

Learning Spoken Language Representations with Neural Lattice Language Modeling

Chao-Wei Huang Yun-Nung (Vivian) Chen

Department of Computer Science and Information Engineering

National Taiwan University, Taipei, Taiwan

r07922069@ntu.edu.tw y.v.chen@ieee.org

Abstract

Pre-trained language models have achieved huge improvement on many NLP tasks. However, these methods are usually designed for written text, so they do not consider the properties of spoken language. Therefore, this paper aims at generalizing the idea of language model pre-training to lattices generated by recognition systems. We propose a framework that trains neural lattice language models to provide contextualized representations for spoken language understanding tasks. The proposed two-stage pre-training approach reduces the demands of speech data and has better efficiency. Experiments on intent detection and dialogue act recognition datasets demonstrate that our proposed method consistently outperforms strong baselines when evaluated on spoken inputs.¹

1 Introduction

The task of spoken language understanding (SLU) aims at extracting useful information from spoken utterances. Typically, SLU can be decomposed with a two-stage method: 1) an accurate automatic speech recognition (ASR) system transcribes the input speech into texts, and then 2) language understanding techniques are applied to the transcribed texts. These two modules can be developed separately, so most prior work developed the backend language understanding systems based on manual transcripts (Yao et al., 2014; Guo et al., 2014; Mesnil et al., 2014; Goo et al., 2018).

Despite the simplicity of the two-stage method, prior work showed that a tighter integration between two components can lead to better performance. Researchers have extended the ASR 1-best results to n-best lists or word confusion networks in order to preserve the ambiguity of the transcripts.

¹The source code is available at: <https://github.com/MiuLab/Lattice-ELMo>.




Figure 1: Illustration of a lattice.

(Tur et al., 2002; Hakkani-Tür et al., 2006; Henderson et al., 2012; Tür et al., 2013; Masumura et al., 2018). Another line of research focused on using lattices produced by ASR systems. Lattices are directed acyclic graphs (DAGs) that represent multiple recognition hypotheses. An example of ASR lattice is shown in Figure 1. Ladhak et al. (2016) introduced LatticeRNN, a variant of recurrent neural networks (RNNs) that generalize RNNs to lattice-structured inputs in order to improve SLU. Zhang and Yang (2018) proposed a similar idea for Chinese name entity recognition. Sperber et al. (2019); Xiao et al. (2019); Zhang et al. (2019) proposed extensions to enable the transformer model (Vaswani et al., 2017) to consume lattice inputs for machine translation. Huang and Chen (2019) proposed to adapt the transformer model originally pre-trained on written texts to consume lattices in order to improve SLU performance. Buckman and Neubig (2018) also found that utilizing lattices that represent multiple granularities of sentences can improve language modeling.

With recent introduction of large pre-trained language models (LMs) such as ELMo (Peters et al., 2018), GPT (Radford, 2018) and BERT (Devlin et al., 2019), we have observed huge improvements on natural language understanding tasks. These models are pre-trained on large amount of written texts so that they provide the downstream tasks with high-quality representations. However, applying these models to the spoken scenarios poses

several discrepancies between the pre-training task and the target task, such as the domain mismatch between written texts and spoken utterances with ASR errors. It has been shown that fine-tuning the pre-trained language models on the data from the target tasks can mitigate the domain mismatch problem (Howard and Ruder, 2018; Chronopoulou et al., 2019). Siddhant et al. (2018) focused on pre-training a language model specifically for spoken content with huge amount of automatic transcripts, which requires a large collection of in-domain speech.

In this paper, we propose a novel spoken language representation learning framework, which focuses on learning contextualized representations of lattices based on our proposed lattice language modeling objective. The proposed framework consists of two stages of LM pre-training to reduce the demands for lattice data. We conduct experiments on benchmark datasets for spoken language understanding, including intent classification and dialogue act recognition. The proposed method consistently achieves superior performance, with relative error reduction ranging from 3% to 42% compare to pre-trained sequential LM.

