

# Prior Knowledge Driven Label Embedding for Slot Filling in Natural Language Understanding

Su Zhu , *Student Member, IEEE*, Zijian Zhao, Rao Ma , and Kai Yu , *Senior Member, IEEE*

**Abstract**—Traditional slot filling in natural language understanding (NLU) predicts a one-hot vector for each word. This form of label representation lacks semantic correlation modeling, which leads to severe data sparsity problem, especially when adapting an NLU model to a new domain. To address this issue, a novel label embedding based slot filling framework is proposed in this article. Here, distributed label embedding is constructed for each slot using prior knowledge. Three encoding methods are investigated to incorporate different kinds of prior knowledge about slots: *atomic concepts, slot descriptions, and slot exemplars*. The proposed label embeddings tend to share text patterns and reuses data with different slot labels. This makes it useful for adaptive NLU with limited data. Also, since label embedding is independent of NLU model, it is compatible with almost all deep learning based slot filling models. The proposed approaches are evaluated on three datasets. Experiments on single domain and domain adaptation tasks show that label embedding achieves significant performance improvement over traditional one-hot label representation as well as advanced zero-shot approaches.

**Index Terms**—Natural language understanding, slot filling, label embedding, prior knowledge, domain adaptation.

## I. INTRODUCTION

NOWDAYS, the development of mobile internet and smart devices has led to the tremendous growth of conversational dialogue systems, such as Amazon Alexa, Google Assistant, Apple Siri and Microsoft Cortana. Natural language understanding (NLU) is a key component of these systems, parsing user's utterances into the corresponding semantic representations [1] for certain narrow domains (e.g., *booking restaurant, searching flight*). As a main task of NLU, *slot filling* is typically treated as a sequence labelling problem in which contiguous sequences of words are tagged with semantic labels (slots) [2]–[6].

Deep learning has achieved great success for the slot filling task in NLU [5]–[11], outperforming most traditional approaches [4], [12]. However, deep learning is notorious for performing poorly with limited labelled data. This is called

Manuscript received July 30, 2019; revised October 31, 2019 and February 5, 2020; accepted March 6, 2020. Date of publication March 12, 2020; date of current version May 22, 2020. The work of S. Zhu, Z. Zhao, R. Ma, and K. Yu was supported by the National Key Research and Development Program of China under Grant 2017YFB1002102 and the China NSFC projects under Grant 61573241. Experiments have been carried out on the PI supercomputer at Shanghai Jiao Tong University. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Eric Fosler-Lussier. (Corresponding author: Kai Yu.)

The authors are with the SpeechLab, Department of Computer Science and Engineering, and MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, Shanghai 200240, China (e-mail: paul2204@sjtu.edu.cn; 1248uu@sjtu.edu.cn; rm1031@sjtu.edu.cn; kai.yu@sjtu.edu.cn).

Digital Object Identifier 10.1109/TASLP.2020.2980152

the *data sparsity* problem, which usually refers to limited data samples of feature-label pairs. In many supervised learning tasks (e.g., part-of-speech tagging, named-entity recognition), the data sparsity problem mainly lies on feature space since their label spaces are fixed. However, the NLU slot-filling task is domain-specific, and the output semantic labels are defined for all possible semantic scopes of the domain. This brings about *label space data sparsity*, which is an even more serious problem in NLU slot-filling task:

- The label space of slot-filling task defined in each domain is distinct from others. Therefore, it is a challenge to exploit data of different domains efficiently.
- Narrow domains are increasingly popular in commercial dialogue systems. Growing narrow domains lead to fine-grained semantic labels (slots). For example, domain WEATHER concerns `city_name`, while domain TRAVEL makes a distinction between `arrival_city` and `departure_city`.

To mitigate the label space data sparsity problem, a label embedding based method for slot filling in NLU is proposed. The label embedding is driven by prior knowledge about slot definitions, which implicates relations of different semantic labels and promotes the adaptive training of slot filling. It changes the representation of each semantic label from the one-hot vector to a distributed one. For example, slots `arrival_city` and `departure_city` are orthogonal in one-hot vector space, but partially the same in the `city` dimension. Label embedding implicating relations of slots could share patterns and reuse the data of associated slots.

Unlike input word embeddings [13], [14] which could be trained with a large number of unlabelled text corpora, label embeddings involve prior knowledge from domain experts or developers. Three kinds of human knowledge and the corresponding encoding methods are investigated in this paper: atomic concept, slot description encoding and slot exemplar encoding. The atomic concept refers to a pulverized slot, which is smaller semantic unit than the slot, e.g., the atomic concepts of slot `arrival_city` can be `destination` and `city_name`. The slot description literally interprets each slot in natural language, e.g., slot `arrival_city` can be described as “city name of the destination”. The slot exemplar is an annotation example for each slot.

The atomic concept is initially proposed in our previous work [15], while label is not represented as an embedding vector there. Here, we exploit label embedding independent of NLU model, which can be compatible with most deep learning-based

Words	Show	flights	from	Boston	to	New	York	today
Tags	O	O	O	B-FromCity	O	B-ToCity	I-ToCity	B-Date

Fig. 1. An example of annotation for slot filling in domain *searching flight*.

slot-filling models. To further model the temporal dependencies of labels, a new label transition module based on label embedding is also proposed within the conditional random field (CRF) [16]. Our proposed framework is evaluated on DSTC 2&3 [17] and SNIPS [18] benchmarks and a custom dataset (AICar), outperforming all baselines significantly.

The following section discusses how our method relates to the previous works. We also introduce the slot filling task in Section III. Then, we describe the details of label embeddings for slot filling in Section IV and V. The results of the extensive experiments are given in Section VI. We conclude and give some future research directions in Section VII.

## II. RELATED WORK

Standard approaches to solving the data sparsity problem of slot filling include semi-supervised learning [19]–[22], domain adaptation [23]–[31] and zero-shot learning [32]–[37]. The semi-supervised learning and domain adaptation (mostly parameters sharing) methods focus on improving the coverage of feature space while not label space. The zero-shot learning methods involve ontological descriptions to enhance the slot filling. However, Ferreira *et al.* [32] and Yazdani *et al.* [33] made an assumption that all possible values of each slot are known, which is not practical for real-world applications. Our methods seek to interpret semantic labels (slots) in multiple dimensions where relations of slots can be inferred implicitly. It helps data reuse for different labels and domains.

Label embeddings for SLU have been introduced in several works. Chen *et al.* [34] exploit label embedding for intent detection but not suitable for slot filling. Some works choose to extract label embeddings from data samples [38]–[41] by exploiting words of values tagged with a semantic label, but they only focus on the corresponding values. We consider both the value and context information in the slot exemplar encoding. Concept tagger [35] and zero-shot adaptive transfer [36] also incorporate slot embeddings from slot descriptions for the slot filling. However, they proposed a slot-independent conditional layer to predict the existence for each slot one by one, which would take more time for training and inference. Moreover, it is not possible for modelling the time series dependence of different slots in those methods. We propose label embeddings for slot filling without changing the basic workflow, which overcomes the above issues.

## III. SLOT FILLING IN NLU

Slot filling is a major task of natural language understanding (NLU) in task-oriented dialogue systems, which automatically extracts a set of attributes or “slots,” and the corresponding values. It is typically treated as a sequence labelling problem. An example of data annotation is provided in Fig. 1. It follows the popular inside/outside/beginning (IOB) schema, where *Boston*

and *New York* are the departure and arrival cities specified as the slot values in the user’s utterance, respectively. In this paper, we use the word *tag* to refer to the semantic label in IOB schema, which is different from *slot*.

In this section, we will sequentially introduce the slot filling in NLU for single domain and domain adaptation.

### A. Single Domain

Let  $\mathbf{w} = (w_1, w_2, \dots, w_N)$  denote an input sequence (words), and  $\mathbf{t} = (t_1, t_2, \dots, t_N)$  denote its output sequence (tags) in domain  $d$ , where  $N$  is the sequence length and  $t_i$  is one of the predefined tags, i.e.  $t_i \in \mathcal{T}_d$ . Let  $\mathcal{T}_d$  denote all possible labels of domain  $d$ .  $t_i$  would be represented as a one-hot vector,  $\mathbf{o}(t_i)$ , in conventional methods. The slot filling in NLU is to estimate  $p(\mathbf{t}|\mathbf{w})$ , the posterior probability of output sequence  $\mathbf{t}$  given input  $\mathbf{w}$ .

Recently, motivated by a number of successful neural network and deep learning methods in natural language processing, many neural network architectures have been applied in this task, such as vanilla recurrent neural network (RNN) [5], [42]–[44], convolutional neural network (CNN) [8], [45], long short-term memory (LSTM) [7], [46]–[48], encoder-decoder [6], [9], [10], [49]. Bi-directional LSTM based RNN (BLSTM) is adopted in this paper. Every input word is mapped to a vector via  $\mathbf{w}_i = \mathbf{W}_{in}\mathbf{o}(w_i)$ , where  $\mathbf{W}_{in}$  is an embedding matrix and  $\mathbf{o}(w_i)$  a one-hot vector. Hidden vectors are recursively computed at the  $i$ -th time step via:

$$\mathbf{h}_i = [\vec{\mathbf{h}}_i; \overleftarrow{\mathbf{h}}_i]; \vec{\mathbf{h}}_i = \mathbf{f}_{\text{lstm}}(\mathbf{w}_i, \vec{\mathbf{h}}_{i-1}); \overleftarrow{\mathbf{h}}_i = \mathbf{b}_{\text{lstm}}(\mathbf{w}_i, \overleftarrow{\mathbf{h}}_{i+1})$$

where  $[\cdot; \cdot]$  denotes vector concatenation,  $\mathbf{h}_i \in \mathbb{R}^{2n}$  ( $n$  is the hidden size),  $\mathbf{f}_{\text{lstm}}$  and  $\mathbf{b}_{\text{lstm}}$  are the LSTM functions for forward and backward passes respectively. At each time step, a linear output layer is applied to predict tags, i.e.

$$\mathbf{z}_i = \mathbf{W}_o \mathbf{h}_i \quad (1)$$

where  $\mathbf{W}_o \in \mathbb{R}^{|\mathcal{T}_d| \times 2n}$  denotes trainable parameters of the output layer (bias is omitted). In the following sections, prior knowledge driven label embeddings will be included in  $\mathbf{W}_o$ .

