

# Reliable LLM-based User Simulator for Task-Oriented Dialogue Systems

Ivan Sekulić, Silvia Terragni, Victor Guimarães, Nghia Khau,  
Bruna Guedes, Modestas Filipavicius, André Ferreira Manso, Roland Mathis

Telepathy Labs GmbH  
Zürich, Switzerland

firstname.lastname@telepathy.ai

## Abstract

In the realm of dialogue systems, user simulation techniques have emerged as a game-changer, redefining the evaluation and enhancement of task-oriented dialogue (TOD) systems. These methods are crucial for replicating real user interactions, enabling applications like synthetic data augmentation, error detection, and robust evaluation. However, existing approaches often rely on rigid rule-based methods or on annotated data.

This paper introduces *DAUS*, a Domain-Aware User Simulator. Leveraging large language models, we fine-tune *DAUS* on real examples of task-oriented dialogues. Results on two relevant benchmarks showcase significant improvements in terms of user goal fulfillment. Notably, we have observed that fine-tuning enhances the simulator’s coherence with user goals, effectively mitigating hallucinations – a major source of inconsistencies in simulator responses.

## 1 Introduction

The field of dialogue systems has seen a notable surge in the utilization of user simulation approaches, primarily for the evaluation and enhancement of conversational search systems (Owoicho et al., 2023) and task-oriented dialogue (TOD) systems (Terragni et al., 2023). User simulation plays a pivotal role in replicating the nuanced interactions of real users with these systems, enabling a wide range of applications such as synthetic data augmentation, error detection, and evaluation (Wan et al., 2022; Sekulić et al., 2022; Li et al., 2022; Balog and Zhai, 2023; Ji et al., 2022).

The significance of user simulation in the development and evaluation of dialogue systems is undeniable. However, the prevailing methodologies often rely on rudimentary rule- and template-based approaches, which can limit their adaptability and effectiveness (Schatzmann et al., 2007; Schatzmann and Young, 2009). Furthermore, certain user

simulation methods require a substantial amount of annotated data (Lin et al., 2021, 2022, 2023), or a deep understanding of the internal workings of the dialogue system they interact with (Schatzmann et al., 2007; Li et al., 2016).

The rise of generative capabilities of large language models (LLMs) enabled user simulators to generate contextually appropriate responses in natural language, without the need for predefined rules (Terragni et al., 2023; Davidson et al., 2023). This shift offers distinct advantages over traditional approaches: i) no human effort is needed to construct the rules; ii) it introduces lexical diversity into utterance generation to assess the robustness of downstream natural language understanding and enables testing of system’s robustness to different dialogue paths. However, LLMs are susceptible to hallucinations (Ji et al., 2023; Terragni et al., 2023), resulting in inconsistency across dialogue turns or the generation of irrelevant information to the user’s goal.

In this paper, we introduce *DAUS*, a generative user simulator for TOD systems. As depicted in Figure 1, once initialized with the user goal description, *DAUS* engages with the system across multiple turns, providing information to fulfill the user’s objectives. Our aim is to minimize the commonly observed user simulator hallucinations and incorrect responses (right-hand side of Figure 1), with an ultimate objective of enabling detection of common errors in TOD systems (left-hand side of Figure 1). Our approach is straightforward yet effective: we build upon the foundation of LLM-based simulators (Terragni et al., 2023; Owoicho et al., 2023) and extend such approach by fine-tuning the LLM on in-domain dialogues, annotated with their user goals. Notably, *DAUS* does not require insights into the inner-workings of the TOD system, its policy, nor system-specific functionalities, as it interacts with the TOD system strictly through natural language.

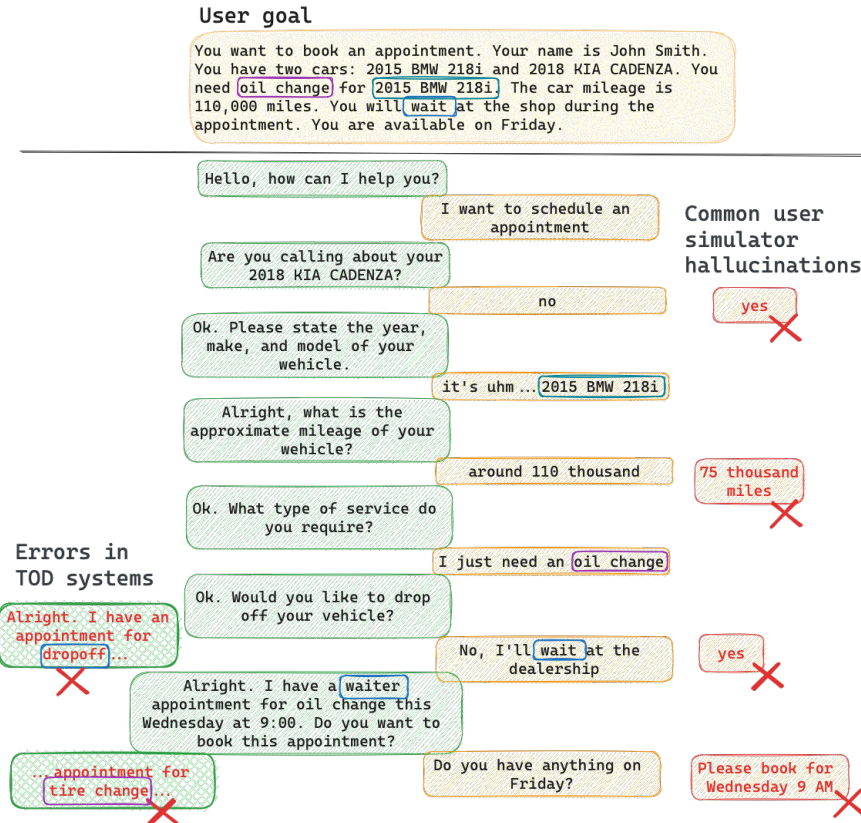


Figure 1: Example conversation between user simulator and TOD system. We aim to minimize common simulator’s hallucinations (right) and thus ease the detection of TOD system failures (left).

We summarize our contributions and findings as follows:

- **Domain-Specific Adaptation:** *DAUS* fine-tunes a pre-trained LLM on domain-specific conversational data, enhancing the simulator’s ability to maintain coherent and contextually relevant dialogues in a specific domain.
- **Reducing Simulator Hallucinations:** *DAUS* mitigates hallucinations originated from in-context learning approaches, which caused inconsistencies and irrelevant information in simulator responses. By fine-tuning on domain-specific data, our approach ensures more coherent and contextually relevant simulated dialogues.
- **Balancing Lexical Diversity in User Simulation:** *DAUS* employs LLMs for user simulation, offering a degree of lexical diversity in generated utterances. While not matching the diversity of in-context learning (partly due to hallucinations), it still provides language variety.

## 2 Related Work

### 2.1 Task-Oriented Dialogue Systems

The field of TOD systems, dedicated to interacting with users to accomplish specific tasks, has recently witnessed notable advancements (Zhang et al., 2020). Given the achievements of LLMs in various natural language processing tasks, there have been efforts to apply them to TOD systems (Raffel et al., 2020; Ouyang et al., 2022). A prominent application involves leveraging LLMs to extract users intents and entities, enhancing the Natural Language Understanding or Dialog State Tracking components (Zhao et al., 2022; Gupta et al., 2022b; Madotto et al., 2021; Madotto and Liu, 2020).

Furthermore, Hudeček and Dušek (2023) suggest that LLMs have the potential to be used off-the-shelf in TOD systems, even without fine-tuning for the specific TOD task, but their performance still lags behind supervised approaches. In response, an alternative approach underscores the benefits of fine-tuning specifically for TOD systems (Bang et al., 2023; Hosseini-Asl et al., 2020; Gupta et al., 2022a). This line of research reveals that fine-tune

LLMs can play a crucial role in enhancing the capabilities of TOD systems.

