

RiSAWOZ: A Large-Scale Multi-Domain Wizard-of-Oz Dataset with Rich Semantic Annotations for Task-Oriented Dialogue Modeling

Jun Quan^{†*}, Shian Zhang^{†*}, Qian Cao[†], Zizhong Li[‡] and Deyi Xiong^{††}

[†]School of Computer Science and Technology, Soochow University, Suzhou, China

[‡]College of Intelligence and Computing, Tianjin University, Tianjin, China

{terryqj0107, sazhang1996, caoqian0905}@gmail.com,
zizhong1208@126.com, dyxiong@tju.edu.cn

Abstract

In order to alleviate the shortage of multi-domain data and to capture discourse phenomena for task-oriented dialogue modeling, we propose **RiSAWOZ**, a large-scale multi-domain Chinese Wizard-of-Oz dataset with Rich Semantic Annotations. RiSAWOZ contains 11.2K human-to-human (H2H) multi-turn semantically annotated dialogues, with more than 150K utterances spanning over 12 domains, which is larger than all previous annotated H2H conversational datasets. Both single- and multi-domain dialogues are constructed, accounting for 65% and 35%, respectively. Each dialogue is labeled with comprehensive dialogue annotations, including dialogue goal in the form of natural language description, domain, dialogue states and acts at both the user and system side. In addition to traditional dialogue annotations, we especially provide linguistic annotations on discourse phenomena, e.g., ellipsis and coreference, in dialogues, which are useful for dialogue coreference and ellipsis resolution tasks. Apart from the fully annotated dataset, we also present a detailed description of the data collection procedure, statistics and analysis of the dataset. A series of benchmark models and results are reported, including natural language understanding (intent detection & slot filling), dialogue state tracking and dialogue context-to-text generation, as well as coreference and ellipsis resolution, which facilitate the baseline comparison for future research on this corpus.¹

1 Introduction

In recent years, we have witnessed that a variety of datasets tailored for task-oriented dialogue have been constructed, such as MultiWOZ (Budzianowski et al., 2018), SGD (Rastogi et al.,

2019a) and CrossWOZ (Zhu et al., 2020), along with the increasing interest in conversational AI in both academia and industry (Gao et al., 2018). These datasets have triggered extensive research in either end-to-end or traditional modular task-oriented dialogue modeling (Wen et al., 2019; Dai et al., 2020). Despite of substantial progress made based on these newly built corpora, more efforts in creating challenging datasets in terms of size, multiple domains, semantic annotations, multilinguality, etc., are still in demand (Wen et al., 2019).

Among the existing datasets, the majority of them are not large in size, e.g., ATIS (Hemphill et al., 1990), WOZ 2.0 (Wen et al., 2017), FRAMES (El Asri et al., 2017) and KVRET (Eric et al., 2017), which might not well support data-hungry neural dialogue models. Very large task-oriented dialogue datasets can be created in a machine-to-machine fashion, such as M2M (Shah et al., 2018) and SGD (Rastogi et al., 2019a). Datasets collected in this way need to simulate both user and system and contain unnatural conversations.

MultiWOZ (Budzianowski et al., 2018), probably the most promising and notable dialogue corpus collected in a Wizard-of-Oz (i.e., Human-to-Human) way recently, is one order of magnitude larger than the aforementioned corpora collected in the same way. However, it contains noisy system-side state annotations and lacks user-side dialogue acts² (Eric et al., 2019; Zhu et al., 2020). Yet another very recent dataset CrossWOZ (Zhu et al., 2020), the first large-scale Chinese H2H dataset for task-oriented dialogue, provides semantic annotations on both user and system side although it is relatively smaller than MultiWOZ. The number of domains in both MultiWOZ and CrossWOZ is

*Equal Contributions.

¹The corpus is publicly available at <https://github.com/terryqj0107/RiSAWOZ>.

²In MultiWOZ 2.1, Eric et al. (2019) re-annotate utterances to fix the noisy state annotation problem via crowdsourced workers.

Type	Single-domain					Multi-domain			
Dataset	DSTC2	WOZ 2.0	FRAMES	KVRET	M2M	MultiWOZ	SGD	CrossWOZ	RiSAWOZ (ours)
Language	EN	EN	EN	EN	EN	EN	EN	ZH	ZH
Speakers	H2M	H2H	H2H	H2H	M2M	H2H	M2M	H2H	H2H
Domains	1	1	1	3	2	7	16	5	12
Dialogues	1,612	600	1,369	2,425	1,500	8,438	16,142	5,012	10,000
Turns	23,354	4,472	19,986	12,732	14,796	115,424	329,964	84,692	134,580
Avg. turns	14.5	7.5	14.6	5.3	9.9	13.7	20.4	16.9	13.5
Slots	8	4	61	13	14	25	214	72	159
Values	212	99	3,871	1,363	138	4,510	14,139	7,871	4,061
Linguistic annotation	No	No	No	No	No	No	No	No	Yes

