

Dialogue State Distillation Network with Inter-slot Contrastive Learning for Dialogue State Tracking

Jing Xu^{1*}, Dandan Song^{1†}, Chong Liu², Siu Cheung Hui³,
Fei Li², Qiang Ju², Xiaonan He², Jian Xie^{2†}

¹School of Computer Science & Technology, Beijing Institute of Technology, Beijing, China

²Baidu Inc., Beijing, China

³Nanyang Technological University, Singapore

{xujing, sdd}@bit.edu.cn, {liuchong06, lifei21, juqiang, hexiaonan, xiejian01}@baidu.com, asschui@ntu.edu.sg

Abstract

In task-oriented dialogue systems, Dialogue State Tracking (DST) aims to extract users’ intentions from the dialogue history. Currently, most existing approaches suffer from error propagation and are unable to dynamically select relevant information when utilizing previous dialogue states. Moreover, the relations between the updates of different slots provide vital clues for DST. However, the existing approaches rely only on predefined graphs to indirectly capture the relations. In this paper, we propose a Dialogue State Distillation Network (DSDN) to utilize relevant information of previous dialogue states and migrate the gap of utilization between training and testing. Thus, it can dynamically exploit previous dialogue states and avoid introducing error propagation simultaneously. Further, we propose an inter-slot contrastive learning loss to effectively capture the slot co-update relations from dialogue context. Experiments are conducted on the widely used MultiWOZ 2.0 and MultiWOZ 2.1 datasets. The experimental results show that our proposed model achieves the state-of-the-art performance for DST.

Introduction

As a key component of task-oriented dialogue systems, Dialogue State Tracking (DST) aims to extract users’ goals or intents over the continuation of multiple turns of a dialogue, which are described as a compact dialogue state. Specifically, a DST model aims to extract the corresponding values of pre-defined slots given a dialogue history. Thus, the extracted dialogue state can be represented as a set of (slot, value) pairs. A DST example is shown in Figure 1.

As previous dialogue states can provide effective and explicit information of a dialogue history (Kim et al. 2020), recent approaches attempt to utilize the previous dialogue state information for DST. Most of these approaches (Zhu et al. 2020; Zeng and Nie 2020a,b) concatenate the dialogue context and the corresponding previous dialogue state as input. However, the golden dialogue states can only be obtained during training. Therefore, utilizing the predicted dialogue states may lead to error propagation during the testing

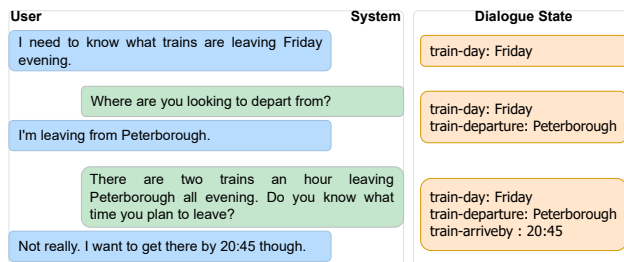


Figure 1: An example of DST. Utterances on the left side are from the user and system over the continuation of the multiple turns of a dialogue. The corresponding dialogue state for each turn is on the right side. Specifically, “train-day: Friday” represents that the value of the slot “day” in the domain “train” is “Friday”.

process. To alleviate this problem, gated recurrent units are used in a graph attention network to learn to decide which dialogue states are kept (Chen et al. 2020a). In addition, a uniform scheduled sampling strategy is also used to make the DST model adapt to the data noise of dialogue states, thus improving the robustness of the model (Zhou et al. 2021). Unfortunately, these attempts are still unable to tackle the problems of the unavailable golden dialogue states and the potential error of predicted dialogue states, though they strive to reduce the negative impact of wrongly predicted dialogue states during the testing process. Moreover, for different current dialogue contexts, the relevant information of previous dialogue states should be dynamically obtained in contrast to missing in the current approaches. Therefore, it is a challenging problem on how to dynamically incorporate previous dialogue states without error propagation for DST.

Further, we notice that State Operation Prediction (SOP), which aims to predict whether the value of a given slot needs to be updated compared with the previous turn, is an important auxiliary task for DST (Zhu et al. 2020; Zhang et al. 2020a; Zhou et al. 2021). Most existing approaches conduct the prediction of each slot independently. However, slot co-update relations provide a vital clue for the prediction. For example, if users provide their check-in time, they will also tend to provide their hotel stay days when booking a hotel.

*Work done during internship at Baidu Inc.

†Corresponding authors.

Thus, the values of the slots “book day” and “book stay” are usually updated at the same time in the domain “hotel”. Although several recent approaches incorporate the slot relations across domains to indirectly capture the slot co-update relations for DST (Chen et al. 2020a; Zeng and Nie 2020b; Zhu et al. 2020; Feng et al. 2022), they rely on predefined schema graphs or designed graphs and fail to capture the relations from dialogue context automatically. Recent works (Ye et al. 2021; Wang et al. 2022b) also attempt to use self-attention to capture slot dependency. However, they mainly provide mutual guidance among slots in an implicit manner. Therefore, how to explicitly capture the slot co-update relations without predefined graphs is another important issue for DST.

In this paper, we propose an effective model equipped with a Dialogue State Distillation Network (DSDN) to dynamically exploit previous dialogue states without introducing error propagation. In DSDN, a combination of turn-level and dialogue-level slot relevant information is first encoded by an encoder. Then, a teacher network with golden previous dialogue states tackles SOP while a student network with slot only concatenation learns to imitate the behavior of the teacher network. As a result, the student network can utilize the previous dialogue states without extra input. Moreover, the rich interactions between the encoder and the teacher/student network are conducted to push the teacher/student network to dynamically decide the utilization of previous dialogue states. Finally, the well-learned student network can help our model tackle DST more effectively.

Additionally, an inter-slot contrastive learning loss is also proposed to explicitly capture the slot co-update relations. In details, we construct the contrastive learning loss between slot representations from both turn-level and dialogue-level according to the labels of SOP. In this way, the model can directly capture the label correlations of SOP between slots automatically and then help improve DST performance. In summary, the main contributions of this paper are as follows:

- We propose an effective Dialogue State Distillation Network to dynamically exploit previous dialogue states without introducing error propagation. To the best of our knowledge, this study is the first to exploit knowledge distillation for DST.
- We propose an inter-slot contrastive learning loss to capture the slot co-update relations from dialogue context automatically.
- Extensive experiments are conducted on the widely used MultiWOZ 2.0 and MultiWOZ 2.1 datasets. Our proposed model achieves the state-of-the-art performance on MultiWOZ 2.1 and competitive performance on MultiWOZ 2.0.

