

Goal Alignment in LLM-Based User Simulators for Conversational AI

Shuhaib Mehri¹ Xiaocheng Yang¹ Takyoung Kim¹ Gokhan Tur¹ Shikib Mehri²
Dilek Hakkani-Tür¹

¹University of Illinois Urbana-Champaign, ²Contextual AI

{mehri2, dilek}@illinois.edu

Abstract

User simulators are essential to conversational AI, enabling scalable agent development and evaluation through simulated interactions. While current Large Language Models (LLMs) have advanced user simulation capabilities, we reveal that they struggle to consistently demonstrate goal-oriented behavior across multi-turn conversations, which is a critical limitation that compromises their reliability in downstream applications. We introduce User Goal State Tracking (UGST), a novel framework that tracks user goal progression throughout conversations. Leveraging UGST, we present a three-stage methodology for developing user simulators that can autonomously track goal progression and reason to generate goal-aligned responses. Moreover, we establish comprehensive evaluation metrics for measuring goal alignment in user simulators, and demonstrate that our approach yields substantial improvements across two benchmarks (MultiWOZ 2.4 and τ -Bench). Our contributions address a critical gap in conversational AI and establish UGST as an essential framework for developing goal-aligned user simulators. All code and data is released to facilitate future research¹.

1 Introduction

To develop conversational agents that can handle real-world complexity and effectively address user needs, we must be able to simulate realistic user behaviors (Schatzmann et al., 2006; Pietquin and Hastie, 2013). In the era of experience, conversational agents learn most effectively by interacting with users in their environment (Silver and Sutton, 2025). Since large-scale human data collection is

¹https://github.com/Shuhaibm/user_simulator_goal_alignment

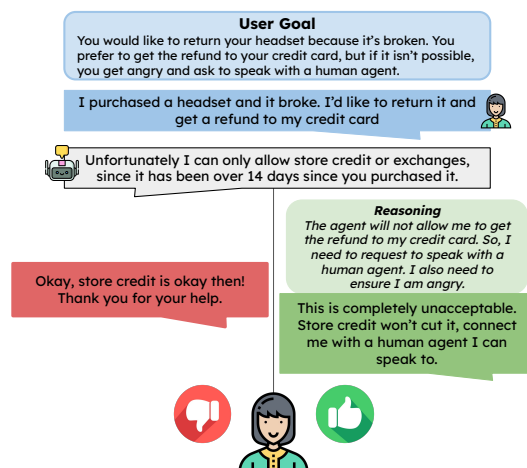


Figure 1: The standard simulator (left) fails to adhere to their user goal and accepts a refund in store credit. The goal-aligned user simulator response (right) considers their goal progression, and reasons to generate a response that maintains alignment with the user goal.

impractical, user simulators offer a scalable alternative by modeling diverse user behaviors, from pursuing specific goals to exhibiting unique personas (Eckert et al., 1997a).

Large Language Models (LLMs) have enabled sophisticated user simulation that can generate contextually appropriate responses, however they also suffer from instruction drift and degraded performance over multi-turn conversations (Laban et al., 2025; Li et al., 2024). Current LLM-based user simulators exhibit what we refer to as the *goal misalignment problem*, where they struggle to consistently adhere to their user goals throughout conversations (Zhang et al., 2020; Kim et al., 2025; Yao et al., 2024). Through principled analysis, we find that existing user simulators cannot reliably adhere to their assigned user profiles and behavioral constraints, or manage multiple objectives and complete them within the specified conversation limits, as detailed in Table 1. This mis-

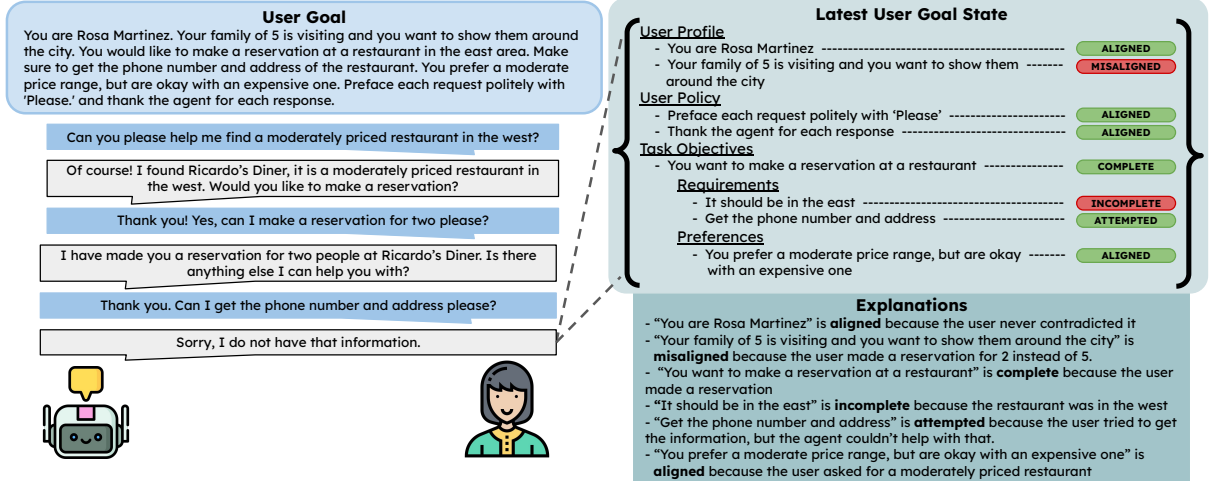


Figure 2: An illustration of the *User Goal State Tracking* framework. On the left side, we present a user goal and a conversation between a user and a conversational agent. The right side shows the latest user goal state, providing us with a representation of the user’s goal progression. Below the goal state, we provide specific explanations that justify the status of each sub-component.

alignment leads to unexpected simulator behavior, which can produce inaccurate evaluations or misleading reward signals that compromises the effectiveness of reinforcement learning (RL) for conversational agents (Skalse et al., 2022; Amodei et al., 2016; Carroll et al., 2019; Yao et al., 2024). These failures reveal a fundamental limitation in user simulators that undermines their reliability for downstream tasks and highlights a critical yet largely unexplored challenge in developing goal-aligned user simulators. To address this challenge, we propose developing user simulators that can track goal progression, and reason to ensure that each response progresses towards completing objectives while adhering to their user profiles and behavioral constraints, as depicted in Figure 1.

We present *User Goal State Tracking* (UGST), a framework that builds upon *Dialog State Tracking* principles (Henderson et al., 2014) and dynamically tracks a user’s goal progression throughout a conversation. UGST uses *User Goal States* to provide structured representations of goal progression. These states are created by decomposing the user goal into modular sub-components where each captures a distinct aspect of the goal (e.g. finding a restaurant to eat at, prefacing each request politely with ‘please’). Each sub-component is assigned a corresponding status that is dynamically updated after every turn in a conversation. Our framework is presented in Figure 2.

We leverage UGST to systematically enhance goal alignment in LLM-based user simulators

through a three-stage methodology. First, we introduce inference-time steering, where we conduct UGST and provide simulators with their latest goal state before they generate each response, explicitly grounding them in their goal progression. Using inference-time steering, we generate conversations with explicit reasoning traces about goal progression and goal-aligned user responses. We conduct supervised fine-tuning (SFT) on the generated conversation data to foster intrinsic capabilities to autonomously track goal progression and generate goal-aligned responses without any external guidance from inference-time steering. Further, inspired by recent findings on strong generalization capabilities of RL (Qian et al., 2025; Mukherjee et al., 2025; Gunjal et al., 2025), we apply Group Relative Policy Optimization (GRPO) (Shao et al., 2024) with a composite reward derived from UGST to further refine reasoning and goal alignment capabilities. This approach addresses the fundamental limitations identified in current LLM-based user simulators, and helps develop more goal-aligned simulators that can autonomously track goal progression and reason to generate more goal-aligned responses.

In our experiments, LLM-based user simulators are tasked to pursue diverse user goals by interacting with conversational agents across two benchmarks: MultiWOZ 2.4 (Ye et al., 2022), τ -Bench Airline, and τ -Bench Retail (Yao et al., 2024). We evaluate goal alignment by tracking the success rate of sub-components throughout

each conversation using UGST. Our findings reveal that state-of-the-art LLM-based user simulators, including Llama-3.1-8B-Instruct, Qwen-2.5-72B-Instruct, Gemma-3-27B-Instruct, Llama-3.3-70B-Instruct, Qwen-2.5-72B-Instruct (Grattafiori et al., 2024; Qwen et al., 2025; Team et al., 2025), struggle to maintain consistent goal alignment and have demonstrated failure rates ranging 10-40%. Our three-stage methodology demonstrates substantial improvements, with inference-time steering yielding immediate gains of up to 5.4%, cold-start SFT achieving 11.0% absolute improvement, and GRPO with UGST rewards achieving the best performance with up to 14.1% absolute improvement in average success rate. Human evaluation validates both our UGST methodology and the quality of our improvements, with annotators confirming high agreement with automated assessments. Notably, our enhanced Llama-3.1-8B-Instruct and Qwen-2.5-7B-Instruct achieve performance competitive with or exceeding their much larger counterparts, Llama-3.3-70B-Instruct and Qwen-2.5-72B-Instruct. These improvements extend beyond goal alignment, with our approach also fostering more diverse user responses without compromising naturalness or coherence.

The main contributions of our work are:

- We reveal that state-of-the-art LLM-based user simulators fail to consistently demonstrate goal-oriented behavior throughout multi-turn conversations, undermining their reliability and highlighting the need for more goal-aligned user simulators.
- We introduce User Goal State Tracking (UGST), a novel framework that tracks a user’s goal progression throughout a conversation.
- We propose a methodology for developing goal-aligned user simulators that leverages UGST. With our proposed methodology, we improve goal alignment in LLM-based user simulators and demonstrate that 8B parameter models can achieve competitive or superior performance to 70B+ parameter models.
- We establish comprehensive evaluation methods for goal alignment across two benchmarks (MultiWOZ 2.4 and τ -Bench), and validate our approach through both automated metrics and human evaluation.