Table 1: Comparison of our dataset to other task-oriented dialogue datasets (training set). H2H, H2M, M2M represent human-to-human, human-to-machine, machine-to-machine respectively.

fewer than 10. MultiWOZ dialogues cover 7 domains. However, the distribution of dialogues over these domains is imbalanced. Dialogues from two domains (*hospital*, *police*) account for less than 6% in the training set and don't appear in either the development or test set. CrossWOZ involves 5 domains and dialogue goal descriptions for the domain *taxi* and *metro* are relatively simple than those from other domains. Neither MultiWOZ nor CrossWOZ provide linguistic annotations to capture discourse phenomena which are ubiquitous in multi-turn dialogues and are important in dialogue modeling (Quan et al., 2019; Su et al., 2019; Rastogi et al., 2019b; Zhang et al., 2019a)

In order to alleviate the aforementioned issues, we propose **RiSAWOZ**, a large-scale Chinese multi-domain Wizard-of-Oz task-oriented dialogue dataset with rich semantic annotations. Compared with existing datasets (particularly MultiWOZ and CrossWOZ), our contributions can be summarized as follows:

- RiSAWOZ is to date the largest fully annotated human-to-human task-oriented dialogue dataset to our knowledge. It contains 11,200 multi-turn dialogues with more than 150K utterances ranging over 12 domains, namely *Attraction*, *Restaurant*, *Hotel*, *Flight*, *Train*, *Weather*, *Movie*, *TV*, *Computer*, *Car*, *Hospital* and *Education* (particularly after-school remedial courses), several of which are not covered in previous datasets. Compared with other Wizard-of-Oz datasets (e.g. MultiWOZ and CrossWOZ), RiSAWOZ offers a wider and more balanced coverage on both domains and slots, making it suitable for not only practical use in industrial scenarios of related domains but also research on domain adaptation, few/zero-shot learning in task-oriented dia-



Figure 1: A dialogue example spanning over multiple domains. We show dialogue annotations and necessary linguistic annotations (green words for ellipsis resolution and green box for coreference clusters) for each user utterance (in yellow callout) and system utterance (in blue callout). Better viewed in color.

logue. Detailed comparison of RiSAWOZ to existing task-oriented dialogue datasets is shown in Table 1.

- We provide richer manual semantic annotations on the crowd-sourced dialogues, including both dialogue annotations (i.e., various structured semantic labels for dialogue modeling) and linguistic annotations that are not available in previous Wizard-of-Oz datasets (e.g., MultiWOZ or CrossWOZ). Figure 1

shows a dialogue example that demonstrates semantic annotations in RiSAWOZ. User goal description, domain label, dialogue states and dialogue acts at both user and system side are annotated for each dialogue. In order to facilitate the study of ellipsis and coreference in dialogue, we also provide two kinds of linguistic annotations collected in two different ways. Annotations for unified ellipsis/coreference resolution via utterance rewriting are more comprehensive and at least one order of magnitude larger than existing datasets with such annotations (Quan et al., 2019; Su et al., 2019; Zhang et al., 2019a; Rastogi et al., 2019b). Coreference clusters in each dialogue are also manually annotated, providing a new large-scale coreference dataset on dialogue, which is complementary to previous coreference datasets on documents (Pradhan et al., 2012). In a nutshell, RiSAWOZ integrates human-to-human conversations, dialogue annotations and linguistic annotations on ellipsis/coreference into a single unified dataset.

- We use RiSAWOZ as a new benchmark testbed and report benchmark results on 5 tasks for future comparison study and tracking progress on this dataset. The 5 tasks are NLU, DST, Dialogue Context-to-Text Generation, Coreference Resolution and Unified Generative Ellipsis and Coreference Resolution. We discuss the usability of the dataset for other tasks, e.g., Dialogue Policy Learning, Natural Language Generation, User Simulator, Dialogue Summarization, etc. The dataset and the benchmark models will be publicly available soon.

2 Related Work

We follow Budzianowski et al. (2018) to roughly categorize existing task-oriented dialogue datasets into three groups: machine-to-machine, human-to-machine, and human-to-human. From the perspective of domain quantity and data scale, most existing datasets cover only one single or a few domains while large-scale multi-domain datasets are not widely available. As suggested by Wen et al. (2019), task-oriented dialogue datasets in languages other than English are few. To the best of our knowledge, there has been no large-scale dialogue datasets with linguistic annotations aiming at ubiquitous discourse phenomena (e.g., ellipsis

and coreference) in dialogue. Although some recent works have proposed datasets with utterance completion annotation for ellipsis or coreference in dialogue (Quan et al., 2019; Su et al., 2019), these datasets are at small scale and with simple dialogue goals. No dialogue datasets provide annotations of coreference clusters.