1) *BLSTM-Softmax*: The posterior probability is factorized:

$$p(\mathbf{t}|\mathbf{w}) = \prod_{i=1}^N p(t_i|\mathbf{h}_i) = \prod_{i=1}^N \text{softmax}(\mathbf{z}_i)^\top \mathbf{o}(t_i) \quad (2)$$

where  $\mathbf{o}(t_i)$  is a one-hot vector with a dimension for tag  $t_i$ , so  $\text{softmax}(\mathbf{z}_i)^\top \mathbf{o}(t_i)$  is precisely the  $t_i$ -th element of the distribution defined by the *softmax* function.

2) *BLSTM-CRF*: To model the time series dependence of output labels, several methods based on encoder-decoder architectures [6], [9], [10], [49] are proposed for slot filling. Besides that, CRF output layer can jointly decode the best chain of labels for a given input sentence [50]–[52]. The CRF output layer calculates the posterior probability at the sentence level:

$$p(\mathbf{t}|\mathbf{w}) = \frac{\exp(\psi(\mathbf{w}, \mathbf{t}))}{\sum_{\mathbf{t}'} \exp(\psi(\mathbf{w}, \mathbf{t}'))} \quad (3)$$

TABLE I  
EXAMPLES ABOUT CONFLICTS OF DATA ANNOTATION FROM DIFFERENT  
DOMAINS. WE USE [VALUE: SLOT] FOR ANNOTATION

domain	data samples
DSTC2	I'm looking for a [Thai:Food] restaurant.
DSTC3	I'm looking for a [Thai:Food] [restaurant:Type].
Weather	I'm going to [Beijing:City], and what's the weather?
Travel	I'm going to [Beijing:ToCity], and what's the weather?

The score of a sentence  $w$  along with a path of tags  $\mathbf{t}$  is given by the sum of transition scores and network scores:

$$\psi(w, \mathbf{t}) = \sum_{i=1}^N ([\mathbf{A}]_{t_{i-1}, t_i} + \mathbf{z}_i^\top \mathbf{o}(t_i)) \quad (4)$$

where  $\mathbf{A} \in \mathbb{R}^{|\mathcal{T}_d| \times |\mathcal{T}_d|}$  is a transition matrix, and its element  $[\mathbf{A}]_{m,n}$  modelling the transition from the  $m$ -th to the  $n$ -th label for a pair of consecutive time steps.

The maximum likelihood estimation is used for training, and the negative log-likelihood is given by

$$Loss = -\log p(\mathbf{t}|w)$$

. Decoding is to search for the label sequence  $\mathbf{t}^*$  with the highest posterior probability, i.e.

$$\mathbf{t}^* = \underset{\mathbf{t}'}{\operatorname{argmax}} p(\mathbf{t}'|w)$$

The searching is time independent for the *softmax* output function, while it can be solved efficiently by adopting the Viterbi algorithm for the CRF output layer.

## B. Domain Adaptation

Domain adaptation can alleviate the data sparsity problem of slot filling by taking advantage of data from other domains, such as feature augmentation [23], parameters sharing [24]–[26], and bag of experts [27], [28].

Sharing parameters for different domains is a popular approach to transferring knowledge learned from source domains to a target domain. Suppose we have a set of source domains,  $\mathcal{D}_{src}$ , and a target domain  $td$ . We mix all data of  $\mathcal{D}_{src}$  to pre-train a slot filling model, then fine-tune the model by using data of  $td$ . The pre-training procedure can be implemented only once for different target domains.

In the pre-training stage, different domains in  $\mathcal{D}_{src}$  may not be able to share the output layer fully. Because each narrow domain has a closed and limited semantic space which is different from others. It may lead to conflicts of data annotation. Several examples of disputes between different domains (e.g., DSTC2 V.S. DSTC3) are shown in Table I.

Therefore, the output layer should be masked at different positions, according to which domain the training data is from. For all the source domains, we have a union of labels,  $\mathcal{T}_{src} = \bigcup_{d \in \mathcal{D}_{src}} \mathcal{T}_d$ . Therefore, the output matrix is scaled up to  $\mathbf{W}_o^{src} \in \mathbb{R}^{|\mathcal{T}_{src}| \times 2n}$ , and the transition matrix is  $\mathbf{A}_{src} \in \mathbb{R}^{|\mathcal{T}_{src}| \times |\mathcal{T}_{src}|}$ . For each domain  $d \in \mathcal{D}_{src}$ , there is mask vector  $\mathbf{m}_d \in \mathbb{R}^{|\mathcal{T}_{src}|}$ :

$$[\mathbf{m}_d]_{t_j} = \begin{cases} 0, & \text{if } t_j \in \mathcal{T}_d \\ -\infty, & \text{otherwise} \end{cases} \quad (5)$$

where each  $t_j \in \mathcal{T}_{src}$  is covered. It is used to control the output for domain  $d$  by ignoring all labels out of this domain and replacing Equation (1) with:

$$\mathbf{z}_i = \mathbf{W}_o^{src} \mathbf{h}_i + \mathbf{m}_d \quad (6)$$

At the beginning of the fine-tuning stage of the target domain  $td$ , parameters are copied from the model pre-trained in the source domains for initialization. The word embedding matrix and BLSTM parameters are retained, while the output and transition matrices are copied partially:

$$[\mathbf{W}_o^{td}]_{t_j} = \begin{cases} [\mathbf{W}_o^{src}]_{t_j}, & \text{if } t_j \in \mathcal{T}_{src} \cap \mathcal{T}_{td} \\ \text{random init.}, & \text{otherwise} \end{cases} \quad (7)$$

$$[\mathbf{A}_{td}]_{t_j, t_k} = \begin{cases} [\mathbf{A}_{src}]_{t_j, t_k}, & \text{if } \{t_j, t_k\} \subset \mathcal{T}_{src} \cap \mathcal{T}_{td} \\ \text{random init.}, & \text{otherwise} \end{cases} \quad (8)$$

where  $t_j \in \mathcal{T}_{td}$ ,  $t_k \in \mathcal{T}_{td}$ ,  $\mathbf{W}_o^{td} \in \mathbb{R}^{|\mathcal{T}_{td}| \times 2n}$ ,  $\mathbf{A}_{td} \in \mathbb{R}^{|\mathcal{T}_{td}| \times |\mathcal{T}_{td}|}$ ,  $[\cdot]_j$  denotes the  $j$ -th row of a matrix or vector.

## IV. WHY LABEL EMBEDDING?

As illustrated in Equation (1), there is a score estimates how possible each label  $t_j$  is to appear at the  $i$ -th time step,

$$[\mathbf{z}_i]_{t_j} = (\mathbf{W}_o \mathbf{h}_i)^\top \mathbf{o}(t_j), \quad (9)$$

before the *softmax* or CRF function.

If we consider an alternate low-rank decomposition of  $\mathbf{W}_o$ , i.e.  $\mathbf{W}_o = \mathbf{B}^\top \mathbf{C}$  where  $\mathbf{W}_o \in \mathbb{R}^{|\mathcal{T}_d| \times 2n}$ ,  $\mathbf{B} \in \mathbb{R}^{K \times |\mathcal{T}_d|}$ , and  $\mathbf{C} \in \mathbb{R}^{K \times 2n}$ . Equation (9) can be rewritten as

$$[\mathbf{z}_i]_{t_j} = (\mathbf{B}^\top \mathbf{C} \mathbf{h}_i)^\top \mathbf{o}(t_j) = (\mathbf{C} \mathbf{h}_i)^\top (\mathbf{B} \mathbf{o}(t_j)) \quad (10)$$

$\mathbf{B}$  is a matrix corresponding to the lookup table of label embeddings, and  $\mathbf{C}$  is a matrix of linear transformation, which respectively map the label  $t_i$  and features  $\mathbf{h}_i$  into a joint space.

Imagining that there are already pre-trained or pre-defined label embeddings which reflect relations (similarities and differences) among labels, the model learned from labelled data may even be adapted to labels never exist in the data. It can help reduce the size of required in-domain data, enabling developers to bootstrap new domains quickly.

## V. LABEL EMBEDDING FOR SLOT FILLING

In this section, we will introduce the proposed slot-filling method based on label embeddings, representing each output tag as a distributed vector. However, unlike word embeddings [13], [14] which can be pre-trained with a large number of unlabelled text corpora, data with semantic labels is very limited and in the face of the label space sparsity issue. We derive label embeddings from semantic definitions of certain domains. Three kinds of prior knowledge about slot definitions are investigated as well as the corresponding encoding methods: atomic concept, slot description and slot exemplar.

Since the label (tag) of slot filling is in IOB format, we make a simple notation to split tag into IOB and slot. For each tag  $t$  (e.g., “B-from\_city”), IOB( $t$ ) gets the IOB symbol (e.g., “B”), and SLOT( $t$ ) returns the slot name (e.g., “from\_city”). Note that SLOT(“O”) returns null.

