

JMultiWOZ: A Large-Scale Japanese Multi-Domain Task-Oriented Dialogue Dataset

Atsumoto Ohashi*, Ryu Hirai*, Shinya Iizuka, Ryuichiro Higashinaka

Graduate School of Informatics, Nagoya University, Japan
{ohashi.atsumoto.c0, hirai.ryu.k6, iizuka.shinya.a8}@s.mail.nagoya-u.ac.jp
higashinaka@i.nagoya-u.ac.jp

Abstract

Dialogue datasets are crucial for deep learning-based task-oriented dialogue system research. While numerous English language multi-domain task-oriented dialogue datasets have been developed and contributed to significant advancements in task-oriented dialogue systems, such a dataset does not exist in Japanese, and research in this area is limited compared to that in English. In this study, towards the advancement of research and development of task-oriented dialogue systems in Japanese, we constructed JMultiWOZ, the first Japanese language large-scale multi-domain task-oriented dialogue dataset. Using JMultiWOZ, we evaluated the dialogue state tracking and response generation capabilities of the state-of-the-art methods on the existing major English benchmark dataset MultiWOZ2.2 and the latest large language model (LLM)-based methods. Our evaluation results demonstrated that JMultiWOZ provides a benchmark that is on par with MultiWOZ2.2. In addition, through evaluation experiments of interactive dialogues with the models and human participants, we identified limitations in the task completion capabilities of LLMs in Japanese.

Keywords: Multi-domain task-oriented dialogue, Dialogue state tracking, Response generation

1. Introduction

Methods based on deep learning have been actively introduced in the research of task-oriented dialogue systems (Gao et al., 2018; Zhang et al., 2020b), which have greatly improved their performance on task completion (Zhang et al., 2020a; Hosseini-Asl et al., 2020; He et al., 2022). Task-oriented dialogue datasets are essential for developing these neural models, and a number of single-domain task-oriented dialogue datasets in English have been developed previously (Henderson et al., 2014; Wen et al., 2017; Eric et al., 2017; Shah et al., 2018).

MultiWOZ is a large-scale dialogue corpus that was developed to address more complex multi-domain dialogues (Budzianowski et al., 2018). It includes dialogues spanning seven domains, namely, tourist attractions, hotels, restaurants, taxis, trains, police stations, and hospitals, and led to the development of subsequent dialogue models. Following the introduction of MultiWOZ, numerous large-scale dialogue datasets have been constructed (Rastogi et al., 2020; Mosig et al., 2020; Chen et al., 2021), and task-oriented dialogue models using these as benchmarks have been actively researched. Cross-WOZ (Zhu et al., 2020) and other large-scale multi-domain dialogue datasets have been introduced for Chinese as well (Quan et al., 2020; Dai et al., 2022), promoting research on Chinese task-oriented dialogue systems.

However, due to the high cost of constructing task-oriented dialogue corpora, there are fewer

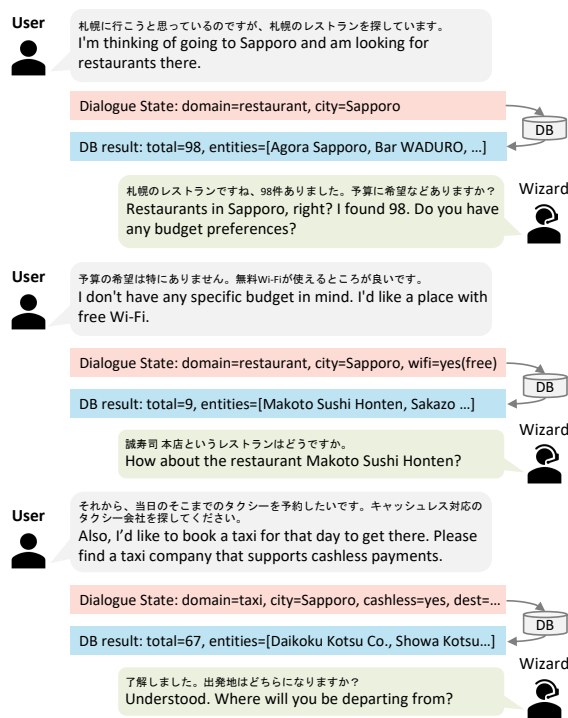


Figure 1: An example of dialogue across two domains: restaurants and taxis. The gray and green message bubbles represent the utterances of the user and the wizard, respectively. The red and blue boxes indicate the annotation of the dialogue state and the database results, respectively. The bubble for each utterance contains both the original Japanese utterance and its English translation by the authors.

* Equal contribution

multi-domain task-oriented dialogue corpora in other languages compared to English and Chinese (Hung et al., 2022). In this study, we focus on Japanese, in which the research and development of task-oriented dialogue models based on deep learning have been limited thus far.

Toward the research and development of task-oriented dialogue systems in Japanese, we have constructed the Japanese Multi-Domain Wizard of Oz (JMultiWOZ), the first Japanese multi-domain task-oriented dialogue dataset.¹ JMultiWOZ contains a total of 4,246 conversations spanning six travel-related domains (tourist attractions, accommodation, restaurants, shopping facilities, taxis, and weather). An example of a dialogue is shown in Figure 1. JMultiWOZ provides the dialogue state at each turn and the database of each domain for implementing and benchmarking task-oriented dialogue models.

In this paper, we outline the procedure for constructing the JMultiWOZ dataset and introduce its statistics. We evaluated the dataset on the two main tasks of task-oriented dialogue enabled by JMultiWOZ, i.e., dialogue state tracking (DST) and response generation (RG), using the state-of-the-art (SOTA) methods (Bang et al., 2023) and the latest LLM-based methods (Hudeček and Dusek, 2023). For further validation, the end-to-end dialogue capability of these dialogue models was evaluated by interactions with human participants. The contributions of this study are threefold:

- We constructed JMultiWOZ, the first large-scale Japanese multi-domain task-oriented dialogue dataset.
- We evaluated the dataset on DST and RG tasks using existing SOTA models and the latest LLMs, and demonstrated that JMultiWOZ can provide a Japanese benchmark of complexity comparable to that of the major English dataset MultiWOZ2.2.
- We conducted a human evaluation experiment, which showed that, even with the latest LLMs, there remain challenges concerning the capabilities of task-oriented dialogue in Japanese.

