

SPACE-2: Tree-Structured Semi-Supervised Contrastive Pre-training for Task-Oriented Dialog Understanding

Wanwei He^{1,2,✉}, Yinpei Dai^{3,✉}, Binyuan Hui^{3,✉}, Min Yang^{1*}, Zheng Cao³,
Jianbo Dong³, Fei Huang³, Luo Si³, Yongbin Li^{3*}

¹Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences, China

²University of Chinese Academy of Sciences, China

³DAMO Academy, Alibaba Group

{ww.he, min.yang}@siat.ac.cn

{yinpei.dyp, f.huang, luo.si, shuide.lyb}@alibaba-inc.com

Abstract

Pre-training methods with contrastive learning objectives have shown remarkable success in dialog understanding tasks. However, current contrastive learning solely considers the self-augmented dialog samples as positive samples and treats all other dialog samples as negative ones, which enforces dissimilar representations even for dialogs that are semantically related. In this paper, we propose **SPACE-2**, a tree-structured pre-trained conversation model, which learns dialog representations from limited labeled dialogs and large-scale unlabeled dialog corpora via semi-supervised contrastive pre-training. Concretely, we first define a general semantic tree structure (STS) to unify the inconsistent annotation schema across different dialog datasets, so that the rich structural information stored in all labeled data can be exploited. Then we propose a novel multi-view score function to increase the relevance of all possible dialogs that share similar STSs and only push away other completely different dialogs during supervised contrastive pre-training. To fully exploit unlabeled dialogs, a basic self-supervised contrastive loss is also added to refine the learned representations. Experiments show that our method can achieve new state-of-the-art results on the DialogGLUE benchmark consisting of seven datasets and four popular dialog understanding tasks. For reproducibility, we release the code and data at <https://github.com/AlibabaResearch/DAMO-ConvAI/tree/main/space-2>.

1 Introduction

Task-oriented dialog (TOD) systems aim to help users to accomplish specific tasks through natural language interactions, *e.g.*, restaurant booking, hotel reserving, and movie searching (Young et al., 2013). In order to fulfill a user goal, the dialog

systems must be capable of extracting structured semantics from the dialog utterances, which is often referred to as *dialog understanding*. Common tasks of dialog understanding include: 1) intent prediction, where the system classifies the user intention given the current utterance, 2) slot filling, where the system extracts the value spans for pre-defined slots from user utterances, and 3) dialog state tracking, where slot-value pairs are predicted and updated according to multi-turn dialog history.

Recent methods address the dialog understanding problems mainly by task-adaptive training with pre-trained language models (PLMs) (Henderson et al., 2020; Mehri et al., 2020; Zhang et al., 2022). By fine-tuning the entire parameters of PLMs on dialog corpora, the versatile knowledge acquired from large-scale corpora can be adapted to handle various dialog understanding tasks. For example, TOD-BERT (Wu et al., 2020) continually pre-trained a BERT on large-scale task-oriented dialog corpora with similar objectives to grasp more knowledge in conversations and achieved good results on a wide range of tasks. To further enhance the representations of dialogs, contrastive learning is widely used for dialog pre-training via pulling semantically similar sentences together and pushing apart dis-similar ones, which has shown strong ability in few-shot dialog understanding (Zhang et al., 2020, 2021a).

Although contrastive learning (CL) is effective in pre-training dialog models, it still suffers from several problems. First, the current practice of CL often neglects the semantic structures of dialog examples and only utilizes shallow self-supervised information. Figure 1 illustrates the semantic frames of two dialog examples, which indicate the hierarchical structures with semantic meanings in terms of intentions and slots. The semantic structure is often summarized as the turn-level structural annotation and is prominent in multi-domain task-oriented dialogs. Exploiting such structure during

* Corresponding authors.

✉ Equal contribution.

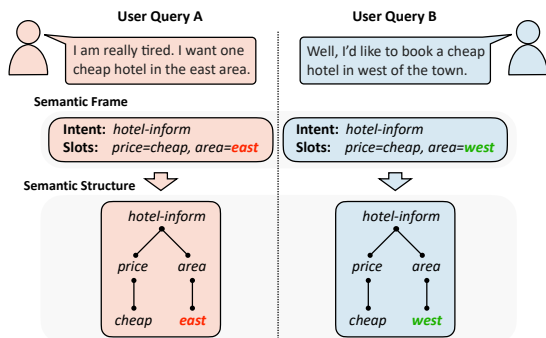


Figure 1: Two dialog examples that have similar semantic contents from the MULTIWOZ dataset. The difference is highlighted in red and green colors.

task-adaptive pre-training has been proved to be effective in dialog understanding tasks (Yu et al., 2020). However, In conventional CL, only the examples with the exact same annotations are deemed as positive samples, while all other examples are considered as negative ones (Gao et al., 2021), ignoring the fact that different examples may share similar semantics to some extent as shown in Figure 1. Second, in previous dialog pre-training methods, no annotations have been exploited to learn better pre-trained representations. But in computer vision, there have already been many works demonstrating that labeled data can accelerate the pre-training procedure as well as improve the model performance (Assran et al., 2020; Khosla et al., 2020; Dai et al., 2021b). Therefore, we argue that it is crucial to combine both the labeled and unlabeled dialog data to learn more powerful pre-trained dialog models.

Built on these motivations, we propose a novel tree-structured pre-trained conversation model (SPACE-2), which fully exploits the structural information in labeled data to improve the dialog pre-training on large-scale unlabeled corpora via semi-supervised contrastive learning. To be specific, we first define a general semantic tree structure (STS), which is suitable for TOD and compatible with the discrepancy of annotation schema in varied dialog datasets. Then we propose a novel multi-view scoring function to measure the similarity among different labeled dialogs. Given two dialogs with STSs, the scoring function is able to calculate their relations hierarchically and aggregate the output multiple scores from either global or local aspects. In this way, all labeled data are viewed as positive samples with soft scores instead of hard scores (0 or 1) in traditional CL. Thus, more subtle semantic structures in training samples can be taken into ac-

count, and dialog representations can be learned in a semi-supervised manner.

Our contributions are summarized as follows:

- We explore tree-structured semi-supervised contrastive pre-training for task-oriented dialog understanding. To the best of our knowledge, this is the first study to inject structural information and exploit labeled semantic frames in a pre-trained conversation model.
- We propose a multi-view score function for similarity measurement among different labeled dialogs, which could calculate relations hierarchically and aggregate multiple judgements from either global or local aspects.
- We conduct extensive experiments on the DialogGLUE benchmark with seven different datasets. Empirical results show that our SPACE-2 consistently performs better than strong competitors on four dialog understanding tasks (i.e., intent prediction, slot filling, semantic parsing, and dialog state tracking).

2 Related Work

2.1 Pre-trained Models for Dialog Understanding

Recent advances in pre-trained language models (PLMs) have spurred success in natural language understanding in task-oriented dialog (TOD) systems. Since general-purpose PLMs such as BERT (Devlin et al., 2019) are not able to capture dialog-oriented knowledge in TOD, such as explicit user intent, many studies proposed to perform pre-training on dialog corpora. For example, ConvRT (Henderson et al., 2020) explored the pre-trained language transformers for TOD via learning a dialog response selection loss. ConvBERT (Mehri et al., 2020) fine-tuned BERT on a large open-domain dialog corpus using a masked language modeling objective so that more semantically meaningful representations of dialogs can be learned. TOD-BERT (Wu et al., 2020) incorporated role tokens in the input and employed a contrastive objective for response selection. Zhang et al. (2021b) proposed to use masked utterance modeling and response contrastive loss to produce DialogueBERT and showed good performance on intent, emotion, and entity recognition. Apart from training with common objectives, some work (Yu et al., 2020; He et al., 2020b; Hui et al., 2021; Dai et al., 2021a) also

Statistics	<i>AnPreDial</i>	<i>UnPreDial</i>
# Datasets	32	21
# Dialogs	459,465	3,217,058
# Turns	3,366,479	19,578,028
Avg. tokens per turn	13.9	14.5
Avg. tokens per dialog	101.8	88.2
Total unique tokens	46.8M	283.7M

Table 1: Statistics of our pre-trained dialog dataset.

proposed to learn relations between schema and keywords of utterances during pre-training, which benefits in certain downstream tasks that require parsing semantics. The work of (Sun et al., 2019; Liu et al., 2020a) tried to integrate structural information of knowledge graphs to enhance language representations. Unlike these methods, we are the first to inject the semantic tree structure into pre-training models for dialog understanding.

