

# TOD-BERT: Pre-trained Natural Language Understanding for Task-Oriented Dialogue

Chien-Sheng Wu, Steven Hoi, Richard Socher, and Caiming Xiong

Salesforce Research

[wu.jason, shoi, rsocher, cxiong]@salesforce.com

## Abstract

The underlying difference of linguistic patterns between general text and task-oriented dialogue makes existing pre-trained language models less useful in practice. In this work, we unify nine human-human and multi-turn task-oriented dialogue datasets for language modeling. To better model dialogue behavior during pre-training, we incorporate user and system tokens into the masked language modeling. We propose a contrastive objective function to simulate the response selection task. Our pre-trained task-oriented dialogue BERT (TOD-BERT) outperforms strong baselines like BERT on four downstream task-oriented dialogue applications, including intention recognition, dialogue state tracking, dialogue act prediction, and response selection. We also show that TOD-BERT has a stronger few-shot ability that can mitigate the data scarcity problem for task-oriented dialogue.

## 1 Introduction

Pre-trained models with self-attention encoder architectures (Devlin et al., 2018; Liu et al., 2019) have been commonly used in many NLP applications. Such models are self-supervised based on a massive scale of general text corpora, such as English Wikipedia or books (Zhu et al., 2015). By further fine-tuning these representations, breakthroughs have been continuously reported for various downstream tasks, especially natural language understanding.

However, previous work (Rashkin et al., 2018; Wolf et al., 2019) shows that there are some deficiencies in the performance to apply fine-tuning on conversational corpora directly. One possible reason could be the intrinsic difference of linguistic patterns between human conversations and writing text, resulting in a large gap of data distributions (Bao et al., 2019). Therefore, pre-training dialogue

language models using chit-chat corpora from social media, such as Twitter or Reddit, has been recently investigated, especially for dialogue response generation (Zhang et al., 2019) and retrieval (Henderson et al., 2019b). Although these open-domain dialogues are diverse and easy-to-get, they are usually short, noisy, and without specific chatting goals.

On the other hand, a task-oriented dialogue has explicit goals (e.g. restaurant reservation or ticket booking) and many conversational interactions. But each dataset is usually small and scattered because obtaining and labeling such data is time-consuming. Moreover, a task-oriented dialogue has explicit user and system behaviors where a user has his/her goal, and a system has its belief and database information, which makes the language understanding component and dialogue policy learning more important than those chit-chat scenarios.

This paper aims to prove this hypothesis: self-supervised language model pre-training using task-oriented corpora can learn better representations than existing pre-trained models for task-oriented downstream tasks. We emphasize that what we care about the most is not whether our pre-trained model can achieve state-of-the-art results on each downstream task since most of the current best models are built on top of pre-trained models, and ours can easily replace them. We avoid adding too many additional components on top of the pre-training architecture when fine-tuning in our experiments.

We collect and combine nine human-human and multi-turn task-oriented dialogue corpora to train a task-oriented dialogue BERT (TOD-BERT). In total, there are around 100k dialogues with 1.4M utterances across over 60 different domains. Like BERT (Devlin et al., 2018), TOD-BERT is formulated as a masked language model and uses a deep bidirectional Transformer (Vaswani et al., 2017)

encoder as its model architecture. Unlike BERT, TOD-BERT incorporates two special tokens for user and system to model the corresponding dialogue behavior. A contrastive objective function of response selection task is combined during pre-training stage to capture response similarity. We select BERT because it is the most widely used model in NLP research recently, and our unified datasets can be easily applied to pre-train any existing language models.

We test TOD-BERT on task-oriented dialogue systems on four core downstream tasks, including intention recognition, dialogue state tracking, dialogue act prediction, and response selection. What we observe is: TOD-BERT outperforms BERT and other strong baselines such as GPT-2 (Radford et al., 2019) and DialoGPT (Zhang et al., 2019) on all the selected downstream tasks, which further confirms its effectiveness for improving dialogue language understanding. We find that response contrastive learning is beneficial, but it is currently overlooked not well-investigated in dialogue pre-training research. More importantly, TOD-BERT has a stronger few-shot ability than BERT on each task, suggesting that it can reduce the need for expensive human-annotated labels. TOD-BERT can be easily leveraged and adapted to a new task-oriented dialogue dataset. Our source code and data processing are released to facilitate future research on pre-training and fine-tuning of task-oriented dialogue <sup>1</sup>.

## 2 Related Work

**General Pre-trained Language Models**, which are trained on massive general text such as Wikipedia and BookCorpus, can be roughly divided into two categories: uni-directional or bi-directional attention mechanisms. GPT (Radford et al., 2018) and GPT-2 (Radford et al., 2019) are representatives of uni-directional language models using a Transformer decoder, where the objective is to maximize left-to-right generation likelihood. These models are commonly applied in natural language generation tasks. On the other hand, BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019), and their variances are pre-trained using a Transformer encoder with bi-directional token prediction. These models are usually evaluated on classification tasks such as GLUE benchmark (Wang et al., 2018) or span-based question answering tasks (Ra-

ipurkar et al., 2016).