2. Related Work

2.1. Task-Oriented Dialogue Corpora

Many task-oriented dialogue corpora have been created in English thus far. Previously, single-domain dialogues, where only one domain appears in a dialogue, were predominant, such as WOZ2.0 (Wen

et al., 2017), Frames (El Asri et al., 2017), and KVRET (Eric et al., 2017). All of these dialogues were conducted using the Wizard of Oz (WOZ) method (Kelley, 1984), with a human user and another person acting as the system (i.e., the wizard). There are also human-to-machine (Henderson et al., 2014) and machine-to-machine dialogue corpora simulated between machines (Shah et al., 2018). Multi-domain dialogue corpora, in which multiple domains appear in a single dialogue, are being increasingly constructed to address more complex requirements. MultiWOZ (Budzianowski et al., 2018) is a representative example, being a large-scale corpus with over 10,000 dialogues covering seven travel-related domains. Other existing large-scale multi-domain dialogue corpora include Schema-Guided Dialogue (Rastogi et al., 2020), STAR (Mosig et al., 2020), and ABCD (Chen et al., 2021).

There are also several large-scale multi-domain dialogue corpora in Chinese. CrossWOZ (Zhu et al., 2020) is the first multi-domain dialogue corpus in Chinese and contains about 6,000 travel-related dialogues. The RiSAWOZ corpus (Quan et al., 2020) introduced more domains and dialogues, and subsequently, CGoDial (Dai et al., 2022) was devised as an extension of other dialogue corpora including RiSAWOZ. In addition, BiTOD (Lin et al., 2021b) was constructed to develop bilingual multi-domain dialogue models in both English and Chinese.

Constructing task-oriented dialogue corpora is generally costly, and there are few large-scale multi-domain datasets outside of English and Chinese (Hung et al., 2022). SCUD (Hayashibe, 2022) is a Japanese single-domain task-oriented dialogue corpus related to accommodation search but it only contains 210 dialogues.

2.2. Translation-based Corpora

Given the high cost of constructing task-oriented dialogue corpora, efforts are being made to construct dialogue corpora in other languages by translating readily available English corpora. For instance, GlobalWOZ (Ding et al., 2022) is a dialogue corpus that expanded MultiWOZ into 17 languages through machine translation. Among the 17 languages, high quality has been achieved in three languages (Chinese, Spanish, and Indonesian) through post-editing by professional translators for some dialogues in the test set. However, the quality is not guaranteed for other languages, including Japanese. Other corpora constructed from machine translation and manual post-editing of MultiWOZ include AllWOZ (Zuo et al., 2021) and Multi²WOZ (Hung et al., 2022), neither of which includes Japanese.

¹Our data and code are publicly available at <https://github.com/nu-dialogue/jmultiwoz>

Domain	Slots
restaurant	city, name, genre, area, pricerange, station, wifi, parking, opentime, phone, address, accesstime, closed, priceinfo, reference, people, day, time
accommodation	city, name, genre, area, pricerange, station, wifi, parking, withrestaurant, phone, address, accesstime, checkintime, checkouttime, priceinfo, reference, people, day, stay
attraction	city, name, genre, area, station, wifi, parking, opentime, phone, address, accesstime, closed, adultfee, childfee, priceinfo
shopping	city, name, genre, area, station, parking, opentime, phone, address, accesstime, closed
taxi	city, name, cashless, jumbo, phone, reference, day, time, departurepoint, arrivalpoint
weather	city, area, day, weather, mintemperature, maxtemperature

Table 1: List of slots defined in the ontology for each domain

Problems due to poor translations have been reported in translation-based dialogue corpora (e.g., ‘translationese,’ lack of cultural adaptation) (Majewska et al., 2023), which may prevent the models’ practical performance from being evaluated accurately (Hu et al., 2023). Therefore, this study aims to construct a realistic dataset in the Japanese context by collecting dialogues from scratch.

3. Data Collection

JMultiWOZ is a corpus containing dialogues of travelers planning a trip to one of nine cities in Japan (Sapporo, Sendai, Tokyo, Yokohama, Nagoya, Kyoto, Osaka, Fukuoka, and Naha) while collecting tourist information. The six domains include tourist attractions, accommodations, restaurants, shopping facilities, taxis, and weather. Using the WOZ method, each dialogue was conducted by two human interlocutors, one as a traveler (user) and the other as an information provider (wizard).

The following sections describe the five steps of constructing this corpus: (1) ontology definition, (2) construction of the backend database that the wizard uses to obtain travel information, (3) design and creation of user goals, (4) dialogue collection, and (5) annotation of the full dialogue state.

3.1. Definition of Ontology

In a task-oriented dialogue, the *ontology* represents the structure of the backend database. Specifically, it defines attributes such as name, address, and phone number for each *entity* in the database. An entity is a unit of record in the database, such as a specific tourist attraction or restaurant, and its attributes are called *slots*. In this study, the ontology of each domain was defined with reference to existing studies (Budzianowski et al., 2018; Zhu et al., 2020) while considering the characteristics of Japanese culture so as not to be unnatural. For instance, dialogues related to the police and hospital

domains that exist in MultiWOZ are rarely encountered in Japanese dialogue travel centers. In view of this, we eliminated police and hospital domains, and instead introduced more culturally appropriate domains such as shopping and weather. The ontology for all domains is shown in Table 1.

3.2. Construction of Backend Database

Based on the ontology, a backend database that the wizard uses to retrieve entities and travel information during the dialogue was constructed for each domain. To enhance the realism of the dialogue using real entities, lists of facilities publicly available from governments and municipalities of various cities were used to construct databases for tourist spots, accommodations, restaurants, shopping facilities, and taxis. From this list, only facilities that have publicly accessible websites were selected to be included in the database, and information for all slots of each entity was manually obtained from the website. For the taxi domain, the unit of entities was set as taxi companies.

The final number of entities contained in the database for each domain were as follows: 447 for tourist spots, 884 for accommodations, 952 for restaurants, 445 for shopping facilities, and 167 for taxis. For the construction of the weather domain database, the unit of entity was set as the date, and 365 days’ worth of weather information was artificially created for each city.