Related Work

Dialogue State Tracking In task-oriented dialogue systems, as natural language understanding (Wang et al. 2022a; Xing and Tsang 2022) and natural language generation (Mi, Wang, and Li 2022) are standalone and less interactive with each other, recent works focus more on dialogue management such as dialogue state tracking (Zhang et al.

2020c). Existing approaches can be divided into two categories: ontology-based and ontology-free. The ontology-based approaches (Zhong, Xiong, and Socher 2018; Ramadan, Budzianowski, and Gasic 2018; Ye et al. 2021) require an ontology including all possible predefined values for each slot and are simplified into a multi-class classification task. Conversely, the ontology-free approaches (Xu and Hu 2018; Hosseini-Asl et al. 2020; Gao et al. 2019; Feng et al. 2022) generate the values of slots from dialogue context or vocabularies. In this paper, our work is related to the ontology-based approaches.

Recently, the utilization of previous dialogue states has attracted much attention for DST. Several works (Kim et al. 2020; Zhu et al. 2020; Zeng and Nie 2020a,b) use the previous dialogue states as part of the input. Once a slot in the dialogue states is wrongly predicted, the error information will continue propagating until the end of the dialogue. To improve the utilization effectiveness, the previous dialogue state information is controlled by a recurrent graph attention network (Chen et al. 2020a). A multi-level fusion gate and a uniform scheduled sampling strategy are also used to reduce the negative impact of wrong predictions (Zhou et al. 2021).

As for slot co-update relations, most existing approaches (Zhu et al. 2020; Chen et al. 2020a; Wu et al. 2020; Ouyang et al. 2020; Lin, Tseng, and Byrne 2021; Feng et al. 2022) use graph neural networks based on predefined schema graphs to exploit the prior knowledge of domain or slot dependency, thus capturing them indirectly. Among these models, the co-update relations between slots are only explicitly defined by Feng et al. (2022). They fuse schema graphs and dialogue context to obtain diverse dynamic slot relations which include the slot co-update relations. The obtained slot relations further enhance the final predictions.

Knowledge Distillation Knowledge distillation (Hinton, Vinyals, and Dean 2015) aims to distill knowledge from teacher networks and transfer it to student networks. Traditionally, it pushes the student networks to imitate the feature representations or probability distributions of the teacher networks. Recently, knowledge distillation has been widely used in natural language processing (Sun et al. 2020; Zhang et al. 2020b; Wei et al. 2021; Zhou, Xu, and McAuley 2022). In this paper, we propose a novel knowledge distillation network to migrate the gap of utilizing previous dialogue states between the training and testing processes.

Contrastive Learning Contrastive learning is first used in computer vision with self-supervised learning (Chen et al. 2020b), which increases the similarity between the representations of original data samples and differently augmented data samples. Further, supervised contrastive learning (Khosla et al. 2020) extends it and generates positive samples from not only data augmentation but also training instances with the same label. In natural language processing, contrastive learning has been widely used in diverse fields, such as pre-trained language model and information extraction (Wang et al. 2021a,b; Das et al. 2022; Saha, Yadav, and Bansal 2022). As for DST, we are the first to introduce contrastive learning to exploit the slot co-update relations from the whole dialogue automatically.

Proposed Model

For DST, we use $X = \{(U_1, R_1), \dots, (U_T, R_T)\}$ to represent a dialogue with T turns and $S = \{s_1, \dots, s_J\}$ to represent a pre-defined slot set, where U_t is the user utterance at the t -th dialogue turn, R_t denotes the corresponding system response and J denotes the total number of pre-defined slots. For each turn t , given $X = \{(U_1, R_1), \dots, (U_t, R_t)\} (1 \leq t \leq T)$, the goal of DST is to extract the dialogue state which can be represented as $\mathcal{B} = \{(s_1, v_1^t), \dots, (s_J, v_J^t)\}$, where v_j^t denotes the value of s_j for the t -th turn. Note that if a slot is not mentioned, its value will be annotated as “none”. As for any slot, we follow (Wang et al. 2022b) to use a concatenation of its domain name and slot name to represent it, thus integrating both domain information and slot information.

In this section, we propose a novel model to tackle DST. Figure 2 shows the proposed model which consists of four modules: Encoder, Dialogue State Distillation, Inter-slot Contrastive Learning and Decoder. Except the Inter-slot Contrastive Learning module, the other three modules constitute the Dialogue State Distillation Network. Next, we will introduce each module in details.

Encoder

Based on previous works (Lee, Lee, and Kim 2019; Ye et al. 2021; Wang et al. 2022b), we aim to combine turn-level and dialogue-level slot relevant information in this module. First, we use a pre-trained BERT (Kenton and Toutanova 2019) to encode dialogue context, slots and slot values respectively. Then, we retrieve slot relevant context from dialogue at turn-level and dialogue-level. Finally, a max-pooling operation is used to fuse the retrieved information.

Base Encoder For the dialogue context at the t -th turn, we concatenate the user utterance and the system utterance, and encode them into hidden representations:

$$\mathbf{R}_t = \text{BERT}_{\text{finetune}}([\text{CLS}] \oplus U_t \oplus [\text{SEP}] \oplus R_t \oplus [\text{SEP}]) \quad (1)$$

where $[\text{CLS}]$ and $[\text{SEP}]$ denote special tokens in BERT, and $\text{BERT}_{\text{finetune}}$ denotes that the used BERT will be finetuned during the training process. For any given slot $s_j \in S$ or the corresponding value v_j^t at the given t -th turn, we use another BERT to encode it, whose parameters are fixed during the training process. Unlike encoding dialogue context, we use the hidden representation of the special token $[\text{CLS}]$ to represent s_j or v_j^t :

$$\mathbf{h}^{S_j} = \text{BERT}_{\text{fixed}}([\text{CLS}] \oplus s_j \oplus [\text{SEP}]) \quad (2)$$

$$\mathbf{h}^{V_j^t} = \text{BERT}_{\text{fixed}}([\text{CLS}] \oplus v_j^t \oplus [\text{SEP}]) \quad (3)$$