Machine-to-Machine Collecting data of this type requires to create a user and system simulator to generate multi-turn dialogue outlines, which are further transformed into natural language utterances via paraphrasing with predefined rules or crowdsourced workers (Shah et al., 2018; Rastogi et al., 2019a). Despite of less human effort required in this approach, the diversity and complexity of created dialogues greatly depend on the quality of user and system simulators. It’s also difficult to avoid mismatch between machine-created dialogues and real human conversations.

Human-to-Machine In this method, humans converse with an existing dialogue system to collect dialogue data. The Dialogue State Tracking Challenges (DSTC) has provided several datasets created in this way (Williams et al., 2013; Henderson et al., 2014a,b). Generally, the quality of human-to-machine data heavily relies on the performance of the given dialogue system.

Human-to-Human To collect data of this type, crowdsourced workers talk to each other according to given dialogue goals to create diverse and natural dialogues. ATIS (Hemphill et al., 1990), WOZ 2.0 (Wen et al., 2017), FRAMES (El Asri et al., 2017) and KVRET (Eric et al., 2017) are small-scale datasets built in this way. In contrast, MultiWOZ Budzianowski et al. (2018) and CrossWOZ (Zhu et al., 2020) are two large-scale H2H datasets proposed recently.

Coreference Resolution Coreference is ubiquitous in dialogue. However, there is no available dialogue dataset with labeled coreference clusters. Generally, coreference datasets are created on text paragraphs or documents. The OntoNotes 5.0 dataset (Pradhan et al., 2012) is one of the most widely-used document-level dataset on coreference resolution from the CoNLL-2012 shared task.

Generative Ellipsis and Coreference Resolution In recent years, ellipsis and coreference resolution in dialogue has been treated as an end-to-end generative task. Su et al. (2019) propose to rewrite dialogue utterances to recover all co-referred and omitted information with an annotated Chinese

Domain	Attraction	Restaurant	Hotel	Flight	Train	Weather	Movie	TV	Computer	Car	Hospital	Courses
Entities	40	50	70	665	2,016	70	100	105	50	52	17	90
Slots	12	11	11	9	10	6	10	11	20	23	16	20

Table 2: Database statistics on entities and slots.

Domain	Informable Slots	Requestable Slots
Attraction	name, area, type, consumption, metro station	metro station, ticket price, phone number, address, score, opening hours, features
Restaurant	name, area, cuisine, pricerange, metro station	metro station, per capita consumption, address, phone, score, business hours, dishes
Hotel	name, area, star, pricerange, hotel type, room type, parking	parking, room charge, address, phone, score
Flight	departure, destination, date, class cabin	flight information, departure time, arrival time, ticket price, punctuality rate
Train	departure, destination, date, train type, seat type	train number, duration, departure time, arrival time, ticket price
Weather	city, date	weather condition, temperature, wind, UV intensity
Movie	production country or area, type, decade, movie star	director, movie title, name list, release date, film length, Douban score
TV	production country or area, type, decade, tv star	director, TV title, name list, premiere time, episodes, episode length, Douban score
Computer	brand, computer type, usage, memory capacity, screen size, CPU category, pricerange, series	product name, operating system, game performance, CPU model, GPU category, GPU model, features, colour, standby time, hard disk capacity, weight
Car	series, classification, size, number of seats, brand, hybrid, power level, 4WD, fuel consumption, price range	parking assist system, cruise control system, heated seats, ventilated seats
Hospital	name, level, type, public or private, area, general or specialized, key departments, metro station	address, phone, registration time, service time, bus routes, CT, 3.0T MRI, DSA
Class	grade, subject, level, Day, Time, area	campus, start date, end date, start time, end time, times, hours, classroom, teacher, price

Table 3: All informable and requestable slots in each domain.

chit-chat dialogue dataset. [Quan et al. \(2019\)](#) annotate an off-the-shelf task-oriented dialogue dataset CamRest676 ([Wen et al., 2016a,b](#)) with coreference and ellipsis information and propose an end-to-end generative resolution model for both ellipsis and coreference in a single unified framework. This task is also treated as an auxiliary module to improve dialogue understanding ([Zhang et al., 2019a](#)) and dialogue state tracking ([Rastogi et al., 2019b](#)). However, the scale of these built or used dialogue datasets is relatively small.

3 Dataset Creation

The whole process of data collection consists of database and ontology construction, goal generation, dialogue collection and two rounds of annotations.

3.1 Database and Ontology Construction

We crawl 3,325 unique entities with their attributes from several Chinese public websites. Statistics on entities and slots across 12 domains are shown in Table 2. An ontology is constructed over these entities and attributes, which defines all possible slots for each domain and all possible values for each

slot. Slots in the dataset can be divided into two categories: informable and requestable, as shown in Table 3. Informable slots are attributes that allow the user to constrain the search into the database. Requestable slots represent specific attributes that the user wants to know about an entity.