Some language models can support both uni-directional and bi-directional attention, such as UniLM (Dong et al., 2019). Conditional language model pre-training is also proposed. For example, CTRL (Keskar et al., 2019) is a conditional Transformer model, trained to condition on control codes that govern style, content, and task-specific behavior. Recently, multi-task language model pre-training with unified sequence-to-sequence generation is proposed. Text-to-text Transformer (T5) (Raffel et al., 2019) unifies multiple text modeling tasks and achieves the promising results in various NLP benchmarks.

**Dialogue Pre-trained Language Models** are mostly trained on open-domain conversational data from Reddit or Twitter for dialogue response generation. Transfertransfo (Wolf et al., 2019) achieves good performance on ConvAI-2 dialogue competition using GPT-2. DialoGPT (Zhang et al., 2019) is an extension of GPT-2 that is pre-trained on Reddit data for open-domain response generation. ConveRT (Henderson et al., 2019a) pre-trained a dual transformer encoder for response selection task on large-scale Reddit (input, response) pairs. PLATO (Bao et al., 2019) uses both Twitter and Reddit data to pre-train a dialogue generation model with discrete latent variables. All of them are designed to cope with the response generation task for open-domain chatbots.

Pretraining for task-oriented dialogues, on the other hand, has few related works. Budzianowski and Vulić (2019) first apply the GPT-2 model to train on response generation task, which takes system belief, database result, and last dialogue turn as input to predict next system responses. It only uses one dataset to train its model because few public datasets have database information available. Henderson et al. (2019b) pre-trained a response selection model for task-oriented dialogues. They first pre-train on Reddit corpora and then fine-tune on target dialogue domains, but their training and fine-tuning code is not released. Peng et al. (2020) focus on the natural language generation (NLG) task, which assumes dialogue acts and slot-tagging results are given to generate a natural language response. Pre-training on a set of annotated NLG corpora can improve conditional generation quality using a GPT-2 model.

<sup>1</sup>[github.com/jasonwu0731/ToD-BERT](https://github.com/jasonwu0731/ToD-BERT)

Name	# Dialogue	# Utterance	Avg. Turn	# Domain
MetaLWOZ (Lee et al., 2019)	37,884	432,036	11.4	47
Schema (Rastogi et al., 2019)	22,825	463,284	20.3	17
Taskmaster (Byrne et al., 2019)	13,215	303,066	22.9	6
MWOZ (Budzianowski et al., 2018)	10,420	71,410	6.9	7
MSR-E2E (Li et al., 2018)	10,087	74,686	7.4	3
SMD (Eric and Manning, 2017)	3,031	15,928	5.3	3
Frames (Asri et al., 2017)	1,369	19,986	14.6	3
WOZ (Mrkšić et al., 2016)	1,200	5,012	4.2	1
CamRest676 (Wen et al., 2016)	676	2,744	4.1	1

Table 1: Data statistics for task-oriented dialogue datasets.

### 3 Method

This section discusses each dataset used in our task-oriented pre-training and how we process the data. Then we introduce the selected pre-training base model and its objective functions.

#### 3.1 Datasets

We collect nine different task-oriented datasets which are English, human-human and multi-turn. In total, there are 100,707 dialogues, which contain 1,388,152 utterances over 60 domains. Dataset statistics is shown in Table 1.

- **MetaLWOZ** (Lee et al., 2019): Meta-Learning Wizard-of-Oz is a dataset designed to help develop models capable of predicting user responses in unseen domains. This large dataset was created by crowdsourcing 37,884 goal-oriented dialogs, covering 227 tasks in 47 domains. The MetaLWOZ dataset is used as the fast adaptation task for DSTC8 (Kim et al., 2019) dialogue competition.
- **Schema** (Rastogi et al., 2019): Schema-guided dialogue has 22,825 dialogues and provides a challenging testbed for several tasks, in particular, dialogue state tracking. Each schema is a set of tracking slots, and each domain could have multiple possible schemas. This allows a single dialogue system to support many services and facilitates the simple integration of new services without requiring much training data. The Schema dataset is used as the dialogue state tracking task for DSTC8 (Kim et al., 2019) dialogue competition.
- **Taskmaster** (Byrne et al., 2019): This dataset includes 13,215 dialogues comprising six do-

main, including 5,507 spoken and 7,708 written dialogs created with two distinct procedures. One is a two-person Wizard of Oz approach that one person acts like a robot, and the other is a self-dialogue approach in which crowdsourced workers wrote the entire dialog themselves. It has 22.9 average conversational turns in a single dialogue, which is the longest among all task-oriented datasets listed.