## 2.2 User Simulation

The state of the art in user simulation for TOD systems has evolved significantly in the recent years. Initially, Eckert et al. (1997) proposed the Bigram model, which estimates a user action conditioned on the system actions. Although efficient, this model does not account for the user goal coherence. Rule-based methods like Agenda-based (Schatzmann et al., 2007; Schatzmann and Young, 2009; Keizer et al., 2010) addresses the coherence issue but relies on the manual definition of rules.

Data-driven approaches, leveraging deep learning models (Gür et al., 2018; Asri et al., 2016; Lin et al., 2021, 2022, 2023), overcome the rule-based constraints but require significant computational resources and annotated data. These methods mandate dialog annotation for user goal fulfillment at each turn. In-context learning approaches (Terragni et al., 2023) have recently gained traction, designing prompts using snippets of example dialogues, the user’s goal (expressed in natural language as in Terragni et al. (2023), or structured format as in Davidson et al. (2023)), and the dialog history. While these approaches demand fewer resources than fine-tuning methods and eschew manual annotation, they underscore limitations of LLMs, including hallucinations, repetitions, and incomplete user goal fulfillment.

## 3 Generative User Simulator

In this section, we define the task of generative user simulation for TOD systems. Moreover, we describe our approach, based on fine-tuned LLMs.

### 3.1 Background

When interacting with a TOD system, users aim to fulfill their goal, e.g., book a flight, or cancel their reservation in a restaurant. Therefore, a user simulator ( $U$ ), designed to imitate a real user, interacts with the TOD system ( $S$ ) with a given user goal  $\mathcal{G}$ . Formally, interactions are a sequence of utterances, where the system’s utterances  $s$  and the user’s utterances  $u$  take turns, forming a dialogue history  $\mathcal{H} = [s_1, u_1, \dots, s_t, u_t, \dots, s_N, u_N]$ , with  $s_t$  and  $u_t$  corresponding to system’s and user’s utterance at turn  $t$ , respectively, and  $N$  being the total number of exchanged utterances.

We define the user goal  $\mathcal{G}$  as all the information

the user requires to achieve their aim. An example of user goal is the following: *You want to try an Indian restaurant. The restaurant must be cheap and in the center. Book a table for 2 people at 8PM.* At the end of the dialogue, we expect the user simulator to have fulfilled  $\mathcal{G}$ . While the goal  $\mathcal{G}$  can be represented either in structural format (Davidson et al., 2023) or in natural language (Terragni et al., 2023), in this work we focus on  $\mathcal{G}$  represented in natural language.  $\mathcal{G}$  is usually defined by a domain expert or randomly sampled.

### 3.2 Our Approach

We propose Domain-Aware User Simulator (DAUS), a model that relies on learning the specifics of interactions with a TOD system from conversational data. The data needs to contain the goal  $\mathcal{G}$  and the dialogue history  $\mathcal{H}$ . Typically, such datasets are derived from user conversations with production TOD systems, or created and curated through crowd-sourcing or user studies.

We cast the above-described problem of simulator’s goal fulfillment to an utterance-level generation task. Specifically, the main task of  $U$  is to generate the next utterance  $u_t$  by modeling:

$$u_t = \phi(\mathcal{G}, \mathcal{H}) \quad (1)$$

where  $\phi$  is the function to generate a user utterance. The  $u_t$  needs be aligned with  $\mathcal{G}$  and  $\mathcal{H}$ , i.e., it needs to be faithful towards the given goal, as well as coherent with the dialogue so far.

Given that both  $\mathcal{G}$  and  $\mathcal{H}$  are in natural language, we model  $\phi$  from Eq. 1 with a language modeling-based approach. Specifically, we first construct a prompt to feed an LLM, by combining  $\mathcal{G}$  and  $\mathcal{H}$ . We further employ the LLM to generate the  $u_t$  in auto-regressive fashion:

$$p_{LLM}(\mathbf{u}_t | \mathcal{G}, \mathcal{H}) = \prod_{i=1}^n p_{LLM}(x_t^i | x_t^{<i}, \mathcal{G}, \mathcal{H}) \quad (2)$$

where  $x_t^i$  is the  $i$ -th token of the utterance at turn  $t$ . We break down the dialogue from the data by turn, yielding  $N$  data points for each conversation.

Regarding the interaction between our fine-tuned LLM and a TOD system, we follow the same paradigm from Terragni et al. (2023). DAUS receives a fresh prompt, which comprises the user’s goal for the ongoing dialogue and the cumulative dialogue history. Unlike Terragni et al. (2023), we do not provide any example dialogues to serve as shots. We additionally post-process the generated

utterance to ensure that a clean message is passed to TOD systems (i.e., removal of special characters and trailing tokens).

## 4 Experimental Setting

In this section, we describe datasets, implementation details, and experimental setting for simulator-system interactions.

### 4.1 Data sources

Table 1: Dataset Statistics (after pre-processing).

Dataset	Avg # Turns	Avg # Words per User Utterance	Avg # Words per TOD Utterance
MultiWOZ	5.86	13.13	14.86
<i>AutomotiveData</i>	11.20	3.44	12.06

We consider two data sources to evaluate our approach. First, we experiment on internal dialogue data of user-TOD system phone call interactions within the automotive industry, dubbed *AutomotiveData*.<sup>1</sup> Second, we use the well-established dataset of multi-domain TOD systems – MultiWOZ 2.1 (Eric et al., 2019). Both data sources contain user goal  $\mathcal{G}$  in natural language and multi-turn dialogues (compliant with Section 3.1). For each dataset, we randomly sample 2,500 dialogues for training, 300 for testing and 300 for validation. The statistics of the resulting datasets are reported in Table 1.

### 4.2 TOD Systems

*DAUS* communicates with TOD systems through natural language, making it system-agnostic. For our user simulator fine-tuned on *AutomotiveData*, we employ an internal TOD system. To evaluate *DAUS* fine-tuned on MultiWOZ, we use the ConvLab2 framework (Zhu et al., 2020), extended by Terragni et al. (2023), which integrates LLM-based few-shot user simulators.<sup>2</sup> We use the same TOD the authors used in their original work. We identify a challenge with the default stopping criteria that prematurely end dialogues when users express gratitude. This does not always signify

<sup>1</sup>In order to protect our users’ privacy, we do not release any user data nor models fine-tuned on user’s data. Examples presented throughout the paper are synthetically constructed, whilst preserving realistic user goals. Users have been informed about and have consented to data collection.

<sup>2</sup><https://github.com/telepathylabsai/prompt-based-user-simulator>.

the end of the interaction as users may continue with their goals (e.g., “*Thanks for booking my flight. I also need a hotel*” would terminate the conversation). Therefore, we modify the criteria to exclude termination on “*thanks*” intent. We consequently re-run the experiments presented in Terragni et al. (2023). Moreover, we publicly release the updated framework and the user simulator fine-tuned on MultiWOZ 2.1 at <https://github.com/telepathylabsai/finetuned-user-simulator>.

### 4.3 User Goal Settings

For the MultiWOZ data within ConvLab2 framework, we follow the previous work for construction of the user goals (Zhu et al., 2020; Terragni et al., 2023). Specifically, the user goals are randomly sampled, conditioned on the domains and entities frequency in the training data. We generate 100 dialogues per user goal.

For evaluation on our internal TOD system, a domain expert manually defined user goals for 8 test cases, detailed in Appendix A. The test cases vary depending on the complexity and the main task that the simulator has to fulfill. As such, we label the test cases accordingly: *B* for *book* appointment task, *C* for *cancel* appointment task, *R* for *reschedule* appointment task. Moreover, each label is associated with a graded difficulty indicator, i.e., *easy* or *hard*. We generate 100 dialogues per test case (i.e., per user goal).

### 4.4 Fine-tuning Details

We conduct our experiments with the recently released open-source LLM — Llama-2 (Touvron et al., 2023). The prompt, mentioned in Section 3.2, is constructed by concatenating the task description, user goal  $\mathcal{G}$ , and the dialog history  $\mathcal{H}$ . Moreover, we separate every utterance with a special “<end-turn>” token.

