

OPAL: Ontology-Aware Pretrained Language Model for End-to-End Task-Oriented Dialogue

Zhi Chen¹, Yuncong Liu¹, Lu Chen^{1*}, Su Zhu², Mengyue Wu¹, Kai Yu^{1*}

¹X-LANCE Lab, Department of Computer Science and Engineering

MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University

State Key Lab of Media Convergence Production Technology and Systems, Beijing, China

²AISpeech Co., Ltd., Suzhou, China

{zhenchi713, chenlusz, kai.yu}@sjtu.edu.cn

Abstract

This paper presents an ontology-aware pre-trained language model (OPAL) for end-to-end task-oriented dialogue (TOD). Unlike chit-chat dialogue models, task-oriented dialogue models fulfill at least two task-specific modules: Dialogue state tracker (DST) and response generator (RG). The dialogue state consists of the domain-slot-value triples, which are regarded as the user's constraints to search the domain-related databases. The large-scale task-oriented dialogue data with the annotated structured dialogue state usually are inaccessible. It prevents the development of the pre-trained language model for the task-oriented dialogue. We propose a simple yet effective pretraining method to alleviate this problem, which consists of two pretraining phases. The first phase is to pretrain on large-scale contextual text data, where the structured information of the text is extracted by the information extracting tool. To bridge the gap between the pretraining method and downstream tasks, we design two pretraining tasks: ontology-like triple recovery and next-text generation, which simulates the DST and RG, respectively. The second phase is to fine-tune the pretrained model on the TOD data. The experimental results show that our proposed method achieves an exciting boost and obtains competitive performance even without any TOD data on CamRest676 and MultiWOZ benchmarks.

1 Introduction

A task-oriented dialogue (TOD) system aims to assist users in accomplishing a specific task by interacting with natural language, for example, reserving a hotel or booking flight tickets. With the popularity of the industrial dialogue system, the

task-oriented dialogue system attracts extensive attention in research.

The existing task-oriented dialogue system can be classified into two categories: pipeline format and end-to-end format. The pipeline TOD system (Uttes et al., 2017; Weisz et al., 2018) is composed of four modules: natural language understanding (NLU) (Quirk et al., 2015), dialogue state tracking (DST) (Xu et al., 2020; Chen et al., 2020c), dialogue policy (DP) (Chen et al., 2018, 2019, 2020b), and natural language generation (NLG) (Wen et al., 2015; Li et al., 2016; Zhao et al., 2017). Since each module of the system is trained separately and executes sequentially, it faces two serious issues: error accumulation and high annotation cost. Thus, the end-to-end dialogue system (Lee et al., 2019b; Zhao et al., 2019) gradually becomes the research focus, which formulates the task-oriented dialogue as a sequence-to-sequence task. The dialogue state, database (DB) state, and the corresponding system response are directly concatenated together and flattened as a token sequence. The DB state is the status of the domain-related database searched with the dialogue state, as shown in Figure 1.

Thanks to the success of pretraining language models (Kenton and Toutanova, 2019; Raffel et al., 2020), effective application has shed light on open-domain (chit-chat) dialogues (Bao et al., 2020; Adiwardana et al., 2020). Nevertheless, utilizing such pretrained language models on TOD systems remains challenging due to the limited TOD data with annotated dialogue state. Unlike the open-domain dialogue, TOD is restricted by a dialogue ontology, which defines the dialogue domains, the slots and their candidate values. The TOD system needs to predict the dialogue state and feedback the DB content to accomplish a task. The dialogue state is structured information

*The corresponding authors are Lu Chen and Kai Yu.

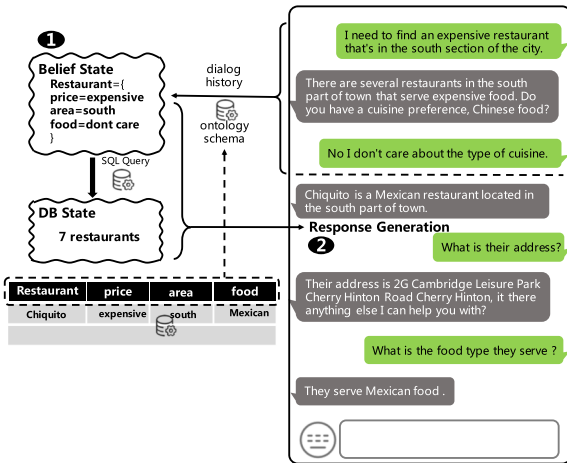


Figure 1: A task-oriented dialogue example. The dialogue model needs to infer the dialogue state based on the dialogue history and ontology schema. The DB state is searched by the generated dialogue state. The last step is to generate system response.

extracted from the dialogue context, which is a set of domain-slot-value triples.

Recently, some works (Hosseini-Asl et al., 2020; Lin et al., 2020) try to directly leverage the pretrained language models, e.g., GPT-2 (Radford et al., 2019) and BART (Lewis et al., 2020), in the end-to-end TOD system. Such models (Mehri et al., 2019) are pretrained on the large-scale contextual text with the general self-supervised method, e.g., language modeling and language denoising. However, in the task-oriented dialogue task, the dialogue state is structured information rather than a contextual text. The inconsistency between the pretrained and downstream tasks will impact the performance of the PLMs on the TOD benchmarks. To alleviate this problem, SOLOIST (Peng et al., 2020a) fine-tunes the pretrained GPT-2 with the existing annotated TOD data and then transfers it to the other task-oriented dialogue generation tasks. Similarly, NCM (Liu et al., 2021) first warm-ups the Transformer-based model with large-scale Reddit¹ (Völske et al., 2017) data and then fine-tunes the model on the TOD data. However, the existing TOD data is too limited to pretrain a large-scale language model.

To alleviate the problems above and advance pretrained language model research, especially its application on TOD, we propose an **Ontology-aware Pretrained Language model (OPAL)**.

¹<http://files.pushshift.io/reddit/>.

From the high-level perspective, we can abstract the end-to-end TOD task into two sub-tasks: **ontology-like triple recovery** and **next-text generation**, which corresponds to dialogue state tracking task and response generating task. The ontology-like triple recovery in the TOD means to predict the corresponding value given the domain and the slot. The next-text generation is easy to design for the contextual text, which directly fulfills with masking the last sentence. The challenge is how to design the ontology-like triple recovery task, which needs to obtain the structured information from the contextual text. In this paper, we utilize the external OpenIE tools (Angeli et al., 2015; Kolluru et al., 2020)² to extract the relation triples (subject-relation-object) from the contextual text as the structured information. In most cases, the domain-slot-value triple can be regarded as relation triple, for example, *train-arrive-12:30*. The relation triples extracted from the contextual text can be regarded as the ontology-like triples. We design self-supervised **ontology-like triple recovery** task and **next-text generation** task to pretrain the model.

The main contributions of this paper are summarized below:

- We leverage the external tool OpenIE to generate large amounts of TOD-like data, which is important for the development of pretrained language models in the TOD community.
- To the best of our knowledge, this is the first work to design *self-supervised* tasks for end-to-end TOD tasks. It bridges the gap between pretrained language models and end-to-end TOD models.
- The experimental results show that our proposed pretrained model OPAL can get competitive performance even without any annotated TOD data in the pretraining process.
- Further fine-tuned on the annotated TOD data, our proposed method obtains exciting performance gain on CamRest676 and MultiWOZ datasets.

