

Improving Slot Filling by Utilizing Contextual Information

Amir Pouran Ben Veyseh^{*1}, Franck Dernoncourt²,
and Thien Huu Nguyen¹

¹Department of Computer and Information Science,
University of Oregon, Eugene, Oregon, USA

²Adobe Research, San Jose, CA, USA

{apouranb, thien}@cs.uoregon.edu
franck.dernoncourt@adobe.com

Abstract

Slot Filling (SF) is one of the sub-tasks of Spoken Language Understanding (SLU) which aims to extract semantic constituents from a given natural language utterance. It is formulated as a sequence labeling task. Recently, it has been shown that contextual information is vital for this task. However, existing models employ contextual information in a restricted manner, e.g., using self-attention. Such methods fail to distinguish the effects of the context on the word representation and the word label. To address this issue, in this paper, we propose a novel method to incorporate the contextual information in two different levels, i.e., representation level and task-specific (i.e., label) level. Our extensive experiments on three benchmark datasets on SF show the effectiveness of our model leading to new state-of-the-art results on all three benchmark datasets for the task of SF.

1 Introduction

Slot Filling (SF) is the task of identifying the semantic constituents expressed in a natural language utterance. It is one of the sub-tasks of spoken language understanding (SLU) and plays a vital role in personal assistant tools such as Siri, Alexa, and Google Assistant. This task is formulated as a sequence labeling problem. For instance, in the given sentence “*Play Signe Anderson chant music that is newest.*”, the goal is to identify “*Signe Anderson*” as “*artist*”, “*chant music*” as “*music-item*” and “*newest*” as “*sort*”.

Early work on SF has employed feature engineering methods to train statistical models, e.g., Conditional Random Field (Raymond and Riccardi, 2007). Later, deep learning emerged as a promising approach for SF (Yao et al., 2014; Peng et al., 2015;

Liu and Lane, 2016). The success of deep models could be attributed to pre-trained word embeddings to generalize words and deep learning architectures to compose the word embeddings to induce effective representations. In addition to improving word representation using deep models, Liu and Lane (2016) showed that incorporating the context of each word into its representation could improve the results. Concretely, the effect of using context in word representation is two-fold: (1) **Representation Level:** As the meaning of the word is dependent on its context, incorporating the contextual information is vital to represent the true meaning of the word in the sentence (2) **Task Level:** For SF, the label of the word is related to the other words in the sentence and providing information about the other words, in prediction layer, could improve the performance. Unfortunately, the existing work employs the context in a restricted manner, e.g., via attention mechanism, which obfuscates the model about the two aforementioned effects of the contextual information.

In order to address the limitations of the prior work to exploit the context for SF, in this paper, we propose a multi-task setting to train the model. More specifically, our model is encouraged to explicitly ensure the two aforementioned effects of the contextual information for the task of SF. In particular, in addition to the main sequence labeling task, we introduce new sub-tasks to ensure each effect. Firstly, in the representation level, we enforce the consistency between the word representations and their context. This enforcement is achieved via increasing the Mutual Information (MI) between these two representations. Secondly, in the task level, we propose two new sub-tasks: (1) To predict the label of the word solely from its context and (2) To predict which labels exist in the given sentence in a multi-label classification setting. By doing so, we encourage the model to encode task-

^{*} This work was done when the first author was an intern at Adobe Research.

specific features in the context of each word. Our extensive experiments on three benchmark datasets, empirically prove the effectiveness of the proposed model leading to new the state-of-the-art results on all three datasets.