2.2 Contrastive Learning

Contrastive learning (CL) is one of the pre-training methods that leverage large-scale unlabeled data to learn meaningful sentence representations. The key idea of CL is to narrow the distance between two semantically similar sentence representations, meanwhile, push away the representations of dissimilar sentences (Gao et al., 2021; Wu et al., 2021; Yan et al., 2021). There are also several studies focused on applying CL to dialog understanding tasks. For example, Zhang et al. (2021a) used CL as self-supervised pre-training and adapted the pre-trained models on few-shot intent prediction. The studies (Mehri and Eric, 2021; Vulić et al., 2021) utilized supervised CL at the fine-tuning stages for intent prediction, where the samples from the same class are all regarded as positives. He et al. (2020a) proposed to use adversarial samples as difficult negative instances in CL for zero-shot cross-domain slot-filling. Different from previous work, we use CL in both pre-training and fine-tuning stages. In addition, we train dialog models with both labeled and unlabeled data in a semi-supervised manner.

2.3 Semi-Supervised Contrastive Learning

Semi-supervised contrastive learning has proved to be an effective pre-training method in many research field, such as image recognition (Yuan et al., 2021; Park et al., 2021; Li et al., 2021; Kim et al., 2021a), image segmentation (Alonso et al., 2021; Zhou et al., 2021) and speech recognition (Xiao et al., 2021; Inoue and Goto, 2020). It can

combine both labeled and unlabeled data together to train effective neural models for downstream tasks. He et al. (2022b) proposed to use consistency regularization loss to learn dialog policy from labeled and unlabeled dialog corpora via a semi-supervised manner. However, leveraging such a semi-supervised contrastive learning paradigm in task-oriented dialog understanding remains unexplored. In this paper, we explore semantic structures in TOD, and propose a new contrastive learning approach that calculates semantic tree-structure similarity among all possible labeled dialog data.

3 Pre-training Data

In this section, we first describe the dialog datasets (He et al., 2022a) used for pre-training, including a labeled dialog dataset (*AnPreDial*) and a large-scale unlabeled dialog corpus (*UnPreDial*). Then, we elaborate on the semantic tree structure (STS) for the task-oriented dialog data.

3.1 Dataset Description

To provide sufficient high-quality dialog corpora to train our pre-trained conversation model, we use a labeled dialog dataset called *AnPreDial*, which contains 32 existing labeled TOD datasets, ranging from single-turn question answering to multi-turn dialogs. We also use a large-scale unlabeled dialog corpus called *UnPreDial* with 21 dialog datasets, ranging from open-domain online forums to document-grounded dialogs. The statistics of our pre-trained dialog corpora are shown in Table 1 and more details can be found in Appendix A.

3.2 Semantic Tree Structure

Task-oriented dialogs often contain several user goals, and the system needs to help fulfill these goals through interacting with users. To this end, the user utterances and system responses are usually composed of rich semantics such as intents and slots. In different datasets, the annotation taxonomy used to describe the semantic scopes varies. For example, in some intent recognition datasets such as BANKING77 (Casanueva et al., 2020) and CLINIC150 (Larson et al., 2019), they only possess a pre-defined intent set and formulate the task as a classification problem. In other datasets like Taskmaster (Byrne et al., 2019a), only the slot labels are provided; for instance, given a user utterance “I need a table for tonight at 7 pm for 8 people”, the label is “time.reservation=7




Figure 2: Different semantic tree structures in *AnPreDial*, where the blue, red, green and purple circles indicate DISV, ISV, IS, SV sub-structures.

pm, num.guests=8" without any intent information. Therefore, in order to utilize all possible data in *AnPreDial* to pre-train our model, we adopt a unified schema called semantic tree structure (STS) that is suitable for every TOD sample.

Typically, the STS contains four layers: *domain* layer, *intent* layer, *slot* layer, and *value* layer. Every layer is composed of nodes, which are referred to as respective elements of labeled semantic frames. If no matched annotations are provided, the nodes of the corresponding layer are set as empty (denoted as NULL). The first layer consists of domain nodes as successors of the root node, preceded by the intent nodes. As children of intent nodes, slot nodes occupy the third layer of the semantic tree. As leaf nodes, the value nodes take the slot nodes as parents. Figure 2 shows three tree structures of different dialog samples from MULTIWOZ (Eric et al., 2020), BANKING77 and Taskmaster. In MULTIWOZ, the labels follow the dialog act schema in Young (2007), and all layers contain certain semantic elements. But the other two datasets do not follow the schema and lack respective elements in some layers.

4 Methods

In this section, we first introduce the model architecture. Then we expound on the pre-training procedure with the proposed tree-structured semi-supervised contrastive learning (CL) paradigm.

4.1 Model Architecture

As illustrated in Figure 3, we build our SPACE-2 model based on the bidirectional Transformer architecture (Vaswani et al., 2017). Different from the vanilla input representations as in BERT (Devlin et al., 2019), we set our input embeddings consisting of four elements: tokens, roles, turns, and positions. Role embeddings are used to segment

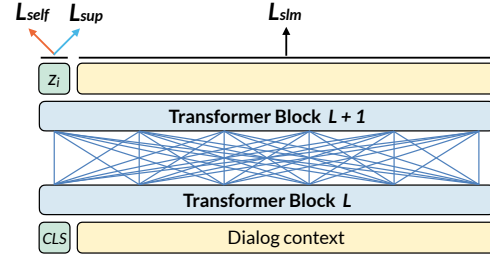


Figure 3: Model Framework. L_{self} is the tree-structured self-supervised contrastive loss for unlabeled data, and L_{sup} is the tree-structured supervised contrastive loss for labeled data. L_{slm} is the span MLM loss for all data.

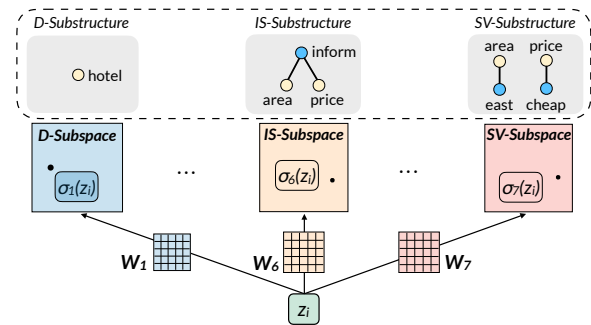


Figure 4: Sub-spaces for the multi-view scoring objective. The dashed black line denotes the annotations of different sub-structures, only used in our supervised CL.

which role the current token belongs to either user or system. Turn embeddings are assigned to each token according to its turn number in the dialog. Position embeddings are assigned to each token according to its relative position within its belonging sentence. Appendix B shows more details.

4.2 Tree-Structured Semi-Supervised Contrastive Learning Paradigm

We aim to leverage semi-supervised pre-training to learn better pre-trained representations from both the labeled and unlabeled data. Concretely, we adopt a tree-structured supervised contrastive objective on the labeled dataset *AnPreDial*, while a tree-structured self-supervised contrastive objective on the unlabeled dataset *UnPreDial*.

