

LLM Roleplay: Simulating Human-Chatbot Interaction

Hovhannes Tamoyan, Hendrik Schuff, and Iryna Gurevych

Ubiquitous Knowledge Processing Lab (UKP Lab)
Department of Computer Science and Hessian Center for AI (hessian.AI)
Technical University of Darmstadt
www.ukp.tu-darmstadt.de

Abstract

The development of chatbots requires collecting a large number of human-chatbot dialogues to reflect the breadth of users’ sociodemographic backgrounds and conversational goals. However, the resource requirements to conduct the respective user studies can be prohibitively high and often only allow for a narrow analysis of specific dialogue goals and participant demographics. In this paper, we propose LLM Roleplay, the first comprehensive method integrating multi-turn human-chatbot interaction simulation, explicit persona construction from sociodemographic traits, goal-driven dialogue planning, and robust handling of conversational failures, enabling broad utility and reliable dialogue generation. To validate our method, we collect natural human-chatbot dialogues from different sociodemographic groups and conduct a user study to compare these with our generated dialogues. We evaluate the capabilities of state-of-the-art LLMs in maintaining a conversation during their embodiment of a specific persona and find that our method can simulate human-chatbot dialogues with a high indistinguishability rate.¹

1 Introduction

Collecting human-chatbot dialogues requires recruiting and managing a large number of human annotators, which can pose prohibitive obstacles to researchers who aim to develop conversational AI agents (i.e., chatbots). To circumvent the latter and the limitations of publicly available data, numerous methods employing chatbots to generate dialogues have been introduced lately (Xu et al., 2023b; Kim et al., 2023; Zhu et al., 2023; Ding et al., 2023; Svikhnushina and Pu, 2023; Zhao et al., 2024b). These methods can generate dialogues much faster and more cost-effectively while still approximating the quality and variety of human annotators (Zhang et al., 2024).

¹<https://github.com/UKPLab/llm-roleplay>

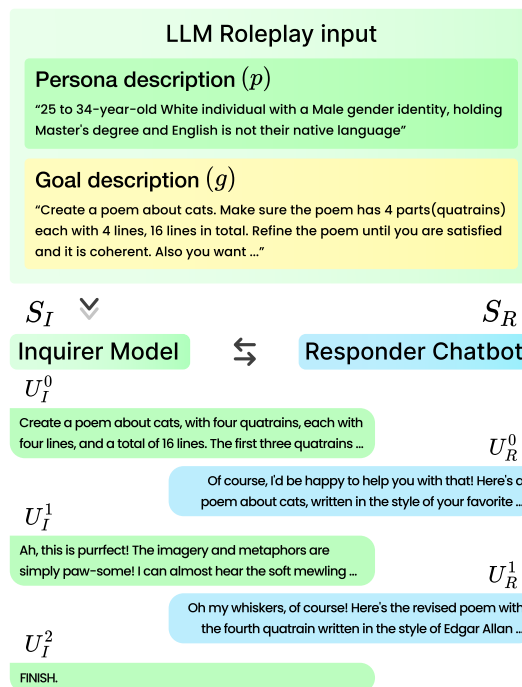


Figure 1: Schematic illustration of our method: A textual description of a persona and a goal (top) is used to instruct the inquirer (S_I) model to embody the given persona (left) and engage in a dialogue with the responder (S_R) chatbot (right). We show that dialogues simulated by the inquirer LLM and the responder chatbot can effectively simulate human-chatbot interaction.

Recently, several studies have leveraged such synthetic data generated through distillation and self-improvement techniques to enhance the capabilities of large language models (LLMs) (Zhang et al., 2024). By employing instruction tuning and fine-tuning techniques, these efforts aim to ensure that LLMs generate useful and safe responses that adhere to provided instructions (Xu et al., 2023a; Taori et al., 2023; Mukherjee et al., 2023; Li et al., 2023b; Peng et al., 2023; Almazrouei et al., 2023).