2 Neural Lattice Language Model

The two-stage framework that learns contextualized representations for spoken language is proposed and detailed below.

2.1 Problem Formulation

In the SLU task, the model input is an utterance X containing a sequence of words $X = [x_1, x_2, \dots, x_{|X|}]$, and the goal is to map X to its corresponding class y . The inputs can also be stored in a lattice form, where we use edge-labeled lattices in this work. A lattice $L = \{N, E\}$ is defined by a set of $|N|$ nodes $N = \{n_1, n_2, \dots, n_{|N|}\}$ and a set of $|E|$ transitions $E = \{e_1, e_2, \dots, e_{|E|}\}$. A weighted transition is defined as $e = \{prev[e], next[e], w[e], P(e)\}$, where $prev[e]$ and $next[e]$ denote the previous node and next node respectively, $w[e]$ denotes the associated word, and $P(e)$ denotes the transition probability. We use $in[n]$ and $out[n]$ to denote the sets of incoming and outgoing transitions of a node n . $L_{<n} = \{N_{<n}, E_{<n}\}$ denotes the sub-lattice which consists of all paths between the starting node and a node n .

2.2 LatticeRNN

The LatticeRNN (Ladhak et al., 2016) model generalizes sequential RNN to lattice-structured inputs. It traverses the nodes and transitions of a lattice in a topological order. For each transition e , LatticeRNN takes $w[e]$ as input and the representation of its previous node $h[prev[e]]$ as the previous hidden state, and then produces a new hidden state of e , $h[e]$. The representation of a node $h[n]$ is obtained by pooling the hidden states of the incoming transitions. In this work, we employ the *WeightedPool* variant proposed by Ladhak et al. (2016), which computes the node representation as

$$h[n] = \sum_{e \in in[n]} P(e) \cdot h[e].$$

Note that we can represent any sequential text as a linear-chain lattice, so LatticeRNN can be seen as a strict generalization of RNNs to DAG-like structures. This property enables us to initialize the weights in a LatticeRNN with the weights of a RNN as long as they use the same recurrent cell.

2.3 Lattice Language Modeling

Language models usually estimate $p(X)$ by factorizing it into

$$p(X) = \prod_{t=0}^{|X|} p(x_t | X_{<t}),$$

where $X_{<t} = [x_1, \dots, x_{t-1}]$ denotes the previous context. Training a LM is essentially asking the model to predict a distribution of the next word given the previous words. We extend the sequential LM analogously to *lattice language modeling*, where the model is expected to predict the next transitions of a node n given $L_{<n}$. The ground truth distribution is therefore defined as:

$$p(w | L_{<n}) = \begin{cases} P(e), & \text{if } \exists e \in out[n] \text{ s.t. } w[e] = w \\ 0, & \text{otherwise.} \end{cases}$$

LatticeRNN is adopted as the backbone of our lattice language model. Since the node representation $h[n]$ encodes all information of $L_{<n}$, we pass $h[n]$ to a linear decoder to obtain the distribution of next transitions:

$$p_{\theta}(w | h[n]) = \text{softmax}(W^T h[n]),$$

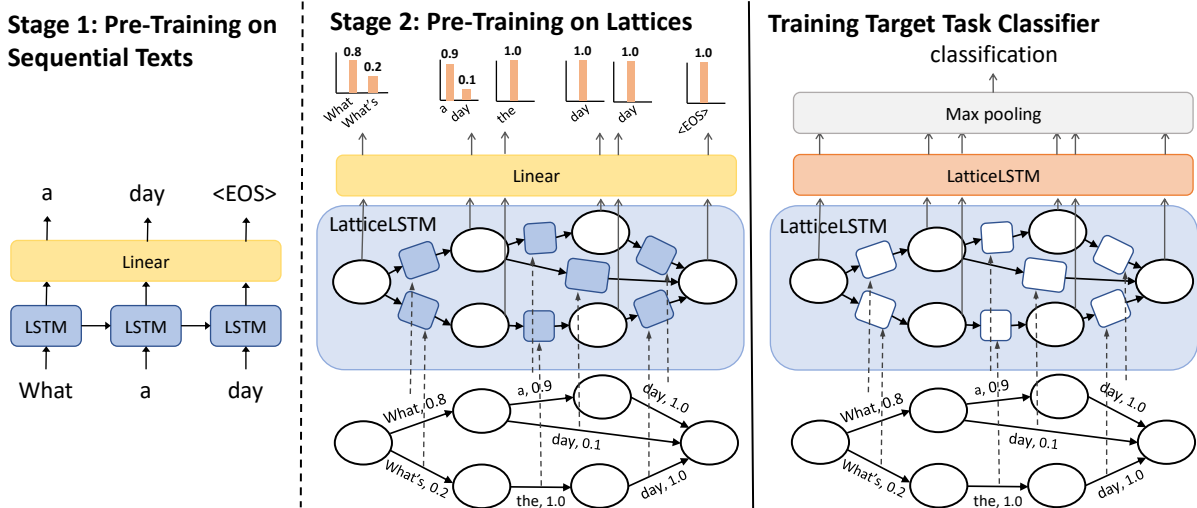


Figure 2: Illustration of the proposed framework. The weights of the pre-trained LatticeLSTM LM are fixed when training the target task classifier (shown in white blocks), while the weights of the newly added LatticeLSTM classifier are trained from scratch (shown in colored block).

where θ denotes the parameters of the LatticeRNN and W denotes the trainable parameters of the decoder. We train our lattice language model by minimizing the KL divergence between the ground truth distribution $p(w | L_{<n})$ and the predicted distribution $p_{\theta}(w | h[n])$.

Note that the objective for training sequential LM is a special case of the lattice language modeling objective defined above, where the inputs are linear-chain lattices. Hence, a sequential LM can be viewed as a lattice LM trained on linear-chain lattices only. This property inspires us to pre-train our lattice LM in a 2-stage fashion described below.

2.4 Two-Stage Pre-Training

Inspired by ULMFiT (Howard and Ruder, 2018), we propose a two-stage pre-training method to train our lattice language model. The proposed method is illustrated in Figure 2.

- Stage 1: Pre-train on sequential texts
In the first stage, we follow the recent trend of pre-trained LMs by pre-training a bidirectional LSTM (Hochreiter and Schmidhuber, 1997) LM on general domain text corpus. Here the cell architecture is the same as ELMo (Peters et al., 2018).
- Stage 2: Pre-train on lattices
In this stage, we use a bidirectional LatticeLSTM with the same cell architecture as the LSTM pre-trained in the previous stage. Note that in the backward direction we use reversed

lattices as input. We initialize the weights of the LatticeLSTM with the weights of the pre-trained LSTM. The LatticeLSTM is further pre-trained on lattices from the training set of the target task with the lattice language modeling objective described above.

We consider this two-stage method more *approachable* and *efficient* than directly pre-training a lattice LM on large amount of lattices because 1) general domain written data is much easier to collect than lattices which require spoken data, and 2) LatticeRNNs are considered less efficient than RNNs due to the difficulty of parallelization in computing.

2.5 Target Task Classifier Training

After pre-training, our model is capable of providing representations for lattices. Following (Peters et al., 2018), the pre-trained lattice LM is used to produce contextualized node embeddings for downstream classification tasks, as illustrated in the right part of Figure 2. We use the same strategy as Peters et al. (2018) to linearly combine the hidden states from different layers into a representation for each node. The classifier is a newly added 2-layer LatticeLSTM, which takes the node representations as input, followed by max-pooling over nodes, a linear layer and finally a softmax layer. We use the cross entropy loss to train the classifier on each target classification tasks. Note that the parameters of the pre-trained lattice LM are fixed during this stage.