2 End-to-End Task-Oriented Dialogue

As previously introduced, the pipeline dialogue system consists of four modules. The NLU

²<https://github.com/dair-iitd/openie6>.

module is to recognize the user’s intents and the corresponding slot values. The DST module combines the previous state and the results of the NLU to update the current dialogue state. The DP module chooses the discrete dialogue acts according to the dialogue state and the database state to respond to the user. The NLG module generates the natural language based on the chosen dialogue acts. There are at least four kinds of annotation in such systems: the user’s intent, the slot value, the dialogue state, and the dialogue act. The heavy annotation labor enormously increases the cost of building a pipeline system. Its poor scalability further influences the pipeline dialogue system development.

Compared with the pipeline system, this paper’s end-to-end task-oriented dialogue system only requires the annotated dialogue state. The end-to-end TOD system is fed with the dialogue context \mathbf{c} and generates the dialogue state \mathbf{b} and delexicalized response \mathbf{r} , where the database (DB) state \mathbf{d} is retrieved from the results searched with \mathbf{b} . The delexicalized response means that the specific slot values are replaced with the corresponding slot placeholders. The lexicalized response is recovered from the delexicalized one with the generated dialogue state and DB state. The training sample at each dialogue turn of the end-to-end TOD model is defined as:

$$\mathbf{x} = (\mathbf{c}, \mathbf{b}, \mathbf{d}, \mathbf{r}). \quad (1)$$

For the task-oriented dialogue, the dialogue context not only consists of the dialogue history \mathbf{h} but also includes the dialogue ontology schema \mathbf{s} , which is usually ignored by the existing end-to-end models. The ontology can be seen as prior knowledge designed by the dialogue expert, which defines the dialogue domain, slots, and candidate values. The end-to-end TOD model needs to fulfill two sub-tasks: Dialogue state tracking (DST) and response generation (RG). Formally, the learning goal of the TOD model is to maximize the joint probability $p_{\theta}(\mathbf{x})$, which can be factorized in an auto-regressive manner as:

$$p_{\theta}(\mathbf{x}) = p(\mathbf{c}, \mathbf{b}, \mathbf{d}, \mathbf{r}), \quad (2)$$

$$= p(\mathbf{h}, \mathbf{s}, \mathbf{b}, \mathbf{d}, \mathbf{r}), \quad (3)$$

$$= \underbrace{p(\mathbf{r}|\mathbf{b}, \mathbf{d}, \mathbf{h}, \mathbf{s})}_{\text{RG}} \underbrace{p(\mathbf{b}|\mathbf{h}, \mathbf{s})}_{\text{DST}} p(\mathbf{h}, \mathbf{s}), \quad (4)$$

where the factorization from (3) to (4) is based on the fact that the database-lookup operation is a deterministic process. The $p(\mathbf{h}, \mathbf{s})$ is the prior probability of the paired dialogue and ontology (as the input of the model), which depends on the distribution of the (pre-)training data and is independent on the model. The dialogue state tracker intrinsically extracts the ontology-related constraints demanded by the user, where the ontology schema is given in advance.

3 Ontology-Aware Pretraining Method

The existing task-oriented dialogue data with the given ontology is limited to pretrain the language model. To increase the scale of the pretraining data, we divide the pretraining process into two phases. The first phase pretrains the model on the large-scale contextual text. The triples of the text are extracted by the latest neural-based OpenIE6 (Kolluru et al., 2020). There is still a glaring discrepancy between the contextual text and the dialogue. For example, the dialogue always contains co-reference and information ellipsis (Iyyer et al., 2017). We pretrain the model on the smaller TOD data at the second phase to further decrease the gap between the pretrained model and the downstream tasks. The two phases are complementary to each other introduced as below:

Phase-1: Pretrained on Contextual Text In traditional dialogue pretrained models (Zhang et al., 2020c), the crawled Reddit data is popular to be used as pretrained corpus. However, Reddit data contain lots of the co-reference and information ellipsis, which seriously impact the performance of the external information extraction tool. Different from the dialogue data, the co-reference and information ellipsis are infrequent in the contextual text of the Wikipedia.³ More details are shown in Section 5.1 to validate the effects of pretrained corpora. We use the neural-based OpenIE6 to extract the ontology-like knowledge of contextual text automatically. We directly simulate the extracted subject-relation-object triples as the domain-slot-value triples. As shown in Figure 2, the object values in the extracted ontology are masked during the pretraining process. One of our designed pretraining tasks is to recover the

³<https://dumps.wikimedia.org/enwiki/latest/enwiki-latest-pages-articles.xml.bz2>.

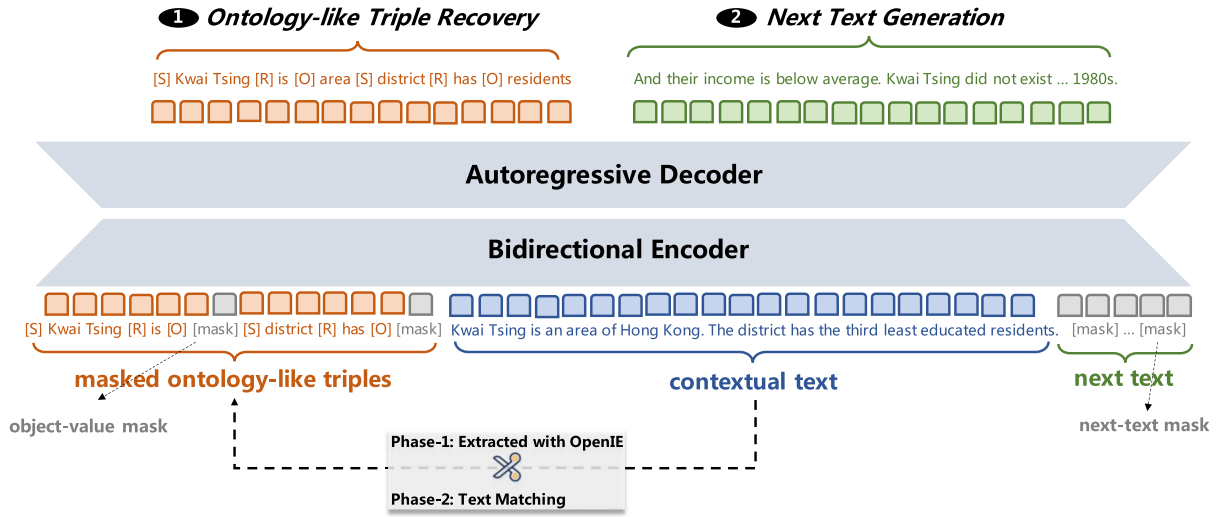


Figure 2: The ontology-aware pretraining method contains two masking strategies: object-value mask and next-text mask. The corresponding self-supervised learning methods are ontology-like triple recovery and next-text generation. The ontology-like triples of the contextual text are extracted by the external tool OpenIE at the pretraining phase-1 and matched with the given whole ontology at phase-2.

ontology-like triples (named **ontology-like triple recovery** [OR]), which is similar to the DST task. To increase the inference ability of the pretrained model, we mask the next text (one or two sentences, which are randomly chosen) and push the model to infer the next text (named **next-text generation** [NTG]), which is similar to the RG task. Thus, the pretraining sample is composed of four elements: masked ontology-like triples \hat{s} , the masked document context \hat{h} , ontology-like triples \hat{b} , and the next text \hat{r} . Similar to Equation 4, the goal of the pretraining model is to maximize the joint probability:

$$p(\hat{h}, \hat{s}, \hat{b}, \hat{r}) = \underbrace{p(\hat{r}|\hat{b}, \hat{h}, \hat{s})}_{\text{NTG}} \underbrace{p(\hat{b}|\hat{h}, \hat{s})}_{\text{OR}} p(\hat{h}, \hat{s}). \quad (5)$$

To obtain the qualified triples of a sentence using OpenIE6, we remove all the stopwords in the triples and filter the triples in which one of the triple components is a blank space. It is also the main reason that we do not choose the Reddit at this pretraining phase. There are many pronouns in the text, with which is hard to extract qualified triples. This pretraining phase vastly increases the scale of the pretraining data. There are four steps to filter the triples of the sentence:

- Remove all the stopwords in the triples and filter the triples in which one of triple component is a blank space.