Multi-level Slot Attention Based on the hidden representations of dialogue context and slots, we further utilize turn-level and dialogue-level slot attention to obtain slot relevant information. A multi-head attention mechanism (Vaswani et al. 2017) is first used for the turn-level slot attention and we use $\text{MultiHead}(\mathbf{Q}, \mathbf{K}, \mathbf{V})$ to represent it, where \mathbf{Q} denotes a query matrix or vector, \mathbf{K} and \mathbf{V} denote a key matrix and a value matrix respectively. Thus, the turn-level slot attention is denoted as follows:

$$\mathbf{r}_t^{S_j} = \text{MultiHead}(\mathbf{h}^{S_j}, \mathbf{R}_t, \mathbf{R}_t) \quad (4)$$

As the hidden representations of current dialogue context at the t -th turn encoded by Base Encoder lack of previous dialogue context information, we use a transformer encoder (Vaswani et al. 2017) to aggregate contextual information before dialogue-level slot attention. It is composed of N stacked layers and each layer consists of two sub-layers including a multi-head attention mechanism and a position-wise fully connected feed-forward network. We use the output of the last layer as the contextual representations of each slot, which is denoted as follows:

$$\mathbf{D}_t^N = \text{TransformerEncoder}(\mathbf{r}_1^{S_j}, \dots, \mathbf{r}_t^{S_j}) \quad (5)$$

After that, the dialogue-level slot attention is applied to the contextual representations \mathbf{D}_t^N , thus summarizing the slot relevant information from the up-to-now utterances for the slot s_j . Similarly, we use another multi-head attention mechanism as follows:

$$\mathbf{d}_t^{S_j} = \text{MultiHead}(\mathbf{h}^{S_j}, \mathbf{D}_t^N, \mathbf{D}_t^N) \quad (6)$$

Fusion To fuse the turn-level and dialogue-level information of slot s_j , we employ a max-pooling operation to obtain the fusion information $\mathbf{f}_t^{S_j}$ as follows:

$$[\mathbf{f}_t^{S_j}]_k = \max \left\{ [\mathbf{r}_t^{S_j}]_k, [\mathbf{d}_t^{S_j}]_k \right\} \quad (7)$$

where $[\cdot]_k$ represents the k -th value of a vector.

Dialogue State Distillation

In this module, we aim to learn to distill previous dialogue states from dialogue context without error propagation. To achieve this goal, we first use two BERTs (called the teacher network and the student network respectively) to encode different inputs: golden previous dialogue states and slot only concatenation. Then, we use a multi-head attention mechanism to retrieve relevant information of dialogue states or slots. After that, we train the teacher network to tackle SOP and drive the student network to imitate it simultaneously. Thus, the student network can dynamically utilize relevant information of previous dialogue states without the gap between training and testing.

Base Encoder For the teacher network, we concatenate all golden pairs $(s_j, v_j^{t-1}, 1 \leq j \leq J)$ of the $(t-1)$ -th turn as the input for the t -th turn and use a BERT to encode it:

$$\text{SV}_j^t = [\text{SLOT}_j^{\text{tea}}] \oplus s_j \oplus - \oplus v_j^{t-1}, 1 \leq j \leq J \quad (8)$$

$\mathbf{H}^{\text{tea}} = \text{BERT}_{\text{finetune}}([\text{CLS}] \oplus \text{SV}_1^t, \dots, \text{SV}_J^t \oplus [\text{SEP}])$ (9) where $[\text{SLOT}_j^{\text{tea}}] (1 \leq j \leq J)$ represents special tokens and “-” is a separation symbol. Based on \mathbf{H}^{tea} , we use the representation at the position of $[\text{SLOT}_j^{\text{tea}}] (1 \leq j \leq J)$ to represent (s_j, v_j^{t-1}) , denoted as $\mathbf{h}_{j,t}^{\text{tea}}$.

For the student network, we use slot only concatenation as the input at any turn and use another BERT to encode it:

$$\mathbf{S}_j = [\text{SLOT}_j^{\text{stu}}] \oplus s_j, 1 \leq j \leq J \quad (10)$$

$\mathbf{H}^{\text{stu}} = \text{BERT}_{\text{finetune}}([\text{CLS}] \oplus \mathbf{S}_1, \dots, \mathbf{S}_J \oplus [\text{SEP}])$ (11) where $[\text{SLOT}_j^{\text{stu}}] (1 \leq j \leq J)$ denotes special tokens. Then, we use the representation at the position of $[\text{SLOT}_j^{\text{stu}}] (1 \leq j \leq J)$ to represent s_j based on \mathbf{H}^{stu} , denoted as $\mathbf{h}_j^{\text{stu}}$.

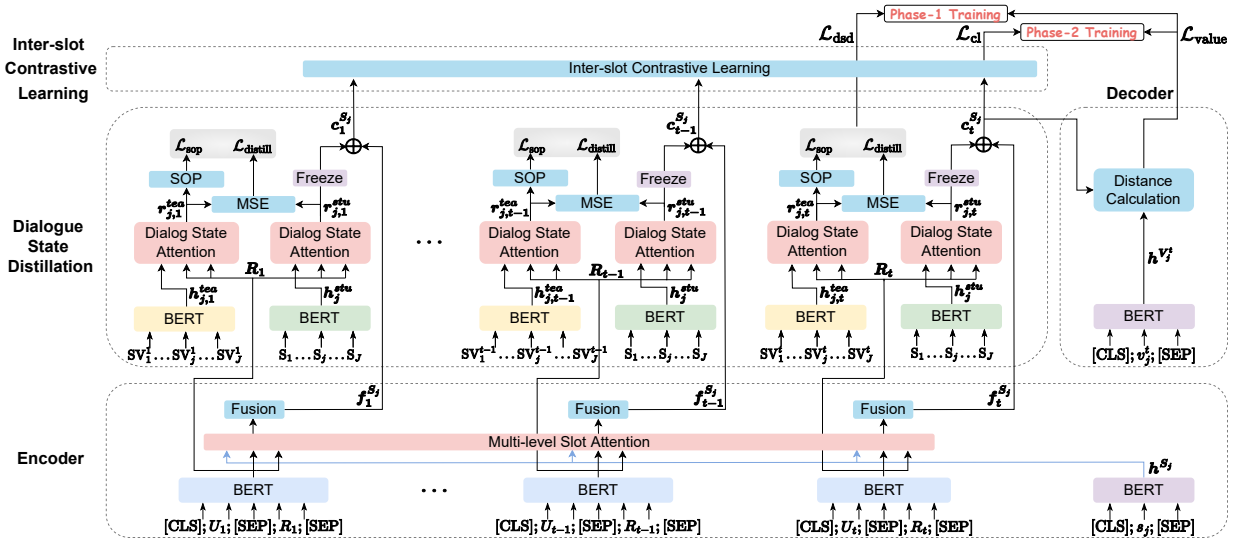


Figure 2: The architecture of our proposed model.