4.2.1 Tree-Structured Supervised CL

For labeled data, we treat every dialog sample as a positive sample assigned with similarity scores to each other. Between two semantic tree structures (STSs), we define a score function to calculate similarity scores from different views. Let D , I , S , V denote the sets of node in domain, intent, slot and value layer. For the sample from MULTIWOZ dataset in Figure 2, its $D=\{restaurant\}$, $I=\{inform,$

request}, $S=\{food, area, name\}$ and $V=\{indian, south\}$. For the sample from BANKING77 dataset, its $D=\{\}$, $I=\{card_arrival\}$, $S=\{\}$ and $V=\{\}$. We also build the sets of each possible path between two nodes on the STS, which we denote as DI, IS, SV, DIS, ISV, DISV. Thus for the MULTIWOZ sample in Figure 2, $ISV=\{inform \rightarrow food \rightarrow indian, inform \rightarrow area \rightarrow south\}$, and for BANKING77, $ISV=\{\}$. The total super set of all above node-sets and path-sets is denoted as $\mathcal{S}_{all} = \{D, I, S, V, DI, IS, SV, DIS, ISV, DISV\}$, where the size $K = 10$. Given a pair of STSs (denoted as T_i, T_j), then we calculate the Jaccard similarity coefficient for each set s_k in \mathcal{S}_{all} to get the score f_k respectively as follows:

$$f_k(T_i, T_j) = J(s_k^i, s_k^j) \quad (1)$$

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|} \quad (2)$$

where we use i, j to discriminate s_k and \mathcal{S}_{all} for two different STSs: $s_k^i \in \mathcal{S}_{all}^i, s_k^j \in \mathcal{S}_{all}^j$. J is the function for the Jaccard similarity coefficient, which takes the ratio of Intersection over Union.

Given the dialog context, our model output a pooled representation at the [CLS] position as the sentence embedding z of the whole context. Thus for any sample pair i, j , the output sentence embeddings are denoted as z_i, z_j . We use $f_{i,j}^k$ to denote the k -th similarity score $f_k(T_i, T_j)$ for simplicity. Suppose the size of the current batch is N , we follow the common practice in Gao et al. (2021) to duplicate the data via dropout-based data augmentation. Then we acquire a $2N$ -length new batch. Let $I = \{1 \dots 2N\}$ be the index set of the new batch, and $C(i) \equiv I/\{i\}$. Since we compute multiple scores $f_{i,j}^k$ for each STS pair, different types of scoring functions can be used to construct the CL loss. In this paper, We propose both single-view and multi-view scoring functions to build the tree-structured supervised CL objectives as follows and find that the multi-view method performs better.

Single-View Scoring Objective. Similar to the common practice in the current CL, we simply average the K scores into a single value to weigh the semantic similarity among samples:

$$L_{sup}^{single} = - \sum_{i \in I} \sum_{j \in C(i)} \left(\frac{1}{K} \sum_{k=1}^K f_{i,j}^k \right) \times \log \frac{\exp(\sigma(z_i) \cdot \sigma(z_j)/\tau)}{\sum_{l \in C(i)} \exp(\sigma(z_i) \cdot \sigma(z_l)/\tau)} \quad (3)$$

$$\sigma(z) = \text{Norm}(Wz + b) \quad (4)$$

where $\sigma(\cdot)$ is a normalized linear mapping to transform z to a unit vector. W is a learnable square weight matrix and b is a bias vector. $\text{Norm}(\cdot)$ is the normalization operation and $\tau \in [0, 1]$ is a temperature hyper-parameter.

Multi-View Scoring Objective. Figure 4 demonstrates our multi-view scoring objective. Instead of summing up as one score, we can separately utilize multiple scores to better exploit sub-structures from different semantic spaces (Zhu et al., 2019) as the following loss:

$$L_{sup}^{multi} = - \sum_{i \in I} \sum_{j \in C(i)} \sum_{k=1}^K \left(\frac{f_{i,j}^k}{\sum_{m \in C(i)} f_{i,m}^k} \times \log \frac{\exp(\sigma_k(z_i) \cdot \sigma_k(z_j)/\tau)}{\sum_{l \in C(i)} \exp(\sigma_k(z_i) \cdot \sigma_k(z_l)/\tau)} \right) \quad (5)$$

$$\sigma_k(z) = \text{Norm}(W_k z + b_k) \quad (6)$$

where $\sigma_k(\cdot)$ denotes linear mappings with different parameters. As illustrated in Figure 5, our methods are quite distinct from vanilla supervised CL (Vulić et al., 2021), where only augmented data or data from the same label are considered as positive samples ($f_{i,j}^k = 1$), and all other data are considered as negative samples ($f_{i,j}^k = 0$). Therefore, our method can be viewed as a generalized case where all labeled data are regarded as positive samples assigned with soft scores.

4.2.2 Tree-Structured Semi-Supervised CL

For unlabeled data, since there are no available labels, we adopt a tree-structured self-supervised contrastive objective in a similar way in Gao et al. (2021), where only the augmented data by dropout is deemed as a positive sample. In particular, under the multi-view setting, the objective L_{self}^{multi} still shares the same K sub-spaces as L_{sup}^{multi} as follows:

$$L_{self}^{single} = - \sum_{i \in I} \sum_{j \in C(i)} \mathbf{1}(j = i^+) \times \log \frac{\exp(\sigma(z_i) \cdot \sigma(z_j)/\tau)}{\sum_{l \in C(i)} \exp(\sigma(z_i) \cdot \sigma(z_l)/\tau)} \quad (7)$$

$$L_{self}^{multi} = - \sum_{i \in I} \sum_{j \in C(i)} \sum_{k=1}^K \mathbf{1}(j = i^+) \times \log \frac{\exp(\sigma^k(z_i) \cdot \sigma^k(z_j)/\tau)}{\sum_{l \in C(i)} \exp(\sigma^k(z_i) \cdot \sigma^k(z_l)/\tau)} \quad (8)$$

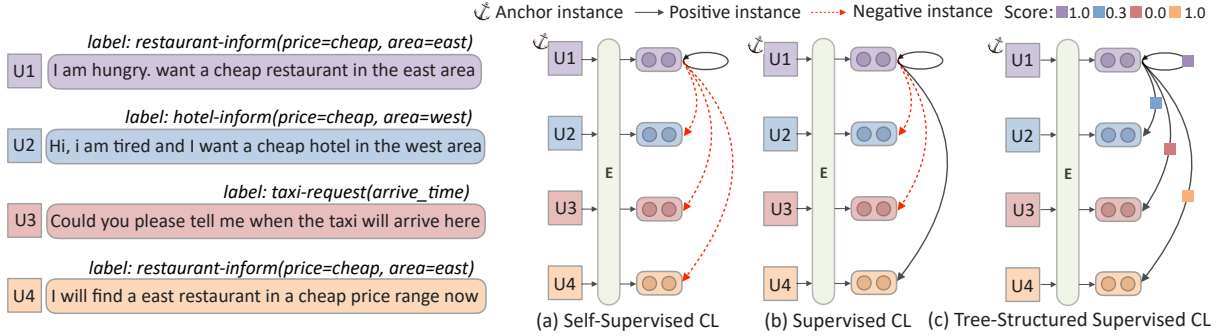


Figure 5: (a) Self-supervised CL only predicts augmented itself from in-batch negatives, with different dropout masks applied. (b) Supervised CL considers samples of the exact same label as positives. (c) Tree-structured supervised CL considers all in-batch samples as positives with soft scores. Only scores of ISV set are depicted here.

where $\mathbf{1}(j = i^+)$ means it only takes 1 when the data sample j is the augmented data i^+ of data sample i , otherwise 0.

In order to maintain the language modeling ability and extract the slot values better, we also apply the spanMLM loss L_{slm} (Joshi et al., 2020) to the output representations on top of the dialog context. Therefore, our final loss is computed as:

$$L_{total}^* = L_{sup}^* + L_{self}^* + L_{slm} \quad (9)$$

where $*$ can be either *single* or *multi*.

5 Experimental Setup

5.1 Task Formulation

DialoGLUE benchmark consists of seven different datasets, including four tasks as follows:

Intent Prediction. This is a classification task for models to predict which intent class the user conveys in the utterance. There are three datasets accordingly: 1) BANKING77 (Casanueva et al., 2020), a single turn dataset in the banking domain with 77 intent labels and 13k utterances; 2) CLINIC150 (Larson et al., 2019), a multi-domain dataset that contains 23.7k utterances, 150 intents, 10 domains; and 3) HWU64 (Liu et al., 2021) includes 25.7k utterances for 64 intents spanning 21 domains. To conduct the downstream task, a linear layer neural network is added to the pooled output z to classify. The evaluation metric is accuracy.

