Article

Elastic CRFs for Open-Ontology Slot Filling

Yinpei Dai [1,2,†] , **Yichi Zhang** [1,2,†], **Hong Liu** [1,2], **Zhijian Ou** [1,2,3,*] , **Yi Huang** [2,4] and **Junlan Feng** [2,4]

1   Speech Processing and Machine Intelligence (SPMI) Lab, Tsinghua University, Beijing 100084, China; daiyinpei123@163.com (Y.D.); zhangyic17@tsinghua.org.cn (Y.Z.); hliu17@mails.tsinghua.edu.cn (H.L.)
2   Tsinghua University-China Mobile Communications Group Co., Ltd. Joint Institute, Beijing 100084, China; huangyi@chinamobile.com (Y.H.); fengjunlan@chinamobile.com (J.F.)
3   Beijing National Research Center for Information Science and Technology, Beijing 100084, China
4   China Mobile Research Institute, Beijing 100053, China
*   Correspondence: ozj@tsinghua.edu.cn
†   These authors contributed equally to this work.

Abstract: Slot filling is a crucial component in task-oriented dialog systems that is used to parse (user) utterances into semantic concepts called slots. An ontology is defined by the collection of slots and the values that each slot can take. The most widely used practice of treating slot filling as a sequence labeling task suffers from two main drawbacks. First, the ontology is usually pre-defined and fixed and therefore is not able to detect new labels for unseen slots. Second, the one-hot encoding of slot labels ignores the correlations between slots with similar semantics, which makes it difficult to share knowledge learned across different domains. To address these problems, we propose a new model called elastic conditional random field (eCRF), where each slot is represented by the embedding of its natural language description and modeled by a CRF layer. New slot values can be detected by eCRF whenever a language description is available for the slot. In our experiment, we show that eCRFs outperform existing models in both in-domain and cross-domain tasks, especially in predicting unseen slots and values.

Keywords: open ontology; slot filling; conditional random fields; dialog systems

1. Introduction

Slot filling [1,2] is a crucial component in task-oriented dialog systems and parses (user) utterances into semantic concepts in terms of a set of named entities called slots. The example in Figure 1 contains the slots `time` and `movie`. In parsing, some span in the utterance is identified as the slot value for some slot; e.g., here, *"6 pm"* is marked as the slot `time`. An ontology, which describes the scope of semantics that the dialog system can process, is defined by the collection of slots and the values that each slot can take. A widely used practice for slot filling is to introduce IOB tags [3] and assign a label to each token in the utterance. A label, e.g., `B-time`, is a combination of the slot name and one of the IOB tags. These labels are then used to identify the values for different slots from the utterance. In this manner, slot filling is treated as a sequence labeling task, as illustrated in Figure 1, for which the two dominant classes of methods are based on recurrent neural networks (RNNs) [1] and conditional random fields (CRFs) [4], respectively. This practice has been widely employed for slot filling [2,5] and many other similar sequence labeling problems [6]. However, this practice suffers from two drawbacks.

First, currently, most slot-filling methods are unable to predict new labels for unseen slots. The ontology is usually pre-defined and fixed. It is difficult to accommodate new semantic concepts (slots) in slot filling. However, users may often add new semantic concepts in a domain and dialog systems are expected to work across an increasingly wide range of domains. Thus, it is highly desirable for slot-filling models to be able to handle new slots, whether in-domain or cross-domain, with the least expense being incurred after

training on a certain domain. In this paper, we are interested in developing such open-ontology slot filling, which means that the collection of slots and values is open-ended for slot filling. Second, in current slot-filling models [5,7], slot labels are generally encoded as one-hot vectors. However, slot labels are not merely discrete classes. There are natural language descriptions for each slot, e.g., the description "number of people" for the slot `#people`. This one-hot encoding ignores the semantic meanings and relations for slots, which are implicit in their natural language descriptions and useful for slot filling.

Utterance	6	pm	for	the	movie	called	avatar
Slot labels	B-time	I-time	O	O	O	O	B-movie

Figure 1. An example of slot filling in the movie domain.

There are prior efforts to address the above two drawbacks. The difficulty of transferring between domains could be partly alleviated with multi-task learning [8–10], by performing joint learning on multiple domains. Practically, varying only the last output layer for different domains and sharing the parameters of the rest layers has shown to be a successful approach [11]. In this approach, the slot-filling model can leverage all available multi-domain data and transfer them to handle those slots with sparse training data. However, basically, this multi-task learning approach is unable to predict labels for zero-shot slots (namely those slots that are unseen in training data and whose values are unknown). It can be seen that this difficulty is also related to the drawback of one-hot encoding slot labels, which hinders the exploitation of semantic relations and shared statistical properties between different slots. A recent work [12] proposes utilizing slot label descriptions towards zero-shot slot filling by introducing slot encodings from natural language descriptions. Basically, they use RNN-based sequence labeling, taking the slot encoding vector as an additional conditional input and outputting the IOB tags in each position. Sequence labeling is carried out independently for all slots. Though yielding promising results, there are two shortcomings. First, independent sequence labeling may make conflicting predictions. Second, interactions between slots are ignored in sequence labeling.