3.3. Design of User Goal

A user goal is the objective the user aims to achieve through the dialogue with the wizard, and one goal is set for each dialogue. An example of a user goal is shown in Table 2. Each goal covers one or more domains, and each domain in the goal consists of one or more *informable* slots, which are the search criteria for the entity the user is looking for, such as the desired budget or destination, and one or more *requestable* slots, which are the attribute in-




Figure 2: Web UI of wizard for dialogue collection: (A): A search form for an entity from the backend database, (B): An interface displaying detailed information about the selected entity and a reservation form for the entity, (C): An interface for chatting with the user.

Domain	Slot
attraction	info. city = Nagoya parking = yes(free) wifi = yes(free)
	reqt. area, station, closed
accommodation	info. pricerange = expensive wifi = yes(free)
	reqt. accesstime, genre

Table 2: Example of a user goal. “info.” and “reqt.” in the “Slot” column indicate informable and requestable slots, respectively.

formation of the entity the user needs to obtain, such as phone numbers or addresses. In domains where reservations are commonly made in real life (for instance, after finding the desired restaurant or accommodation, reservations are often made with conditions such as the date and number of people), several conditions for reservation are randomly added as *book* slots.

The composition and diversity of user goals are directly linked to the naturalness and diversity of the dialogues in the corpus. To introduce diversity in dialogue length and complexity, 1-3 domains are randomly selected for each user goal. Next, slots to be included in each domain are randomly chosen. From the ontology defined in Section 3.1, 2-7 slots, including informable, requestable, and possibly bookable slots, are selected for each domain. The informable slot “city”, which indicates the user’s tourist destination city, is shared among the domains within a user goal. To enhance the realism of the dialogue, following (Budzianowski

et al., 2018), some goals were set with tasks to change the value of an informable slot and/or book slot amid the dialogue (for instance, change the originally communicated reservation condition of 5 p.m. to 6 p.m.). Ahead of subsequent dialogue collection, a total of 5,000 unique user goals were created.

3.4. Dialogue Collection

Dialogues were collected using a backend database and randomly generated dialogue goals. Based on the dialogue collection platform² of (Zhu et al., 2020), the dialogue web UIs for the user and wizard were implemented. The web UI used by the wizard is depicted in Figure 2. Crowd workers for both the user and wizard roles were recruited via Lancers³, a major Japanese crowdsourcing service. Only those who consented to the publication of data obtained from the dialogue collection participated.

The workers read the instructions in the dialogue manual, watched a demonstration video that explained how to operate the web UI, and learned the workflow before participating in the dialogues. To ensure diverse user utterances, each user could participate in a maximum of 100 dialogues, and they could engage with the same wizard for a maximum of 20 dialogues. Meanwhile, because it is preferable for the wizard to behave consistently, there was no limit to the number of dialogues, and the same wizard was allowed to engage in conversations repeatedly. In the end, 65 users and 18 operators participated in the dialogue collection.

²<https://github.com/thu-coai/CrossWOZ>

³<https://www.lancers.jp>

1. You are planning a trip to **Nagoya**.
2. Find a **tourist attraction** to visit for one day. Choose a place where the parking is **free**. The place should also have **free** Wi-Fi.
3. Once you find a desired tourist spot, ask about its **area, nearest station, and closed days**.
4. Find a **place to stay** for the night. The budget is on the **higher** side. The place should also offer **free** Wi-Fi.
5. Once you find a suitable accommodation, ask about the **time it takes from the nearest station and the type of accommodation**.

Table 3: Explanation of the user goal shown in Table 2. The **blue text** indicates informable slot values, while the **green text** indicates requestable slot values.

The tasks of the user and wizard, and the quality control for the wizards are explained in detail in the following paragraphs.

User’s Task The user aims to properly convey the informable slots in the user goal to the wizard, obtain information about the requestable slots from the wizard, and in some cases, reserve the entity with the conditions of booking slots. Instead of a user goal formatted as in Table 2, users comprehend the goal by reading template sentences that describe each slot. Template sentences that explain the user goal in Table 2 are shown in Table 3.

Wizard’s Task The wizard searches for entities in the backend database that match the constraints conveyed by the user and provides the user with information about the found entities. The wizard’s web UI (Figure 2) has interfaces for (A) searching for entities based on the user’s criteria and (B) checking detailed information or making reservations for the found entities in addition to (C) the panel for chatting with the user. The database search query (DB query) input by the wizard in each turn is recorded as part of the dialogue state (details can be found in Section 3.5).

Quality Control of Wizard Some studies (Eric et al., 2020; Zang et al., 2020) have reported issues such as inconsistencies in the slot value notation within the dialogue state (e.g., the value “18:00” for the time slot appears in several ways: “1800”, “0600pm”, and “6 PM”) in existing corpora. Such inconsistencies can confuse or underestimate the capabilities of the dialogue model, making it impossible to provide an appropriate benchmark. Therefore, input values were selected from dropdown menus to prevent wizards from manually entering DB queries. Additionally, to enhance the quality of the wizard, 3-5 practice dialogues were conducted

	Total	Train	Dev	Test
Dialogues	4,246	3,654	300	300
Turns	61,186	52,405	4,346	4,435
Tokens	1,102,658	943,653	78,683	80,322
Vocab.	11,121	10,309	2,854	2,971

Table 4: Statistics of JMultiWOZ

Context	User:	I’m looking for a hotel in Sapporo, and it would be helpful if it has free wifi.
	Wizard:	How about JR INN Sapporo for example?
	User:	Good, what are the prices like?
DB query		city=Sapporo, wifi=yes
Dialogue State		city=Sapporo, wifi=yes, name=JR INN Sapporo

Table 5: Examples of a DB query obtained during dialogue collection and a dialogue state obtained from the additional annotation

beforehand. The workers received feedback from the authors on errors. Only workers who no longer had issues after repeated feedback participated in the actual dialogues.

Through the above procedure, a total of 4,508 dialogues were collected. Dialogues that would be considered noise in the dataset were then revised or excluded through the following two procedures:

1. At the end of each dialogue, workers completed questionnaires to report any issues that occurred during the dialogue. We manually checked dialogues with reported issues and removed those containing major errors, such as when workers misinterpreted the user goal.
2. We deleted dialogues in which the DB query at the end of the dialogue did not match the informable slots in the user goal.

The resulting corpus after this modification contained a total of 4,246 dialogues. Here, the dev and test sets consist of 300 dialogues each, randomly selected from the 4,246 dialogues, and the remaining 3,646 dialogues are used as the train set (refer to Table 4 for the statistics of each set).