Slot Filling. This is a sequential labeling task in which a model tries to predict the IOB tags upon the sequence of dialog context. There are two datasets: 1) REST8K (Coope et al., 2020), which contains 8.2k utterances and 5 slots (date, time, people, first name, last name) from a restaurant domain; 2) a transformed DSTC8 (Rastogi et al., 2020) dataset comprising of 5.5k utterances and slot annotations

in 4 domains (buses, events, homes, rental cars). Both datasets are single-turn formed and every token in the utterance is predicted as either being the beginning of a slot value (B-), inside a slot value (I-), or not belonging to a slot value (O). A linear layer neural network is added to the representation of the final output layer to predict correct tags (e.g., “B-time”, “I-people”). The evaluation metric is the macro-averaged F1 score.

Semantic Parsing. The original TOP (Gupta et al., 2018) dataset targeting at hierarchical parsing is transformed into a joint task of intent prediction and slot-filling, which contains 44k utterances. The metric is the exact-match as in Mehri et al. (2020).

Dialog State Tracking. This task is multi-turn slot-filling, where the model interprets the user meanings during the course of a dialog to maintain a dialog state in terms of slot-value pairs. MULTIWOZ 2.1 (Eric et al., 2020) is the chosen dataset version with 10k dialogs. Following Mehri et al. (2020), we report joint goal accuracy and use TripPy (Heck et al., 2020) as the downstream model.

5.2 Baselines

Our compared baselines can be divided into two categories. The first type is the pre-trained model adapted for general dialog understanding tasks in DialoGLUE. We choose BERT (Devlin et al., 2019), ConvBERT (Mehri et al., 2020) and its variants, and TOD-BERT (Wu et al., 2020) as our baselines. The second type is the pre-trained model designed only for specific dialog tasks. In this paper, we choose Example+Observer (Mehri and Eric, 2021), DNNC (Zhang et al., 2020), CPFT (Zhang et al., 2021a); and ConvFit (Vulić et al., 2021) for the intent prediction task. We also choose SpanConveRT (Coope et al., 2020), ConVEx (Hender-

Full data training								
Model	<i>average</i>	BANKING77	HWU64	CLINIC150	REST8K	DSTC8	TOP	MULTIWOZ
BERT [†]	86.08	93.02	89.87	95.93	95.53	90.05	81.90	56.30
ConvBERT [†]	86.01	92.95	90.43	97.07	95.90	87.58	82.13	56.00
ConvBERT+MLM [†]	86.89	93.44	92.38	97.11	95.44	91.20	82.08	56.56
ConvBERT-DG [†]	82.90	93.21	91.64	96.96	93.44	74.54	72.22	58.57
ConvBERT-DG+MLM [†]	85.34	92.99	91.82	97.11	94.34	86.49	76.36	58.29
TOD-BERT _{adapt}	86.95	93.61	92.95	97.34	94.23	90.61	82.33	57.63
ConvBERT+MLM _{adapt}	87.42	94.01	93.24	97.40	95.55	91.12	82.64	58.01
SPACE-2 ^{single}	87.83	94.68	94.05	97.56	95.58	91.20	82.66	59.08
SPACE-2 ^{multi}	88.10	94.77	94.33	97.80	96.20	91.38	82.74	59.51
<i>w/o tree</i>	87.57	94.19	93.49	97.44	96.08	91.04	82.23	58.52
<i>w/o annotation</i>	87.33	93.93	93.23	97.33	95.99	90.93	82.02	57.87
<i>w/o UnPreDial</i>	87.77	94.35	93.77	97.64	96.14	91.18	82.30	58.99
10-shot data training								
Model	<i>average</i>	BANKING77	HWU64	CLINIC150	REST8K	DSTC8	TOP	MULTIWOZ
BERT [†]	66.07	79.87	81.69	89.52	87.28	45.05	74.38	4.69
ConvBERT [†]	68.03	83.63	83.77	92.10	86.90	49.08	74.86	5.90
ConvBERT+MLM [†]	68.22	83.99	84.52	92.75	86.17	48.40	78.84	6.87
ConvBERT-DG [†]	73.75	84.42	85.17	92.87	87.65	41.94	75.27	48.94
ConvBERT-DG+MLM [†]	73.80	85.06	85.69	93.06	87.58	44.36	72.01	48.89
TOD-BERT _{adapt}	75.70	85.99	86.74	93.07	87.62	50.19	77.77	48.54
ConvBERT+MLM _{adapt}	76.51	86.73	87.29	93.65	87.72	52.48	78.88	48.85
SPACE-2 ^{single}	77.43	88.08	88.29	94.80	88.52	53.27	79.01	50.07
SPACE-2 ^{multi}	77.98	88.31	88.85	95.22	88.85	54.41	79.55	50.70
<i>w/o tree</i>	77.31	87.41	87.55	94.76	88.30	54.26	78.95	49.97
<i>w/o annotation</i>	77.20	87.60	87.64	94.51	87.93	54.01	79.00	49.73
<i>w/o UnPreDial</i>	77.65	88.15	88.38	94.71	88.61	54.12	79.33	50.28

Table 2: Total results of seven datasets from the DialoGLUE benchmark. [†] denotes original results from DialoGLUE. _{adapt} denotes re-implementation via adaptively pre-training on our corpora. ‘w/o’ denotes ablation study in Sec 6.4

son and Vulić, 2021) and GenSF (Mehri and Eskenazi, 2021) for the slot-filling task. More details about baselines can be found in Appendix C.

5.3 Settings

In our experiment, we initialize SPACE-2 with ConvBERT. The input length is bound to 256 and the batch size is 128. AdamW is used for optimization with an initial learning rate of 1e-5. The dropout rate is 0.2. For semi-supervised pre-training, we combine the labeled and unlabeled data to form a batch in a 1:1 ratio. For each few-shot experiment, we **exclude** the training data in the *AnPreDial* accordingly to avoid unfair data use. For instance, if the target few-shot task is HWU64, we only use the few-shot training data of HWU64 and all other datasets as pre-training resources, so no extra training data (dev&test) in HWU64 is used. For all downstream tasks, we average 5 seeds for final results.

6 Experimental Results

6.1 Overall Performance on DialoGLUE

We follow the original settings in Mehri et al. (2020) and obtain complete results for our model with single-view scoring objectives, denoted as SPACE-2^{single}, and with multi-view scoring objectives, denoted as SPACE-2^{multi}, on all seven datasets of DialoGLUE. As shown in Table 2, our SPACE-2^{multi} model achieves state-of-the-art performance across all tasks, improving the average score of previous best results by **1.39%** (86.89%→88.10%) and **5.66%** (73.80%→77.98%) on full-data training and 10-shot training settings, respectively. This indicates that our semi-supervised contrastive learning can acquire better pre-trained representations to adapt efficiently to downstream dialog tasks by utilizing both labeled and unlabeled data. Although the superior results of SPACE-2^{single} over other baselines, it is worse than SPACE-2^{multi}, suggesting that multi-view scoring is crucial to improve the overall performance. To eliminate the discrepancy brought by

Model	BANKING77		CLINIC150		HWU64	
	Few	Full	Few	Full	Few	Full
Example+Observer	85.95	93.83	93.97	97.31	86.28	93.03
DNNC	86.71	-	93.76	-	84.72	-
CPFT	87.20	-	94.18	-	87.13	-
ConvFiT	87.38	94.16	92.89	97.34	85.32	92.42
SPACE-2 ^{multi}	88.31	94.77	95.22	97.80	88.85	94.33

Table 3: Results on intent prediction tasks, including 10-shot training (Few) and full-set training (Full) settings.

the pre-training dialog corpora, we choose two similar models, TOD-BERT and ConvBERT+MLM, to continually pre-train on *AnPreDial* and *UnPreDial* for fairness, but they are still not comparable to SPACE-2.

6.2 More Comparison on Intent Prediction

The experimental results of three intent prediction datasets are shown in Table 3. As we can see, SPACE-2^{multi} achieves SOTA results on all datasets under all settings. Specifically, it outperforms the previous best model ConvFiT by 0.93%, 2.33% 3.53% on BANKING77, CLINIC150, HWU64 under the few-data training setting, respectively. It also improves ConvFiT by 0.61%, 0.46%, and 1.91% on the above three datasets under the full-shot training setting. The improvements indicate that our method has a better ability to discriminate similar intents via tree-structure enhanced contrastive learning than vanilla pre-training on ConvFiT.