However, current dialogue generation methods have two critical limitations. First, the set of possible user responses for a given chatbot utterance

is large while the existing datasets typically only cover a single or few annotated user responses. Previously generated datasets are constrained by the number of dialogues and turns, making it difficult to extend them to novel domains or conversational goals. Second, existing methods rarely account for subjective response behavior and implicitly assume an average user. This can create a gap in representation (Rottger et al., 2022) and pose the risk of evaluating and improving chatbots for a specific sociodemographic group while neglecting others.

The quality of existing datasets is primarily assessed through model performance in a supervised fine-tuning (SFT) setup, which is influenced by various properties of the generated dataset. For example, Chen et al. (2023) argue that the initial prompt plays a crucial role in the quality of the generated dialogue. On the other hand, Zhao et al. (2024a) demonstrate that the length of the response is a key factor, and with significantly fewer samples, it outperforms methods that focus solely on quality. Meanwhile, Shen (2024) emphasizes the importance of mimicking human-style interactions in the dialogues.

To mitigate the shortcomings of existing dialogue generation methods and respond to the needs identified in prior work for SFT, we introduce LLM Roleplay, the first comprehensive method integrating multi-turn human-chatbot interaction simulation, explicit persona construction from sociodemographic traits, goal-driven dialogue planning, and robust handling of conversational failures, enabling broad utility and reliable dialogue generation. Our method is the first to instruct an LLM (inquirer) to adopt a specific persona and prompt a chatbot (responder) to achieve a given conversational goal, thereby eliciting realistic human-AI interactions with a particular LLM. Figure 1 illustrates an exemplary application of our method.

In this work, we address the following two research questions: (i) To what extent can we simulate real human-chatbot dialogues using LLM-chatbot dialogues? (ii) How do various LLMs perform as inquirers within this setup? To answer those research questions, we conduct two user studies. First, we collect real human-chatbot dialogues by asking participants to reach various conversational goals and link the collected dialogues with participants’ sociodemographic backgrounds. Next, we use LLM Roleplay to simulate dialogues with the same set of personas and goals. Second, we conduct a human evaluation with another pool

of participants to assess how well the generated dialogues mimic the collected ones. Specifically, we present participants with two dialogues: one collected from the first study and one simulated using LLM Roleplay, both involving the same persona and the conversation goal.

We then ask the participants to identify which of the dialogues was simulated. We find that LLM Roleplay approximates natural human-chatbot dialogues with a high level of indistinguishability.

Overall, this paper contributes: (i) a novel method to simulate human-chatbot dialogues with an arbitrary choice of personas and conversational goals; (ii) a human evaluation confirming our method’s potential to closely resemble real dialogues; (iii) a dataset of goal-oriented human-chatbot and model-chatbot dialogues using our hand-crafted multi-hop goals involving four state-of-the-art LLMs, (iv) an in-depth comparison of open-source and proprietary LLMs in maintaining conversations while embodying specific personas, (v) an open-source implementation of our plug-and-play method, readily applicable for simulating dialogues with any combination of model and chatbot across various conversational goals and personas.

2 Large Language Model Roleplay

In the following, we introduce the notation that we are following throughout this paper. Refer to Figure 1 for an annotated example.

Persona (\mathcal{P}) is the composition of sociodemographic features. We conceptualize persona as a written description of an individual from a specific sociodemographic group. We follow Kumar et al. (2021) to define the sociodemographic features and the options of the latter.

Goal (\mathcal{G}) is the textual representation of the conversational goal.

Subject (\mathcal{S}) is a dialogue participant. We denote the **inquirer** as $\mathcal{S}_{\mathcal{I}}$, the entity that asks questions, and the **responder** as $\mathcal{S}_{\mathcal{R}}$, the entity that answers the given questions. In our setup, the inquirer ($\mathcal{S}_{\mathcal{I}}$) is an LLM and the responder ($\mathcal{S}_{\mathcal{R}}$) is a conversational agent or chatbot (not necessarily an LLM). The human inquirer will be denoted as $\mathcal{S}_{\mathcal{I}}^h$.

