



Unsupervised and Efficient Vocabulary Expansion for Recurrent Neural Network Language Models in ASR

Yerbolat Khassanov, Eng Siong Chng

Rolls-Royce@NTU Corporate Lab, Nanyang Technological University, Singapore

yerbolat002@edu.ntu.sg, aseschn@ntu.edu.sg

Abstract

In automatic speech recognition (ASR) systems, recurrent neural network language models (RNNLM) are used to rescore a word lattice or N -best hypotheses list. Due to the expensive training, the RNNLM's vocabulary set accommodates only small shortlist of most frequent words. This leads to suboptimal performance if an input speech contains many out-of-shortlist (OOS) words.

An effective solution is to increase the shortlist size and retrain the entire network which is highly inefficient. Therefore, we propose an efficient method to expand the shortlist set of a pretrained RNNLM without incurring expensive retraining and using additional training data. Our method exploits the structure of RNNLM which can be decoupled into three parts: input projection layer, middle layers, and output projection layer. Specifically, our method expands the word embedding matrices in projection layers and keeps the middle layers unchanged. In this approach, the functionality of the pretrained RNNLM will be correctly maintained as long as OOS words are properly modeled in two embedding spaces. We propose to model the OOS words by borrowing linguistic knowledge from appropriate in-shortlist words. Additionally, we propose to generate the list of OOS words to expand vocabulary in unsupervised manner by automatically extracting them from ASR output.

Index Terms: vocabulary expansion, recurrent neural network, language model, speech recognition, word embedding

1. Introduction

The language model (LM) plays an important role in automatic speech recognition (ASR) system. It ensures that recognized output hypotheses obey the linguistic regularities of the target language. The LMs are employed at two different stages of the state-of-the-art ASR pipeline: decoding and rescoring. At the decoding stage, a simple model such as count-based n -gram [1] is used as a background LM to produce initial word lattice. At the rescoring stage, this word lattice or N -best hypotheses list extracted from it is rescored by a more complex model such as recurrent neural network language model (RNNLM) [2, 3, 4].

Due to the simplicity and efficient implementation, the count-based n -gram LM is trained over a large vocabulary set, typically in the order of hundreds of thousands words. On the other hand, computationally expensive RNNLM is usually trained with a small subset of most frequent words known as in-shortlist (IS) set, typically in the order of tens of thousands words, whereas the remaining words are deemed out-of-shortlist (OOS) and jointly modeled by single node $\langle unk \rangle$ [5]. Since the probability mass of $\langle unk \rangle$ node is shared by many words, it will poorly represent properties of the individual words leading to unreliable probability estimates of OOS words. Moreover, these estimates tend to be very small which makes RNNLM biased in favor of hypotheses mostly

comprised of IS words. Consequently, if an input speech with many OOS words is supplied to ASR system, the performance of RNNLM will be suboptimal.

An effective solution is to increase the IS set size and retrain the entire network. However, this approach is highly inefficient as training RNNLM might take from several days up to several weeks depending on the scale of application [6]. Moreover, additional textual data containing training instances of OOS words would be required, which is difficult to find for rare domain-specific words. Therefore, the effective and efficient methods to expand RNNLM's vocabulary coverage is of great interest [7].

In this work, we propose an efficient method to expand the vocabulary of pretrained RNNLM without incurring expensive retraining and using additional training data. To achieve this, we exploit the structure of RNNLM which can be decoupled into three parts: 1) input projection layer, 2) middle layers and 3) output projection layer as shown in figure 1. The input and output projection layers are defined by input and output word embedding matrices that perform linear word transformations from high to low and low to high dimensions, respectively. The middle layers are a non-linear function used to generate high-level feature representation of contextual information. Our method expands the vocabulary coverage of RNNLM by inserting new words into input and output word embedding matrices, and keeping the parameters of middle layers unchanged. This method keeps the functionality of pretrained RNNLM intact as long as new words are properly modeled in input and output word embedding spaces. We propose to model the new words by borrowing linguistic knowledge from other "similar" words present in word embedding matrices.