CRFs have been shown to be one of the most successful approaches for sequence labeling, especially for capturing the interactions between labels. A widely used method is to implement a CRF layer on top of features generated by a RNN [1]. These recent neural CRFs are different from conventional CRFs, which mainly use discrete indicator features. However, these recent CRFs still work with a closed set of labels. In this paper, we propose a novel neural CRF model, called elastic CRF (eCRF), for open-set sequence labeling, by leveraging label descriptions inspired from [12]. The key idea of eCRFs is to use slot descriptions to create semantically meaningful IOB tags [3], which are further used for a new calculation of potential functions in the CRF framework. Compared to traditional fixed IOB tags in original CRFs, our eCRFs are able to process new slots unseen during training without retraining the model. Such flexibility is the motivation for calling it an "elastic" CRF model.

The eCRFs are powerful models for open-ontology slot filling. Intuitively, the node potentials of eCRFs combine the neural features of both the utterance and the slot descriptions, and the edge potentials model the interactions between different slots. In the experiments, we make use of the Google simulated dataset [13], and re-split the dataset according to the in-domain task and the cross-domain task, which focus on the challenge of handling unseen values and unseen slots, respectively. The results show that eCRFs significantly outperform not only a BiLSTM baseline but also the concept tagger (CT) in [12] for both tasks, especially in predictions of unseen slots and values.

In Section 2, we discuss related work. The new eCRF model is detailed in Section 3. Section 4 describes the dataset and task formulations. Section 5 presents the experiments, followed by the conclusion in Section 6.

2. Related Work

One line of related work is zero-shot slot-filling learning [14]. The term *open ontology* referred in this paper is a different name for zero-shot slot filling in spoken language understanding (SLU) for dialog systems. Zero-shot learning has been applied in various of SLU tasks. The authors of [15] leverage the intent embeddings to detect new intent labels which are not included in the training data. Additionally, ref. [12] exploits the slot label descriptions to parse the novel semantic frames for domain scaling and [16] extends the natural language generation module to generalize the responses into an unseen domain via latent action matching. The authors of [17] propose utilizing both the slot description and a small number of examples of slot values to enhance model robustness. In [18], the authors focus on multi-turn zero-shot slot filling in conversation. These studies have utilized the natural language descriptions of the labels, and by constructing the semantic encoder to take the label descriptions as inputs, any new labels in the testing phrase can still be predicted by the model. Our eCRFs also use this semantic encoder structure. However, unlike processing each label description separately in [12], eCRFs are trained and tested by jointly exploiting all possible slot descriptions at one time. Thus, they could capture relations between slot labels and relieve the burden of adjusting the oversampling ratio.

Another line of related work is models for slot filling. CRFs have been extensively applied in traditional slot-filling tasks [19,20], but are restricted by a fixed set of labels. With the progress of deep learning, state-of-art slot-filling methods usually utilize BiLSTM networks [9,21,22]. Extended models, such as encoder–decoder [5] and memory network [23] designs, are explored. More recently, ref. [24] proposes a coarse-to-fine approach (Coach) for cross-domain slot filling, which detects the value span boundary first and then predicts the specific fine types for the slot entities. With the advance of pre-trained models [25], there are also many work [26–28] that adapt the well-studied machine reading comprehension (MRC) framework to solve open-ontology slot filling. Motivated by the BiLSTM-CRF architecture [19,29,30], our eCRFs combine the representation power of deep neural networks and dependency modeling ability of CRFs, together with a newly designed potential function.