Utterance (\mathcal{U}) is the output of a Subject (\mathcal{S}), i.e. either the output of the inquirer ($\mathcal{U}_{\mathcal{I}}$) or the responder ($\mathcal{U}_{\mathcal{R}}$). For example, $\mathcal{U}_{\mathcal{I}}^i$ will denote the inquirer’s i -th utterance. We refer to two consecu-

tive utterances of two different subjects as a turn: $\mathcal{T}^i = [\mathcal{U}_I^i; \mathcal{U}_R^i]$.

Dialogue (\mathcal{D}) is a sequence of one or more turns. The dialogue of t turns is denoted as: $\mathcal{D}^t := \{\mathcal{T}^0, \mathcal{T}^1, \dots, \mathcal{T}^t\}$. We denote the maximum number of turns by max_t , i.e. $t \leq max_t$ which we discuss in more detail in Section 3.1.

At a high level, LLM Roleplay consists of three main steps: (i) Initial conditioning of the inquirer model with persona-specific text alongside the goal description. (ii) Subsequently, the inquirer’s output is provided to the responder model. (iii) Lastly, the output of the responder is returned to the inquirer by asking it to either output a follow-up question or a pre-defined token to terminate the dialogue.

We begin with assembling prompting templates for both of the subjects. We create a system prompt template (SYS_I) for the inquirer LLM. This template prompts the LLM to embody the given persona, provide a prompt to address the designated goal and output the specified termination token if it considers the goal accomplished. As for the responder, we use a default system prompt (SYS_R) to promote it to be a helpful and honest assistant. Additionally, we develop a response forwarder template (INTER_I) for the inquirer. This template prompts the inquirer to assess the conclusiveness of the responder’s answer, determining whether to output a subsequent question or the termination token. We provide the system and forwarder prompt templates in Table 5 in Appendix B. The algorithm for the LLM Roleplay is shown in Algorithm 1.

To obtain a dialogue for a given persona and a goal we create a prompt by passing the persona (\mathcal{P}) and the goal (\mathcal{G}) to the inquirer system prompt template and generate an output based on inquirer LLM distribution:

$$\mathcal{U}_I^0 \sim \mathcal{S}_I(\text{SYS_I}(\mathcal{P}, \mathcal{G})). \quad (1)$$

A deterministic prompt extraction function, `extract_prompt`, that looks for a string in double quotes in the response, is applied to the latter to extract the prompt from the response:

$$\mathcal{U}_I^0 := \text{extract_prompt}(\mathcal{U}_I^0). \quad (2)$$

The responder receives the prompt of the inquirer, and generates an output:

$$\mathcal{U}_R^0 \sim \mathcal{S}_R(\mathcal{U}_I^0). \quad (3)$$

Given the output of the responder, we condition the inquirer using the output forwarder template:

$$\mathcal{U}_I^1 \sim \mathcal{S}_I([\mathcal{U}_I^0; \text{INTER_I}(\mathcal{U}_R^0)]). \quad (4)$$

If the output of the inquirer begins or ends with the termination token, the stopping condition, `stop` is met, and consequently, the algorithm terminates. Otherwise, the prompt is extracted using `extract_prompt`, and the process persists for the coming turns until it reaches the maximum number of turns: $t = max_t$. For the t -th turn, the utterances will be:

$$\mathcal{U}_I^t \sim \mathcal{S}_I(\mathcal{T}^0, \mathcal{T}^1, \dots, \mathcal{T}^{t-1}) \quad (5)$$

$$\mathcal{U}_R^t \sim \mathcal{S}_R(\mathcal{T}^0, \mathcal{T}^1, \dots, \mathcal{T}^{t-1}, \mathcal{U}_I^t) \quad (6)$$

Algorithm 1: LLM Roleplay

Input : \mathcal{P}, \mathcal{G} : persona and goal
Output : $\mathcal{D}(\mathcal{S}_I, \mathcal{S}_R)$