TABLE II  
EXAMPLES OF SLOT REPRESENTATION BY ATOMIC CONCEPTS

slot	atomic concepts
city	{ <i>city_name</i> }
from_city	{ <i>city_name</i> , <i>from_location</i> }
depart_city	{ <i>city_name</i> , <i>from_location</i> }
arrive_airport	{ <i>airport_name</i> , <i>to_location</i> }
city_of_birth	{ <i>city_name</i> , <i>birth</i> }
date_of_birth	{ <i>date</i> , <i>birth</i> }
deny-to_city	{ <i>city_name</i> , <i>to_location</i> , <i>deny</i> }

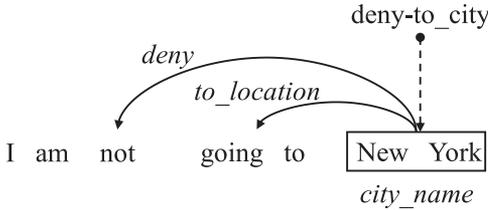


Fig. 2. An example of context dependency of slot `deny-to_city`. *deny* and *to\_location* are context-aware atomic concepts, and *city\_name* focuses on the value “New York”.

### A. Atomic Concept

Atomic concept assumes that there are substructures with smaller semantic units for semantic slots. It is a pure prior knowledge that each slot is represented as a set of atoms. The atomic concepts are pulverized slots and much smaller semantic units than the slots. Each slot is composed of different atomic concepts, e.g., slot `from_city` can be defined as a set of atoms {*from\_location*, *city\_name*}, and slot `date_of_birth` can be represented as {*date*, *birth*}, as shown in Table II.

Atomic concepts are not automatically defined, which are designed manually by domain experts or developers. Atomic concepts can be classified into two categories, one is value-aware, and the other is context-aware. The value-aware atom is often a category name of values, e.g., “New York” is *city\_name*. Context-aware atoms reflect implicit dependencies around the targeted value, e.g., *to\_location* depends on context like “going to,” “fly to,” “arrive at,” as shown in Fig. 2. Modelling on the atomic concept helps find out the linguistic patterns of related slots by atom sharing, and it can even decrease the required amount of data.

1) *Single Domain*: We represent each slot  $s$  of domain  $d$  as a set of atomic concepts, i.e.  $\text{atoms}(s) = \{a_1, \dots, a_{M_s}\}$ ,  $a_l \in \mathcal{A}_d$ , where  $\mathcal{A}_d$  is the vocabulary of atomic concepts and  $M_s$  is the number of atomic concepts belong to this slot. Thus, each slot can be viewed as a binary vector,  $\mathbf{b}(\text{atoms}(s))$  with length of  $|\mathcal{A}_d|$ , while  $\mathbf{b}(\text{atoms}(\text{SLOT}(\text{“O”}))) = \mathbf{0}$ . Due to the IOB schema, the label embedding consists of two parts, 3-dimensional IOB one-hot vector and a binary vector of atoms, as illustrated in Fig. 3(a). The label embedding of tag  $t$  is

$$\text{emb}(t) = [\mathbf{o}(\text{IOB}(t)); \mathbf{b}(\text{atoms}(\text{SLOT}(t)))] \quad (11)$$

where  $\text{emb}(t) \in \mathbb{R}^{3+|\mathcal{A}_d|}$ ,  $t \in \mathcal{T}_d$ .

Finally, a label embedding matrix for domain  $d$  is created, i.e.,  $\mathbf{B}_d \in \mathbb{R}^{(3+|\mathcal{A}_d|) \times |\mathcal{T}_d|}$ , where  $[\mathbf{B}_d^\top]_t = \text{emb}(t)$ . Equation (1) can

be rewritten as

$$\mathbf{z}_i = \mathbf{B}_d^\top \mathbf{C}_d \mathbf{h}_i \quad (12)$$

where  $\mathbf{C}_d \in \mathbb{R}^{(3+|\mathcal{A}_d|) \times 2n}$  denotes trainable parameters of the output layer, but  $\mathbf{B}_d$  is fixed as a kind of prior knowledge.

2) *Domain Adaptation*: Similar to Subsection III-B, the domain adaptation with atomic concepts is also introduced as follows. For all the source domains, we have a union of atoms,  $\mathcal{A}_{src} = \bigcup_{d \in \mathcal{D}_{src}} \mathcal{A}_d$ . Therefore, the output matrix is scaled up to  $\mathbf{C}_{src} \in \mathbb{R}^{(3+|\mathcal{A}_{src}|) \times 2n}$ , and  $\mathbf{B}_d \in \mathbb{R}^{(3+|\mathcal{A}_{src}|) \times |\mathcal{T}_d|}$ . For each domain  $d \in \mathcal{D}_{src}$ , there is mask vector  $\mathbf{m}_d \in \mathbb{R}^{3+|\mathcal{A}_{src}|}$ ,

$$[\mathbf{m}_d]_{a_j} = \begin{cases} 0, & \text{if } a_j \in \mathcal{A}_d \text{ or } a_j \in \text{IOB} \\ -\infty, & \text{otherwise} \end{cases} \quad (13)$$

for each  $a_j \in \mathcal{A}_{src} \cup \text{IOB}$ . It is used to ignore the atomic concepts out of domain  $d$  and replace Equation (6) with:

$$\mathbf{z}_i = \mathbf{B}_d^\top (\mathbf{C}_{src} \mathbf{h}_i + \mathbf{m}_d) \quad (14)$$

For fine-tuning in the target domain  $td$ , the word embedding matrix and BLSTM parameters are still retained. The initialization of  $\mathbf{C}_{td} \in \mathbb{R}^{(3+|\mathcal{A}_{td}|) \times 2n}$  for  $td$  is

$$[\mathbf{C}_{td}]_{a_j} = \begin{cases} [\mathbf{C}_{src}]_{a_j}, & \text{if } a_j \in (\mathcal{A}_{src} \cap \mathcal{A}_{td}) \cup \text{IOB} \\ \text{Random init.}, & \text{otherwise} \end{cases}$$

for each  $a_j \in \mathcal{A}_{td} \cup \text{IOB}$ . Thus, we calculate output scores of the target domain as  $\mathbf{z}_i = \mathbf{B}_{td}^\top \mathbf{C}_{td} \mathbf{h}_i$ , where  $\mathbf{B}_{td} \in \mathbb{R}^{(3+|\mathcal{A}_{td}|) \times |\mathcal{T}_{td}|}$ .

Composition of atomic concepts for each slot provides an explicit interpretation about slot relations. However, it requires a large number of human efforts to design atomic concepts carefully, and it is hard to keep them universal for different domains.

### B. Slot Description Encoding

In this method, we assume that there is a textual description for each slot in natural language. It is also pre-defined by domain experts or developers but with less human labors than the atomic concept. For example, `deny-to_city` can be described as “not this arrival city”. The descriptions of relative slots may include same, synonymous and even antonymous words, which can help reveal the relation between different slots. Thus, we can extract label embeddings literally from slot descriptions.

Unlike the label embedding bound to atomic concepts, an encoder is required to convert literal slot descriptions of variable lengths into fixed dimension vectors, as delineated in Fig. 3(b). Suppose each slot  $s$  has a description,  $\text{desc}(s) = x_1 \cdots x_{L_s}$ , where  $L_s$  is the sequence length. Three encoders are applied: simple mean operation of word embeddings, BLSTM and CNN models. We will introduce them in a single domain, then clarify how to make domain adaptation with slot description-based label embeddings.

1) *Mean Encoder of Word Embeddings*: Every description word is mapped to an  $m$ -dimensional vector via  $\mathbf{x}_i = \mathbf{W}_{in} \mathbf{o}(x_i)$ , where  $\mathbf{W}_{in}$  is the word embedding matrix and  $\mathbf{o}(x_i)$  a one-hot vector. The embedding of slot  $s$  is simply the mean of

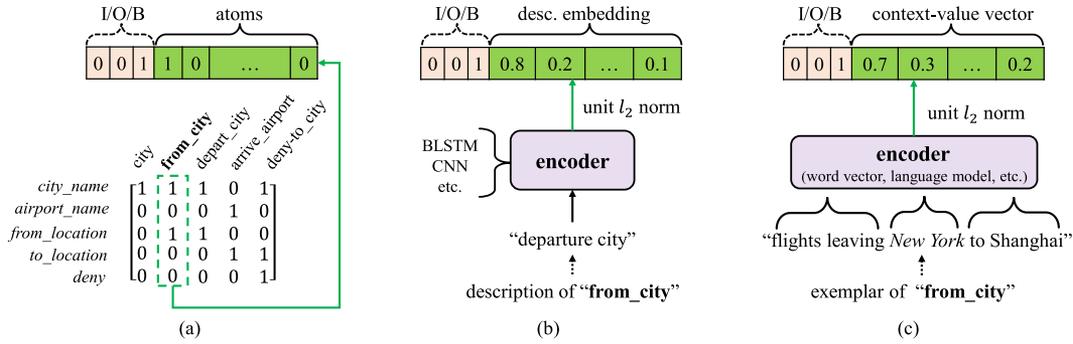


Fig. 3. Three different label embeddings consisted of IOB symbol and slot vector. (a) Atomic concept: each slot is represented as finer-grained semantic units. (b) Slot description encoding: each slot has a textual description in natural language which can be encoded into a vector by several sequence models. (c) Slot exemplar encoding: given a number of exemplars for each slot, the exemplars are encoded into vectors which contain the context and value information.

word embeddings of all words in  $\text{desc}(s)$ , i.e.

$$\text{mean}(\text{desc}(s)) = \text{unit} \left( \frac{1}{L_s} \sum_{l=1}^{L_s} \mathbf{x}_l \right) \quad (15)$$

where  $\text{unit}(\mathbf{v}) = \frac{\mathbf{v}}{\|\mathbf{v}\|}$  gets the unit  $l_2$  norm of vector  $\mathbf{v}$ , which enforces slot embeddings in a limited scale.