Furthermore, the list of OOS words to expand vocabulary can be generated either in supervised or unsupervised manners. For example, in supervised manner, they can be manually collected by the human expert. Whereas in unsupervised manner, a subset of most frequent words from OOS set can be selected. In this work, we propose to generate the list of OOS words in unsupervised manner by automatically extracting them from ASR output. The motivation is that the background LM usually covers much larger vocabulary, and hence, during the decoding stage it will produce a word lattice which will contain the most relevant OOS words that might be present in test data.

We evaluate our method by rescoring N -best list output from the state-of-the-art TED¹ talks ASR system. The experimental results show that vocabulary expanded RNNLM achieves 4% relative word error rate (WER) improvement over the conventional RNNLM. Moreover, 7% relative WER improvement is achieved over the strong Kneser-Ney smoothed 5-gram model used to rescore the word lattice. Importantly, all these improvements are achieved without using additional training data and by incurring very little computational cost.

¹<https://www.ted.com/>

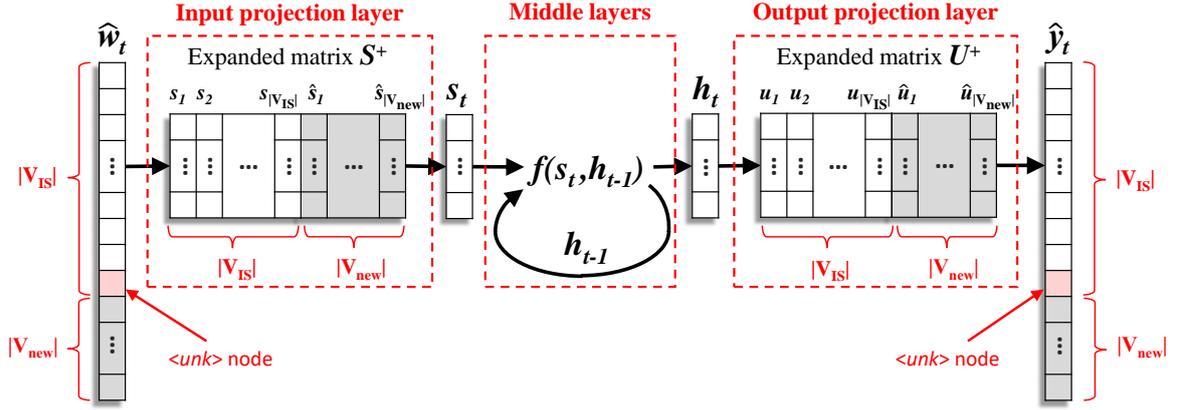


Figure 1: RNNLM architecture after vocabulary expansion.

The rest of the paper is organized as follows. The related works on vocabulary coverage expansion of RNNLMs are reviewed in section 2. Section 3 briefly describes the RNNLM architecture. Section 4 presents the proposed methodology to increase the IS set of RNNLM by expanding the word embedding matrices. In section 5, the experiment setup and obtained results are discussed. Lastly, section 6 concludes the paper.

2. Related works

This section briefly describes popular approaches to expand the vocabulary coverages of RNNLMs. These approaches mostly focus on intelligently redistributing the probability mass of $\langle unk \rangle$ node among OOS words, optimizing the training speed for large-vocabulary models or training sub-word level RNNLM. These approaches can also be used in combination.

Redistributing probability mass of $\langle unk \rangle$: Park et al. [5] proposed to expand the vocabulary coverage by gathering all OOS words under special node $\langle unk \rangle$ and explicitly modeling it together with IS words, see figure 1. This is a standard scheme commonly employed in the state-of-the-art RNNLMs.

The probability mass of $\langle unk \rangle$ node is then can be redistributed among OOS words by using statistics of simpler LMs such as count-based n -gram model as follows:

$$\tilde{P}_R(w_{t+1}|h_t) = \begin{cases} P_R(w_{t+1}|h_t) & w_{t+1} \in V_{IS} \\ \beta(w_{t+1}|h_t) P_R(\langle unk \rangle|h_t) & otherwise \end{cases} \quad (1)$$

$$\beta(w_{t+1}|h_t) = \frac{P_N(w_{t+1}|h_t)}{\sum_{w \notin V_{IS}} P_N(w|h_t)} \quad (2)$$

where $P_R()$ and $P_N()$ are conditional probability estimates according to RNNLM and n -gram LM respectively, for some word w_{t+1} given context h_t . The n -gram model is trained with whole vocabulary set V , whereas RNNLM is trained with smaller in-shortlist subset $V_{IS} \subset V$. The $\beta()$ is a normalization coefficient used to ensure the sum-to-one constraint of obtained probability function $\tilde{P}_R()$.