3.2 Dialogue Goal

First of all, we design dialogue goal templates with placeholders representing slots and values. We have designed 80 dialogue goal templates for 12 domains, including 52 single-domain goals and 28 multi-domain goals. Then we randomly sample actual slots and values from the ontology to fill in the placeholders in the templates to generate dialogue goal instances. We finally generate 5,600 dialogue goal instances. An example of dialogue goal (i.e. user goal) is shown at the top of Figure 1. A dialogue goal is a natural language description that only the user can see. The user needs to talk with the system step by step according to the given goal description until the task is finished. We assign each dialogue goal to two different pairs of workers to accomplish. In this way, we can collect two different dialogues from one dialogue goal. The

User action type	Inform, request, greeting, bye, general
System action type	Inform, recommend, request, no-offer, greeting, bye, general

Table 4: Dialogue act types.

total of dialogues is therefore $2 * 5,600 = 11,200$. This setting can ensure the diversity of collected dialogues while the amount of cost for crowdsourced workers is under the budget.

3.3 Dialogue Collection and Annotation

In order to collect high-quality coherent dialogues and annotations, we develop a collecting platform based on the Client-Server architecture, including two versions of client platform for user and system side respectively. Crowdsourced workers can choose to play the role of either user or system. They work in pairs and enter the chat room to construct dialogues. To ensure the quality of dialogue collection, we hire workers via our in-house crowdsourcing platform and train the workers strictly in advance. Finally, we select 92 well-trained workers to participate in our dialogue collection and annotation. At this stage, we collect task-oriented dialogue data with dialogue annotations including domain labels, dialogue states and acts.

3.3.1 User Side

During dialogue collection, a user first reads through the natural language description of a given dialogue goal to understand the task that is required to finish. After that, the user communicates with the system step by step to accomplish the given goal. We encourage the users to follow their own personalized language style in communication and train different workers to play the role of user, which makes created dialogues more complex and diverse, and more similar to the spoken conversations in our daily life. According to the instructions described in the given dialogue goal, the user should provide specific constraints to the system step by step and request the corresponding information. The user can terminate the dialogue when he/she believes the task has been accomplished.

3.3.2 System Side

The worker who plays the role of system (i.e. wizard) provides consulting services in various domains to users. When receiving an utterance from the user side, the wizard needs to first determine which domain the user is talking about and convert

the user utterance into structured user acts. A dialogue act for both user and system side consists of the act type (i.e. intent) and slot-value pairs such as *inform (area=Gusu District)*. All the dialogue act types are shown in Table 4. We define 5 different user act types. **Inform** represents that a user provides specific constraints to information search from the database. **Request** denotes that a user asks for the values of specific slots. **Greeting** and **bye** are to express greeting and farewell. **General** represents other behaviors that are not covered above.

By understanding the goal of a user utterance, the wizard needs to annotate the constraints the user wants to provide and the slots requested by the user. The constraints are called belief states which are a set of slot-value pairs. The belief state is persistent across turns and is used to query the database. The wizard then retrieves the database according to the constraints. Considering both the results of database retrieval and the dialogue context, the wizard should send a natural language response to the user. In addition, the wizard needs to convert the natural language system response into structured system acts.

Similar to the user acts, 7 different system act types are predefined. **Inform** represents that the system informs the user about the attribute values of specific entities. **Recommend** denotes that the system recommends required entities to the user. **Request** represents that the system asks the user whether there are special constraints for the slots in question. **No offer** means that the system tells the user there is no matched entity. The remaining three act types, **greeting**, **bye** and **general**, are the same as those user act types described above.

Although the tasks for the system side look complex at the first glance, we design a simple and easy-to-operate graphical user interface (GUI). The wizard only needs to follow the prompts to perform simple operations, such as checking the multi-check box, picking the drop-down box, filling in the input box, clicking specific buttons, etc., which can be easily done by well-trained workers.

In this way, the information of domains, belief states, dialogue acts for both user and system side can be annotated during the process of collecting dialogues. A dialogue example with these annotations are demonstrated in Figure 1.

