram	2.5 G	(210, 230)	(191, 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	(76, 88)	(73, 84)
char GCNN-20B	1.3 G	(66, 76)	(64, 75)

els, 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  $\leq 3$  times were dropped.

**ConvLMs** As a baseline for ConvLMs, we use the ‘GCNN-14B’ word-level LM architecture from [23], 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* [28] 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.

ConvLM were trained with the *fairseq* toolkit<sup>1</sup> [29], using Nesterov accelerated gradient descent [30] with fixed learning rate. Gradient clipping and weight normalization are used following [23].

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

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

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  $\mathcal{V}$ :

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

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, [31]) 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  $\mathcal{V}$ ).<sup>2</sup>

**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 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 language models trained in Section 3. Both acoustic model training and decoding were performed with the *wav2letter++* open source library<sup>3</sup> [32]. The decoder was adapted to support character-level LMs, as well as lexicon-free decoding, alleviating the need for a word lexicon while decoding.

**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 accord-

<sup>1</sup><https://github.com/facebookresearch/fairseq>

<sup>2</sup>Only about 20 word occurrences for *nov93dev* WSJ and around 200 (300) for *clean (other)* Librispeech.

<sup>3</sup><https://github.com/facebookresearch/wav2letter>

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 (Single) <sup>4</sup>	yes	3.0		8.3		3.6		8.6	
CAPIO (Ensemble) [33]	yes	2.7		7.6		3.2		7.6	
DeepSpeech2 <sup>5</sup>	yes					5.8		12.7	
Sequence-to-sequence [17]	yes	3.5		11.5		3.8		12.8	
Learnable front-end [18]	yes	3.2		10.1		3.4		11.2	
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	?	?	?	?
char GCNN-20B	no	3.4	1.4	10.8	5.7	3.5	1.3	11.7	6.1

ing 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 [34] and weight normalization [35]. 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 language model weight  $\alpha$  was randomly sampled from the interval  $(0, 5)$ , the word  $\beta$  and silence  $\gamma$  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.

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

Language Model	Lexicon	nov93dev		nov92	
		WER	CER	WER	CER
EE-LF-MMI (char LM) <sup>6</sup>	no			5.4	
EE-LF-MMI (word LM) <sup>6</sup>	yes			4.1	
CNN-BLSTM-HMM <sup>7</sup>	yes	6.6		3.5	
DeepSpeech2 <sup>5</sup>	yes	5.0		3.6	
Learnable front-end [18]	yes	6.8		3.5	
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

**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 language models.

<sup>4</sup>Speaker adaptation; pronunciation lexicon [33]<sup>5</sup>12k hours AM train set and common crawl LM [14]<sup>6</sup>Data augmentation;  $n$ -gram LM [19]<sup>7</sup>Speaker adaptation; 3k acoustic states [36]

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	4.44	1.65	13.46	6.76	11.16	3.84	22.16	10.84
char 15-gram*	yes	4.61	1.72	14.01	7.01	11.43	3.83	22.76	10.92
char 20-gram*	yes	4.49	1.68	13.87	6.88	11.78	3.74	22.40	10.83
char 15-gram*	no	4.70	1.74	14.09	7.01	9.37	3.24	21.38	10.36
char 20-gram*	no	4.50	1.68	13.93	6.84	9.41	3.18	21.69	10.27
word ConvLM*	yes	2.74	1.05	10.09	5.21	9.94	3.52	19.39	9.41
char ConvLM-20B	yes	?	?	?	?	10.38	3.41	18.84	9.56
char ConvLM-20B	no	3.14	1.21	10.96	5.69	7.74	2.66	18.34	9.38

Table 6: Librispeech transcription examples with  $n$ -gram LMs: T — target, W — word-level LM decoder with lexicon, C — character-level LM decoder with lexicon, F — lexicon-free decoder with character-level LM. Underlined words are not presented in the lexicon.

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




Figure 1: Distribution of the effective beam size for the beam-search decoder with different LMs on ‘other’ part of Librispeech validation set. Beam size 8000 is used for all cases.

**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 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 the 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. References

