

# InstructTODS: Large Language Models for End-to-End Task-Oriented Dialogue Systems

Willy Chung, Samuel Cahyawijaya, Bryan Wilie, Holy Lovenia, Pascale Fung

Hong Kong University of Science and Technology

Clear Water Bay, Hong Kong

whcchung@connect.ust.hk

## Abstract

Large language models (LLMs) have been used for diverse tasks in natural language processing (NLP), yet remain under-explored for task-oriented dialogue systems (TODS), especially for end-to-end TODS. We present InstructTODS, a novel off-the-shelf framework for zero-shot end-to-end task-oriented dialogue systems that can adapt to diverse domains without fine-tuning. By leveraging LLMs, InstructTODS generates a proxy belief state that seamlessly translates user intentions into dynamic queries for efficient interaction with any KB. Our extensive experiments demonstrate that InstructTODS achieves comparable performance to fully fine-tuned TODS in guiding dialogues to successful completion without prior knowledge or task-specific data. Furthermore, a rigorous human evaluation of end-to-end TODS shows that InstructTODS produces dialogue responses that notably outperform both the gold responses and the state-of-the-art TODS in terms of helpfulness, informativeness, and humanness. Moreover, the effectiveness of LLMs in TODS is further supported by our comprehensive evaluations on TODS subtasks: dialogue state tracking, intent classification, and response generation. Code and implementations could be found here<sup>1</sup>.

## 1 Introduction

LLMs have consistently pushed new frontiers in natural language processing (NLP) in terms of performance across a variety of benchmarks, such as MMLU (Hendrycks et al., 2020), BIG-Bench (Lewkowycz et al., 2022) and HELM (Bommasani et al., 2022), achieving state-of-the-art results in both natural language understanding (NLU) and generation (NLG) tasks (Bang et al., 2023). Various applications of LLMs have also been adopted in the industry, most prominently ChatGPT<sup>2</sup> and

<sup>1</sup><https://github.com/WillyHC22/InstructTODS/>

<sup>2</sup><http://chatgpt.openai.com/>

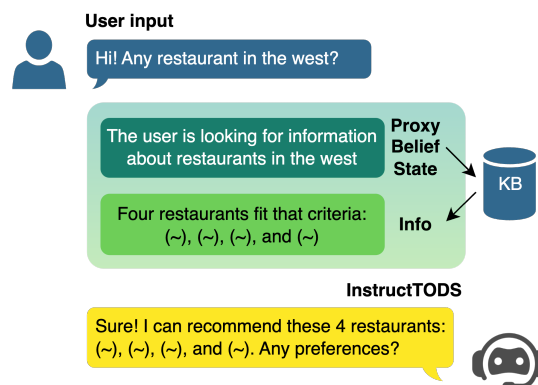


Figure 1: InstructTODS is the first zero-shot end-to-end task-oriented dialogue system that requires no task-specific annotations, and ontology while generating more human-preferred responses.

GPT-4<sup>3</sup>, which can provide a natural answer to a diverse range of questions fluently and coherently.

Among the manifold tasks in NLP, task-oriented dialogue systems (TODS) represent a crucial domain. In general, TODS can be categorized into: the pipelined approach (Ham et al., 2020; Hosseini-Asl et al., 2020a; Ohashi and Higashinaka, 2022), relying on multiple sequential modules and heavy annotations for dialogue states and system actions, and the end-to-end approach (Banerjee and Khapra, 2019; Qin et al., 2020; He et al., 2022), where the systems generate responses directly from the user input and the KB. Both approaches lack adaptability to unseen domains. This adaptability often requires domain-specific structures (ontology), and data for TODS is notoriously expensive to collect and annotate (Eric et al., 2020). In this regard, LLMs present great potential thanks to their extensive pre-trained knowledge, enabling them to adapt to contextual information without any parameter updates or additional task-specific data.

However, utilizing LLMs for tasks requiring knowledge grounding, such as TODS, poses a criti-

<sup>3</sup><https://openai.com/gpt-4>

Setting	Model	Banking77				CLINC150				OOS*
		Single	top-1	Multi		Single	top-1	Multi		
				top-2	top-3			top-2	top-3	
<b>Few-Shot Baseline</b>	RoBERTa <sub>large</sub>	78.99	–	–	–	89.89	–	–	–	–
	ICDA <sub>XL</sub>	83.90	–	–	–	92.62	–	–	–	–
	DNNC	80.40	–	–	–	91.02	–	–	–	–
	CPFT	80.86	–	–	–	92.34	–	–	–	–
<b>Zero-Shot Baseline</b>	BART <sub>large</sub> -MNLI	35.91	35.91	49.09	56.14	26.44	26.44	36.46	43.13	0.2
	BART <sub>large</sub> -MNLI <sub>corr</sub>	42.24	42.24	55.19	62.2	40.16	40.16	51.88	57.66	0.9
<b>Zero-Shot LLM</b>	Modular (GPT-3.5)	65.45	<b>64.51</b>	77.69	83.18	64.91	63.22	72.20	82.42	10.9
	Modular (GPT-4)	<b>74.09</b>	64.06	<b>80.75</b>	<b>86.33</b>	<b>73.91</b>	<b>69.90</b>	<b>81.88</b>	<b>90.33</b>	<b>62.1</b>

Table 1: Performance comparison on **intent classification**. LLMs outperform most baselines in our benchmark. Best performances in each section are in **bold**. \*Out-of-scope intent of CLINC150.

cal challenge that calls for thorough investigation and exploration. TODS requires dialogue systems to adeptly complete a specific goal by interacting with a user in natural language according to a certain set of bounded functions, ontology, and knowledge within the corresponding domain. Nevertheless, naively feeding all of the knowledge to the LLMs in TODS could lead to the generation of misleading and unfaithful information, i.e., hallucination (Ji et al., 2023; Azamfirei et al., 2023).

In this work, we first investigate the capability of LLMs to perform three key TODS objectives in zero-shot settings, specifically dialogue state tracking (DST), intent classification (IC), and response generation (RG). While LLMs demonstrate impressive capabilities and understanding of these tasks individually, a closer examination of their shortcomings reveals that the modular approach is not the most suitable for effectively using LLMs in TODS due to its restrictiveness. Rather than confining interactions within predefined elements like slots, values, or system actions, it is more advantageous to harness the emergent abilities of LLMs to process unstructured information, which also enables the system to easily adapt to new domains.

From these observations, we propose InstructTODS, a fully off-the-shelf framework to perform end-to-end unified TODS in a zero-shot setting using LLMs. InstructTODS is adaptable to any KB and does not require any ontologies or task-specific data. Instead of using predefined slot values, InstructTODS generates an unstructured proxy belief state from the dialogue context. Then, an action thought is generated to query the KB dynamically in natural language using an LLM. The retrieved information is then given to generate the response.

In summary, our contributions are as follows:

- We provide an extensive evaluation and com-

prehensive analysis of LLMs’ zero-shot performance in several TODS subtasks, notably intent classification, dialogue state tracking, and response generation.

- We introduce InstructTODS, a fully off-the-shelf framework to leverage instruction-tuned LLMs in zero-shot setting for end-to-end unified task-oriented dialogue, with the benefit of being effectively adaptable to any knowledge base (KB) while alleviating the need for any additional form of task-relevant data, such as intent, belief state, system action, etc.
- We provide valuable insights from the TODS experiments on the more general advantages and failure cases of LLMs to perform complex zero-shot NLP tasks.

## 2 Evaluating LLMs on Zero-Shot Task-Oriented Dialogue Subtasks

As an intermediary step in exploring the potential of end-to-end TODS solutions, we first investigate how well the performance of state-of-the-art LLMs (we presented the comparison of different LLMs over multiple tasks in Appendix A), i.e., GPT-3.5 and GPT-4, in performing various modular task-oriented objectives in their respective settings.

### 2.1 TODS Subtasks

Let us define a dialogue set  $\mathcal{D}_n = \{u_1, r_1, u_2, r_2, \dots, u_n, r_n\}$  where  $u_i$  and  $r_i$  denotes the user utterance and the system reply at turn  $i$ , respectively.

**Intent Classification (IC)** For IC, we have the set of labels  $\mathcal{C} = \{c_1, c_2, \dots, c_t\}$ , from which we build the input for the LLM as  $x_i^{ic} = \mathbb{P}^{ic}(\mathcal{I}^{ic}, \text{Concat}(c_j)_{j=0}^t, u_i)$  where  $\mathbb{P}^{ic}(\cdot)$  is the IC

Model	Attraction		Hotel		Restaurant		Taxi		Train		Average	
	JGA	Slot-F1	JGA	Slot-F1	JGA	Slot-F1	JGA	Slot-F1	JGA	Slot-F1	JGA	Slot-F1
TRADE	20.06	–	14.20	–	12.59	–	59.21	–	22.39	–	25.69	–
MA-DST	22.46	–	16.28	–	13.56	–	59.27	–	22.76	–	26.87	–
TransferQA	31.25	–	22.72	–	26.28	–	61.87	–	36.72	–	35.77	–
T5Dep	37.83	–	26.50	–	27.05	–	<b>69.23</b>	–	40.27	–	40.18	–
Modular (GPT-3.5, w/ all slots)	30.23	65.38	26.77	76.28	48.28	82.90	56.22	75.33	53.75	83.64	42.02	78.60
Modular (GPT-3.5, w/ domain slot)	39.53	74.89	27.03	79.78	51.72	85.06	63.24	83.98	52.50	84.84	44.48	82.53
Modular (GPT-4, w/ all slots)	39.53	78.99	31.23	<b>84.07</b>	55.86	88.23	63.24	82.71	<b>59.83</b>	89.72	48.16	85.62
Modular (GPT-4, w/ domain slot)	<b>46.51</b>	<b>81.13</b>	<b>31.76</b>	83.42	<b>56.90</b>	<b>88.47</b>	<b>65.96</b>	<b>84.33</b>	52.50	<b>89.73</b>	<b>48.35</b>	<b>85.82</b>

Table 2: Performance comparison on **zero-shot DST benchmark**. LLMs outperform all baselines in our benchmark. Baseline results are directly taken from their respective works. The best performances in each section are in **bold**.

input template,  $\mathcal{I}^{ic}$  refers to the natural language instruction for IC and  $\text{Concat}(c_j)$  is the concatenation of all labels. We evaluate two generation settings, a single output setting where we query the model for the inferred intent, and a multi-output setting where we query the model for the top-3 intents given the user query by simply changing the instruction  $\mathcal{I}^{ic}$ . As such, we recast the classification task in a text-generation manner and compare our results with state-of-the-art IC baselines.