Later, [3] proposed to uniformly redistribute the probability mass of $\langle unk \rangle$ token among OOS words as follows:

$$\tilde{P}_R(w_{t+1}|h_t) = \begin{cases} P_R(w_{t+1}|h_t) & w_{t+1} \in V_{IS} \\ \frac{P_R(\langle unk \rangle|h_t)}{|V \setminus V_{IS}| + 1} & otherwise \end{cases} \quad (3)$$

where ‘\’ symbol is the set difference operation. In this way, the vocabulary coverage of RNNLM is expanded to the full vocabulary size $|V|$ without relying on the statistics of simpler LMs.

Training speed optimization: Rather than expanding vocabulary of the pretrained model, this group of studies focuses on speeding-up the training of large-vocabulary RNNLMs.

One of the most effective ways to speed up the training of RNNLMs is to approximate the softmax function. The softmax function is used to normalize obtained word scores to form a probability distribution, hence, it requires scores $y_{w'}$ of every word in the vocabulary:

$$\text{softmax}(y_w) = \frac{\exp(y_w)}{\sum_{w' \in V_{IS}} \exp(y_{w'})} \quad (4)$$

Consequently, its computational cost is proportional to the number of words in the vocabulary and it dominates the training of the whole model which is the network’s main bottleneck [8].

Many techniques have been proposed to approximate the softmax computation. The most popular ones include hierarchical softmax [9, 10, 11], importance sampling [12, 13] and noise contrastive estimation [8, 14]. The comparative study of these techniques can be found in [7, 15]. Other techniques, besides softmax function approximation, to speed up the training of large-vocabulary models can be found in [16].

Sub-word level RNNLM: Another effective method to expand the vocabulary coverage is to train a sub-word level RNNLM. Different from standard word-level RNNLMs, they model finer linguistic units such as characters [17] or syllables [18], hence, a larger range of words will be covered. Furthermore, character-level RNNLM doesn’t suffer from the OOS problem, though, it performs worse than word-level models² [18]. Recently, there has been a lot of research effort aiming to train the hybrid of word and sub-word level models where promising results are obtained [19, 6, 20].

3. RNNLM architecture

The conventional RNNLM architecture can be decoupled into three parts: 1) input projection layer, 2) middle layers and 3) output projection layer, as shown in figure 1. The input projection layer is defined by input word embedding matrix $S \in \mathbb{R}^{d_s \times |V_{IS}|}$ used to transform the one-hot encoding representation of word $w_t \in \mathbb{R}^{|V_{IS}|}$ at time t into lower dimensional continuous space vector $s_t \in \mathbb{R}^{d_s}$, where d_s is input word embedding vector dimension:

$$s_t = S w_t \quad (5)$$

This vector s_t and compressed context vector from previous time step $h_{t-1} \in \mathbb{R}^{d_h}$ are then merged by non-linear middle

²At least for English.

layer, which can be represented as function $f()$, to produce a new compressed context vector $h_t \in \mathbb{R}^{d_h}$, where d_h is context vector dimension:

$$h_t = f(s_t, h_{t-1}) \quad (6)$$

The function $f()$ can be simple activation units such as sigmoid and hyperbolic tangent, or more complex units such as LSTM [3] and GRU [21]. The middle layer can also be formed by stacking several such functions.

The compressed context vector h_t is then supplied to output projection layer where it is transformed into higher dimension vector $y_t \in \mathbb{R}^{|V_{vis}|}$ by output word embedding matrix $U \in \mathbb{R}^{d_h \times |V_{vis}|}$:

$$y_t = U^T h_t \quad (7)$$

The entries of output vector y_t represent the scores of words to follow the context h_t . These scores are then normalized by softmax function to form probability distribution (eq. (4)).