3. Proposed Model

Our new model presents an extension from existing neural CRFs [29,30]. Existing neural CRFs in many other sequence labeling tasks are restricted by a fixed set of labels, e.g., *PERSON, LOCATION, ORGANIZATION, MISC* in the name entity recognition (NER) task, and thus can not be applied for open-ontology slot filling. To overcome this shortcoming, we propose a novel framework called elastic conditional random field (eCRF), which consists of three parts. (1) A slot description encoder is employed to encode the slot descriptions into semantic embeddings, then (2) a BiLSTM is used to extract contextual neural features, and finally (3) the outputs of both the slot description encoder and the BiLSTM are combined to define a novel potential function in the CRF. The main framework of eCRF is illustrated in Figure 2 and each part is detailed in the following subsections.
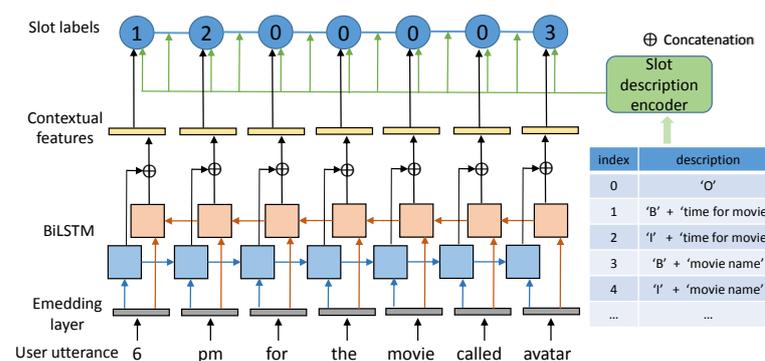

Figure 2. The architecture of the elastic CRF (eCRF) model.

3.1. Slot Description Encoder

Let $X = (x_1, x_2, ..., x_n)$ denote the input user utterance and $D^i = (d_1^i, d_2^i, ...)$ denote the description of slot s^i. In our experiment, slot descriptions are simple complementary phrases, e.g., 'number of people' for the slot #people, 'theatre name' for the slot theatre_name, but other richer expression can be used. The goal of our task is to find all possible text spans in X as values for each s^i. We adapted the IOB tagging scheme as in [3]. Traditionally, the IOB tags are made up three type, 'B', 'I', and 'O', which indicate the beginning position of a value span, the intermediate and ending positions of the value span and the rest position belonging to no values. To be specific, if a word is predicted to have the 'B' tag or multiple words are predicted to have '$B, I, .., I$' tags, the word span is the value of a slot. Instead of using a combination of the slot name and one of the IOB tags as in Figure 1, we used the combination of the slot description and one of the IOB tags in order to leverage the semantic meanings of slots. As shown in Figure 2, the slot description encoder takes all slot descriptions as input, and outputs are distributed representations for all possible combinations of the IOB tags and the slot descriptions, such as 'O', '$B + D^1$', '$I + D^1$', '$B + D^2$', '$I + D^2$', ... The set of these new combined slot labels is denoted as $\mathcal{S}$. We use indexes of these labels to suggest the corresponding positions within the utterance. For example, in Figure 2, '6', 'pm' and 'avatar' are predicted as the positions of '$B+$ time for movie', '$I+$ time for movie' and '$B+$ movie name', which means that '6 pm' is the value of slot movie_time and 'avatar' is the value of slot movie_name. A function $e(\cdot) \in \mathbb{R}^d$ is used to denote the output vector from the slot description encoder as follows:

$$e(B + D^i) = FC\Big(f(D^i) \oplus emb(B)\Big) \tag{1}$$

$$e(I + D^i) = FC\Big(f(D^i) \oplus emb(I)\Big) \tag{2}$$

$$e(O) = FC\Big(\overrightarrow{\mathbf{0}} \oplus emb(O)\Big) \tag{3}$$

where $FC(\cdot)$ denotes a one-hidden-layer fully connected network and $f(\cdot)$ denotes an encoder that maps the descriptions into semantic embeddings. In this paper, we use a simple averaging function of all word embeddings in D^i as in [12]. $emb(\cdot)$ is an embedding lookup function for the IOB tags and $\oplus$ denotes the concatenation operation. Note that for $e(O)$, we use a zero vector $\overrightarrow{\mathbf{0}}$ with the same size as the output vector of $f(\cdot)$ since the 'O' tag should be independent of any D^i. A difference between our slot description encoder and that in [12] is that we leverage the embeddings of the IOB tags so that the dependencies between tags in different slot labels are modeled.

3.2. BiLSTM Feature Extractor

Bidirectional long short-term memory (BiLSTM) has been widely utilized in sequence models to capture the contextual semantic feature of input sentences [19,29]. In eCRF, we also exploit BiLSTMs to extract the contextual neural features. Through concatenating the hidden states from both forward and backward passes, we acquire the distributed representations of contextual features $H = (h_1, h_2, ...h_n)$, in which each $h_i \in \mathbb{R}^d$.