**Dialogue State Tracking (DST)** For DST, we define the total set of slots  $S = \{s_{1,D_1}, s_{2,D_1}, \dots, s_{k,D_l}\}$  where  $s_{i,D_j}$  is the  $i$ -th slot associated to domain  $D_j$ . We give a singular hand-crafted exemplar distinct from the dataset to guide the generation format directly as JSON. We build the input  $x_i^{dst} = \mathbb{P}^{dst}(\mathcal{I}^{dst}, f^{dst}(S), \mathcal{D}_i)$  by providing the entire dialogue context, where  $\mathbb{P}^{dst}(\cdot)$  is the DST input template,  $\mathcal{I}^{dst}$  denotes the instruction for DST  $f^{dst}(S)$  refers to a textual transformation of the set of slots. We evaluate two settings with different slot transformations: one by providing all slots and another with only the active domain slots.

**Response Generation (RG)** For RG, given a dialogue  $\mathcal{D}$ , we define the set of oracle system actions  $A = \{a_{1,1}, a_{1,2}, \dots, a_{n,m}\}$  where  $a_{i,j}$  denotes the  $j$ -th system action of turn  $i$ . We construct the input  $x_i^{rg} = \mathbb{P}^{rg}(\mathcal{I}^{rg}, f^{rg}(a_{i,1}, a_{i,2}, \dots, a_{i,m}), \mathcal{D}_i)$  where  $\mathbb{P}^{rg}(\cdot)$  is RG input template,  $\mathcal{I}^{rg}$  denotes the instruction for RG and  $f^{rg}(\cdot)$  refers to a textual transformation of the set of system actions. We evaluate the capability of LLMs to leverage a structured system action while addressing the dialogue context to generate a response to the user.

## 2.2 Experiment Settings

**Dataset** For the dialogue state tracking, we evaluate the LLMs’ capability on MultiWOZ 2.1

(MWOZ) (Eric et al., 2020). For intent classification, we evaluate two datasets: Banking77 (Casanueva et al., 2020), a fine-grained intent dataset in the banking domain, and CLINC150 (Larson et al., 2019), coarse-grained intents classification datasets covering over 10 different domains. The main challenge of the CLINC150 dataset is on inferring out-of-scope intent, which is particularly challenging without any model training.

**Evaluation** We evaluate dialogue state tracking with Joint Goal Accuracy (JGA) and Slot-F1. We compute JGA using exact matching instead of fuzzy matching, with minor typo fixes in MWOZ following prior works (Hosseini-Asl et al., 2020a; Su et al., 2022). For intent classification, we evaluate the accuracy when predicting only one intent (single) and the top-3 intents (multi) in a text generation setting. BLEU (Papineni et al., 2002), Inform, and Success (Eric et al., 2020) are used for response generation. In addition to these metrics, we also compare lexical diversity (Shen, 2022), i.e., HDD (McCarthy and Jarvis, 2010), MATTR (Covington and McFall, 2010), MTLT (McCarthy, 2005), and VOCD (McCarthy and Jarvis, 2007), fluency through perplexity, and human-likability using USL-H (Phy et al., 2020).

**Baseline** For intent classification, we compare with various few-shot fine-tuned baselines: RoBERTa (Liu et al., 2019), ICDA (Lin et al., 2023), DNNC (Zhang et al., 2020a), and CPFT (Zhang et al., 2021). While for zero-shot baseline, we employ MNLI (Williams et al., 2018) fine-tuned BART<sub>large</sub> models (Lewis et al., 2020) by framing intent classification as an NLI task. For dialogue state tracking, we compare with multiple strong zero-shot baselines in the single-domain setting: TRADE (Wu et al., 2019a), MA-DST (Kumar et al., 2020), TransferQA (Lin et al., 2021a) and T5Dep (Wang et al., 2022).

Model	Reference-based			Fluency	Lexical diversity				Human pref.
	BLEU	Inf.	Succ.	PPL	HDD	MATTR	MTLD	VOCD	USL-H
LAVA	11.33	95.8	<b>94.9</b>	<b>25.45</b>	65.35	74.84	30.72	25.84	59.68
SFN	14.11	<b>97.7</b>	91.6	51.97	70.68	78.67	34.25	36.02	65.41
DAMD	14.94	78	68.7	58.41	71.45	78.09	29.08	37.57	62.62
MARCO	16.5	95.3	91.1	36.00	73.40	83.39	44.48	42.78	70.35
MinTL	18.39	85	80.8	49.77	71.31	82.76	38.99	37.26	65.36
BORT	16.75	91.1	88.3	53.45	70.94	81.82	38.41	36.28	66.00
HDSA	<b>20.02</b>	95.8	90.2	43.37	71.71	82.95	42.04	38.02	68.36
RSTOD	15.98	91.6	86.9	76.05	73.11	82.41	42.08	41.88	68.54
Modular (GPT-4)	6.12	86.42	78.48	36.63	<b>80.59</b>	<b>89.56</b>	<b>66.64</b>	<b>70.13</b>	<b>89.66</b>

Table 3: Performance comparison on **response generation**. Although lower in BLEU, responses by the LLM-powered modular TODS are more human-preferred. The reported results for the baselines are taken from their respective work. The best performances in each group are in **bold**.

For the response generation, we compare with modular and non-unified end-to-end TODS—e.g., having split decoder modules for DST and response generation—including SFN (Mehri et al., 2019), LAVA (Lubis et al., 2020), DAMD (Zhang et al., 2020b), MARCO (Wang et al., 2020), MinTL (Lin et al., 2020), HDSA (Santra et al., 2021), RSTOD (Cholakov and Kolev, 2022), and BORT (Sun et al., 2022).

### 2.3 Key Takeaways

The evaluation results for intent classification, DST, and response generation are shown in Table 1, Table 2, and Table 3, respectively. We summarize the key insights as follows:

**LLMs outperform most baselines.** LLMs show significant improvements in intent classification and DST tasks compared to other zero-shot and few-shot baselines and perform almost comparably to few-shot models in the intent classification task.

**LLMs offer better generalization and adaptable solutions to TOD.** Unlike fine-tuned models, LLMs approach all tasks in an autoregressive generation manner, allowing greater flexibility and scalability to adapt to other tasks and domains.

**LLMs generate responses that better reflect human preference.** Unlike other fine-tuned approaches, LLMs generate responses that are distinct from the gold responses resulting in lower BLEU scores. Nevertheless, the responses from modular LLMs are more fluent, diverse, and human-likable compared to baselines, while having competitive Inform and Success rates.

**LLMs do not solve multi-domain DST problems** Despite their strong performance, LLMs often over-

predict active slots, leading to errors in the *all slots* setting. LLMs tend to mix up parallel slots over different domains, especially for either temporal or spatial information, e.g., destination and departure, leave time and arrival time, etc.

### 3 InstructTODS: An Instruction-Based Zero-shot End-to-End TODS

By leveraging the insights from solving TODS subtasks on §2.3, we develop the first zero-shot end-to-end framework that operates without any domain information (ontology) and requires no task-specific annotations such as dialogue state, system act, intent, etc. This method is not only cost-efficient but also alleviates the ontology constraint of LLMs in the modular DST task and promotes the strength of LLMs in generating better and more human-preferred responses. Let us define a structured knowledge base (KB) as a set of tuples  $\mathcal{K} = \{(v_1^{a_1}, \dots, v_1^{a_k}), \dots, (v_p^{a_1}, \dots, v_p^{a_k})\}$  where  $(a_i)_{i=0}^k$  are the attributes of the KB, and  $(v_j^{a_i})_{j=0}^p$  are all the values associated to the attribute  $a_i$ .

We first define a naive modular LLM response generation approach that serves as a baseline, denoted as  $\mathbf{RG}_{naive}$ .<sup>4</sup>  $\mathbf{RG}_{naive}$  generates the user response by taking the entire KB along with the dialogue context as input. In this approach, we rely on the ability of the LLM to parse the entire KB during inference while processing the dialogue context, in order to perform in-context retrieval and response generation at the same time. As such, we build the input  $x_i^{RG} = \mathbb{P}^{RG}(\mathcal{I}^{RG}, f^{RG}(\mathcal{K}), \mathcal{D}_i)$  where  $\mathbb{P}^{RG}(\cdot)$  is the response generation input template,  $\mathcal{I}^{RG}$  denotes the instruction for response generation and  $f^{RG}(\mathcal{K})$  refers to a textual transformation

<sup>4</sup> $\mathbf{RG}_{naive}$  is based on Bang et al. (2023)’s limited observation of zero-shot end-to-end TODS with ChatGPT (see §6.2).

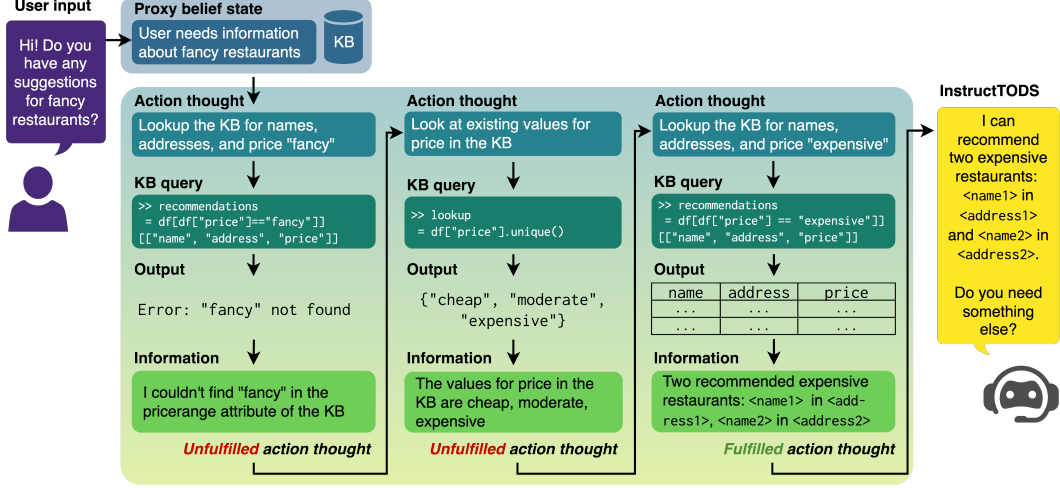


Figure 2: Overview of **InstructTODS**, a framework to utilize LLM for zero-shot end-to-end task-oriented dialogue.

of the KB where we filter unnecessary information and values that are too long as they are not needed to accomplish the user goal. In this approach, the bottleneck resides in the context window limit of the LLMs. Unlike other approaches, InstructTODS aims to make the best use of the LLM abilities to perform end-to-end tasks in zero-shot settings without the need for additional modular NLU and DST models, allowing zero-cost adaptation to various domains with no parameter update.

In general, in order to process the dialogue history and interact with the KB, InstructTODS introduces two concepts, i.e., proxy belief state and action thought. The results from KB and the dialogue history are then fed as a context to the LLM for generating the user response.

In the following paragraphs, we describe each component of InstructTODS in more detail.