2) *BLSTM Encoder*: The hidden vectors of BLSTM encoder are recursively computed at the  $l$ -th time step via:

$$\vec{\mathbf{v}}_l = f_{\text{lstm}}(\mathbf{x}_l, \vec{\mathbf{v}}_{l-1}), l = 1, \dots, L_s \quad (16)$$

$$\overleftarrow{\mathbf{v}}_l = b_{\text{lstm}}(\mathbf{x}_l, \overleftarrow{\mathbf{v}}_{l+1}), l = L_s, \dots, 1 \quad (17)$$

where  $f_{\text{lstm}}$  and  $b_{\text{lstm}}$  are the LSTM functions for forward and backward passes respectively. We concatenate the final hidden vectors of bi-direction to be the slot embedding, i.e.

$$\text{BLSTM}(\text{desc}(s)) = \text{unit}([\vec{\mathbf{v}}_{L_s}; \overleftarrow{\mathbf{v}}_1]) \quad (18)$$

3) *CNN Encoder*: CNN is also a typical model to encode a sequence of words into a fixed-dimensional vector [53]. A linear transformation is applied on a window of  $k$  words to produce a set of new features (we set  $k = 3$  in this work). To avoid that the window size is larger than the sequence length, zero-padding is used on both sides of the input sequence. After processing all possible windows of words, we get a sequence of new feature vectors. A max-over-time pooling operation [54] is applied over new features, which takes the maximum value of each row to form as the final features  $\mathbf{c}$ . Thus,

$$\text{CNN}(\text{desc}(s)) = \text{unit}(\mathbf{c}) \quad (19)$$

4) *Single Domain and Domain Adaptation*: With one of these three encoders, we can encode the description of a slot as a vector. Thus, the label embedding of tag  $t \in \mathcal{T}_d$  is

$$\text{emb}(t) = [\mathbf{o}(\text{IOB}(t)); \text{ENCODER}(\text{desc}(\text{SLOT}(t)))] \quad (20)$$

where ENCODER can be mean, BLSTM and CNN,  $\text{emb}(t) \in \mathbb{R}^K$ ,  $K$  is the embedding length. Note that  $\text{emb}(\text{"O"}) = \mathbf{0}$ .

For any domain  $d \in \mathcal{D}_{src} \cup \{td\}$ , we get a label embedding matrix,  $\mathbf{B}_d \in \mathbb{R}^{K \times |\mathcal{T}_d|}$ , and  $[\mathbf{B}_d^\top]_t = \text{emb}(t)$  for any  $t \in \mathcal{T}_d$ . Both Equation (1) and (6) can be rewritten as

$$\mathbf{z}_i = \mathbf{B}_d^\top \mathbf{C} \mathbf{h}_i \quad (21)$$

where  $\mathbf{z}_i \in \mathbb{R}^{|\mathcal{T}_d|}$ ,  $\mathbf{C} \in \mathbb{R}^{K \times 2n}$  denotes trainable parameters of the output layer, but  $\mathbf{B}_d$  is computed by the slot description encoder on domain  $d$ . Since the encoder is domain-independent,  $\mathbf{C}$  and the encoder can also be fully shared across domains, as well as the sentence BLSTM which produces  $\mathbf{h}_i$ .

This method requires an unstructured short-text to describe each slot, which is easier to be obtained, while the involved label encoders may take more time for computation.

### C. Slot Exemplar Encoding

In this method, we propose to extract label embeddings from slot exemplars which contain values of each slot and their neighbor contexts, as shown in Fig. 3(c). An exemplar of slot  $s$  can be a sentence  $e = e_1 \cdots e_{|e|}$  and the *(start, end)* positions  $(sp, ep)$  for annotation of the slot, i.e.  $e_{sp:ep}$  is a value of  $s$ . For example, one exemplar of slot *from\_city* is "flights leaving New York to Shanghai," and "New York" at positions (3, 4) is its value. Two methods are introduced to obtain label embedding by encoding slot exemplars, which utilize words embeddings and language models.

1) *Word Embeddings*: Every word in the exemplar sentence is mapped to a vector via  $\mathbf{e}_j = \mathbf{W}_e \mathbf{o}(e_j)$ , where  $\mathbf{W}_e$  is a word embedding matrix and  $\mathbf{o}(e_j)$  is a one-hot vector. To specify the value and context information of the slot, we split the exemplar sentence into three segmentations, left context, value and right context, as illustrated in Fig. 4(a). Then, word embeddings in these three parts are respectively averaged and concatenated, composing the slot embedding

$$\text{emb}(s) = \text{unit}([\text{ave}(\mathbf{e}_{<sp}); \text{ave}(\mathbf{e}_{sp:ep}); \text{ave}(\mathbf{e}_{>ep})]) \quad (22)$$

where  $\text{ave}(\mathbf{e}_{<sp})$  and  $\text{ave}(\mathbf{e}_{>ep})$  are zero vectors when  $sp = 1$  and  $ep = |e|$  respectively.

2) *Language Models*: Different from word embeddings, LSTM based bidirectional language models (biLMs) [55] directly encode context information into hidden states. The input sentence  $e = e_1 \cdots e_{|e|}$  is converted into two sequences of hidden vectors  $(\vec{\mathbf{u}}_{1:|e|}, \overleftarrow{\mathbf{u}}_{1:|e|})$  by forward and backward LMs respectively, as shown in Fig. 4(b). Since hidden states already contain information of history and current word, we use the hidden vectors at positions ranging from  $sp$  to  $ep$  to compose

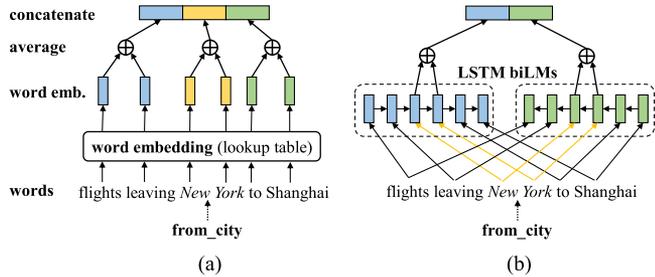


Fig. 4. Two methods to encode data exemplar of a slot, e.g., *from\_city*. (a) Word embeddings: word embeddings in the left context, value and right context are averaged and concatenated respectively. (b) Language models: a biLMs is used to encode the exemplar into context-aware vectors, and we exploit the vectors at the positions of the slot annotation.

the slot embedding:

$$\text{emb}(s) = \text{unit}([\text{ave}(\vec{\mathbf{u}}_{sp:ep}); \text{ave}(\overleftarrow{\mathbf{u}}_{sp:ep})]) \quad (23)$$

3) *Single Domain and Domain Adaptation*: Now, we can encode the exemplars of any slot as a vector (if there are multiple exemplars for one slot, an average is taken). Thus, the label embedding of tag  $t \in \mathcal{T}_d$  is

$$\text{emb}(t) = [\mathbf{o}(\text{IOB}(t)); \text{emb}(\text{SLOT}(t))] \quad (24)$$

where  $\text{emb}(t) \in \mathbb{R}^K$ ,  $K$  is the embedding length, and  $\text{emb}(\text{SLOT}(\text{"O"})) = \mathbf{0}$ .

For any domain  $d \in \mathcal{D}_{src} \cup \{td\}$ , we get the label embedding matrix,  $\mathbf{B}_d \in \mathbb{R}^{K \times |\mathcal{T}_d|}$ , and  $[\mathbf{B}_d^\top]_t = \text{emb}(t)$  for any  $t \in \mathcal{T}_d$ . Both Equation (1) and (6) can be rewritten as

$$\mathbf{z}_i = \mathbf{B}_d^\top \mathbf{C} \mathbf{h}_i \quad (25)$$

where  $\mathbf{z}_i \in \mathbb{R}^{|\mathcal{T}_d|}$ ,  $\mathbf{C} \in \mathbb{R}^{K \times 2n}$  denotes trainable parameters of the output layer which can be shared across domains. However,  $\mathbf{B}_d$  is computed by using pre-trained word embeddings or biLMs in this paper, which will be fixed.

This method requires minimal prior knowledge which is some kind of labeled data sample, while it may heavily rely on the quality of pre-trained word embeddings or biLMs.

#### D. Label Embedding-Based CRF Layer

The typical linear-chain CRF provides the transition between a pair of labels as a *scalar* score, as shown in Equation (4). However, prior knowledge about slots is not exploited in the traditional transition matrix of CRF layer. We propose a new transition model of CRF Layer depending on label embeddings for slot filling. It calculates the transition score using a *bilinear* model. The transition score of two tags  $t_{i-1}$  and  $t_i$  in Equation (4) is computed via:

$$[\mathbf{A}_d]_{t_{i-1}, t_i} = \text{emb}(t_{i-1})^\top \mathbf{W}_a \text{emb}(t_i) \quad (26)$$

where  $t_{i-1} \in \mathcal{T}_d$ ,  $t_i \in \mathcal{T}_d$ ,  $\mathbf{A}_d \in \mathbb{R}^{|\mathcal{T}_d| \times |\mathcal{T}_d|}$  is a transition matrix for domain  $d$ . The parameters of  $\mathbf{W}_a$  are domain independent and can be fully shared across domains. In the inference stage,  $\mathbf{A}_d$  is computed only once per domain.