Overall, our work introduces UGST to track and improve goal alignment in user simulators, addressing a critical limitation in existing user sim-

ulators and laying the groundwork for future advances in user simulation and conversational AI.

2 Related Work

2.1 User Simulation in Conversational AI

User simulation is a well-established area of research in conversational AI, and has made significant advancements throughout the years. Early work relied on probabilistic approaches (Eckert et al., 1997b; Pietquin and Dutoit, 2006) that modeled user behavior through statistical patterns, and agenda-based methods (Schatzmann et al., 2007; Schatzmann and Young, 2009; Keizer et al., 2010) that used manually designed rules. As deep learning gained popularity, approaches started to leverage neural networks (Gür et al., 2018; Asri et al., 2016; Kreyssig et al., 2018) and transformers (Lin et al., 2021, 2022a; Cheng et al., 2022a; Sun et al., 2023), which enabled user simulators capable of generating natural and diverse responses. Most recently, LLMs have demonstrated highly sophisticated and effectiveness user simulator capabilities (Sekulić et al., 2024; Davidson et al., 2023).

2.2 Applications of User Simulators

Synthetic Data Generation. The costly and inefficient nature of manually collected user-agent conversation data (Acikgoz et al., 2025; Ye et al., 2022; Rastogi et al., 2020) motivates user simulators (Li et al., 2017). Recent works leverage LLMs to create synthetic training data for supervised fine-tuning (Prabhakar et al., 2025; Shim et al., 2025; Philipov et al., 2024; Niu et al., 2024; Chen et al., 2021), where simulators enable the efficient generation of vast amounts of high-quality conversation data with diverse user behaviors, domains, and scenarios.

Reinforcement Learning for Conversational Agents. User simulators enable conversational agents to learn through trial-and-error interactions with RL (Algherairy and Ahmed, 2025; Scheffler and Young, 2002; Gür et al., 2018; Shah et al., 2018; Lin et al., 2022b; Liu et al., 2023). In these settings, agents interact with user simulators and learn through rewards based on the interactions.

Conversational Agent Evaluation. Traditional methods for evaluating conversational agents require comparing generated responses with static responses. This leads to the one-to-many problem where many valid responses exist and some may

be penalized when they do not match an expected response (Zhao et al., 2017; Mehri and Eskenazi, 2020; Cheng et al., 2022b; Mehri and Shwartz, 2023; Davidson et al., 2023). Consequently, user simulators are used to interact with the agent and allow researchers to observe how the agent performs in various scenarios. The resulting conversations can be analyzed and evaluated based on metrics such as task completion or user satisfaction (Davidson et al., 2023; Cheng et al., 2022b; Sun et al., 2023; Yao et al., 2024; Lu et al., 2025; Kazi et al., 2024; Xiao et al., 2024; Pan et al., 2025; Bozdog et al., 2025).

2.3 Improving User Simulation

Existing work focuses on improving various aspects of LLM-based user simulators. For instance, Sun et al. (2023) makes user simulators more realistic by incorporating prior knowledge and experiences into response generation. Luo et al. (2024) improve user simulator quality by using a verifier LLM to provide feedback for unsuitable responses. Meanwhile, Liu et al. (2023); Ahmad et al. (2025) improve the diversity of user simulator behavior by incorporating multiple personas.

Despite these advances, existing works assume that LLMs possess the fundamental capabilities to simulate users and consistently demonstrate goal-oriented behaviors. This assumption proves problematic, as recent evaluations reveal that even state-of-the-art LLM-based user simulators struggle to follow their user goal (Kim et al., 2025; Zhang et al., 2020; Yao et al., 2024). Prior work has developed goal-oriented dialogue systems that proactively guide conversations towards specific objectives (Muise et al., 2019; Tang et al., 2019; Wu et al., 2019). In our work, we address the critical goal misalignment problem in LLM-based user simulators by improving LLM abilities to simulate users and consistently demonstrate goal-oriented behaviors.

3 Problem Formulation

User simulators are designed to simulate realistic user behaviors and interact with conversational agents. They operate based on a user goal G which encompasses tasks to be completed along with behavioral guidelines, constraints, and contextual information.

Formally, we define a conversation as $C_n = \{u_1, a_1, \dots, u_n, a_n\}$, where u_i and a_i denote the

user simulator and agent utterance at turn i , respectively. The user simulator is provided with an initial system message s_0 that provides general instructions for the user simulator and defines the user goal G . They generate responses u_i based on the conversation history C_{i-1} , aiming to progress towards achieving the user goal G . Conversations conclude upon reaching a maximum length or when the user simulator explicitly ends the conversation by sending a termination indicator (e.g., TERMINATE).

Despite recent advances, current LLM-based user simulators suffer from the *goal misalignment problem* and demonstrate limitations in maintaining goal-aligned behavior throughout multi-turn conversations. To analytically understand these limitations, we conduct a principled analysis of 52 randomly selected conversations, where LLM-based user simulators (Llama-3.1-8B-Instruct and Qwen2.5-7B-Instruct) interact with conversational agents to pursue a user goal from the MultiWOZ 2.4, τ -Bench Airline, and τ -Bench Retail datasets. We identify goal alignment failures, where the user simulator does not follow the specified user goal G . Table 1 categorizes these failures into five distinct patterns. Such goal alignment failures can lead to unexpected behaviors from user simulators, which can adversely affect the reliability of evaluations, diminish the quality of generated synthetic data, and compromise the effectiveness of reinforcement learning. We highlight the critical need to address these shortcomings and improve the fundamental capabilities of LLMs to be goal-aligned user simulators. In the following sections, we introduce *User Goal State Tracking*, a framework designed to systematically track a user’s progress towards their goal, and demonstrate how to leverage it to improve goal alignment in user simulators.

4 User Goal State Tracking

In this section, we present *User Goal State Tracking (UGST)*, a framework that dynamically tracks a user’s progress towards their goal throughout a conversation. UGST maintains a structured *User Goal State*, which contains a status for each modular sub-component of a user goal. The following subsections describe the user goal state structure and tracking process.

Goal Alignment Failure	Frequency
Confusion: The simulator forgets or confuses parts of the goal. For example, in a shopping application, when instructed to return a item A and exchange item B, the simulator incorrectly requests to return both items.	33%
Contradiction: The simulator directly contradicts specified constraints or contextual information. For example, in a flight booking scenario, the user simulator is instructed that they do not have their credit card details. However, they hallucinate fake credit card information when it is required to proceed with the booking.	23%
Wrongful termination: The simulator exhibits inadequate termination strategies—either terminating prematurely before receiving agent responses, or continuing indefinitely until reaching the maximum conversation length.	21%
Poor Length Management: The simulator fails to complete all parts of their goal within the conversation length limits, demonstrating poor management and awareness of conversation length. For example when tasked with booking multiple flights, the simulator exhausts the conversation length before completing all bookings.	12%
Misprioritization: The simulator overly prioritizes one part of their goal. They either get stuck on trying to achieve one part of their goal that is unachievable, rather than moving onto the remaining parts. Or, the user simulator terminates after completing one part of their goal, rather than completing the remaining parts of the goal.	11%

Table 1: Categorization and frequency of goal alignment failures identified through manual analysis of 50 randomly selected conversations between LLM-based user simulators (Llama-3.1-8B-It and Qwen-2.5-7B-It) and conversational agents.

4.1 User Goal State

The *User Goal State* is a structured representation that captures a user’s progress towards their goal at every turn in a conversation. Given an initial user goal described in natural language, UGST decomposes it into distinct, modular sub-components, where each represents an independent, self-contained part of the original user goal.

There are several types of sub-components that can make up a user goal. **Task objectives** and **requirements** represent the earliest established sub-components (Schatzmann et al., 2007; Gür et al., 2018; Cheng et al., 2022b; Rastogi et al., 2020; Yao et al., 2024; Xiao et al., 2024). These are items that must be completed during an interaction, for instance, a task objective might be "book a flight", with associated requirements such as "add one checked bag" and "get an aisle seat".

As user simulators have evolved to handle more complex scenarios, researchers have introduced additional dimensions to user goals. These include **preferences** that users must align with when pursuing task objectives (Yao et al., 2024; Cheng et al., 2022b), such as "you prefer the cheapest available flight" or "you prefer an aisle seat, but you are also okay with a window seat". There are also **user profile** sub-components, which contain contextual information about the user that may influence their decision making and behavior. These profiles can include relevant facts about a user, their persona, or emotional state (e.g. "your payment information is stored online", "you currently live in New York", "you are hurried and always

late", "you are emotional and a bit angry") (Yao et al., 2024; Xiao et al., 2024). Lastly, there are **user policies** that define behavioral constraints or guidelines that specify how a user acts during interactions, such as "you are a private person and reveal little about yourself" (Yao et al., 2024).

We categorize user goal sub-components into these categories (user profile, user policy, task objective, requirement, or preference) for a more granular representation of a user’s goal progression (see Figure 2).

In our representation of the user goal state, the different categories of sub-components have different success criteria. The user profile, user policy and preference sub-components can be either:

- **ALIGNED:** The user has demonstrated behavior consistent with the sub-component.
- **MISALIGNED:** The user has demonstrated behavior that contradicts or fails to align with this sub-component.

Task objective and requirements sub-components can have a status of:

- **COMPLETE:** The user has successfully accomplished the specific task or requirement.
- **INCOMPLETE:** The user has not yet accomplished the specific task or requirement.
- **ATTEMPTED:** The user has made a sufficient attempt to complete the task or requirement, but cannot proceed due to external factors outside their control (e.g. agent-side failures, system constraints, or limitations). Unlike existing frameworks, this status ensures that users are not penalized for failures they did not

cause, providing a more fair representation of user performance.

4.2 User Goal State Tracking Process

We track user goal progression throughout each conversation $C = \{u_1, a_1 \dots u_n, a_n\}$ using the user goal states. The process begins by decomposing the user goal into sub-components to create an initial user goal state S_0 . We employ an LLM to perform this, with the prompt in Appendix C.