**Proxy Belief State** We generate a proxy belief state  $\tilde{B}_i = \mathbb{P}^{BS}(\mathcal{D}_i)$  from the dialogue history where  $\mathbb{P}^{BS}(\cdot)$  denotes the prompt template and  $\mathcal{D}_i$  the dialogue context.  $\tilde{B}_i$  encapsulates everything that the user is looking for in natural language at this point of the dialogue. Note that, the proxy belief state does not need any prior knowledge about the domain nor any ontology to operate (e.g. domain, trackable slots, values, types of information, etc.). The proxy belief state is directly used to interact with the KB in a multi-turn fashion.

**KB Interaction** To interact with the KB, we generate an Action thought  $A = \mathbb{P}^{act}(\tilde{B}_i, (a_i)_{i=0}^k)$  where  $\mathbb{P}^{act}(\cdot)$  is the template for action generation and  $(a_i)_{i=0}^k$  the attributes of the KB. By providing

the existing attributes of the KB at this step, we ground the LLM to accurately translate the belief state into information that can be queried from the KB, while filtering out unnecessary data. The action thought serves as an intermediary to leverage the code generation ability of LLM by generating a query  $Q = \mathbb{P}^{KB}(A, \mathcal{K})$  where  $\mathbb{P}^{KB}(\cdot)$  is the template for code generation. The output from the KB is then parsed by the LLM to extract relevant information, denoted as  $I$ , presented in natural language, which provides a summary of the KB interaction. It also determines whether the current action thought has been fulfilled. If it remains unanswered, a new action thought is generated based on the extracted information, and the process repeats until a stopping criterion is reached indicating that no relevant knowledge is found in the KB.

**Response Generation** Once the KB interaction concludes, the final information, together with the original dialogue context, is passed to the model to generate the response  $Y = \mathbb{P}^{RG}(I, \mathcal{D}_i)$  where  $\mathbb{P}^{RG}$  represents the response generation template and  $I$  the final information from the KB interaction. In the case where no knowledge is found in the KB, the LLM prompts the user to provide additional information. We provide the prompt template in Appendix D. The depiction of how the InstructTODS framework works is presented in Figure 2.

## 4 Experiment settings

**Baselines** Our framework is compared to other end-to-end unified TODS approaches that perform end-to-end TODS using a unified text-to-text paradigm through a single generalized text

Model	Attraction		Hotel		Restaurant		Taxi		Train		All		
	Inf.	Succ.	Inf.	Succ.	Inf.	Succ.	Inf.	Succ.	Inf.	Succ.	BLEU	Inf.	Succ.
SOLOIST	<b>100</b>	<b>90.90</b>	<b>90.00</b>	<b>85.00</b>	78.30	70.00	<b>100</b>	<b>100</b>	81.80	78.80	13.58	88.80	84.30
UBAR	<b>100</b>	<b>90.90</b>	85.00	70.00	91.70	83.30	<b>100</b>	90.00	90.90	84.80	15.05	91.90	82.10
AUGPT	90.90	81.80	71.70	60.00	81.70	73.30	<b>100</b>	84.00	<b>97.00</b>	<b>93.90</b>	15.56	86.10	76.20
GALAXY	90.90	72.70	81.70	76.70	91.70	83.80	<b>100</b>	<b>100</b>	93.90	<b>93.90</b>	<b>18.10</b>	91.00	<b>86.10</b>
PPTOD	81.80	81.80	71.70	71.70	86.70	<b>86.70</b>	<b>100</b>	<b>100</b>	89.20	84.80	16.44	89.20	84.80
RG <sub>naive</sub>	81.80	36.37	<b>90.00</b>	83.33	<b>96.70</b>	83.33	<b>100</b>	89.80	<b>97.00</b>	63.67	3.95	<b>94.90</b>	82.16
InstructTODS	72.70	54.55	85.00	75.00	91.70	73.33	<b>100</b>	89.80	90.90	72.73	3.94	90.70	76.20

Table 4: **Task completion** performance comparison. InstructTODS have competitive Inform and Success rates compared to other end-to-end fine-tuned TODS baselines. **Bold** represents the highest score in each column.

generation model, i.e., SimpleTOD (Hosseini-Asl et al., 2020b), PPTOD (Su et al., 2022), Soloist (Peng et al., 2021), UBAR (Yang et al., 2021), AuGPT (Kulhánek et al., 2021), and Galaxy (He et al., 2022). In addition, as described in §3, we add the naive version of the LLM response generation approach which is fed by the full KB (RG<sub>naive</sub>), as an additional baseline to better evaluate the effectiveness of our framework.

**Datasets** We evaluate the end-to-end zero-shot capability on MultiWOZ 2.1 (MWOZ) (Eric et al., 2020; Lewkowycz et al., 2022). We split the evaluation into two settings, i.e., single-domain and multi-domain evaluation settings, where we show the capability of LLMs to tackle more complex TODS tasks in zero-shot end-to-end settings.

**Automatic Evaluation** For evaluating the end-to-end framework, we measure the per domain Inform and Success rates, and the BLEU (Papineni et al., 2002), Inform rate, and Success rate (Eric et al., 2020) for all domains. The evaluation metric is computed on the delexicalized responses to avoid favoring models that provide more information than others and focus solely on the vocabulary used for the response generation. Additionally, we also incorporate an automatic human-likability score, namely USL-H (Phy et al., 2020).

**Human Evaluation** We conduct an extensive human evaluation to measure the capability of LLMs in conducting zero-shot end-to-end unified TOD. Specifically, we conduct two human evaluations, which measure: 1) the informativeness, helpfulness, and humanness of the generated responses, and 2) the information correctness and hallucination rate of our InstructTODS. For evaluating informativeness, helpfulness, and humanness, we ask 3 annotators to rate the quality of the response using a 4-point Likert scale (see Appendix B). The sys-

tem is helpful if it answers the user’s request while pushing the conversation towards goal completion, informative if the system provides enough related information while answering the user, and human if the generated answer is fluent and human-preferred. For measuring the incorrectness and the hallucination rate, the metrics are evaluated by a single TOD expert. The incorrectness and hallucination rate are measured by manually checking the ratio of correct, incorrect, and hallucinated entities provided in the generated responses. We conduct the human evaluation by taking 50 generated responses from all the models and the gold responses.

## 5 Results and Analysis

### 5.1 Automatic Evaluation

Our automatic evaluation is shown in Table 4. In general, we find a similar trend with the modular LLMs where LLMs produce lower BLEU scores— $\sim 4$  BLEU against  $\sim 15$  BLEU—with competitive Inform and Success rates compared to other end-to-end unified TODS baselines. Note that, as mentioned in §2.3, LLMs often generate completely different responses to the gold knowledge, hence producing low automatic evaluation scores. Nevertheless, the low automatic evaluation scores do not sufficiently reflect the capability of InstructTODS. We will further elaborate on this in §5.2, raising a question of the sufficiency of evaluating TODS quality using only a single gold response. Some comparative generation samples between the different models can be found in Appendix C.

### 5.2 Informativeness, Helpfulness, and Humanness of InstructTODS

The results for our human evaluation are shown in Figure 3 for InstructTODS in comparison with the naive approach, the gold responses, and the two best-performing baselines in task completion (i.e.,

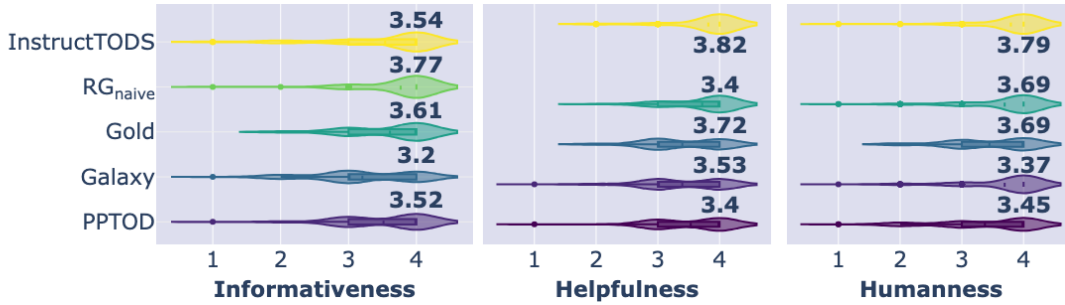


Figure 3: Human evaluation comparison on informativeness (left), helpfulness (center), and humanness (right).

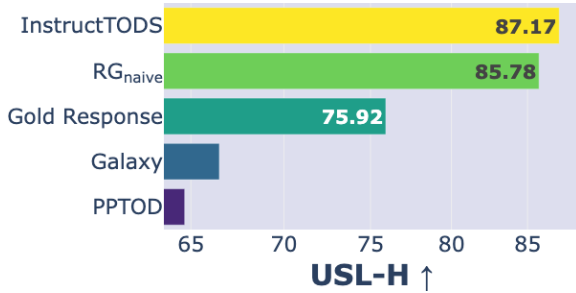


Figure 4: InstructTODS have higher human preference scores than the gold responses and baselines.

Galaxy and PPTOD). From the results, we show that InstructTODS is more informative, helpful, and human-like than the two fine-tuned end-to-end baselines by a noticeable margin. For both helpfulness and humanness, InstructTODS also outperforms RG<sub>naive</sub> and the gold response. Aligning with the human evaluation results, the generated responses by our framework also have higher humanness scores as shown in Figure 4, even higher than the gold responses. RG<sub>naive</sub> is the most informative, which is expected as the model processes the entire KB for information, however, the quality of the information greatly differs as shown in §5.3.

### 5.3 Incorrectness and Hallucination

We show the results for incorrectness and hallucination for the LLM-generated responses in Figure 5. While a sample can be incorrect, e.g., if the LLM database interaction fails, the LLMs do not necessarily generate unfaithful information. InstructTODS is more robust than naively employing the LLMs, improving the correctness by 15% and showing 11% of hallucination, half the amount of the RG<sub>naive</sub>. We observe that some types of information are more prone to hallucination, notably time and address. This bias towards temporal and spatial information aligns with our observation of LLMs’ performance in DST (§2.3).

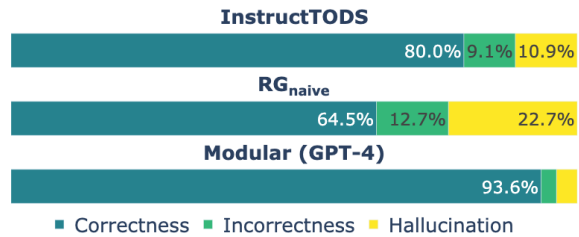


Figure 5: Human evaluation on correctness, incorrectness, and hallucination for RG<sub>naive</sub> and InstructTODS.