3.3. Elastic CRF (eCRF) Labeler

Let $Y = (y_1, y_2, ...y_n)$ denote the output sequence of slot labels, where $y_i \in \mathcal{S}$. Then the potential function of our elastic neural CRF is defined as follows:

$$\Psi(Y, W) = \sum_{i=1}^{n} e(y_i)^T h_i + \sum_{i=1}^{n-1} e(y_i)^T W e(y_{i+1}) \tag{4}$$

where $W \in \mathbb{R}^{d \times d}$ is a learnable matrix. The potential function consists of two items. The first term, called the node potential, calculates semantic similarity of the slot descriptions and the extracted contextual features. The second term, called the edge potential, captures

interactions between the slot labels through a bilinear calculation. Then, the likelihood of eCRF is defined as follows:

$$p(Y|X, D) = \frac{\exp(\Psi(Y, X))}{\sum_{Y'} \exp(\Psi(Y', X))} \tag{5}$$

The eCRF is trained by conditional maximum likelihood (CML), and we used Viterbi decoding for inferences as follows:

$$\hat{Y} = \underset{y'_1, \ldots, y'_n \in \mathcal{S}}{\mathrm{argmax}} \; p(y'_1, \ldots, y'_n | X, D) \tag{6}$$

In our experiment, we employed the pre-train trick [31] to speed up model learning. Namely, we first masked the edge potential term and trained only with the node potential term for a certain number of training steps, and then added the edge potentials in training. More details can be found in Section 5.2.

4. Dataset and Tasks

In the experiments, we used the recent Google simulated dataset (accessed from https://github.com/google-research-datasets/simulated-dialogue on 1 June 2018) as our main dataset. It is collected by the machines talking to machines (M2M) self-play schema [13]. Two domains, restaurant and movie, were chosen. There are two common slots, i.e., time and date, in both domains, and an around 40% out-of-vocabulary (OOV) rate in the test sets. However, since this dataset was not originally built for the open-ontology slot filling, the number of unseen values in the testing set is very limited. In order to properly use this dataset for the study, we designed two different tasks, the in-domain task and the cross-domain task, and accordingly re-split the whole dataset into new training and testing sets.

In the in-domain task, we aimed to evaluate various models for handling unknown values given all known slots. For each domain, we re-split the whole dataset by fixing the ratio between the number of types of values in training and testing. Suppose the sets of all values occurred in the training set and testing set are V_{train} and V_{test}, respectively; we defined the value ratio between training and testing as $|V_{train}| : |V_{test} - V_{train}|$. Three value ratios were chosen for model evaluations, that is, 75:25, 50:50 and 25:75.

For the cross-domain task, we aimed to evaluate various models for handling unknown slots. Similar to the zero-shot multi-domain learning [12], we trained the model on one domain and evaluated it on the other domain. The common slots of the two domains are treated as known slots while the other slots were treated as unknown slots.

After determining the training and testing sets, a validation set is randomly extracted from the training set, satisfying two conditions: (1) the ratio between the total number of utterances in the new training set and validation set is 4:1, and (2) around 50% of the validation set contains unseen slots or values with respect to the new training set. In this way, a reasonable validation set is constructed so that model training can be monitored for stopping for open-ontology prediction.

5. Experiments

5.1. Baselines

In this paper, we compare our eCRF model with the concept tagging model proposed in [12] and a simple BiLSTM-based tagging model.

As shown in Figure 3, the Concept tagging (CT) model employs a slot description encoder that takes the slot descriptions as input without the IOB tags. A one-layer BiLSTM is used to extract the contextual features of user utterances. The contextual features and the description encoder outputs are concatenated and sent to a feedforward neural network (FNN). This is followed by another one-layer BiLSTM. Finally, a softmax layer is used to calculate the distribution over slot labels. Since the slot descriptions are already used as conditional inputs, the output slot label set only consists of three labels, i.e., 'I', 'B', 'O'. In

both training and testing, the descriptions of each slot are iteratively fed into the model and evaluated separately.
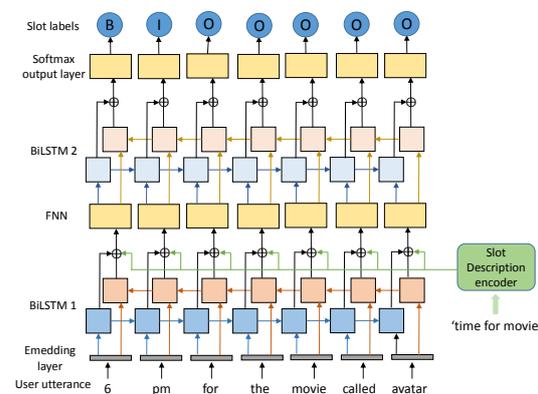

Figure 3. The architecture of the concept Tagging (CT) model [12].