3.5. Annotation of Full Dialogue State

The dialogue state is the information about the conditions of the entity that the user seeks, known up to each turn and recorded as a set of slot-value pairs. The DB query entered by the wizard at each turn can be used as a part of the dialogue state annotation. However, the DB query does not contain any non-explicitly communicated values. For

	# slot-value pairs
DB search query	155,274
Newly annotated dialogue state	58,745
Final dialogue state	214,019

Table 6: Number of slot-value pairs annotated in JMultiWOZ

	MultiWOZ	CrossWOZ	JMultiWOZ
Language	English	Chinese	Japanese
Domains	7	5	6
Slots	57	72	79
Dialogues	8,438	5,012	3,646
Turns	115,424	84,692	52,405
Avg. domains	1.8	3.24	2.02
Avg. turns	13.7	16.9	14.4
Synchronous	No	Yes	Yes

Table 7: Comparison of multi-domain task-oriented dialogue datasets on the train set; “Synchronous” indicates whether one user and one wizard actually conducted the dialogue synchronously during the dialogue collection.

example, as shown by Table 5, if the user accepts the entity name proposed by the system (e.g., wizard: “How about JR INN Sapporo?”, and user: “OK, what are the prices like?”), the entity name is not searched by the wizard; hence it is not recorded automatically.

To build the complete dialogue state, we recruited additional crowd workers to annotate this non-explicit value. After reading the manual, each worker annotated the values using a dedicated UI. As in the case of dialogue collection, the values were input by selection in order to suppress the perturbation of the slot value notations. To ensure the dialogue quality, each worker annotated ten dialogues for training and received feedback from the authors on errors. After repeated feedback, only those workers who no longer had issues participated in the annotation process.

In the end, six workers shared the annotation tasks for a total of 30,593 wizard turns.⁴ Table 6 shows the statistics of the number of slots in the dialogue states. A total of 58,745 slots (about 37.8% of the slots recorded in the DB query) were added to the dialogue states. Here, 59 randomly selected dialogues were annotated by the authors, and the match rate between these annotations and those annotated by the workers was 94.1%, indicating that the annotations were sufficiently high quality.

⁴Of the total 61,186 turns in the dataset (see Table 4), 30,593 were wizard turns, and the dialogue states were annotated for them.




Figure 3: The distribution of dialogue lengths, divided into dialogues containing only one domain (single-domain) and dialogues containing two or more domains (multi-domain).

3.6. Statistics


Table 7 shows the statistics of JMultiWOZ compared with the major multi-domain task-oriented dialogue datasets in English and Chinese, MultiWOZ and CrossWOZ, respectively. Given that the number of domains, the number of slots, the average number of domains, and the average number of turns are roughly equivalent, the complexity of dialogues in JMultiWOZ can be considered to be on par with the existing datasets. Figure 3 illustrates the distribution of the number of dialogue turns when all 4,246 dialogues are divided into dialogues containing only one domain (single-domain) and dialogues containing multiple domains (multi-domain). A wide variety of length and complexity can be seen in both types of dialogues.

4. Benchmark


JMultiWOZ provides benchmarks for two common tasks in task-oriented dialogues, dialogue state tracking (DST) and response generation (RG). DST is the task of estimating the dialogue state at each turn, while RG is the task of generating the next system response based on the dialogue context at each turn. To demonstrate that JMultiWOZ can provide benchmarks equivalent to those in existing English dialogue corpora, we evaluate the aforementioned two tasks in JMultiWOZ using the SOTA methods from MultiWOZ2.2 (Zang et al., 2020) and the latest LLM-based methods. Note that MultiWOZ2.2 is a version in which various annotation errors in the original MultiWOZ have been corrected.

4.1. Baseline Models

The most recent task-oriented dialogue models can be generally categorized into two approaches: (1) fine-tuning with supervised learning on medium-sized pretrained language models (Supervised Fine-Tuning; SFT) (Lin et al., 2021a; Su et al., 2022), and (2) generating responses in a zero-/few-



(a) T5 Pipeline (Bang et al., 2023). (1) First, the dialogue state is predicted from the given dialogue context, and (2) the result is added to the input to generate the final response. Both (1) and (2) are performed on the same model.



(b) LLM pipeline (Hudeček and Dusek, 2023) in zero-shot setting. First, (1) the current active domain is estimated from the dialogue context. Next, (2) the dialogue state is tracked, and (3) the response is generated for that domain using a prompt focused on that domain.

Figure 4: Task-oriented dialogue model pipeline for DST and RG

shot setting using an LLM (Hudeček and Dusek, 2023; Raposo et al., 2023). In this experiment, we validate both methods using the following models. Note that the model-predicted dialogue states were used as input for RG, instead of the ground-truth dialogue states. This is in line with the settings used in previous studies (Bang et al., 2023; Hudeček and Dusek, 2023), allowing for comparison with the scores reported in those studies.

Supervised fine-tuning For SFT, we use the T5 (Raffel et al., 2020)-based model TOATOD (Bang et al., 2023), which is the SOTA model in MultiWOZ2.2. TOATOD uses T5-base pretrained on a large number of task-oriented dialogue corpora as its backbone model, but such corpora do not exist in Japanese. Therefore, for the evaluation of JMultiWOZ, we simply use T5-base/large⁵ pretrained on Japanese data as backbone models.

Figure 4a depicts the overview of response generation using T5. First, the dialogue state is estimated from the dialogue context (i.e., DST). Then, the database search results, obtained by using the dialogue state as a DB query, are combined with the context and dialogue state to generate the final response (i.e., RG). During fine-tuning, the mapping of ground truth inputs and outputs at each step is learned using maximum likelihood estimation (MLE). Because the same model is used to perform both DST and RG, a prefix indicating the task is added to the input. Specific examples of the input-output sequences used in the training of T5 can be found in Section A in the appendix.

⁵<https://huggingface.co/retrieva-jp/t5-large-long>

In the training of T5-base and T5-large, the batch size was set to 32, and they were trained for five epochs. The AdamW (Loshchilov and Hutter, 2019) optimizer was used, and the learning rate was initially set to 5e-5 and then linearly decayed according to the number of steps. For evaluation, the checkpoint at the final step was used, and for inference in DST and RG, greedy search was adopted in both cases.