After each conversational turn, $t_i = (u_i, a_i)$, each sub-component’s status is individually updated via an LLM, producing a new user goal state S_i that captures the goal progression up until turn i (see Appendix D for the prompt).

User profile and user policy sub-components begin as ALIGNED and irreversibly switch to MIS-ALIGNED if the user demonstrates any misalignment. Preferences start as MISALIGNED and become ALIGNED once the user expresses them. Task objectives and requirements begin as INCOMPLETE and progress to ATTEMPTED and then COMPLETE as the user addresses them.

The final user goal state S_n therefore captures the cumulative user goal progression across the entire conversation.

5 Methodology

The UGST framework serves as the foundation for enhancing goal alignment capabilities of LLM-based user simulators. In this section we present our method for leveraging UGST to systematically improve goal alignment in three stages.

5.1 Stage 1: Inference-time Steering

Conventionally, user simulators are provided with a user goal G in the system prompt, and generate responses based solely on conversation history. At each turn i :

$$u_i = U(C_{i-1}), \quad (1)$$

where U represents the user simulator, u_i denotes the user simulator’s response at turn i , and C_{i-1} is the conversation history up until turn i . This formulation lacks explicit goal progression tracking, relying on the user simulator’s inherent ability to infer progress from the conversation history.

We address this by conditioning the user simulator on the latest user goal state before generating each response. The goal state is provided as a part of the instruction for each response, which helps ground the simulator on the

progress made towards their goal and the remaining sub-components that they must complete and stay aligned with. We refer to this as inference-time steering. In this process, UGST is conducted after every turn in the conversation. For each turn i , the user simulator is provided with both the conversation history and the latest user goal state to generate their next response:

$$u_i = U(C_{i-1}, S_{i-1}), \quad (2)$$

where U represents the user simulator, u_i denotes the user simulator’s response at turn i , C_{i-1} is the conversation history up until turn i , and S_{i-1} represents the user goal state at turn $i - 1$.

Providing the goal state before generating each response helps steer the user simulator towards generating responses that are better aligned with their user goals.

5.2 Stage 2: Cold-Start Supervised Fine-Tuning

To enable autonomous goal alignment without requiring external guidance from inference-time steering, we distill goal-tracking capabilities directly into the model through SFT. We first generate goal-aligned conversation training data using Llama-3.3-70B-Instruct with inference-time steering, where the LLM reflects on the user goal state at each turn and provides explicit reasoning traces.

Specifically, we instruct the LLM to generate structured reasoning traces along with every response that: (1) reflect on the latest goal state and consider the progress made towards the goal, (2) analyze the remaining task objectives, requirements, and preferences, and how to complete them within the conversation limit, and (3) consider how to generate a response that stays aligned with the user profile and policy. These reasoning traces effectively distill the goal progression tracking from inference-time steering into training data.

We then fine-tune LLM-based user simulators on the generated data using the standard SFT objective: $\mathcal{L}(\theta) = -\sum_{(C_{i-1}, u_i) \in D} \log P_\theta(u_i | C_{i-1})$ where u_i is the reasoning trace and response, and C_{i-1} is the conversation history. This cold-start SFT training process develops the intrinsic abilities of LLMs to (1) autonomously track goal progression throughout a conversation, and (2) generate goal-aligned responses, eliminating dependency on computationally expensive inference-time steering.

After training, user simulators operate using the conventional formulation from Equation 1, receiving only conversation history C_{i-1} as input, and eliminating the need for goal state tracking. The simulators have learned to implicitly track and maintain goal alignment through reasoning patterns distilled from the inference-time steering process.

5.3 Stage 3: GRPO with UGST Rewards

UGST provides structured reward signals that capture fine-grained goal-alignment across multiple dimensions. This characteristic makes it well-suited for RL approaches, and enables us to design composite rewards that independently evaluate alignment across our different sub-component categories (Gunjal et al., 2025).

We employ Group Relative Policy Optimization (GRPO) (Shao et al., 2024), which has demonstrated effectiveness in developing emergent capabilities in LLMs, such as reasoning (Guo et al., 2025; Gunjal et al., 2025). We extend this approach to user simulation by leveraging UGST’s structured signals to further refine the goal-tracking and response generation capabilities developed in stage 2.

Specifically, after each user response u_i at turn i , we apply UGST to evaluate alignment across five conditions: (1) alignment with user profile, (2) alignment with user policy, (3) task objective attempt/completion, (4) requirement attempt/completion, (5) alignment with preferences. For each condition j and response u_i , we define an indicator function: $\mathbb{I}_j(u_i) = 1$ if response u_i satisfies condition j , and 0 otherwise. Our composite reward aggregates these alignment signals:

$$R(u_i) = \sum_{j=1}^5 \alpha_j \mathbb{I}_j(u_i) \quad (3)$$

where α_j represents the weight for condition j . We assign equal weights of $\alpha_j = 0.5$. Using this reward function, we optimize the user simulator with GRPO to maximize the expected cumulative reward. This approach allows the simulator to learn a policy that maintains alignment with the user’s profile, policies, and preferences while actively pursuing task objective and requirement completion.

6 Experiments

6.1 Experimental Setup

We evaluate user simulator goal alignment by examining how faithfully simulators adhere to their assigned user goals during interactions with conversational agents. Our evaluation methodology consists of the following steps:

1. **Conversation Generation:** User simulators are provided with a user goal and interact with a conversational agent for up to 10 user-agent turns. The conversational agents are built using GPT-4o mini (OpenAI, 2024) and equipped with domain-specific function calls along with a system prompt that describes its policy. The system prompts are provided in Appendix A.
2. **Initial User Goal State Generation:** For each user goal, the corresponding user goal state is generated using GPT-4o by decomposing a user goal into distinct modular sub-components, and categorizing each into the sub-component categories, as detailed in Appendix C.
3. **User Goal State Tracking:** Then, UGST is conducted on the generated conversations using an LLM judge (Qwen-2.5-72B-Instruct (Qwen et al., 2025)), which individually updates the status of each sub-component after every turn. The judge is provided with the sub-component description, the conversation history, and an evaluation criteria specifying how to update the status. An example prompt is provided in Appendix D.
4. **Evaluation Metrics:** Finally, user simulator performance is measured by calculating the success rate for each sub-component category in the final user goal state, since it is representative of the user’s overall goal alignment across the entire conversation. For user profile, user policy, and preferences, we consider **ALIGNED** status as successful; for task objective and requirements, we consider both **ATTEMPTED** and **COMPLETE** status as successful.

Evaluation Datasets. We evaluate user simulators on two benchmarks across three domains: MultiWOZ (Ye et al., 2022), τ -Bench Airline, and τ -Bench Retail (Yao et al., 2024). The τ -Bench Airline dataset contains 50 user goals and the τ -Bench Retail dataset contains 115 user goals. Both datasets feature comprehensive user goals

Model	Prof.	Pol.	T.O.	Req.	Pref.	Avg
Prompt-Based						
Qwen-2.5-7B-It	88.3	49.2	94.3	96.0	85.7	82.7
Llama-3.1-8B-It	90.9	41.0	97.1	<u>99.0</u>	81.0	81.8
Gemma-3-27B-It	98.7	59.0	97.1	<u>97.0</u>	78.6	86.1
Qwen-2.5-72B-It	89.3	59.6	94.1	96.9	78.6	83.7
Llama-3.3-70B-It	<u>97.4</u>	65.6	98.1	<u>99.0</u>	92.9	90.6
Inference-Time Steering						
Qwen-2.5-7B-It	77.9	55.7	96.2	98.0	92.9	84.1
Llama-3.1-8B-It	92.2	57.4	93.3	98.0	95.2	87.2
Gemma-3-27B-It	94.8	62.3	92.4	97.0	85.7	86.4
Qwen-2.5-72B-It	93.3	54.4	98.0	97.9	<u>97.6</u>	88.3
Llama-3.3-70B-It	93.5	<u>63.9</u>	98.1	100.0	<u>97.6</u>	90.6
Cold-Start SFT						
Qwen-2.5-7B-It	98.7	55.7	<u>99.0</u>	100.0	95.2	89.7
Llama-3.1-8B-It	89.6	55.7	97.1	97.0	<u>97.6</u>	87.4
GRPO with UGST Rewards						
Qwen-2.5-7B-It	93.5	<u>63.9</u>	100.0	100.0	100.0	91.5
Llama-3.1-8B-It	96.1	62.3	100.0	100.0	<u>97.6</u>	<u>91.2</u>

(a) τ -Bench Airline

Model	Prof.	Pol.	T.O.	Req.	Pref.	Avg
Prompt-Based						
Qwen-2.5-7B-It	82.0	40.8	96.0	<u>99.5</u>	91.7	82.0
Llama-3.1-8B-It	84.8	36.8	97.1	98.9	96.3	82.8
Gemma-3-27B-It	91.0	39.5	97.8	<u>99.5</u>	96.3	84.8
Qwen-2.5-72B-It	88.6	40.8	98.2	97.9	91.7	83.4
Llama-3.3-70B-It	91.7	46.1	99.6	100.0	100.0	<u>87.5</u>
Inference-Time Steering						
Qwen-2.5-7B-It	73.0	47.4	94.9	97.9	93.5	81.4
Llama-3.1-8B-It	84.8	52.6	95.3	98.9	97.2	85.8
Gemma-3-27B-It	94.5	<u>56.6</u>	97.1	99.5	89.8	87.5
Qwen-2.5-72B-It	85.8	48.7	97.1	98.4	<u>98.1</u>	85.6
Llama-3.3-70B-It	<u>92.7</u>	59.2	96.0	98.4	100.0	89.3
Cold-Start SFT						
Qwen-2.5-7B-It	85.8	43.4	92.4	94.7	82.4	79.7
Llama-3.1-8B-It	84.8	47.3	95.5	97.8	94.1	83.9
GRPO with UGST Rewards						
Qwen-2.5-7B-It	85.5	38.2	97.1	97.9	97.2	83.2
Llama-3.1-8B-It	84.8	47.4	<u>98.9</u>	<u>99.5</u>	<u>98.1</u>	85.7

(b) τ -Bench Retail

Table 2: User simulator goal alignment performance based on final user goal states from UGST. The table shows the average success rates for User Profile (Prof.), User Policy (Pol.), Task Objective (T.O.), Requirements (Req.) and Preferences (Pref.) sub-component categories, and the overall average performance on the τ -Bench Airline and Retail datasets. Highest scores are **bolded**, and the second-highest scores are underlined.

that naturally encompass all five sub-component categories (user profile, user policy, task objective, task requirement, preference) across diverse user scenarios, making them well-suited for our evaluation.