Different from the data collecting way that multiple workers contribute to one dialogue adopted

	Train	Dev	Test	Single-domain	Multi-domain	All
Dialogues	10,000	600	600	7,261	3,939	11,200
Turns	134,580	8,116	9,286	88,618	63,364	151,982
Tokens	1,462,727	87,886	108,032	954,298	704,347	1,658,645
Vocab.	11,486	4,514	4,927	10,065	8,031	11,971
Avg. turns	13.46	13.53	15.48	12.20	16.09	13.57
Avg. tokens	10.87	10.83	11.63	10.77	11.12	10.91
Avg. acts	1.46	1.46	1.54	1.43	1.51	1.46
Avg. u-acts	1.79	1.79	1.85	1.73	1.88	1.79
Avg. s-acts	1.13	1.12	1.23	1.13	1.14	1.13
Coref. clusters	17,561	1,082	1,230	10,982	10,269	20,915
Avg. c-clusters	1.76	1.8	2.05	1.42	2.45	1.77

Table 5: Data statistics. The average numbers of turns and coreference clusters are for each dialogue. The average numbers of tokens and dialogue acts are for each turn. U-acts and s-acts represent user and system acts respectively.

by MultiWOZ (Budzianowski et al., 2018), in our dataset, the construction and annotation for each dialogue are completed by a pair of well-trained workers. This is to guarantee the coherence and consistency of each created dialogue and the accuracy of the annotation for them. Moreover, we train each worker to play different roles to diversify dialogue utterances.

3.4 Linguistic Annotation

3.4.1 Coreference Clusters Annotation

We develop a toolkit with easy-to-operate GUI for annotating coreference clusters. With this annotation toolkit, well-trained annotators read through a dialogue to locate all entity mentions. They then group each of these mentions into an appropriate cluster. As shown at the bottom of Figure 1, entity mentions in each cluster are co-referential to one another.

3.4.2 Ellipsis and Coreference Annotation via Utterance Rewriting

As shown in Figure 1, both referenced and absent information can be recovered by rewriting an incomplete utterance into a complete version. In this way, we can reformulate ellipsis and coreference resolution as sentence rewriting in a unified framework. The merit of such rewriting is to help the dialogue model better understand the goal of a user utterance in context.

In order to facilitate such task reformulation, we provide the second type of linguistic annotation on RiSAWOZ: utterance rewriting for ellipsis and coreference resolution. We train crowd-sourced workers to accomplish this annotation task and develop an annotation toolkit for them. Each annotator needs to read an entire dialogue sentence

	User Utterances	Rate (%)
Ellipsis	23,181	30.50
Coreference	19,993	26.31
Both	3,582	4.71
Neither	29,235	38.47
Total	75,991	100

Table 6: Statistics on utterances containing ellipsis and coreference.

by sentence, detecting ellipsis or coreference phenomena in user utterances. For an utterance with ellipsis or coreference, the annotator rewrites the utterance into its complete version with recovered referenced/absent information according to dialogue context. If none of them occurs in the user utterance, the original utterance is kept. Both cases are presented in the example in Figure 1.

4 Our Dialogue Dataset

Our dataset contains not only single-domain dialogues, but also a great amount of multi-domain dialogues where domains are naturally connected. For example, a user wants to travel from one place to another. After checking the air ticket or train ticket, she wants to ask for the local weather information as well. In this section, we will introduce our dataset from two aspects: data structure and data statistics.

4.1 Data Structure

As shown in Figure 1, each dialogue in our dataset consists of a user goal description in natural language, a label of dialogue domain, multiple user and system turns and a set of coreference clusters. In each user turn, the user act and dialogue state are annotated over the user utterance. We also label whether there are ellipsis or coreference phe-

Task	Model	Metrics	Single-domain	Multi-domain	All data
Natural Language Understanding	BERT + context	Dialogue act F ₁	82.64	81.68	82.15
			84.63	83.61	84.10
Dialogue State Tracking	TRADE (rand init) TRADE (fastText) MLCSG (fastText)	Joint Accuracy	65.35	50.49	58.19
			68.55	50.94	60.07
			73.04	58.77	66.16
Context-to-Text Generation	DAMD	Inform	79.19	65.00	73.73
		Success	55.66	54.68	55.18
		BLEU	32.90	22.00	27.90
		Combined Score	100.33	81.84	92.36
Coreference Resolution	e2e-coref	MUC F ₁	86.62	83.00	84.68
		B ³ F ₁	83.75	79.03	81.24
		CEAF _{φ₄} F ₁	83.11	79.78	81.29
		Avg. F ₁	84.49	80.60	82.41
Unified Generative Ellipsis and Coreference Resolution	GECOR	EM	61.65	54.61	58.26
		BLEU	87.75	86.19	87.50
		Resolution_F ₁	77.00	79.15	78.14

Table 7: Performance of benchmark models on single-domain, multi-domain and all dialogues of test set.

nomena in each user utterance. If so, a complete version of the user utterance is provided. In each system turn, the system utterance is labeled with the corresponding system acts.

4.2 Data Statistics

Dialogue Statistics: We reshuffle all created dialogues and divide them into the training/dev/test sets which maintain approximately the same distribution on domains. As shown in Table 5, the training set contains 10,000 dialogues with 134,580 turns. The development and test set contain 600 dialogues with more than 8K and 9K turns respectively. The 5th column of Table 5 shows the statistics on single-domain dialogues. Multi-domain dialogues (the 6th column of Table 5) cover 8 domains excluding *Computer*, *Car*, *Hospital* and *Education*.