LLM zero-/few-shot We use the zero-/few-shot response generation pipeline using an LLM introduced by Hudeček and Dusek (2023). Figure 4b shows the 3-step flow of zero-shot response generation. First, the LLM estimates the current active domain from the context. Then, focusing on that domain, it estimates the dialogue state. Based on the state, it searches the DB and uses all of the results to generate the final response. In the few-shot setting, two examples retrieved from the train set are mixed into each prompt in the pipeline. These examples are retrieved from the train set based on the similarity between the dialogue context embeddings of the retrieved examples and the current context embeddings. The Japanese sentence-transformers⁶ were used to create embedding vectors using two consecutive turns of utterances as the dialogue context.

We used APIs of OpenAI GPT-3.5 (gpt-3.5-turbo) and GPT-4 (gpt-4) for the LLM-based zero-/few-shot settings. GPT-3.5 is the highest performing model for DST and RG on MultiWOZ, as reported by Hudeček and Dusek (2023). GPT-

⁶<https://huggingface.co/cl-nagoya/sup-simcse-ja-large>

Model	Few-shot	MultiWOZ 2.2			JMultiWOZ		
		JGA	Slot-F1	BLEU	JGA	Slot-F1	BLEU
T5-base [†]	✗	0.64	0.94	17.04	0.59	0.95	42.31
T5-large	✗	–	–	–	0.77	0.98	49.68
GPT-3.5 [‡]	✗	0.13	0.40	4.17	0.16	0.70	5.31
GPT-4	✗	–	–	–	0.31	0.86	8.87
GPT-3.5 [‡]	✓	0.27	0.51	6.77	0.25	0.82	12.91
GPT-4	✓	–	–	–	0.36	0.89	15.76

Table 8: Automatic evaluation results of the baseline model. Dagger [†] indicates that the score of T5-base in MultiWOZ2.2 is cited from the evaluation results of TOATOD reported in (Bang et al., 2023). Double-dagger [‡] indicates that the score of GPT-3.5 in MultiWOZ2.2 is cited from the reported value in (Hudeček and Dusek, 2023).

4 is the latest model provided by OpenAI’s API and is known for its high capabilities in Japanese.⁷ The Japanese prompts were created based on the prompts in (Hudeček and Dusek, 2023). See Section B in the Appendix for examples of the prompts.

4.2. Evaluation Metrics

To evaluate DST, we used the most standard metric, joint goal accuracy (JGA). JGA is a measure of whether the estimated dialogue state perfectly matches the ground truth dialogue state at each turn, resulting in a binary outcome (0/1). Additionally, we also utilized Slot-F1 to evaluate for each turn the match rate between the set of slot values in the estimated dialogue state and the set of slot values in the ground truth state, in terms of F1 score. To evaluate RG, we used the BLEU score, which indicates the similarity between the generated response and the ground truth response.

4.3. Results

Table 8 shows a comparison between the evaluation results reported in previous studies on MultiWOZ2.2 and those of JMultiWOZ. On JMultiWOZ, the SFT method, namely T5-base/large, generally had the highest performance for both DST and RG, and there seemed to be certain limitations to the performance of LLMs. This trend is consistent with the results of MultiWOZ2.2.

For the DST metrics, Slot-F1 was higher in JMultiWOZ. As described in Section 3.4, in the construction of JMultiWOZ, we chose the selection-based input for the DB search query by the wizard to reduce variations in dialogue state reported as issues in MultiWOZ. This enabled the dialogue models to predict exact values without confusion, resulting in higher slot-F1.

⁷Based on the results of the Japanese LLM leaderboard (<https://api.wandb.ai/links/wandb-japan/6ff86bp3> as of October 2023).

For JGA, the primary metric for DST, the difference between MultiWOZ2.2 and JMultiWOZ was minimal. This suggests that the complexity and annotation accuracy of JMultiWOZ are comparable to that of MultiWOZ2.2, demonstrating its capacity to provide benchmarks equivalent to the existing datasets. Notably, there is a difference of about 5% in JGA for T5-base across both corpora. This gap can be attributed to the fact that in MultiWOZ2.2, besides supervised learning, T5-base was boosted via reinforcement learning with the reward functions where JGA was incorporated.

For the RG metric, BLEU, JMultiWOZ yielded significantly higher scores compared to MultiWOZ2.2. This is likely because the wizard’s utterances in JMultiWOZ are more consistent than in MultiWOZ2.2. During the dialogue collection for JMultiWOZ, the dialogues were conducted synchronously, ensuring that the wizards did not switch in the middle of each dialogue. Moreover, we provided ample training for the wizards. Through such quality control measures, we believe that consistent system responses could be attained.

5. Human Evaluation

The model’s actual task completion capability in real dialogues must be assessed by actual people when evaluating task-oriented dialogue models as well as by automatic evaluation (Iizuka et al., 2023). In this section, we evaluate the end-to-end dialogue capabilities of the four dialogue models acquired in Section 4, namely, T5-base, T5-large, GPT-3.5, and GPT-4, by having them engage in conversations with crowd workers. Note that both GPT-3.5 and GPT-4 used a few-shot setting.

5.1. Settings

The setting of this evaluation experiment followed that of Iizuka et al. (2023), where they evaluated the end-to-end dialogue performance of the TOA-

Model	MultiWOZ 2.2						JMultiWOZ					
	N	Success	Turn	Und.	App.	Sat.	N	Success	Turn	Und.	App.	Sat.
T5-base [†]	36	66.67	12.56	3.83	3.81	3.72	38	65.79	10.74	3.92	3.71	3.55
T5-large	–	–	–	–	–	–	40	75.00	10.10	4.05	3.98	3.82
GPT-3.5 [†]	42	57.14	11.55	3.79	3.98	4.05	41	24.39	11.05	2.90	2.24	1.95
GPT-4 [†]	42	76.19	11.88	4.26	4.36	4.00	42	57.14	9.55	3.93	3.52	3.02

Table 9: Human evaluation results. “N” indicates the number of participants who interacted with each model. “Und.,” “App.,” and “Sat.” indicate the worker’s subjective evaluation of the system’s understanding, appropriateness of the system response, and satisfaction with the dialogue, respectively. Dagger [†] indicates that the score in MultiWOZ2.2 is cited from the reported value in (Iizuka et al., 2023).