4. Vocabulary expansion

This section describes our proposed method to expand the vocabulary coverage of pretrained RNNLM. Our method is based on the observation that input and output projection layers learn the word embedding matrices, and middle layers learn the mapping from the input word embedding vectors to compressed context vectors. Thus, by modifying the input and output word embedding matrices to accommodate new words, we can expand the vocabulary coverage of RNNLM. Meanwhile, the parameters of middle layers are kept unchanged which allows us to avoid expensive retraining. This approach will preserve the linguistic regularities encapsulated within pretrained RNNLM as long as the new words are properly modeled in input and output embedding spaces. To model the new words, we will use word embedding vectors of “similar” words present in V_{IS} set.

The proposed method has three main challenges: 1) how to find relevant OOS words for vocabulary expansion, 2) criteria to select “similar” candidate words to model a target OOS word and 3) how to expand the word embedding matrices. The details are discussed in section 4.1, 4.2 and 4.3, respectively.

4.1. Finding relevant OOS words

The first step to vocabulary expansion is finding relevant OOS words. This step is important as expanding vocabulary with irrelevant words absent in the input test data is ineffective. The relevant OOS words can be found either in supervised or unsupervised manners. For example, in supervised manner, they can be manually collected by human expert. In unsupervised manner, the subset of most frequent OOS words can be selected.

In this work, we employed an unsupervised method where relevant OOS words are automatically extracted from the ASR output. The reason is that at the decoding stage a background LM covering very large vocabulary set is commonly employed. Subsequently, the generated word lattice will contain the most relevant OOS words that might be present in the input test data.

4.2. Selecting candidate words

Given a list of relevant OOS words, let’s call it set V_{new} , the next step is to select candidate words that will be used to model each of them. The selected candidates must be present in V_{IS} set and should be similar to the target OOS word in both semantic meaning and syntactic behavior. Selecting inadequate candidates might deteriorate the linguistic regularities incorporated within pretrained RNNLM, thus, they should be carefully

inspected. In natural language processing, many effective techniques exist that can find appropriate candidate words satisfying conditions mentioned above [22, 23, 24].

4.3. Expanding word embedding matrices

This section describes our proposed approach to expand the word embedding matrices S and U .

The matrix S : This matrix holds input word embedding vectors $S = [s_1 \dots s_{|V_{IS}|}]$ and it’s used to transform words from discrete form into lower dimensional continuous space. In this space, vectors of “similar” words are clustered together [25]. Moreover, these vectors have been shown to capture meaningful semantic and syntactic features of language [26]. Subsequently, if two words have similar semantic and syntactic roles, their embedding vectors are expected to belong to the same cluster. As such, words in a cluster can be used to approximate a new word that belongs to the same cluster.

For example, let’s consider a scenario where we want to add a new word Astana, the capital of Kazakhstan, to the vocabulary set of an existing RNNLM. Here, we can select candidate words from V_{IS} with similar semantic and syntactic roles such as London and Paris. Specifically, we extract the input word embedding vectors of selected candidates $C = \{s_{London}, s_{Paris}\}$ and combine them to form a new input word embedding vector $\hat{s}_{astana} \in \mathbb{R}^{d_s}$ as follows:

$$\hat{s}_{astana} = \frac{\sum_{s \in C} m_s s}{|C|} \quad (8)$$

where m_s is some normalized metric used to weigh the candidates. We repeat this procedure for all words in set V_{new} .

Obtained new input word embedding vectors are then used to form matrix $\hat{S} = [\hat{s}_1 \dots \hat{s}_{|V_{new}|}]$ where $\hat{S} \in \mathbb{R}^{d_s \times |V_{new}|}$, which is appended to the initial matrix $S \in \mathbb{R}^{d_s \times |V_{IS}|}$ to form the expanded matrix:

$$S^+ = [S \ \hat{S}] \quad (9)$$

where $S^+ \in \mathbb{R}^{d_s \times (|V_{IS}| + |V_{new}|)}$.

The input word vector $w_t \in \mathbb{R}^{|V_{IS}|}$ should be also expanded to accommodate the new words from V_{new} , which results in the new input vector $\hat{w}_t \in \mathbb{R}^{|V_{IS}| + |V_{new}|}$. The input vector and input word embedding matrix after expansion are depicted in figure 1.