- **MWOZ** (Budzianowski et al., 2018): Multi-Domain Wizard-of-Oz dataset contains 10,420 dialogues over seven domains, and it has multiple domains in a single dialogue. It has a detailed description of the data collection procedure, user goal, system act, and dialogue state labels. Different from most of the existing corpora, it also provides full database information.
- **MSR-E2E** (Li et al., 2018): Microsoft end-to-end dialogue challenge has 10,087 dialogues in three domains, movie-ticket booking, restaurant reservation, and taxi booking. It also includes an experiment platform with built-in simulators in each domain.
- **SMD** (Eric and Manning, 2017): Stanford multi-domain dialogue is an in-car personal assistant dataset, comprising 3,301 dialogues and three domains: calendar scheduling, weather information retrieval, and point-of-interest navigation. It is designed to smoothly interface with knowledge bases, where a knowledge snippet is attached with each dialogue as a piece of simplified database information.
- **Frames** (Asri et al., 2017): This dataset comprises 1,369 human-human dialogues with an average of 14.6 turns per dialogue, where users

are given some constraints to book a trip and assistants who search a database to find appropriate trips. Unlike other datasets, it has labels to keep track of different semantic frames, which is the decision-making behavior of users throughout each dialogue.

- **WOZ** (Mrkšić et al., 2016) and **Cam-Rest676** (Wen et al., 2016): These two corpora use the same data collection procedure and same ontology from DSTC2 (Henderson et al., 2014). They are one of the first task-oriented dialogue datasets that use Wizard of Oz style with text input instead of speech input, which improves the model’s capacity for the semantic understanding instead of its robustness to automatic speech recognition errors.

### 3.2 TOD-BERT

We train our TOD-BERT based on BERT architecture using two loss functions: masked language modeling (MLM) loss and response contrastive loss (RCL). Note that the datasets we used can be used to pre-train any existing language model architecture, and here we select BERT because it is the most widely used model in NLP research. We use the BERT-base uncased model, which is a transformer self-attention encoder (Vaswani et al., 2017) with 12 layers and 12 attention heads with its hidden size  $d_B = 768$ .

To capture speaker information and the underlying interaction behavior in dialogue, we add two special tokens, [USR] and [SYS], to the byte-pair embeddings (Mrkšić et al., 2016). We prefix the special token to each user utterance and system response, and concatenate all the utterances in the same dialogue into one flat sequence, as shown in Figure 1. For example, for a dialogue  $D = \{S_1, U_1, \dots, S_n, U_n\}$ , where  $n$  is the number of dialogue turns and each  $S_i$  or  $U_i$  contains a sequence of words, the input of the pre-training model is processed as “[SYS]  $S_1$  [USR]  $U_1 \dots$ ” with standard positional embeddings and segmentation embeddings.

**Masked language modeling** is a common pre-training strategy for BERT-like architectures, in which a random sample of tokens in the input sequence is selected and replaced with the special token [MASK]. The MLM loss function is the cross-entropy loss on predicting the masked tokens. In the original implementation, random masking and

replacement are performed once in the beginning and saved for the training duration. Here we conduct token masking dynamically during batch training. TOD-BERT is initialized from BERT, a good starting parameter set, then is further pre-trained on those task-oriented corpora. The MLM loss function is

$$L_{mlm} = - \sum_{m=1}^M \log P(x_m), \quad (1)$$

where  $M$  is the total number of masked tokens and  $P(x_m)$  is the predicted probability of the token  $x_m$  over the vocabulary size.

**Response contrastive loss** can also be used for dialogue language modeling since it does not require any additional human annotation. Pre-training with RCL can bring us several advantages: 1) we can learn a better representation for the [CLS] token, as it is essential for all the downstream tasks, and 2) we encourage the model to capture underlying dialogue sequential order, structure information, and response similarity.

Unlike the original next sentence prediction (NSP) objective in BERT pre-training, which concatenates two segments  $A$  and  $B$  to predict whether they are consecutive text with binary classification, we apply a dual-encoder approach (Henderson et al., 2019a) and simulate multiple negative samples. We first draw a batch of dialogues  $\{D_1, \dots, D_b\}$  and split each dialogue at a randomly selected turn  $t$ . For example,  $D_1$  will be separated into two segments, one is the context  $\{S_1^1, U_1^1, \dots, S_t^1, U_t^1\}$  and the other is the response  $\{S_{t+1}^1\}$ . We use TOD-BERT to encode all the contexts and their corresponding responses separately.

Afterwards, we have a context matrix  $C \in \mathbb{R}^{b \times d_B}$  and a response matrix  $R \in \mathbb{R}^{b \times d_B}$  by taking the output [CLS] representations from the  $b$  dialogues. We treat other responses in the same batch as randomly selected negative samples. The RCL objective function is

$$L_{rcl} = - \sum_{i=1}^b \log M_{i,i}, \quad (2)$$

$$M = \text{Softmax}(CR^T) \in \mathbb{R}^{b \times b}.$$

Increasing batch size to a certain amount can obtain better performance on downstream tasks, especially for the response selection. The Softmax function normalizes the vector per row. In our setting, increasing batch size also means changing the positive and negative ratio in the contrastive learning. Batch size is a hyper-parameter that may be






Figure 1: Dialogue pre-training based on Transformer encoder with user and system special tokens. Two objective functions are used: masked language modeling and response contrastive learning.

limited by hardware. We also try different negative sampling strategies during pre-training such as local sampling (Saeidi et al., 2017), but do not observe significant change compared to random sampling.