			ATIS	SNIPS	SWDA	MRDA
Manual	(a)	biLSTM	-	97.00	71.19	79.99
	(b)	(a) + ELMo	-	96.80	72.18	81.48
Lattice oracle	(c)	biLSTM	92.97	94.02	63.92	70.49
	(d)	(c) + ELMo	96.21	95.14	65.14	73.34
ASR 1-Best	(e)	biLSTM	91.60	91.89	60.54	67.35
	(f)	(e) + ELMo	94.99	91.98	61.65	68.52
	(g)	BERT-base	95.97	93.29	61.23	67.90
Lattices	(h)	biLatticeLSTM	91.69	93.43	61.29	69.95
	(i)	Proposed	95.84	95.37	62.88	72.04
	(j)	(i) w/o Stage 1	94.65	95.19	61.81	71.71
	(k)	(i) w/o Stage 2	95.35	94.58	62.41	71.66
	(l)	(i) evaluated on 1-best	95.05	92.40	61.12	68.04

Table 2: Results of our experiments in terms of accuracy (%). Some audio files in ATIS are missing, so the testing sets of manual transcripts and ASR transcripts are different. Hence, we do not report the results for ATIS using manual transcripts. The best results obtained by using ASR output for each dataset are marked in bold.

	ATIS	SNIPS	SWDA	MRDA
Train	4,478	13,084	103,326	73,588
Valid	500	700	8,989	15,037
Test	869	700	15,927	14,800
#Classes	22	7	43	5
WER(%)	15.55	45.61	28.41	32.04
Oracle WER	9.19	18.79	17.15	21.53

Table 1: Data statistics.

3 Experiments

In order to evaluate the quality of the pre-trained lattice LM, we conduct the experiments for two common tasks in spoken language understanding.

3.1 Tasks and Datasets

Intent detection and dialogue act recognition are two common tasks about spoken language understanding. The benchmark datasets used for intent detection are ATIS (Airline Travel Information Systems) (Hemphill et al., 1990; Dahl et al., 1994; Tur et al., 2010) and SNIPS (Coucke et al., 2018). We use the NXT-format of the Switchboard (Stolcke et al., 2000) Dialogue Act Corpus (SWDA) (Calhoun et al., 2010) and the ICSI Meeting Recorder Dialogue Act Corpus (MRDA) (Shriberg et al., 2004) for benchmarking dialogue act recognition. The SNIPS corpus only contains written text, so we synthesize a spoken version of the dataset using a commercial text-to-speech service. We use an ASR system trained on WSJ (Paul and Baker, 1992) with Kaldi (Povey et al., 2011) to transcribe ATIS, and an ASR system released by Kaldi to transcribe other datasets. The statistics of datasets are summarized in Table 1. All tasks are evaluated

with overall classification accuracy.

3.2 Model and Training Details

In order to conduct fair comparison with ELMo (Peters et al., 2018), we directly adopt their pre-trained model as our pre-trained sequential LM. The hidden size of the LatticeLSTM classifier is set to 300. We use adam as the optimizer with learning rate 0.0001 for LM pre-training and 0.001 for training the classifier. The checkpoint with the best validation accuracy is used for evaluation.

3.3 Results

The results in terms of the classification accuracy are shown in Table 2. All reported numbers are averaged over at least three training runs. Rows (a) and (b) can be considered as the performance upperbound, where we use manual transcripts to train and evaluate the models. We also use BERT-base (Devlin et al., 2019) as a strong baseline, which takes ASR 1-best as the input (row (g)). Compare with the results on manual transcripts, using ASR results largely degrades the performance due to recognition errors, as shown in rows (e)-(g). In addition, adding pre-trained ELMo embeddings brings consistent improvement over the biLSTM baseline, except for SNIPS when using manual transcripts (row (b)). The baseline models trained on ASR 1-best are also evaluated on lattice oracle paths. We report the results as the performance upperbound for the baseline models (rows (c)-(d)).