The matrix U : This matrix holds the output word embedding vectors $U = [u_1 \dots u_{|V_{IS}|}]$ where $u \in \mathbb{R}^{d_h}$. These vectors are compared against context vector h_t using the dot product to determine the score of the next possible word w_{t+1} [8]. Intuitively, for a given context h_t , the interchangeable words with similar semantic and syntactic roles should have similar scores to follow it. Therefore, in the output word embedding space, interchangeable words should belong to the same cluster. Subsequently, we can use the same procedure and candidates which were used to expand matrix S to model the new words in output word embedding space. However, this time we operate in the column space of matrix U . The output vector and output word embedding matrix after expansion are depicted in figure 1.

5. Experiment

This section describes experiments conducted to evaluate the efficacy of proposed vocabulary expansion method for pretrained RNNLMs on ASR task. The ASR system is built by Kaldi [27] speech recognition toolkit on the first release of TED-LIUM [28] speech corpus. To highlight the importance of vocabulary expansion, we train the LMs on generic-domain text corpus One Billion Word Benchmark (OBWB) [29].

Table 1: *The characteristics of TED-LIUM corpus*

Characteristics	Train	Dev	Test
No. of talks	774	8	11
No. of words	1.3M	17.7k	27.5k
Total duration	122hr	1.5hr	2.5hr

As the baseline LMs, we trained two state-of-the-art models, namely the modified Kneser-Ney smoothed 5-gram (KN5) [30] and recurrent LSTM network (LSTM) [3]. We call our system VE-LSTM which is constructed by expanding the vocabulary of the baseline LSTM. The performance of these three models is evaluated using both perplexity and WER.

Experiment setup: The TED-LIUM corpus is comprised of monologue talks given by experts on specific topics, its characteristics are given in table 1. Its train set was used to build the acoustic model with the ‘nnet3+chain’ setup of Kaldi including the latest developments. Its dev set was used to tune hyper-parameters such as the number of candidates to use to model the new words, word insertion penalty and LM scale. The test set was used to compare the performance of proposed VE-LSTM and two baseline models. Additionally, the TED-LIUM corpus has a pre-designed pronunciation lexicon of 150k words which was also used as a vocabulary set for baseline LMs.

The OBWB corpus consists of text collected from various domains including the news and parliamentary speeches. Its train set contains around 700M words and is used to train both baseline LMs. Its validation set of size 141k words was used to stop the training of LSTM model.

The baseline KN5 was trained using SRILM [31] toolkit with 150k vocabulary. It was used to rescore the word lattice and 300-best list. Its pruned³ version KN5_pruned was used as a background LM during the decoding stage.

The baseline LSTM was trained as a four-layer network similar to [3] using our own implementation in PyTorch [32]. The LSTM explicitly models only the 10k most frequent words⁴ of 150k vocabulary set. The remaining 140k words are modeled by uniformly distributing the probability mass of <unk> node using equation (3). Thus, the input w_t and output y_t vector sizes are $10k + 1$ which we call as V_{IS} set. Hence, the baseline LSTM theoretically models the same 150k vocabulary set as KN5. The OOS rate with respect to dev and test sets are 6.8% and 5.5%, respectively. The input d_s and output d_h word embedding vector dimensions were set to 300 and 1500, respectively. The parameters of the model are learned by truncated backpropagation through time algorithm (BPTT) [33] unrolled for 10 steps. For regularization, we applied 50% dropout on the non-recurrent connections as suggested by [34].

The VE-LSTM model is obtained by expanding the vocabulary of baseline LSTM with OOS words extracted from the ASR output. For example, to construct the VE-LSTM model for the test set, we collect the list of OOS words V_{new} from the recognized hypotheses of the test set. For each OOS word in V_{new} , we then select the appropriate set of candidate words C . The selection criteria will be explained later. Next, selected candidates are used to model the new input and output word embedding vectors of target OOS words as in equation (8). For simplicity, we didn’t weigh the selected candidates. Lastly, these generated new vectors are appended to the input and output word embedding matrices of baseline LSTM model, see figure 1. Consequently, the obtained VE-LSTM will

³The pruning coefficient is 10^{-7} .

⁴Plus the beginning <s> and end of sentence </s> symbols.