**Overall** pre-training loss function is the weighted-sum of  $L_{mlm}$  and  $L_{rcl}$ , and in our experiments, we simply sum them up. We gradually reduce the learning rate without a warm-up period. We train TOD-BERT with AdamW (Loshchilov and Hutter, 2017) optimizer with a dropout ratio of 0.1 on all layers and attention weights. GELU activation functions (Hendrycks and Gimpel, 2016) is used. Models are early-stopped using perplexity scores of a held-out development set, with mini-batches containing 32 sequences of maximum length 512 tokens. Experiments are conducted on two NVIDIA Tesla V100 GPUs.

## 4 Downstream Tasks

We care the most in this paper whether TOD-BERT, a pre-trained language model using aggregated task-oriented corpora, can show any advantage over BERT. Therefore, we avoid adding too many additional components on top of its architecture when fine-tuning on each downstream task. Also, we use the same architecture with a similar number of parameters for a fair comparison. All the model parameters are updated with a gradient clipping to 1.0 using the same hyper-parameters during fine-tuning. We select four crucial task-oriented downstream tasks to evaluate: intent recognition, dialogue state tracking, dialogue act prediction, and response selection. All of them are core components in modularized task-oriented systems (Wen et al., 2016).

**Intent recognition** task is a multi-class classification problem, where we input a sentence  $U$  and

models predict one single intent class over  $I$  possible intents.

$$P_{int} = \text{Softmax}(W_1(F(U))) \in \mathbb{R}^I, \quad (3)$$

where  $F$  is a pre-trained language model and we use its [CLS] embeddings as the output representation.  $W_1 \in \mathbb{R}^{I \times d_B}$  is a trainable linear mapping. The model is trained with cross-entropy loss between the predicted distributions  $P_{int}$  and the true intent labels.

**Dialogue state tracking** can be treated as a multi-class classification problem using a predefined ontology. Unlike intent, we use dialogue history  $X$  (a sequence of utterances) as input and a model predicts slot values for each (domain, slot) pair at each dialogue turn. Each corresponding value  $v_i^j$ , the  $i$ -th value for the  $j$ -th (domain, slot) pair, is passed into a pre-trained model and fixed its representation during training.

$$S_i^j = \text{Sim}(G_j(F(X)), F(v_i^j)) \in \mathbb{R}^1, \quad (4)$$

where  $\text{Sim}$  is the cosine similarity function, and  $S^j \in \mathbb{R}^{|v^j|}$  is the probability distribution of the  $j$ -th (domain, slot) pair over its possible values.  $G_j$  is the slot projection layer of the  $j$  slot, and the number of layers  $|G|$  is equal to the number of (domain, slot) pairs. The model is trained with cross-entropy loss summed over all the pairs.

**Dialogue act prediction** is a multi-label classification problem because a system response may contain multiple dialogue acts, e.g., request and inform at the same time. Model take dialogue history as input and predict a binary result for each possible dialogue act:

$$A = \text{Sigmoid}(W_2(F(X))) \in \mathbb{R}^N, \quad (5)$$

where  $W_2 \in \mathbb{R}^{d_B \times N}$  is a trainable linear mapping,  $N$  is the number of possible dialogue acts, and each value in  $A$  is between  $[0, 1]$  after a Sigmoid layer. The model is trained with binary cross-entropy loss and the  $i$ -th dialogue act is considered as a triggered dialogue act if  $A_i > 0.5$ .

**Response selection** is a ranking problem, aiming to retrieve the most relative system response from a candidate pool. We use a dual-encoder strategy (Henderson et al., 2019b) and compute similarity scores between source  $X$  and target  $Y$ ,

$$r_i = \text{Sim}(F(X), F(Y_i)) \in \mathbb{R}^1, \quad (6)$$

where  $Y_i$  is the  $i$ -th response candidate and  $r_i$  is its cosine similarity score. Source  $X$  can be truncated, and we limit the context lengths to the most recent 256 tokens in our experiments. We randomly sample several system responses from the corpus as negative samples. Although it may not be a true negative sample, it is common to train a ranker and evaluate its results (Henderson et al., 2019a).

## 5 Evaluation Datasets

We pick up several datasets, OOS, DSTC2, GSIM, and MWOZ, for downstream evaluation. The first three corpora are not included in the pre-trained task-oriented datasets. For MWOZ, to be fair, we do not include its test set dialogues during the pre-training stage. Details of each evaluation dataset are discussed in the following:

- **OOS** (Larson et al., 2019): The out-of-scope intent dataset is one of the largest annotated intent datasets, including 15,100/3,100/5,500 samples for the train, validation, and test sets, respectively. It covers 151 intent classes over ten domains, including 150 in-scope intent and one out-of-scope intent. The out-of-scope intent means that a user utterance that does not fall into any of the predefined intents. Each of the intents has 100 training samples.
- **DSTC2** (Henderson et al., 2014): DSTC2 is a human-machine task-oriented dataset that may include a certain system response noise. It has 1,612/506/1117 dialogues for train, validation, and test sets, respectively. We follow Paul et al. (2019) to map the original dialogue act labels to universal dialogue acts, which results in 9 different system dialogue acts.
- **GSIM** (Shah et al., 2018a): GSIM is a human-rewrote machine-machine task-oriented corpus, including 1500/469/1039 dialogues for the train, validation, and test sets, respectively. We combine its two domains, movie and restaurant domains, into one single corpus. It is collected by Machines Talking To Machines (M2M) (Shah et al., 2018b) approach, a functionality-driven process combining a dialogue self-play step and a crowdsourcing step. We map its dialogue act labels to universal dialogue acts (Paul et al., 2019), resulting in 6 different system dialogue acts.
- **MWOZ** (Budzianowski et al., 2018): MWOZ is the most common benchmark for task-oriented

dialogues, especially for dialogue state tracking. It has 8420/1000/1000 dialogues for train, validation, and test sets, respectively. Across seven different domains, in total, it has 30 (domain, slot) pairs that need to be tracked in the test set. We use its revised version MWOZ 2.1, which has the same dialogue transcripts but with cleaner state label annotation.

## 6 Results

For each downstream task, we first conduct the experiments using the whole dataset, and then we simulate the few-shot setting to show the strength of our TOD-BERT. We run at least three times with different random seeds for each few-shot experiment to reduce data sampling variance, and we report its mean and standard deviation for these limited data scenarios. We investigate two versions of TOD-BERT; one is TOD-BERT-mlm that only uses MLM loss during pre-training, and the other is TOD-BERT-jnt, which is jointly trained with the MLM and RCL objectives. We compare TOD-BERT with BERT and other baselines, including two other strong pre-training models GPT-2 (Radford et al., 2019) and DialoGPT (Zhang et al., 2019). For a GPT-based model, we use mean pooling of its hidden states as its output representation, which we found it is better than using only the last token.

### 6.1 Linear Probe

Before fine-tuning each pre-trained models, we first investigate their feature extraction ability by probing their output representations. Probing methods are proposed to determine what information is carried intrinsically by the learned embeddings (Tenney et al., 2019). We probe the output representation using one single-layer perceptron on top of a “fixed” pre-trained language model and only fine-tune that layer for a downstream task with the same hyper-parameters. Table 3 shows the probing results of domain classification on MWOZ, intent identification on OOS, and dialogue act prediction on MWOZ. TOD-BERT-jnt achieves the highest performance in this setting, suggesting its representation contains the most useful information.

### 6.2 Intent Recognition

TOD-BERT outperforms BERT and other strong baselines in one of the largest intent recognition

	Model	Acc (all)	Acc (in)	Acc (out)	Recall (out)
1-Shot	BERT	29.3% ± 3.4%	35.7% ± 4.1%	81.3% ± 0.4%	0.4% ± 0.3%
	TOD-BERT-mlm	38.9% ± 6.3%	47.4% ± 7.6%	81.6% ± 0.2%	<b>0.5%</b> ± 0.2%
	TOD-BERT-jnt	<b>42.5%</b> ± 0.1%	<b>52.0%</b> ± 0.1%	<b>81.7%</b> ± 0.1%	0.1% ± 0.1%
10-Shot	BERT	75.5% ± 1.1%	88.6% ± 1.1%	84.7% ± 0.3%	16.5% ± 1.7%
	TOD-BERT-mlm	76.6% ± 0.8%	90.5% ± 1.2%	84.3% ± 0.2%	14.0% ± 1.3%
	TOD-BERT-jnt	<b>77.3%</b> ± 0.5%	<b>91.0%</b> ± 0.5%	<b>84.5%</b> ± 0.4%	<b>15.3%</b> ± 2.1%
Full (100-Shot)	FastText*	-	89.0%	-	9.7%
	SVM*	-	91.0%	-	14.5%
	CNN*	-	91.2%	-	18.9%
	GPT2	83.0%	94.1%	87.7%	32.0%
	DialoGPT	83.9%	95.5%	87.6%	32.1%
	BERT	84.9%	95.8%	88.1%	35.6%
	TOD-BERT-mlm	85.9%	96.1%	89.5%	<b>46.3%</b>
	TOD-BERT-jnt	<b>86.6%</b>	<b>96.2%</b>	<b>89.9%</b>	43.6%

Table 2: Intent recognition results on the OOS dataset, one of the largest intent corpus. Models with \* are reported from Larson et al. (2019).

	Domain (acc)	Intent (acc)	Dialogue Act (F1-micro)
GPT2	63.5%	74.7%	85.7%
DialoGPT	63.0%	65.7%	84.2%
BERT	60.5%	71.1%	85.3%
TOD-BERT-mlm	63.9%	70.7%	83.5%
TOD-BERT-jnt	<b>68.7%</b>	<b>77.8%</b>	<b>86.2%</b>

Table 3: Probing results of different pre-trained language models using a single-layer perceptron.

datasets, as shown in Table 2. We evaluate accuracy on all the data, the in-domain intents only, and the out-of-scope intent only. Note that there are two ways to predict out-of-scope intent, one is to treat it as an additional class, and the other is to set a threshold for prediction confidence. Here we report the results of the first setting. TOD-BERT-jnt achieves the highest in-scope and out-of-scope accuracy. Besides, we conduct 1-shot and 10-shot experiments by randomly sampling one and ten utterances from each intent class in the training set. TOD-BERT-jnt has 13.2% all-intent accuracy improvement and 16.3% in-domain accuracy improvement compared to BERT in the 1-shot setting.