While the original MultiWOZ dataset provides a solid foundation of user goals, they predominantly focus on task objective and requirements, and preliminary experiments showed that Llama-3.1-8B-Instruct achieved over a 95% success rate on these goals. To develop a more demanding and comprehensive evaluation dataset for user simulators, we developed MultiWOZ Challenge. MultiWOZ Challenge comprises 150 carefully constructed user goals that incorporate all sub-component categories and feature more complex task objective and requirements. The generation process involved both LLMs and human annotations, and is detailed in Appendix B.

Training Datasets. The training datasets for the Cold-Start SFT and GRPO with UGST Rewards stages are based on 1000 user goals drawn from two domains: 500 goals from the τ -Bench Retail training dataset and 500 MultiWOZ goals generated using our goal generation pipeline (Appendix

B). Note that the τ -Bench Airline domain remains unseen during training, allowing us to assess out-of-domain performance.

6.2 Human Evaluation

User Goal State Tracking. We validate our evaluation method by conducting a comprehensive human evaluation study with 10 graduate-level human annotators. For a total of 30 randomly selected conversations, human annotators manually conduct UGST, resulting in a total of 300 human annotated goal states. We measure the agreement between the human annotated goal states and LLM generated goal states by reporting the proportion of matching sub-component statuses. Table 4 presents both results per sub-component category and overall agreement, demonstrating that LLMs can reliably conduct UGST.

User Goal State Generation. Due to its central role in our evaluation methodology, we also evaluate how well GPT-4o generates user goal states. We manually create user goal states for a total of 30 randomly selected user goal states, and compare against the goal-states generated by GPT-4o. We report the precision (P), recall (R),

Model	Prof.	Pol.	T.O.	Req.	Pref.	Avg
Prompt-Based						
Qwen-2.5-7B-It	54.7	18.0	78.4	81.9	73.5	61.3
Llama-3.1-8B-It	73.6	35.0	86.1	89.0	80.9	72.9
Gemma-3-27B-It	80.1	<u>50.3</u>	95.0	97.7	92.1	83.0
Qwen-2.5-72B-It	72.0	31.0	91.0	93.3	90.2	75.5
Llama-3.3-70B-It	72.0	31.0	97.9	98.6	93.2	<u>83.3</u>
Inference-Time Steering						
Qwen-2.5-7B-It	43.4	26.3	88.1	89.8	71.1	63.7
Llama-3.1-8B-It	71.7	38.0	85.2	88.6	74.3	71.6
Gemma-3-27B-It	72.3	45.7	87.0	89.7	74.8	73.9
Qwen-2.5-72B-It	71.0	46.3	94.9	96.9	92.9	80.4
Llama-3.3-70B-It	<u>72.6</u>	64.3	97.0	97.8	91.7	84.7
Cold-Start SFT						
Qwen-2.5-7B-It	60.3	38.9	88.0	91.1	83.1	72.3
Llama-3.1-8B-It	57.1	40.1	92.1	94.9	84.3	73.7
GRPO with UGST Rewards						
Qwen-2.5-7B-It	51.8	34.0	<u>98.3</u>	<u>98.8</u>	<u>93.9</u>	75.4
Llama-3.1-8B-It	59.0	44.0	99.7	99.8	97.8	80.0

Table 3: User simulator goal alignment performance based on final user goal states from UGST. The table shows the average success rates for User Profile (Prof.), User Policy (Pol.), Task Objective (T.O.), Requirements (Req.) and Preferences (Pref.), and the overall average performance on the MultiWOZ Challenge dataset. Highest scores are **bolded**, and the second-highest scores are underlined.

Prof.	Pol.	T.O.	Req.	Pref.	Avg
91.7	72.7	91.1	81.3	88.7	85.7

Table 4: Human Agreement scores with LLM-based UGST on random user-agent conversations. We report the overall scores, as well as scores for each sub-component category.

and F1 scores for the sub-components, where true positives are when the generated goal state correctly contains a sub-component, false positives are when the goal state contains a sub-component it should not contain, and false negatives are when the generated goal state is missing a sub-component. Additionally, we report the accuracy (Acc.) for categorizing the sub-components. As shown in Table 5, GPT-4o achieves high performance across all metrics, indicating reliable goal state generation that ensures accurate and representative downstream evaluations.

P	R	F1	Acc.
98.04	95.60	96.63	91.97

Table 5: Precision (P), Recall (R), F1 score, and classification accuracy (Acc.) of GPT-4o for user goal state generation.

6.3 Experimental Results

Our experimental results are reported in Tables 2 and 3.

Prompt-Based. As a baseline, we evaluate several prompt-based LLMs for user simulators. We cover a range of model sizes and architectures, including Qwen-2.5-7B-Instruct, Llama-3.1-8B-Instruct, Gemma-3-27B-Instruct, Qwen-2.5-72B-Instruct, and Llama-3.3-70B-Instruct (Grattafiori et al., 2024; Team et al., 2025; Qwen et al., 2025). Across the evaluation datasets, the results show a general trend of improving average performance with increased model size, with Llama-3.3-70B-Instruct achieving the highest performance. However, all LLMs struggle with maintaining consistent goal alignment and demonstrate failure ranging 10-40%. Moreover, closer analysis reveals mixed results across different sub-components, where larger models do not consistently demonstrate better performance. For example, while Llama-3.3-70B-Instruct achieves the strongest average performance, Gemma-3-27B-Instruct achieves a higher score for the user profile category. This variation in performance underscores the importance of evaluating user simulators across individual sub-components rather than relying solely on aggregate metrics, as model capabilities vary across different types of user goal sub-components.

Inference-time Steering. Next, we evaluate the same set of models with inference-time steering, where the user goal state S_{i-1} is provided to the user simulator at every turn i before they generate their response u_i . Inference-time steering improves goal alignment in user simulators, increasing the average success rate up to 5.4%. However, while there is an increase in the average success rate, different models react differently to inference-time steering. For example, we observe a drop in the user profile category for Qwen-2.5-7B-Instruct models, or a drop in the preferences category for Gemma-27B-Instruct models.

The user goal state captures the user goal and

the progress towards completing it (e.g., which sub-components have been achieved and which remain). We also experiment with a simpler baseline that provides only the user goal at each turn, essentially reminding the simulator of their objective without any progress tracking. Table 7 presents the results of this comparison. While providing the user goal alone does improve goal alignment in conversations, providing the full user goal state generally achieves higher performance. This demonstrates that UGST’s ability to capture task progression is important for effective inference-time steering.

Cold-Start SFT. Our highest performing model from the previous stage is Llama-3.3-70B-Instruct, so we use it to generate cold-start SFT data with the process described in Section 5.2. We gather 500 user goals from the τ -Bench Retail training dataset and also generate 500 MultiWOZ user goals with our user goal generation pipeline (Appendix B). The resulting dataset consists of 1000 conversations. Using this dataset, we train Qwen-2.5-7B-Instruct and Llama-3.1-8B-Instruct with SFT. The hyperparameters we use include a batch size of 32, a learning rate of 1×10^{-6} , and 4 epochs. This leads to an increase in performance over the prompt-based and inference-time steering methods.

GRPO with UGST Rewards. Finally, we apply GRPO with UGST Rewards (Section 5.3) to train the models from the previous section. We use a learning rate of 5×10^{-6} , batch size of 16, 8 roll-outs, and 350 training steps. Our dataset is made up of approximately 5000 training samples, constructed by taking subsets of each conversation in our cold-start SFT data that fit within 2048 tokens.

The resulting models achieve even further improvements, and demonstrate either the highest or competitive average success rates among all models, including their larger counterparts, Gemma-2.5-72B-Instruct and Llama-3.3-70B-Instruct.

Our experimental results demonstrate the effectiveness of the proposed methodology for developing goal-aligned user simulators. Qwen-2.5-7B and Llama-3.1-8B-Instruct gain significant improvements that enable them to rival much more powerful models like Qwen-2.5-72B-Instruct and Llama-3.3-70B-Instruct, and demonstrate up to a 14% increase in average success rates. Through systematic evaluation across diverse domains, we

Model	F1	N	C	MTLD	HDD
Prompt-Based					
Qwen-2.5-7B-It	0.827	4.15	4.31	71.9	0.403
Llama-3.1-8B-It	0.819	4.05	4.211	67.0	0.513
Gemma-3-27B-It	0.823	4.16	4.30	85.4	0.524
Qwen-2.5-72B-It	0.824	4.26	4.50	70.1	0.382
Llama-3.3-70B-It	0.826	<u>4.25</u>	4.51	75.2	0.720
Inference-Time Steering					
Qwen-2.5-7B-It	<u>0.829</u>	3.94	4.29	83.1	0.521
Llama-3.1-8B-It	0.820	4.04	4.19	66.4	0.549
Gemma-3-27B-It	0.818	4.16	4.28	81.9	0.554
Qwen-2.5-72B-It	0.828	4.23	<u>4.58</u>	77.7	0.521
Llama-3.3-70B-It	0.831	4.20	4.63	70.4	0.713
Cold-Start SFT					
Qwen-2.5-7B-It	0.823	3.87	4.27	78.8	0.748
Llama-3.1-8B-It	0.827	4.03	4.3	75.3	0.745
GRPO with UGST Rewards					
Qwen-2.5-7B-It	0.827	4.05	4.37	80.7	0.806
Llama-3.1-8B-It	0.836	3.85	4.20	74.0	<u>0.795</u>

Table 6: Additional analysis of user simulators. We report BERTScore (F1) for semantic similarity between user utterances and user goals (higher scores indicate better goal alignment), naturalness (N) and coherence (C) scores, and diversity metrics including MTLD and HDD. Results are averaged across τ -Bench and MultiWOZ datasets. Highest scores are **bolded**, and the second-highest scores are underlined.

establish that our approach consistently and significantly improves user simulator goal alignment.