## 5.4 LLMs on Multi-Domain TOD

While it is possible to use InstructTODS in multi-domain with distinct KBs per domain, as we see in Figure 6, the performance degrades quickly for Success and slightly less for Inform as the number of domains increases. While fine-tuned end-to-end baselines operate with only one KB at a given turn by tracking the active domain through either state changes (Peng et al., 2021; Yang et al., 2021) or slot names (Kulhánek et al., 2021), our zero-shot framework does not assume any external knowledge nor ontology information. As such, all KBs are provided at each turn, and due to different KBs attributes overlapping in MWOZ, InstructTODS often queries incompatible information from the proxy belief state (e.g., "food" and "destination" at the same time), which are in different KBs. Hence, multi-domain degradation is largely due to the KB interaction failure.

## 6 Related Work

### 6.1 Task-Oriented Dialogue System

Task-oriented dialogue systems (TODS) can be broadly classified into two categories (Chen et al., 2017; Gao et al., 2018; Zhang et al., 2020c) which include pipelined approaches (Nagata and Morimoto, 1994; Levin et al., 1997, 2000; Hurtado et al., 2005; Williams and Young, 2007; Hori

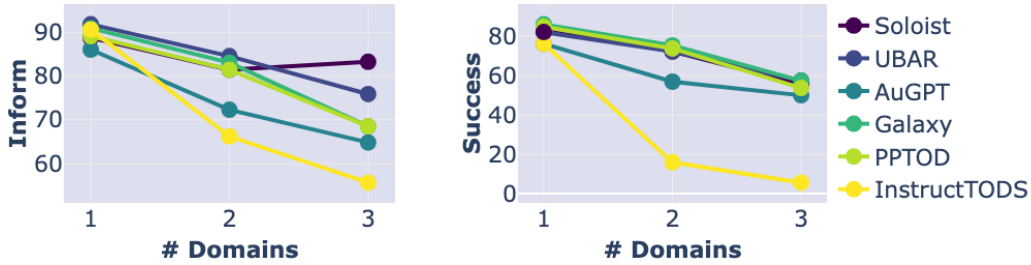


Figure 6: End-to-end TODS performance degrades as the number of active domains in the dialogue increases.

et al., 2009; Lee et al., 2009) and end-to-end approaches (Madotto et al., 2018; Wu et al., 2019b; Madotto et al., 2020b; Raghu et al., 2019; Qin et al., 2020; Hosseini-Asl et al., 2020a; Lei et al., 2018; Hosseini-Asl et al., 2020b; Lin et al., 2020; He et al., 2022; Kulhánek et al., 2021; Peng et al., 2021; Su et al., 2022; Yang et al., 2021). The pipelined approach utilizes multiple modules in order to generate the system responses. While the end-to-end approach directly generates responses from the user input and the KB in an end-to-end manner.

**End-to-end TODS** Early approaches for end-to-end TODS employ template responses in a retrieval or generation setting (Zhao et al., 2017; Eric et al., 2017; Wu et al., 2019b). While other approaches inject KBs directly into the model to perform end-to-end generation (Madotto et al., 2018, 2020a). A more recent end-to-end TODS tackles end-to-end response generation in a single sequence prediction problem (Hosseini-Asl et al., 2020a; Yang et al., 2021; Peng et al., 2021) with an autoregressive model. These approaches still mostly leverage TOD data (belief states, system acts, etc.) during generation. As general pre-trained LMs were shown to be effective for TODS (Mehri et al., 2019; Lubis et al., 2020; Lin et al., 2020), several subsequent works have explored pre-training approaches directly tailored towards TODS (Zhang et al., 2020b; Su et al., 2022; He et al., 2022). To the best of our knowledge, prior works require a structured format of dialogue states, system acts, and/or template responses, whereas InstructTODS alleviates such needs by incorporating an unstructured proxy belief state, which requires no domain-specific knowledge nor ontology to operate, allowing zero-shot adaptation to various TOD domains.

## 6.2 Zero-Shot Generalization of LLMs

LLMs have shown remarkable zero-shot generalization capabilities in various NLP tasks (Brown et al., 2020; Scao et al., 2022; Chowdhery et al., 2022;

Thoppilan et al., 2022). This is further improved through instruction tuning (Wei et al., 2021; Sanh et al., 2021; Wei et al., 2022; Chung et al., 2022; Longpre et al., 2023; Cahyawijaya et al., 2023b), which enables a better generalization to unseen tasks, and reinforcement learning with human feedback (Christiano et al., 2017; Ouyang et al., 2022; Bai et al., 2022), which enables a better alignment of human preferences. The zero-shot generalization ability of LLMs has also been explored in more specific cases, e.g., multiple choice question answering (Robinson and Wingate, 2023), biomedical NLP (Fries et al., 2022), reasoning (Bang et al., 2023), low-resource languages (Cahyawijaya et al., 2023a,b; Asai et al., 2023), code-switching (Yong et al., 2023; Zhang et al., 2023).

**LLMs for TODS** Recent works explore the applicability of LLMs in solving modular TOD tasks (Bang et al., 2023; Hudeček and Dušek, 2023) and a pipeline manner (Hosseini-Asl et al., 2020b; Su et al., 2022; Peng et al., 2021; Yang et al., 2021; Kulhánek et al., 2021; He et al., 2022). Additionally, Bang et al. (2023) inspect ChatGPT’s capability for zero-shot end-to-end TODS, however, it is limited to only  $\sim 1\%$  of the test set available. Therefore, to the best of our knowledge, our work is the first to comprehensively study the utilization of LLMs for zero-shot end-to-end TODS.

## 7 Conclusion

In this paper, we introduce InstructTODS, an off-the-shelf framework to effectively perform end-to-end TODS in zero-shot utilizing LLMs. We compare InstructTODS to several state-of-the-art fully fine-tuned end-to-end TODS and show that InstructTODS manages to guide the conversation towards goal completion similarly to the fine-tuned systems on MWOZ while generating answers that are more informative, helpful, and human-like than previous approaches. Furthermore, we investigate the capability of LLMs in performing various TOD



subtasks in zero-shot settings, demonstrating better diversity and human preference on response generation, and state-of-the-art zero-shot results on dialogue state tracking and intent classification.

## 8 Limitation

**Generalization to Other Datasets** In our work, we only assess the effectiveness of InstructTODS on MultiWoZ 2.1 dataset, whose size is a magnitude higher than other TODS datasets (Eric et al., 2020). We conjecture that the generalization to other datasets will follow the same trend as described in §5, where it excels in the single-domain setting while still struggling in the multi-domain setting. We expect future work to extend the assessment on InstructTODS to other datasets and domains.

**Generalization to Other Languages** In recent years, various task-oriented dialogue systems in languages other than English have been introduced, such as CrossWoZ (Zhu et al., 2020), BiTOD (Lin et al., 2021b), GlobalWoZ (Ding et al., 2022), and COD (Majewska et al., 2023). As suggested in prior works evaluating LLMs in low-resource languages (Bang et al., 2023; Asai et al., 2023; Cahyawijaya et al., 2023b,a; Workshop et al., 2023; Kabra et al., 2023; Zhang et al., 2023), we conjecture that the performance in other languages follow the general trend in LLMs where the performance in low-resource languages will be lower compared to the high-resource languages. Future work might explore and further extend methods for improving the generalization of InstructTODS to other languages.

**Generalization to Other LLMs** In this work, we only explore two proprietary LLMs which display strong performance on various NLP tasks, i.e., GPT-3.5 and GPT-4. Despite the lack of transparency of these models, we expect that when other publicly available LLMs achieve the same performance as these proprietary LLMs, a similar capability of zero-shot end-to-end TODS will emerge. We expect future work to explore the generalization of InstructTODS and its improvement in other LLMs.

## 9 Ethics Statement

Our research endeavors to develop an off-the-shelf framework for zero-shot end-to-end Task-Oriented Dialogue Systems (TODS) using Large Language

Models (LLMs). It is important to note that this study does not involve the use of any sensitive data and the experimental evaluation is conducted on publicly available datasets. To ensure the quality of our results, we have employed crowdsourcing for the human evaluation of the generated dialogue responses. While our study does not raise any ethical concerns regarding privacy, confidentiality, or bias, we recognize that the use of LLMs in dialogue systems may have ethical implications related to potential biases in the training data and the generated responses. Therefore, we emphasize the importance of ongoing research toward developing ethical guidelines and best practices for the use of LLMs in dialogue systems. In line with our commitment to transparency and reproducibility, we will be releasing our code publicly. We believe that this will encourage open and collaborative research towards the development of more ethical and effective dialogue systems.

## References

- Anonymous. 2023. Nusawrites: Constructing high-quality corpora for underrepresented and extremely low-resource languages. Anonymous preprint under review.
- Akari Asai, Sneha Kudugunta, Xinyan Velocity Yu, Terra Blevins, Hila Gonen, Machel Reid, Yulia Tsvetkov, Sebastian Ruder, and Hannaneh Hajishirzi. 2023. [Buffet: Benchmarking large language models for few-shot cross-lingual transfer](#).
- Razvan Azamfirei, Sapna R Kudchadkar, and James Fackler. 2023. Large language models and the perils of their hallucinations. *Critical Care*, 27(1):1–2.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#).
- Suman Banerjee and Mitesh M Khapra. 2019. Graph convolutional network with sequential attention for goal-oriented dialogue systems. *Transactions of the Association for Computational Linguistics*, 7:485–500.
- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, et al. 2023. A multi-task, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity. *arXiv preprint arXiv:2302.04023*.
- Rishi Bommasani, Percy Liang, and Tony Lee. 2022. Holistic evaluation of language models. *Annals of the New York Academy of Sciences*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Samuel Cahyawijaya, Holy Lovenia, Alham Fikri Aji, Genta Indra Winata, Bryan Wilie, Rahmad Mahendra, Christian Wibisono, Ade Romadhony, Karissa Vincentio, Fajri Koto, Jennifer Santoso, David Moeljadi, Cahya Wirawan, Frederikus Hudi, Ivan Halim Parnangan, Ika Alfina, Muhammad Satrio Wicaksono, Ilham Firdausi Putra, Samsul Rahmadani, Yulianti Oenang, Ali Akbar Septiandri, James Jaya, Kaustubh D. Dhole, Arie Ardiyanti Suryani, Rifki Afina Putri, Dan Su, Keith Stevens, Made Nindyatama Nityasya, Muhammad Farid Adilazuarda, Ryan Ignatius, Ryandito Diandaru, Tiezheng Yu, Vito Ghifari, Wenliang Dai, Yan Xu, Dyah Damapusita, Cuk Tho, Ichwanul Muslim Karo Karo, Tirana Noor Fatyanosa, Ziwei Ji, Pascale Fung, Graham Neubig, Timothy Baldwin, Sebastian Ruder, Herry Sujaini, Sakriani Sakti, and Ayu Purwarianti. 2023a. [Nusacrowd: Open source initiative for indonesian nlp resources](#).
- Samuel Cahyawijaya, Holy Lovenia, Tiezheng Yu, Willy Chung, and Pascale Fung. 2023b. [Instruct-align: Teaching novel languages with llms through alignment-based cross-lingual instruction](#).
- Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. 2020. Efficient intent detection with dual sentence encoders. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 38–45.
- Hongshen Chen, Xiaorui Liu, Dawei Yin, and Jiliang Tang. 2017. [A survey on dialogue systems: Recent advances and new frontiers](#). *SIGKDD Explor. Newsl.*, 19(2):25–35.
- Radostin Cholakov and Todor Kolev. 2022. [Efficient task-oriented dialogue systems with response selection as an auxiliary task](#). In *Proceedings of the 5th International Conference on Natural Language and Speech Processing (ICNLSP 2022)*, pages 12–18, Trento, Italy. Association for Computational Linguistics.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.
- Paul F. Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 4302–4310, Red Hook, NY, USA. Curran Associates Inc.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- Michael A. Covington and Joe D. McFall. 2010. [Cutting the gordian knot: The moving-average type–token ratio \(MATTR\)](#). *Journal of Quantitative Linguistics*, 17(2):94–100.
- 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 (Volume 1: Long Papers)*, pages 1639–1657, Dublin, Ireland. Association for Computational Linguistics.

- Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyang Gao, Adarsh Kumar, Anuj Goyal, Peter Ku, and Dilek Hakkani-Tur. 2020. Multiwoz 2.1: A consolidated multi-domain dialogue dataset with state corrections and state tracking 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 SIGdial Meeting on Discourse and Dialogue*, pages 37–49.
- Jason Fries, Leon Weber, Natasha Seelam, Gabriel Al-tay, Debajyoti Datta, Samuele Garda, Sunny Kang, Rosaline Su, Wojciech Kusa, Samuel Cahyawijaya, Fabio Barth, Simon Ott, Matthias Samwald, Stephen Bach, Stella Biderman, Mario Sanger, Bo Wang, Alison Callahan, Daniel Le3n Perian, Th3o Gigant, Patrick Haller, Jenny Chim, Jose Posada, John Giorgi, Karthik Rangasai Sivaraman, Marc Pamies, Marianna Nezhurina, Robert Martin, Michael Cullan, Moritz Freidank, Nathan Dahlberg, Shubhan-shu Mishra, Shamik Bose, Nicholas Broad, Yanis Labrak, Shlok Deshmukh, Sid Kiblawi, Ayush Singh, Minh Chien Vu, Trishala Neeraj, Jonas Golde, Albert Villanova del Moral, and Benjamin Beilharz. 2022. [Bigbio: A framework for data-centric biomedical natural language processing](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 25792–25806. Curran Associates, Inc.
- 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. ACM.
- Donghoon Ham, Jeong-Gwan Lee, Youngsoo Jang, and Kee-Eung Kim. 2020. End-to-end neural pipeline for goal-oriented dialogue systems using gpt-2. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 583–592.
- Wanwei He, Yinpei Dai, Yinhe Zheng, Yuchuan Wu, Zheng Cao, Dermot Liu, Peng Jiang, Min Yang, Fei Huang, Luo Si, et al. 2022. Galaxy: A generative pre-trained model for task-oriented dialog with semi-supervised learning and explicit policy injection. *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. In *International Conference on Learning Representations*.
- Chiori Hori, Kiyonori Ohtake, Teruhisa Misu, Hideki Kashioka, and Satoshi Nakamura. 2009. Statistical dialog management applied to wfst-based dialog systems. In *IEEE International Conference on Acoustics, Speech and Signal Processing, 2009. ICASSP 2009.*, pages 4793–4796. IEEE.
- Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. 2020a. A simple language model for task-oriented dialogue. *Advances in Neural Information Processing Systems*, 33:20179–20191.
- Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. 2020b. [A simple language model for task-oriented dialogue](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 20179–20191. Curran Associates, Inc.
- Vojt3ch Hudeek and Ondr3ej Duek. 2023. Are llms all you need for task-oriented dialogue? *arXiv preprint arXiv:2304.06556*.
- Llus F Hurtado, David Griol, Emilio Sanchis, and Encarna Segarra. 2005. A stochastic approach to dialog management. In *IEEE Workshop on Automatic Speech Recognition and Understanding, 2005.*, pages 226–231. IEEE.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38.
- Anubha Kabra, Emmy Liu, Simran Khanuja, Al-ham Fikri Aji, Genta Indra Winata, Samuel Cahyawijaya, Anuoluwapo Aremu, Perez Ogayo, and Graham Neubig. 2023. [Multi-lingual and multi-cultural figurative language understanding](#).
- Jonas Kulhanek, Vojt3ch Hudeek, Tomas Nekvinda, and Ondr3ej Duek. 2021. [AuGPT: Auxiliary tasks and data augmentation for end-to-end dialogue with pre-trained language models](#). In *Proceedings of the 3rd Workshop on Natural Language Processing for Conversational AI*, pages 198–210, Online. Association for Computational Linguistics.
- Adarsh Kumar, Peter Ku, Anuj Goyal, Angeliki Metallinou, and Dilek Hakkani-Tur. 2020. Ma-dst: Multi-attention-based scalable dialog state tracking. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 8107–8114.
- Stefan Larson, Anish Mahendran, Joseph J Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K Kummerfeld, Kevin Leach, Michael A Laurenzano, Lingjia Tang, et al. 2019. An evaluation dataset for intent classification and out-of-scope prediction. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1311–1316.
- Cheongjae Lee, Sangkeun Jung, Seokhwan Kim, and Gary Geunbae Lee. 2009. Example-based dialog modeling for practical multi-domain dialog system. *Speech Communication*, 51(5):466–484.

- Wenqiang Lei, Xisen Jin, Min-Yen Kan, Zhaochun Ren, Xiangnan He, and Dawei Yin. 2018. Sequicity: Simplifying task-oriented dialogue systems with single sequence-to-sequence architectures. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1437–1447.
- Esther Levin, Roberto Pieraccini, and Wieland Eckert. 1997. Learning dialogue strategies within the markov decision process framework. In *1997 IEEE Workshop on Automatic Speech Recognition and Understanding Proceedings*, pages 72–79. IEEE.
- Esther Levin, Roberto Pieraccini, and Wieland Eckert. 2000. A stochastic model of human-machine interaction for learning dialog strategies. *IEEE Transactions on speech and audio processing*, 8(1):11–23.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880.
- Aitor Lewkowycz, Ambrose Slone, Anders Andreassen, Daniel Freeman, Ethan S Dyer, Gaurav Mishra, Guy Gur-Ari, Jaehoon Lee, Jascha Sohl-dickstein, Kristen Chiafullo, et al. 2022. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models.
- Yen-Ting Lin, Alexandros Papangelis, Seokhwan Kim, Sungjin Lee, Devamanyu Hazarika, Mahdi Namazifar, Di Jin, Yang Liu, and Dilek Hakkani-Tur. 2023. Selective in-context data augmentation for intent detection using pointwise V-information. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1463–1476, Dubrovnik, Croatia. Association for Computational Linguistics.
- Zhaojiang Lin, Bing Liu, Andrea Madotto, Seungwhan Moon, Zhenpeng Zhou, Paul A Crook, Zhiguang Wang, Zhou Yu, Eunjoon Cho, Rajen Subba, et al. 2021a. Zero-shot dialogue state tracking via cross-task transfer. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7890–7900.
- Zhaojiang Lin, Andrea Madotto, Genta Indra Winata, and Pascale Fung. 2020. MinTL: Minimalist transfer learning for task-oriented dialogue systems. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3391–3405, Online. Association for Computational Linguistics.
- 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*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach.
- Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V Le, Barret Zoph, Jason Wei, et al. 2023. The flan collection: Designing data and methods for effective instruction tuning. *arXiv preprint arXiv:2301.13688*.
- Nurul Lubis, Christian Geishhauser, Michael Heck, Hsien-chin Lin, Marco Moresi, Carel van Niekerk, and Milica Gasic. 2020. LAVA: Latent action spaces via variational auto-encoding for dialogue policy optimization. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 465–479, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Andrea Madotto, Samuel Cahyawijaya, Genta Indra Winata, Yan Xu, Zihan Liu, Zhaojiang Lin, and Pascale Fung. 2020a. Learning knowledge bases with parameters for task-oriented dialogue systems. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2372–2394.
- Andrea Madotto, Zhaojiang Lin, Chien-Sheng Wu, Jamin Shin, and Pascale Fung. 2020b. Attention over parameters for dialogue systems.
- Andrea Madotto, Chien-Sheng Wu, and Pascale Fung. 2018. Mem2seq: Effectively incorporating knowledge bases into end-to-end task-oriented dialog systems. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1468–1478.
- 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*, 11:139–156.
- Philip M. McCarthy. 2005. *An assessment of the range and usefulness of lexical diversity measures and the potential of the measure of textual, lexical diversity (MTLD)*. Ph.D. thesis, The University of Memphis.
- Philip M McCarthy and Scott Jarvis. 2007. vocd: A theoretical and empirical evaluation. *Language Testing*, 24(4):459–488.
- Philip M. McCarthy and Scott Jarvis. 2010. MTLT, vocd-d, and HD-d: A validation study of sophisticated approaches to lexical diversity assessment. *Behavior Research Methods*, 42(2):381–392.
- Shikib Mehri, Tejas Srinivasan, and Maxine Eskenazi. 2019. Structured fusion networks for dialog. In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 165–177, Stockholm, Sweden. Association for Computational Linguistics.