We utilize LoRA (Hu et al., 2021) – a parameter-efficient fine-tuning technique, capable of reaching performances comparable to fully fine-tuned models, whilst requiring only a fraction of the computational resources. We adhere to the hyperparameter recommendations and instructions of the recent work on the topic (Hu et al., 2021; He et al., 2021) and use the following LoRA hyperparameters throughout the experiments: rank  $r$  of 64,  $\alpha$  of 32, and dropout of 0.05. Moreover, we optimize attention layers (query and key matrices) of the Llama-2 model. We use the 13B Llama-2 ver-

sion for the main experiments, and the 7B version for comparison and the generalization study. We perform hyperparameter grid search for learning rate on the dev sets of our datasets. We settle for  $lr = 3e^{-5}$  and the batch size of 12 and 32 for the 13B and 7B versions, respectively.

#### 4.5 Baselines

We compare our Llama-2 fine-tuned model with several pre-trained models in zero-shot or few-shot fashion, following (Terragni et al., 2023; Davidson et al., 2023). In particular, we consider the following pre-trained models:

- Llama 2 with 13B parameters.
- GPT-3.5 Turbo4 (Chat-GPT), version 0613 (Brown et al., 2020a). For data privacy reasons, we employ this model only for the MultiWOZ experiments.
- Flan-T5 (Chung et al., 2022) with 3B parameters (XL), to reproduce results of Terragni et al. (2023).

In addition to the LLM-based models, we consider an agenda-based simulator (ABUS) (Wen et al., 2015), designed specifically for MultiWOZ within ConvLab2 framework, thus requiring the knowledge of TOD system’s policy. We include two variants of ABUS: the first with template-based NLG and the second with data-driven NLG, dubbed ABUS-T and ABUS-D, respectively. Let us notice that ABUS is a strong baseline, as it is tailored for communicating with the MultiWOZ-based TOD from ConvLab2, therefore it is included as a reference of the potential upper-bound for user goal fulfillment performance. We follow Terragni et al. (2023) and set the temperature for inference to 0.9 for all MultiWOZ experiments, and 0.7 for internal experiments (value chosen through grid search).

## 5 Evaluation

We comprehensively evaluate our method, aiming to assess its ability to achieve designated user goals in dialogues and its impact on lexical diversity when aligning with real user language patterns. Moreover, we perform qualitative analysis of simulated dialogues via human evaluation. In this section, we detail these evaluation procedures.

Additionally, we examine utterance-level metrics, comparing generated utterances with those in the target dataset, using both general natural language generation and domain-specific entity-based

metrics. However, we found that these metrics poorly correlate with the simulator’s task completion. Detailed information about these metrics and their results can be found in Appendix C.

### 5.1 Goal Fulfillment Evaluation Metrics

Our objective is to evaluate the goal fulfillment at the end of the dialogue. For MultiWOZ experiments, we consider well-known metrics such as Success, Completion and Book rate. These metrics aim to capture how successful was the dialogue in terms of fulfilling specific subtasks from the user goal (e.g., whether the restaurant is booked). We also compute the average precision ( $P$ ), recall ( $R$ ) and  $F_1$  scores by matching the entities expressed through the simulated dialogue to the ones in the initial user goal. These metrics aim to assess the simulator’s faithfulness and consistency of entities with the user goal (e.g., whether the correct restaurant *type* was booked). For a comprehensive understanding of the metric definitions, please refer to Zhu et al. (2020) and Terragni et al. (2023).

Regarding our in-house TOD, it is worth noting that we do not differentiate between *book*, *inform* and *request* entities. Therefore, we adapt the mentioned metrics, except for the Book Rate, while considering all entities as *inform* entities. Moreover, we compute several metrics specific to automotive domain: *user subtask* indicating whether the subtask (*book*, *cancel*, or *reschedule* the appointment) matches the one given in the user goal; *caller info* and *car info* indicating whether user information (name, phone number) and vehicle information (car year, make, and model) match the ones in the goal, respectively; *transport type* assessing the chosen transport type (e.g., dropping of the vehicle, waiting for the service in the dealership).

### 5.2 Lexical Diversity of Generated Utterances

Lexical diversity (LD) is a measure of word variability and vocabulary size of a given text corpus, in our case, the set of generated user utterances from 100 conversations. We report MTLN scores (McCarthy, 2005), and a number of unigram words (Unig) and average user utterance length (UttLen). LD results are reported in Section 6.2.

### 5.3 Qualitative Analysis

During the analysis of the generated simulated dialogues, we observed several re-occurring issues. We categorize them as the simulator’s failure (*hallucination*, *incomplete user goal* fulfillment, or *loop-*

Table 2: Results of goal fulfillment task in simulator interaction with the internal TOD system. The results are averaged across the eight user goals.

Model	Num Shots	Compl Rate	Succ Rate	$P$	$R$	$F_1$	User Subtask	Caller Info	Car Info	Transport Type	UttLen	Unig	MTLD
FlanT5-XL (Terragni et al., 2023)	2	0.46	0.27	0.72	0.86	0.76	70.9	85.5	65.6	39.2	<b>2.8</b>	<b>209</b>	<b>23.4</b>
Llama-2	0	0.35	0.13	0.62	0.87	0.69	50.4	88.8	72.2	12.8	2.4	161	15.5
	1	0.37	0.12	0.67	0.89	0.74	65.6	89.1	81.6	8.0	2.0	149	14.5
	2	0.36	0.15	0.66	0.91	0.74	68.9	90.3	80.2	8.0	2.0	129	13.7
<i>DAUS</i>	0	<b>0.51</b>	<b>0.40</b>	<b>0.91</b>	<b>0.92</b>	<b>0.91</b>	<b>99.5</b>	<b>98.5</b>	<b>99.0</b>	<b>80.7</b>	1.7	112	16.5

ing/repeating utterances across turns) or TOD system’s failure (*NLU misclassification* due to missing user’s intent or entities, *forcing end of dialogue*, or *looping/repeating* utterances). Our aim is to assess the prevalence of these patterns and identify potential limitations of LLM-based user simulators. To this end, we employ three annotators to annotate 45 dialogues generated with an LLM-based baseline and 45 dialogues generated with *DAUS* within ConvLab2 framework. The annotators are domain-experts and employees of the authors’ institution. We provide guidelines for each of the categories and go through an on-boarding process with the annotators. The labels for each of the dialogues are determined by majority vote. Annotators reach moderate to good agreement, as measured by Fleiss’  $\kappa$ , detailed in Appendix D.

## 6 Results

In this section, we examine our study’s findings across three main threads. First, we investigate the impact of fine-tuning LLMs with domain-specific data on goal fulfillment in dialog interactions (Section 6.1). Next, we explore the link between fine-tuning and the lexical diversity of generated utterances (Section 6.2). Finally, we assess whether the adaptability of LLM-based user simulators to unseen user tasks is influenced by the diversity of subtask types in their training data (Section 6.3).

### 6.1 Goal Fulfillment

**Internal TOD System.** Table 2 shows results on the goal fulfillment task of *DAUS* and the baselines detailed in Section 4.5, averaged across different user goals. We present the results per each of the eight specific user goals, detailed in Section 4.3, in Appendix B for space-saving purposes.

As a first remark, *DAUS* outperforms all of the baselines across all the goal fulfillment metrics. We observe the largest improvements for domain-

specific metrics, e.g., precision and recall of relevant entities and accuracy of the transport type. This indicates that fine-tuning on in-domain data improves simulator’s knowledge of the domain-specific terminology. We further expand on this observation in Section 7.1.

Regarding the baselines, FlanT5, employing 2 shots as examples, is the second best model. As observed in (Terragni et al., 2023) as well, this instruction fine-tuned model outperforms Llama-2 with 2 shots in most of the cases.

**MultiWOZ Data within ConvLab2.** We show the goal fulfillment performance of *DAUS* and the baselines in interaction with ConvLab2’s TOD system on MultiWOZ 2.1 in Table 3. As in Section 6.1, we observe strong performance of *DAUS*. Specifically, *DAUS* outperforms all of the in-context learning approaches in terms of goal fulfillment, including prior state-of-the-art (Terragni et al., 2023). Moreover, our method outperforms few-shot GPT-3.5, a model significantly larger than ours (estimated 175 billion parameters vs 13 billion). This further suggests the benefits of fine-tuning LLMs on domain-specific conversational data, as stronger performance can be achieved with significantly smaller LLMs, thus reducing the computational requirements of the simulator.