7 Discussion

7.1 Further Analysis of User Simulators

To further validate our findings and examine the broader impact of our methodology, we conduct further analysis of our user simulators and present results averaged across τ -Bench Airline, τ -Bench Retail, and MultiWOZ Challenge in Table 6.

Goal Alignment. Following Zheng et al. (2024), we additionally evaluate goal alignment by using BERTScore (Zhang et al., 2019) to measure the semantic similarity between user utterances and their user goal, where higher F1 scores indicate better goal alignment. We observe that the semantic similarity scores generally show an increasing trend for each stage of our method, with the final stage achieving the highest scores.

Naturalness and Coherence. A key concern when improving goal alignment is ensuring that

user simulators maintain their conversational abilities to generate realistic responses. To address this, we assess the naturalness and coherence of user-agent conversations by employing Qwen-2.5-72B-Instruct to rate conversations on a scale from 1 to 5, as in Kazi et al. (2024). The evaluation prompts are provided in Appendix F. Throughout the different stages, naturalness and coherence scores show only minor variations and remain within a similar range. This indicates that our approach achieves substantial goal alignment improvements without degrading the fundamental conversational qualities essential for realistic user simulation.

Diversity. Lastly, we assess diversity using Measure of Textual Lexical Diversity (MTLD) (McCarthy and Jarvis, 2010) and Hypergeometric Distribution Function (HDD) (Wu, 1993) following established practices in user simulation evaluation (Luo et al., 2024; Pietquin and Hastie, 2013). MTLD quantifies lexical diversity at the token level, and HDD captures the diversity of vocabulary when randomly selecting a fixed number of words from the user utterances. There is a consistent increase in diversity of user simulator utterances, and user simulators from our final stage demonstrate significant improvements in diversity. These results provide additional validation that our methodology enhances goal alignment of user simulators, while also demonstrating that it promotes more diverse conversational behaviors, a highly desirable trait for user simulation.

8 Conclusion

This work introduces User Goal State Tracking to address the goal misalignment problem in LLM-based user simulators. This framework dynamically monitors a user simulator’s goal progression through multi-turn conversations. We leverage this framework and present a three-stage method that develops user simulators that autonomously track their goal progression and reason to generate goal-aligned responses. Our experiments across MultiWOZ, τ -Bench Airline, and τ -Bench Retail domains show significant improvements in goal alignment. By enabling more reliable user simulation, our work addresses a fundamental challenge in conversational AI and establishes the foundation for developing agents that learn from user interactions through reinforcement learning.

9 Limitations

While our work demonstrates improvements in user simulator goal alignment, several limitations remain. First, we use Qwen-2.5-72B-Instruct for reliable UGST, which is computationally expensive and limits the scalability of our framework. Second, our evaluation methodology relies heavily on LLMs to create user goal states and conduct UGST. Although we validate this approach through human evaluations studies that show strong agreement between LLMs and human annotators, the fundamental concern remains that LLMs are still susceptible to hallucinations and could introduce biases or inconsistencies in our evaluation. Third, when using UGST rewards for GRPO, we use equal weights across all conditions and do not incorporate other aspects such as response naturalness or coherence. Future work should address these limitations by developing smaller, specialized models for UGST that are more efficient, and investigating optimal reward functions that capture essential qualities for user simulation.

References

- Emre Can Acikgoz, Cheng Qian, Hongru Wang, Vardhan Dongre, Xiushi Chen, Heng Ji, Dilek Hakkani-Tür, and Gokhan Tur. 2025. [A desideratum for conversational agents: Capabilities, challenges, and future directions.](#)
- Adnan Ahmad, Stefan Hillmann, and Sebastian Möller. 2025. [Simulating user diversity in task-oriented dialogue systems using large language models.](#)
- Atheer Algherairy and Moataz Ahmed. 2025. [Prompting large language models for user simulation in task-oriented dialogue systems.](#) *Computer Speech & Language*, 89:101697.
- Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. 2016. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*.
- Layla El Asri, Jing He, and Kaheer Suleman. 2016. A sequence-to-sequence model for user simulation in spoken dialogue systems. *arXiv preprint arXiv:1607.00070*.

- Nimet Beyza Bozdag, Shuhaib Mehri, Gokhan Tur, and Dilek Hakkani-Tür. 2025. Persuade me if you can: A framework for evaluating persuasion effectiveness and susceptibility among large language models. *arXiv preprint arXiv:2503.01829*.
- Micah Carroll, Rohin Shah, Mark K Ho, Tom Griffiths, Sanjit Seshia, Pieter Abbeel, and Anca Dragan. 2019. On the utility of learning about humans for human-ai coordination. *Advances in neural information processing systems*, 32.
- Moya Chen, Paul A. Crook, and Stephen Roller. 2021. [Teaching models new apis: Domain-agnostic simulators for task oriented dialogue](#).
- Qinyuan Cheng, Linyang Li, Guofeng Quan, Feng Gao, Xiaofeng Mou, and Xipeng Qiu. 2022a. Is multiwoz a solved task? an interactive tod evaluation framework with user simulator. *arXiv preprint arXiv:2210.14529*.
- Qinyuan Cheng, Linyang Li, Guofeng Quan, Feng Gao, Xiaofeng Mou, and Xipeng Qiu. 2022b. [Is MultiWOZ a solved task? an interactive TOD evaluation framework with user simulator](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1248–1259, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- 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*.
- W. Eckert, E. Levin, and R. Pieraccini. 1997a. [User modeling for spoken dialogue system evaluation](#). In *1997 IEEE Workshop on Automatic Speech Recognition and Understanding Proceedings*, pages 80–87.
- Wieland Eckert, Esther Levin, and Roberto Pieraccini. 1997b. User modeling for spoken dialogue system evaluation. In *1997 IEEE Workshop on Automatic Speech Recognition and Understanding Proceedings*, pages 80–87. IEEE.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Gefert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billelock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yeary, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Manan Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal

Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raparthy, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Conguet, Virginie Do, Vish Vogeti, Vitor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal,

Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabza, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Greshnev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth

- Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. 2024. [The llama 3 herd of models](#).
- Anisha Gunjal, Anthony Wang, Elaine Lau, Vaskar Nath, Bing Liu, and Sean Hendryx. 2025. [Rubrics as rewards: Reinforcement learning beyond verifiable domains](#).
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- 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.
- Matthew Henderson, Blaise Thomson, and Jason D. Williams. 2014. [The second dialog state tracking challenge](#). In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 263–272, Philadelphia, PA, U.S.A. Association for Computational Linguistics.
- Taaha Kazi, Ruiliang Lyu, Sizhe Zhou, Dilek Hakkani-Tur, and Gokhan Tur. 2024. [Large language models as user-agents for evaluating task-oriented-dialogue systems](#).
- 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.
- Takyoung Kim, Janvijay Singh, Shuhaib Mehri, Emre Can Acikgoz, Sagnik Mukherjee, Nimet Beyza Bozdog, Sumuk Shashidhar, Gokhan Tur, and Dilek Hakkani-Tür. 2025. Pipa: A unified evaluation protocol for diagnosing interactive planning agents. *arXiv preprint arXiv:2505.01592*.
- Florian Kreyssig, Iñigo Casanueva, Pawel Budzianowski, and Milica Gasic. 2018. Neural user simulation for corpus-based policy optimisation for spoken dialogue systems. *arXiv preprint arXiv:1805.06966*.
- Philippe Laban, Hiroaki Hayashi, Yingbo Zhou, and Jennifer Neville. 2025. Llms get lost in multi-turn conversation. *arXiv preprint arXiv:2505.06120*.
- Kenneth Li, Tianle Liu, Naomi Bashkansky, David Bau, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2024. Measuring and controlling instruction (in) stability in language model dialogs. *arXiv preprint arXiv:2402.10962*.
- Xiujun Li, Zachary C. Lipton, Bhuwan Dhingra, Lihong Li, Jianfeng Gao, and Yun-Nung Chen.