Compared with the traditional linear-chain CRF whose transition scores are initialized as independent scalars and learned

TABLE III  
DATASET STATISTICS. "AL(desc.\*)" DENOTES THE AVERAGE LENGTH OF SLOT DESCRIPTIONS

Dataset	Domain	Train	Test	#Slot	#Atom	AL(desc.)
DSTC 2&3	DSTC2	15611	9890	17	11	2.2
	DSTC3	109	18715	30	16	2.4
AICar	Map	12000	20661	30	16	4.3
	Weather	10150	7236	36	20	3.4
	Flight	1500	4407	22	18	4.0
SNIPS	-	13784	700	39	-	1.8

from scratch, we compute transition scores by matching label embeddings of start and end tags. Our model could transfer the knowledge about label transition by utilizing implicit slot relations in label embeddings. Besides the bilinear model, there are other alternatives, like attention mechanism [56], [57], multi-head attention [58].

## VI. EXPERIMENTS

In this section, we first introduce the datasets and baselines used. Then, Section VI-C gives the details of experimental setup. In Section VI-D, we compare the performance of our proposed methods with traditional slot filling methods.

### A. Datasets

We carry out experiments on the following NLU datasets.

- **DSTC 2&3** [17]: It contains two domains, DSTC2 (source domain) and DSTC3 (target domain). DSTC2 comprises of dialogues from the restaurant information domain in Cambridge. DSTC3 introduces a tourist information domain about restaurants, pubs and coffee shops in Cambridge, which is an extension of DSTC2 but has only a seed set for training. We use manual transcripts as input by ignoring speech recognition errors. Several slots of DSTC2 are reused in DSTC3.
- **AICar**: The AICar dataset is a custom dataset collected from a commercial in-car spoken dialogue system which has thousands of users. It consists of three domains: MAP navigation, WEATHER forecast and FLIGHT information. There is almost no slot overlap among these domains, yet the slots are partially associated, e.g., MAP has slot *starting\_point*, WEATHER has *city\_name*, and FLIGHT has *departure\_city*.
- **SNIPS**: It is a public NLU dataset [18] of crowdsourced user utterances across 7 intents and about 2000 training samples per intent. It contains natural language corpus collected in a crowdsourced fashion to benchmark the performance of voice assistants.

All of the datasets are annotated with slot labels.<sup>1</sup> The detailed statistics of each dataset are shown in Table III. The language of DSTC 2&3 and SNIPS is English, and AICar is in Chinese. To avoid Chinese word segmentation errors [59], the slot filling annotation of AICar is at Chinese-character level.

<sup>1</sup>Notice that because the original semantic annotations are not aligned with words in DSTC 2&3, we simply use string matching for alignment, which converts the semantic annotations into the slot filling format.

## B. Baselines

We compare with one traditional method, two label embedding methods and two zero-shot learning based approaches:

- One-hot: A 2-layer BLSTM model and one linear output layer (CRF layer is optional) for slot filling, as shown in Section III. Each label is discrete and represented as a one-hot vector in this method.
- Canonical Correlation Analysis (CCA) label embedding [38]: CCA is exploited to induce slot vector by analysis cooccurrences of slots and words, i.e. words tagged with the corresponding slots.
- Value word embedding [39]: We take the average of pre-trained embeddings of the value words for slot embedding, while ignoring the context words.
- Concept Tagger (CT) [35] and Zero-shot Adaptive Transfer (ZAT) [36]: These two zero-shot methods also condition slot filling on textual slot descriptions, which have a contextual BLSTM layer and slot dependent BLSTM layer. However, they predict an IOB sequence for each slot, which may lead to overlapped segmentation and cost more time.

## C. Experimental Setup

1) *Label Embeddings of Slot Filling*: Three different approaches to extracting label embeddings are investigated in this paper, via atomic concepts, slot descriptions and slot exemplars, which involve different degrees of human knowledge. Table III also provides the detailed statistics of label representations. For atomic concepts, we act as domain experts to manually split each slot into small units for all five domains, e.g., `deny_starting_point` is split into “deny,” “from\_location” and “poi\_name”. As an extension of DSTC2, DSTC3 contains 6 new atoms that constitute about 25% of slot tags. For slot descriptions, we leverage tokenized slot names since each slot name is already a meaningful phrase in the datasets, e.g., `deny_starting_point` is described as “deny starting point”. The training data can be directly exploited as data exemplars for the corresponding slots.

2) *Network Hyper-Parameters*: The experimental setup was kept the same for all three datasets with minor adjustments in hyper-parameters. We adopt a 2-layer BLSTM model to encode sentences with 200 hidden units. For training, the network parameters are randomly initialized under the uniform distribution  $[-0.2, 0.2]$ . We use optimizer Adam [60] with learning rate 0.001 for all experiments. The *dropout* with a probability of 0.5 is applied to the non-recurrent connections during the training stage. The batch size is 20 for domain DSTC2, MAP and WEATHER, 5 for domain DSTC3 and FLIGHT, and 10 for SNIPS. The maximum norm for gradient clipping is set to 5, and we use weight decay regularization on all weights with  $\lambda=1e-6$  to avoid over-fitting.

3) *Network Input Embeddings*: We exploit pre-trained word embeddings to initialize the input embedding layers. For AICar, we train an LSTM based bidirectional language models (biLMs) at Chinese-character level by using zhwiki<sup>2</sup> corpus. The embedding dimension is 200. For DSTC 2&3 and SNIPS, we directly

adopt a pre-trained biLMs of ELMo [55] in English.<sup>3</sup> The word embedding dimension is 1024.

4) *Encoder of Slot Exemplars*: The pre-trained input embeddings and BiLMs are also exploited to compute label embeddings based on slot exemplars, as shown in Section V-C. In this paper, we choose to fix the encoder of slot exemplars to reduce training cost. Therefore, the performance of slot filling based on slot exemplars may heavily rely on qualities of the pre-trained BiLMs.

5) *Data Settings of Domain Adaptation*: To evaluate the efficiency of different methods in leveraging data, we set up a domain adaptation problem. In DSTC 2&3, DSTC2 is source domain, and DSTC3 is target domain. In AICar, we choose each domain as the target one and the rest two domains as source domains. Moreover, 1, 2.5, 5, 10, 20, 40, 60, 80 percent of the training data in the target domain are randomly selected to simulate low resource settings of the target domain. Each data selection is repeated 10 times by using different random seeds, and the final results of the evaluation are averaged. For SNIPS, we keep one intent as the target domain and the rest intents as the source domain. We pre-train the slot filling model on the source domains, then fine-tune and evaluate it on the target domain.

We randomly selected 80% of the training data for model training and the remaining 20% for validation. We keep the learning rate for 50 epochs and save the parameters that give the best performance on the validation set. Finally, we report the  $F_1$ -score of the semantic slots on the test set with parameters that have achieved the best  $F_1$ -score on the validation set. The  $F_1$ -score is calculated using CoNLL evaluation script.<sup>4</sup>

6) *Significance Test*: We use McNemar’s test to establish the statistical significance of a method over another ( $p < 0.05$ ).

## D. Main Results

We first compare different label embeddings by exploiting BLSTM-*softmax* models for slot filling, except for CT and ZAT. BLSTM-CRF models with the *scalar* and *biliner* transition modules are then investigated in Section VI-D4.

The results on DSTC 2&3 dataset are shown in Table IV. By comparing the column “DSTC3\_train” with “+DSTC2\_train,” we can see performances on the target domain are improved dramatically due to that the parameters are pre-trained on the source domain. Row 1-2 show that sharing the output weights of reused labels (see Equation (7)) can help transfer knowledge from DSTC2 to DSTC3. By comparing the baselines with our proposed methods, most of our methods can outperform all baselines on both single domain and domain adaptation tasks. Notably, by using label embeddings extracted from slot exemplars with the pre-trained biLMs of ELMo, the slot filling model can achieve the best results (78.58% and 87.55%).

Table V shows the performances of domain adaptation on AICar dataset with three domains. We take each domain as the target domain one by one and leave the rest two as the source domains. We evaluate the baselines and our methods with parts

<sup>3</sup>[Online]. Available: [https://github.com/allenai/allennlp/blob/master/tutorials/how\\_to/elmo.md](https://github.com/allenai/allennlp/blob/master/tutorials/how_to/elmo.md)

<sup>4</sup>[Online]. Available: <https://www.clips.uantwerpen.be/conll2000/chunking/output.html>

<sup>2</sup>[Online]. Available: <https://dumps.wikimedia.org/zhwiki/latest>

TABLE IV  
SLOT F<sub>1</sub> SCORES ON DSTC3 TEST SET WITH OR WITHOUT THE TRAINING DATA OF DSTC2. THE BOLD NUMBER INDICATES A STATISTICALLY SIGNIFICANT GAIN OVER THE BEST BASELINE

Method	DSTC3_train (single domain)	+DSTC2_train (domain adapt.)
One-hot	75.52	83.14
One-hot, w/o output layer sharing	75.52	81.22
CCA label embedding	76.25	84.10
value word embedding	75.32	83.97
CT	76.20	83.68
ZAT	75.76	84.29
atomic concept	76.41	<b>86.03</b>
slot description (BLSTM)	<b>77.41</b>	<b>87.16</b>
slot description (CNN)	76.23	<b>85.89</b>
slot description (emb mean)	76.56	84.05
slot exemplar (word embedding)	<b>77.19</b>	<b>86.10</b>
slot exemplar (BiLMs)	<b>78.58</b>	<b>87.55</b>

of the training data on the target domain, which range from 1 to 100 percent. We also test the pre-trained models even without the target training data (zero-shot). From Table V, we can find that:

- Two zero-shot baselines (CT and ZAT) perform better than the One-hot, which benefit from slot descriptions. However, these slot-independent models increase the computation cost sharply since they need to predict an IOB sequence for each slot independently. This could also lead to overlapped slot segmentations and lose the time series dependence of output slots.
- Even without the target training data (0%), our methods outperform all the baselines. Especially, atomic concepts can achieve the best performances. Because atomic concepts give an explicit interpretation of slots, which can recombine associated atoms into a new slot. While slot descriptions as word sequences are more complex and may contain synonyms implicitly. Without the target training data, there is no exemplar available for each slot.
- With the increasing training data of the target domain, slot F<sub>1</sub> scores of all methods are improved concurrently.
- At all data partitions, our proposed methods can achieve the best results and achieve statistically significant gain over the best baseline in most cases. The best performances always lie on atomic concepts and slot descriptions, since they involve more human knowledge than slot exemplars which are just the available training samples. Even so, the slot exemplar-based method can also outperform the baselines in most cases.