6.3 More Comparison on Slot-Filling

We also conduct slot-filling experiments, particularly in the few-shot setting, to probe whether our model is also suitable for slot extraction. We follow the same data split and settings in GenSF (Mehri and Eskenazi, 2021) for a fair comparison. As shown in Table 4, SPACE-2^{multi} achieves state-of-the-art results across all experimental settings on DSTC8 and REST8K. Note that in Mehri and Eskenazi (2021), DSTC8 is evaluated in every single domain (Bus, Event, Home, RentalCar) and REST8K is evaluated with varied proportions (1/8, 1/16, 1/32) of the training set. It shows that our SPACE-2 model can achieve superior results without any specific designs only for slot-filling tasks.

6.4 Ablation Study

To figure out which factor contributes mainly to our semi-supervised pre-training paradigm, we conduct ablation studies on DialogLUE as shown in

Dataset	Span-ConveRT	ConVEx	GenSF	SPACE-2 ^{multi}
DSTC8-Bus	84.0	86.7	90.5	91.6
DSTC8-Event	82.2	87.2	91.2	92.4
DSTC8-Home	95.4	94.5	93.7	96.2
DSTC8-RentalCar	83.0	87.4	86.7	88.3
REST8K -1/8	88.5	90.6	91.8	93.9
REST8K -1/16	81.1	86.4	89.7	92.6
REST8K -1/32	63.8	81.8	82.1	84.7

Table 4: F1 scores across all slots for evaluation on the DSTC8 single-domain datasets in the few-shot setting.

the bottom part of each setting in Table 2. We first investigate the effect of the tree-structured supervised contrastive learning by replacing it with the traditional supervised contrastive learning (denote as ‘w/o tree’). As we can see, vanilla supervised contrastive learning only considers self-augmented data or data with the exact same labels as positive samples, so it can not capture a more detailed semantic structure than our tree-structured method and degrade from 88.10% to 87.57% in an average score. We also try to remove the label information from *AnPreDial* and treat all pre-training data as unlabeled data (denote as ‘w/o annotation’), or remove the unlabeled data (denote as ‘w/o *UnPreDial*’) to see the impact. It shows that both the labeled data and unlabeled data are crucial in our semi-supervised pre-training scheme but labeled data contributes more to the overall performance.

6.5 Visualization Analysis

In this part, we investigate why explicitly decomposing the semantic structure into sub-structures for different views can benefit dialog understanding. Figure 6 illustrates the 2D t-SNE visualization of the output unit vectors $\sigma_k(z)$ for test dialog samples from the MULTIWOZ dataset. Due to the limited space, we only show the sub-spaces of D, I, S, V here. As we can see, The hidden representations of SPACE-2^{multi} are able to differentiate the similar and dissimilar parts in different semantic sub-spaces. The learned latent sub-space is highly correlated with the dialog annotations of *domain*, *intent*, *slot* or *value*, which confirms our assumption.

7 Conclusion

In this work, we propose a new pre-trained conversation model named SPACE-2, which learns dialog representations from both labeled and unlabeled corpora via tree-structured semi-supervised contrastive learning (CL). To be specific, we utilize a semantic tree structure (STS) to unify the incon-

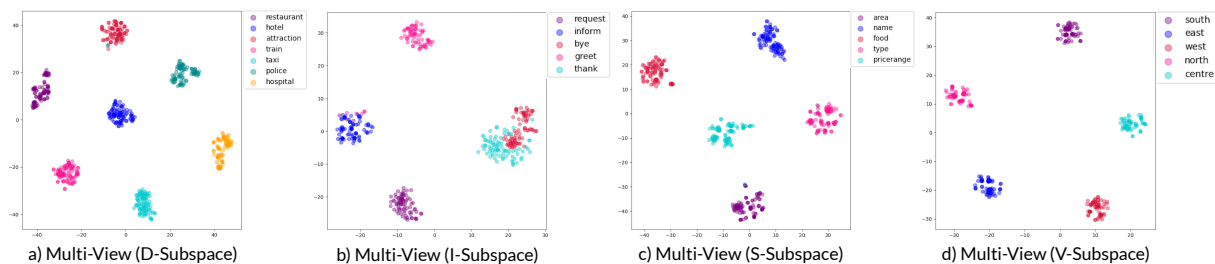


Figure 6: T-SNE visualization of the $\sigma_k(z)$ in multi-view scoring function for MULTIWOZ test dialog samples. Four sub-spaces, including Domain-, Intent-, Slot- and Value- Subspace, are shown here.

sistent annotation schema for all datasets. Then, we propose a novel multi-view scoring function to compute the semantic similarity based on STS’s sub-structures for labeled dialogs. Different from traditional CL, our method increases the relevance of representations for all similarly labeled dialogs and only pushes away completely different dialogs. A self-supervised CL loss is also integrated to exploit unlabeled dialogs. Extensive experiments on the DialoGLUE benchmark with seven datasets demonstrate that our model achieves the best results on all dialog understanding tasks, under the full-data and the few-shot settings.

Acknowledgement

This work was partially supported by National Natural Science Foundation of China (No. 61906185), Youth Innovation Promotion Association of CAS China (No. 2020357), Shenzhen Science and Technology Innovation Program (No. KQTD20190929172835662), Shenzhen Basic Research Foundation (No. JCYJ20210324115614039 and No. JCYJ20200109113441941). This work was supported by Alibaba Group through Alibaba Research Intern Program.

References

- Inigo Alonso, Alberto Sabater, David Ferstl, Luis Montesano, and Ana C Murillo. 2021. Semi-supervised semantic segmentation with pixel-level contrastive learning from a class-wise memory bank. *arXiv preprint arXiv:2104.13415*.
- Mahmoud Assran, Nicolas Ballas, Lluís Castrejon, and Michael Rabbat. 2020. Supervision accelerates pre-training in contrastive semi-supervised learning of visual representations. *arXiv preprint arXiv:2006.10803*.
- Bill Byrne, Karthik Krishnamoorthi, Chinnadhurai Sankar, Arvind Neelakantan, Daniel Duckworth, Semih Yavuz, Ben Goodrich, Amit Dubey, Andy Cedilnik, and Kyu-Young Kim. 2019a. Taskmaster-1: Toward a realistic and diverse dialog dataset. *arXiv preprint arXiv:1909.05358*.
- Bill Byrne, Karthik Krishnamoorthi, Chinnadhurai Sankar, Arvind Neelakantan, Ben Goodrich, Daniel Duckworth, Semih Yavuz, Amit Dubey, Kyu-Young Kim, and Andy Cedilnik. 2019b. [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, Hong Kong, China. Association for Computational Linguistics.
- Giovanni Campagna, Agata Foryciarz, Mehrad Moradshahi, and Monica Lam. 2020. [Zero-shot transfer learning with synthesized data for multi-domain dialogue state tracking](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 122–132, Online. Association for Computational Linguistics.
- Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. 2020. [Efficient intent detection with dual sentence encoders](#). In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 38–45, Online. Association for Computational Linguistics.
- Derek Chen, Howard Chen, Yi Yang, Alex Lin, and Zhou Yu. 2021. Action-based conversations dataset: A corpus for building more in-depth task-oriented dialogue systems. *arXiv preprint arXiv:2104.00783*.
- Samuel Coope, Tyler Farghly, Daniela Gerz, Ivan Vulić, and Matthew Henderson. 2020. [Span-CONVERT: Few-shot span extraction for dialog with pretrained conversational representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 107–121, Online. Association for Computational Linguistics.
- Alice Coucke, Alaa Saade, Adrien Ball, Théodore Bluche, Alexandre Caulier, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Calta-girone, Thibaut Lavril, et al. 2018. Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces. *arXiv preprint arXiv:1805.10190*.