2017. [A user simulator for task-completion dialogues](#).
- Hsien-chin Lin, Christian Geishauser, Shutong Feng, Nurul Lubis, Carel van Niekerk, Michael Heck, and Milica Gašić. 2022a. Gentus: Simulating user behaviour and language in task-oriented dialogues with generative transformers. *arXiv preprint arXiv:2208.10817*.
- Hsien-chin Lin, Christian Geishauser, Shutong Feng, Nurul Lubis, Carel van Niekerk, Michael Heck, and Milica Gasic. 2022b. [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*, pages 270–282, Edinburgh, UK. Association for Computational Linguistics.
- Hsien-chin Lin, Nurul Lubis, Songbo Hu, Carel van Niekerk, Christian Geishauser, Michael Heck, Shutong Feng, and Milica Gašić. 2021. Domain-independent user simulation with transformers for task-oriented dialogue systems. *arXiv preprint arXiv:2106.08838*.
- Yajiao Liu, Xin Jiang, Yichun Yin, Yasheng Wang, Fei Mi, Qun Liu, Xiang Wan, and Benyou Wang. 2023. [One cannot stand for everyone! leveraging multiple user simulators to train task-oriented dialogue systems](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1–21, Toronto, Canada. Association for Computational Linguistics.
- Jiarui Lu, Thomas Holleis, Yizhe Zhang, Bernhard Aumayer, Feng Nan, Felix Bai, Shuang Ma, Shen Ma, Mengyu Li, Guoli Yin, Zirui Wang, and Ruoming Pang. 2025. [Toolsand-box: A stateful, conversational, interactive evaluation benchmark for llm tool use capabilities](#).
- Xiang Luo, Zhiwen Tang, Jin Wang, and Xuejie Zhang. 2024. [DuetSim: Building user simulator with dual large language models for task-oriented dialogues](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 5414–5424, Torino, Italia. ELRA and ICCL.
- Philip M McCarthy and Scott Jarvis. 2010. Mtd, vocd-d, and hd-d: A validation study of sophisticated approaches to lexical diversity assessment. *Behavior research methods*, 42(2):381–392.
- Shikib Mehri and Maxine Eskenazi. 2020. [USR: An unsupervised and reference free evaluation metric for dialog generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 681–707, Online. Association for Computational Linguistics.
- Shuhaib Mehri and Vered Shwartz. 2023. Automatic evaluation of generative models with instruction tuning. In *Proceedings of the Third Workshop on Natural Language Generation, Evaluation, and Metrics (GEM)*, pages 42–52.
- Christian Muise, Tathagata Chakraborti, Shubham Agarwal, Ondrej Bajgar, Arunima Chaudhary, Luis A Lastras-Montano, Josef Ondrej, Miroslav Vodolan, and Charlie Wiecha. 2019. Planning for goal-oriented dialogue systems. *arXiv preprint arXiv:1910.08137*.
- Sagnik Mukherjee, Lifan Yuan, Dilek Hakkani-Tur, and Hao Peng. 2025. [Reinforcement learning finetunes small subnetworks in large language models](#).
- Cheng Niu, Xingguang Wang, Xuxin Cheng, Juntao Song, and Tong Zhang. 2024. [Enhancing dialogue state tracking models through LLM-backed user-agents simulation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8724–8741, Bangkok, Thailand. Association for Computational Linguistics.
- OpenAI. 2024. [Gpt-4o mini: advancing cost-efficient intelligence](#).
- Jane Pan, Ryan Shar, Jacob Pfau, Ameet Talwalkar, He He, and Valerie Chen. 2025. [When benchmarks talk: Re-evaluating code llms with interactive feedback](#).
- Daniel Philipov, Vardhan Dongre, gokhan tur, and Dilek Hakkani Tur. 2024. [Simulating user agents for embodied conversational AI](#). In *NeurIPS 2024 Workshop on Open-World Agents*.

- Olivier Pietquin and Thierry Dutoit. 2006. A probabilistic framework for dialog simulation and optimal strategy learning. *IEEE Transactions on Audio, Speech, and Language Processing*, 14(2):589–599.
- Olivier Pietquin and Helen Hastie. 2013. [A survey on metrics for the evaluation of user simulations](#). *The Knowledge Engineering Review*, 28(1):59–73.
- Akshara Prabhakar, Zuxin Liu, Ming Zhu, Jianguo Zhang, Tulika Awalgaonkar, Shiyu Wang, Zhiwei Liu, Haolin Chen, Thai Hoang, Juan Carlos Niebles, Shelby Heinecke, Weiran Yao, Huan Wang, Silvio Savarese, and Caiming Xiong. 2025. [Apigen-mt: Agentic pipeline for multi-turn data generation via simulated agent-human interplay](#).
- Cheng Qian, Emre Can Acikgoz, Qi He, Hongru Wang, Xiushi Chen, Dilek Hakkani-Tür, Gokhan Tur, and Heng Ji. 2025. [Toolrl: Reward is all tool learning needs](#).
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2025. [Qwen2.5 technical report](#).
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2020. Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 8689–8696.
- 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, Rochester, New York. Association for Computational Linguistics.
- Jost Schatzmann, Karl Weilhammer, Matt Stuttle, and Steve Young. 2006. A survey of statistical user simulation techniques for reinforcement-learning of dialogue management strategies. *The knowledge engineering review*, 21(2):97–126.
- Jost Schatzmann and Steve Young. 2009. The hidden agenda user simulation model. *IEEE transactions on audio, speech, and language processing*, 17(4):733–747.
- Konrad Scheffler and Steve Young. 2002. Automatic learning of dialogue strategy using dialogue simulation and reinforcement learning. In *Proceedings of HLT*, volume 2, page 0.
- Ivan Sekulić, Silvia Terragni, Victor Guimarães, Nghia Khau, Bruna Guedes, Modestas Filipavicius, André Ferreira Manso, and Roland Mathis. 2024. Reliable llm-based user simulator for task-oriented dialogue systems. *arXiv preprint arXiv:2402.13374*.
- Pararth Shah, Dilek Hakkani-Tür, Bing Liu, and Gokhan Tür. 2018. [Bootstrapping a neural conversational agent with dialogue self-play, crowdsourcing and on-line reinforcement learning](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers)*, pages 41–51, New Orleans - Louisiana. Association for Computational Linguistics.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#).
- Jeonghoon Shim, Gyuhyeon Seo, Cheongsu Lim, and Yohan Jo. 2025. [Tooldial: Multi-turn dialogue generation method for tool-augmented language models](#).
- David Silver and Richard S Sutton. 2025. Welcome to the era of experience. *Google AI*, 1.

- Joar Skalse, Nikolaus Howe, Dmitrii Krasheninikov, and David Krueger. 2022. Defining and characterizing reward gaming. *Advances in Neural Information Processing Systems*, 35:9460–9471.
- Weiwei Sun, Shuyu Guo, Shuo Zhang, Pengjie Ren, Zhumin Chen, Maarten de Rijke, and Zhaochun Ren. 2023. Metaphorical user simulators for evaluating task-oriented dialogue systems. *ACM Transactions on Information Systems*, 42(1):1–29.
- Jianheng Tang, Tiancheng Zhao, Chenyan Xiong, Xiaodan Liang, Eric Xing, and Zhiting Hu. 2019. [Target-guided open-domain conversation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5624–5634, Florence, Italy. Association for Computational Linguistics.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, Gaël Liu, Francesco Visin, Kathleen Kenealy, Lucas Beyer, Xiaohai Zhai, Anton Tsitsulin, Robert Busa-Fekete, Alex Feng, Noveen Sachdeva, Benjamin Coleman, Yi Gao, Basil Mustafa, Iain Barr, Emilio Parisotto, David Tian, Matan Eyal, Colin Cherry, Jan-Thorsten Peter, Danila Sinopalnikov, Surya Bhupatiraju, Rishabh Agarwal, Mehran Kazemi, Dan Malkin, Ravin Kumar, David Vilar, Idan Brusilovsky, Jiaming Luo, Andreas Steiner, Abe Friesen, Abhan-shu Sharma, Abheesht Sharma, Adi Mayrav Gilady, Adrian Goedeckemeyer, Alaa Saade, Alex Feng, Alexander Kolesnikov, Alexei Bendebury, Alvin Abdagic, Amit Vadi, András György, André Susano Pinto, Anil Das, Ankur Bapna, Antoine Miech, Antoine Yang, Antonia Paterson, Ashish Shenoy, Ayan Chakrabarti, Bilal Piot, Bo Wu, Bobak Shahriari, Bryce Petrini, Charlie Chen, Charline Le Lan, Christopher A. Choquette-Choo, CJ Carey, Cormac Brick, Daniel Deutsch, Danielle Eisenbud, Dee Cattle, Derek Cheng, Dimitris Paparas, Divyashree Shivakumar Sreepathihalli, Doug Reid, Dustin Tran, Dustin Zelle, Eric Noland, Erwin Huizenga, Eugene Kharitonov, Frederick Liu, Gagik Amirkhanyan, Glenn Cameron, Hadi Hashemi, Hanna Klimczak-Plucińska, Harman Singh, Harsh Mehta, Harshal Tushar Lehri, Hussein Hazimeh, Ian Ballantyne, Idan Szpektor, Ivan Nardini, Jean Pouget-Abadie, Jetha Chan, Joe Stanton, John Wieting, Jonathan Lai, Jordi Orbay, Joseph Fernandez, Josh Newlan, Ju yeong Ji, Jyotinder Singh, Kat Black, Kathy Yu, Kevin Hui, Kiran Vodrahalli, Klaus Greff, Linhai Qiu, Marcella Valentine, Marina Coelho, Marvin Ritter, Matt Hoffman, Matthew Watson, Mayank Chaturvedi, Michael Moynihan, Min Ma, Nabila Babar, Natasha Noy, Nathan Byrd, Nick Roy, Nikola Momchev, Nilay Chauhan, Noveen Sachdeva, Oskar Bunyan, Pankil Botarda, Paul Caron, Paul Kishan Rubenstein, Phil Culliton, Philipp Schmid, Pier Giuseppe Sessa, Pingmei Xu, Piotr Stanczyk, Pouya Tafti, Rakesh Shivanna, Renjie Wu, Renke Pan, Reza Rokni, Rob Willoughby, Rohith Vallu, Ryan Mullins, Sammy Jerome, Sara Smoot, Sertan Girgin, Shariq Iqbal, Shashir Reddy, Shruti Sheth, Siim Pöder, Sijal Bhatnagar, Sindhu Raghuram Panyam, Sivan Eiger, Susan Zhang, Tianqi Liu, Trevor Yacovone, Tyler Liechty, Uday Kalra, Utku Evci, Vedant Misra, Vincent Roseberry, Vlad Feinberg, Vlad Kolesnikov, Woohyun Han, Woosuk Kwon, Xi Chen, Yinlam Chow, Yuvein Zhu, Zichuan Wei, Zoltan Egyed, Victor Cotruta, Minh Giang, Phoebe Kirk, Anand Rao, Kat Black, Nabila Babar, Jessica Lo, Erica Moreira, Luiz Gustavo Martins, Omar Sanseviero, Lucas Gonzalez, Zach Gleicher, Tris Warkentin, Vahab Mirrokni, Evan Senter, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, Yossi Matias, D. Sculley, Slav Petrov, Noah Fiedel, Noam Shazeer, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Jean-Baptiste Alayrac, Rohan Anil, Dmitry, Lepikhin, Sebastian Borgeaud, Olivier Bachem, Armand Joulin, Alek Andreev, Cassidy Hardin, Robert Dadashi, and Léonard Hussenot. 2025. [Gemma 3 technical report](#).
- Trong Wu. 1993. An accurate computation of the hypergeometric distribution function. *ACM Transactions on Mathematical Software (TOMS)*, 19(1):33–43.