As a general remark, results on both benchmarks, i.e., the ConvLab2 and our internal one, show significant improvements across multiple goal fulfillment metrics. Thus, we conclude that *DAUS* indeed does lead to more consistent, reliable, and faithful LLM-based user simulators. We will discuss these results more in depth in our qualitative analysis in Section 7.1.

### 6.2 Lexical Diversity

Lexical diversity (LD) of generated user utterances from internal TOD system and MultiWOZ experiments is presented in the last 3 columns of Tables 2

Table 3: Performance on MultiWOZ 2.1 within ConvLab2 framework.

Model	Num Shots	Compl Rate	Succ Rate	Book Rate	$P$	$R$	$F_1$	UttLen	Unig	MTLD
ABUS-T (Wen et al., 2015)	-	0.93	0.83	0.85	0.84	0.94	0.86	<b>17.4</b>	527	46.9
ABUS-D (Wen et al., 2015)	-	0.86	0.60	0.75	0.87	0.90	0.87	9.8	327	28.0
FlanT5-XL (Terragni et al., 2023)	2	0.19	0.13	0.46	0.45	0.39	0.39	13.7	<b>888</b>	41.2
Llama-2	0	0.07	0.04	0.13	0.31	0.21	0.23	8.1	697	30.7
	2	0.09	0.08	0.30	0.46	0.34	0.39	10.0	765	38.8
GPT-3.5	2	0.35	0.19	0.34	0.49	0.52	0.48	16.3	626	38.1
<i>DAUS</i>	0	<b>0.41</b>	<b>0.29</b>	<b>0.66</b>	<b>0.69</b>	<b>0.69</b>	<b>0.67</b>	10.6	789	<b>54.9</b>

and 3. We observe a drop in LD, as measured by the length of the generated utterances and the total number of unigrams, when *DAUS* is fine-tuned on *AutomotiveData*. This suggests a limited vocabulary in the training data, which is expected due to the real users often responding with one or two words, especially in the *cancel* task. *DAUS* had a relatively high MTLT score, because of the correctly generated caller, car and transport entities, which usually have unique values. However, a low unigram score is due to averaging metrics over 8 user tasks, where only 3 of them are the entity-rich *book* task. Meanwhile, the higher LD of FlanT5-based method is due to its prevalent hallucinations, thus falsely inflating the LD scores by generating out-of-context content (see Section 7.1).

In MultiWOZ-based experiments, results indicate higher LD than ABUS baselines, as measured by MTLT, while the generated utterances are slightly shorter compared to FlanT5. As such, *DAUS* does not seem to lose LD during fine-tuning on MultiWOZ, while fine-tuning on *AutomotiveData* seems to reduce it slightly. This can be explained by the fact that *AutomotiveData* contains both specific vocabulary and utterances from real product users, which makes it hard for in-context learning approaches to imitate. On the other hand, fine-tuning procedure enables the model to learn the particulars of such interactions.

### 6.3 Generalization to Unseen User Tasks

Table 4 shows the percentage of successful subtask identifications for four variants of our model: *DAUS* fine-tuned on the full dataset described in Section 4.1, and *DAUS* fine-tuned on modified datasets by removing certain subtasks (*book* ( $B$ ), *cancel* ( $C$ ), or *reschedule* ( $R$ )) from the training sets. With this experiment, we aim to assess the

generalization abilities of our approach.

Table 4: Percentage of dialogues with successfully identified subtask types across the test cases, with models fine-tuned on specific combinations of subtask types.

	<i>DAUS</i> ( $C+R+B$ )	$C+R$	$B+R$	$B+C$
$B_{easy}$	99	100	100	99
$B_{hard1}$	93	29	85	99
$B_{hard2}$	99	86	94	97
$C_{easy}$	96	100	75	99
$C_{hard}$	100	100	77	96
$R_{easy}$	88	100	98	34
$R_{hard1}$	97	50	69	0
$R_{hard2}$	86	84	56	0

Results show a decrease in performance when a model is not shown the specific subtask during training. For example, when we fine-tune *DAUS* on the combination of *book* and *reschedule* subtasks, we observe a considerable drop in performance on the *cancel* subtask. However, the largest drop is observed in the most complex subtask type, *reschedule*, where the model fine-tuned on  $B+C$  data completely fails to successfully communicate its goal for both  $R_{hard}$  test cases.

We can conclude that *DAUS* does not generalize well to unseen user goal subtasks. Nevertheless, the overall performance of the fine-tuned models across all of the subtasks is still comparable to the performance of few-shot based models (e.g.,  $B+C$  correctly predicts the subtask type, on average, in 66% of the dialogues, while Llama-2 2-shot does it in 69% of the dialogues, on average).

## 7 Qualitative Analysis

In this section, we detail and discuss the findings of our qualitative analysis of simulated dialogues.

Table 5: Percentage of the observed patterns per sample annotated in simulated dialogues in MultiWOZ.

Label	<i>FlanT5</i>	<i>DAUS</i>
Hallucination	73%	36%
Looping simulator	69%	6%
Incomplete goal	78%	53%
Looping system	20%	22%
NLU misclassification	60%	40%
Forced end	27%	27%

## 7.1 Human Evaluation of Generated Dialogues

Table 5 presents the prevalence of patterns, described in Section 5.3, observed through manual annotation of the simulated MultiWOZ dialogues. We observe consistent decrease in hallucinations, reduced number of dialogues with incomplete goal fulfillment, as well as reduced repetition of utterances in dialogues generated by *DAUS*, compared to *FlanT5*-based simulator. Below, we report the main findings from our analysis.

**Hallucinations.** The percentage of dialogues containing hallucinations drops from 73% for *FlanT5*-based simulator to 36% with *DAUS*. We observe that *FlanT5* frequently experiences severe failures, mostly because it generates non-specified pieces of information, such as defining a random range of time for a taxi pickup, inventing a location for an attraction or referring to a restaurant that have not been previously mentioned. Such hallucinations lead to dialogue failures, without possibility to recover the conversation. On the other hand, *DAUS* does not hallucinate nor misinterpret entities from the user goal and the dialogue, but rather sometimes asks for additional information that is not required by the user goal (e.g., asking restaurant’s phone number, even though it is not strictly specified in the goal). Thus, we conclude that not only the prevalence of the hallucinations is reduced, but also their severity.

**TOD system failures.** We observe that the dialogue frequently stops before all of the subtasks from the user goal are completed (i.e., if the goal requires to book a restaurant and a taxi, perhaps only the restaurant is booked). While this is partly due to the simulator missing certain part of the user goal, it often happens due to the TOD system failures. We observe several issues with the TOD system within the ConvLab2 framework. First, as mentioned in the previous paragraph, the system

aims to prematurely complete the dialogue, thus offering no opportunity to the user (simulator) to fully express their goal. Second, the system’s NLU component often misclassifies certain entities, leading to incorrectly booked restaurants or similar. An example dialogue of both prematurely ending the dialogue and NLU misclassification:

*CUSTOMER: Can you help me with a robbery that happened? I am the victim.*

*ASSISTANT: Thank you for using our services . (forcing premature end of conversation)*

*CUSTOMER: It happened at 9:10*

*ASSISTANT: Can you tell me which day you 'd like to travel on? What is your destination? Where will you be departing from? (NLU misclassification)*

Finally, we observe loops in system’s policy, where the system is stuck in policy states and is not able to continue the conversation.

## 7.2 Domain-specific Vocabulary

During our analysis of conversations generated using our internal TOD system, we observed that *DAUS* acquires domain-specific vocabulary, which we believe enhances its performance. For instance, when scheduling a vehicle service appointment, customers specify the preferred transport type, usually *dropoff* (the customer drops off the vehicle and returns when the service is complete) or *waiter* (the customer waits at the dealership during the service). This phenomenon, supported by performance improvements for *Transport Type* (see Table 2), is evident through qualitative dialogue analysis. Our initial experiments indicate that in-context learning approaches aiming to explain these terms in the prompt do not consistently capture their nuances.