2 Related Work

In the literature, Slot Filling (SF), is categorized as one of the sub-tasks of spoken language understanding (SLU). Early work employed feature engineering for statistical models, e.g., Conditional Random Field (Raymond and Riccardi, 2007). Due to the lack of generalisation ability of feature based models, deep learning based models superseded them (Yao et al., 2014; Peng et al., 2015; Kurata et al., 2016; Hakkani-Tür et al., 2016). Also, joint models to simultaneously predict the intent of the utterance and to extract the semantic slots has also gained a lot of attention (Guo et al., 2014; Liu and Lane, 2016; Zhang and Wang, 2016; Wang et al., 2018; Goo et al., 2018; Qin et al., 2019; E et al., 2019). In addition to the supervised settings, recently other setting such as progressive learning (Shen et al., 2019) or zero-shot learning has also been studied (Shah et al., 2019). To the best of our knowledge, none of the existing work introduces a multi-task learning solely for the SF to incorporate the contextual information in both representation and task levels.

3 Model

Our model is trained in a multi-task setting in which the main task is slot filling to identify the best possible sequence of labels for the given sentence. In the first auxiliary task we aim to increase consistency between the word representation and its context. The second auxiliary task is to enhance task specific information in contextual information. In this section, we explain each of these tasks in more details.

3.1 Slot Filling

Formally, the input to a SF model is a sequence of words $X = [x_1, x_2, \dots, x_n]$ and our goal is to predict the sequence of labels $Y = [y_1, y_2, \dots, y_n]$. In our model, the word x_i is represented by vector e_i which is the concatenation of the pre-trained word embedding and POS tag embedding of the word x_i . In order to obtain a more abstract representation of the words, we employ a Bi-directional Long Short-Term Memory (BiLSTM) over the word rep-

resentations $E = [e_1, e_2, \dots, e_n]$ to generate the abstract vectors $H = [h_1, h_2, \dots, h_n]$. The vector h_i is the final representation of the word x_i and is fed into a two-layer feed forward neural net to compute the label scores s_i for the given word, $s_i = FF(h_i)$. As the task of SF is formulated as a sequence labeling task, we exploit a conditional random field (CRF) layer as the final layer of SF prediction. More specifically, the predicted label scores $S = [s_1, s_2, \dots, s_n]$ are provided as emission score to the CRF layer to predict the label sequence $\hat{Y} = [\hat{y}_1, \hat{y}_2, \dots, \hat{y}_n]$. To train the model, the negative log-likelihood is used as the loss function for SF prediction, i.e., \mathcal{L}_{pred} .

3.2 Consistency between Word and Context

In this sub-task we aim to increase the consistency between the word representation and its context. To obtain the context of each word, we use max pooling over the outputs of the BiLSTM for all words of the sentence excluding the word itself, $h_i^c = MaxPooling(h_1, h_2, \dots, h_n/h_i)$. We aim to increase the consistency between vectors h_i and h_i^c . To this end, we propose to maximize the Mutual Information (MI) between the word representation and its context. In information theory, MI evaluates how much information we know about one random variable if the value of another variable is revealed. Formally, the mutual information between two random variable X_1 and X_2 is obtained by:

$$MI(X_1, X_2) = \int_{X_1} \int_{X_2} P(X_1, X_2) \cdot \log \frac{P(X_1, X_2)}{P(X_1)P(X_2)} dX_1 dX_2 \quad (1)$$

Using this definition of MI, we can reformulate the MI equation as KL-Divergence between the joint distribution $P_{X_1 X_2} = P(X_1, X_2)$ and the product of marginal distributions $P_{X_1} \otimes_{X_2} = P(X_1)P(X_2)$:

$$MI(X_1, X_2) = D_{KL}(P_{X_1 X_2} || P_{X_1} \otimes_{X_2}) \quad (2)$$

Based on this understanding of MI, if the two random variables are dependent then the mutual information between them (i.e. the KL-Divergence in Equation 2) would be the highest. Consequently, if the representations h_i and h_i^c are encouraged to

have large mutual information, we expect them to share more information.