- Remove the triples in which one of the triple components contains more than 4 words.
- For the triples that have the same subject-relation pair, randomly select one of the triples and remove the others.
- Randomly select two triples from the rest of triples, if their length is larger than two. This is to extract no more than two triples in a sentence.

Phase-2: Pretrained on TOD Data To further decrease the gap between the pretrained language model and the end-to-end model, we leverage the smaller task-oriented data in the pretraining process. Instead of extracting the ontology-like triples with OpenIE6, the TOD ontology is designed by the dialogue experts. We directly use the text matching method to extract the domain-slot-value triples from the dialogue context with the given ontology. Note that the extracted triples with text matching operation are not the dialogue state. In this pretraining phase, the system-mentioned ontology triples also have to be recovered, which is consistent with the previous pretraining process. In other words, different from SOLOIST (Peng et al., 2020a) and NCM (Liu et al., 2021), we do not need to use the annotated dialogue state and only utilize the given dialogue ontology to match the ontology-related triples. This attribute increases the generalization of the proposed

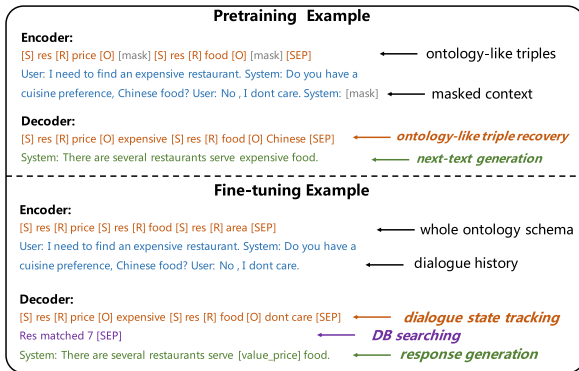


Figure 3: A toy example to show the differences between the pretraining data and the fine-tuning data.

ontology-aware pretraining methods, where the ontology is much easier to be obtained than the dialogue state annotation. We share a toy example to distinguish the usage of the pretraining TOD data and the fine-tuning data of the end-to-end TOD task in Figure 3. During the pretraining process, the ontology is extracted from the context, which is just a part of the given ontology. The ontology recovery is to recover all the ontology-related triples, for example, the triple res-food-Chinese is not in the dialogue state. During fine-tuning process, there is an extra database searching step.

4 Experiments

We evaluate our proposed pretrained model OPAL on dialogue state tracking tasks and end-to-end TOD tasks. To further validate the effectiveness of the proposed OPAL, we conduct the ablation study to analyze the effects of the different pretraining ingredients. Last but not least, we design the resource-limited experiments to figure out the sample efficiency of the proposed OPAL on the end-to-end TOD task and show some cases to study the strength of the proposed OPAL.

4.1 Corpora

At phase-1 of the proposed OPAL, we use the Wikipedia corpus to pretrain the model. There are 72.24 million samples collected from Wikipedia. We have used five task-oriented dialogue datasets in the experiments, shown in Table 1, where the Schema (Rastogi et al., 2020) and the TaskMaster (Byrne et al., 2019) are leveraged in the phase-2

Dataset	#Dialogue	#Domain	#Slot	X-Domain	Usage
Schema	22,825	17	123	✓	P
TaskMaster	17,304	7	281	✗	P
MultiWOZ	10,438	7	46	✓	F
WOZ	1,200	1	4	✗	F
CamRest676	676	1	4	✗	F

Table 1: The five task-oriented dialogue datasets used in this paper. The X-domain (cross-domain) means that a dialogue can contain different dialogue domains. The usages of the datasets are grouped into Pretraining (named as P) and Fine-tuning (named as F), which means that the corresponding dataset is used in the pretraining phase and the fine-tuning phase.

of the pretraining process and the rests are the downstream benchmarks. The WOZ (Mrkšić et al., 2017) and the CamRest676 (Wen et al., 2016) are the single-domain task-oriented dialogue corpora, which are the well-studied DST benchmark and end-to-end TOD benchmark, respectively. MultiWOZ is a kind of multi-domain dialogue corpus, which is challenging due to its multi-domain setting and diverse language styles. There are two versions of the MultiWOZ dataset used in the experiments: MultiWOZ2.0 (Budzianowski et al., 2018) and MultiWOZ2.1 (Eric et al., 2019), where MultiWOZ2.1 fixes most of DST annotation errors in MultiWOZ2.0. To fairly compare to the other baselines, we run the end-to-end TOD tasks on the MultiWOZ2.0 and run the DST tasks on the MultiWOZ2.0 and MultiWOZ2.1.

4.2 Metrics

For the dialogue state tracking task, we use the joint goal accuracy (*JGA*) to evaluate the models. Only if all the predicted slot values at each turn are exactly matched with the golden, does it confirm the successful prediction of the DST model. For the end-to-end TOD task, there are three reported scores: *Inform*, *Success*, and *BLEU*. *Inform* measures whether the system response has provided the right entity. *Success* reports whether the system response has provided all the requested slots. *BLEU* evaluates the naturalness of the generated system response. Following Budzianowski et al. (2018), the combined score (*Combined*) is also reported using $Combined = (Inform + Success) \times 0.5 + BLEU$.

4.3 Experimental Setup

We implement the proposed OPAL with HuggingFace’s Transformers (Wolf et al., 2020) and BART, which is a pretrained denoising autoencoder. To validate the generalization of the proposed pretraining method, we set the base version and large version (BART_L) of the BART as the backbone of the proposed OPAL, named OPAL and OPAL_L, respectively. The learning rates of the pretraining and fine-tuning are both 1e-5. The optimizer is AdamW. At phase-1 of the pretraining process, the total training steps is 280,000 and the batch size is 256. It is pretrained on four P100 GPUs (16G memory for each). This pretraining process costs 260 hours (one epoch on Wikipedia). Similar to NCM (Liu et al., 2021), we pretrain 100,000 steps at the phase-2. At the fine-tuning process of the downstream tasks, the batch size is 32. We conduct significant tests (paired t-test) (Koehn, 2004) with five different seeds on the end-to-end TOD task, where the final results are trained with the default seed 42.

4.4 Baselines

We compare the proposed OPAL with the strong baselines, which hold the state-of-the-art (SOTA) performance on the DST and end-to-end TOD.

The DST models can be divided into two categories: classification method and generation method. The classification methods rely on the optional slot values of the ontology and select the value from it. Their scalability is a severe problem for the practical dialogue system. The generation methods directly extract the values from the dialogue context, which are comparable to the proposed OPAL.