- 1 $\mathcal{D} \leftarrow \{\}$;
- 2 $\mathcal{U}_I^0 \leftarrow \mathcal{S}_I(\text{SYS_I}(\mathcal{P}, \mathcal{G}))$
- 3 $\mathcal{U}_I^0 := \text{extract_prompt}(\mathcal{U}_I^0)$
- 4 **if** $\mathcal{U}_I^0 = \emptyset$ **then break**;
- 5 $\mathcal{U}_R^0 \leftarrow \mathcal{S}_R(\mathcal{U}_I^0)$
- 6 **for** $t = 1 \rightarrow max_t$ **do**
- 7 $\mathcal{U}_I^t \leftarrow \mathcal{S}_I([\mathcal{U}_I^{t-1}; \text{INTER_I}(\mathcal{U}_R^{t-1})])$;
- 8 **if** `stop`(\mathcal{U}_I^t) **then**
- 9 | **break**;
- 10 **end**
- 11 $\mathcal{U}_I^t := \text{extract_prompt}(\mathcal{U}_I^t)$
- 12 **if** $\mathcal{U}_I^t = \emptyset$ **then**
- 13 | **break**;
- 14 **end**
- 15 $\mathcal{U}_R^t \leftarrow \mathcal{S}_R(\mathcal{D}, \mathcal{U}_I^t)$;
- 16 $\mathcal{D} \leftarrow \{\mathcal{D}, [\mathcal{U}_I^t; \mathcal{U}_R^t]\}$;
- 17 **end**

3 Experiments

In the first study, we collect dialogues between human inquirers and model responders ($\mathcal{D}(\mathcal{S}_I^h, \mathcal{S}_R)$) by engaging participants with various personas (\mathcal{P}) in interactions with a chat-tuned LLM intending to achieve a given goal (\mathcal{G}). Utilizing the same set of personas (\mathcal{P}) and goals (\mathcal{G}) we generate dialogues between model inquirers and the same responder ($\mathcal{D}(\mathcal{S}_I, \mathcal{S}_R)$) by employing the LLM Roleplay method (Section 3.1). We report the statistics of the generated dialogues such as the average number of turns per dialogue, the number of tokens per

	Llama-2	Mixtral	Vicuna	GPT4
Avg. # Turns per Dialogue	2.62 (1.54)	3.86 (2.19)	7.12 (3.79)	7.60 (3.08)
Avg. # Tokens per Prompt	77.77 (46.20)	50.82 (26.47)	75.19 (92.43)	68.13 (60.37)
Avg. # Tokens per Response	347.50 (142.14)	302.47 (151.30)	228.29 (163.92)	267.69 (147.26)
No-prompt	6.82% (1.48%)	0.97% (0.57%)	7.90% (0.54%)	0.17% (0.04%)
Multiple Prompts	8.79% (0.95%)	8.40% (1.52%)	6.08% (0.41%)	39.21% (2.38)
Incoherent Response	3.12% (0.30%)	0.13% (0.09%)	0.79% (0.10%)	0.03% (0.04%)
Number of Self-Replies	5.50% (3.25%)	5.99% (1.43)	69.39% (5.77%)	5.48% (0.09%)
Incoherent Response (Responder)	0.56% (0.57%)	1.16% (0.19%)	8.01% (0.93%)	7.50% (0.35%)

Table 1: Analysis of persona-specific dialogue collection (top) and failure cases (bottom) conducted for Llama-2, Mixtral, Vicuna, and GPT4. The results are averaged over runs with three different seeds. We show that GPT4 is successful at holding dialogues with longer utterances and having relatively fewer failure cases. Meanwhile, Mixtral is better at providing short on-point prompts. The number of utterances in dialogues is preferred to be larger, while for other metrics, smaller values are better. The standard deviation is indicated in parentheses.

prompt (inquirer output), and the number of tokens per response (responder output). Please refer to Appendix B for details on the generation parameters.

Furthermore, we investigate the failure cases of the inquirer LLMs, in following the given instruction. We depict the most common failure cases, report their statistics, and describe their detection mechanisms (Section 3.3).