- Yinpei Dai, Hangyu Li, Yongbin Li, Jian Sun, Fei Huang, Luo Si, and Xiaodan Zhu. 2021a. [Preview, attend and review: Schema-aware curriculum learning for multi-domain dialogue state tracking](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 879–885, Online. Association for Computational Linguistics.
- Zhigang Dai, Bolun Cai, Yugeng Lin, and Junying Chen. 2021b. Unimoco: Unsupervised, semi-supervised and full-supervised visual representation learning. *arXiv preprint arXiv:2103.10773*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2018. Wizard of wikipedia: Knowledge-powered conversational agents. *arXiv preprint arXiv:1811.01241*.
- 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](#). In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 207–219, Saarbrücken, Germany. Association for Computational Linguistics.
- Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyang Gao, Adarsh Kumar, Anuj Goyal, Peter Ku, and Dilek Hakkani-Tur. 2020. [MultiWOZ 2.1: A consolidated multi-domain dialogue dataset with state corrections and state tracking base-lines](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 422–428, Marseille, France. European Language Resources Association.
- 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, Saarbrücken, Germany. Association for Computational Linguistics.
- Joachim Fainberg, Ben Krause, Mihai Dobre, Marco Damonte, Emmanuel Kahembwe, Daniel Duma, Bonnie Webber, and Federico Fancellu. 2018. Talking to myself: self-dialogues as data for conversational agents. *arXiv preprint arXiv:1809.06641*.
- Song Feng, Hui Wan, Chulaka Gunasekara, Siva Patel, Sachindra Joshi, and Luis Lastras. 2020. [doc2dial: A goal-oriented document-grounded dialogue dataset](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8118–8128, Online. Association for Computational Linguistics.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Karthik Gopalakrishnan, Behnam Hedayatnia, Qinglang Chen, Anna Gottardi, Sanjeev Kwatra, Anu Venkatesh, Raefer Gabriel, Dilek Hakkani-Tür, and Amazon Alexa AI. 2019. Topical-chat: Towards knowledge-grounded open-domain conversations. In *INTERSPEECH*, pages 1891–1895.
- Mansi Gupta, Nitish Kulkarni, Raghuv eer Chanda, Anirudha Rayasam, and Zachary C Lipton. 2019. Amazonqa: A review-based question answering task. *arXiv preprint arXiv:1908.04364*.
- Sonal Gupta, Rushin Shah, Mrinal Mohit, Anuj Kumar, and Mike Lewis. 2018. [Semantic parsing for task oriented dialog using hierarchical representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2787–2792, Brussels, Belgium. Association for Computational Linguistics.
- Keqing He, Jinchao Zhang, Yuanmeng Yan, Weiran Xu, Cheng Niu, and Jie Zhou. 2020a. Contrastive zero-shot learning for cross-domain slot filling with adversarial attack. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1461–1467.
- Wanwei He, Yinpei Dai, Min Yang, Jian Sun, Fei Huang, Luo Si, and Yongbin Li. 2022a. Space-3: Unified dialog model pre-training for task-oriented dialog understanding and generation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 187–200.
- Wanwei He, Yinpei Dai, Yinhe Zheng, Yuchuan Wu, Zheng Cao, Dermot Liu, Peng Jiang, Min Yang, Fei Huang, Luo Si, et al. 2022b. Space: A generative pre-trained model for task-oriented dialog with semi-supervised learning and explicit policy injection. *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Wanwei He, Min Yang, Rui Yan, Chengming Li, Ying Shen, and Ruifeng Xu. 2020b. Amalgamating knowledge from two teachers for task-oriented dialogue system with adversarial training. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3498–3507.

- Michael Heck, Carel van Niekerk, Nurul Lubis, Christian Geisbauer, 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, 1st virtual meeting. Association for Computational Linguistics.
- Charles T. Hemphill, John J. Godfrey, and George R. Doddington. 1990. [The ATIS spoken language systems pilot corpus](#). In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*.
- Matthew Henderson, Iñigo Casanueva, Nikola Mrkšić, Pei-Hao Su, Tsung-Hsien Wen, and Ivan Vulić. 2020. [ConveRT: Efficient and accurate conversational representations from transformers](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2161–2174, Online. Association for Computational Linguistics.
- Matthew Henderson and Ivan Vulić. 2021. [ConVEx: Data-efficient and few-shot slot labeling](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3375–3389, Online. Association for Computational Linguistics.
- Binyuan Hui, Ruiying Geng, Qiyu Ren, Binhua Li, Yongbin Li, Jian Sun, Fei Huang, Luo Si, Pengfei Zhu, and Xiaodan Zhu. 2021. Dynamic hybrid relation exploration network for cross-domain context-dependent semantic parsing. In *AAAI*.
- Nakamasa Inoue and Keita Goto. 2020. Semi-supervised contrastive learning with generalized contrastive loss and its application to speaker recognition. In *2020 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pages 1641–1646. IEEE.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. [SpanBERT: Improving pre-training by representing and predicting spans](#). *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Endri Kacupaj, Joan Plepi, Kuldeep Singh, Harsh Thakkar, Jens Lehmann, and Maria Maleshkova. 2021. [Conversational question answering over knowledge graphs with transformer and graph attention networks](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 850–862, Online. Association for Computational Linguistics.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. *arXiv preprint arXiv:2004.11362*.
- Byoungjip Kim, Jinho Choo, Yeong-Dae Kwon, Seongho Joe, Seungjai Min, and Youngjune Gwon. 2021a. Selfmatch: Combining contrastive self-supervision and consistency for semi-supervised learning. *arXiv preprint arXiv:2101.06480*.
- Seokhwan Kim, Yang Liu, Di Jin, Alexandros Papanagelis, Karthik Gopalakrishnan, Behnam Hedayatnia, and Dilek Hakkani-Tur. 2021b. ["how robust r u?": Evaluating task-oriented dialogue systems on spoken conversations](#).
- Satwik Kottur, Seungwhan Moon, Alborz Geramifard, and Babak Damavandi. 2021. Simmc 2.0: A task-oriented dialog dataset for immersive multimodal conversations. *arXiv preprint arXiv:2104.08667*.
- Stefan Larson, Anish Mahendran, Joseph J. Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K. Kummerfeld, Kevin Leach, Michael A. Laurenzano, Lingjia Tang, and Jason Mars. 2019. [An evaluation dataset for intent classification and out-of-scope prediction](#). 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 1311–1316, Hong Kong, China. Association for Computational Linguistics.
- Junnan Li, Caiming Xiong, and Steven CH Hoi. 2021. Comatch: Semi-supervised learning with contrastive graph regularization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9475–9484.
- 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. *arXiv preprint arXiv:2010.12850*.
- Xiujun Li, Yu Wang, Siqi Sun, Sarah Panda, Jingjing Liu, and Jianfeng Gao. 2018. Microsoft dialogue challenge: Building end-to-end task-completion dialogue systems. *arXiv preprint arXiv:1807.11125*.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. Dailydialog: A manually labelled multi-turn dialogue dataset. *arXiv preprint arXiv:1710.03957*.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, Peng Xu, Feijun Jiang, Yuxiang Hu, Chen Shi, and Pascale Fung. 2021. Bitod: A bilingual multi-domain dataset for task-oriented dialogue modeling. *arXiv preprint arXiv:2106.02787*.
- Jingjing Liu, Panupong Pasupat, Scott Cyphers, and Jim Glass. 2013. Asgard: A portable architecture for multilingual dialogue systems. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 8386–8390. IEEE.
- Weijie Liu, Peng Zhou, Zhe Zhao, Zhiruo Wang, Qi Ju, Haotang Deng, and Ping Wang. 2020a. K-bert: Enabling language representation with knowledge graph.