For the end-to-end TOD tasks, the existing end-to-end TOD systems can be grouped into modular systems and sequential systems. The modular systems use multiple decoders to generate the downstream outputs independently and are trained in an end-to-end manner. The sequential systems formulate the end-to-end TOD as a single sequence prediction problem. Sequicity (Lei et al., 2018) proposes a two-stage CopyNet method to generate the dialogue state and the system response. HRED-TS (Peng et al., 2019) proposes a teacher-student framework with a hierarchical recurrent encoder-decoder backbone. DAMD (Zhang et al., 2020b) designs a domain-aware multi-decoder network with the multi-action data

augmentation method. DSTC8 Winner Ham et al., 2020 and SimpleTOD (Hosseini-Asl et al., 2020) successfully leverage the pretrained language model GPT-2 for the end-to-end TOD modeling in the unified way. Inspired by SimpleTOD, SOLOIST (Peng et al., 2020a) fine-tunes GPT-2 with out-of-domain TOD data and obtains excellent transferability. MinTL-BART (Lin et al., 2020) and UBAR (Yang et al., 2021) improve the end-to-end TOD system by changing the input content without extra assumptions. HTER (Santra et al., 2021) improves the end-to-end TOD system by a hierarchical dialogue modeling mechanism. NCM (Liu et al., 2021) improves the decoder with the noisy channel model and proposes a two-stage pretraining method to warm up the Transformer-based model, where the model first pretrains on the Reddit corpus and then on the task-oriented dialogues. NCM is the closest method to our proposed method. We mainly compare our proposed method with this method.

4.5 Results on End-to-End TOD

We first fine-tune our pretrained models OPAL and OPAL_L on two well-studied end-to-end TOD datasets: MultiWOZ2.0 and CamRest676, as shown in Table 2 and Table 3. We compare our models with strong baselines in the end-to-end dialogue learning setting.

To validate the generalization of our proposed ontology-aware pretraining method, we set the base-version and large-version BART as the backbones of the pretraining models. Compared with the performance fine-tuned on the original BARTs, the proposed OPAL and OPAL_L achieve 7.32 and 10.94 overall performance gains on the MultiWOZ2.0 dataset and absolute 7.04 point gains on the CamRest676 dataset. SOLOIST (Peng et al., 2020a) and NCM (Liu et al., 2021) are the two closest methods to OPAL, which both leverage the out-of-domain TOD in pretraining the Transformer-based models. Different from our methods, these two approaches rely on DST annotation. Our proposed models can still obtain the best task completion (Inform and Success) and have lower BLEU scores than NCM barely. Compared with overall baselines, our proposed models reach the new SOTA overall performance (Combined) on both two datasets. The large-version model OPAL_L outperforms the base-version OPAL with a 2.53 performance gain on the combined score. To fairly compare to

Model	Model Size	Dialogue Act	Inform	Success	BLEU	Combined
Sequicity (Lei et al., 2018)	–	✗	66.40	45.30	15.54	71.39
HRED-TS (Peng et al., 2019)	–	✓	70.00	58.00	17.50	81.50
DSTC8 Winner (Ham et al., 2020)	124M	✓	73.00	62.40	16.00	83.50
DAMD (Zhang et al., 2020b)	–	✓	76.40	60.40	16.60	85.00
SimpleTOD (Hosseini-Asl et al., 2020)	117M	✓	84.40	70.10	15.01	92.26
SOLOIST (Peng et al., 2020a)	117M	✗	85.50	72.90	16.54	95.74
MinTL-BART (Lin et al., 2020)	406M	✗	84.88	74.91	17.89	97.78
UBAR (Yang et al., 2021)	82M	✗	88.20	79.50	16.43	100.28
NCM _B (Liu et al., 2021)	116M	✓	85.90	74.80	19.76	100.11
NCM _L (Liu et al., 2021)	292M	✓	86.90	76.20	20.58	102.13
HTER (Santra et al., 2021)	–	✓	91.72	75.80	19.05	102.81
<hr/>						
BART	139M	✗	87.50	72.20	16.67	96.53
OPAL	139M	✗	89.40	81.10	18.60	103.85
BART _L	406M	✗	86.20	70.30	17.01	95.26
OPAL _L	406M	✗	88.00	82.80	20.80	106.20

Table 2: End-to-end response generation results on MultiWOZ2.0. ✓ and ✗ denote whether the dialogue act annotation is used in the training process. We list all the model sizes of the Transformer-based end-to-end TOD models. Notice that we directly use the UBAR result provided by Liu et al. According to the released code of the UBAR, they have not used the standard evaluation metric, which is unfair to compare to other methods. We also run their code with released model checkpoint, whose combined score is even worse than the result provided by Liu et al. Results are significant ($p < 0.01$) comparing the OPAL model and BART model as the initialized TOD model.

Model	Inform	Success	BLEU	Combined
Sequicity	92.30	85.03	21.40	110.20
SOLOIST	94.70	87.10	25.50	116.40
NCM _B	94.30	85.20	25.98	115.73
NCM _L	95.40	85.30	26.89	117.24
<hr/>				
BART	96.31	79.41	24.74	112.61
OPAL	96.32	89.86	26.56	119.65

Table 3: End-to-end response generation results on CamRest676.

other baselines, we only report the base-version OPAL’s performance in the next experiments.

Compared with NCM_B, our proposed OPAL has higher task-completion (revealed by $\text{Inform} + \text{Success}$) $\times 0.5$ performance. However, BLEU score of OPAL is lower than BLEU of NCM_B. Figure 4 shows the correlation between BLEU score and task-completion ability. The fine-tuned model tried to balance between BLEU score and task-completion ability. With the progress of training process, the BLEU score is descending and the task-completion ability is enhanced. The main reason is that there are different expres-




Figure 4: The correlation between BLEU score and task-completion ability at first 20 fine-tuning epochs. They are the average evaluation results on MultiWOZ2.0 with different five seeds.

sions on the same system intention, which is the typical one-to-many mapping problem (Zhao and Eskenazi, 2018) in the dialogue generation. The final fine-tuned model has stronger task-completion ability but sacrifices the dialogue diversity. In the evaluation, we choose the model with the highest combination score.

Model	JGA	
	MultiWOZ2.0	MultiWOZ2.1
FJST (Eric et al., 2017)	40.20	38.00
HyST (Goel et al., 2019)	44.24	–
SUMBT (Lee et al., 2019a)	46.65	–
TOD-BERT (Wu et al., 2020)	–	48.00
DST-Picklist (Zhang et al., 2020a)	–	53.30
SST (Chen et al., 2020a)	51.17	55.23
TripPy (Heck et al., 2020)	–	55.29
FPDSC (Zhou et al., 2021)	53.17	59.07
TRADE (Wu et al., 2019)	48.62	45.60
COMER (Ren et al., 2019)	48.79	–
NADST (Le et al., 2020)	50.52	49.04
DSTQA (Zhou and Small, 2019)	51.44	51.17
SOM-DST (Kim et al., 2020)	51.38	52.57
MinTL-BART (Lin et al., 2020)	52.10	53.62
SimpleTOD (Hosseini-Asl et al., 2020)	–	55.72
UBAR (Yang et al., 2021)	52.59	56.20
SOLOIST (Peng et al., 2020a)	53.20	56.85
OPAL	54.10	57.05

Table 4: Dialogue state tracking results on MultiWOZ2.0 and MultiWOZ2.1. The upper part is for classification-based models and the lower part belongs to generation-based models.