1) *Input Embeddings With Pre-Trained Language Model*: Besides using pre-trained word embeddings, some advanced pre-trained language models (e.g. ELMo [55], BERT [61]) can also be used to get input embeddings. However, it is orthogonal to the investigation of label embeddings. Table VI shows results on DSTC3 test set using limited seed instances, indicating that our method can get improvements over different types of input embeddings simultaneously. For BERT, *bert-base-cased* model<sup>5</sup> is utilized, while it performs worse than ELMo. The reason may

<sup>5</sup>[Online]. Available: <https://github.com/google-research/bert#pre-trained-models>

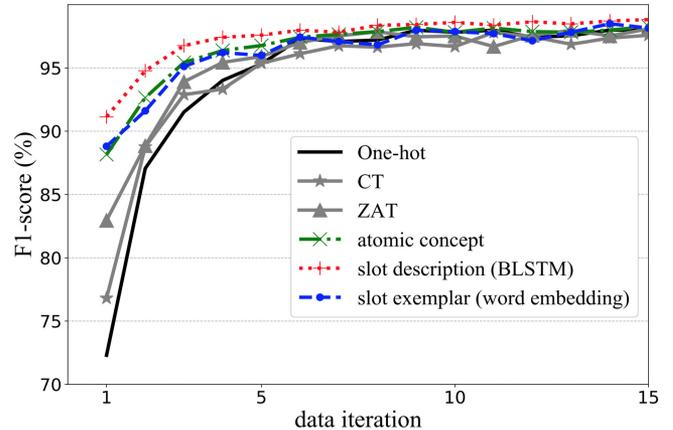


Fig. 5. Learning curves of the baselines and our methods when FLIGHT is the target domain and 100% training data is used for fine-tuning. It shows performances on the validation set at different data iterations.

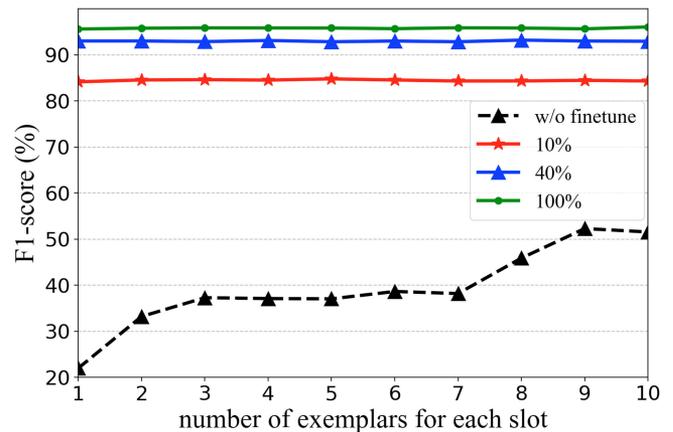


Fig. 6. Performances of the “slot exemplar (BiLMs)” method when FLIGHT is the target domain and {none, 10%, 40%, 100%} training data is used for fine-tuning. The number of exemplars for each slot ranges from 1 to 10.

be that ELMo could get better generalization for the dataset by using char-based word embeddings.

2) *Learning Curve*: Fig. 5 shows learning curves of different methods when the target domain is FLIGHT and 100% training data is used for fine-tuning. It shows slot F<sub>1</sub> scores on the validation set at increasing numbers of data iteration. We can see that our methods can converge much faster than the baselines without label embeddings. Especially, the “slot description (BLSTM)” method converges fastest.

3) *Number of Slot Exemplars*: Although we exploit all data samples of the training set as data exemplars for the corresponding slots, it is meaningful to investigate that how many exemplars for each slot are enough. Fig. 6 shows the number of slot exemplars used versus performance on FLIGHT domain. If there is no training data for fine-tuning, more exemplars usually give better performance. If there is an amount of labelled data for training, the performance is dominated by the size of training data in the target domain.

TABLE V

SLOT F<sub>1</sub> SCORES ON AICAR DATASET WITH THREE DOMAINS. WE TRAIN AND EVALUATE ON EACH TARGET DOMAIN AND PRE-TRAIN THE SLOT FILLING MODEL ON THE REST TWO SOURCE DOMAINS. COLUMNS REPRESENT DIFFERENT PROPORTIONS OF THE TRAINING DATA IN THE TARGET DOMAIN, DEFINED IN SECTION VI-C5. THE RESULTS IN BOLD BLACK ARE THE BEST F<sub>1</sub> SCORES. TIME PER BATCH REFERS TO TRAINING AND TESTING TIMES ON GeForce GTX 1080 Ti GPU

Target Domain	Method	Time per batch	0%	1%	2.5%	5%	10%	20%	40%	60%	80%	100%
Map	One-hot	9.2ms	0	40.81	60.24	72.67	82.63	87.94	90.99	92.92	93.61	94.79
	CT	94.2ms	0.22	40.93	62.23	75.20	82.79	87.58	91.28	92.88	94.02	94.80
	ZAT	118ms	0.05	45.07	63.74	73.79	82.10	87.39	91.33	92.40	93.65	94.62
	atomic concept	9.5ms	<b>1.55</b>	46.34	66.57	77.20	84.74	89.27	92.04	93.76	<b>94.81</b>	<b>95.78</b>
	slot description (BLSTM)	11.9ms	0.89	<b>46.55</b>	67.26	77.11	84.83	88.68	92.02	93.81	94.56	95.65
	slot description (CNN)	10.4ms	0.78	43.44	67.41	77.88	84.85	88.79	91.86	93.77	94.41	95.31
	slot description (emb mean)	9.9ms	0.32	45.20	<b>70.12</b>	<b>79.95</b>	<b>85.92</b>	<b>89.37</b>	<b>92.48</b>	<b>93.90</b>	94.67	95.47
	slot exemplar (word embedding)	9.2ms	0	42.54	64.34	76.29	84.16	88.45	91.61	93.40	94.10	94.71
	slot exemplar (BiLMs)	9.2ms	0	45.49	66.17	76.74	84.45	88.80	91.91	93.61	94.42	95.72
Weather	One-hot	7.5ms	0	32.91	39.35	49.23	62.40	72.96	80.57	84.13	86.50	88.60
	CT	140ms	4.97	39.47	45.46	52.26	62.67	73.39	81.42	84.85	86.96	87.85
	ZAT	187ms	15.73	39.78	47.03	54.10	65.28	73.60	80.60	84.19	86.75	88.61
	atomic concept	8.0ms	<b>24.69</b>	41.73	<b>48.38</b>	56.51	67.39	75.67	82.48	85.08	87.35	88.79
	slot description (BLSTM)	10.8ms	6.81	<b>42.14</b>	47.41	<b>56.67</b>	66.84	75.72	82.70	<b>85.89</b>	<b>88.23</b>	89.43
	slot description (CNN)	9.3ms	16.74	40.08	46.47	54.73	<b>67.41</b>	75.85	<b>82.74</b>	85.37	87.86	<b>89.44</b>
	slot description (emb mean)	8.4ms	15.75	41.72	48.32	55.71	66.54	<b>76.06</b>	82.47	85.56	87.90	89.42
	slot exemplar (word embedding)	7.7ms	0	41.41	46.59	55.09	65.73	75.28	82.44	85.31	87.56	88.72
	slot exemplar (BiLMs)	7.8ms	0	39.35	44.28	52.41	64.50	74.61	81.62	84.75	87.56	88.00
Flight	One-hot	6.2ms	0	55.69	65.84	73.22	80.68	87.58	92.26	94.14	94.92	95.26
	CT	69.5ms	2.07	59.31	71.01	75.36	81.40	87.74	92.32	93.84	94.97	95.74
	ZAT	99.2ms	1.89	65.91	70.44	74.80	82.03	87.81	92.37	93.93	94.77	95.34
	atomic concept	6.7ms	<b>52.37</b>	62.72	72.74	78.58	83.42	88.76	92.79	93.99	95.02	95.27
	slot description (BLSTM)	9.3ms	2.95	<b>70.49</b>	75.93	79.93	85.49	89.96	92.94	<b>94.99</b>	95.33	95.47
	slot description (CNN)	8.4ms	2.41	64.78	73.15	78.45	84.92	89.58	93.17	94.69	<b>95.41</b>	95.83
	slot description (emb mean)	7.0ms	8.45	69.97	<b>76.02</b>	<b>80.77</b>	<b>85.79</b>	<b>90.33</b>	<b>93.33</b>	94.60	95.29	95.76
	slot exemplar (word embedding)	6.4ms	0	68.49	73.73	79.13	83.82	88.93	92.56	94.18	94.95	95.30
	slot exemplar (BiLMs)	6.3ms	0	66.58	72.96	77.85	82.93	88.66	92.92	94.39	95.17	<b>95.98</b>

TABLE VI

SLOT F<sub>1</sub> SCORES ON DSTC3 TEST SET WITHOUT THE TRAINING DATA OF DSTC2, BUT WITH DIFFERENT PRE-TRAINED LANGUAGE MODELS

Method	pre. word emb.	ELMo	BERT
One-hot	75.52	78.13	76.64
slot description (BLSTM)	77.41	80.03	78.13

TABLE VII

SLOT F<sub>1</sub> SCORES OF DIFFERENT MODELS WITH AND WITHOUT CRF LAYERS ON AICAR DATASET, WHEN FLIGHT IS THE TARGET DOMAIN AND THE REST TWO ARE THE SOURCE DOMAINS. 10% TRAINING DATA OF FLIGHT IS USED FOR FINE-TUNING, AND RESULTS ARE AVERAGED OVER 10 RANDOM DATA SPLITS

Method	w/o CRF	with CRF	
		scalar	bilinear
One-hot	80.68	82.74	85.49 <sup>a</sup>
atomic concept	83.42	84.26	86.95
slot description (BLSTM)	85.49	<b>86.83</b>	<b>88.46</b>
slot description (CNN)	84.92	85.56	87.24
slot description (emb mean)	<b>85.79</b>	86.76	87.49
slot exemplar (word embedding)	83.82	85.61	87.10
slot exemplar (BiLMs)	82.93	84.45	86.83

<sup>a</sup>weights of the linear output layer act as label embeddings in the *bilinear* transition model.