Table 2: *The perplexity and WER results on dev and test sets of TED-LIUM corpus*

LM	Perplexity		WER (%)		
	Dev	Test	Rescore	Dev	Test
KN5_pruned	384	341	-	13.05	12.82
KN5	237	218	Lattice 300-best	10.85 11.03	10.49 10.79
LSTM	211	185	300-best	10.63	10.18
VE-LSTM	163	150	300-best	10.44	9.77

explicitly model $10k + 1 + |V_{new}|$ words, whereas the remaining $140k - |V_{new}|$ words are modeled by uniformly distributing the probability mass of <unk> node using equation (3).

To select candidate words we used the classical skip-gram model [23]. The skip-gram model is trained with default parameters on OBWB corpus covering all unique words. Typically, when presented with a target OOS word, the skip-gram model returns a list of “similar” words. From this list, we only select top eight⁵ words which are present in V_{IS} set.

Results: The experiment results are given in table 2. We evaluated LMs on perplexity and WER measure. The perplexity was computed on the reference data and the OOS words for vocabulary expansion were extracted from the reference data as well. The perplexity results computed on the test set show that VE-LSTM significantly outperforms KN5 and LSTM models by 31% and 18% relative, respectively.

For WER experiment, the KN5 is evaluated on both lattice and 300-best rescoring tasks. The LSTM and VE-LSTM are evaluated only on 300-best rescoring task. We tried to extract the OOS words for vocabulary expansion from different N -best lists. Interestingly, the best result is achieved when they are extracted from the 1-best. The reason is that 1-best hypothesis list contains high confidence score words, hence, OOS words extracted from it will be reliable. Whereas using other N -best lists will result in unreliable OOS words which confuse the VE-LSTM model. The VE-LSTM outperforms the baseline LSTM model by 4% relative WER. Compared to KN5 used to rescore the word lattice, 7% relative WER improvement is achieved. Such improvements suggest that the proposed vocabulary expansion method is effective.

6. Conclusion

In this paper, we have proposed an efficient vocabulary expansion method for pretrained RNNLM. Our method which modifies the input and output projection layers, while keeping the parameters of middle layers unchanged was shown to be feasible. It was found that extracting OOS words for vocabulary expansion from the ASR output is effective when high confidence words are selected. Our method achieved significant perplexity and WER improvements on the state-of-the-art ASR system over two strong baseline LMs. Importantly, the expensive re-training was avoided and no additional training data was used. We believe that our approach of manipulating input and output projection layers is general enough to be applied to other neural network models with similar architectures.

7. Acknowledgements

This work was conducted within the Rolls-Royce@NTU Corporate Lab with support from the National Research Foundation (NRF) Singapore under the Corp Lab@University Scheme.

⁵This number was tuned on the dev set.