Dialogue State Attention For the given dialogue context at the t -th turn and the slot s_j , we use $h_{j,t}^{tea}$ and $h_{j,t}^{stu}$ to retrieve the current dialogue context respectively by two different multi-head attention mechanisms as follows:

$$r_{j,t}^{tea} = \text{MultiHead}(h_{j,t}^{tea}, R_t, R_t) \quad (12)$$

$$r_{j,t}^{stu} = \text{MultiHead}(h_{j,t}^{stu}, R_t, R_t) \quad (13)$$

Distillation To push the teacher network to hold the slot information of the t -th turn, we train it to tackle State Operation Prediction (SOP). SOP aims to identify whether a slot updates its value compared with previous turns. Different from (Zhang et al. 2020a; Zeng and Nie 2020a), we follow (Zhou et al. 2021) to simplify it into a binary classification task. In SOP, we use “1” to denote that the slot updates its value compared with previous turns, while “0” is used if otherwise. We obtain its corresponding loss \mathcal{L}_{sop} as follows:

$$p_{j,t}^{sop} = \sigma(\mathbf{W}^1 \tanh(\mathbf{W}^2 r_{j,t}^{tea})) \quad (14)$$

$$\mathcal{L}_{sop} = \frac{1}{T \cdot J} \sum_{t=1}^T \sum_{j=1}^J -(y_{j,t}^{sop} \cdot \log p_{j,t}^{sop} + (1 - y_{j,t}^{sop}) \cdot \log(1 - p_{j,t}^{sop})) \quad (15)$$

where σ denotes a sigmoid activation function, $y_{j,t}^{sop}$ denotes the golden label of slot s_j for SOP, $\mathbf{W}^1 \in \mathbb{R}^{1 \times d_{out}}$ and $\mathbf{W}^2 \in \mathbb{R}^{d_{out} \times d_{out}}$ are trainable matrices. Note that d_{out} is the dimension of $r_{j,t}^{tea}$.

Meanwhile, we use a mean square error loss (MSE) to calculate the similarity of the learned representations between the teacher network and the student network for each slot. Thus, the student network can learn to distill previous dialogue state information from the teacher network by minimizing the loss. The distillation loss $\mathcal{L}_{distill}$ is obtained as follows:

$$\mathcal{L}_{distill} = \frac{1}{T \cdot J} \sum_{t=1}^T \sum_{j=1}^J \text{MSE}(r_{j,t}^{tea}, r_{j,t}^{stu}) \quad (16)$$

After that, we formulate the overall loss \mathcal{L}_{dsd} in the module as follows:

$$\mathcal{L}_{dsd} = \alpha \mathcal{L}_{sop} + (1 - \alpha) \mathcal{L}_{distill} \quad (17)$$

where α ($0 < \alpha < 1$) is a balancing coefficient.

Inter-slot Contrastive Learning

In this module, we construct contrastive learning between slots to capture the slot co-update relations automatically. As our goal is to capture the co-update relations of slots, we only consider the slots whose labels of SOP are “1” in our contrastive learning. Assume that the SOP label of slot s_j at the t -th turn (denoted as s_j^t) is “1” (i.e., $y_{j,t}^{sop} = 1$), we introduce the strategy of positive sample selection and negative sample selection, and the calculation of contrastive learning loss as follows:

Positive Sample Selection For the slot s_j at the t -th turn, we select the slots whose labels of SOP are also “1” at the same turn as positive samples and the selected set $\mathcal{P}(s_j^t)$ is shown as follows:

$$\mathcal{P}(s_j^t) = \{s_p^t : (y_{p,t}^{sop} = 1) \wedge (p \neq j)\} \quad (18)$$

As contrastive learning pulls representations of the given instance s_j^t and the positive samples together, the similarity between the learned representations of slots which update in the same turn will be increased, thus helping the model capture the co-updated relations between different slots.

Negative Sample Selection For the slot s_j at the t -th turn, we select negative samples from both turn-level and dialogue-level. For turn-level, we select the slots whose labels of SOP are “0” in the same turn as negative samples. And for dialogue-level, we select the slot s_j at the other turns as negative samples. The selected contrastive learning set $\mathcal{I}(s_j^t)$ is shown as follows:

$$\mathcal{N}_{turn}(s_j^t) = \{s_n^t : (y_{n,t}^{sop} = 0) \wedge (n \neq j)\} \quad (19)$$

$$\mathcal{N}_{dialogue}(s_j^t) = \{s_j^{t'} : t' \neq t\} \quad (20)$$

$$\mathcal{I}(s_j^t) = \mathcal{P}(s_j^t) \cup \mathcal{N}_{turn}(s_j^t) \cup \mathcal{N}_{dialogue}(s_j^t) \quad (21)$$

where $\mathcal{N}_{turn}(s_j^t)$ and $\mathcal{N}_{dialogue}(s_j^t)$ denote the turn-level and dialogue-level selected sets respectively. Contrastive learning pushes representations of the given instance s_j^t and the negative samples in $\mathcal{N}_{turn}(s_j^t)$ away. Thus, the similarity of the learned representations between the given instance s_j^t and the negative samples in $\mathcal{N}_{turn}(s_j^t)$ will be reduced. It helps capture the co-updated relations between slots at the same turn. On the other hand, we reduce the similarity of the learned representations between the same slot at different dialogue turns. As a result, distinguishing these subtle representation differences for the same slot can help our model perceive the state changes of each slot better as the dialogue continues.