Computing the KL-Divergence in equation 2 could be prohibitively expensive (Belghazi et al., 2018), so we need to estimate it. To this end, we exploit the adversarial method introduced in (Hjelm et al., 2019). In this method, a discriminator is employed to distinguish between samples from the joint distribution and the product of the marginal distributions to estimate the KL-Divergence in Equation 2. In our case, the sample from joint distribution is the concatenation $[h_i : h_i^c]$ and the sample from the product of the marginal distribution is the concatenation $[h_i : h_j^c]$ where h_j^c is a context vector randomly chosen from the words in the mini-batch. Formally:

$$\mathcal{L}_{disc} = \frac{1}{n} \sum_{i=1}^n - (\log(D[h_i, h_i^c]) + \log(1 - D([h_i, h_j^c]))) \quad (3)$$

Where D is the discriminator. This loss is added to the final loss function of the model.

3.3 Prediction by Contextual Information

In addition to increasing consistency between the word representation and its context representation, we aim to increase the task-specific information in contextual representations. To this end, we train the model on two auxiliary tasks. The first one aims to use the context of each word to predict the label of that word. The goal of the second auxiliary task is to use the global context information to predict sentence level labels. We describe each of these tasks in more details in the following subsections.

Predicting Word Label

In this sub-task, we use the context representations of each word to predict its label. It will increase the information encoded in the context of the word about the label of the word. We use the same context vector h_i^c for the i -th word as described in the previous section. This vector is fed into a two-layer feed forward neural network with a softmax layer at the end to output the probabilities for each class, $P_i(\cdot | \{x_1, x_2, \dots, x_n\} / x_i) = \text{softmax}(FF(h_i^c))$. Finally, we use the following negative log-likelihood as the loss function to be optimized during training:

$$\mathcal{L}_{wp} = \frac{1}{n} \sum_{i=1}^n - \log(P_i(y_i | \{x_1, x_2, \dots, x_n\} / x_i)) \quad (4)$$

Predicting Sentence Labels

The word label prediction enforces the context of each word to contain information about its label but it lacks a global view about the entire sentence. In order to increase the global information about the sentence in the representation of the words, we aim to predict the labels existing in a sentence from the representations of its words. More specifically, we introduce a new sub-task to predict which labels exists in the given sentence. We formulate this task as a multi-label classification problem. Formally, for each sentence, we predict the binary vector $Y^s = [y_1^s, y_2^s, \dots, y_{|L|}^s]$ where L is the set of all possible word labels. In the vector Y^s , y_i^s is 1 if the sentence X contains i -th label from the label set L otherwise it is 0.

To predict vector Y^s , we first compute the representation of the sentence. This representation is obtained by max pooling over the outputs of the BiLSTM, $H = \text{MaxPooling}(h_1, h_2, \dots, h_n)$. Afterwards, the vector H is fed into a two-layer feed forward neural net with a sigmoid activation function at the end to compute the probability distribution of Y^s (i.e., $P_k(\cdot | x_1, x_2, \dots, x_n) = \sigma_k(FF(H))$ for k -th label in L). Note that since this task is a multi-label classification, the number of neurons at the final layer is equal to $|L|$. We optimize the following binary cross-entropy loss:

$$\mathcal{L}_{sp} = \frac{1}{|L|} \sum_{k=1}^{|L|} - (y_k^s \cdot \log(P_k(y_k^s | x_1, x_2, \dots, x_n)) + (1 - y_k^s) \cdot \log(1 - P_k(y_k^s | x_1, x_2, \dots, x_n))) \quad (5)$$

Finally, to train the entire model we optimize the following combined loss:

$$\mathcal{L} = \mathcal{L}_{pred} + \alpha \mathcal{L}_{disc} + \beta \mathcal{L}_{wp} + \gamma \mathcal{L}_{sp} \quad (6)$$

where α , β and γ are the trade-off parameters to be tuned based on the development set performance.