4.6 Results on DST

The classification-based DST models and generation-based DST models are shown in the upper part and lower part of the Table 4 and Table 5, respectively. Table 4 reports the DST results on the MultiWOZ2.0 and MultiWOZ2.1 datasets. Our proposed OPAL can obtain the highest JGA among all the generation-based baselines on both datasets. Compared with the classification-based SOTA model FPDSC (Zhou et al., 2021), OPAL can even achieve 0.93% JGA improvement on the MultiWOZ2.0 dataset. Table 5 shows the DST results on WOZ, which is a single-domain dataset and has only 4 slots. The computational complexity of the classification-based models is proportional to the number of the candidate slot values. The classification-based models have the advantage of predicting slot values from valid candidates on the simpler dialogue domain. It is the main reason that the classification-based models are more popular on the single-domain WOZ dataset. Compared with the well-designed classification-based model BERT-DST (Lai et al., 2020), OPAL has a 0.7% JGA gain. OPAL gets 6.7% higher JGA over the novel generation-based model TRADE (Wu et al., 2019). Notice that we do not compare the pro-

Model	JGA WOZ
NBT (Mrkšić et al., 2017)	84.4
GLAD (Zhong et al., 2018)	88.1
GCE (Nouri and Hosseini-Asl, 2018)	88.5
G-SAT (Balaraman and Magnini, 2019)	88.7
StateNet (Ren et al., 2018)	88.9
BERT-DST (Lai et al., 2020)	90.5
TRADE (Wu et al., 2019) [†]	84.5
OPAL	91.2

Table 5: Dialogue state tracking results on the single-domain WOZ. The upper part is classification-based models and the lower part belongs to generation-based model. [†] represents that the result is produced by us from the released code.

posed model with variants (Yu et al., 2020; Li et al., 2020; Dai et al., 2021) of the data augmentation methods based on TripPy (Heck et al., 2020). In this paper, we pay more attention on the end-to-end task-oriented dialogue generation task. Our model is completely compatible with these data augmentation methods. In the future, we will try these augmentation methods on our model.

5 Analysis

The analysis experiments evaluate the proposed OPAL on the end-to-end TOD tasks to answer three main questions: **Q1**: What role do the different pretraining corpora (Wikipedia and out-of-domain TOD) play? **Q2**: What is the main factor that affects the pretrained model? **Q3**: Does OPAL have a higher sample efficiency than the original BART in the limited-resource setting?

5.1 Ablation Study

Table 6 reports the ablation study of the proposed OPAL, which has two pretraining phases. The phase-1 of OPAL pretrains on the contextual texts and phase-2 pretrains on the task-oriented dialogues with the ontology-aware pretraining method. To evaluate the effects of these two corpora, we separately pretrain the backbones (BART) only on the pretraining data contextual texts or task-oriented dialogues, where the pretrained models are named as WIKI and TOD,

Model	MultiWOZ2.0			
	<i>Inform</i>	<i>Success</i>	<i>BLEU</i>	<i>Combined</i>
OPAL	89.40	81.10	18.60	103.85
Effect of Pretrained Corpora				
WIKI	88.40	79.50	18.28	102.23
TOD	89.00	78.20	17.55	101.15
REDD	86.90	77.10	16.93	98.93
Effect of Pretrained Tasks				
w/o NTG	87.00	80.80	16.88	100.79
w/o OR	85.20	79.50	17.52	99.88
Effect of IE Tools				
OpenIE-Stanford	88.40	79.20	17.34	101.14
BART	87.50	72.20	16.67	96.52

Table 6: Ablation study on MultiWOZ2.0. There are three types of ablation study. The first is to analyze the effects of the pretrained data. The second is to validate the effects of the designed pretrained tasks. The last is to figure out the effects of IE tools. Results are significant ($p < 0.01$) comparing the OPAL model and BART model as the initialized TOD model.

respectively. The pretrained models WIKI and TOD still outperform the original BART by a large margin. It indicates the efficiency of the proposed ontology-aware pretraining method. Especially, the pretrained model WIKI that does not see any TOD data at the pretraining phase can result in competitive performance with the NCM_L . Compared with NCM_B with a similar parameter scale to our model, WIKI has apparent advantages on both end-to-end TOD datasets. The WIKI has the better performance with TOD. We know that WIKI suffers from the unseen TOD data and TOD suffers from the scale of the pretraining data. Our proposed OPAL adopts a two-stage pretraining method to solve the above problem, a classic example of ‘‘one plus one greater than two’’. The two-stage pretrained model OPAL outperforms the separated one with a 1.65 and 2.70 upper combined score on MultiWOZ2.0. This indicates that the ontology-aware contextual text corpus and ontology-aware TOD data are complementary.

To further compare Wikipedia to the Reddit corpus, we also use the same scale of Reddit data to conduct the Phase-1 pretraining, named REDD. WIKI is ahead of REDD in all the automatic metrics (BLEU and task-completion). To deeply analyze the effect factor, we calculate the occupation rate of the extracted triples that contained the pronouns as subject or object. As shown in Figure 5, 31.0% of triples in Reddit data contain pronouns. The highest frequency of pronouns

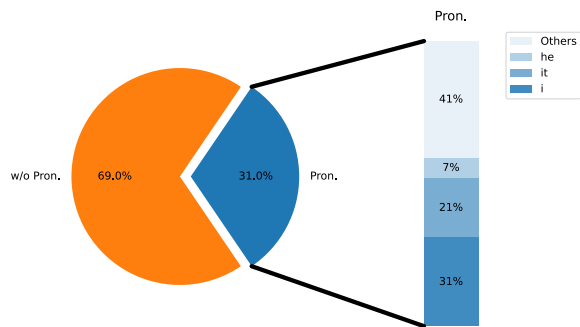


Figure 5: The occupation rate of the extracted triples that contained pronouns as subject or object in the Reddit corpus with OpenIE6.

is ‘‘i’’, which occupies 31%. Only 0.7% triples contain pronouns in Wikipedia. In the TOD, the domains and slot values in the dialogue states are specific entities, which are not pronouns. The meaningless pronouns increase the gap between the pretraining model and the TOD model. The co-reference and information ellipsis in Reddit seriously hurt the performance of the external information extraction tool. It is the main reason that we choose Wikipedia as the pretraining corpus.

We also evaluate the effects of the pretrained tasks: ontology-like triple recovery (OR) and next-text generation (NTG). We directly remove the extracted triples in the input in ‘‘w/o OR’’ study. The ‘‘w/o NTG’’ means that the model only needs to recover the masked triples. The results show that the OR task and NTG task benefit the task completion and the contextual consistency, respectively. In the complex dialogue domain, the single-task pretrained methods cannot achieve comparable performance with OPAL. It indicates that the two designed tasks are both significant to reduce the gap between pretrained model and TOD model.

We further validate the effects of different OpenIE tools. In our main experiments, we use the latest neural-based IE tool OpenIE6. There is also a very popular rule-based IE tool OpenIE-Stanford. Compared with OpenIE-Stanford, neural-based OpenIE6 achieves promising performance improvement on well-studied IE benchmarks (Kolluru et al., 2020). As shown in Table 6, WIKI with OpenIE6 is also better than OpenIE-Stanford tool in all the metrics. However, the improvement of neural-based OpenIE6 is limited, which indicates that the proposed pretraining method is not sensitive about IE accuracy.




Figure 6: Resource-limited response generation results on MultiWOZ2.0. 1% (80 dialogues), 5% (400 dialogues), 10% (800 dialogues), and 20% (1600 dialogues) of training data are used to train each model.

5.2 Sample Efficiency

Under the different resource-limited settings, the proposed OPAL can get all the best performance in terms of task completion (Inform and Success), response naturalness (BLEU), and overall performance among the baselines, as shown in Figure 6. It indicates the sample efficiency of the proposed ontology-aware pretraining method. When the training data is extremely limited (only 80 dialogues), TOD can improve overall performance by a large margin (absolute 3.2 point improvement) than WIKI. This improvement comes from the task completion ability, which indicates the TOD data can increase the generalization of the pretrained model for end-to-end TOD tasks. With the training data increase, WIKI pretrained on the large-scale context text data has the larger performance gain than TOD. When the number of the training data reaches 1600 dialogues, WIKI obtains absolute 4.8 point gains over TOD. It indicates that the scale of the pre-training data influences the growth potential of the pretrained model. On the other hand, TOD outperforms over the WIKI in three of four data limitation cases on task-completion ability. However, WIKI achieves better performance on fluent statement (revealed by *BLEU*). It indicates that WIKI benefits the task-completion ability and TOD facilitates fluency and context consistency.