Contrastive Learning Loss We first integrate the output of Encoder and Dialogue State Distillation, thus utilizing the previous dialogue state information:

$$\mathbf{c}_t^{S_j} = \mathbf{f}_t^{S_j} + \mathbf{r}_{j,t}^{stuu} \quad (22)$$

Then, we follow (Chen et al. 2020b; Khosla et al. 2020) to map $\mathbf{c}_t^{S_j}$ into a new space, where contrastive learning is used to improve the quality of learning, as follows:

$$\mathbf{z}_t^{S_j} = \mathbf{W}^3 \text{ReLU}(\mathbf{W}^4 \mathbf{c}_t^{S_j}) \quad (23)$$

where $\mathbf{W}^3 \in \mathbb{R}^{d_1 \times d_2}$ and $\mathbf{W}^4 \in \mathbb{R}^{d_2 \times d_{out}}$ are trainable matrices. Note that d_1 and d_2 are the dimensional parameters. After that, we obtain the contrastive learning loss \mathcal{L}_{cl} based on Normalized Temperature-scaled Cross Entropy (NT-Xent) (Chen et al. 2020b) as follows:

$$\mathcal{L}_{cl}^{j,t} = \sum_{s_p^t \in \mathcal{P}(s_j^t)} \log \frac{\exp(\text{sim}(\mathbf{z}_t^{S_j}, \mathbf{z}_t^{S_p}) / \tau)}{\sum_{s_n^t \in \mathcal{I}(s_j^t) \setminus s_j^t} \exp(\text{sim}(\mathbf{z}_t^{S_j}, \mathbf{z}_t^{S_n}) / \tau)} \quad (24)$$

$$\mathcal{L}_{cl} = \frac{1}{T \cdot J} \sum_{t=1}^T \sum_{j=1}^J \frac{1}{|\mathcal{P}(s_j^t)|} \mathcal{L}_{cl}^{j,t} \quad (25)$$

where $\text{sim}(\mathbf{z}_t^{S_j}, \mathbf{z}_t^{S_p}) = \mathbf{z}_t^{S_j \top} \mathbf{z}_t^{S_p} / \|\mathbf{z}_t^{S_j}\|_2 \|\mathbf{z}_t^{S_p}\|_2$ and τ is a temperature parameter.

Decoder

In this module, we aim to predict the value for each slot. Given the slot s_j at the t -th turn, the integrated representation $\mathbf{c}_t^{S_j}$ is first fed into a linear layer with layer normalization as follows:

$$\mathbf{o}_t^{S_j} = \text{LayerNorm}(\mathbf{W}^5 \mathbf{c}_t^{S_j}) \quad (26)$$

where $\mathbf{W}^5 \in \mathbb{R}^{d_{out} \times d_{out}}$ is a trainable matrix.

After that, following (Ren et al. 2018; Wang et al. 2022b), we use L2 norm to obtain the distances between $\mathbf{o}_t^{S_j}$ and candidate values of s_j . We obtain the loss \mathcal{L}_{value} as follows:

$$p_{j,t}^{value} = \frac{\exp(-\|\mathbf{o}_t^{S_j} - \mathbf{h}^{V_j^t}\|_2)}{\sum_{v \in \mathcal{V}(s_j)} \exp(-\|\mathbf{o}_t^{S_j} - \mathbf{h}^v\|_2)} \quad (27)$$

$$\mathcal{L}_{value} = \frac{1}{T \cdot J} \sum_{t=1}^T \sum_{j=1}^J -\log p_{j,t}^{value} \quad (28)$$

where $\mathbf{h}^{V_j^t}$ is the representation of golden value for s_j and $\mathcal{V}(s_j)$ denotes the candidate value set of s_j .

Optimization

To push each module to focus on its optimization objective, we employ a two-phase optimization strategy. First, we optimize the joint objective of \mathcal{L}_{phase1} as follows:

$$\mathcal{L}_{phase1} = \mathcal{L}_{value} + \mathcal{L}_{dsd} \quad (29)$$

After that, we fix the parameters of the Dialogue State Distillation module based on development datasets and then optimize the joint objective of \mathcal{L}_{phase2} as follows:

$$\mathcal{L}_{phase2} = \mathcal{L}_{value} + \mathcal{L}_{cl} \quad (30)$$

Performance Evaluation

Experimental Setup

Dataset We evaluate our proposed model on the widely used MultiWOZ 2.0 (Budzianowski et al. 2018) and MultiWOZ 2.1 (Eric et al. 2020) datasets. MultiWOZ 2.0 includes over 10,000 dialogues and 35 slots across 7 domains. MultiWOZ 2.1 is the modified version of MultiWOZ 2.0 and corrects previous annotation errors of MultiWOZ 2.0.

Metrics Following previous works on the two datasets, we use Joint Goal Accuracy (Joint GA) as our evaluation metrics. It is calculated as the proportion of dialogue turns for which the value of each slot is correctly predicted.

Training Details For all modules, we follow (Wang et al. 2022b; Feng et al. 2022) to use the pre-trained BERT-base-uncased model for Base Encoder. The number of heads is set to 4 for the multi-head attention mechanism. In Encoder, the number of stacked layers in the transformer encoder is set to 6. In Dialogue State Distillation, we set the balancing coefficient α to 0.8 and 0.6 for MultiWOZ 2.0 and MultiWOZ 2.1 respectively. In Inter-slot Contrastive Learning, we set d_1 to 512, d_2 to 512 and the temperature parameter τ to 0.01. In the two-phase optimization strategy, the epoch of the first phase training is set to 100 and it stops early when the validation loss is not decreased for 15 consecutive epochs. As for the second phase training, we set the epoch to 15. Moreover, we set the batch size to 8 and 16 on four NVIDIA Tesla V100 GPUs during the first phase and second phase training respectively. Adam is used as the optimizer with learning rates of 1×10^{-4} and 1×10^{-5} during the first phase and second phase training respectively, and the warmup proportion is set to 0.1 during all training processes.