4 Experiments

4.1 Dataset and Parameters

We evaluate our model on three SF datasets. Namely, we employ ATIS (Hemphill et al., 1990), SNIPS (Coucke et al., 2018) and EditMe (Manuvinaurike et al., 2018). ATIS and SNIPS are two widely adopted SF dataset and EditMe is a SF dataset for editing images with four slot labels (i.e., *Action*, *Object*, *Attribute* and *Value*). The

statistics of the datasets are presented in the Appendix A. Based on the experiments on EditMe development set, the following parameters are selected: GloVe embedding with 300 dimensions to initialize word embedding ; 200 dimensions for the all hidden layers in LSTM and feed forward neural net; 0.1 for trade-off parameters α , β and γ ; and Adam optimizer with learning rate 0.001. Following previous work, we use F1-score to evaluate the model.

4.2 Baselines

We compare our model with other deep learning based models for SF. Namely, we compare the proposed model with Joint Seq (Hakkani-Tür et al., 2016), Attention-Based (Liu and Lane, 2016), Sloted-Gated (Goo et al., 2018), SF-ID (E et al., 2019), CAPSULE-NLU (Zhang et al., 2019), and SPTID (Qin et al., 2019). Note that we compare our model with the single-task version of these baselines. We also compare our model with other sequence labeling models which are not specifically proposed for SF. Namely, we compare the model with CVT (Clark et al., 2018) and GCDT (Liu et al., 2019). CVT aims to improve input representation using improving partial views and GCDT exploits contextual information to enhance word representations via concatenation of context and word representation.

4.3 Results

Table 1 reports the performance of the model and baselines. The proposed model outperforms all baselines in all datasets. Among all baselines, GCDT achieves best results on two out of three datasets. This superiority shows the importance of explicitly incorporating the contextual information into word representation for SF. However, the proposed model improves the performance substantially on all datasets by explicitly encouraging the consistency between a word and its context in representation level and task-specific (i.e., label) level. Also, Table 1 shows that EditMe dataset is more challenging than the other datasets, despite having fewer slot types. This difficulty could be explained by the limited number of training examples and more diversity in sentence structures in this dataset.

4.4 Ablation Study

Our model consists of three major components: (1) **MI**: Increasing mutual information between word

Model	SNIPS	ATIS	EditMe
Joint Seq(2016)	87.3	94.3	-
Attention-Based(2016)	87.8	94.2	-
Sloted-Gated(2018)	89.2	95.4	84.9
SF-ID(2019)	90.9	95.5	85.2
CAPSULE-NLU(2019)	91.8	95.2	84.6
SPTID(2019)	90.8	95.1	85.3
CVT(2018)	91.4	94.8	85.4
GCDT(2019)	92.0	95.1	85.6
Ours	93.6	95.8	87.2

Table 1: Performance of the model and baselines on the Test sets.

Model	SNIPS	ATIS	EditMe
Full	93.6	95.8	87.2
Full - MI	92.9	95.3	84.2
Full - WP	91.7	94.9	83.2
Full - SP	92.5	95.2	84.1

Table 2: Test F1-score for the ablated models

and its context representations (2) **WP**: Predicting the label of the word using its context to increase word level task-specific information in the word context (3) **SP**: Predicting which labels exist in the given sentence in a multi-label classification to increase sentence level task-specific information in word representations. In order to analyze the contribution of each of these components, we also evaluate the model performance when we remove one of the components and retrain the model. The results are reported in Table 2. This Table shows that all components are required for the model to have its best performance. Among all components, the word level prediction using the contextual information has the major contribution to the model performance. This fact shows that contextual information trained to be informative about the final task is necessary to obtain the representations which could boost the performance.

5 Conclusion

In this work, we introduced a new deep model for the task of Slot Filling (SF). In a multi-task setting, our model increases the mutual information between the word representation and its context, improves label information in the context and predicts which concepts are expressed in the given sentence. Our experiments on three benchmark datasets show the effectiveness of our model by achieving the state-of-the-art results on all datasets for the SF task.