- Masaaki Nagata and Tsuyoshi Morimoto. 1994. First steps towards statistical modeling of dialogue to predict the speech act type of the next utterance. *Speech communication*, 15(3-4):193–203.
- Atsumoto Ohashi and Ryuichiro Higashinaka. 2022. Post-processing networks: Method for optimizing pipeline task-oriented dialogue systems using reinforcement learning. In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 1–13.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#).
- 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.
- Baolin Peng, Chunyuan Li, Jinchao Li, Shahin Shayan-deh, Lars Liden, and Jianfeng Gao. 2021. [Soloist: Building task bots at scale with transfer learning and machine teaching](#). *Transactions of the Association for Computational Linguistics*, 9:807–824.
- Vitou Phy, Yang Zhao, and Akiko Aizawa. 2020. [Deconstruct to reconstruct a configurable evaluation metric for open-domain dialogue systems](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4164–4178, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Libo Qin, Xiao Xu, Wanxiang Che, Yue Zhang, and Ting Liu. 2020. Dynamic fusion network for multi-domain end-to-end task-oriented dialog. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6344–6354.
- Dinesh Raghu, Nikhil Gupta, et al. 2019. Disentangling language and knowledge in task-oriented dialogs. 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 1239–1255.
- Joshua Robinson and David Wingate. 2023. [Leveraging large language models for multiple choice question answering](#). In *The Eleventh International Conference on Learning Representations*.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Stella Biderman, Leo Gao, Tali Bers, Thomas Wolf, and Alexander M. Rush. 2021. [Multi-task prompted training enables zero-shot task generalization](#).
- Bishal Santra, Potnuru Anusha, and Pawan Goyal. 2021. [Hierarchical transformer for task oriented dialog systems](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5649–5658, Online. Association for Computational Linguistics.
- Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*.
- Lucas Shen. 2022. [LexicalRichness: A small module to compute textual lexical richness](#).
- 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 (Volume 1: Long Papers)*, pages 4661–4676.
- Haipeng Sun, Junwei Bao, Youzheng Wu, and Xiaodong He. 2022. [BORT: Back and denoising reconstruction for end-to-end task-oriented dialog](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2156–2170, Seattle, United States. Association for Computational Linguistics.
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, YaGuang Li, Hongrae Lee, Huaixiu Steven Zheng, Amin Ghafouri, Marcelo Menegali, Yanping Huang, Maxim Krikun, Dmitry Lepikhin, James Qin, Dehao Chen, Yuanzhong Xu, Zhifeng Chen, Adam Roberts, Maarten Bosma, Vincent Zhao, Yanqi Zhou, Chung-Ching Chang, Igor Krivokon, Will Rusch, Marc Pickett, Pranesh Srinivasan, Laichee Man, Kathleen Meier-Hellstern, Meredith Ringel Morris, Tulsee Doshi, Renelito Delos Santos, Toju Duke, Johnny Soraker, Ben Zevenbergen, Vinodkumar Prabhakaran, Mark Diaz, Ben Hutchinson, Kristen Olson, Alejandra Molina, Erin Hoffman-John, Josh Lee, Lora Aroyo, Ravi Rajakumar, Alena Butryna, Matthew Lamm, Viktoriya Kuzmina, Joe Fenton, Aaron Cohen, Rachel Bernstein, Ray Kurzweil, Blaise Aguerre-Arcas, Claire Cui, Marian Croak, Ed Chi, and Quoc Le. 2022. [Lamda: Language models for dialog applications](#).

- Kai Wang, Junfeng Tian, Rui Wang, Xiaojun Quan, and Jianxing Yu. 2020. [Multi-domain dialogue acts and response co-generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7125–7134, Online. Association for Computational Linguistics.
- Qingyue Wang, Yanan Cao, Piji Li, Yanhe Fu, Zheng Lin, and Li Guo. 2022. Slot dependency modeling for zero-shot cross-domain dialogue state tracking. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 510–520.
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022. [Chain of thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Jason D Williams and Steve Young. 2007. Partially observable markov decision processes for spoken dialog systems. *Computer Speech & Language*, 21(2):393–422.
- BigScience Workshop, :, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Lucicioni, François Yvon, Matthias Gallé, Jonathan Tow, Alexander M. Rush, Stella Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas Muennighoff, Albert Villanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas Bekman, Angelina McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz Suarez, Victor Sanh, Hugo Laurençon, Yacine Jernite, Julien Launay, Margaret Mitchell, Colin Raffel, Aaron Gokaslan, Adi Simhi, Aitor Soroa, Alham Fikri Aji, Amit Alfassy, Anna Rogers, Ariel Kreisberg Nitzav, Canwen Xu, Chenghao Mou, Chris Emezue, Christopher Klamm, Colin Leong, Daniel van Strien, David Ifeoluwa Adelani, Dragomir Radev, Eduardo González Ponferrada, Efrat Levkovizh, Ethan Kim, Eyal Bar Natan, Francesco De Toni, Gérard Dupont, Germán Kruszewski, Giada Pistilli, Hady Elsahar, Hamza Benyamina, Hieu Tran, Ian Yu, Idris Abdulmumin, Isaac Johnson, Itziar Gonzalez-Dios, Javier de la Rosa, Jenny Chim, Jesse Dodge, Jian Zhu, Jonathan Chang, Jörg Frohberg, Joseph Tobing, Joydeep Bhattacharjee, Khalid Al-mubarak, Kimbo Chen, Kyle Lo, Leandro Von Werra, Leon Weber, Long Phan, Loubna Ben allal, Ludovic Tanguy, Manan Dey, Manuel Romero Muñoz, Maraim Masoud, María Grandury, Mario Šaško, Max Huang, Maximin Coavoux, Mayank Singh, Mike Tian-Jian Jiang, Minh Chien Vu, Mohammad A. Jauhar, Mustafa Ghaleb, Nishant Subramani, Nora Kassner, Nurulaqilla Khamis, Olivier Nguyen, Omar Espejel, Ona de Gibert, Paulo Villegas, Peter Henderson, Pierre Colombo, Priscilla Amuok, Quentin Lhoest, Rheza Harliman, Rishi Bommasani, Roberto Luis López, Rui Ribeiro, Salomey Osei, Sampo Pyysalo, Sebastian Nagel, Shamik Bose, Shamsuddeen Hassan Muhammad, Shanya Sharma, Shayne Longpre, Somaieh Nikpoor, Stanislav Silberberg, Suhas Pai, Sydney Zink, Tiago Timponi Torrent, Timo Schick, Tristan Thrush, Valentin Danchev, Vassilina Nikoulina, Veronika Laippala, Violette Lepercq, Vrinda Prabhu, Zaid Alyafeai, Zeerak Talat, Arun Raja, Benjamin Heinzerling, Chenglei Si, Davut Emre Taşar, Elizabeth Salesky, Sabrina J. Mielke, Wilson Y. Lee, Abheesht Sharma, Andrea Santilli, Antoine Chaffin, Arnaud Stiegler, Debajyoti Datta, Eliza Szczechla, Gunjan Chhablani, Han Wang, Harshit Pandey, Hendrik Strobelt, Jason Alan Fries, Jos Rozen, Leo Gao, Lintang Sutawika, M Saiful Bari, Maged S. Al-shaibani, Matteo Manica, Nihal Nayak, Ryan Teehan, Samuel Albanie, Sheng Shen, Srulik Ben-David, Stephen H. Bach, Taewoon Kim, Tali Bers, Thibault Fevry, Trishala Neeraj, Urmish Thakker, Vikas Raunak, Xiangru Tang, Zheng-Xin Yong, Zhiqing Sun, Shaked Brody, Yallow Uri, Hadar Tojarieh, Adam Roberts, Hyung Won Chung, Jaesung Tae, Jason Phang, Ofir Press, Conglong Li, Deepak Narayanan, Hatim Bourfoune, Jared Casper, Jeff Rasley, Max Ryabinin, Mayank Mishra, Minjia Zhang, Mohammad Shoeybi, Myriam Peyrounette, Nicolas Patry, Nouamane Tazi, Omar Sanseviero, Patrick von Platen, Pierre Cornette, Pierre François Lavallée, Rémi Lacroix, Samyam Rajbhandari, Sanchit Gandhi, Shaden Smith, Stéphane Reuena, Suraj Patil, Tim Dettmers, Ahmed Baruwa, Amanpreet Singh, Anastasia Cheveleva, Anne-Laure Ligozat, Arjun Subramonian, Aurélie Névél, Charles Lovering, Dan Garrette, Deepak Tunuguntla, Ehud Reiter, Ekaterina Taktasheva, Ekaterina Voloshina, Eli Bogdanov, Genta Indra Winata, Hailey Schoelkopf, Jan-Christoph Kalo, Jekaterina Novikova, Jessica Zosa Forde, Jordan Clive, Jungo Kasai, Ken Kawamura, Liam Hazan, Marine Carpuat, Miruna Clinciu, Najeon Kim, Newton Cheng, Oleg Serikov, Omer Antverg, Oskar van der Wal, Rui Zhang, Ruochen Zhang, Sebastian Gehrmann, Shachar Mirkin, Shani Pais, Tatiana Shavrina, Thomas Scialom, Tian Yun, Tomasz Limisiewicz, Verena Rieser, Vitaly Protasov, Vladislav Mikhailov, Yada Pruksachatkun, Yonatan Belinkov, Zachary Bamberger, Zdeněk Kasner, Alice Rueda, Amanda Pestana, Amir Feizpour, Ammar Khan, Amy Faranak, Ana Santos, Anthony Hevia, Antigona Undrea, Arash Aghagol, Arezoo Abdollahi, Aycha Tammour, Azadeh HajiHosseini, Bahareh Behroozi, Benjamin Ajibade, Bharat