After Chinese word segmentation via Jieba,³ there are 1,658,645 tokens in total in RiSAWOZ. On average, there are 10.91 tokens in each turn and 13.57 turns in each dialogue. Multi-domain dialogues have more turns and utterances are longer than those in single-domain dialogues.

Annotation Statistics: As shown in Table 5, each dialogue contains an average of 1.46 dialogue acts per turn. Each user and system turn have an average of 1.79 and 1.13 dialogue acts respectively. The richness of dialogue acts also make our data set a new challenge. On average, there are 1.77 coreference clusters in each dialogue. As multi-domain dialogues are more complex, each dialogue has an average of 2.45 coreference clusters. Regarding utterance rewriting for ellipsis and coreference res-

olution, 75,991 user utterances are reformulated, as shown in Table 6. Only 38.47% of the user utterances have neither ellipsis nor coreference phenomena, and the remaining 61.53% have at least one of them.

5 RiSAWOZ as a New Benchmark

The large size and rich semantic annotations of RiSAWOZ make it a suitable testbed for various benchmark tasks. We conduct five different evaluation tasks with the benchmark models and in-depth analyses on RiSAWOZ in this section. We also discuss the applicability of RiSAWOZ for other tasks. Results of the 5 tasks are reported in Table 7.

5.1 Natural Language Understanding

Task Definition: In task-oriented dialogue system, the NLU module aims to convert the user utterance into the representation that computer can understand, which includes intent and dialogue act (slot & value) detection.

Model: We adapt BERT (Devlin et al., 2019) for the NLU task (intent detection and slot filling). We initialize BERT with the Chinese pre-trained BERT model (Cui et al., 2019) and then finetune it on RiSAWOZ. To take dialogue history into account, we employ the same BERT to model previous dialogue context. We also experiment on the situation without context. For fine-tuning BERT on RiSAWOZ, we set the learning rate to 0.00003 and the dropout rate to 0.1.

Results: From Table 7, we can clearly find that the model using dialogue context performs better than not. Also, the model obtains lower F₁ scores

³<https://github.com/fxsjy/jieba>

on multi-domain dialogues than single-domain dialogues.

5.2 Dialogue State Tracking

Task Definition: Dialogue State Tracking (DST) is a core component in task-oriented dialogue systems, which extracts dialogue states (user goals) embedded in dialogue context. It has progressed toward open-vocabulary or generation-based DST where state-of-the-art models can generate dialogue states from dialogue context directly.

Model: To report the benchmark results of the DST task, we implement the TRADE model (Wu et al., 2019) and the MLCSG model (Quan and Xiong, 2020) which improves long context modeling through a multi-task learning framework based on TRADE and achieves the state-of-the-art joint accuracy on the MultiWOZ 2.0 dataset (Budzianowski et al., 2018). We train the models with a learning rate of 0.001 and a weight decay rate of 0.5. Early stopping and dropout are also used to prevent overfitting. The dropout rate is set to 0.2.

Results: As illustrated in Table 7, we show the joint accuracy results for the two models under two different word embedding initialization settings: random and fastText (Grave et al., 2018) initialization. When we use randomly initialized word embeddings of 100 dimensions, TRADE achieves 65.35%, 50.49% and 58.19% joint accuracy on single-domain, multi-domain and all data respectively. While using 300 dimensional pretrained word vectors from fastText, TRADE performs a little better. Under the same setting, MLCSG achieves the higher 73.04%, 58.77% and 66.16% joint accuracy. In general, the performances of the two DST models significantly drop on multi-domain dialogues.

5.3 Dialogue Context-to-Text Generation

Task Definition: We recast dialogue response generation a sequence-to-sequence problem: encoding dialogue context to decode system response.

Model: To this task, we use the Domain-Aware Multi-Decoder (DAMD) model (Zhang et al., 2019b) which achieves state-of-the-art performance on the MultiWOZ 2.0 dataset (Budzianowski et al., 2018). It’s an end-to-end model proposed to handle the multi-domain response generation problem, which uses one encoder to encode dialogue context and three decoders to decode the belief span, system action

and system response. We set the vocabulary size to 8,000 and randomly initialize 50-dimensional word embeddings. The size of hidden states is set to 100. We train the model with a learning rate of 0.005 and a decay rate of 0.5.

Results: As illustrated in Table 7, we report inform rate, success rate, BLEU (Papineni et al., 2002) and combined score for this task. The inform rate measures the percentage that the output contains the appropriate entity the user asks for, and the success rate estimates the proportion that all the requested attributes have been answered. The combined score is calculated via $(inform + success) * 0.5 + BLEU$ as an overall quality (Zhang et al., 2019b). Still, multi-domain dialogues exhibit a high difficulty level.