TOD (Bang et al., 2023) and Hudeček and Dusek (2023)’s LLM pipeline models built based on the MultiWOZ2.2 by using crowdsourcing.

Specifically, each worker was first given dialogue instructions and a user goal for each dialogue session. These user goals were randomly sampled from the test set. Then, they engaged in a dialogue with one of the four models. Each dialogue was set to a maximum of 20 turns (one turn consists of one user utterance and one system response), and workers judged whether they achieved the user goal within 20 turns. After the dialogue, the workers evaluated (1) the system’s language understanding ability, (2) the appropriateness of the system’s responses, and (3) their overall satisfaction with the dialogue, each on a 5-point scale. The workers were restricted to converse with only one system.

5.2. Results

Table 9 presents the evaluation results on JMultiWOZ, as well as the results reported by Iizuka et al. (2023) on MultiWOZ2.2. For the models trained with SFT, namely T5-base/large, there was no significant difference in Success, compared to T5-base (TOATOD) on MultiWOZ2.2. These results show that using JMultiWOZ allows for the development and evaluation of Japanese end-to-end dialogue models with performance comparable to that of MultiWOZ2.2.

The performance of the LLM models, i.e., GPT-3.5 and GPT-4, significantly declined compared to that of MultiWOZ2.2. This may be because even the latest LLMs, such as GPT-4, are not capable of handling dynamically changing dialogue contexts in Japanese. Specifically, since the model’s predicted dialogue states and system responses are accumulated in the dialogue history, errors propagate gradually, making it difficult to maintain multi-turn conversations. GPT-4’s Japanese language ability is not as high as that for English (OpenAI, 2023), and this difference in ability is likely reflected in the performance of task-oriented dialogue systems. This limitation of dialogue skills in LLMs in non-English languages should be addressed with multilingual resources in the future, and we anticipate

that JMultiWOZ can contribute to such studies.

6. Conclusion

We have presented JMultiWOZ, the first large-scale Japanese multi-domain task-oriented dialogue dataset. This corpus contains 4,246 WOZ dialogues spanning six travel-related domains and provides benchmarks for dialogue state tracking and response generation. Using JMultiWOZ, we evaluated existing SOTA methods with the T5-based models and the latest LLM-based methods, namely GPT-3.5/4, and demonstrated that JMultiWOZ performs comparably to the major benchmark dataset in English, MultiWOZ. Furthermore, we confirmed that the capabilities of GPT-3.5/4 in particular are limited in Japanese.

For future work, we would like to conduct comprehensive evaluation using more diverse models, including other LLMs with high multilingual abilities (Team et al., 2023). We hope that the development of JMultiWOZ will lead to further research on Japanese dialogue systems, including the improvement of LLMs’ task-oriented dialogue capabilities in Japanese and the development of a multilingual task-oriented dialogue model.

7. Limitations

In this study, to collect native dialogues specific to Japan, we collected online conversations through Japanese crowd workers. However, in prior research, multilingual corpora were built through machine translations of MultiWOZ, and GlobalWOZ (Ding et al., 2022) includes Japanese dialogue data, albeit being entirely machine-translated. Therefore, in the future, it will be necessary to investigate the advantages of JMultiWOZ, which we collected from scratch, compared to GlobalWOZ. Additionally, it may be possible to further improve performance by using GlobalWOZ for pre-training and then fine-tuning with our JMultiWOZ; this will need to be validated in the future.

JMultiWOZ provides benchmarks for two main tasks, DST and RG, through dialogue state anno-

tations, enabling the construction and evaluation of end-to-end task-oriented dialogue models. However, it does not yet have dialogue act (DA) annotations (Budzianowski et al., 2018; Eric et al., 2020). This means that it does not provide a benchmark for another major task in task-oriented dialogue, policy optimization. Therefore, to further enhance its utility, DA annotations should be added in the future.

8. Ethical Considerations

Our ethical considerations span across data sources, dialogue collection, human evaluation, and the implications of using pretrained language models (PLMs).

Backend Database We exclusively employed information sources free from copyright constraints. Our list of entities was primarily derived from websites operated by the government or municipalities. Specific information for each entity was solely extracted from their respective official websites, ensuring authenticity and credibility. We consciously abstained from using tourism sites, or any other source potentially encumbered by copyright issues.

Dialogue Collection and Human Evaluation

Prior to our data collection and evaluation experiments, ethical approval was obtained from the affiliated organization. In the data collection and evaluations, we engaged only crowd workers who explicitly consented to abstain from (1) unsolicited disclosure of personal information during dialogues. Additionally, these workers agreed to (2) relinquish copyright claims over data and artifacts produced during the dialogue collection phase and (3) publication of collected dialogue data. The data to be released will not contain any personally identifiable information of the workers, such as their names.

Dialogue Modeling with PLMs The pre-training data of the PLMs uses a vast amount of textual data from Internet information. Therefore, our dialogue models based on these PLMs may produce potentially harmful or discriminatory responses.

9. Acknowledgments

This work was supported by JST Moonshot R&D Grant number JPMJMS2011. We used the computational resources of the supercomputer “Flow” at the Information Technology Center, Nagoya University.