### 6.3 Dialogue State Tracking

Two evaluation metrics are commonly used in dialogue state tracking task: joint goal accuracy and slot accuracy. The joint goal accuracy compares the predicted dialogue states to the ground truth at each dialogue turn. The ground truth includes slot values for all the possible (domain, slot) pairs. The output is considered as a correct prediction if and only if all the predicted values exactly match its ground truth values. On the other hand, the slot

accuracy individually compares each (domain, slot, value) triplet to its ground truth label.

In Table 5, we compare BERT to TOD-BERT-jnt on the MWOZ 2.1 dataset and find the latter has 2.4% joint goal accuracy improvement. Since the original ontology provided by Budzianowski et al. (2018) is not complete (some labeled values are not included in the ontology), we create a new ontology of all the possible annotated values. We also list several well-known dialogue state trackers as reference, including DSTReader (Gao et al., 2019), HyST (Goel et al., 2019), TRADE (Wu et al., 2019), and ZSDST (Rastogi et al., 2019). We also report the few-shot experiments using 1%, 5%, 10%, and 25% data. Note that 1% of data has around 84 dialogues. TOD-BERT outperforms BERT in all the setting, which further show the strength of task-oriented dialogue pre-training.

### 6.4 Dialogue Act Prediction

We conduct experiments on three different datasets and report micro-F1 and macro-F1 scores for the dialogue act prediction task, a multi-label classification problem. For the MWOZ dataset, we remove the domain information from the original system dialogue act labels. For example, the “taxi-inform” will be simplified to “inform”. This process reduces the number of possible dialogue acts from 31 to 13. For DSTC2 and GSIM corpora, we follow Paul et al. (2019) to apply universal dialogue act mapping that maps the original dialogue act labels to a general dialogue act format, resulting in 9 and 6 unique system dialogue acts in DSTC2 and GSIM, respectively. We run two other baselines, MLP and RNN, to further show the strengths of BERT-based





		MWOZ		DSTC2		GSIM	
		1-to-100	3-to-100	1-to-100	3-to-100	1-to-100	3-to-100
<b>1% Data</b>	BERT	7.8% ± 2.0%	20.5% ± 4.4%	3.7% ± 0.6%	9.6% ± 1.3%	4.0% ± 0.4%	10.3% ± 1.1%
	TOD-BERT-mlm	13.0% ± 1.1%	34.6% ± 0.4%	12.5% ± 6.7%	24.9% ± 10.7%	7.2% ± 4.0%	15.4% ± 8.0%
	TOD-BERT-jnt	-	-	<b>37.5% ± 0.6%</b>	<b>55.9% ± 0.4%</b>	<b>12.5% ± 0.9%</b>	<b>26.8% ± 0.8%</b>
<b>10% Data</b>	BERT	20.9% ± 2.6%	45.4% ± 3.8%	8.9% ± 2.3%	21.4% ± 3.1%	9.8% ± 0.1%	24.4% ± 1.2%
	TOD-BERT-mlm	22.3% ± 3.2%	48.7% ± 4.0%	19.0% ± 16.3%	33.8% ± 20.4%	11.2% ± 2.5%	26.0% ± 2.7%
	TOD-BERT-jnt	-	-	<b>49.7% ± 0.3%</b>	<b>66.6% ± 0.1%</b>	<b>23.0% ± 1.0%</b>	<b>42.6% ± 1.0%</b>
<b>Full Data</b>	GPT2	47.5%	75.4%	53.7%	69.2%	39.1%	60.5%
	DialoGPT	35.7%	64.1%	39.8%	57.1%	16.5%	39.5%
	BERT	47.5%	75.5%	46.6%	62.1%	13.4%	32.9%
	TOD-BERT-mlm	48.1%	74.3%	50.0%	65.1%	36.5%	60.1%
	TOD-BERT-jnt	<b>65.8%</b>	<b>87.0%</b>	<b>56.8%</b>	<b>70.6%</b>	<b>41.0%</b>	<b>65.4%</b>

Table 6: Response selection evaluation results on three corpora for 1%, 10% and full data setting. We report 1-to-100 and 3-to-100 accuracy, which is similar to recall1 and recall@3 given 100 candidates.

metric is computed using a random batch of 100 examples so that responses from other examples in the same batch can be used as random negative candidates. This allows us to compute the metric across many examples in batches efficiently. While it is not guaranteed that the random negatives will indeed be “true” negatives, the 1-of-100 metric still provides a useful evaluation signal. During inference, we run five different random seeds to sample batches and report the average results.