5.4 Coreference Resolution

Task Definition: We predict coreference clusters where all mentions are referring to the same entity for each dialogue.

Model: We use the e2e-coref model (Lee et al., 2017), which is the first end-to-end coreference resolution model, as the benchmark model for this task. The model predicts coreference clusters from texts end-to-end without using any auxiliary syntactic parser or hand-engineered mention detector. It considers all spans in a text as potential mentions and learn distributions over all possible antecedents for each mention. The whole process contains two steps: scoring potential entity mentions by calculating embedding representations of corresponding spans and estimating the score for an antecedent from pairs of span representations. The 300 dimensional word vectors from fastText (Grave et al., 2018) are used for the e2e-coref model. We set the size of hidden states to 200 and the number of layers to 3. The model is trained with a learning rate of 0.001 and a decay rate of 0.999. The dropout rate is set to 0.2.

Results: We report the standard MUC, B³ and CEAF_{φ₄} F₁ metrics using the official CoNLL-2012 evaluation scripts and an average F₁ score of the three metrics. As shown in Table 7, the e2e-coref model achieves 84.49%, 80.60%, 82.41% average F₁ score on single-domain, multi-domain and all data respectively. The model performs the worst on multi-domain dialogues where coreference links may cross different domains.

5.5 Unified Generative Ellipsis and Coreference Resolution

Task Definition: This is a new task reformulated recently (Su et al., 2019; Quan et al., 2019). It usually takes the current user utterance and dialogue context as input. If there exists ellipsis or coreference phenomena in the user utterance, a complete version of the utterance is generated according to the dialogue context. Otherwise, the original user utterance is kept.

Model: We adopt the GECOR model (Quan et al., 2019) which is an end-to-end generative ellipsis and coreference resolution model with two encoders and one decoder which can produce a pragmatically complete user utterance via generation and copying. We set both the size of hidden states and word embeddings to 300. We use 300 dimensional fastText (Grave et al., 2018) word vectors to initialize word embeddings in the embedding layer. We train the model with a learning rate of 0.003 and a decay rate of 0.5. Early stopping is used and the dropout rate is 0.5.

Results: We follow Quan et al. (2019) to use the exact match rate (EM), BLEU (Papineni et al., 2002) and Resolution_{F1} as the evaluation metrics for this task. EM measures whether the generated utterances exactly match the ground-truth utterances. Resolution_{F1} is calculated by comparing machine-generated words with ground-truth words only from the ellipsis / co-reference part of user utterances. As shown in Table 7, the GECOR model achieves 58.26% EM score, 87.50% BLEU score and 78.14% Resolution_{F1} score on all data, which are much lower than the results reported on CamRest676 by Quan et al. (2019).

5.6 Other Tasks

Apart from the five evaluation tasks introduced above, RiSAWOZ can also facilitate the research of many other tasks. For example, the text of dialogues, as well as the annotation of dialogue states and acts, can support the study of dialog policy learning (DPL), natural language generation (NLG) and user simulator. Dialogue act, text and goal description can be potentially used for the task of dialogue summarization (Goo and Chen, 2018). RiSAWOZ is also suitable for domain adaptation, zero-shot and few-shot learning for multi-domain task-oriented dialogue modeling due to its wide domain coverage. Rich linguistic annotations of RiSAWOZ would also promote the deep integra-

tion of discourse and dialogue. We leave these tasks for our future work.

6 Conclusion

In this paper, we have presented RiSAWOZ, to date the largest human-to-human multi-domain dataset annotated with rich semantic information for task-oriented dialogue modeling. We manually label each dialogue in RiSAWOZ not only with comprehensive dialogue annotations for various sub-tasks of task-oriented dialogue systems (e.g., NLU, DST, response generation), but also linguistic annotations over ellipsis and coreference in multi-turn dialogue. In addition, the process of data creation and annotation is described in detail. We also report the benchmark models and results of five evaluation tasks on RiSAWOZ, indicating that the dataset is a challenging testbed for future work. RiSAWOZ is featured with large scale, wide domain coverage, rich semantic annotation and functional diversity, which can facilitate the research of task-oriented dialogue modeling from different aspects.

Acknowledgments

The present research was supported by the National Key Research and Development Project (Grant No.2019QY1802). We would like to thank the anonymous reviewers for their insightful comments. The corresponding author is Deyi Xiong (dyxiong@tju.edu.cn).