The BiLSTM tagging (BT) model is a simplified version of the CT model, created by removing the second BiLSTM layer. As shown in the following experimental results, this second BiLSTM layer plays an important role in transforming the contextual features and slot label features, which largely improves the performance.

5.2. Experimental Setup

In our experiment, the vocabulary size is 1264. We use the open tool (accessed from https://github.com/stanfordnlp/GloVe on 25 October 2015) to train the GloVe embeddings on the whole dataset. The dimension of all word embeddings and the IOB tags are set as 50. The concatenated hidden size of all BiLSTMs are set as 100. The FNNs in the CT and BT models consist of one hidden layer with 100 units. For the pre-training of eCRFs, the edge potential is added in training after 2000 steps. All models are trained with the Adam [32] optimization method with a learning rate of 0.001. Early-stopping is employed on the validation set to prevent over-fitting. For both the CT and BT models, we leveraged oversampling, which sets the ratio of positive and negative samples as 1:1 and trains the model with a minibatch size of 10. For eCRFs, we set the minibatch size as 1. All the codes were implemented with Tensorflow [33].

5.3. In-Domain Task Results

As described in Section 4, for the in-domain tasks, we re-organized the whole dataset into three different new datasets with increasing prediction difficulties, by setting the value ratios between training and testing as 75:25, 50:50 and 25:75. Table 1 shows the average exact-matching accuracies for known values, unknown values, and total values on the testing set for each model.

Table 1. Results for the in-domain tasks: average exact matching accuracies for known values, unknown values and total values for three models. Models are BiLISM tagging (BT) model, concept tagging (CT) model [12] and elastic CRF (eCRF). Sim-R and sim-M are the domains of restaurant and movie respectively. For each domain, three ratios between the number of types of values in training and testing are chosen to re-split the whole dataset to train models. Bold numbers mean the best results among three compared models.

Domain	Value-Ratio	Average Accuracy for Known Values			Average Accuracy for Unknown Values			Average Accuracy for Total Values		
	Train: Test	**BT**	**CT**	**eCRF**	**BT**	**CT**	**eCRF**	**BT**	**CT**	**eCRF**
sim-R	75:25	0.959 ± 0.020	**0.993 ± 0.005**	0.982 ± 0.007	0.555 ± 0.122	0.753 ± 0.108	**0.791 ± 0.047**	0.765 ± 0.069	0.862 ± 0.060	**0.875 ± 0.026**
	50:50	0.968 ± 0.017	**0.994 ± 0.002**	0.984 ± 0.011	0.361 ± 0.083	0.474 ± 0.066	**0.618 ± 0.058**	0.639 ± 0.048	0.677 ± 0.042	**0.754 ± 0.035**
	25:75	0.967 ± 0.041	**0.999 ± 0.001**	0.985 ± 0.009	0.365 ± 0.034	0.441 ± 0.035	**0.516 ± 0.036**	0.554 ± 0.016	0.575 ± 0.030	**0.624 ± 0.027**
sim-M	75:25	0.951 ± 0.034	0.982 ± 0.005	**0.984 ± 0.003**	0.843 ± 0.009	0.876 ± 0.066	**0.905 ± 0.011**	0.914 ± 0.018	0.930 ± 0.037	**0.953 ± 0.005**
	50:50	0.941 ± 0.028	**0.982 ± 0.009**	0.975 ± 0.017	0.655 ± 0.024	0.723 ± 0.076	**0.841 ± 0.024**	0.803 ± 0.014	0.840 ± 0.040	**0.910 ± 0.017**
	25:75	0.948 ± 0.024	**0.991 ± 0.003**	0.988 ± 0.005	0.519 ± 0.034	0.611 ± 0.030	**0.682 ± 0.035**	0.662 ± 0.027	0.718 ± 0.021	**0.784 ± 0.023**

The results demonstrate that eCRFs clearly outperform the BT models in all conditions. Though slightly worse than the CT models on known values, eCRFs achieve much better results than the CT models in terms of accuracies for unknown values. And the superiority becomes larger as the value-ratio in testing set becomes higher. Therefore, in terms of accuracies for total values, eCRFs achieve the best overall performances.

5.4. Cross-Domain Task Results

For the cross-domain tasks, we train models on one domain and test on the other. The common slots such as `time`, `date` are treated as known slots while the rest as unknown slots, such as `theatre_name`, `restaurant_name`. The evaluation metrics are the average exact-matching accuracies for values from known slots, unknown slots and total slots on the target domain. As shown in Table 2, eCRFs outperform other models in all conditions. In the cross-domain tasks, although there are some overlapping between the known slots on the two domains, the user utterances are different in expressing those slots and values. These results demonstrate that our eCRFs have greater generalization ability.