In Table 6, we conduct response selection experiments on three datasets, MWOZ, DSTC2, and GSIM. TOD-BERT-jnt achieves 65.8% 1-to-100 accuracy and 87.0% 3-to-100 accuracy on MWOZ, which surpasses BERT by 18.3% and 11.5%, respectively. The similar results are also consistently observed in DSTC2 and GSIM datasets, and the advantage of the TOD-BERT-jnt is more evident in the few-shot scenario. We do not report TOD-BERT-jnt for MWOZ few-shot setting because it is not fair to compare them with others as the full MWOZ training set is used for response contrastive learning during pre-training stage. The response selection results are sensitive to the training batch size since the larger the batch size the harder the prediction. In our experiments, we set batch size equals to 25 for all the models.

## 7 Visualization

In Figure 2, we visualize the embeddings of BERT, TOD-BERT-mlm, and TOD-BERT-jnt given the same input from the MWOZ test set. Each sample point is a system response representation, which is passed through a pre-trained model and reduced its high-dimension features to a two-dimension point using the t-distributed stochastic neighbor embedding (tSNE) for dimension reduction. Since we know the true domain and dialogue act labels for

each utterance, we use different colors to represent different domains and dialogue acts. As one can observe, TOD-BERT-jnt has more clear group boundaries than TOD-BERT-mlm, and two of them are better than BERT.

To analyze the results quantitatively, we run K-means, a common unsupervised clustering algorithms, on top of the output embeddings of BERT and TOD-BERT. We set K for K-means equal to 10 and 20. After the clustering, we can assign each utterance in the MWOZ test set to a predicted class. We then compute the normalized mutual information (NMI) between the clustering result and the actual domain label for each utterance. Here is what we observe: TOD-BERT consistently achieves higher NMI scores than BERT. For K=10, TOD-BERT has a 0.143 NMI score, and BERT only has 0.094. For K=20, TOD-BERT achieves a 0.213 NMI score, while BERT has 0.109.

## 8 Conclusion

We propose task-oriented dialogue BERT (TOD-BERT) trained on nine human-human and multi-turn task-oriented datasets across over 60 domains. TOD-BERT outperforms BERT on four dialogue downstream tasks, including intention classification, dialogue state tracking, dialogue act prediction, and response selection. It also has a clear advantage in the few-shot experiments when only limited labeled data is available. TOD-BERT is easy-to-deploy and will be open-sourced, allowing the NLP research community to apply or fine-tune any task-oriented conversational problem.


## References

Layla El Asri, Hannes Schulz, Shikhar Sharma, Jeremie Zumer, Justin Harris, Emery Fine, Rahul


- Mehrotra, and Kaheer Suleman. 2017. Frames: A corpus for adding memory to goal-oriented dialogue systems. *arXiv preprint arXiv:1704.00057*.
- Siqi Bao, Huang He, Fan Wang, and Hua Wu. 2019. Plato: Pre-trained dialogue generation model with discrete latent variable. *arXiv preprint arXiv:1910.07931*.
- Paweł Budzianowski and Ivan Vulić. 2019. Hello, it's gpt-2—how can i help you? towards the use of pre-trained language models for task-oriented dialogue systems. *arXiv preprint arXiv:1907.05774*.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Inigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. Multiwoz-a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. *arXiv preprint arXiv:1810.00278*.
- Bill Byrne, Karthik Krishnamoorthi, Chinnadhurai Sankar, Arvind Neelakantan, Daniel Duckworth, Semih Yavuz, Ben Goodrich, Amit Dubey, Andy Cedilnik, and Kyu-Young Kim. 2019. Taskmaster-1: Toward a realistic and diverse dialog dataset. *arXiv preprint arXiv:1909.05358*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. In *Advances in Neural Information Processing Systems*, pages 13042–13054.
- Mihail Eric and Christopher D Manning. 2017. Key-value retrieval networks for task-oriented dialogue. *arXiv preprint arXiv:1705.05414*.
- Shuyang Gao, Abhishek Sethi, Sanchit Aggarwal, Tagyoung Chung, and Dilek Hakkani-Tür. 2019. Dialog state tracking: A neural reading comprehension approach. *arXiv preprint arXiv:1908.01946*.
- Rahul Goel, Shachi Paul, and Dilek Hakkani-Tür. 2019. Hyst: A hybrid approach for flexible and accurate dialogue state tracking. *arXiv preprint arXiv:1907.00883*.
- Matthew Henderson, Iñigo Casanueva, Nikola Mrkšić, Pei-Hao Su, Ivan Vulić, et al. 2019a. Convert: Efficient and accurate conversational representations from transformers. *arXiv preprint arXiv:1911.03688*.
- Matthew Henderson, Blaise Thomson, and Jason D. Williams. 2014. [The second dialog state tracking challenge](#). In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 263–272, Philadelphia, PA, U.S.A. Association for Computational Linguistics.
- Matthew Henderson, Ivan Vulić, Daniela Gerz, Iñigo Casanueva, Paweł Budzianowski, Sam Coope, Georgios Spithourakis, Tsung-Hsien Wen, Nikola Mrkšić, and Pei-Hao Su. 2019b. [Training neural response selection for task-oriented dialogue systems](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5392–5404, Florence, Italy. Association for Computational Linguistics.
- Dan Hendrycks and Kevin Gimpel. 2016. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. 2019. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*.
- Seokhwan Kim, Michel Galley, Chulaka Gunasekara, Adam Atkinson Sungjin Lee, Baolin Peng, Hannes Schulz, Jianfeng Gao, Jinchao Li, Mahmoud Adada, Minlie Huang, Luis Lastras, Jonathan K. Kummerfeld, Walter S. Lasecki, Chiori Hori, Anoop Cherian, Tim K. Marks, Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, and Raghav Gupta. 2019. The eighth dialog system technology challenge. *arXiv preprint*.
- Stefan Larson, Anish Mahendran, Joseph J Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K Kummerfeld, Kevin Leach, Michael A Laurenzano, Lingjia Tang, et al. 2019. An evaluation dataset for intent classification and out-of-scope prediction. *arXiv preprint arXiv:1909.02027*.
- Sungjin Lee, Hannes Schulz, Adam Atkinson, Jianfeng Gao, Kaheer Suleman, Layla El Asri, Mahmoud Adada, Minlie Huang, Shikhar Sharma, Wendy Tay, and Xiujun Li. 2019. [Multi-domain task-completion dialog challenge](#). In *Dialog System Technology Challenges 8*.
- Xiujun Li, Sarah Panda, JJ (Jingjing) Liu, and Jianfeng Gao. 2018. [Microsoft dialogue challenge: Building end-to-end task-completion dialogue systems](#). In *SLT 2018*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Nikola Mrkšić, Diarmuid O Séaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve Young. 2016. Neural belief tracker: Data-driven dialogue state tracking. *arXiv preprint arXiv:1606.03777*.
- Shachi Paul, Rahul Goel, and Dilek Hakkani-Tür. 2019. Towards universal dialogue act tagging for task-oriented dialogues. *arXiv preprint arXiv:1907.03020*.