- Saxena, Carlos Muñoz Ferrandis, Danish Contractor, David Lansky, Davis David, Douwe Kiela, Duong A. Nguyen, Edward Tan, Emi Baylor, Ezinwanne Ozoani, Fatima Mirza, Frankline Onon-iwu, Habib Rezanejad, Hessie Jones, Indrani Bhat-tacharya, Irene Solaiman, Irina Sedenko, Isar Ne-jadgholi, Jesse Passmore, Josh Seltzer, Julio Bonis Sanz, Livia Dutra, Mairon Samagaio, Maraim El-badri, Margot Mieskes, Marissa Gerchick, Martha Akinlolu, Michael McKenna, Mike Qiu, Muhammed Ghauri, Mykola Burynok, Nafis Abrar, Nazneen Ra-jani, Nour Elkott, Nour Fahmy, Olanrewaju Samuel, Ran An, Rasmus Kromann, Ryan Hao, Samira Al-izadeh, Sarmad Shubber, Silas Wang, Sourav Roy, Sylvain Viguier, Thanh Le, Tobi Oyebade, Trieu Le, Yoyo Yang, Zach Nguyen, Abhinav Ramesh Kashyap, Alfredo Palasciano, Alison Callahan, Anima Shukla, Antonio Miranda-Escalada, Ayush Singh, Benjamin Beilharz, Bo Wang, Caio Brito, Chenxi Zhou, Chirag Jain, Chuxin Xu, Clémentine Fourrier, Daniel León Perinián, Daniel Molano, Dian Yu, Enrique Manjavac-as, Fabio Barth, Florian Fuhrmann, Gabriel Altay, Giyaseddin Bayrak, Gully Burns, Helena U. Vrabec, Imane Bello, Ishani Dash, Jihyun Kang, John Giorgi, Jonas Golde, Jose David Posada, Karthik Ranga-sai Sivaraman, Lokesh Bulchandani, Lu Liu, Luisa Shinzato, Madeleine Hahn de Bykhovetz, Maiko Takeuchi, Marc Pàmies, Maria A Castillo, Mari-anna Nezhurina, Mario Sängler, Matthias Samwald, Michael Cullan, Michael Weinberg, Michiel De Wolf, Mina Mihaljcic, Minna Liu, Moritz Freidank, Myungsun Kang, Natasha Seelam, Nathan Dahlberg, Nicholas Michio Broad, Nikolaus Muellner, Pascale Fung, Patrick Haller, Ramya Chandrasekhar, Renata Eisenberg, Robert Martin, Rodrigo Canalli, Rosaline Su, Ruisi Su, Samuel Cahyawijaya, Samuele Garda, Shlok S Deshmukh, Shubhanshu Mishra, Sid Ki-blawi, Simon Ott, Sinee Sang-aroonsiri, Srishti Ku-mar, Stefan Schweter, Sushil Bharati, Tanmay Laud, Théo Gigant, Tomoya Kainuma, Wojciech Kusa, Ya-nis Labrak, Yash Shailesh Bajaj, Yash Venkatraman, Yifan Xu, Yingxin Xu, Yu Xu, Zhe Tan, Zhongli Xie, Zifan Ye, Mathilde Bras, Younes Belkada, and Thomas Wolf. 2023. [Bloom: A 176b-parameter open-access multilingual language model](#).
- Chien-Sheng Wu, Andrea Madotto, Ehsan Hosseini-Asl, Caiming Xiong, Richard Socher, and Pascale Fung. 2019a. 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.
- Chien-sheng Wu, Richard Socher, and Caiming Xiong. 2019b. Global-to-local memory pointer networks for task-oriented dialogue. In *7th International Conference on Learning Representations, ICLR 2019*.
- Minghao Wu, Abdul Waheed, Chiyu Zhang, Muham-mad Abdul-Mageed, and Alham Fikri Aji. 2023. [Lamini-lm: A diverse herd of distilled models from large-scale instructions](#).
- Yunyi Yang, Yunhao Li, and Xiaojun Quan. 2021. [UBAR: Towards fully end-to-end task-oriented di-alog system with GPT-2](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(16):14230–14238.
- Zheng-Xin Yong, Ruochen Zhang, Jessica Zosa Forde, Skyler Wang, Samuel Cahyawijaya, Holy Lovenia, Genta Indra Winata, Lintang Sutawika, Jan Chris-tian Blaise Cruz, Long Phan, Yin Lin Tan, and Al-ham Fikri Aji. 2023. [Prompting multilingual large language models to generate code-mixed texts: The case of south east asian languages](#).
- Jianguo Zhang, Trung Bui, Seunghyun Yoon, Xiang Chen, Zhiwei Liu, Congying Xia, Quan Hung Tran, Walter Chang, and Philip Yu. 2021. [Few-shot intent detection via contrastive pre-training and fine-tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1906–1912, Online and Punta Cana, Dominican Re-public. Association for Computational Linguistics.
- Jianguo Zhang, Kazuma Hashimoto, Wenhao Liu, Chien-Sheng Wu, Yao Wan, Philip Yu, Richard Socher, and Caiming Xiong. 2020a. [Discrimina-tive nearest neighbor few-shot intent detection by transferring natural language inference](#). In *Proceed-ings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5064–5082, Online. Association for Computational Linguistics.
- Ruochen Zhang, Samuel Cahyawijaya, Jan Chris-tian Blaise Cruz, and Alham Fikri Aji. 2023. [Mul-tilingual large language models are not \(yet\) code-switchers](#).
- Yichi Zhang, Zhijian Ou, and Zhou Yu. 2020b. [Task-oriented dialog systems that consider multiple appro-priate responses under the same context](#). *Proceed-ings of the AAAI Conference on Artificial Intelligence*, 34(05):9604–9611.
- Zheng Zhang, Ryuichi Takanobu, Qi Zhu, MinLie Huang, and XiaoYan Zhu. 2020c. [Recent advances and challenges in task-oriented dialog systems](#). *Science China Technological Sciences*, 63(10):2011–2027.
- Tiancheng Zhao, Allen Lu, Kyusong Lee, and Maxine Eskenazi. 2017. Generative encoder-decoder models for task-oriented spoken dialog systems with chatting capability. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 27–36.
- Qi Zhu, Kaili Huang, Zheng Zhang, Xiaoyan Zhu, and Minlie Huang. 2020. Crosswoz: A large-scale chi-nese cross-domain task-oriented dialogue dataset. *Transactions of the Association for Computational Linguistics*, 8:281–295.

## A Comparison of LLMs over Various NLP Tasks

We show the performance comparison of various LLMs on both NLU and NLG tasks in Figure A1. The data are collected from various prior works focusing on benchmarking the capabilities of LLMs (Bang et al., 2023; Cahyawijaya et al., 2023a; Anonymous, 2023; Asai et al., 2023; OpenAI, 2023; Wu et al., 2023).

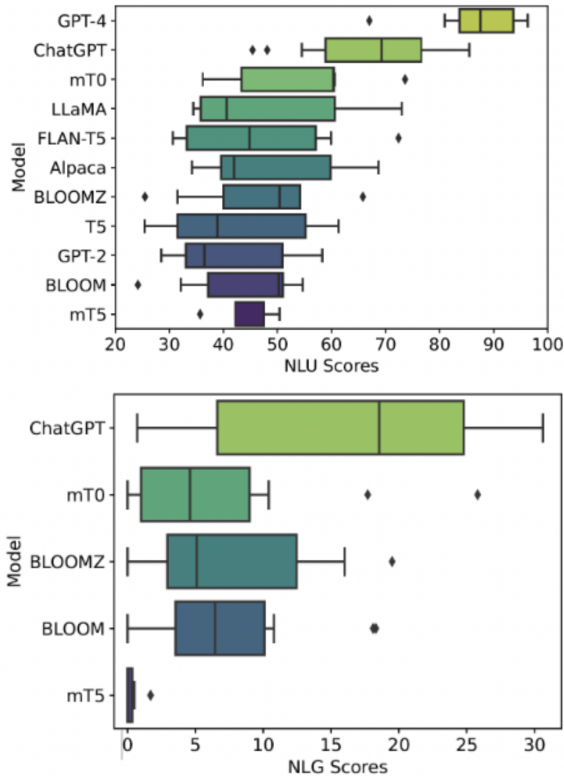


Figure A1: Average LLM performance comparison on various (left) NLG and (right) NLU tasks.

## B Human Evaluation

We give additional details concerning the human evaluation in this section. The instructions for each metric given to the evaluators are defined as follow:

**Informativeness** Amount of information that the system provides while answering the user’s utterance.

1. The response has no information at all
2. The response provides at least one piece of information, but clearly not enough.
3. The response provides several pieces of information but more could be provided

4. The response gives all the information you would expect in that turn

**Helpfulness** The system answers the user’s utterance and pushes the conversation towards completion

1. The response is doing neither
2. The response is just pushing the conversation towards completion without answering the question
3. The response is just answering the question but not pushing the conversation towards completion
4. The response is doing both

**Humanness** The system answers the user’s utterance in a human-like manner

1. The response is completely machine-like
2. The response contains human-like acknowledgement, but still sound machine-like
3. The response has human acknowledgement and you could imagine it being said by a human
4. The response is completely indiscernible from a human answer

## C Generation Samples

We show a few generation samples of InstructTODS as well as the modular evaluation of TODS subtasks. In table A1, we show different delexicalized responses of the fine-tuned baseline compared to the generation by InstructTODS. For the modular subtasks, we show a sample of correct prediction as well as difference failure cases. The samples for dialogue state tracking are in Table A2, for intent classification in Table A3 and for response generation in Table A4.

## D Prompts

In table A5, we give the templates and samples of the prompts used in  $RG_{naive}$ , and in the end-to-end setting for the proxy belief state, the initial knowledge base interaction and the response generation using the extracted information.



ID	SNG0991.json
Dialogue Context	Alright, let's try for a cheap hotel.
Gold	Do you want a cheap hotel or a cheap guesthouse?
SOLIST	I have [value_count] [value_pricerange] hotels. [value_count] in the [value_area] and [value_count] in the [value_area]. Which area do you prefer ?
UBAR	What area would you like to stay in?
AUGPT	I have [value_count] [value_price range] [value_type] available. Do you have a preference for the area?
GALAXY	[value_name] meets your criteria. Would you like me to book you a room ?
PPTOD	There are [value_choice] [value_price] [value_type] with free parking. Do you have a preference for the area ?
InstrucTODS (Zero-Shot)	Great, I found [value_name] located at [value_address] [value_postcode]. It has free parking and falls within the [value_pricerange] price range. Would you like me to confirm a booking for you?
ID	MUL2320.json
Dialogue Context	No, that's everything I needed today. Thank you!
Gold	Ok, thank you for calling. Goodbye.
SOLIST	You are welcome. Have a great day !
UBAR	Thank you for using our system!
AUGPT	You're welcome. Have a great day!
GALAXY	You are welcome. Have a great day !
PPTOD	You are welcome. Have a great day !
InstrucTODS (Zero-Shot)	Great, let me know if you need any further assistance. Have a great day!
ID	MUL1489.json
Dialogue Context	USER: Yes, please. I need a table for 8 on Friday at 17:30, please. SYSTEM: Ok, you are all set for cote on Friday, table for 8 at 17:30. Can I help with anything else? USER: Can I have the reference number for the reservation please? SYSTEM: Booking was unsuccessful. Can you try another time slot? USER: What about 16:30?
Gold	One moment while I try to make the reservation of table for 8, Friday at 16:30.
SOLIST	I have made that reservation for you. Your reference number is [value_reference]. Is there anything else I can help you with ?
UBAR	I was able to book your table for [value_day] at [value_time]. Your reference number is [value_reference].
AUGPT	Yes, the reference number is [value_reference].
GALAXY	I am sorry, but I was unable to book that for you. Would you like to find another restaurant?
PPTOD	Your reference number is [value_reference]. Is there anything else I can help you with ?
InstrucTODS (Zero-Shot)	Sure your reference number for the reservation at [value_name] on [value_day] at [value_time] for a table of [value_people] is [value_reference]. Is there anything else I can assist you with?

Table A1: Samples of different system responses for both the proposed E2E framework in comparison with the gold response. InstrucTOD's response are delexicalized after generation, while the fine-tuned systems generate the delexicalized response directly.