10. Bibliographical References

- Namo Bang, Jeehyun Lee, and Myoung-Wan Koo. 2023. [Task-Optimized Adapters for an End-to-End Task-Oriented Dialogue System](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 7355–7369.
- 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.
- Derek Chen, Howard Chen, Yi Yang, Alexander Lin, and Zhou Yu. 2021. [Action-Based Conversations Dataset: A Corpus for Building More In-Depth Task-Oriented Dialogue Systems](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3002–3017.
- Yinpei Dai, Wanwei He, Bowen Li, Yuchuan Wu, Zheng Cao, Zhongqi An, Jian Sun, and Yongbin Li. 2022. [CGoDial: A large-scale benchmark for Chinese goal-oriented dialog evaluation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4097–4111.
- Bosheng Ding, Junjie Hu, Lidong Bing, Mahani Aljunied, Shafiq Joty, Luo Si, and Chunyan Miao. 2022. [GlobalWoZ: Globalizing MultiWoZ to develop multilingual task-oriented dialogue systems](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, pages 1639–1657.
- 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 Meeting of the Special Interest Group on Discourse and Dialogue*, pages 207–219.
- 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 baselines](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 422–428.
- 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 Meeting of the Special Interest Group on Discourse and Dialogue*, pages 37–49.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2018. Neural Approaches to Conversational AI. In *Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 1371–1374.
- Yuta Hayashibe. 2022. Self-Contained Utterance Description Corpus for Japanese Dialog. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 1249–1255.
- Wanwei He, Yinpei Dai, Yinhe Zheng, Yuchuan Wu, Zheng Cao, Dermot Liu, Peng Jiang, Min Yang, Fei Huang, Luo Si, Jian Sun, and Yongbin Li. 2022. GALAXY: A Generative Pre-trained Model for Task-Oriented Dialog with Semi-supervised Learning and Explicit Policy Injection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 10, pages 10749–10757.
- Matthew Henderson, Blaise Thomson, and Jason D. Williams. 2014. [The Second Dialog State Tracking Challenge](#). In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 263–272.
- Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. 2020. [A simple language model for task-oriented dialogue](#). In *Proceedings of Advances in Neural Information Processing Systems*, pages 20179–20191.
- Songbo Hu, Han Zhou, Mete Hergul, Milan Gritta, Guchun Zhang, Ignacio Iacobacci, Ivan Vulić, and Anna Korhonen. 2023. [Multi 3 WOZ: A multilingual, multi-domain, multi-parallel dataset for training and evaluating culturally adapted task-oriented dialog systems](#). *Transactions of the Association for Computational Linguistics*, pages 1396–1415.
- Vojtěch Hudeček and Ondrej Dusek. 2023. [Are Large Language Models All You Need for Task-Oriented Dialogue?](#) In *Proceedings of the 24th Meeting of the Special Interest Group on Discourse and Dialogue*, pages 216–228.
- Chia-Chien Hung, Anne Lauscher, Ivan Vulić, Simone Ponzetto, and Goran Glavaš. 2022. [Multi2WOZ: A robust multilingual dataset and conversational pretraining for task-oriented dialog](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3687–3703.
- Shinya Iizuka, Shota Mochizuka, Atsumoto Ohashi, Sanae Yamashita, Ao Guo, and Ryuichiro Higashinaka. 2023. [Clarifying the Dialogue-Level Performance of GPT-3.5 and GPT-4 in Task-Oriented and Non-Task-Oriented Dialogue Systems](#). In *Proceedings of AAAI Fall Symposium on Artificial Intelligence for Human-Robot Interaction*, pages 182–186.
- J. F. Kelley. 1984. An iterative design methodology for user-friendly natural language office information applications. *ACM Transactions on Information Systems*, 2:26–41.
- Zhaojiang Lin, Bing Liu, Seungwhan Moon, Paul Crook, Zhenpeng Zhou, Zhiguang Wang, Zhou Yu, Andrea Madotto, Eunjoon Cho, and Rajen Subba. 2021a. [Leveraging Slot Descriptions for Zero-Shot Cross-Domain Dialogue State Tracking](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5640–5648.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, Peng Xu, Feijun Jiang, Yuxiang Hu, Chen Shi, and Pascale Fung. 2021b. [BiTOD: A bilingual multi-domain dataset for task-oriented dialogue modeling](#). *arXiv preprint arXiv:2106.02787*.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Olga Majewska, Evgeniia Razumovskaia, Edoardo M. Ponti, Ivan Vulić, and Anna Korhonen. 2023. [Cross-Lingual Dialogue Dataset Creation via Outline-Based Generation](#). *Transactions of the Association for Computational Linguistics*, pages 139–156.
- Johannes EM Mosig, Shikib Mehri, and Thomas Kober. 2020. [STAR: A Schema-Guided Dialog Dataset for Transfer Learning](#). *arXiv preprint arXiv:2010.11853*.
- OpenAI. 2023. [GPT-4 Technical Report](#). *ArXiv*, abs/2303.08774.
- Jun Quan, Shian Zhang, Qian Cao, Zizhong Li, and Deyi Xiong. 2020. [RiSAWOZ: A large-scale multi-domain Wizard-of-Oz dataset with rich semantic annotations for task-oriented dialogue modeling](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 930–940.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena,

- Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, pages 1–67.
- Gonçalo Raposo, Luisa Coheur, and Bruno Martins. 2023. [Prompting, retrieval, training: An exploration of different approaches for task-oriented dialogue generation](#). In *Proceedings of the 24th Meeting of the Special Interest Group on Discourse and Dialogue*, pages 400–412.
- 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.
- 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*.
- Yixuan Su, Lei Shu, Elman Mansimov, Arshit Gupta, Deng Cai, Yi-An Lai, and Yi Zhang. 2022. [Multi-task pre-training for plug-and-play task-oriented dialogue system](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, pages 4661–4676.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. [Gemini: a family of highly capable multimodal models](#). *arXiv preprint arXiv:2312.11805*.
- 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*, pages 438–449.
- Xiaoxue Zang, Abhinav Rastogi, Srinivas Sunkara, Raghav Gupta, Jianguo Zhang, and Jindong Chen. 2020. [MultiWOZ 2.2 : A dialogue dataset with additional annotation corrections and state tracking baselines](#). In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 109–117.
- Yichi Zhang, Zhijian Ou, and Zhou Yu. 2020a. [Task-Oriented Dialog Systems That Consider Multiple Appropriate Responses under the Same Context](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9604–9611.
- Zheng Zhang, Ryuichi Takanobu, Qi Zhu, Minlie Huang, and XiaoYan Zhu. 2020b. [Recent advances and challenges in task-oriented dialog systems](#). *Science China Technological Sciences*, pages 1–17.
- 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*, 8:281–295.
- Lei Zuo, Kun Qian, Bowen Yang, and Zhou Yu. 2021. [AllWOZ: Towards Multilingual Task-Oriented Dialog Systems for All](#). *arXiv preprint arXiv:2112.08333*.

A. Example of Input-Output Data for T5

Figures 5 and 6 show specific examples of the input-output sequences used in the training of T5.