Additionally, we noticed that, when fine-tuned on *AutomotiveData* containing phone call conversations with real users, *DAUS* tends to generate filler words like “uhm” and “yeah”.

## 8 Conclusions

The use of a domain-aware LLM-based user simulator, such as *DAUS*, shows promising results in multi-turn interactions with TOD systems. *DAUS* can fulfill user goals by generating consistent and faithful utterances. Compared to previous LLM-based approaches (Terragni et al., 2023), our method has demonstrated superior performance, as measured by multiple metrics designed to capture the fulfillment of the given goal, as well as



faithfulness across the dialogue. This indicates that *DAUS* is capable of effectively simulating user behavior and can serve as a valuable tool for testing and evaluating TOD systems. Moreover, our approach requires relatively small training dataset and imposes modest computational demands, thanks to parameter-efficient fine-tuning. This discovery aligns with findings in related research that contrasts in-context learning with parameter-efficient fine-tuning (Mosbach et al., 2023; Liu et al., 2022). Consequently, our approach emerges as a pragmatic choice for broader adoption within the NLP and Conversational AI community.

The potential applications of LLM-based user simulators are synthetic data augmentation (Li et al., 2022), supporting reinforcement learning approaches (Shi et al., 2019), and TOD system evaluation (Terragni et al., 2023; Zhu et al., 2020). *DAUS*'s reliability and consistency to the user goal make it particularly suitable for TOD system evaluation. As we have seen previously, an incomplete user goal can mainly imply two scenarios: a user simulator who hallucinates or a TOD system that is not able to understand the user's requirements. Therefore, the presence of a reliable user simulator is crucial: it allows us to identify the TOD system's errors with high accuracy.

Moreover, we stress that at the center of our approach is an LLM, leading to potentially different generations given the same input, depending on the sampling method. This means that *DAUS* is more flexible than certain agenda-based simulators, which usually rely on template-based responses. As such, we are able to simulate a dialogue with the same user goal multiple times, which results in multiple different attempts of the simulator to fulfill its goal, going through potentially different conversational paths. Therefore, we are able to test the robustness of the TOD system to different expressions of the same user goal.

## 9 Limitations

The approach employed in our study has several inherent limitations, primarily stemming from the use of LLMs. Most notably, GPT-3.5, the model we utilized in our experiments, is not open-source and freely available, which can hinder replicability of the experiments. Another limitation is related to the opaqueness of the model's training and fine-tuning processes. These models undergo pre-training and fine-tuning on diverse datasets, the

specifics of which are often undisclosed. Consequently, it is challenging to ascertain whether these models have been exposed to specific datasets, such as MultiWOZ 2.1, or datasets with similar characteristics, which could raise concerns about models performance and potential biases.

Furthermore, our experiments were conducted exclusively on two English-language datasets. While LLMs are known for their transfer learning capabilities, allowing for the potential extension of results to other datasets, there is no guarantee of their generalizability across various domains or low-resource languages. The effectiveness of these models in domains distinct from the ones they were trained on remains uncertain and should be approached with caution.

In our analysis, we also observed instances where LLMs exhibit hallucinations. Despite being superior to in-context learning approaches like (Terragni et al., 2023), we still encountered cases of LLM responses that deviated from the expected or coherent output. These hallucinations may lead to unpredictable and potentially inappropriate responses in certain conversational contexts, raising concerns about the reliability and safety of such systems.

We also noticed a decrease in performance when certain user subtasks are omitted from the training dataset when we fine-tune *DAUS*, although the overall performance remains comparable to that of few-shot models. In our analysis, we did not investigate if providing one or two dialog shots would address this performance decrease.

Finally, the methodology relies on conversational data for fine-tuning LLMs. This reliance introduces additional limitations. Firstly, obtaining suitable conversational data may be challenging or even unfeasible in some scenarios. Researchers may resort to crowd-sourcing tools to gather dialogue examples or use LLMs themselves to generate synthetic data, which could introduce biases or inaccuracies. Secondly, the quality of the conversational data used for fine-tuning plays a pivotal role in the model's performance. In our study, we utilized well-curated conversational data, but we did not investigate the impact of using noisier or less meticulously curated data. The use of lower-quality data sources may affect the model's performance and raise questions about its reliability and robustness in real-world applications.

## 10 Ethics Statement

The use of LLMs for user simulation raises ethical considerations. We acknowledge the potential for perpetuating biases and stereotypes present in the data used to train these models (Brown et al., 2020b; Lucy and Bamman, 2021; Bender et al., 2021). While we have not implemented specific measures to mitigate these risks in this paper, we recognize their importance and urge the research community to address these challenges.

It is essential to note that we have used the user simulator solely to evaluate the performance of a dialogue system. However, LLMs can be used in a reinforcement learning setting to train dialog systems (Shi et al., 2019). In such cases, it is crucial to use these models judiciously because of their unpredictable and potentially inappropriate responses.

In addition to ethical considerations, it is crucial to acknowledge the significant environmental impact of LLMs. Their training and deployment consumes a considerable amount of energy, leading to environmental issues (Strubell et al., 2019). We should also be aware of the significant carbon footprint while fine-tuning the LLMs and using them for inference.

## Acknowledgement

Our gratitude to Damián Pascual for streamlining the implementation of the fine-tuning framework, saving us valuable time. Special thanks to the reviewers, Diana Nicoleta Popa, Vijeta Avijeet, and our colleagues at Telepathy Labs in Zürich for their constructive feedback and insightful discussions.

## References

Layla El Asri, Jing He, and Kaheer Suleman. 2016. A sequence-to-sequence model for user simulation in spoken dialogue systems. In *Interspeech 2016, 17th Annual Conference of the International Speech Communication Association*, pages 1151–1155. ISCA.

Krisztian Balog and ChengXiang Zhai. 2023. User simulation for evaluating information access systems. *arXiv preprint arXiv:2306.08550*.

Namo Bang, Jeehyun Lee, and Myoung-Wan Koo. 2023. Task-optimized adapters for an end-to-end task-oriented dialogue system. *arXiv preprint arXiv:2305.02468*.

Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models

be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623.

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020a. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020b. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- 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*.
- Sam Davidson, Salvatore Romeo, Raphael Shu, James Gung, Arshit Gupta, Saab Mansour, and Yi Zhang. 2023. User simulation with large language models for evaluating task-oriented dialogue. *arXiv preprint arXiv:2309.13233*.
- Wieland Eckert, Esther Levin, and Roberto Pieraccini. 1997. User modeling for spoken dialogue system evaluation. *1997 IEEE Workshop on Automatic Speech Recognition and Understanding Proceedings*, pages 80–87.
- Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyag Gao, and Dilek Hakkani-Tur. 2019. Multiwoz 2.1: Multi-domain dialogue state corrections and state tracking baselines. *arXiv preprint arXiv:1907.01669*.
- Prakhar Gupta, Cathy Jiao, Yi-Ting Yeh, Shikib Mehri, Maxine Eskenazi, and Jeffrey Bigham. 2022a. *InstructDial: Improving zero and few-shot generalization in dialogue through instruction tuning*. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 505–525, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Raghav Gupta, Harrison Lee, Jeffrey Zhao, Yuan Cao, Abhinav Rastogi, and Yonghui Wu. 2022b. *Show, don't tell: Demonstrations outperform descriptions for schema-guided task-oriented dialogue*. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages

- 4541–4549, Seattle, United States. Association for Computational Linguistics.
- Izzeddin Gür, Dilek Hakkani-Tür, Gokhan Tür, and Pararth Shah. 2018. User modeling for task oriented dialogues. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 900–906. IEEE.
- Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. 2021. Towards a unified view of parameter-efficient transfer learning. *arXiv preprint arXiv:2110.04366*.
- Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. 2020. A simple language model for task-oriented dialogue. *Advances in Neural Information Processing Systems*, 33:20179–20191.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Vojtěch Hudeček and Ondřej Dušek. 2023. Are llms all you need for task-oriented dialogue? *arXiv preprint arXiv:2304.06556*.
- Tianbo Ji, Yvette Graham, Gareth Jones, Chenyang Lyu, and Qun Liu. 2022. Achieving reliable human assessment of open-domain dialogue systems. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL’22, pages 6416–6437.
- 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 Comput. Surv.*, 55(12).
- Simon Keizer, Milica Gasic, Filip Jurcicek, François Mairesse, Blaise Thomson, Kai Yu, and Steve Young. 2010. Parameter estimation for agenda-based user simulation. In *Proceedings of the SIGDIAL 2010 Conference*, pages 116–123.
- Alon Lavie and Abhaya Agarwal. 2007. Meteor: An automatic metric for mt evaluation with high levels of correlation with human judgments. In *Proceedings of the Second Workshop on Statistical Machine Translation, StatMT ’07*, page 228–231, USA. Association for Computational Linguistics.
- Xiujun Li, Zachary C Lipton, Bhuwan Dhingra, Lihong Li, Jianfeng Gao, and Yun-Nung Chen. 2016. A user simulator for task-completion dialogues. *arXiv preprint arXiv:1612.05688*.
- Zekun Li, Wenhui Chen, Shiyang Li, Hong Wang, Jing Qian, and Xifeng Yan. 2022. Controllable Dialogue Simulation with In-context Learning. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4330–4347. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Hsien-Chin Lin, Shutong Feng, Christian Geishauer, Nurul Lubis, Carel van Niekerk, Michael Heck, Benjamin Ruppik, Renato Vukovic, and Milica Gasic. 2023. Emous: Simulating user emotions in task-oriented dialogues. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2526–2531.
- Hsien-Chin Lin, Christian Geishauer, Shutong Feng, Nurul Lubis, Carel van Niekerk, Michael Heck, and Milica Gasic. 2022. GenTUS: Simulating User Behaviour and Language in Task-oriented Dialogues with Generative Transformers. In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue, SIGDIAL 2022*, pages 270–282. Association for Computational Linguistics.
- Hsien-Chin Lin, Nurul Lubis, Songbo Hu, Carel van Niekerk, Christian Geishauer, Michael Heck, Shutong Feng, and Milica Gasic. 2021. Domain-independent User Simulation with Transformers for Task-oriented Dialogue Systems. In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue, SIGdial 2021*, pages 445–456. Association for Computational Linguistics.
- Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohata, Tenghao Huang, Mohit Bansal, and Colin Raffel. 2022. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. *Advances in Neural Information Processing Systems*, 35:1950–1965.
- Li Lucy and David Bamman. 2021. Gender and representation bias in gpt-3 generated stories. In *Proceedings of the Third Workshop on Narrative Understanding*, pages 48–55.
- Andrea Madotto, Zhaojiang Lin, Genta Indra Winata, and Pascale Fung. 2021. Few-shot bot: Prompt-based learning for dialogue systems. *arXiv preprint arXiv:2110.08118*.
- Andrea Madotto and Zihan Liu. 2020. Language models as few-shot learner for task-oriented dialogue systems. *ArXiv*, abs/2008.06239.
- 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.
- Marius Mosbach, Tiago Pimentel, Shauli Ravfogel, Dietrich Klakow, and Yanai Elazar. 2023. **Few-shot fine-tuning vs. in-context learning: A fair comparison and evaluation**. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 12284–12314, Toronto, Canada. Association for Computational Linguistics.

- 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](#).
- Paul Owoicho, Ivan Sekulic, Mohammad Aliannejadi, Jeffrey Dalton, and Fabio Crestani. 2023. Exploiting simulated user feedback for conversational search: Ranking, rewriting, and beyond. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 632–642.
- 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.
- 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. *J. Mach. Learn. Res.*, 21(1).
- Jost Schatzmann, Blaise Thomson, Karl Weilhammer, Hui Ye, and Steve Young. 2007. Agenda-based user simulation for bootstrapping a pomdp dialogue system. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Companion Volume, Short Papers*, pages 149–152.
- Jost Schatzmann and Steve J. Young. 2009. The hidden agenda user simulation model. *IEEE Transactions on Audio, Speech, and Language Processing*, 17:733–747.
- Ivan Sekulić, Mohammad Aliannejadi, and Fabio Crestani. 2022. Evaluating mixed-initiative conversational search systems via user simulation. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, pages 888–896.
- Weiyan Shi, Kun Qian, Xuwei Wang, and Zhou Yu. 2019. How to build user simulators to train rl-based dialog systems. 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 1990–2000.
- Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. Energy and policy considerations for deep learning in nlp. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3645–3650.
- Silvia Terragni, Modestas Filipavicius, Nghia Khau, Bruna Guedes, André Manso, and Roland Mathis. 2023. [In-context learning user simulators for task-oriented dialog systems](#). In *Proceedings of the 1st Workshop on Foundations and Applications in Large-scale AI Models -Pre-training, Fine-tuning, and Prompt-based Learning*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Dazhen Wan, Zheng Zhang, Qi Zhu, Lizi Liao, and Minlie Huang. 2022. [A unified dialogue user simulator for few-shot data augmentation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3788–3799, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Tsung-Hsien Wen, Milica Gašić, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015. Semantically conditioned LSTM-based natural language generation for spoken dialogue systems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1711–1721. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Zheng Zhang, Ryuichi Takanobu, Qi Zhu, Minlie Huang, and XiaoYan Zhu. 2020. Recent advances and challenges in task-oriented dialog systems. *Science China Technological Sciences*, 63(10):2011–2027.
- Jeffrey Zhao, Raghav Gupta, Yuan Cao, Dian Yu, Mingqiu Wang, Harrison Lee, Abhinav Rastogi, Izhak Shafran, and Yonghui Wu. 2022. Description-driven task-oriented dialog modeling. *arXiv preprint arXiv:2201.08904*.
- Qi Zhu, Zheng Zhang, Yan Fang, Xiang Li, Ryuichi Takanobu, Jinchao Li, Baolin Peng, Jianfeng Gao, Xiaoyan Zhu, and Minlie Huang. 2020. Convlab-2: An open-source toolkit for building, evaluating, and diagnosing dialogue systems. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.

## A User tasks

Description of eight different test cases (user goals) are provided in Table 6. We additionally add comparisons with FLanT5-XXL.

Table 6: Description of user goals with subtask types.

#	User subtask	Difficulty	User goal details
1	Book	Easy	New customer; Available: today 4PM; Transport_type: waiter; Service: check engine.
2	Book	Hard	Known customer with 1 appointment and 2 cars; Available: Wednesday; Transport_type: dropoff; Unknown Service.
3	Book	Hard	Known customer with 3 appointments and 2 cars; Available: Wednesday; Transport_type: dropoff; Two services: engine over-heating and oil change.
4	Cancel	Easy	Known customer with 1 appointment.
5	Cancel	Hard	Known customer with 3 appointments.
6	Reschedule	Easy	Known customer with 1 appointment; Available: 10 AM; Transport_type: dropoff; Unknown service.
7	Reschedule	Hard	Known customer with 1 appointment; Available: afternoon; Transport_type: waiter; Service: oil change.
8	Reschedule	Hard	Known customer from unknown phone number; With 3 appointments; Available: Wednesday; Transport_type: loaner; Two services: Oil change and engine check

## B Results per Tasks

Table 8 shows the breakdown of the results of baselines and DAUS per specific user goal.

## C Utterance-Level Metrics

In addition to dialogue-level metrics detailed in Section 5, we consider a number of utterance-level metrics. Such metrics are based on comparisons of generated utterances to the target utterance in the test set of the appropriate dataset,

Table 7: Inter-Annotator Agreement, as measured by Fleiss’  $\kappa$  for samples from DAUS and FLanT5-XL.

	DAUS	FLanT5-XL
Hallucination	0.365	0.499
Incomplete Goal	0.585	0.754
Looping Simulator	0.319	0.687
NLU Misclassification	0.356	0.308
Forces end of dialogue	0.314	0.367
Looping System	0.640	0.084

described in Section 4.1. We consider two main types of utterance-level metrics: 1) natural language generation (NLG) metrics; and 2) natural language understanding-based (NLU) metrics. We compute several well-known NLG metrics: BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), BERTScore (Zhang et al., 2019), METEOR (Lavie and Agarwal, 2007), as well as cosine similarity between embedded generated and target utterances.