5.3 Case Study

Our proposed pretrained model OPAL has improved the performance on task completion and contextual consistency over the original BART. As shown in Figure 7, we can see that the dialogue model fine-tuned from BART misses responding a request (address) to the user. Instead, our proposed OPAL accurately provides all the requested information to the user. As shown in

dialogue_id: SNG02115 turn:3	
User:	can you give me their phone number , address and price range , please ?
GT Response:	they are in the [value_price] price range . [value_phone] is their number . you can find them on [value_address]
BART:	it is in the [value_price] price range and the phone number is [value_phone] . [value_address] Missing → Success = 0
OPAL:	[value_name] is in the [value_price] price range and is located at:[value_address] . the phone number is [value_phone] → Success = 1

Figure 7: Third dialogue turn in the dialogue session SNG02115 from MultiWOZ2.0 development set. The oracle response is represented as GT Response. BART and OPAL means that the responses are generated by the corresponding models.

dialogue_id: SNG0921 turn: 1-2	
turn-1	User: hello ! i need a place to bed down for the night that offers free parking and wifi . can you help me ?
	GT Response: i can definitely help with that . do you have a preference in location ?
	BART: i sure can . first i 'll need to know what day you 're arriving and the number of nights you 'll be staying .
	OPAL: i can definitely help with that . do you have a preference in location ?
turn-2	User: i would like a moderate priced hotel in any part of town please .
	GT Response: there are [value_stars] [value_type] in the [value_area] area with [value_price] price -s . [value_name] and [value_name] . both are 2 stars , would you me to book you for either 1 ?
	BART: that narrows it down to [value_choice] options . is there a star rating you prefer ? [value_name] Missing → Inform = 0
	OPAL: i have a [value_name] in the [value_area] . it is a [value_price] [value_type] . would you like to book a room ? Inform = 1

Figure 8: The first two dialogue turns in the dialogue session SNG921 from the MultiWOZ2.0 development set. The oracle response is represented as GT Response. BART and OPAL means that the responses are generated by the corresponding models.

Figure 8, at the first turn, we can see that our proposed OPAL can provide the more similar response as the oracle than BART. It indicates that OPAL has the better performance on the response

prediction. At the second turn, the dialogue system needs to provide the correct entity to the user. The original BART model chooses to miss it. Our proposed OPAL recommends an entity to the user in time. Compared with the original BART, the proposed OPAL has a obvious advantage in modeling the task-oriented dialogue, which not only generates the precise response but also completes the dialogue task successfully. This performance improvement comes from the two-stage ontology-aware pretraining method on the large-scale contextual text with the hand-crafted ontology-like triples and the small task-oriented dialogue data with given ontology.

6 Related Work

End-to-End TOD Systems Early studies for end-to-end task-oriented dialogue systems either design a neural network-based model or propose a reinforcement learning method to use the reward signal to update the whole system. In these systems, the modules in the pipeline TOD system still exist and need their separated annotation. These systems usually can get promising performance on one specific task but have poor transferability. With the emergence of the multi-domain TOD benchmark, like MultiWOZ, the generative DST method has replaced the classification method as the mainstream over recent years due to its better generalization ability. It encourages formulating the end-to-end TOD as a text-to-text task. Lei et al. (2018) propose a two-stage CopyNet to generate the dialogue state and response jointly with a single seq2seq architecture. Zhang et al. (2020b) design a data augmentation method to increase the response diversity. The dialogue state, dialogue act, and the response are generated with a shared encoder and the different decoders. Note that our proposed model does not use the annotated dialogue acts. Recently, some work (Hosseini-Asl et al., 2020; Peng et al., 2020a; Lin et al., 2020; Yang et al., 2021) directly leverages the pretrained language models (like GPT-2 and BART) as the end-to-end TOD model in a unified way. Liu et al. (2021) propose a Transformer-based noisy channel method to model the response prior and use the Reddit data and TOD data to warm up the TOD model. Most recently, Su et al. (2021) formulate all the end-to-end TOD tasks as the unified generation tasks, which learns in a multitask learning man-

ner. He et al. (2021) propose a semi-supervised method to explicitly learn dialogue policy from limited labeled dialogues. Our proposed pretrained method is compatible with these end-to-end TOD training strategies.

Self-supervised Learning for Dialogue System

Recent advances in supervised learning have witnessed the success of the pretrained language models (PLMs) on language understanding and generation tasks. Since the large-scale comment data in Reddit can be regarded as a kind of chit-chat dialogue, the self-supervised methods have been used in the chit-chat systems first. DialoGPT (Zhang et al., 2020c) adapts the pretrained GPT-2 in the large-scale dialogue data. PLATO (Bao et al., 2020) proposes a discrete latent variable pretraining method to solve the one-to-many problem of the dialogue system. Meena (Adiwardana et al., 2020) pretrains a large-scale model with the dialogue data and demonstrates its conversation ability. SC-GPT (Peng et al., 2020b) uses a pretrained language model to convert a dialog act to a natural language response. For the task-oriented dialogue, the large-scale domain-specific dialogue data is inaccessible. The TOD models (Jiang et al., 2020; Wu et al., 2020; Yu et al., 2020) are usually pretrained on the chit-chat dialogues (Reddit) first and then fine-tuned on the smaller released or synthetic TOD data. Different from the above PLMs, we pretrain the TOD model directly with the large-scale contextual text. We extract relation triples of the contextual text as the grounded ontology-like knowledge and design adaptive self-supervised learning tasks for the end-to-end TOD.

Knowledge-grounded PLMs Recently, there is an important branch of PLM to study how to integrate the knowledge into the PLM. ERNIE (Zhang et al., 2019) utilizes the external knowledge graph to recognize the type of the mentioned entity. There is a entity type embedding layer as one of input representation. To enhance the knowledge-related representation, they improve the mask mechanism by masking a whole entity directly. Similarly, Rosset et al. (2020) proposes an knowledge-aware language model (KALM), which is decoder-only Transformer-based architecture, like GPT. KALM proposes an entity tokenizer to directly segment popular entities as a single token. Some fields, like medicine, include

considerable proprietary information, and it is crucial to integrate the proprietary knowledge into the pretrained model. SMedBERT (Zhang et al., 2021) incorporates deep structured semantics knowledge from neighbors of linked-entity. In this paper, we aim to utilize the external tool OpenIE6 to produce lots of TOD-like data to bridge the gap between pretrained task and end-to-end TOD system. The proposed ontology-like triple recovery task only masks the object values in the extracted triples, rather than randomly masking mentioned entities.

7 Conclusion and Future Work

In this paper, we propose an ontology-aware pretraining method for modeling the end-to-end task-oriented dialogue. The scale of the existing task-oriented dialogue data is far from the need for the pretrained model. Thus, we leverage the external tool OpenIE6 in extracting the ontology-like knowledge of the large-scale contextual texts. To bridge the gap between the pretrained and end-to-end TOD models, we design two adaptive self-supervised learning tasks: ontology-like triple recovery and next-text generation. The pretraining process is divided into two phases, where phase-1 pretrains on the large-scale ontology-aware contextual texts and phase-2 pretrains on the ontology-aware TOD data. Our proposed OPAL achieves excellent performance on the end-to-end TOD tasks and dialogue state tracking tasks. In the future, we will evaluate the effect of the different ontology-building methods.