4) *CRF Layer*: We also apply a CRF layer to model the time series dependence of output labels, which CT and ZAT cannot model. We want to know whether label embeddings can improve different slot filling models. Table VII shows the results of using different label embeddings with or without the CRF layer, in the case that FLIGHT is the target domain and 10% training data is used for fine-tuning. From the table, we can

TABLE VIII

PERFORMANCES OF THE “SLOT DESCRIPTION (BLSTM)” METHOD WITH VARIED SLOT DESCRIPTIONS ON FLIGHT. FLIGHT IS THE TARGET DOMAIN, AND ITS 10% TRAINING DATA IS USED FOR FINE-TUNING. “AL(DESC.)” DENOTES THE AVERAGE LENGTH OF SLOT DESCRIPTIONS

AL(desc.)	4.0 (ori.)	7.7	7.4	7.0	6.4	6.2
SLOT F <sub>1</sub>	85.49	85.84	84.66	85.80	86.43	85.08

see that CRF layers give improvements for the baseline and our methods. Compared to the traditional *scalar* transition model (as shown in Equation (4)), our proposed *bilinear* transition model (as shown in Section V-D) can improve significantly. The *bilinear* model calculates the transition score of two labels depending on their label embeddings, which shares transition patterns of related labels. It helps mitigate the data sparsity problem.

5) *Discussions About Varied Slot Descriptions*: Different from atomic concepts or slot exemplars which can be in a united form, slot descriptions in natural language would be varied due to different domain experts or developers. Five participants were asked to describe each slot in the FLIGHT domain. The average length of slot descriptions from different participants varies from 6.2 to 7.7. If we use these slot descriptions in the “slot description (BLSTM)” method when FLIGHT is target domain and 10% training data can be used for fine-tuning, slot F<sub>1</sub> scores range from 84.66% to 86.43%. As shown in Table VIII, our method seems robust to varied slot descriptions.

TABLE IX  
SLOT  $F_1$  SCORES ON SNIPS TEST SET IN THE SETTING OF DOMAIN  
ADAPTATION WITH ONLY 50 INSTANCES OF EACH TARGET INTENT FOR  
FINE-TUNING

Method	Average
One-hot	85.63
CT	83.79
ZAT	86.10
slot description (BLSTM)	86.94
slot description (CNN)	86.99
slot description (emb mean)	86.57
slot exemplar (word embedding)	85.67
slot exemplar (BiLMs)	<b>87.33</b>

6) *Results on SNIPS Dataset:* We also report results on the SNIPS public dataset where one intent are kept as the target domain and the rest intents serve as the source domain. Only 50 train instances of the target domain can be utilized for fine-tuning, which are chosen in random. All 7 intents play a role of the target domain one by one, and the slot  $F_1$  scores are averaged over seven target domains. As shown in Table IX, our methods can also outperform the baselines. Especially, the “slot exemplar (BiLMs)” method achieves the best slot  $F_1$  score.

## VII. CONCLUSION

This paper has proposed prior knowledge-driven label embedding for the NLU slot filling task to address the data sparsity problem (especially label space data sparsity). It requires human knowledge about slot definitions to consider slot relations into the label embeddings. We investigate three different kinds of prior knowledge about slots and the corresponding encoding methods: atomic concepts, slot descriptions and slot exemplars. These label embeddings are fused with different traditional slot filling models (e.g., BLSTM-softmax and BLSTM-CRF). Notably, we propose a novel label transition model based on label embeddings for CRF. The proposed approaches are evaluated on three datasets (DSTC 2&3, AICar and SNIPS) in the settings of single domain and especially domain adaptation. From the experimental results, we find that these label embeddings involving prior knowledge can significantly improve the performances and efficiencies over traditional slot filling models.

The proposed label embedding based slot filling framework shows promising perspectives of future improvements:

- The slot exemplar encoding relies on a pre-trained and high qualified BiLM which is fixed. In future work, we will develop a trainable encoder for slot exemplars.
- There are a few other kinds of prior knowledge, like the hierarchical dependency of *domain-intent-slot* [40]. Graph neural networks (GNN) have been applied to encode structured data [62], [63]. Applying GNN to encode structured knowledge is one of the future work.
- Nowadays, developing NLU systems for certain narrow domains independently is the mainstream approach. It costs much to collect data, while it is convenient for domain experts or developers to provide a few domain definitions (prior knowledge). Combination of deep learning-based slot filling and prior knowledge of narrow domains will be a promising research direction.

## REFERENCES

- [1] Y.-Y. Wang, L. Deng, and A. Acero, “Spoken language understanding—an introduction to the statistical framework,” *IEEE Signal Process. Mag.*, vol. 22, no. 5, pp. 16–31, Jan. 2005.
- [2] Y. Wang, L. Deng, and A. Acero, “Semantic frame-based spoken language understanding,” *Spoken Lang. Understanding: Syst. Extracting Semantic Inf. Speech*, New York, NY, USA: Wiley, 2011, ch. 3, pp. 35–80.
- [3] Y. He and S. Young, “A data-driven spoken language understanding system,” in *Proc. IEEE Workshop Autom. Speech Recognit. Understanding*, 2003, pp. 583–588.
- [4] C. Raymond and G. Riccardi, “Generative and discriminative algorithms for spoken language understanding,” in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2007, pp. 1605–1608.
- [5] G. Mesnil *et al.*, “Using recurrent neural networks for slot filling in spoken language understanding,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 23, no. 3, pp. 530–539, Mar. 2015.
- [6] B. Liu and I. Lane, “Attention-based recurrent neural network models for joint intent detection and slot filling,” in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2016, pp. 685–689.
- [7] K. Yao, B. Peng, Y. Zhang, D. Yu, G. Zweig, and Y. Shi, “Spoken language understanding using long short-term memory neural networks,” in *Proc. IEEE Spoken Lang. Technol. Workshop*, 2014, pp. 189–194.
- [8] N. T. Vu, “Sequential convolutional neural networks for slot filling in spoken language understanding,” in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2016, pp. 3250–3254.
- [9] G. Kurata, B. Xiang, B. Zhou, and M. Yu, “Leveraging sentence-level information with encoder lstm for semantic slot filling,” in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2016, pp. 2077–2083.
- [10] S. Zhu and K. Yu, “Encoder-decoder with focus-mechanism for sequence labelling based spoken language understanding,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2017, pp. 5675–5679.
- [11] C. Li, L. Li, and J. Qi, “A self-attentive model with gate mechanism for spoken language understanding,” in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2018, pp. 3824–3833.
- [12] L. S. Zettlemoyer and M. Collins, “Online learning of relaxed CCG grammars for parsing to logical form,” in *Proc. Conf. Empirical Methods Natural Lang. Process.-CoNLL*, 2007, pp. 678–687.
- [13] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, “Distributed representations of words and phrases and their compositionality,” in *Proc. Conf. Neural Inf. Process. Syst.*, 2013, pp. 3111–3119.
- [14] T. Mikolov, K. Chen, G. Corrado, and J. Dean, “Efficient estimation of word representations in vector space,” in *Proc. Int. Conf. Learn. Represent.*, 2013, *arXiv:1301.3781*.
- [15] S. Zhu and K. Yu, “Concept transfer learning for adaptive language understanding,” in *Proc. SIGDIAL*, 2018, pp. 391–399.
- [16] J. D. Lafferty, A. McCallum, and F. C. N. Pereira, “Conditional random fields: Probabilistic models for segmenting and labeling sequence data,” in *Proc. Int. Conf. Mach. Learn.*, 2001, pp. 282–289.
- [17] M. Henderson, B. Thomson, and J. Williams, “Dialog state tracking challenge 2&3,” 2013. [Online]. Available: <http://camdial.org/mh521/dstc/>
- [18] A. Coucke *et al.*, “Snips Voice Platform: An embedded Spoken Language Understanding system for private-by-design voice interfaces,” *CoRR*, 2018, *arXiv:1805.10190*.
- [19] G. Tur, D. Hakkani-Tür, and R. E. Schapire, “Combining active and semi-supervised learning for spoken language understanding,” *Speech Commun.*, vol. 45, no. 2, pp. 171–186, 2005.
- [20] A. Celikyilmaz, R. Sarikaya, D. Hakkani-Tür, X. Liu, N. Ramesh, and G. Tür, “A new pre-training method for training deep learning models with application to spoken language understanding,” in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2016, pp. 3255–3259.
- [21] O. Lan, S. Zhu, and K. Yu, “Semi-supervised training using adversarial multi-task learning for spoken language understanding,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2018, pp. 6049–6053.
- [22] S. Zhu, O. Lan, and K. Yu, “Robust spoken language understanding with unsupervised asr-error adaptation,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2018, pp. 6179–6183.
- [23] H. Daumé, III, “Frustratingly easy domain adaptation,” 2009, in *Proc. ACL*, 2007, pp. 256–263.
- [24] A. Jaech, L. Heck, and M. Ostendorf, “Domain adaptation of recurrent neural networks for natural language understanding,” in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2016, pp. 690–694.
- [25] Y.-B. Kim, K. Stratos, and R. Sarikaya, “Frustratingly easy neural domain adaptation” in *Proc. Int. Conf. Comput. Linguist.*, 2016, pp. 387–396.