Baselines We use the following models as baselines for comparison with our proposed model: TripPy (Heck et al. 2020) combines three copy strategies for slot value predictions; DST-Picklist (Zhang et al. 2020a) first predicts slot updating information and then extracts the values of all slots via category classification; SST (Chen et al. 2020a) integrates schema graph and dialogue context, and controls the slot updating information from previous dialogue

Model	MultiWOZ		Total
	2.0	2.1	
TripPy (2020)	53.51	55.32	108.83
DST-Picklist (2020a)	54.39	53.30	107.69
SST (2020a)	51.17	55.23	106.40
Transformer-DST (2020a)	54.64	55.35	109.99
Graph-DST (2020b)	52.78	53.85	106.63
FPDSC (turn-level) (2021)	55.03	57.88	112.91
FPDSC (dual-level) (2021)	53.17	59.07	112.24
STAR (2021)	54.53	56.36	110.89
DSGFNet (2022)	-	56.70	-
PPTOD (2022)	53.89	57.45	111.34
LUNA (2022b)	55.31	57.62	112.93
Our Model	55.30	59.25	114.55

Table 1: Experimental results in Joint GA (%) based on MultiWOZ 2.0 and MultiWOZ 2.1.

states by graph attention networks; Transformer-DST (Zeng and Nie 2020a) jointly optimizes slot updating information prediction and slot value prediction for DST; Graph-DST (Zeng and Nie 2020b) constructs a dialogue state graph for each turn of dialogues to replace static schema graphs; FPDSC (Zhou et al. 2021) proposes a multi-level fusion of dialogue context and previous dialogue state information for DST; STAR (Ye et al. 2021) captures slot correlations by a self-attention mechanism; DSGFNet (Feng et al. 2022) fuses dialogue context and slot-domain relations to generate dynamic slot relations for DST; PPTOD (Su et al. 2022) learns task-oriented dialogue tasks with a multi-task learning style; LUNA (Wang et al. 2022b) aligns each slot with the utterance which is the most relevant to the slot to enhance DST. Besides, we do not compare our proposed model with TripPy+CoCoAug (Li et al. 2021), TripPy+SaCLog (Dai et al. 2021), DiCoS-DST (Guo et al. 2022) and ASSIST (Ye, Feng, and Yilmaz 2022) which have benefited from extra resources or more powerful pre-trained language models.

Performance Results

The experimental results in Joint GA are shown in Table 1. From the results, we observe that our model achieves the state-of-the-art performance on MultiWOZ 2.1 and almost the same performance as the state-of-the-art model LUNA on MultiWOZ 2.0. Further, we observe that previous baseline models cannot achieve the most competitive performance on both MultiWOZ 2.0 and MultiWOZ 2.1. Therefore, to compare the comprehensive performance among these models, we sum the performance on the two datasets for each model. As can be seen, our model outperforms the latest state-of-the-art model LUNA by 1.62%. Moreover, only our model and Graph-DST utilize both previous dialogue states and slot relations. However, our model outperforms Graph-DST by 7.92% in the summed Joint GA. We infer that Graph-DST relies on predicted dialogue states for enhancing input and capturing slot relations, thus the effectiveness of utilization is limited. Besides, as FPDSC employs fusion gates to reduce error information from predicted dialogue states, it achieves superior perfor-




Figure 3: Per-turn performance comparison in Joint GA (%) on MultiWOZ 2.1.

mance compared with Graph-DST and Transformer-DST, which also utilize previous dialogue states. It demonstrates the importance of reducing error information when utilizing previous dialogue states and our model uses knowledge distillation to tackle it.

Additionally, we compare the per-turn performance between STAR, LUNA and our model on MultiWOZ 2.1¹. From the results shown in Figure 3, we observe that our model is superior to LUNA and STAR in almost all dialogue turns. It further demonstrates the effectiveness of our model regardless of the length of dialogue.

Ablation Study

We conduct an ablation study on the MultiWOZ 2.0 and MultiWOZ 2.1 datasets to analyze the effectiveness of each module in our model. As shown in Table 3, we observe that both the Dialogue State Distillation and Inter-slot Contrastive Learning modules contribute significantly to the performance of our proposed model regardless of the used dataset. Here, we focus the discussion on MultiWOZ 2.0. Firstly, if we remove the Dialogue State Distillation and Inter-slot Contrastive Learning modules separately, the performance of the proposed model will be dropped by 1.77% and 1.05% respectively in Joint GA. It demonstrates that both the Dialogue State Distillation and Inter-slot Contrastive Learning modules can boost the performance of the proposed model. Secondly, if we train with only $\mathcal{L}_{\text{value}}$, the performance of the proposed model will be decreased significantly by 3.37% in Joint GA. It further demonstrates the effectiveness of distilling previous dialogue states from dialogue context and capturing the slot co-update relations.

Further Analysis

In this section, we study the effectiveness of capturing the slot co-updated relations in our proposed model. First, we remove both the Dialogue State Distillation and Inter-slot Contrastive Learning modules, denoted as “Baseline”. Based on the “Baseline”, we use different SOP training losses, which include “Cross Entropy Loss” and “Contrastive Learning Loss”, to replace the “Contrastive

¹As both STAR and LUNA do not present their per-turn performance on MultiWOZ 2.0, we follow them to only present the results on MultiWOZ 2.1 for comparison.

Turn	Dialogue Context	Our model w/o Distillation	Our model
1	[System]: There’s only one cheap hotel in town, the cambridge belfry, located in the west part of town. Do you need to book a room? [User]: Does it have a star of 3?	hotel-name: cambridge belfry hotel-pricerange: cheap hotel-star: 3 hotel-type: hotel	hotel-name: cambridge belfry hotel-pricerange: cheap hotel-star: 3 hotel-type: hotel
2	[System]: No, it has 4 star rating. [User]: Do you have any expensive hotels with a 3 star rating?	hotel-name: cambridge belfry hotel-pricerange:expensive hotel-star: 3 hotel-type: hotel	hotel-name: none hotel-pricerange:expensive hotel-star: 3 hotel-type: hotel

Table 2: A case study on the MultiWOZ 2.1 dataset. The wrong predictions of slots are in bold.

Model	MultiWOZ	
	2.0	2.1
Our model	55.30	59.25
- Dialogue State Distillation	53.53 (-1.77)	57.28 (-1.97)
- Inter-slot Contrastive Learning	54.25 (-1.05)	58.59 (-0.66)
- above two modules	51.93 (-3.37)	56.10 (-3.15)

Table 3: Ablation studies on the MultiWOZ 2.0 and MultiWOZ 2.1 datasets with Joint GA (%).