Moreover, we design several domain-specific NLU-based metrics. TOD systems are composed of multiple modules, with NLU module, that aims to understand and parse the given user utterance, being one of the essential modules. Thus, we employ NLU component of the TOD systems to extract user (simulator) intent and mentioned entities. Similarly to NLG-metrics, we compare the intent and entities extracted from the generated utterance, to those in the target utterance. Specifically, we design the following metrics:

- Cosine similarity between the embedded **intents** extracted from the generated utterance and the target utterance. Intents are embedded with RoBERTa model.
- Cosine similarity between the generated and the target utterance, in which the entities were masked. Utterances are embedded with RoBERTa model.
- Precision, Recall, and  $F_1$  of **entities** between the generated and the target utterances.

Table 9 shows the results across the described metrics.

## D Qualitative Analysis Details

Table 7 shows the Inter-Annotator agreement per model and per type of pattern.

Table 10 shows three examples of failed simulated dialogues.

## **E Computing Infrastructure**

We ran the experiments on a machine equipped with two AMD® EPYC 7763 64-Core Processors, and 10 NVIDIA RTX A6000 GPUs with 48GB RAM each, CUDA v11.6, Driver Version 510.54. All the experiments ran on a single GPU. As detailed earlier, we use Llama-2 (7B and 13B parameters versions), as well as FlanT5 (3B and 11B versions). Fine-tuning of a single Llama-2 model requires approximately 12 GPU hours. We estimate all of the experiments to require several hundred GPU hours.

## **F Use of AI assistants for writing**

ChatGPT was used for rephrasing certain sections of this work to enhance clarity and coherence. It was not involved in generating new content such as tables, citations, or equations. The authors' first language is not English, and the assistance from ChatGPT aimed to improve readability.

Table 8: Results of selected baselines and *DAUS* (the main method based on Llama-2 13B, as well as the 7B version) per specific user goal.

Subtask	Model	N shots	User Task	Compl Rate	Succ Rate	P	R	F1	Service Info	Transport	Car Info	Caller Info
$C_{hard}$	Llama-2-13b	0	43	100	43	0.74	0.79	0.76				99.5
	Llama-2-13b	2	52	100	44	0.77	0.86	0.8				100
	FlanT5-xxl	0	61	100	57	0.83	0.84	0.83				99.5
	FlanT5-xxl	2	65	100	63	0.84	0.9	0.85				99
	FlanT5-xl	0	67	100	64	0.85	0.89	0.86				98
	FlanT5-xl	2	75	100	73	0.89	0.94	0.9				100
	DAUS-7b	0	93	100	93	0.96	0.96	0.96				100
	DAUS	0	100	100	100	1	1	1				100
$B_{hard2}$	Llama-2-13b	0	94	23	1	0.59	0.89	0.67	31	4	73	57.5
	Llama-2-13b	2	98	27	1	0.62	0.87	0.71	44	7	84	64
	FlanT5-xxl	0	81	64	19	0.78	0.91	0.81	74	45	78.3	78
	FlanT5-xxl	2	91	72	15	0.77	0.86	0.8	83	42	86	84
	FlanT5-xl	0	81	18	4	0.37	0.77	0.44	36	18	29.3	22.5
	FlanT5-xl	2	95	58	6	0.66	0.81	0.7	75	37	74.6	68.5
	DAUS-7b	0	99	76	29	0.87	0.84	0.85		64	100	88.5
	DAUS	0	100	89	50	0.93	0.85	0.88	90	93	100	99
$B_{easy}$	Llama-2-13b	0	97	43	23	0.77	0.91	0.82	59	49	77	100
	Llama-2-13b	2	100	43	4	0.76	0.93	0.83	51	15	92.3	100
	FlanT5-xxl	0	90	65	46	0.85	0.93	0.86	70	63	90	99
	FlanT5-xxl	2	98	57	50	0.89	0.88	0.88	60	80	84	100
	FlanT5-xl	0	94	14	14	0.73	0.86	0.78	34	91	40.6	99.5
	FlanT5-xl	2	97	23	22	0.81	0.85	0.82	26	94	54.3	100
	DAUS-7b	0	96	55	22	0.92	0.87	0.89		99	98.7	100
	DAUS	0	100	37	15	0.93	0.89	0.91	38	98	100	98
$B_{hard1}$	Llama-2-13b	0	65	1	0	0.59	0.9	0.69	19	6	64	100
	Llama-2-13b	2	83	0	0	0.62	0.9	0.71	4	1	71.3	100
	FlanT5-xxl	0	80	10	0	0.82	0.84	0.81	16	71	84	99.5
	FlanT5-xxl	2	56	9	0	0.69	0.86	0.73	35	44	69.3	100
	FlanT5-xl	0	40	2	0	0.6	0.84	0.67	25	28	61.6	100
	FlanT5-xl	2	24	1	0	0.48	0.91	0.6	62	10	48.3	100
	DAUS-7b	0	78	2	0	0.81	0.82	0.8		80	86	100
	DAUS	0	99	15	0	0.84	0.84	0.83	17	84	94	95.5
$C_{easy}$	Llama-2-13b	0	39	100	37	0.76	0.78	0.76				100
	Llama-2-13b	2	67	100	61	0.85	0.89	0.86				100
	FlanT5-xxl	0	75	100	74	0.91	0.89	0.89				100
	FlanT5-xxl	2	94	100	93	0.98	0.97	0.98				100
	FlanT5-xl	0	73	100	71	0.89	0.87	0.87				100
	FlanT5-xl	2	97	100	97	0.99	0.99	0.99				100
	DAUS-7b	0	100	100	100	1	1	1				100
	DAUS	0	100	100	100	1	1	1				100
$R_{easy}$	Llama-2-13b	0	14	1	0	0.51	0.91	0.63	2	7	77.6	100
	Llama-2-13b	2	38	2	2	0.54	0.93	0.67	9	6	80	98.5
	FlanT5-xxl	0	60	3	2	0.78	0.91	0.83	16	80	98.3	99.5
	FlanT5-xxl	2	76	28	8	0.84	0.91	0.86	45	82	99.3	100
	FlanT5-xl	0	44	8	3	0.75	0.88	0.8	26	79	92	100
	FlanT5-xl	2	71	30	3	0.81	0.91	0.85	54	70	98.6	100
	DAUS-7b	0	99	10	10	0.97	0.91	0.94		99	100	100
	DAUS	0	99	6	5	0.91	0.93	0.91	9	100	100	100
$R_{hard1}$	Llama-2-13b	0	25	7	0	0.55	0.88	0.66	27	5	79	93.5
	Llama-2-13b	2	55	13	6	0.61	0.93	0.72	42	15	79.6	91
	FlanT5-xxl	0	20	14	5	0.67	0.87	0.75	88	47	86	85.5
	FlanT5-xxl	2	34	29	10	0.68	0.8	0.73	95	31	81.6	80.5
	FlanT5-xl	0	15	12	3	0.43	0.61	0.49	72	16	40.3	41.5
	FlanT5-xl	2	61	53	11	0.71	0.83	0.74	82	21	74.6	74.5
	DAUS-7b	0	48	22	21	0.72	0.87	0.78		77	84.3	96
	DAUS	0	100	62	46	0.9	0.94	0.91	98	99	100	100
$R_{hard2}$	Llama-2-13b	0	26	1	0	0.48	0.89	0.59	20	6	62.3	60
	Llama-2-13b	2	58	1	0	0.53	0.94	0.65	23	4	74	69
	FlanT5-xxl	0	26	7	4	0.67	0.93	0.75	67	50	81	79
	FlanT5-xxl	2	42	8	1	0.64	0.83	0.7	49	33	82.3	68.5
	FlanT5-xl	0	9	0	0	0.21	0.43	0.26	17	5	19.3	13.5
	FlanT5-xl	2	47	2	0	0.42	0.69	0.47	14	3	43	41.5
	DAUS-7b	0	36	6	5	0.67	0.83	0.73		67	98	69.5
	DAUS	0	98	1	0	0.78	0.93	0.84	5	10	100	95.5

Table 9: NLG- and NLU-based utterance-level metrics.