In the second study, we conduct a human evaluation wherein another set of participants compare the natural $\mathcal{D}(\mathcal{S}_I^h, \mathcal{S}_R)$ and the simulated $\mathcal{D}(\mathcal{S}_I, \mathcal{S}_R)$ dialogues to discern the simulated counterpart $\mathcal{D}(\mathcal{S}_I, \mathcal{S}_R)$ (Section 3.4). We report the total and per-model undetectability rates, the utterance number on which the dialogue was detected, and the distribution of the duration and confidence choices. We define a dialogue as detected if a human evaluator correctly identifies it as artificial when presented with a real and a synthetic dialogue pair, and undetected otherwise. We leave task-success and factual-correctness evaluation for future work, as these metrics emphasize task-specific or knowledge-grounded objectives Sajjad et al. (2022), which diverge from our focus on assessing generation quality and variety through indistinguishability/undetectability and lexical diversity.

Moreover, we use generalized linear models to analyze the detection rate of the simulated dialogue, the detection utterances number, and the duration users spent making a choice to analyze how different inquirer LLMs behave.

For our experiments, we used a single NVIDIA A100 GPU with 80GB memory for Llama-2 and Vicuna. We utilized up to 92% of the memory.

3.1 Persona-Specific Dialogue Collection

For this study, participants were instructed to interact with Llama-2² (\mathcal{S}_R) to accomplish a designated goal. We additionally ask participants to provide sociodemographic information (\mathcal{P}) such as age group, gender, race, level of education, and whether they identify as native English speaker. We base the choice of sociodemographic features and the options of the features on previous work by Kumar et al. (2021). Please note that the method is not limited to these traits; future work can experiment with any set of persona features. We provide a detailed screenshot of the persona information form interface in Figure 9, and a screenshot of our chat interface in Figure 10 in the Appendix B. In order to cover different conversational goals, we design ten handcrafted multi-hop goals (\mathcal{G}) spanning three domains: "Math", "Coding", and "General Knowledge".

We conduct a study involving 20 participants each engaged in tackling 10 goals, resulting in the generation of 200 natural human-chatbot interaction dialogues ($\mathcal{D}(\mathcal{S}_I^h, \mathcal{S}_R)$). We provide the full sociodemographic distribution of the participants in Figure 3 to Figure 7 in Appendix B.

Subsequently, we generate dialogues utilizing our LLM Roleplay method with a set of three state-of-the-art LLMs and one proprietary conversational agent as inquirers (\mathcal{S}_I): llama-2-13B-Chat (Llama-2) (Touvron et al., 2023), Mixtral-8x7B-Instruct-v0.1 (Mixtral) (Jiang et al., 2024), vicuna-13b-v1.5-16k (Vicuna) (Peng et al., 2023), and GPT4 (OpenAI et al., 2024). See sample natural and gener-

²Due to limited deployment resources quantized Llama-2-13B-chat-GGUF is utilized (Touvron et al., 2023), employing the Q5_K_M quantization method with 5 bits (Frantar et al., 2023)

SD Information	TTR	dist-1	dist-2
None	0.281 (0.004)	0.284 (0.003)	0.625 (0.007)
Age	0.572 (0.010)	0.576 (0.010)	0.892 (0.007)
Race	0.582 (0.006)	0.587 (0.005)	0.880 (0.002)
Gender	0.411 (0.003)	0.415 (0.003)	0.747 (0.002)
Education	0.579 (0.015)	0.578 (0.016)	0.869 (0.005)
Is Native EN Speaker	0.394 (0.003)	0.394 (0.003)	0.721 (0.005)
All	0.606 (0.019)	0.605 (0.019)	0.872 (0.006)

Table 2: Ablation study on the impact of sociodemographic (SD) information on the lexical diversity of the simulated user’s language measured via mean Type-Token Ratio (TTR) and Distinct-N (dist-n) along with corresponding variance reported in parentheses. We observe enhanced lexical diversity with the addition of individual sociodemographic features.

ated dialogues using Mixtral inquirer and Llama-2 responder from Figure 14 to Figure 18 in Appendix E.

In total, the dialogue collection and generation results in the creation of 200 natural human-chatbot ($\mathcal{D}(\mathcal{S}_I^h, \mathcal{S}_R)$) and 800 (with 4 inquirers) simulated ($\mathcal{D}(\mathcal{S}_I, \mathcal{S}_R)$) dialogue pairs.