In the lattice setting, the baseline bidirectional LatticeLSTM (Ladhak et al., 2016) (row (h)) con-

sistently outperforms the biLSTM with 1-best input (row (e)), demonstrating the importance of taking lattices into account. Our proposed method achieves the best results on all datasets except for ATIS (row(i)), with relative error reduction ranging from 3.2% to 42% compare to biLSTM+ELMo (row(f)). The proposed method also achieves performance comparable to BERT-base on ATIS. We perform ablation study for the proposed two-stage pre-training method and report the results in rows (j) and (k). It is clear that skipping either stage degrades the performance on all datasets, demonstrating that both stages are crucial in the proposed framework. We also evaluate the proposed model on 1-best results (row (l)). The results show that it is still beneficial to use lattice as input after fine-tuning.

4 Conclusion

In this paper, we propose a spoken language representation learning framework that learns contextualized representation of lattices. We introduce the lattice language modeling objective and a two-stage pre-training method that efficiently trains a neural lattice language model to provide the downstream tasks with contextualized lattice representations. The experiments show that our proposed framework is capable of providing high-quality representations of lattices, yielding consistent improvement on SLU tasks.

Acknowledgement

We thank reviewers for their insightful comments. This work was financially supported from the Young Scholar Fellowship Program by Ministry of Science and Technology (MOST) in Taiwan, under Grant 109-2636-E-002-026.

References

Jacob Buckman and Graham Neubig. 2018. Neural lattice language models. *Transactions of the Association for Computational Linguistics*, 6:529–541.

Sasha Calhoun, Jean Carletta, Jason M Brenier, Neil Mayo, Dan Jurafsky, Mark Steedman, and David Beaver. 2010. The next-format switchboard corpus: a rich resource for investigating the syntax, semantics, pragmatics and prosody of dialogue. *Language resources and evaluation*, 44(4):387–419.

Alexandra Chronopoulou, Christos Baziotis, and Alexandros Potamianos. 2019. An embarrassingly

simple approach for transfer learning from pre-trained language models. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2089–2095. ACL.

Alice Coucke, Alaa Saade, Adrien Ball, Théodore Bluche, Alexandre Caulier, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Caltagirone, Thibaut Lavril, et al. 2018. Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces. *arXiv preprint arXiv:1805.10190*.

Deborah A Dahl, Madeleine Bates, Michael Brown, William Fisher, Kate Hunicke-Smith, David Pallett, Christine Pao, Alexander Rudnicky, and Elizabeth Shriberg. 1994. Expanding the scope of the atis task: The atis-3 corpus. In *Proceedings of the workshop on Human Language Technology*, pages 43–48.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186. ACL.

Chih-Wen Goo, Guang Gao, Yun-Kai Hsu, Chih-Li Huo, Tsung-Chieh Chen, Keng-Wei Hsu, and Yun-Nung Chen. 2018. Slot-gated modeling for joint slot filling and intent prediction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 753–757. ACL.

Daniel Guo, Gokhan Tur, Wen-tau Yih, and Geoffrey Zweig. 2014. Joint semantic utterance classification and slot filling with recursive neural networks. In *2014 IEEE Spoken Language Technology Workshop*, pages 554–559.

Dilek Hakkani-Tür, Frédéric Béchet, Giuseppe Riccardi, and Gokhan Tur. 2006. Beyond ASR 1-best: Using word confusion networks in spoken language understanding. *Computer Speech & Language*, 20(4):495–514.

Charles T. Hemphill, John J. Godfrey, and George R. Doddington. 1990. The ATIS spoken language systems pilot corpus. In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*.

Matthew Henderson, Milica Gašić, Blaise Thomson, Pirros Tsiakoulis, Kai Yu, and Steve Young. 2012. Discriminative spoken language understanding using word confusion networks. In *2012 IEEE Spoken Language Technology Workshop*, pages 176–181.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*.

- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339. ACL.
- Chao-Wei Huang and Yun-Nung Chen. 2019. Adapting pretrained transformer to lattices for spoken language understanding. In *Proceedings of 2019 IEEE Automatic Speech Recognition and Understanding Workshop*, pages 845–852.
- Faisal Ladhak, Ankur Gandhe, Markus Dreyer, Lambert Mathias, Ariya Rastrow, and Björn Hoffmeister. 2016. LatticeRNN: Recurrent neural networks over lattices. In *Proceedings of INTERSPEECH*, pages 695–699.
- Ryo Masumura, Yusuke Ijima, Taichi Asami, Hirokazu Masataki, and Ryuichiro Higashinaka. 2018. Neural confnet classification: Fully neural network based spoken utterance classification using word confusion networks. In *Proceedings of 2018 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 6039–6043.
- Grégoire Mesnil, Yann Dauphin, Kaisheng Yao, Yoshua Bengio, Li Deng, Dilek Hakkani-Tur, Xiaodong He, Larry Heck, Gokhan Tur, Dong Yu, and Geoffrey Zweig. 2014. Using recurrent neural networks for slot filling in spoken language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(3):530–539.
- Douglas B. Paul and Janet M. Baker. 1992. The design for the wall street journal-based csr corpus. In *Proceedings of the Workshop on Speech and Natural Language*, HLT '91.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237. ACL.
- Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, et al. 2011. The Kaldi speech recognition toolkit. Technical report.
- Alec Radford. 2018. Improving language understanding by generative pre-training.
- Elizabeth Shriberg, Raj Dhillon, Sonali Bhagat, Jeremy Ang, and Hannah Carvey. 2004. The ICSI meeting recorder dialog act (MRDA) corpus. In *Proceedings of the 5th SIGdial Workshop on Discourse and Dialogue at HLT-NAACL 2004*, pages 97–100, Cambridge, Massachusetts, USA. ACL.
- Aditya Siddhant, Anuj Goyal, and Angeliki Metallinou. 2018. Unsupervised transfer learning for spoken language understanding in intelligent agents. *arXiv preprint arXiv:1811.05370*.
- Matthias Sperber, Graham Neubig, Ngoc-Quan Pham, and Alex Waibel. 2019. Self-attentional models for lattice inputs. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1185–1197. ACL.
- Andreas Stolcke, Klaus Ries, Noah Coccaro, Elizabeth Shriberg, Rebecca Bates, Daniel Jurafsky, Paul Taylor, Rachel Martin, Carol Van Ess-Dykema, and Marie Meteer. 2000. Dialogue act modeling for automatic tagging and recognition of conversational speech. *Computational Linguistics*, 26(3):339–374.
- Gökhan Tür, Anoop Deoras, and Dilek Z. Hakkani-Tür. 2013. Semantic parsing using word confusion networks with conditional random fields. In *Proceedings of INTERSPEECH*.
- Gokhan Tur, Dilek Hakkani-Tür, and Larry Heck. 2010. What is left to be understood in ATIS? In *Proceedings of 2010 IEEE Spoken Language Technology Workshop (SLT)*, pages 19–24.
- Gokhan Tur, Jerry Wright, Allen Gorin, Giuseppe Riccardi, and Dilek Hakkani-Tür. 2002. Improving spoken language understanding using word confusion networks. In *Seventh International Conference on Spoken Language Processing*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6000–6010. Curran Associates Inc.
- Fengshun Xiao, Jiangtong Li, Hai Zhao, Rui Wang, and Kehai Chen. 2019. Lattice-based transformer encoder for neural machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3090–3097. ACL.
- Kaisheng Yao, Baolin Peng, Yu Zhang, Dong Yu, Geoffrey Zweig, and Yangyang Shi. 2014. Spoken language understanding using long short-term memory neural networks. In *2014 IEEE Spoken Language Technology Workshop*, pages 189–194.
- Pei Zhang, Niyu Ge, Boxing Chen, and Kai Fan. 2019. Lattice transformer for speech translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6475–6484. ACL.
- Yue Zhang and Jie Yang. 2018. Chinese NER using lattice LSTM. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1554–1564. ACL.