Table 2. Results for the cross-domain tasks: average exact matching accuracies for values from known slots, unknown slots and total slots on test domain for three models. Bold numbers mean the best results among three compared models.

Train Domain	Test Domain	Average Accuracy for Known Slots			Average Accuracy for Unknown Slots			Average Accuracy for Total Slots		
		BT	CT	eCRF	BT	CT	eCRF	BT	CT	eCRF
sim-M	sim-R	0.980 ± 0.025	0.974 ± 0.009	$\mathbf{0.988 \pm 0.004}$	0.136 ± 0.045	0.121 ± 0.077	$\mathbf{0.243 \pm 0.009}$	0.502 ± 0.036	0.491 ± 0.044	$\mathbf{0.566 \pm 0.007}$
sim-R	sim-M	0.814 ± 0.064	0.915 ± 0.013	$\mathbf{0.926 \pm 0.024}$	0.165 ± 0.040	0.246 ± 0.017	$\mathbf{0.377 \pm 0.031}$	0.508 ± 0.035	0.599 ± 0.006	$\mathbf{0.667 \pm 0.020}$

Figures 4–6 show the prediction results for the same utterance on the movie domain with the eCRF and CT models. Figure 4 illustrates the predicted scores with only node potentials for eCRFs, while Figure 5 gives the predicted scores with both node and edge potentials. It can be seen that the boundaries of slot labels for some slots are mistakenly placed in Figure 4, e.g., the value "*lincoln square cinemas*" for the unknown slot `theatre_name` is falsely predicted as two values "*lincoln*" and "*square cinemas*". When taking both node and edge potentials into account, correct predictions are obtained for all the three slots, as shown in Figure 5. The output probabilities of slot labels for the CT model are shown in Figure 6. Although the CT model gives the right prediction for the known slot `date` and unknown slot `#tickets`, it mistakenly predicts the value for the unknown slot `theatre_name` as "*lincoln square*", as it fails to learn the semantic relations between slot labels.
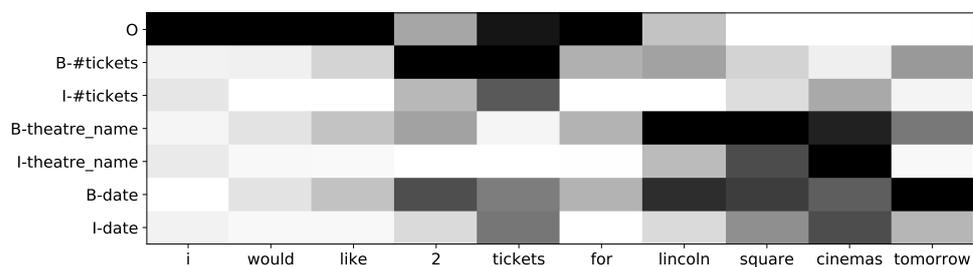

Figure 4. Potential scores with only node potentials in eCRFs for the cross-domain task. The darker the color, the higher the potential score.
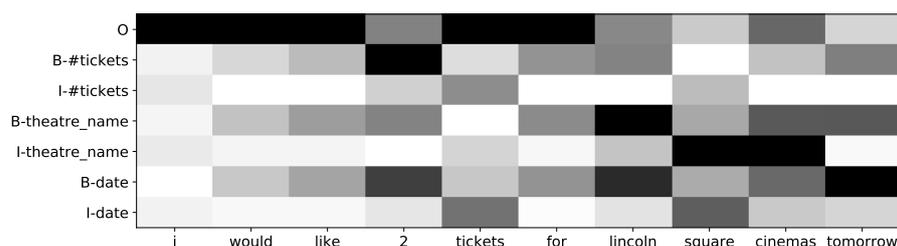
Figure 5. Potential scores with both node and edge potentials in eCRFs for the cross-domain task.
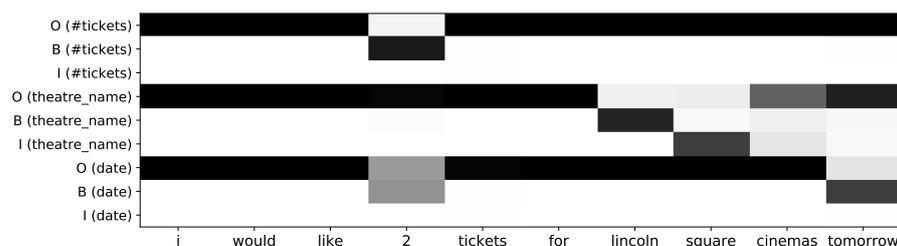

Figure 6. Probabilities of the IOB labels for each slot in the CT model.