- [1] X. Zhang, J. Zhao, and Y. LeCun, "Character-level convolutional networks for text classification," in *Advances in neural information processing systems*, 2015, pp. 649–657.
- [2] G. Potamianos and F. Jelinek, "A study of n-gram and decision tree letter language modeling methods," *Speech Communication*, vol. 24, no. 3, pp. 171–192, 1998.
- [3] B. Carpenter, "Scaling high-order character language models to gigabytes," in *Proceedings of the Workshop on Software*. Association for Computational Linguistics, 2005, pp. 86–99.
- [4] Y. Kim, Y. Jernite, D. Sontag, and A. M. Rush, "Character-aware neural language models," in *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [5] D. Vilar, J.-T. Peter, and H. Ney, "Can we translate letters?" in *Proceedings of the Second Workshop on Statistical Machine Translation*, ser. StatMT'07. Stroudsburg, PA, USA: Association for Computational Linguistics, 2007, pp. 33–39. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1626355.1626360>
- [6] P. Nakov and J. 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*. Association for Computational Linguistics, 2012, pp. 301–305.
- [7] R. Sennrich, B. Haddow, and A. Birch, "Neural machine translation of rare words with subword units," *arXiv preprint arXiv:1508.07909*, 2015.
- [8] W. Ling, I. Trancoso, C. Dyer, and A. W. Black, "Character-based neural machine translation," *arXiv preprint arXiv:1511.04586*, 2015.
- [9] M. R. Costa-Jussa and J. A. Fonollosa, "Character-based neural machine translation," *arXiv preprint arXiv:1603.00810*, 2016.
- [10] J. Chung, K. Cho, and Y. Bengio, "A character-level decoder without explicit segmentation for neural machine translation," *arXiv preprint arXiv:1603.06147*, 2016.
- [11] A. Graves and N. Jaitly, "Towards end-to-end speech recognition with recurrent neural networks," in *International Conference on Machine Learning*, 2014, pp. 1764–1772.
- [12] 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.
- [13] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, "End-to-end attention-based large vocabulary speech recognition," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 4945–4949.
- [14] D. Amodei, S. Ananthanarayanan, R. Anubhai, J. Bai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, Q. Cheng, G. Chen *et al.*, "Deep speech 2: End-to-end speech recognition in english and mandarin," in *International Conference on Machine Learning*, 2016, pp. 173–182.
- [15] V. Liptchinsky, G. Synnaeve, and R. Collobert, "Letter-based speech recognition with gated convnets," *CoRR*, vol. abs/1712.09444, vol. 1, 2017.
- [16] Y. Zhang, M. Pezeshki, P. Brakel, S. Zhang, C. L. Y. Bengio, and A. Courville, "Towards end-to-end speech recognition with deep convolutional neural networks," *arXiv preprint arXiv:1701.02720*, 2017.
- [17] A. Zeyer, K. Irie, R. Schlüter, and H. Ney, "Improved training of end-to-end attention models for speech recognition," *arXiv preprint arXiv:1805.03294*, 2018.
- [18] N. Zeghidour, Q. Xu, V. Liptchinsky, N. Usunier, G. Synnaeve, and R. Collobert, "Fully convolutional speech recognition," 12 2018.
- [19] H. Hadian, H. Sameti, D. Povey, and S. Khudanpur, "End-to-end speech recognition using lattice-free mmi," in *Proc. Interspeech*, 2018, pp. 12–16.
- [20] A. Maas, Z. Xie, D. Jurafsky, and A. 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*, 2015, pp. 345–354.
- [21] A. Ahmed, Y. Hifny, K. Shaalan, and S. Toral, "Lexicon free arabic speech recognition recipe," in *International Conference on Advanced Intelligent Systems and Informatics*. Springer, 2016, pp. 147–159.
- [22] P. Smit, S. R. Gangireddy, S. Enarvi, S. Virpioja, M. Kurimo *et al.*, "Character-based units for unlimited vocabulary continuous speech recognition," in *ASRU 2017—IEEE Workshop on Automatic Speech Recognition & Understanding*, 2017.
- [23] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier, "Language modeling with gated convolutional networks," *arXiv preprint arXiv:1612.08083*, 2016.
- [24] R. Collobert, C. Puhersch, and G. Synnaeve, "Wav2letter: an end-to-end convnet-based speech recognition system," *arXiv preprint arXiv:1609.03193*, 2016.
- [25] P. C. Woodland, J. J. Odell, V. Valtchev, and S. J. Young, "Large vocabulary continuous speech recognition using htk," in *Acoustics, Speech, and Signal Processing, 1994. ICASSP-94., 1994 IEEE International Conference on*, vol. 2. Ieee, 1994, pp. II–125.
- [26] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*. IEEE, 2015, pp. 5206–5210.
- [27] K. Heafield, I. Pouzyrevsky, J. H. Clark, and P. 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)*, vol. 2, 2013, pp. 690–696.
- [28] E. Grave, A. Joulin, M. Cissé, H. Jégou *et al.*, "Efficient softmax approximation for gpus," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR.org, 2017, pp. 1302–1310.
- [29] J. Gehring, M. Auli, D. Grangier, D. Yarats, and Y. N. Dauphin, "Convolutional Sequence to Sequence Learning," *ArXiv e-prints*, May 2017.
- [30] Y. Nesterov, "A method for unconstrained convex minimization problem with the rate of convergence  $O(1/k^2)$ ," in *Doklady AN USSR*, vol. 269, 1983, pp. 543–547.
- [31] B. Krause, L. Lu, I. Murray, and S. Renals, "Multiplicative lstm for sequence modelling," *arXiv preprint arXiv:1609.07959*, 2016.
- [32] Q. X. J. C. J. K. G. S. V. L. R. C. Vineel Pratap, Awni Hannun, "wav2letter++: The fastest open-source speech recognition system," *CoRR*, vol. abs/1812.07625, 2018. [Online]. Available: <https://arxiv.org/abs/1812.07625>
- [33] K. J. Han, A. Chandrashekar, J. Kim, and I. Lane, "The capio 2017 conversational speech recognition system," *arXiv preprint arXiv:1801.00059*, 2017.
- [34] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training recurrent neural networks," in *International Conference on Machine Learning*, 2013, pp. 1310–1318.
- [35] T. Salimans and D. P. Kingma, "Weight normalization: A simple reparameterization to accelerate training of deep neural networks," in *Advances in Neural Information Processing Systems*, 2016, pp. 901–909.
- [36] W. Chan and I. Lane, "Deep recurrent neural networks for acoustic modelling," *arXiv preprint arXiv:1504.01482*, 2015.