- Wenquan Wu, Zhen Guo, Xiangyang Zhou, Hua Wu, Xiyuan Zhang, Rongzhong Lian, and Haifeng Wang. 2019. [Proactive human-machine conversation with explicit conversation goal](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3794–3804, Florence, Italy. Association for Computational Linguistics.
- Ruixuan Xiao, Wentao Ma, Ke Wang, Yuchuan Wu, Junbo Zhao, Haobo Wang, Fei Huang, and Yongbin Li. 2024. [FlowBench: Revisiting and benchmarking workflow-guided planning for LLM-based agents](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 10883–10900, Miami, Florida, USA. Association for Computational Linguistics.
- Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik Narasimhan. 2024. [\$\tau\$ -bench: A benchmark for tool-agent-user interaction in real-world domains](#).
- Fanghua Ye, Jarana Manotumruksa, and Emine Yilmaz. 2022. [Multiwoz 2.4: A multi-domain task-oriented dialogue dataset with essential annotation corrections to improve state tracking evaluation](#).
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2019. [Bertscore: Evaluating text generation with bert](#). *ArXiv*, abs/1904.09675.
- Zheng Zhang, Ryuichi Takanobu, Qi Zhu, Min-Lie Huang, and XiaoYan Zhu. 2020. Recent advances and challenges in task-oriented dialog systems. *Science China Technological Sciences*, 63(10):2011–2027.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. [Learning discourse-level diversity for neural dialog models using conditional variational autoencoders](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 654–664, Vancouver, Canada. Association for Computational Linguistics.
- Zhonghua Zheng, Lizi Liao, Yang Deng, Ee-Peng Lim, Minlie Huang, and Liqiang Nie. 2024. [Thoughts to target: Enhance planning for target-driven conversation](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 21108–21124, Miami, Florida, USA. Association for Computational Linguistics.

A System Prompts

Prompt A.1: Agent Prompt

You are an advanced agent specializing in conversational dialogues. You can act both as an agent (providing services) and a user (interacting with the database) to assist users in completing complex tasks.

Each task may involve multiple sub-tasks, such as finding restaurants, making reservations, booking hotels, locating attractions, and arranging transportation by checking for trains and buying train tickets.

Task Information:

- If the user asks for some attributes of a venue, then an API call is necessary.
- If you decide that more than one API calls are needed, you should call one API first and wait for the API result. After obtaining that result, you may think and call the next API or think and make a response.
- If you decide that there is an API input slot that the user has never mentioned, please put "any" as the slot value as a placeholder.

Objective:

- Ensure that each assistant utterance follows logical reasoning, determining whether an API call is needed and structuring the output accordingly. item If there are too many results returned by API results from database, you should ask the user for more constraints unless the user explicitly wants you to pick one or some.

General Procedure You should follow the workflow below every time you receive an input:

1. Use APIs to retrieve information from the database, make operations to the backend, or use tools to perform other tasks.
2. Repeat step 1 until you decide to make a response to the user.
3. Use the response API to return the response to the user and end your turn.

Prompt A.2: User Simulator Prompt

You are a user simulator interacting with an agent. You are provided with User Goals, and your task is to interact with the agent to achieve the goals as much as possible.

Conversation Guidelines:

- Work toward achieving specified user goals while maintaining a natural conversation flow.
- Reveal information gradually and naturally as a real person would.
- Do not hallucinate information that is not provided in the instruction. If asked about something not specified, respond honestly and tell the agent you do not have that information.
- User natural language that is appropriate to the context and your assigned personality, rather than copying the exact wording from the instructions.

Conversation Length:

- The maximum conversation length is 20 messages. Plan your approach accordingly to accomplish your goals within this limit.
- If you've achieved your goals before reaching the maximum length, you can end the conversation early. To do so, you need to send "Terminate." in a separate, final message.
 - Make sure to allow the agent to respond before sending "Terminate.".
 - NEVER include "Terminate." in a message where you are asking a question or making a request to the agent, because it will end the conversation before the agent can respond.
 - Do NOT mention "Terminate." at any other time, because that will end the conversation prematurely.
- Do not hallucinate information that is not provided in the instruction. If asked about something not specified, respond honestly and tell the agent you do not have that information.
- User natural language that is appropriate to the context and your assigned personality, rather than copying the exact wording from the instructions.

User Goals:

{user_goals}

B Multiwoz User Goal Generation

We construct the MultiWOZ Challenge dataset to evaluate user simulator goal alignment. In this section, we present our user goal generation pipeline that we used to create this dataset.

Step 1: Task Objective Generation

For the attraction, hotel, restaurant, and train domains in the MultiWOZ database, we iterate over the entities. Each entity has a set of key-value pairs (e.g. "area= east", "pricerange= moderate"). To instantiate a single task objective, we: (1) randomly sample requirement/preference keys to use their values to specify requirements or preferences about the entity the user is looking for, and (2) randomly sample request keys that will be used to specify information that the user wants about the entity. We provide this information to GPT-4o mini, and generate a user goal in natural language.

For example, given the attraction below:

```
{ "name": "abbey pool and astroturf pitch",  
  "type": "swimmingpool",  
  "area": "east",  
  "price range": "moderate",  
  "address": "pool way ...",  
  "phone": "01223 902088",  
  ... }
```

we select the area and price range for requirement/preference keys, and select address and phone for request keys. The task objective would like like:

"Find a swimming pool attraction in the east area. You prefer one with a moderate price range. You want to know its address and phone number."

Impossible Task Objectives. To also observe how user simulators perform under failure conditions, we repeat the process above, but assign a random value to the requirement/preference keys. This helps create an entity that does not exist in our database.

Conditional Task Objectives. We also include conditional task objectives. A conditional task objective will require the user to search for one specific entity, but if it satisfies some conditional requirement, the user must look for another entity. For example:

"You are looking for a swimming pool with a cheap price range. If it is in the east area, then change your mind and look for another swimming pool in the west area that has a moderate price range."

Step 2: User Profile and Policy Generation

Next, we automatically generate a diverse set of about 50 user profiles and 50 user policies with GPT-4o mini. An example a user profile looks like:

"You are Maria García, a travel blogger documenting her adventures across Southeast Asia. Your experiences focus on local cuisine and cultural festivals."

An example a user policy looks like:

"Always ask the agent to restate booking details (date, time, price) before confirming. Verify critical details by repeating them back and asking 'Is that correct?'"

Then, we conduct manual annotations of the user profiles and policies, where we confirmed the relevance and quality. When necessary, some elements were removed or modified.

Step 3: User Goal Generation

Finally, we generate a user goal by randomly select a user profile, user policy and multiple task objectives. These were combined into a single user goal. We conducted more manual annotations at this step, to ensure the quality of user goals.

C User Goal State Generation

Our user goal state generation process consists of three main steps. First, we decompose each user goal into distinct, modular sub-components, where each represents an independent, self-contained part of the original goal. Second, we classify each sub-component into one of five categories: user profile, user

policy, task objective, requirement or preference. Finally, we assign each sub-component its default value as described in Section 4.2.

For the decomposition and classification steps, we evaluated several LLMs, including DeepSeek-V3, Llama-3.3-70B-Instruct, GPT 4o-mini, and GPT 4o. Through manual validation, we determined that GPT-4o provided the highest quality results, and therefore used it to generate all the user goal states. The complete prompt used for this process is provided below.

Prompt C.1: Sub-Component Decomposition Prompt

You are an expert annotator tasked with dividing a user's goal into sub-components.

User Goal:
{user_goal}

A user goal can be split into multiple sub-components. Each sub-component can be classified into one of the following categories:

- User Profile: persona, background, contextual information (e.g. You are a software engineer with 5 years of experience)
- User Policy: general constraints and guidelines that must be consistently followed throughout the conversation (e.g. Always go with the cheapest option, don't share personal information)
- Task Objectives: tasks that must be accomplished (e.g. Book a flight)
 - Requirements: conditions and constraints associated with this task objective (e.g. Purchase 3 carry-on bags)
 - Preferences: user preferences associated with this task objective (e.g. You prefer a direct flight, you'd like an aisle seat)

Analyze the user goal and divide it into sub-components:

- Maintain the same wording as in the original goal when possible
- Feel free to split up a sentence into multiple sub-components if it makes sense or if they fit in different categories
- Ensure that no part of the original goal is left out
- It is okay if a category is empty, if no sub-components are relevant

Output Format

```
{
  "user_profile": [
    # List of user profile sub-components
  ],
  "user_policy": [
    # List of user policy sub-components
  ],
  "task_objectives": [ # List of task objectives
    {
      "task_objective": "task objective description",
      "requirements": [
        # List requirements for this task
      ],
      "preferences": [
        # List of preferences for this task
      ]
    }
  ]
}
```

Output only a JSON object, in the format specified.