8. References

- [1] J. T. Goodman, "A bit of progress in language modeling," *Computer Speech & Language*, vol. 15, no. 4, pp. 403–434, 2001.
- [2] T. Mikolov, M. Karafiát, L. Burget, J. Cernocký, and S. Khudanpur, "Recurrent neural network based language model," in *Interspeech*, 2010.
- [3] M. Sundermeyer, R. Schlüter, and H. Ney, "Lstm neural networks for language modeling," in *Interspeech*, 2012.
- [4] X. Liu, X. Chen, Y. Wang, M. J. Gales, and P. C. Woodland, "Two efficient lattice rescoring methods using recurrent neural network language models," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 8, pp. 1438–1449, 2016.
- [5] J. Park, X. Liu, M. J. Gales, and P. C. Woodland, "Improved neural network based language modelling and adaptation," in *Interspeech*, 2010.
- [6] R. Jozefowicz, O. Vinyals, M. Schuster, N. Shazeer, and Y. Wu, "Exploring the limits of language modeling," *arXiv preprint arXiv:1602.02410*, 2016.
- [7] W. Chen, D. Grangier, and M. Auli, "Strategies for training large vocabulary neural language models," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, 2016, pp. 1975–1985.
- [8] A. Mnih and Y. W. Teh, "A fast and simple algorithm for training neural probabilistic language models," in *In Proceedings of the International Conference on Machine Learning*, 2012.
- [9] F. Morin and Y. Bengio, "Hierarchical probabilistic neural network language model," in *AISTATS*, 2005, pp. 246–252.
- [10] A. Mnih and G. E. Hinton, "A scalable hierarchical distributed language model," in *Advances in neural information processing systems*, 2009, pp. 1081–1088.
- [11] T. Mikolov, S. Kombrink, L. Burget, J. ernock, and S. Khudanpur, "Extensions of recurrent neural network language model," in *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011, pp. 5528–5531.
- [12] Y. Bengio and J.-S. Senécal, "Quick training of probabilistic neural nets by importance sampling," in *AISTATS*, 2003, pp. 1–9.
- [13] —, "Adaptive importance sampling to accelerate training of a neural probabilistic language model," *IEEE Transactions on Neural Networks*, vol. 19, no. 4, pp. 713–722, 2008.
- [14] X. Chen, X. Liu, M. J. F. Gales, and P. C. Woodland, "Recurrent neural network language model training with noise contrastive estimation for speech recognition," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5411–5415.
- [15] E. Grave, A. Joulin, M. Cissé, D. Grangier, and H. Jégou, "Efficient softmax approximation for gpus," in *In Proceedings of the International Conference on Machine Learning*, 2017.
- [16] T. Mikolov, A. Deoras, D. Povey, L. Burget, and J. ernock, "Strategies for training large scale neural network language models," in *2011 IEEE Workshop on Automatic Speech Recognition Understanding*, 2011, pp. 196–201.
- [17] I. Sutskever, J. Martens, and G. E. Hinton, "Generating text with recurrent neural networks," in *In Proceedings of the International Conference on Machine Learning*, 2011, pp. 1017–1024.
- [18] T. Mikolov, I. Sutskever, A. Deoras, H.-S. Le, S. Kombrink, and J. Cernocký, "Subword language modeling with neural networks," 2011.
- [19] Y. Kim, Y. Jernite, D. Sontag, and A. M. Rush, "Character-aware neural language models," in *AAAI*, 2016, pp. 2741–2749.
- [20] H. Xu, K. Li, Y. Wang, J. Wang, S. Kang, X. Chen, D. Povey, and S. Khudanpur, "Neural network language modeling with letter-based features and importance sampling," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
- [21] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Gated feedback recurrent neural networks," in *In Proceedings of the International Conference on Machine Learning*, 2015, pp. 2067–2075.
- [22] G. Miller and C. Fellbaum, "Wordnet: An electronic lexical database," 1998.
- [23] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.
- [24] J. Pennington, R. Socher, and C. Manning, "Glove: Global vectors for word representation," in *Proceedings of the 2014 conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1532–1543.
- [25] R. Collobert and J. Weston, "A unified architecture for natural language processing: Deep neural networks with multitask learning," in *In Proceedings of the International Conference on Machine Learning*, 2008, pp. 160–167.
- [26] T. Mikolov, W.-t. Yih, and G. Zweig, "Linguistic regularities in continuous space word representations," in *HLT-NAACL*, 2013, pp. 746–751.
- [27] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, "The kaldi speech recognition toolkit," in *IEEE 2011 Workshop on Automatic Speech Recognition and Understanding*, 2011.
- [28] A. Rousseau, P. Deléglise, and Y. Esteve, "Ted-lium: an automatic speech recognition dedicated corpus," in *LREC*, 2012, pp. 125–129.
- [29] C. Chelba, T. Mikolov, M. Schuster, Q. Ge, T. Brants, P. Koehn, and T. Robinson, "One billion word benchmark for measuring progress in statistical language modeling," in *Interspeech*, 2014.
- [30] S. F. Chen and J. Goodman, "An empirical study of smoothing techniques for language modeling," *Computer Speech & Language*, vol. 13, no. 4, pp. 359–394, 1999.
- [31] A. Stolcke, "Srlm - an extensible language modeling toolkit," in *Interspeech*, 2002.
- [32] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," 2017.
- [33] R. J. Williams and J. Peng, "An efficient gradient-based algorithm for on-line training of recurrent network trajectories," *Neural computation*, vol. 2, no. 4, pp. 490–501, 1990.
- [34] W. Zaremba, I. Sutskever, and O. Vinyals, "Recurrent neural network regularization," *CoRR*, 2014.