We observe that, on average, GPT-4 (OpenAI et al., 2024) is successful at holding longer dialogues with 7.60 turns on average (Table 1). Mixtral (Jiang et al., 2024) on the other hand is better at generating shorter and on-point (based on manual analysis) prompts with 50.82 tokens per prompt. The generated dialogues have an average of 5.30 turns, each consisting of an average of 67.97 tokens per prompt, resulting in longer dialogues than those from the previous work (Table 4).

3.2 Impact of Persona-Specific Information

We assess how adding sociodemographic information to the persona description that is provided to the model affects the lexical diversity of the simulated user utterances. We compare three scenarios: (a) using a baseline prompt with no sociodemographic information in the persona description, (b) adding single sociodemographic features, (c) and including all features as described in our method. Table 2 shows the respective Type-Token Ratio (TTR) (Zipf, 2013) and Distinct-N (dist-n) (Li et al., 2016) values quantifying lexical diversity for Mixtral inquirer. We observe that adding single sociodemographic features as well as combined sociodemographic features to the model prompt significantly increases lexical diversity.

3.3 Failure Cases

Since the outputs of the LLMs are free-form, applying deterministic functions to their output proves challenging, resulting in a theoretically infinite number of turns. Thus, there is a need for a set limit on the number of turns: max_t . Nevertheless, we try to capture the algorithm failure cases to have an automatic assessment of the inquirer LLM failure cases. We expect the inquirer LLM to provide the intended prompt enclosed in double quotes as we explicitly request this in our instructions. See a sample expected output and the failure cases examples in Table 8 in Appendix C. All of the LLMs in our experiments face the following issues.

Prompt not in double-quotes. The model fails to provide a prompt within double quotes, thus causing the dialogue to terminate.

Incoherent output. The responder produces a repetitive token sequence. We identify such outputs and preemptively end the dialogue. We utilize the incoherent function to analyze text for incoherent strings (see Algorithm 2 in Appendix C).

Incoherent output of responder. The inquirer model fails to detect the case when the responder outputs incoherent text, leading to unsuccessful dialogues. We spot such outputs and stop the dialog, using the same incoherent function.

Inquirer self-reply. The inquirer model fails to maintain its intended role and answers its own question, resulting in a fully generated dialogue in a single utterance. We detect this deterministically by examining the presence of any special tokens of the responder model within the output. For instance, in the case of Llama-2, this token appears as "[INST]", while for Vicuna, it manifests as "### Human:".

Multiple prompts. The inquirer outputs multiple strings enclosed in double quotes. As sometimes this overlaps with the previous case (inquirer self-reply), we select the first as the prompt.

Dialogue-stopping criterion failure. In the intermediate prompts, we ask the chatbot to output a pre-defined token when it "thinks" the goal assigned to it is achieved. Nonetheless, it follows a limited set of tokens; for instance, "FINISH" was utilized in our experiments.

We attribute these failures to two common issues in LLMs: limited context length and a restricted set of fine-tuned instructions (Kaddour et al., 2023).

Subset		Llama-2	Mixtral	Vicuna	GPT4
<i>total</i>	Undetectability Rate	33.5%	44.0%	22.5%	35.0%
	Confidence: "very confident"	33 (16.50%)	32 (16.00%)	75 (37.50%)	43 (21.50%)
	Confidence: "confident"	108 (54.00%)	83 (41.50%)	75 (37.50%)	97 (48.50%)
	Confidence: "somewhat confident"	59 (42.50%)	85 (25.00%)	50 (20.00%)	60 (30.00%)
<i>detected</i>	Duration	86.10 (157.14)	99.45 (98.81)	94.63 (150.70)	124.88 (227.36)
	Utterance Number	1.72 (0.78)	2.38 (1.35)	2.50 (1.84)	3.72 (2.35)
	Confidence: "very confident"	24 (12.00%)	19 (9.50%)	65 (32.50%)	32 (16.00%)
	Confidence: "confident"	80 (40.00%)	51 (25.50%)	59 (29.50%)	66 (33.00%)
	Confidence: "somewhat confident"	29 (14.50%)	42 (21.00%)	31 (15.50%)	32 (16.00%)
<i>undetected</i>	Duration	120.44 (139.93)	114.93 (157.46)	62.24 (62.37)	94.52 (124.03)
	Confidence: "very confident"	9 (4.50%)	13 (6.50%)	10 (5.00%)	11 (5.50%)
	Confidence: "confident"	28 (14.00%)	32 (16.00%)	16 (8.00%)	31 (15.50%)
	Confidence: "somewhat confident"	30 (15.00%)	43 (21.50%)	19 (9.50%)	28 (14.00%)