References

- Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, R. Devon Hjelm, and Aaron C. Courville. 2018. Mutual information neural estimation. In *ICML*.
- Kevin Clark, Minh-Thang Luong, Christopher D Manning, and Quoc V Le. 2018. Semi-supervised sequence modeling with cross-view training. In *EMNLP*.
- Alice Coucke, Alaa Saade, Adrien Ball, Théodore Bluche, Alexandre Caulier, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Caltagirone, Thibaut Lavril, and et al. 2018. Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces. In *arXiv*.
- Haihong E, Peiqing Niu, Zhongfu Chen, and Meina Song. 2019. A novel bi-directional interrelated model for joint intent detection and slot filling. In *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 *NAACL-HLT*.
- Daniel Guo, Gokhan Tur, Wen-tau Yih, and Geoffrey Zweig. 2014. Joint semantic utterance classification and slot filling with recursive neural networks. In *SLT*.
- Dilek Hakkani-Tür, Gökhan Tür, Asli Celikyilmaz, Yun-Nung Chen, Jianfeng Gao, Li Deng, and Ye-Yi Wang. 2016. Multi-domain joint semantic frame parsing using bi-directional rnn-lstm. In *Inter-speech*.
- 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*.
- R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. 2019. Learning deep representations by mutual information estimation and maximization. In *ICLR*.
- Gakuto Kurata, Bing Xiang, Bowen Zhou, and Mo Yu. 2016. Leveraging sentence-level information with encoder LSTM for semantic slot filling. In *EMNLP*.
- Bing Liu and Ian Lane. 2016. Attention-based recurrent neural network models for joint intent detection and slot filling. In *arXiv*.
- Yijin Liu, Fandong Meng, Jinchao Zhang, Jinan Xu, Yufeng Chen, and Jie Zhou. 2019. GCDT: A global context enhanced deep transition architecture for sequence labeling. In *ACL*.
- Ramesh Manuvinakurike, Jacqueline Brixey, Trung Bui, Walter Chang, Doo Soon Kim, Ron Artstein, and Kallirroi Georgila. 2018. *Edit me: A corpus and a framework for understanding natural language image editing*. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Baolin Peng, Kaisheng Yao, Li Jing, and Kam-Fai Wong. 2015. Recurrent neural networks with external memory for spoken language understanding. In *NLPCC*.
- Libo Qin, Wanxiang Che, Yangming Li, Haoyang Wen, and Ting Liu. 2019. A stack-propagation framework with token-level intent detection for spoken language understanding. In *EMNLP*.
- Christian Raymond and Giuseppe Riccardi. 2007. Generative and discriminative algorithms for spoken language understanding. In *ISCA*.
- Darsh J Shah, Raghav Gupta, Amir A Fayazi, and Dilek Hakkani-Tur. 2019. Robust zero-shot cross-domain slot filling with example values. *arXiv*.
- Yilin Shen, Xiangyu Zeng, and Hongxia Jin. 2019. A progressive model to enable continual learning for semantic slot filling. In *EMNLP*.
- Yu Wang, Yilin Shen, and Hongxia Jin. 2018. A bi-model based RNN semantic frame parsing model for intent detection and slot filling. In *NAACL-HLT*.
- 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 *SLT*.
- Chenwei Zhang, Yaliang Li, Nan Du, Wei Fan, and Philip Yu. 2019. Joint slot filling and intent detection via capsule neural networks. In *ACL*.
- Xiaodong Zhang and Houfeng Wang. 2016. A joint model of intent determination and slot filling for spoken language understanding. In *IJCAI*.

A Dataset Statistics

In our experiments, we employ three benchmark datasets, ATIS, SNIPS and EditMe. Table 3 presents the statistics of these three datasets. Moreover, in order to provide more insight into the EditMe dataset, we report the labels statistics of this dataset in Table 4.

Dataset	Train	Dev	Test
SNIPS	13,084	700	700
ATIS	4,478	500	893
EditMe	1,737	497	559

Table 3: Total number of examples in test/dev/train splits of the datasets

Label	Train	Dev	Test
Action	1,562	448	479
Object	4,676	1,447	1,501
Attribute	1,437	403	462
Value	507	207	155

Table 4: Label Statistics of EditMe dataset