6. Conclusions

In this paper, we propose a novel model, the elastic conditional random field (eCRF), for open-ontology slot-filling task. The natural language descriptions of slots and (user) utterances are encoded into the same semantic embedding space to implement the node and edge potentials. We recompose the Google simulated dataset and demonstrate that eCRFs achieve better performances in both in-domain tasks and cross-domain tasks than existing models.

There are interesting future works to further enhance the parsing ability and adaptation capacity of eCRFs: (1) encoding the descriptions of more semantic labels including the intent labels, domain labels and action labels for better generalization and (2) upgrading the CRF architecture with a slot label language model that can capture long-range dependencies between labels.

Author Contributions: Methodology: Y.D., Y.Z. and Z.O.; software: Y.D. and Y.Z.; validation: Y.D., Y.Z. and H.L.; investigation: Y.D., Y.Z., H.L., Z.O. and Y.H.; writing: Y.D., Y.Z. and Z.O.; project administration: Z.O. and J.F. All authors have read and agreed to the published version of the manuscript.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Wang, Y.Y.; Deng, L.; Acero, A. Spoken language understanding. *Signal Process. Mag. IEEE* **2005**, *22*, 16–31. [CrossRef]
2. Mesnil, G.; Dauphin, Y.; Yao, K.; Bengio, Y.; Deng, L.; Hakkani-Tur, D.; He, X.; Heck, L.; Tur, G.; Yu, D. Using recurrent neural networks for slot filling in spoken language understanding. *IEEE/ACM Trans. Audio Speech Lang. Process.* **2015**, *23*, 530–539. [CrossRef]
3. Ramshaw, L.A.; Marcus, M.P. *Text Chunking Using Transformation-Based Learning*; Springer: Dordrecht, The Netherlands, 1999.
4. Lafferty, J.D.; Mccallum, A.; Pereira, F.C.N. Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. In Proceedings of the ICML, Williamstown, MA, USA, 28 June–1 July 2001.
5. Liu, B.; Lane, I. *Attention-Based Recurrent Neural Network Models for Joint Intent Detection and Slot Filling*; Interspeech: San Francisco, CA, USA, 2016.
6. Sang, E.F.; De Meulder, F. Introduction to the CoNLL-2003 Shared Task: Language-Independent Named Entity Recognition. In Proceedings of the HLT-NAACL, Edmonton, AB, Canada, 27 May 2003.
7. Liu, B.; Lane, I. Recurrent Neural Network Structured Output Prediction for Spoken Language Understanding. In Proceedings of the NIPS Workshop on Machine Learning for Spoken Language Understanding and Interactions, Montreal, QC, Canada, 11 December 2015.

8.   Rastogi, A.; Hakkani-Tür, D.Z.; Heck, L.P. Scalable multi-domain dialogue state tracking. In Proceedings of the 2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU), Okinawa, Japan, 16–20 December 2017; pp. 561–568.

9.   Hakkani-Tür, D.; Tur, G.; Celikyilmaz, A.; Chen, Y.N.; Gao, J.; Deng, L.; Wang, Y.Y. *Multi-Domain Joint Semantic Frame Parsing using Bi-directional RNN-LSTM*; InterSpeech: San Francisco, CA, USA, 2016.

10.   Mrksic, N.; Séaghdha, D.Ó.; Thomson, B.; Gasic, M.; Su, P.H.; Vandyke, D.; Wen, T.H.; Young, S.J. Multi-Domain Dialog State Tracking Using Recurrent Neural Networks. In Proceedings of the ACL, Beijing, China, 26–31 July 2015.

11.   Jaech, A.; Heck, L.P.; Ostendorf, M. *Domain Adaptation of Recurrent Neural Networks for Natural Language Understanding*; Interspeech: San Francisco, CA, USA, 2016.

12.   Bapna, A.; Tür, G.; Hakkani-Tür, D.Z.; Heck, L.P. *Towards Zero-Shot Frame Semantic Parsing for Domain Scaling*; Interspeech: Stockholm, Sweden, 2017.

13.   Shah, P.; Hakkani-Tür, D.Z.; Tür, G.; Rastogi, A.; Bapna, A.; Nayak, N.; Heck, L.P. Building a Conversational Agent Overnight with Dialogue Self-Play. *arXiv* **2018**, arXiv:1801.04871.

14.   Larochelle, H.; Erhan, D.; Bengio, Y. Zero-Data Learning of New Tasks. In Proceedings of the AAAI 2014, Québec City, QC, Canada, 27–31 July 2014.