References

- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. [MultiWOZ - a large-scale multi-domain wizard-of-Oz dataset for task-oriented dialogue modelling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026, Brussels, Belgium. Association for Computational Linguistics.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Ziqing Yang, Shijin Wang, and Guoping Hu. 2019. Pre-training with whole word masking for chinese bert. *arXiv preprint arXiv:1906.08101*.
- Yinpei Dai, Huihua Yu, Yixuan Jiang, Chengguang Tang, Yongbin Li, and Jian Sun. 2020. A survey on dialog management: Recent advances and challenges. *arXiv preprint arXiv:2005.02233*.
- 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.
- 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, Shuyag Gao, and Dilek Hakkani-Tür. 2019. Multiwoz 2.1: Multi-domain dialogue state corrections and state tracking baselines. *arXiv preprint arXiv:1907.01669*.
- Mihail Eric, Lakshmi Krishnan, Francois Charette, and Christopher D. Manning. 2017. **Key-value retrieval networks for task-oriented dialogue**. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 37–49, Saarbrücken, Germany. Association for Computational Linguistics.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2018. Neural approaches to conversational ai. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 1371–1374.
- Chih-Wen Goo and Yun-Nung Chen. 2018. Abstractive dialogue summarization with sentence-gated modeling optimized by dialogue acts. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 735–742. IEEE.
- Edouard Grave, Piotr Bojanowski, Prakhara Gupta, Armand Joulin, and Tomas Mikolov. 2018. Learning word vectors for 157 languages. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*.
- 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, Blaise Thomson, and Jason D. Williams. 2014a. **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, Blaise Thomson, and Jason D. Williams. 2014b. The third dialog state tracking challenge. In *2014 IEEE Spoken Language Technology Workshop (SLT)*, pages 324–329. IEEE.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. End-to-end neural coreference resolution. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 188–197.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **Bleu: a method for automatic evaluation of machine translation**. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. **CoNLL-2012 shared task: Modeling multilingual unrestricted coreference in OntoNotes**. In *Joint Conference on EMNLP and CoNLL - Shared Task*, pages 1–40, Jeju Island, Korea. Association for Computational Linguistics.
- Jun Quan and Deyi Xiong. 2020. Modeling long context for task-oriented dialogue state generation. *arXiv preprint arXiv:2004.14080*.
- Jun Quan, Deyi Xiong, Bonnie Webber, and Changjian Hu. 2019. **GECOR: An end-to-end generative ellipsis and co-reference resolution model for task-oriented dialogue**. 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 4547–4557, Hong Kong, China. Association for Computational Linguistics.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2019a. Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset. *arXiv preprint arXiv:1909.05855*.
- Pushpendre Rastogi, Arpit Gupta, Tongfei Chen, and Mathias Lambert. 2019b. Scaling multi-domain dialogue state tracking via query reformulation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Industry Papers)*, pages 97–105.
- Pararth Shah, Dilek Hakkani-Tür, Gokhan Tür, Abhinav Rastogi, Ankur Bapna, Neha Nayak, and Larry Heck. 2018. Building a conversational agent overnight with dialogue self-play. *arXiv preprint arXiv:1801.04871*.
- Hui Su, Xiaoyu Shen, Rongzhi Zhang, Fei Sun, Pengwei Hu, Cheng Niu, and Jie Zhou. 2019. **Improving multi-turn dialogue modelling with utterance ReWriter**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 22–31, Florence, Italy. Association for Computational Linguistics.

- Tsung-Hsien Wen, Milica Gasic, Nikola Mrkšić, Lina M Rojas Barahona, Pei-Hao Su, Stefan Ultes, David Vandyke, and Steve Young. 2016a. Conditional generation and snapshot learning in neural dialogue systems. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2153–2162.
- Tsung-Hsien Wen, Pei-Hao Su, Paweł Budzianowski, Iñigo Casanueva, and Ivan Vulić. 2019. Data collection and end-to-end learning for conversational ai. 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): Tutorial Abstracts*, Hong Kong, China. Association for Computational Linguistics.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrksic, Milica Gasic, Lina M Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2016b. A network-based end-to-end trainable task-oriented dialogue system. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, pages 438–449.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gašić, Lina M. Rojas-Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017. [A network-based end-to-end trainable task-oriented dialogue system](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 438–449, Valencia, Spain. Association for Computational Linguistics.
- Jason Williams, Antoine Raux, Deepak Ramachandran, and Alan Black. 2013. [The dialog state tracking challenge](#). In *Proceedings of the SIGDIAL 2013 Conference*, pages 404–413, Metz, France. Association for Computational Linguistics.
- 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.
- Xiyuan Zhang, Chengxi Li, Dian Yu, Samuel Davidson, and Zhou Yu. 2019a. Filling conversation ellipsis for better social dialog understanding. *arXiv preprint arXiv:1911.10776*.
- Yichi Zhang, Zhijian Ou, and Zhou Yu. 2019b. Task-oriented dialog systems that consider multiple appropriate responses under the same context. *arXiv preprint arXiv:1911.10484*.
- Qi Zhu, Kaili Huang, Zheng Zhang, Xiaoyan Zhu, and Minlie Huang. 2020. CrossWOZ: A large-scale chinese cross-domain task-oriented dialogue dataset. *Transactions of the Association for Computational Linguistics*.