- In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2901–2908.
- Xingkun Liu, Arash Eshghi, Pawel Swietojanski, and Verena Rieser. 2021. Benchmarking natural language understanding services for building conversational agents. In *Increasing Naturalness and Flexibility in Spoken Dialogue Interaction*, pages 165–183. Springer.
- Zihan Liu, Yan Xu, Tiezheng Yu, Wenliang Dai, Ziwei Ji, Samuel Cahyawijaya, Andrea Madotto, and Pascale Fung. 2020b. Crossner: Evaluating cross-domain named entity recognition. *arXiv preprint arXiv:2012.04373*.
- Shikib Mehri and Mihail Eric. 2021. [Example-driven intent prediction with observers](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2979–2992. Online. Association for Computational Linguistics.
- Shikib Mehri, Mihail Eric, and Dilek Hakkani-Tur. 2020. Dialogue: A natural language understanding benchmark for task-oriented dialogue. *arXiv preprint arXiv:2009.13570*.
- Shikib Mehri and Maxine Eskenazi. 2021. [GenSF: Simultaneous adaptation of generative pre-trained models and slot filling](#). In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 489–498, Singapore and Online. Association for Computational Linguistics.
- Johannes EM Mosig, Shikib Mehri, and Thomas Kober. 2020. Star: A schema-guided dialog dataset for transfer learning. *arXiv preprint arXiv:2010.11853*.
- 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, Vancouver, Canada. Association for Computational Linguistics.
- Jongjin Park, Sukmin Yun, Jongheon Jeong, and Jinwoo Shin. 2021. Opencos: Contrastive semi-supervised learning for handling open-set unlabeled data. *arXiv preprint arXiv:2107.08943*.
- Denis Peskov, Nancy Clarke, Jason Krone, Brigi Fodor, Yi Zhang, Adel Youssef, and Mona Diab. 2019. [Multi-domain goal-oriented dialogues \(MultiDoGO\): Strategies toward curating and annotating large scale dialogue data](#). 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 4526–4536, Hong Kong, China. Association for Computational Linguistics.
- Filip Radlinski, Krisztian Balog, Bill Byrne, and Karthik Krishnamoorthi. 2019. [Coached conversational preference elicitation: A case study in understanding movie preferences](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 353–360, Stockholm, Sweden. Association for Computational Linguistics.
- 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.
- Siva Reddy, Danqi Chen, and Christopher D. Manning. 2019. [CoQA: A conversational question answering challenge](#). *Transactions of the Association for Computational Linguistics*, 7:249–266.
- Amrita Saha, Mitesh Khapra, and Karthik Sankaranarayanan. 2018a. Towards building large scale multimodal domain-aware conversation systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Amrita Saha, Vardaan Pahuja, Mitesh Khapra, Karthik Sankaranarayanan, and Sarath Chandar. 2018b. Complex sequential question answering: Towards learning to converse over linked question answer pairs with a knowledge graph. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Sebastian Schuster, Sonal Gupta, Rushin Shah, and Mike Lewis. 2019. [Cross-lingual transfer learning for multilingual task oriented dialog](#). 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 3795–3805, Minneapolis, Minnesota. Association for Computational Linguistics.
- Pararth Shah, Dilek Hakkani-Tür, Bing Liu, and Gokhan Tür. 2018. [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, New Orleans - Louisiana. Association for Computational Linguistics.
- Igor Shalyminov, Alessandro Sordani, Adam Atkinson, and Hannes Schulz. 2020. Fast domain adaptation for goal-oriented dialogue using a hybrid generative-retrieval transformer. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8039–8043. IEEE.
- Andreas Stolcke, Klaus Ries, Noah Coccaro, Elizabeth Shriberg, Rebecca Bates, Daniel Jurafsky, Paul Taylor, Rachel Martin, Carol Van Ess-Dykema, and Marie Meteer. 2000. [Dialogue act modeling for automatic tagging and recognition of conversational speech](#). *Computational Linguistics*, 26(3):339–374.
- Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Xuyi Chen, Han Zhang, Xin Tian, Danxiang Zhu, Hao

- Tian, and Hua Wu. 2019. Ernie: Enhanced representation through knowledge integration. *arXiv preprint arXiv:1904.09223*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Ivan Vulić, Pei-Hao Su, Sam Coope, Daniela Gerz, Paweł Budzianowski, Iñigo Casanueva, Nikola Mrkšić, and Tsung-Hsien Wen. 2021. Convfit: Conversational fine-tuning of pretrained language models. *arXiv preprint arXiv:2109.10126*.
- Xuwei Wang, Weiyang Shi, Richard Kim, Yoojung Oh, Sijia Yang, Jingwen Zhang, and Zhou Yu. 2019. Persuasion for good: Towards a personalized persuasive dialogue system for social good. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5635–5649, Florence, Italy. Association for Computational Linguistics.
- Wei Wei, Quoc Le, Andrew Dai, and Jia Li. 2018. [Air-Dialogue: An environment for goal-oriented dialogue research](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3844–3854, Brussels, Belgium. Association for Computational Linguistics.
- Myers Will, Etchart Tyler, and Fulda Nancy. 2020. Conversational scaffolding: An analogybased approach to response prioritization in open-domain dialogs. In *Proceedings of the 12th International Conference on Agents and Artificial Intelligence (ICAART)*.
- Jason D Williams, Antoine Raux, and Matthew Henderson. 2016. The dialog state tracking challenge series: A review. *Dialogue & Discourse*, 7(3):4–33.
- Chien-Sheng Wu, Steven C.H. Hoi, Richard Socher, and Caiming Xiong. 2020. [TOD-BERT: Pre-trained natural language understanding for task-oriented dialogue](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 917–929, Online. Association for Computational Linguistics.
- Xing Wu, Chaochen Gao, Liangjun Zang, Jizhong Han, Zhongyuan Wang, and Songlin Hu. 2021. Esimcse: Enhanced sample building method for contrastive learning of unsupervised sentence embedding. *arXiv preprint arXiv:2109.04380*.
- Alex Xiao, Christian Fuegen, and Abdelrahman Mohamed. 2021. Contrastive semi-supervised learning for asr. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3870–3874. IEEE.
- Yuanmeng Yan, Rumei Li, Sirui Wang, Fuzheng Zhang, Wei Wu, and Weiran Xu. 2021. Consert: A contrastive framework for self-supervised sentence representation transfer. *arXiv preprint arXiv:2105.11741*.
- Steve Young. 2007. Cued standard dialogue acts. *Report, Cambridge University Engineering Department, 14th October, 2007*.
- Steve Young, Milica Gašić, Blaise Thomson, and Jason D Williams. 2013. Pomdp-based statistical spoken dialog systems: A review. *Proceedings of the IEEE*, 101(5):1160–1179.
- Tao Yu, Rui Zhang, Heyang Er, Suyi Li, Eric Xue, Bo Pang, Xi Victoria Lin, Yi Chern Tan, Tianze Shi, Zihan Li, Youxuan Jiang, Michihiro Yasunaga, Sungrok Shim, Tao Chen, Alexander Fabbri, Zifan Li, Luyao Chen, Yuwen Zhang, Shreya Dixit, Vincent Zhang, Caiming Xiong, Richard Socher, Walter Lasecki, and Dragomir Radev. 2019. [CoSQL: A conversational text-to-SQL challenge towards cross-domain natural language interfaces to databases](#). 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 1962–1979, Hong Kong, China. Association for Computational Linguistics.
- 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*.
- Xinkai Yuan, Zilinghan Li, and Gaoang Wang. 2021. Activematch: End-to-end semi-supervised active representation learning. *arXiv preprint arXiv:2110.02521*.
- Guangtao Zeng, Wenmian Yang, Zeqian Ju, Yue Yang, Sicheng Wang, Ruisi Zhang, Meng Zhou, Jiaqi Zeng, Xiangyu Dong, Ruoyu Zhang, Hongchao Fang, Penghui Zhu, Shu Chen, and Pengtao Xie. 2020. [MedDialog: Large-scale medical dialogue datasets](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9241–9250, Online. Association for Computational Linguistics.
- Jianguo Zhang, Trung Bui, Seunghyun Yoon, Xiang Chen, Zhiwei Liu, Congying Xia, Quan Hung Tran, Walter Chang, and Philip Yu. 2021a. Few-shot intent detection via contrastive pre-training and fine-tuning. *arXiv preprint arXiv:2109.06349*.
- Jianguo Zhang, Kazuma Hashimoto, Wenhao Liu, Chien-Sheng Wu, Yao Wan, Philip Yu, Richard Socher, and Caiming Xiong. 2020. [Discriminative nearest neighbor few-shot intent detection by transferring natural language inference](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5064–5082, Online. Association for Computational Linguistics.
- Sai Zhang, Yuwei Hu, Yuchuan Wu, Jiaman Wu, Yongbin Li, Jian Sun, Caixia Yuan, and Xiaojie Wang. 2022. [A slot is not built in one utterance: Spoken language dialogs with sub-slots](#). In *Findings of*

the Association for Computational Linguistics: ACL 2022, pages 309–321, Dublin, Ireland. Association for Computational Linguistics.

Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. [Personalizing dialogue agents: I have a dog, do you have pets too?](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.

Zhenyu Zhang, Tao Guo, and Meng Chen. 2021b. Dialoguebert: A self-supervised learning based dialogue pre-training encoder. *arXiv preprint arXiv:2109.10480*.

Kangyan Zhou, Shrimai Prabhumoye, and Alan W Black. 2018. [A dataset for document grounded conversations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 708–713, Brussels, Belgium. Association for Computational Linguistics.

Yanning Zhou, Hang Xu, Wei Zhang, Bin Gao, and Pheng-Ann Heng. 2021. C3-semiseg: Contrastive semi-supervised segmentation via cross-set learning and dynamic class-balancing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7036–7045.

Pengfei Zhu, Binyuan Hui, Changqing Zhang, Dawei Du, Longyin Wen, and Qinghua Hu. 2019. Multi-view deep subspace clustering networks. *arXiv preprint arXiv:1908.01978*.

A Details for Pre-training Dialog Corpora

All dialog corpora we used in our experiments are given in Table 5 and Table 6. There are totally 32 labeled datasets and 21 unlabeled datasets. For the labeled data, each dialog sample is unified into the same data format with tree-structured labels. For the unlabeled data, each sample is performed careful cleaning due to the high noise. The processing methods include: (1) removing the instances where there is a URL in utterances. (2) removing the instances containing word repetitions of at least three words. (3) removing non-English sentences. (4) removing sentences containing special markers such as “[” or “]”, as this could be markup. (5) removing offensive language. (6) replacing the non-unicode characters like emojis.

B Input Representations

The input embeddings consist of four elements: tokens, roles, turns, and positions. Role embeddings, E_u , E_s , are like segmentation embeddings in BERT and are used to differentiate which role the current token belongs to either user or system. Turn embeddings are assigned to each token according to its adverse turn number in the input dialog data. Position embeddings are assigned to each token according to its relative position within its belonging sentence. For the tokens, we use special tokens [BOU], [EOU] to bound the user utterance, and [BOS], [EOS] to bound the system response.

C Details for Compared Baselines

C.1 Baselines for Dialog Understanding

We choose the following baselines: 1) **BERT**, a BERT_{base} model adapted to all downstream tasks; 2) **ConvBERT**, a BERT continually pre-trained on a large open-domain dialogue corpus with MLM (Mehri et al., 2020); 3) **ConvBERT-DG**, a ConvBERT continually pre-trained on all DialoGLUE data before conducting specific downstream tasks with MLM. 4) **ConvBERT+MLM**, a ConvBERT continually trained on training data of specific downstream tasks; 5) **ConvBERT-DG+MLM**, a ConvBERT+MLM pre-trained on all DialoGLUE data before conducting downstream tasks; and 6) **TOD-BERT**, a BERT pre-trained on nine existing task-oriented dialogue corpora using MLM and response selection loss (Wu et al., 2020).

C.2 Baselines for Specific Tasks

In this paper, we focus on intent prediction and slot-filling in particular since there are adequate pre-trained baselines specifically. For the intent prediction task, we choose: 1) **Example+Observer**, a ConvBERT+MLM that uses example-driven training based on similarity matching and observers for transformer attentions (Mehri and Eric, 2021); 2) **USE+ConveRT**, a dual sentence encoder pre-trained on 654M Reddit dialogs (Casanueva et al., 2020); 3) **DNNC**, a discriminative nearest-neighbor model which finds the best-matched sample via similarity matching (Zhang et al., 2020); 4) **CPFT**, which uses self-supervised contrastive learning to perform pre-training and fine-tuning (Zhang et al., 2021a); and 5) **ConvFit**, which uses regular supervised contrastive learning on labeled data after pre-training on chatting data (Vulić et al., 2021). For the slot-filling task, we choose: 1) **Span-ConveRT**, a slot-filling model pre-trained on collected data (Coope et al., 2020); 2) **ConVEx**, a span-ConveRT pre-trained with the pairwise cloze task on Reddit for few-shot slot labeling (Henderson and Vulić, 2021); and 3) **GenSF**, a generative pre-trained slot-filling model (Mehri and Eskenazi, 2021).

Dataset	# Dialog	# Utterance
BANKING77 (Casanueva et al., 2020)	13,083	13,083
CLINIC150 (Larson et al., 2019)	23,700	23,700
HWU64 (Liu et al., 2021)	25,716	25,716
REST8K (Coope et al., 2020)	11,115	11,975
TOP(Gupta et al., 2018)	44,783	44,783
ATIS(Hemphill et al., 1990)	5,817	5,817
SNIPS(Coucke et al., 2018)	14,484	14,484
CrossNER(Liu et al., 2020b)	27,458	27,458
FB_TOD_SF(Schuster et al., 2019)	43,323	43,323
MIT-restaurant(Liu et al., 2013)	9,181	9,181
MIT-movies-eng(Liu et al., 2013)	12,218	12,218
MIT-movies-trival10k13(Liu et al., 2013)	9,769	9,769
MULTIWOZ _coco(Li et al., 2020)	32,062	64,124
MULTIWOZ (Eric et al., 2020)	10,433	142,968
STAR(Mosig et al., 2020)	5,820	98,962
DailyDialog(Li et al., 2017)	13,118	102,979
SGD(Rastogi et al., 2020)	22,825	463,284
Frames(El Asri et al., 2017)	1,369	19,986
MSRe2e(Li et al., 2018)	10,087	74,686
DSTC2(Williams et al., 2016)	4,953	73,228
DSTC3(Williams et al., 2016)	1,980	30,456
SimJoint(Shah et al., 2018)	3,008	27,120
MulDoGo(Peskov et al., 2019)	11,668	194,097
WOZ(Mrkšić et al., 2017)	1,200	8,824
TaskMaster1(Byrne et al., 2019b)	13,210	279,287
TaskMaster2(Byrne et al., 2019b)	17,289	292,830
TaskMaster3(Byrne et al., 2019b)	23,757	477,026
InCar(Eric et al., 2017)	3,031	15,928
MULTIWOZ _synthesis(Campagna et al., 2020)	37,605	401,075
SwDA(Stolcke et al., 2000)	1,434	274,786
BiTOD(Lin et al., 2021)	3,689	72,462
PersuaGOOD(Wang et al., 2019)	300	10,864
Total	459,465	3,366,479

Table 5: Statistics for each labeled dataset in *AnPreDial*.

Name	# Dialog	# Utterance
MulDoGo_un(Peskov et al., 2019)	63,404	1,013,985
ABCD(Chen et al., 2021)	10,042	143,855
AirDialog(Wei et al., 2018)	361,822	4,645,475
CCPE(Radlinski et al., 2019)	502	11,971
MetalWOZ(Shalyminov et al., 2020)	40,203	458,237
CMU_DoG(Zhou et al., 2018)	4,221	134,197
ConvQuestions(Kacupaj et al., 2021)	11,200	112,000
CoQA(Reddy et al., 2019)	7,699	233,260
CoSQL(Yu et al., 2019)	2,458	25,433
doc2dial(Feng et al., 2020)	688	52,688
DSTC10-track2(Kim et al., 2021b)	2,919	55,377
DSTC10-track3(Kottur et al., 2021)	11,244	110,767
MedicalDialog(Zeng et al., 2020)	482	981
Self-Dialog(Fainberg et al., 2018)	24,165	372,720
WOW(Dinan et al., 2018)	22,311	201,999
TopicChat(Gopalakrishnan et al., 2019)	8,628	188,378
Persona-Chat(Zhang et al., 2018)	11,087	81,032
MMD(Saha et al., 2018a)	1,506,129	5,477,066
CSQA(Saha et al., 2018b)	197,001	4,153,092
AmazonQA(Gupta et al., 2019)	923,685	1,847,370
ChitChat(Will et al., 2020)	7,168	258,145
Total	3,217,058	19,578,028

Table 6: Statistics for each unlabeled dataset in *UnPreDial*.