ID	SNG01679.json
Prompt	<p>Generate the dialogue state of the following dialogue between a USER and a task-oriented dialogue SYSTEM. The results should be in a single python dictionary following this format: "domain1-slot1":value1, "domain2-slot2":"value2". Use the provided domain and slots, and nothing else:</p> <p>SLOTS: attraction-area, attraction-name, [...], hotel-day, hotel-stars, [...], train-leaveAt</p> <p>CONTEXT: USER: I need a taxi to take me to Pipasha Restaurant to leave after 01:45.</p> <p>BELIEF STATES:</p>
Gold	{ 'taxi-leaveAt': '01:45', 'taxi-destination': 'pipasha restaurant' }
Modular (GPT-4)	{ "taxi-leaveAt": "01:45", "taxi-destination": "Pipasha Restaurant" }
ID	SNG01936.json
Prompt	<p>Generate the dialogue state of the following dialogue between a USER and a task-oriented dialogue SYSTEM. The results should be in a single python dictionary following this format: "domain1-slot1":value1, "domain2-slot2":"value2". Use the provided domain and slots, and nothing else:</p> <p>SLOTS: attraction-area, attraction-name, [...], hotel-day, hotel-stars, [...], train-leaveAt</p> <p>CONTEXT: USER: I am looking for city centre north b and b SYSTEM: I have found the guesthouse you were wanting. Would you like me to book this for you? USER: Yes, please book it for 1 person and for 5 nights starting Friday.</p> <p>BELIEF STATES:</p>
Gold	{ 'hotel-stay': '5', 'hotel-day': 'friday', 'hotel-people': '1', 'hotel-name': 'city centre north b and b' }
Modular (GPT-4)	{ "hotel-stay": "5", "hotel-day": "Friday", "hotel-people": "1", "hotel-area": "city centre north", "hotel-type": "b and b" }
Comment	In this example, the LLM mistakes the first part of the name "city centre north b and b" as "hotel-area" since it contains "centre north", and the end of the name as a "hotel-type" of hotel "b and b".
ID	SNG0263.json
Prompt	<p>Generate the dialogue state of the following dialogue between a USER and a task-oriented dialogue SYSTEM. The results should be in a single python dictionary following this format: "domain1-slot1":value1, "domain2-slot2":"value2". Use the provided domain and slots, and nothing else:</p> <p>SLOTS: attraction-area, attraction-name, [...], hotel-day, hotel-stars, [...], train-leaveAt</p> <p>CONTEXT: USER: I am looking for a train from Cambridge to Birmingham New Street. SYSTEM: The next train leaving from Cambridge for Birmingham New Street departs Friday at 5:01, and will arrive by 7:44. USER: I need a train that departs after 08:30 on Friday. SYSTEM: I have a train leaving Cambridge arriving at Birmingham New Street on Friday at 9:01. Would you like me to book this for you? USER: Yes please, that sounds perfect.</p> <p>BELIEF STATES:</p>
Gold	{ 'train-leaveAt': '08:30', 'train-destination': 'birmingham new street', 'train-day': 'friday', 'train-departure': 'cambridge' }
Modular (GPT-4)	{ "train-leaveAt": "08:30", "train-destination": "Birmingham New Street", "train-day": "Friday", "train-departure": "Cambridge", "train-arriveBy": "9:01" }
Comment	In this example, the LLM mistakenly parse the time given in the system utterance as "train-arriveBy". This type of mistake happens often in domains with temporal information as the dialogue history gets longer

Table A2: Samples of Zero-Shot Dialogue State Tracking. First example is correct, the other two show different type of LLM mistakes.

Prompt	Select the correct intent that is expressed in the given SENTENCE among the list of INTENTS provided. Generate the correct intent and nothing more:  INTENTS: 0 activate_my_card 1 age_limit 2 apple_pay_or_google_pay [...] 74 why_verify_identity 75 wrong_amount_of_cash_received 76 wrong_exchange_rate_for_cash_withdrawal  SENTENCE: "Is there a way to know when my card will arrive?"
Gold	card_arrival
Modular (GPT-4)	11 card_arrival
Prompt	Select the correct intent that is expressed in the given SENTENCE among the list of INTENTS provided. Generate the correct intent and nothing more:  INTENTS: 0 activate_my_card 1 age_limit 2 apple_pay_or_google_pay [...] 74 why_verify_identity 75 wrong_amount_of_cash_received 76 wrong_exchange_rate_for_cash_withdrawal  SENTENCE: "Can I reactivate my <b>lost</b> card that I found this morning in my jacket pocket?"
Gold	card_linking
Modular (GPT-4)	<b>41 lost_or_stolen_card</b>
Comment	The LLM focuses more on "lost" in the user utterance, although the emphasis to understand the correct intent should be put on "reactivate".
Prompt	Select the correct intent that is expressed in the given SENTENCE among the list of INTENTS provided. Generate the correct intent and nothing more:  INTENTS: 0 activate_my_card 1 age_limit 2 apple_pay_or_google_pay [...] 74 why_verify_identity 75 wrong_amount_of_cash_received 76 wrong_exchange_rate_for_cash_withdrawal  SENTENCE: "Why won't my card show up on the app?"
Gold	card_linking
Modular (GPT-4)	<b>card_not_working</b>
Comment	In the utterance context, the main source of confusion is due to the card linking being implicit (by mentioning the app), while the card not working is more explicit (not show up).

Table A3: Samples of Zero-Shot Intent Classification. First example is correct, the other two show different type of LLM mistakes.

ID	PMUL3027.json
Prompt	<p>In a task oriented dialogue setting, generate a SYSTEM response to the USER query in the conversation provided in CONTEXT. You should follow the information provided in ACT to generate this answer. Do not answer with anything other than what is provided in the dialogue act:</p> <p>CONTEXT:  USER: I am staying in the centre of town for the weekend, what is there to do there?  SYSTEM: We have several things to do! Architecture, colleges, museums...What type of attraction are you most interested in?  USER: It doesn't matter but can you recommend one and give me the entrance fee?  ACT:Recommend the user for free, castle galleries.  SYSTEM:</p>
Gold	I recommend castle galleries and it's free to get in!
Modular (GPT-4)	I recommend visiting the castle galleries, which are a great attraction and the best part is that the entrance is free!
ID	PMUL3027.json
Prompt	<p>In a task oriented dialogue setting, generate a SYSTEM response to the USER query in the conversation provided in CONTEXT. You should follow the information provided in ACT to generate this answer. Do not answer with anything other than what is provided in the dialogue act:</p> <p>CONTEXT:  USER: I want to find a moderately priced restaurant.  SYSTEM: I have many options available for you! Is there a certain area or cuisine that interests you?  USER: Yes I would like the restaurant to be located in the center of the attractions.  SYSTEM: There are 21 restaurants available in the centre of town. How about a specific type of cuisine?  USER: i need to know the food type and postcode and it should also have mutiple sports  SYSTEM: I am sorry I do not understand what you just said. Please repeat in a way that makes sense.  USER: get me the food type and the post code  SYSTEM: My favorite it The Copper Kettle at 4 Kings Parade City Centre CB21SJ. It serves British food. Does that interest you?  USER: Absolutely! Thank you! I also need information on the attractions that have multiple sports in town, in the same area as the restaurant please.  SYSTEM: The only multiple sports attraction is located in the east of town. Would you like more information?  USER: No, I want to be in the centre of town. What about architecture attractions?  SYSTEM: All Saints Church is an architectural attraction. It is in the Centre of town and entrance is free!  USER: Can I get their phone number and address please?  SYSTEM: Their phone number is 01223452587. Is there anything else I can assist you with?  USER: Oh, what type of architecture is it?  ACT:Inform the user that the type is architecture.  SYSTEM:</p>
Gold	It is architecture type. Anything else?
Modular (GPT-4)	All Saints Church features <b>Gothic</b> architecture.
Comment	The type of architecture "Gothic" is hallucinated by the LLM here.

Table A4: Samples of Zero-Shot Response Generation using oracle system action. Second example show hallucination.

Setting	Template	Sample
$RG_{naive}$	<p>Generate the answer of the SYSTEM in the following conversation between a USER and a SYSTEM in a task-oriented dialogue setting. You can either request more details to the user that is available in the knowledge base to complete the goal, or simply answer the user’s request. Do not provide multiple choice for the user to choose, just recommend one, and generate nothing other than the SYSTEM reply. Use the following knowledge base to interact with the user and perform {TASK}:</p> <p>{DATABASE}</p> <p>{DIALOGUE_CONTEXT}</p> <p>SYSTEM:</p>	<p>Generate the answer of the SYSTEM in the following conversation between a USER and a SYSTEM in a task-oriented dialogue setting. You can either request more details to the user that is available in the knowledge base to complete the goal, or simply answer the user’s request. Do not provide multiple choice for the user to choose, just recommend one, and generate nothing other than the SYSTEM reply. Use the following knowledge base to interact with the user and perform restaurant booking:</p> <p>name: ..., address: ..., food: ..., name: ..., address:..., food: ...,</p> <p>USER: Hi, any indian restaurants here? SYSTEM: Yes, we have Indian Express in the intermediate pricerange, would you want to know more? USER: Do you have any restaurant in the cheap pricerange? SYSTEM:</p>
Proxy BS	<p>Suppose you have access to a database with all necessary INFORMATION, what do you need to query to the database in order to reply to the user in the following conversation?</p> <p>INFORMATION: {SLOTS}</p> <p>You can follow these examples: {EXAMPLES}</p> <p>{DIALOGUE_CONTEXT}</p> <p>Need:</p>	<p>Suppose you have access to a database with all necessary INFORMATION, what do you need to query to the database in order to reply to the user in the following conversation?</p> <p>INFORMATION: address, area, name, phone, postcode, pricerange, entrance fee, food, internet, parking, stars</p> <p>You can follow these examples: USER: I need fruits. SYSTEM: Do you have any preferences? USER: Yes, apples if possible. How expensive are they? Need: Information about pricerange for apple</p> <p>USER: Hi, any indian restaurants here? SYSTEM: Yes, we have Indian Express in the intermediate pricerange, would you want to know more? USER: Do you have any restaurant in the cheap pricerange? Need:</p>
Initial KB Interaction	<p>If there are multiple options fitting this criteria, pick a few to propose: {Proxy BS}</p>	<p>If there are multiple options fitting this criteria, pick a few to propose: Information about cheap indian restaurant</p>
E2E RG	<p>In a task-oriented dialogue setting, generate a natural and helpful SYSTEM response to the USER query in the conversation provided in CONTEXT. You should follow the information provided in ACT to generate this answer. Do not mention that you are referring to a dataframe and don’t overload the user with too many choices. You can perform {TASK}.</p> <p>You can follow these examples: {EXAMPLES}</p> <p>CONTEXT: {CONTEXT}</p> <p>ACT: {ACT}</p> <p>SYSTEM:</p>	<p>In a task-oriented dialogue setting, generate a natural and helpful SYSTEM response to the USER query in the conversation provided in CONTEXT. You should follow the information provided in ACT to generate this answer. Do not mention that you are referring to a dataframe and don’t overload the user with too many choices. You can perform restaurant booking.</p> <p>You can follow these examples: USER: I need a place to fish. SYSTEM: Any preference in the type of fish? USER: Preferably salmons, but sardines are also fine. ACT: Blue Lake, 37th Avenue SYSTEM: How about in blue lake, 37th avenue?</p> <p>CONTEXT: USER: Hi, any indian restaurants here? SYSTEM: Yes, we have Indian Express in the intermediate pricerange, would you want to know more? USER: Do you have any restaurant in the cheap pricerange? ACT: Royal Naan SYSTEM:</p>

Table A5: Templates and samples for the prompt used in the  $RG_{naive}$ , Proxy BS, initial KB interaction and E2E RG settings.