Learning Loss” in the Inter-slot Contrastive Learning module. Note that “Contrastive Learning Loss⁻” means that the dialogue-level selected negative sample set $\mathcal{N}_{dialogue}(s_j^t)$ is removed from $\mathcal{N}(s_j^t)$. Table 4 shows the performance comparison between the different variant models on the MultiWOZ 2.0 and MultiWOZ 2.1 datasets. As space is limited, we focus the analysis on MultiWOZ 2.0. From the results, we observe that using “Cross Entropy Loss” to learn SOP can improve the performance of “Baseline” by 0.52%. It shows that pushing the model to predict the updating information of slots has achieved some improvements for DST. Moreover, using “Contrastive Learning Loss⁻” can improve the performance of “Baseline” by 1.22%. It indicates that capturing the slot co-updated relations is more effective on the utilization of SOP labels compared with learning to predict the updating information of each slot independently in “Cross Entropy Loss”. Besides, we observe that using “Contrastive Learning Loss” of our model improves the performance of “Baseline” by 1.60%. It demonstrates that the contrastive learning between slots which are from different turns of dialogues can further enhance the performance for DST.

Moreover, we evaluate the performance difference of SOP when using “Cross Entropy Loss” and “Contrastive Learning Loss” based on “Baseline”. For MultiWOZ 2.0, the performance of SOP in Joint GA is 75.52% when using “Contrastive Learning Loss” and outperforms using “Cross Entropy Loss” by 2.93%. As for MultiWOZ 2.1, we obtain the similar performance difference of SOP. It helps present the effectiveness of our contrastive learning more intuitively.

Case Study

In this section, we discuss a dialogue snippet in the test set of the MultiWOZ 2.1 dataset. Table 2 shows the prediction

Model	MultiWOZ	
	2.0	2.1
Baseline	51.93	56.10
+ Cross Entropy Loss	52.45 (+0.52)	56.76 (+0.66)
+ Contrastive Learning Loss ⁻	53.15 (+1.22)	57.18 (+1.08)
+ Contrastive Learning Loss	53.53 (+1.60)	57.28 (+1.18)

Table 4: Performance comparison in Joint GA (%) between different variant models which utilize the label information of SOP on MultiWOZ 2.0 and MultiWOZ 2.1.

results of DST by using “Our model w/o Distillation” and our model. “Our model w/o Distillation” means that the Dialogue State Distillation module is removed from our model and previous dialogue states are directly concatenated on the dialogue context in Equation (1). In the first turn of the dialogue snippet, the system mentions a hotel name “cambridge belfry” and the user does not show a clear attitude towards this hotel. Maybe the dialogue context given by the system is too salient, both models ignore the ambiguous attitude of the user and predict the value of “hotel-name” wrongly. At the second turn, “Our model w/o Distillation” utilizes the wrongly predicted information of “hotel-name” at the first turn and the information misleads the prediction of “Our model w/o Distillation”. As our model uses the Dialogue State Distillation module to distill previous dialogue state information without introducing error propagation, it avoids the disturbance of the wrongly predicted information and predicts the value of “hotel-name” as “none” correctly.

Conclusion

In this paper, we propose a Dialogue State Distillation Network (DSDN) to utilize relevant information of previous dialogue states and fill the gap of utilization between training and testing, thus dynamically exploiting previous dialogue states without error propagation. In addition, an inter-slot contrastive learning loss is also proposed to effectively capture the slot co-update relations from dialogue context. Experimental results based on the MultiWOZ 2.0 and MultiWOZ 2.1 datasets have shown that our proposed model has achieved promising performance compared with state-of-the-art models for DST.

Acknowledgements

This work is supported by National Key Research and Development Program of China (Grant No. 2020YFC0832606), National Natural Science Foundation of China (Grant Nos. 61976021, U1811262), and Beijing Academy of Artificial Intelligence (BAAI).

References

- Budzianowski, P.; Wen, T.-H.; Tseng, B.-H.; Casanueva, I.; Ultes, S.; Ramadan, O.; and Gasic, M. 2018. MultiWOZ-A Large-Scale Multi-Domain Wizard-of-Oz Dataset for Task-Oriented Dialogue Modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 5016–5026.
- Chen, L.; Lv, B.; Wang, C.; Zhu, S.; Tan, B.; and Yu, K. 2020a. Schema-guided multi-domain dialogue state tracking with graph attention neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 7521–7528.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020b. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Dai, Y.; Li, H.; Li, Y.; Sun, J.; Huang, F.; Si, L.; and Zhu, X. 2021. Preview, Attend and Review: Schema-Aware Curriculum Learning for Multi-Domain Dialogue State Tracking. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, 879–885.
- Das, S. S. S.; Katiyar, A.; Passonneau, R. J.; and Zhang, R. 2022. CONTaiNER: Few-Shot Named Entity Recognition via Contrastive Learning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 6338–6353.
- Eric, M.; Goel, R.; Paul, S.; Sethi, A.; Agarwal, S.; Gao, S.; Kumar, A.; Goyal, A.; Ku, P.; and Hakkani-Tur, D. 2020. MultiWOZ 2.1: A Consolidated Multi-Domain Dialogue Dataset with State Corrections and State Tracking Baselines. In *Proceedings of the 12th Language Resources and Evaluation Conference*, 422–428.
- Feng, Y.; Lipani, A.; Ye, F.; Zhang, Q.; and Yilmaz, E. 2022. Dynamic Schema Graph Fusion Network for Multi-Domain Dialogue State Tracking. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 115–126.
- Gao, S.; Sethi, A.; Agarwal, S.; Chung, T.; and Hakkani-Tur, D. 2019. Dialog State Tracking: A Neural Reading Comprehension Approach. In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, 264–273.
- Guo, J.; Shuang, K.; Li, J.; Wang, Z.; and Liu, Y. 2022. Beyond the Granularity: Multi-Perspective Dialogue Collaborative Selection for Dialogue State Tracking. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2320–2332.
- Heck, M.; van Niekerk, C.; Lubis, N.; Geishausser, C.; Lin, H.-C.; Moresi, M.; and Gasic, M. 2020. TripPy: A Triple Copy Strategy for Value Independent Neural Dialog State Tracking. In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, 35–44.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the Knowledge in a Neural Network. *arXiv preprint arXiv:1503.02531*.
- Hosseini-Asl, E.; McCann, B.; Wu, C.-S.; Yavuz, S.; and Socher, R. 2020. A simple language model for task-oriented dialogue. *Advances in Neural Information Processing Systems*, 33: 20179–20191.
- Kenton, J. D. M.-W. C.; and Toutanova, L. K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL-HLT*, 4171–4186.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33: 18661–18673.
- Kim, S.; Yang, S.; Kim, G.; and Lee, S.-W. 2020. Efficient Dialogue State Tracking by Selectively Overwriting Memory. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 567–582.
- Lee, H.; Lee, J.; and Kim, T.-Y. 2019. SUMBT: Slot-Utterance Matching for Universal and Scalable Belief Tracking. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 5478–5483.
- Li, S.; Yavuz, S.; Hashimoto, K.; Li, J.; Niu, T.; Rajani, N.; Yan, X.; Zhou, Y.; and Xiong, C. 2021. CoCo: Controllable Counterfactuals for Evaluating Dialogue State Trackers. In *International Conference on Learning Representations*.
- Lin, W.; Tseng, B.-H.; and Byrne, B. 2021. Knowledge-Aware Graph-Enhanced GPT-2 for Dialogue State Tracking. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 7871–7881.
- Mi, F.; Wang, Y.; and Li, Y. 2022. Cins: Comprehensive instruction for few-shot learning in task-oriented dialog systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 11076–11084.
- Ouyang, Y.; Chen, M.; Dai, X.; Zhao, Y.; Huang, S.; and Chen, J. 2020. Dialogue state tracking with explicit slot connection modeling. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 34–40.
- Ramadan, O.; Budzianowski, P.; and Gasic, M. 2018. Large-Scale Multi-Domain Belief Tracking with Knowledge Sharing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 432–437.
- Ren, L.; Xie, K.; Chen, L.; and Yu, K. 2018. Towards Universal Dialogue State Tracking. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2780–2786.
- Saha, S.; Yadav, P.; and Bansal, M. 2022. Explanation Graph Generation via Pre-trained Language Models: An Empirical