Acknowledgments

We would like to thank the ACL team and four anonymous reviewers for their insightful comments. This work has been supported by China NSFC Projects (No.62120106006, No.62106142, and No.92048205), Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102), and CCF-Tencent Open Fund and Startup Fund for Youngman Research at SJTU (SFYR at SJTU).

References

Daniel Adiwardana, Minh-Thang Luong, David R. So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, et al. 2020.

Towards a human-like open-domain chatbot. *arXiv preprint arXiv:2001.09977*.

Gabor Angeli, Melvin Jose Johnson Premkumar, and Christopher D. Manning. 2015. Leveraging linguistic structure for open domain information extraction. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 344–354. <https://doi.org/10.3115/v1/P15-1034>

Vevake Balaraman and Bernardo Magnini. 2019. Scalable neural dialogue state tracking. In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 830–837. IEEE.

Siqi Bao, Huang He, Fan Wang, Hua Wu, and Haifeng Wang. 2020. PLATO: Pre-trained dialogue generation model with discrete latent variable. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 85–96. <https://doi.org/10.18653/v1/2020.acl-main.9>

Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Ultes Stefan, Ramadan Osman, and Milica Gašić. 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 (EMNLP)*.

Bill Byrne, Karthik Krishnamoorthi, Chinnadhurai Sankar, Arvind Neelakantan, Ben Goodrich, Daniel Duckworth, Semih Yavuz, Amit Dubey, Kyu-Young Kim, and Andy Cedilnik. 2019. Taskmaster-1: Toward a realistic and diverse dialog dataset. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4516–4525. <https://doi.org/10.18653/v1/D19-1459>

Lu Chen, Cheng Chang, Zhi Chen, Bowen Tan, Milica Gašić, and Kai Yu. 2018. Policy adaptation for deep reinforcement learning-based dialogue management. In *2018 IEEE International Conference on Acoustics, Speech and*

- Signal Processing (ICASSP)*, pages 6074–6078. IEEE.
- Lu Chen, Zhi Chen, Bowen Tan, Sishan Long, Milica Gašić, and Kai Yu. 2019. AgentGraph: Toward universal dialogue management with structured deep reinforcement learning. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(9):1378–1391. <https://doi.org/10.1109/TASLP.2019.2919872>
- Lu Chen, Boer Lv, Chi Wang, Su Zhu, Bowen Tan, and Kai Yu. 2020a. Schema-guided multi-domain dialogue state tracking with graph attention neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7521–7528.
- Zhi Chen, Lu Chen, Xiaoyuan Liu, and Kai Yu. 2020b. Distributed structured actor-critic reinforcement learning for universal dialogue management. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:2400–2411. <https://doi.org/10.1109/TASLP.2020.3013392>
- Zhi Chen, Lu Chen, Zihan Xu, Yanbin Zhao, Su Zhu, and Kai Yu. 2020c. Credit: Coarse-to-fine sequence generation for dialogue state tracking. *arXiv preprint arXiv:2009.10435*.
- Yinpei Dai, Hangyu Li, Yongbin Li, Jian Sun, Fei Huang, Luo Si, and Xiaodan Zhu. 2021. Preview, attend and review: Schema-aware curriculum learning for multi-domain dialog state tracking. *arXiv preprint arXiv:2106.00291*.
- Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyang Gao, and Dilek Hakkani-Tur. 2019. MultiWOZ 2.1: Multi-domain dialogue state corrections and state tracking baselines. *arXiv preprint arXiv:1907.01669*.
- Mihail Eric, Lakshmi Krishnan, Francois Charette, and Christopher D. Manning. 2017. Key-value retrieval networks for task-oriented dialogue. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 37–49.
- Rahul Goel, Shachi Paul, and Dilek Hakkani-Tur. 2019. Hyst: A hybrid approach for flexible and accurate dialogue state tracking. *arXiv preprint arXiv:1907.00883*. <https://doi.org/10.21437/Interspeech.2019-1863>
- Donghoon Ham, Jeong-Gwan Lee, Youngsoo Jang, and Kee-Eung Kim. 2020. End-to-end neural pipeline for goal-oriented dialogue systems using GPT-2. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 583–592.
- Wanwei He, Yinpei Dai, Yinhe Zheng, Yuchuan Wu, Zheng Cao, Dermot Liu, Peng Jiang, Min Yang, Fei Huang, Luo Si, Jian Sun, and Yongbin Li. 2021. GALAXY: A generative pretrained model for task-oriented dialog with semi-supervised learning and explicit policy injection. *arXiv preprint arXiv:2111.14592*.
- Michael Heck, Carel van Niekerk, Nurul Lubis, Christian Geisshauser, Hsien-Chin Lin, Marco Moresi, and Milica Gasic. 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*, pages 35–44.
- Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. 2020. A simple language model for task-oriented dialogue. *arXiv preprint arXiv:2005.00796*.
- Mohit Iyyer, Wen-tau Yih, and Ming-Wei Chang. 2017. Search-based neural structured learning for sequential question answering. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1821–1831. <https://doi.org/10.18653/v1/P17-1167>
- Zi-Hang Jiang, Weihao Yu, Daquan Zhou, Yunpeng Chen, Jiashi Feng, and Shuicheng Yan. 2020. ConvBERT: Improving BERT with span-based dynamic convolution. *Advances in Neural Information Processing Systems*, 33.
- Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. BERT: Pretraining of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Sungdong Kim, Sohee Yang, Gyuwan Kim, and Sang-Woo Lee. 2020. Efficient dialogue state tracking by selectively overwriting memory. In *Proceedings of the 58th Annual Meeting of*

- the Association for Computational Linguistics*, pages 567–582.
- Philipp Koehn. 2004. Statistical significance tests for machine translation evaluation. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. Association for Computational Linguistics.
- Keshav Kolluru, Vaibhav Adlakha, Samarth Aggarwal, Soumen Chakrabarti. 2020. OpenIE6: Iterative grid labeling and coordination analysis for open information extraction. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3748–3761. <https://doi.org/10.18653/v1/2020.emnlp-main.306>
- Tuan Manh Lai, Quan Hung Tran, Trung Bui, and Daisuke Kihara. 2020. A simple but effective BERT model for dialog state tracking on resource-limited systems. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8034–8038. IEEE.
- Hung Le, Richard Socher, and Steven C. H. Hoi. 2020. Non-autoregressive dialog state tracking. In *International Conference on Learning Representations*.
- Hwaran Lee, Jinsik Lee, and Tae-Yoon Kim. 2019a. SUMBT: Slot-utterance matching for universal and scalable belief tracking. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5478–5483.
- Sungjin Lee, Qi Zhu, Ryuichi Takanobu, Zheng Zhang, Yaoqin Zhang, Xiang Li, Jinchao Li, Baolin Peng, Xiujun Li, Minlie Huang, and Jianfeng Gao. 2019b. ConvLab: Multi-domain end-to-end dialog system platform. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 64–69.
- Wenqiang Lei, Xisen Jin, Min-Yen Kan, Zhaochun Ren, Xiangnan He, and Dawei Yin. 2018. Sequicity: Simplifying task-oriented dialogue systems with single sequence-to-sequence architectures. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1437–1447.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pretraining for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880. <https://doi.org/10.18653/v1/2020.acl-main.703>
- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016. Deep reinforcement learning for dialogue generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202.
- Shiyang Li, Semih Yavuz, Kazuma Hashimoto, Jia Li, Tong Niu, Nazneen Rajani, Xifeng Yan, Yingbo Zhou, and Caiming Xiong. 2020. CoCo: Controllable counterfactuals for evaluating dialogue state trackers. In *International Conference on Learning Representations*.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, and Pascale Fung. 2020. MinTI: Minimalist transfer learning for task-oriented dialogue systems. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3391–3405.
- Qi Liu, Lei Yu, Laura Rimell, and Phil Blunsom. 2021. Pretraining the noisy channel model for task-oriented dialogue. *arXiv preprint arXiv:2103.10518*. <https://doi.org/10.1162/tacl.a.00390>
- Shikib Mehri, Evgeniia Razumovskaia, Tiancheng Zhao, and Maxine Eskenazi. 2019. Pretraining methods for dialog context representation learning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3836–3845.
- Nikola Mrkšić, Diarmuid Ó Séaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve Young. 2017. Neural belief tracker: Data-driven dialogue state tracking. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1777–1788. <https://doi.org/10.18653/v1/P17-1163>