15.   Chen, Y.N.; Hakkani-Tür, D.; He, X. Zero-Shot Learning of Intent Embeddings for Expansion by Convolutional Deep Structured Semantic Models. In Proceedings of the ICASSP 2016, Shanghai, China, 20–25 March 2016.

16.   Zhao, T.; Eskénazi, M. *Zero-Shot Dialog Generation with Cross-Domain Latent Actions*; SIGDIAL: Edinburgh, UK, 2018.

17.   Shah, D.; Gupta, R.; Fayazi, A.; Hakkani-Tur, D. Robust Zero-Shot Cross-Domain Slot Filling with Example Values.   In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Florence, Italy, 28 July–2 August 2019; pp. 5484–5490.

18.   Lin, Z.; Liu, B.; Moon, S.; Crook, P.; Zhou, Z.; Wang, Z.; Yu, Z.; Madotto, A.; Cho, E.; Subba, R. Leveraging Slot Descriptions for Zero-Shot Cross-Domain Dialogue StateTracking. In Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Online, 4 November 2021; pp. 5640–5648. [CrossRef]

19.   Chen, T.; Xu, R.; He, Y.; Wang, X. Improving sentiment analysis via sentence type classification using BiLSTM-CRF and CNN. *Expert Syst. Appl.* **2017**, *72*, 221–230. [CrossRef]

20.   Xu, P.; Sarikaya, R. *Convolutional Neural Network Based Triangular CRF for Joint Intent Detection and Slot Filling*; ASRU: Olomouc, Czech Republic, 2014; pp. 78–83.

21.   Kurata, G.; Xiang, B.; Zhou, B.; Yu, M. Leveraging Sentence-Level Information with Encoder LSTM for Semantic Slot Filling. In Proceedings of the EMNLP 2016, Austin, TX, USA, 1–5 November 2016.

22.   Vu, N.T.; Gupta, P.; Adel, H.; Schütze, H. Bi-Directional Recurrent Neural Network with Ranking Loss for Spoken Language Understanding. In Proceedings of the ICASSP 2016, Shanghai, China, 20–25 March 2016.

23.   Chen, Y.N.; Hakkani-Tür, D.; Tur, G.; Gao, J.; Deng, L. *End-to-End Memory Networks with Knowledge Carryover for Multi-Turn Spoken Language Understanding*; InterSpeech: San Francisco, CA, USA, 2016.

24.   Liu, Z.; Winata, G.I.; Xu, P.; Fung, P. Coach: A Coarse-to-Fine Approach for Cross-domain Slot Filling. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Online, 5–10 July 2020; pp. 19–25.

25.   Devlin, J.; Chang, M.W.; Lee, K.; Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv* **2018**, arXiv:1810.04805.

26.   Li, X.; Feng, J.; Meng, Y.; Han, Q.; Wu, F.; Li, J.   A unified MRC framework for named entity recognition.   *arXiv* **2019**, arXiv:1910.11476.

27.   Gao, S.; Agarwal, S.; Chung, T.; Jin, D.; Hakkani-Tur, D. From machine reading comprehension to dialogue state tracking: Bridging the gap. *arXiv* **2020**, arXiv:2004.05827.

28.   Yu, M.; Liu, J.; Chen, Y.; Xu, J.; Zhang, Y. Cross-Domain Slot Filling as Machine Reading Comprehension. In Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21, Montreal, QC, Canada, 19–26 August 2021; Zhou, Z.H., Ed.; pp. 3992–3998.

29.   Ma, X.; Hovy, E. End-to-End Sequence Labeling via Bi-Directional LSTM-CNNs-CRF. In Proceedings of the ACL, Berlin, Germany, 7–12 August 2016.

30.   Lample, G.; Ballesteros, M.; Subramanian, S.; Kawakami, K.; Dyer, C. Neural Architectures for Named Entity Recognition. In Proceedings of the NAACL-HLT 2016, San Diego, CA, USA, 12–17 June 2016.

31.   Belanger, D.; McCallum, A. Structured Prediction Energy Networks. In Proceedings of the ICML 2016, New York City, NY, USA, 19–24 June 2016.

32.   Kingma, D.P.; Ba, J. Adam: A Method for Stochastic Optimization. In Proceedings of the International Conference on Learning Representations, Banff, AB, Canada, 14–16 April 2014.

33.   Abadi, M.; Barham, P.; Chen, J.; Chen, Z.; Davis, A.; Dean, J.; Devin, M.; Ghemawat, S.; Irving, G.; Isard, M. TensorFlow: A System for Large-Scale Machine Learning. In Proceedings of the OSDI 2016, Savannah, GA, USA, 2–4 November 2016.