- [26] B. Liu and I. Lane, "Multi-domain adversarial learning for slot filling in spoken language understanding," in *Proc. Neural Inf. Process. Syst. Workshop Conversational AI*, 2018, pp. 1–6.
- [27] Y.-B. Kim, K. Stratos, and D. Kim, "Domain attention with an ensemble of experts," in *Proc. Ann. Meeting Assoc. Comput. Linguist.*, 2017, pp. 643–653.
- [28] R. Jha, A. Marin, S. Shivaprasad, and I. Zitouni, "Bag of experts architectures for model reuse in conversational language understanding," in *Proc. North Amer. Chapter Assoc. Comput. Linguist.*, 2018, pp. 153–161.
- [29] A. K. Goyal, A. Metallinou, and S. Matsoukas, "Fast and scalable expansion of natural language understanding functionality for intelligent agents," in *Proc. North Amer. Chapter Assoc. Comput. Linguistics*, 2018, pp. 145–152.
- [30] H. Zhang, S. Zhu, S. Fan, and K. Yu, "Joint spoken language understanding and domain adaptive language modeling," in *Proc. ISIDE*, 2018, pp. 311–324.
- [31] Z. Zhao, S. Zhu, and K. Yu, "Data augmentation with atomic templates for spoken language understanding," in *Proc. Conf. Empirical Methods Natural Lang. Process.-Int'l Joint Conf. Natural Lang. Process. Workshops*, 2019, pp. 3628–3634.
- [32] E. Ferreira, B. Jabaian, and F. Lefèvre, "Zero-shot semantic parser for spoken language understanding," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2015, pp. 1403–1407.
- [33] M. Yazdani and J. Henderson, "A model of zero-shot learning of spoken language understanding" in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2015, pp. 244–249.
- [34] Y.-N. Chen, D. H.-Tür, and X. He, "Zero-shot learning of intent embeddings for expansion by convolutional deep structured semantic models," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2016, pp. 6045–6049.
- [35] A. Bapna, G. Tür, D. H.-Tür, and L. P. Heck, "Towards zero-shot frame semantic parsing for domain scaling," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2017, pp. 2476–2480.
- [36] S. Lee and R. Jha, "Zero-shot adaptive transfer for conversational language understanding," in *Proc. Conf. Artif. Intell.*, 2019, pp. 6642–6649.
- [37] D. J. Shah, R. Gupta, A. A. Fayazi, and D. H.-Tür, "Robust zero-shot cross-domain slot filling with example values," in *Proc. Annual Meeting Assoc. Comput. Linguist.*, 2019, pp. 5484–5490.
- [38] Y.-B. Kim, K. Stratos, R. Sarikaya, and M. Jeong, "New transfer learning techniques for disparate label sets," in *Proc. Annual Meeting Assoc. Comput. Linguistics-Int'l Joint Conf. Natural Lang. Process. Workshops*, 2015, pp. 473–482.
- [39] Y. Ma, E. Cambria, and S. Gao, "Label embedding for zero-shot fine-grained named entity typing," in *Proc. Int. Conf. Comput. Linguistics*, 2016, pp. 171–180.
- [40] J. Lee, D. Kim, R. Sarikaya, and Y.-B. Kim, "Coupled representation learning for domains, intents and slots in spoken language understanding," in *Proc. IEEE Spoken Lang. Technol. Workshop*, 2018, pp. 714–719.
- [41] J. Wu, L. F. D'Haro, N. F. Chen, P. Krishnaswamy, and R. E. Banchs, "Joint learning of word and label embeddings for sequence labelling in spoken language understanding," *ASRU*, pp. 800–806, 2019.
- [42] G. Mesnil, X. He, L. Deng, and Y. Bengio, "Investigation of recurrent-neural-network architectures and learning methods for spoken language understanding," in *Annu. Conf. Int. Speech Commun. Assoc.*, 2013, pp. 3771–3775.
- [43] K. Yao, G. Zweig, M.-Y. Hwang, Y. Shi, and D. Yu, "Recurrent neural networks for language understanding," in *Annu. Conf. Int. Speech Commun. Assoc.*, 2013, pp. 2524–2528.
- [44] N. T. Vu, P. Gupta, H. Adel, and H. Schütze, "Bi-directional recurrent neural network with ranking loss for spoken language understanding," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2016, pp. 6060–6064.
- [45] P. Xu and R. Sarikaya, "Convolutional neural network based triangular crf for joint intent detection and slot filling," in *Proc. IEEE IEEE Workshop Autom. Speech Recognit. Understanding*, 2013, pp. 78–83.
- [46] X. Zhang and H. Wang, "A joint model of intent determination and slot filling for spoken language understanding," in *Proc. Int. Joint Conf. Artif. Intell.*, 2016, pp. 2993–2999.
- [47] D. H.-Tür et al., "Multi-domain joint semantic frame parsing using bi-directional RNN-LSTM," in *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, 2016, pp. 715–719.
- [48] N. Reimers and I. Gurevych, "Optimal hyperparameters for deep lstm-networks for sequence labeling tasks," *CoRR*, 2017, *arXiv:1707.06799*.
- [49] F. Zhai, S. Potdar, B. Xiang, and B. Zhou, "Neural models for sequence chunking" in *Proc. Conf. Artif. Intell.*, 2017, pp. 3365–3371.
- [50] K. Yao, B. Peng, G. Zweig, D. Yu, X. Li, and F. Gao, "Recurrent conditional random field for language understanding," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2014, pp. 4077–4081.
- [51] Z. Huang, W. Xu, and K. Yu, "Bidirectional LSTM-CRF models for sequence tagging," *CoRR*, 2015, *arXiv:1508.01991*.
- [52] X. Ma and E. Hovy, "End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF," in *Proc. Ann. Meeting Assoc. Comput. Linguist.*, 2016, pp. 1064–1074.
- [53] Y. Kim, "Convolutional neural networks for sentence classification," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2014, pp. 1746–1751.
- [54] R. Collobert, J. Weston, L. Bottou, M. Karlen, K. Kavukcuoglu, and P. Kuksa, "Natural language processing (almost) from scratch," *J. Mach. Learn. Res.*, vol. 12, pp. 2493–2537, 2011.
- [55] M. Peters et al., "Deep contextualized word representations," in *Proc. North Amer. Chapter Assoc. Comput. Linguist.*, 2018, pp. 2227–2237.
- [56] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *Proc. Int. Conf. Learn. Represent.*, 2015, *arXiv:1409.0473*.
- [57] M.-T. Luong, H. Pham, and C. D. Manning, "Effective approaches to attention-based neural machine translation," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2015, pp. 1412–1421.
- [58] A. Vaswani et al., "Attention is all you need," in *Proc. Neural Inf. Process. Syst.*, 2017, pp. 5998–6008.
- [59] Y. Meng, X. Li, X. Sun, Q. Han, A. Yuan, and J. Li, "Is word segmentation necessary for deep learning of Chinese representations?" in *Proc. Annu. Meeting Assoc. Comput. Linguist.*, 2019, pp. 3242–3252.
- [60] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *Proc. Int. Conf. Learn. Represent.*, 2015, *arXiv:1412.6980*.
- [61] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. North Amer. Chapter Assoc. Comput. Linguist.*, 2019, pp. 4171–4186.
- [62] L. Chen, B. Tan, S. Long, and K. Yu, "Structured dialogue policy with graph neural networks," in *Proc. Int. Conf. Comput. Linguist.*, 2018, pp. 1257–1268.
- [63] L. Chen, Z. Chen, B. Tan, S. Long, M. Gašić, and K. Yu, "AgentGraph: Toward universal dialogue management with structured deep reinforcement learning," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 27, no. 9, pp. 1378–1391, May 2019.



**Su Zhu** (Student Member, IEEE) received the B.Eng. degree in computer science from Xi'an Jiao Tong University, China in 2013, and the M.Sc. degree from the Department of Computer Science, Shanghai Jiao Tong University, Shanghai, China, in 2016. He is currently working toward the Ph.D. degree with the SpeechLab, Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China. His research interests include spoken/natural language understanding, dialogue systems, and structured deep learning.



**Zijian Zhao** received the B.Eng. degree in computer science from Shanghai Jiao Tong University, Shanghai, China, in 2017. He is currently working toward the M.S. degree with the SpeechLab, Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China. His research interests include dialogue system, spoken language understanding, and machine learning.



**Rao Ma** received the B.Eng. degree in computer science and technology from Nanjing University, China in 2018. She is currently working toward the M.S. degree with the SpeechLab, Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China. Her research interests include natural language processing, language model and large vocabulary continuous speech recognition.



**Kai Yu** (Senior Member, IEEE) is a professor at Computer Science and Engineering Department, Shanghai Jiao Tong University, China. He received his B.Eng. and M.Sc. from Tsinghua University, China in 1999 and 2002, respectively. He is a Professor with Computer Science and Engineering Department, Shanghai Jiao Tong University, China. He then joined the Machine Intelligence Lab at the Engineering Department at Cambridge University, U.K., where he received the Ph.D. degree in 2006. His main research interests lie in the area of speech-based human machine interaction including speech recognition, synthesis, language understanding and dialogue management. He is a member of the IEEE Speech and Language Processing Technical Committee.