- Baolin Peng, Chenguang Zhu, Chunyuan Li, Xijun Li, Jinchao Li, Michael Zeng, and Jianfeng Gao. 2020. Few-shot natural language generation for task-oriented dialog. *arXiv preprint arXiv:2002.12328*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2018. Towards empathetic open-domain conversation models: A new benchmark and dataset. *arXiv preprint arXiv:1811.00207*.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2019. Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset. *arXiv preprint arXiv:1909.05855*.
- Marzieh Saeidi, Ritwik Kulkarni, Theodosia Togia, and Michele Sama. 2017. The effect of negative sampling strategy on capturing semantic similarity in document embeddings. In *Proceedings of the 2nd Workshop on Semantic Deep Learning (SemDeep-2)*, pages 1–8.
- Pararth Shah, Dilek Hakkani-Tur, Bing Liu, and Gokhan Tur. 2018a. Bootstrapping a neural conversational agent with dialogue self-play, crowdsourcing and on-line reinforcement learning. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers)*, pages 41–51.
- Pararth Shah, Dilek Hakkani-Tür, Gokhan Tür, Abhinav Rastogi, Ankur Bapna, Neha Nayak, and Larry Heck. 2018b. Building a conversational agent overnight with dialogue self-play. *arXiv preprint arXiv:1801.04871*.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R Bowman, Dipanjan Das, et al. 2019. What do you learn from context? probing for sentence structure in contextualized word representations. *arXiv preprint arXiv:1905.06316*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrksic, Milica Gasic, Lina M Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2016. A network-based end-to-end trainable task-oriented dialogue system. *arXiv preprint arXiv:1604.04562*.
- Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. Transfertransfo: A transfer learning approach for neural network based conversational agents. *arXiv preprint arXiv:1901.08149*.
- 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, Florence, Italy. Association for Computational Linguistics.
- Yinfei Yang, Steve Yuan, Daniel Cer, Sheng-yi Kong, Noah Constant, Petr Pilar, Heming Ge, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. Learning semantic textual similarity from conversations. In *Proceedings of The Third Workshop on Representation Learning for NLP*, pages 164–174, Melbourne, Australia. Association for Computational Linguistics.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2019. Dialogpt: Large-scale generative pre-training for conversational response generation. *arXiv preprint arXiv:1911.00536*.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE international conference on computer vision*, pages 19–27.


## A Appendices



(a) BERT




(b) TOD-BERT-mlm




(c) TOD-BERT-jnt


Figure 3: The tSNE visualization of BERT and TOD-BERT representations of system responses in MWOZ test set. Different colors mean different domains.



(a) BERT




(b) TOD-BERT-mlm




(c) TOD-BERT-jnt

Figure 4: The tSNE visualization of BERT and TOD-BERT representations of system responses in MWOZ test set. Different colors mean different dialogue acts.






(a) BERT



(b) TOD-BERT-mlm



(c) TOD-BERT-jnt

Figure 5: The tSNE visualization of BERT and TOD-BERT representations of system responses in MWOZ test set. Different colors mean different dialogue slots.