Input text 対話から信念状態を推定: <顧客> 東京へ旅行に行くのですが、Wi-Fiが無料で利用できて、レストラン併設の宿泊施設はありますか？ <店員> WiFi無料でレストラン併設の宿泊施設は東京には18件あります <顧客> 安めのところもありますか？ <店員> 値段が安めの条件を追加すると、残念ながら1件もヒットしません。 <顧客> そうですか。じゃあWi-Fiはなくてもいいです。
Output text general active_domain hotel, general city 東京, hotel pricerange 安め, hotel withrestaurant 有り

(a) Original data in Japanese

Input text Dialogue state from the context: <User> I'm planning to travel to Tokyo. Are there any accommodations that offer free Wi-Fi and have an on-site restaurant? <Wizard> There are 18 accommodations in Tokyo that offer free Wi-Fi and have an on-site restaurant. <User> Are there any options in the cheap price range? <Wizard> Unfortunately, when adding the condition of lower prices, there are no hits. <User> I see. In that case, I don't need Wi-Fi.
Output text general active_domain hotel, general city Tokyo, hotel pricerange cheap, hotel withrestaurant yes

(b) English translation by the authors

Figure 5: An example of input/output data used for training the T5 on the DST task. At the beginning of the input sequence, a prefix indicating the DST task is attached. Additionally, each element of the input sequence is prefixed with indicators denoting the speaker.

Input text 対話から応答を生成: <顧客> 東京へ旅行に行くのですが、Wi-Fiが無料で利用できて、レストラン併設の宿泊施設はありますか？ <店員> WiFi無料でレストラン併設の宿泊施設は東京には18件あります <顧客> 安めのところもありますか？ <店員> 値段が安めの条件を追加すると、残念ながら1件もヒットしません。 <顧客> そうですか。じゃあWi-Fiはなくてもいいです。 <信念状態> general active_domain hotel, general city 東京, hotel pricerange 安め, hotel withrestaurant 有り <検索結果> total 1, candidate [ビジネスホテル山百合], selected [city 東京, name ビジネスホテル山百合, genre ビジネスホテル, area 台東区, pricerange 安め, station [上野駅, 京成上野駅], wifi 無し, parking 無し, withrestaurant 有り, phone 0338317759, address 東京都...
Output text 値段が安めでレストラン併設となれば、ビジネスホテル山百合、これ一択です。

(a) Original data in Japanese

Input text Response from the context: <User> I'm planning to travel to Tokyo. Are there any accommodations that offer free Wi-Fi and have an on-site restaurant? <Wizard> There are 18 accommodations in Tokyo that offer free Wi-Fi and have an on-site restaurant. <User> Are there any options in the cheap price range? <Wizard> Unfortunately, when adding the condition of lower prices, there are no hits. <User> I see. In that case, I don't need Wi-Fi. <State> general active_domain hotel, general city Tokyo, hotel pricerange cheap, hotel with_restaurant yes <DB Results> total 1, candidate [Business Hotel Yamayuri], selected [city Tokyo, name Business Hotel Yamayuri, genre business hotel, area Taito-ku, pricerange affordable, station [Ueno Station, Keisei Ueno Station], wifi no, parking no, withrestaurant yes, phone 0338317759, address Tokyo...
Output text If you're looking for an affordable hotel with an on-site restaurant, Business Hotel Yamayuri is the best choice.

(b) English translation by the authors

Figure 6: An example of input/output data used for training the T5 on the RG task. At the beginning of the input sequence, a prefix indicating the RG task is attached. Additionally, each element of the input sequence is prefixed with indicators denoting the speaker, belief states, and the results of database searches.

B. Prompt Examples for LLMs

We created the Japanese prompts for JMultiWOZ based on the prompts for MultiWOZ used by Hudeček and Dusek (2023). Figures 7 and 8 show examples of prompts used for LLMs in a zero-shot setting.

対話文脈から判断できる、レストランに関するスロットと値ペア（信念状態）を抽出してください。信念状態は半角スペースとカンマを使用し、`スロット1 値1, スロット2 値2` という形式で抽出してください。抽出すべきスロットは以下の通りです：

- "name" レストランの名前
- "genre" レストランのジャンル
- "area" レストランのエリア。"~区" や "~市" などの地区名。
- "pricerange" レストランの価格帯。"安め/普通/高め" のいずれか。
- "station" レストランの最寄り駅
- "wifi" レストランのWi-Fiの有無。"有り(無料)/有り(有料)/無し" のいずれか。
- "parking" レストランの駐車場の有無。"有り(無料)/有り(有料)/無し" のいずれか。
- "people" レストランの予約人数
- "day" レストランの予約日
- "time" レストランの予約時間

上記以外の情報は抽出しないでください。
また、文脈中で言及されなかったスロットの値も抽出しないでください。

(a) Original prompt in Japanese

Extract the slot-value pairs (belief state) about the restaurant that can be inferred from the conversation. Use half-width spaces and commas to extract the belief state in the format "slot1 value1, slot2 value2." The slots to be extracted are as follows:

- "name" Name of the restaurant
- "genre" Genre of the restaurant
- "area" Area of the restaurant, such as district names like "-ward" or "-city."
- "pricerange" Price range of the restaurant, either "affordable," "moderate," or "expensive."
- "station" Nearest station to the restaurant
- "wifi" Availability of Wi-Fi at the restaurant, either "free," "paid," or "not available."
- "parking" Availability of parking at the restaurant, either "free," "paid," or "not available."
- "people" Number of people for the restaurant reservation
- "day" Date of the restaurant reservation
- "time" Time of the restaurant reservation

Do not extract any information other than the above.
Also, do not extract values for slots that were not mentioned in the context.

(b) English translation by the authors

Figure 7: Prompt examples used for the DST task in the restaurant domain for the LLMs

あなたは顧客の要望に沿ったレストランを探し出し、予約をするアシスタントです。データベースを用い、エリア、ジャンル、価格帯等からレストランを検索・予約することができます。レストランを見つけたら、その名前、住所、電話番号、その他必要な情報など、顧客から尋ねられた情報を提供してください。
予約が成功したら、その予約番号 (ref) を提供してください

(a) Original prompt in Japanese

You are an assistant that finds and books restaurants based on the customer's preferences. Using a database, you can search for, and book restaurants based on area, genre, price range, etc. Once you find a restaurant, please provide the name, address, phone number, and any other necessary information that the customer has asked for.
If the reservation is successful, please provide the reservation number (ref).

(b) English translation by the authors

Figure 8: Prompt examples used for the RG task in the restaurant domain for the LLMs