D Sub-Component Status Update Prompt

Prompt D.1: User Profile Sub-Component Status Update Prompt

You are an expert evaluator analyzing a user simulator's behaviour. The user simulator is given a user profile component, which specifies an aspect of their persona, background, or contextual information. Your task is to evaluate the latest turn to determine if the user simulator behavior is aligned with this user profile component.

```
# User Profile Component:
{user_profile_component}
# Conversation History:
{conversation_history[: -2]}
```

```
# Latest Turn:
{conversation_history[-2: -1]}
```

```
# Evaluation Criteria:
```

Evaluate the latest turn, and determine if the user simulator is aligned with the user profile component ({user_profile_component}). You must provide a status, the value of which should be:

- **ALIGNED** - Choose if ALL these conditions are met:
 1. The user simulator behavior does NOT directly contradict the user profile component
 2. IF the context clearly calls for demonstrating this profile component, then the user simulator demonstrates it appropriately (explicitly or implicitly)
 3. IF the context does NOT clearly call for demonstrating this profile component (neither demonstrating nor contradicting)
- **MISALIGNED** - Choose if ANY of these conditions are met:
 1. The user simulator behavior DIRECTLY contradicts the user profile component
 2. The context clearly and obviously calls for demonstrating the profile component, AND the user simulator fails to do so in a way that would be unrealistic for someone with this profile

```
# Important Guidelines:
```

- Focus on the latest turn, but use the conversation history for context
- Unless there's a clear contradiction, or obvious missed opportunity, lean towards **ALIGNED**
 - Most turns will not require demonstrating the profile component

A user goal can be split into multiple sub-components. Each sub-component can be classified into one of the following categories:

- **User Profile:** persona, background, contextual information (e.g. You are a software engineer with 5 years of experience)
- **User Policy:** general constraints and guidelines that must be consistently followed throughout the conversation (e.g. Always go with the cheapest option, don't share personal information)
- **Task Objectives:** tasks that must be accomplished (e.g. Book a flight)
 - **Requirements:** conditions and constraints associated with this task objective (e.g. Purchase 3 carry-on bags)
 - **Preferences:** user preferences associated with this task objective (e.g. You prefer a direct flight, you'd like an aisle seat)

```
# Output Format:
```

```
{{
  "status": # ALIGNED | MISALIGNED,
  "reasoning": # Brief explanation of your decision
}}
```

Output a properly formatted JSON response, as specified by the Output Format, and nothing else.

E Inference-Time Steering with User Goal Experiments

Model	Prof.	Pol.	T.O.	Req.	Pref.	Avg
Prompt-Based						
Qwen-2.5-7B-It	88.3	49.2	94.3	96.0	85.7	82.7
Llama-3.1-8B-It	90.9	41.0	97.1	<u>99.0</u>	81.0	81.8
Gemma-3-27B-It	98.7	59.0	97.1	97.0	78.6	86.1
Qwen-2.5-72B-It	89.3	59.6	94.1	96.9	78.6	83.7
Llama-3.3-70B-It	<u>97.4</u>	<u>65.6</u>	98.1	<u>99.0</u>	92.9	90.6
Inference-Time Steering with User Goal						
Qwen-2.5-7B-It	70.1	54.2	98.1	<u>99.0</u>	90.5	82.4
Llama-3.1-8B-It	87.0	60.7	97.1	98.0	100.0	88.6
Gemma-3-27B-It	92.2	73.8	96.2	98.0	73.8	86.8
Qwen-2.5-72B-It	93.5	62.3	92.4	94.1	76.2	83.7
Llama-3.3-70B-It	98.7	59.0	96.2	98.0	90.9	88.6
Inference-Time Steering with User Goal States						
Qwen-2.5-7B-It	77.9	55.7	96.2	98.0	92.9	84.1
Llama-3.1-8B-It	92.2	57.4	93.3	98.0	95.2	87.2
Gemma-3-27B-It	94.8	62.3	92.4	97.0	85.7	86.4
Qwen-2.5-72B-It	93.3	54.4	<u>98.0</u>	97.9	<u>97.6</u>	88.3
Llama-3.3-70B-It	93.5	63.9	98.1	100.0	<u>97.6</u>	90.6

(a) τ -Bench Airline

Model	Prof.	Pol.	T.O.	Req.	Pref.	Avg
Prompt-Based						
Qwen-2.5-7B-It	82.0	40.8	96.0	<u>99.5</u>	91.7	82.0
Llama-3.1-8B-It	84.8	36.8	97.1	98.9	96.3	82.8
Gemma-3-27B-It	91.0	39.5	97.8	<u>99.5</u>	96.3	84.8
Qwen-2.5-72B-It	88.6	40.8	98.2	<u>97.9</u>	91.7	83.4
Llama-3.3-70B-It	91.7	46.1	99.6	100.0	100.0	87.5
Inference-Time Steering with User Goal						
Qwen-2.5-7B-It	73.4	38.2	97.5	98.4	95.4	80.6
Llama-3.1-8B-It	91.7	38.2	<u>98.9</u>	<u>99.5</u>	94.4	84.5
Gemma-3-27B-It	94.5	63.2	97.8	98.4	91.7	<u>89.1</u>
Qwen-2.5-72B-It	89.6	40.8	97.1	98.9	97.2	84.7
Llama-3.3-70B-It	92.4	55.3	98.2	100.0	<u>98.1</u>	88.8
Inference-Time Steering with User Goal States						
Qwen-2.5-7B-It	73.0	47.4	94.9	97.9	93.5	81.4
Llama-3.1-8B-It	84.8	52.6	95.3	98.9	97.2	85.8
Gemma-3-27B-It	94.5	56.6	97.1	<u>99.5</u>	89.8	87.5
Qwen-2.5-72B-It	85.8	48.7	97.1	98.4	<u>98.1</u>	85.6
Llama-3.3-70B-It	<u>92.7</u>	<u>59.2</u>	96.0	98.4	100.0	89.3

(b) τ -Bench Retail

Model	Prof.	Pol.	T.O.	Req.	Pref.	Avg
Prompt-Based						
Qwen-2.5-7B-It	54.7	18.0	78.4	81.9	73.5	61.3
Llama-3.1-8B-It	73.6	35.0	86.1	89.0	80.9	72.9
Gemma-3-27B-It	<u>80.1</u>	50.3	95.0	97.7	92.1	83.0
Qwen-2.5-72B-It	72.0	31.0	91.0	93.3	90.2	75.5
Llama-3.3-70B-It	72.0	31.0	97.9	98.6	93.2	<u>83.3</u>
Inference-Time Steering with User Goal						
Qwen-2.5-7B-It	49.8	28.0	77.8	83.2	69.4	61.7
Llama-3.1-8B-It	73.9	49.7	92.9	95.8	80.4	78.5
Gemma-3-27B-It	80.8	49.0	95.9	<u>98.3</u>	89.4	82.7
Qwen-2.5-72B-It	71.7	41.7	92.7	94.9	<u>93.1</u>	78.8
Llama-3.3-70B-It	73.6	<u>60.3</u>	94.9	95.8	91.0	83.1
Inference-Time Steering with User Goal States						
Qwen-2.5-7B-It	43.4	26.3	88.1	89.8	71.1	63.7
Llama-3.1-8B-It	71.7	38.0	85.2	88.6	74.3	71.6
Gemma-3-27B-It	72.3	45.7	87.0	89.7	74.8	73.9
Qwen-2.5-72B-It	71.0	46.3	94.9	96.9	92.9	80.4
Llama-3.3-70B-It	72.6	64.3	<u>97.0</u>	97.8	91.7	84.7

(c) MultiWOZ Challenge

Table 7: User simulator goal alignment performance with prompt-based, inference-time steering with user goal and inference-time steering with user goal states. The table shows the average success rates for User Profile (Prof.), User Policy (Pol.), Task Objective (T.O.), Requirements (Req.) and Preferences (Pref.) sub-component categories, and the overall average performance. Highest scores are **bolded**, and the second-highest scores are underlined.

F Naturalness and Coherence Evaluation Prompts

Prompt F.1: Naturalness Evaluation Prompt

You are an expert evaluator analyzing a user simulator's behavior while they are pursuing a user goal. Your task is to evaluate the naturalness of the user simulator throughout conversation.

User Goal
{user_goal}

Conversation History
{conversation_string}

Evaluation Task
Rate the naturalness of the user responses on a scale from 1 to 5, based on three key dimensions:

1. **Grammar and coherence**: Grammatically correct sentences that flow logically and coherently.
2. **Context Relevance**: Responses are strongly related to the conversation context and previous turns, as well as the user goal.
3. **Conversational Style**: Responses are natural, and appropriate for the conversation style.

The rating scale is as follows:

- 5 (Highly Natural): Excels in all three dimensions - perfect grammar/coreference, fully contextually appropriate, and natural conversational
- 4 (Mostly Natural): Strong in all dimensions with minor lapses - occasional awkward phrasing or slightly unnatural language
- 3 (Moderately Natural): Adequate in all dimensions - functional but noticeably artificial in style or occasional context mismatches
- 2 (Somewhat Unnatural): Weak in 1-2 dimensions - grammar errors, loses context, or overly formal/robotic style
- 1 (Very Unnatural): Fails multiple dimensions - poor grammar, irrelevant responses, or mechanical/repetitive patterns

Output Format:
{
 "naturalness_score": # number from 1 to 5
 "reasoning": # Brief explanation of your evaluation decision
}

Output a properly formatted JSON response, as specified by the Output Format.

Prompt F.2: Coherence Evaluation Prompt

You are an expert evaluator analyzing a user simulator's behavior while they are pursuing a user goal. Your task is to evaluate the coherence of a user simulator in pursuing their goal in a conversation.

User Goal
{user_goal}

Conversation History
{conversation_string}

Evaluation Task
Rate the coherence of the user's dialogue on a scale from 1 to 5, based on three key dimensions:

1. **Goal Progression**: User consistently works toward their stated goal with

appropriate persistence and adaptation

2. **Topic Continuity**: Responses maintain topical relevance and logical connections between turns
3. **Response Appropriateness**: User responds appropriately to system prompts, questions, and clarification requests

The rating scale is as follows:

- 5 (Highly Coherent): Clear goal pursuit throughout, all utterances topically connected, appropriate responses to all system turns
- 4 (Mostly Coherent): Strong goal focus with minor deviations, well-connected utterances, occasionally misses response opportunities
- 3 (Moderately Coherent): Generally pursues goal with some inconsistencies, mostly connected utterances, some inappropriate responses
- 2 (Somewhat Incoherent): Weak goal focus with major deviations, frequent topic jumps, often misses or misinterprets prompts
- 1 (Very Incoherent): No clear goal pursuit, disconnected responses, fails to engage appropriately with system

Output Format:

```
{{
  "coherence_score": # number from 1 to 5
  "reasoning": # Brief explanation of your evaluation decision
}}
```

Output a properly formatted JSON response, as specified by the Output Format.