- Elnaz Nouri and Ehsan Hosseini-Asl. 2018. Toward scalable neural dialogue state tracking. In *NeurIPS 2018, 2nd Conversational AI workshop*.
- Baolin Peng, Chunyuan Li, Jinchao Li, Shahin Shayandeh, Lars Liden, and Jianfeng Gao. 2020a. SOLOIST: Few-shot task-oriented dialog with a single pre-trained auto-regressive model. *arXiv e-prints*, arXiv–2005.
- Baolin Peng, Chenguang Zhu, Chunyuan Li, Xiujun Li, Jinchao Li, Michael Zeng, and Jianfeng Gao. 2020b. Few-shot natural language generation for task-oriented dialog. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 172–182. <https://doi.org/10.18653/v1/2020.findings-emnlp.17>
- Shuke Peng, Xinjing Huang, Zehao Lin, Feng Ji, Haiqing Chen, and Yin Zhang. 2019. Teacher-student framework enhanced multi-domain dialogue generation. *arXiv preprint arXiv:1908.07137*.
- Chris Quirk, Raymond Mooney, and Michel Galley. 2015. Language to code: Learning semantic parsers for if-this-then-that recipes. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 878–888. <https://doi.org/10.3115/v1/P15-1085>
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2020. Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8689–8696. <https://doi.org/10.1609/aaai.v34i05.6394>
- Liliang Ren, Jianmo Ni, and Julian McAuley. 2019. Scalable and accurate dialogue state tracking via hierarchical sequence generation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1876–1885.
- Liliang Ren, Kaige Xie, Lu Chen, and Kai Yu. 2018. Towards universal dialogue state tracking. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2780–2786.
- Corby Rosset, Chenyan Xiong, Minh Phan, Xia Song, Paul Bennett, and Saurabh Tiwary. 2020. Knowledge-aware language model pretraining. *arXiv preprint arXiv:2007.00655*.
- Bishal Santra, Potnuru Anusha, and Pawan Goyal. 2021. Hierarchical transformer for task oriented dialog systems. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5649–5658. <https://doi.org/10.18653/v1/2021.naacl-main.449>
- Yixuan Su, Lei Shu, Elman Mansimov, Arshit Gupta, Deng Cai, Yi-An Lai, and Yi Zhang. 2021. Multi-task pre-training for plug-and-play task-oriented dialogue system. *arXiv preprint arXiv:2109.14739*.
- Stefan Ultes, Lina M. Rojas Barahona, Pei-Hao Su, David Vandyke, Dongho Kim, Inigo Casanueva, Paweł Budzianowski, Nikola Mrkšić, Tsung-Hsien Wen, Milica Gasic, and Steve Young. 2017. Pydial: A multi-domain statistical dialogue system toolkit. In *Proceedings of ACL 2017, System Demonstrations*, pages 73–78. <https://doi.org/10.18653/v1/P17-4013>
- Michael Völske, Martin Potthast, Shahbaz Syed, and Benno Stein. 2017. Tl; dr: Mining Reddit to learn automatic summarization. In *Proceedings of the Workshop on New Frontiers in Summarization*, pages 59–63. <https://doi.org/10.18653/v1/W17-4508>
- Gellért Weisz, Paweł Budzianowski, Pei-Hao Su, and Milica Gašić. 2018. Sample efficient deep reinforcement learning for dialogue systems

- with large action spaces. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 26(11):2083–2097. <https://doi.org/10.1109/TASLP.2018.2851664>
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrkšić, Lina M. Rojas Barahona, Pei-Hao Su, Stefan Ultes, David Vandyke, and Steve Young. 2016. Conditional generation and snapshot learning in neural dialogue systems. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2153–2162.
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015. Semantically conditioned LSTM-based natural language generation for spoken dialogue systems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1711–1721.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-demos.6>
- Chien-Sheng Wu, Steven CH Hoi, Richard Socher, and Caiming Xiong. 2020. TOD-BERT: Pretrained natural language understanding for task-oriented dialogue. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 917–929.
- Chien-Sheng Wu, Andrea Madotto, Ehsan Hosseini-Asl, Caiming Xiong, Richard Socher, and Pascale Fung. 2019. Transferable multi-domain state generator for task-oriented dialogue systems. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 808–819.
- Zihan Xu, Zhi Chen, Lu Chen, Su Zhu, and Kai Yu. 2020. Memory attention neural network for multi-domain dialogue state tracking. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 41–52. Springer.
- Yunyi Yang, Yunhao Li, and Xiaojun Quan. 2021. UBAR: Towards fully end-to-end task-oriented dialog system with GPT-2. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14230–14238. <https://doi.org/10.1609/aaai.v35i16.17674>
- Tao Yu, Rui Zhang, Alex Polozov, Christopher Meek, and Ahmed Hassan Awadallah. 2020. SCoRe: Pre-training for context representation in conversational semantic parsing. In *International Conference on Learning Representations*.
- Jianguo Zhang, Kazuma Hashimoto, Chien-Sheng Wu, Yao Wang, S Yu Philip, Richard Socher, and Caiming Xiong. 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*, pages 154–167.
- Taolin Zhang, Zerui Cai, Chengyu Wang, Minghui Qiu, Bite Yang, and Xiaofeng He. 2021. SMedBERT: A knowledge-enhanced pre-trained language model with structured semantics for medical text mining. 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)*, pages 5882–5893. <https://doi.org/10.18653/v1/2021.acl-long.457>
- Yichi Zhang, Zhijian Ou, and Zhou Yu. 2020b. Task-oriented dialog systems that consider multiple appropriate responses under the same context. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9604–9611.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and William B. Dolan. 2020c. DIALOGPT: Large-scale generative pre-training for conversational response

- generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278. <https://doi.org/10.18653/v1/2020.acl-demos.30>
- Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. ERNIE: Enhanced language representation with informative entities. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1441–1451. <https://doi.org/10.18653/v1/P19-1139>
- Tiancheng Zhao and Maxine Eskenazi. 2018. Zero-shot dialog generation with cross-domain latent actions. In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*, pages 1–10.
- Tiancheng Zhao, Kaige Xie, and Maxine Eskenazi. 2019. Rethinking action spaces for reinforcement learning in end-to-end dialog agents with latent variable models. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1208–1218.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 654–664.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2018. Global-locally self-attentive encoder for dialogue state tracking. In *ACL*. <https://doi.org/10.18653/v1/P18-1135>
- Jingyao Zhou, Haipang Wu, Zehao Lin, Guodun Li, and Yin Zhang. 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*, pages 228–238.
- Li Zhou and Kevin Small. 2019. Multi-domain dialogue state tracking as dynamic knowledge graph enhanced question answering. *arXiv preprint arXiv:1911.06192*.