- Study with Contrastive Learning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1190–1208.
- Su, Y.; Shu, L.; Mansimov, E.; Gupta, A.; Cai, D.; Lai, Y.-A.; and Zhang, Y. 2022. Multi-Task Pre-Training for Plug-and-Play Task-Oriented Dialogue System. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 4661–4676.
- Sun, H.; Wang, R.; Chen, K.; Utiyama, M.; Sumita, E.; and Zhao, T. 2020. Knowledge Distillation for Multilingual Unsupervised Neural Machine Translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 3525–3535.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, D.; Ding, N.; Li, P.; and Zheng, H. 2021a. CLINE: Contrastive Learning with Semantic Negative Examples for Natural Language Understanding. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2332–2342.
- Wang, Y.; Xu, C.; Sun, Q.; Hu, H.; Tao, C.; Geng, X.; and Jiang, D. 2022a. PromDA: Prompt-based Data Augmentation for Low-Resource NLU Tasks. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 4242–4255.
- Wang, Y.; Zhao, J.; Bao, J.; Duan, C.; Wu, Y.; and He, X. 2022b. LUNA: Learning Slot-Turn Alignment for Dialogue State Tracking. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 3319–3328.
- Wang, Z.; Wang, X.; Han, X.; Lin, Y.; Hou, L.; Liu, Z.; Li, P.; Li, J.; and Zhou, J. 2021b. CLEVE: Contrastive Pre-training for Event Extraction. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 6283–6297.
- Wei, K.; Sun, X.; Zhang, Z.; Zhang, J.; Zhi, G.; and Jin, L. 2021. Trigger is not sufficient: Exploiting frame-aware knowledge for implicit event argument extraction. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 4672–4682.
- Wu, P.; Zou, B.; Jiang, R.; and Aw, A. 2020. GCDST: A graph-based and copy-augmented multi-domain dialogue state tracking. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, 1063–1073.
- Xing, B.; and Tsang, I. W. 2022. Co-guiding Net: Achieving Mutual Guidances between Multiple Intent Detection and Slot Filling via Heterogeneous Semantics-Label Graphs. *arXiv preprint arXiv:2210.10375*.
- Xu, P.; and Hu, Q. 2018. An End-to-end Approach for Handling Unknown Slot Values in Dialogue State Tracking. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1448–1457.
- Ye, F.; Feng, Y.; and Yilmaz, E. 2022. ASSIST: Towards Label Noise-Robust Dialogue State Tracking. In *Findings of the Association for Computational Linguistics: ACL 2022*, 2719–2731.
- Ye, F.; Manotumruksa, J.; Zhang, Q.; Li, S.; and Yilmaz, E. 2021. Slot self-attentive dialogue state tracking. In *Proceedings of the Web Conference 2021*, 1598–1608.
- Zeng, Y.; and Nie, J.-Y. 2020a. Jointly optimizing state operation prediction and value generation for dialogue state tracking. *arXiv preprint arXiv:2010.14061*.
- Zeng, Y.; and Nie, J.-Y. 2020b. Multi-domain dialogue state tracking based on state graph. *arXiv preprint arXiv:2010.11137*.
- Zhang, J.; Hashimoto, K.; Wu, C.-S.; Wang, Y.; Philip, S. Y.; Socher, R.; and Xiong, C. 2020a. Find or Classify? Dual Strategy for Slot-Value Predictions on Multi-Domain Dialog State Tracking. In *Proceedings of the Ninth Joint Conference on Lexical and Computational Semantics*, 154–167.
- Zhang, Z.; Shu, X.; Yu, B.; Liu, T.; Zhao, J.; Li, Q.; and Guo, L. 2020b. Distilling knowledge from well-informed soft labels for neural relation extraction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 9620–9627.
- Zhang, Z.; Takanobu, R.; Zhu, Q.; Huang, M.; and Zhu, X. 2020c. Recent advances and challenges in task-oriented dialogue systems. *Science China Technological Sciences*, 63(10): 2011–2027.
- Zhong, V.; Xiong, C.; and Socher, R. 2018. Global-locally self-attentive encoder for dialogue state tracking. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1458–1467.
- Zhou, J.; Wu, H.; Lin, Z.; Li, G.; and Zhang, Y. 2021. Dialogue State Tracking with Multi-Level Fusion of Predicted Dialogue States and Conversations. In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, 228–238.
- Zhou, W.; Xu, C.; and McAuley, J. 2022. BERT learns to teach: Knowledge distillation with meta learning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 7037–7049.
- Zhu, S.; Li, J.; Chen, L.; and Yu, K. 2020. Efficient Context and Schema Fusion Networks for Multi-Domain Dialogue State Tracking. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, 766–781.