		BLUE	ROUGE	BERTScore	METEOR	utterance similarity	intent similarity	entities_R	entities_P	entities_F <sub>1</sub>
MultiWOZ	Llama-2-7b	0.12	0.18	0.85	0.16	0.31	0.67	0.22	0.35	0.37
	Llama-2-13b	0.13	0.19	0.85	0.16	0.32	0.67	0.26	0.36	0.38
	Llama-2-13b-fine-tuned	0.12	0.19	0.85	0.16	0.31	0.67	0.26	0.35	0.37
<i>AutomotiveData</i>	Llama-2-7b	0.22	0.24	0.88	0.17	0.52	0.58	0.44	0.18	0.47
	Llama-2-13b	0.42	0.43	0.92	0.26	0.68	0.73	0.42	0.25	0.47
	Llama-2-13b-fine-tuned	0.42	0.43	0.92	0.26	0.68	0.73	0.42	0.25	0.47



Table 10: Examples of common failures observed through qualitative analysis.

---

Example 1:  
 REQUIREMENTS: You are looking forward to trying local restaurants. You are looking for a particular restaurant. Its name is called cote. Once you find a restaurant, make sure you get postcode. Once you find the restaurant you want to book a table for 7 people at 12:00 on tuesday. You are looking for a place to stay. The hotel should be in the cheap price range. The hotel should be in the north. The hotel does not need to include free parking. The hotel should be in the type of guesthouse. If there is no such hotel, how about one that is in the moderate price range. Once you find a hotel, make sure you get postcode, phone number. You also want to book a taxi to commute between the two places. **The taxi should arrive at the restaurant from the hotel by 02:45.** Once you find a taxi, make sure you get car type, phone number  
 CONVERSATION:  
 ...  
 ASSISTANT: What time will you be leaving . What time would you like to arrive ?  
 CUSTOMER: I would like to depart at 2 :45 **to return at 5 :15 . (hallucination)**  
 ASSISTANT: Thank you for using our services .  
 CUSTOMER: Thank you .  
 ASSISTANT: I ' m happy to have been able to help you today .  
 CUSTOMER: Thank you for the help the table has been booked and taxi has been booked .  
 ASSISTANT: What time would you like to leave ?  
 CUSTOMER: I would like to leave at 2 : 45 **to return at 5 :15 . (hallucination)**  
 ASSISTANT: Your booking is complete , a tesla will be picking you up . The contact number is 45067427198 .  
 CUSTOMER: Thank you , have a nice day !  
 ASSISTANT: Ok , thank you . Have a good day .  
 CUSTOMER: Can you check cote restaurant please . (looping)  
 ASSISTANT: It is in the centre area .  
 CUSTOMER: Can you check cote restaurant please . (looping)  
 ASSISTANT: They serve french , It is in the centre area .  
 CUSTOMER: Thank you for your help and have a good day !  
 ASSISTANT: Thank you for using our services .  
 CUSTOMER: Can you check the cote restaurant . (looping)  
 ASSISTANT: Okay , how about cote ? That is a french restaurant .  
 CUSTOMER: Can you check cote restaurant please .  
 ASSISTANT: That is a french restaurant . They are in the expensive price range .  
 CUSTOMER: Thank you for your help and have a good day !  
 ASSISTANT: You are more than welcome !.  
 CUSTOMER: Can you please look up cote restaurant for me please . (looping)  
 ...

---

Example 2:  
 REQUIREMENTS: You are looking for a place to stay. The hotel should have a star of 4. The hotel should be in the centre. The hotel should be in the expensive price range. The hotel does not need to include free parking. If there is no such hotel, how about one that has free parking. Once you find the hotel you want to book it for 5 people on sunday for 3 nights  
 CONVERSATION:  
 CUSTOMER: I want to find a hotel with a four star rating.  
 ASSISTANT: I am sorry I have no listings for any with four stars . (NLU misclassification since later it does find criteria)  
 CUSTOMER: How about a guesthouse?  
 ASSISTANT: Unfortunately , I couldn't find anything with four stars . I can't find any guesthouse that fit your criteria , i'm sorry .  
 CUSTOMER: That's fine, but can you try finding an expensive hotel in the centre of town?  
 ASSISTANT: There aren't any that match your criteria in the centre . Any other suggestions ? There is none that is expensive . Would you like to change your criteria ? I am sorry I have no listings for any with four stars . Sorry there is no hotel fitting the description you asked for.  
 CUSTOMER: I am looking for a 4 star hotel that is expensive and in the centre of town with free parking.  
 ASSISTANT: How about university arms hotel ? It is in the centre area . It has wifi . It is a hotel . The hotel is 4 stars .  
 CUSTOMER: Can you book it for me?  
 ASSISTANT: You are more than welcome !  
 CUSTOMER: Ok. Book me for 5 people on sunday for 3 nights.

---

Example 3:  
 REQUIREMENTS: You are also looking for a train. The train should depart from cambridge. The train should leave on sunday. The train should go to stansted airport. The train should arrive by 08:45. Once you find the train you want to make a booking for 1 people. You are looking forward to trying local restaurants. The restaurant should be in the centre. The restaurant should be in the expensive price range. The restaurant should serve modern european food. Once you find a restaurant, make sure you get postcode  
 CONVERSATION:  
 ...  
 CUSTOMER: Thanks. I'm also looking for a restaurant.  
 ASSISTANT: I have 110 options for you !. How about curry prince ?  
 CUSTOMER: I need something in the centre. and I would like to spend some money. Expensive please.  
 ASSISTANT: We have 33 such places . I would suggest cote . It is located in the centre .  
 CUSTOMER: Is Cote a modern European restaurant?  
 ASSISTANT: There no modern european restaurants that I can find right now . Would something else work ? I do n't have anything in the expensive range that fits that criteria . I ' m sorry . i can not find details for cote . I have n't found any in the centre.  
 CUSTOMER: I really need a modern european restaurant, preferably expensive. I'll try a different area, if necessary.  
 ASSISTANT: I do not have anything in that price range for modern european . Another criteria perhaps ? There are no expensive ones . I ' m sorry . i can not find details for cote . There no such restrautns in centre. There are no modern european restaurants . No expensive restaurant. I ' m sorry . i can not find details for cote . There no such restrautns in centre. (looping)  
 CUSTOMER: Can you please try one more time? I really need an expensive restaurant, thank you.  
 ASSISTANT: I am unable to find any modern european restaurants in town . There are none in expensive , perhaps something else ? I ' m sorry . i can not find details for cote . I did not find any restaurants in centre . (looping)  
 CUSTOMER: What about expensive modern european?  
 ASSISTANT: There are no modern european food places , shall I run another search ? There are none in expensive , perhaps something else ? I ' m sorry . i can not find details for cote . Sorry , there are no restaurants like that in the centre . (looping)  
 CUSTOMER: Yes please run another search.  
 ASSISTANT: Thank you for contacting the help desk . Have a great day . (forces premature end of conversation)  
 CUSTOMER: What did you find for expensive modern european?  
 ASSISTANT: There are no modern european restaurants unfortunately . There are none in expensive , perhaps something else ? I ' m sorry . i can not find details for cote . I am sorry there is none even in the centre. (looping)  
 CUSTOMER: You are sorry. I am sorry. You are unhelpful. I will find my own place to eat.  
 ASSISTANT: Have a good day .  
 CUSTOMER: Wait, I'm sorry. What can you suggest for expensive modern European restaurants?  
 ASSISTANT: I do not have anything in that price range for modern european . Another criteria perhaps ? I do n't have anything in the expensive range that fits that criteria . I ' m sorry . i can not find details for cote . I am sorry but there are no restaurants that fit that criteria in the centre . (looping)  
 CUSTOMER: I would really like a restaurant please. Can you search again?  
 ...  
 ASSISTANT: Have a good day .

---