Table 3: Analysis of human-evaluation results for detected, undetected, and total dialogues for Llama-2, Mixtral, Vicuna, and GPT4, showing confidence statistics as occurrences (percentages in parentheses). We exhibit that Mixtral has the highest undetectability rate of 44% (**50%** reflecting absolute indistinguishability). Moreover, it has the lowest confidence choice statistics for not confidently detected and the highest statistics for confidently undetected dialogues. GPT4 has the highest utterance number: 3.72, showing that it is identified in later utterances. The duration (in seconds) and the utterance number standard deviations are in parentheses.

Toxic content detection. To proactively address the generation of potentially harmful content, we employ the Llama-2 Guard model (Team, 2024) to filter out toxic dialogues.

In most cases, GPT4 outperforms other models with lower failure rates: 0.17% for responses without prompts, 0.03% for incoherent responses, and 5.48% for self-replies (Table 1). However, it struggles with providing a single prompt, failing 39.21% of the time. In contrast, Vicuna excels in generating single prompts, achieving a failure rate of 6.08%. Llama-2 receives fewer incoherent outputs, with 0.56%. Mixtral’s performance is intermediate, showing more balanced results across different metrics. Additional plots and detailed data are provided in Appendix B. In these experiments, we detected no unsafe content.

3.4 Human-Evaluation

To answer the question of how well LLM inquirer and responder dialogues $\mathcal{D}(\mathcal{S}_I, \mathcal{S}_R)$ approximate dialogues between human inquirer and responder $\mathcal{D}(\mathcal{S}_I^h, \mathcal{S}_R)$, we conduct a human evaluation study. In each round, participants are shown two dialogues side by side, both featuring the same persona (\mathcal{P}) and solving the same goal (\mathcal{G}).

We allocated different sets of dialogue pairs for each participant, ensuring that no participant encounters multiple dialogues from the same user (from Section 3.1) solving the same goal. A new group of 20 participants was selected, with each participant tasked with reviewing 40 dialogue pairs.

The participants selected for this study represent a wide range of occupational backgrounds. Some have no prior experience with chatbots, while others are industry professionals.

Participants are required to answer three questions for each dialogue pair: (i) Identify which dialogues they perceive as artificial (simulated): Choices include "1st (left)", "2nd (right)", or "Not sure" the latter is considered to be a tie — i.e., the responses are indistinguishable to annotator. (ii) Express the level of confidence in their selection: Options are "Somewhat Confident," "Confident," and "Very Confident". (iii) Identify the specific utterance number that signifies the artificiality within the dialogue: Options depend on the number of utterances of the dialogue pairs. We provide a detailed view of the instructions provided to users in Figure 11 and the interfaces of the applications used for the study Figure 12 in Appendix D.

In conducting the human evaluation, we also track the amount of time participants take to respond to questions. This measure allows us to estimate a proxy measure of the complexity of the dialogue pairs. We assume that the longer it takes for a participant to make a decision, the more challenging the pair is, indicating that the simulated dialogue is difficult to discern.

The study shows that among the 800 samples, the simulated dialogues remained undetected in 33.75% (**50%** reflecting absolute indistinguishability) of the dialogues. Per model, statistics show that Mixtral has the highest undetectability rate of

44.0%, after which is GPT4 with 35.0% followed by Llama-2 and Vicuna with 33.5% and 22.5% respectively (Figure 2). See the per-confidence choice distribution in Figure 13 in Appendix D.

The highest utterance number on the detected set of dialogues has GPT4, meaning that it took more utterances for participants to recognize the simulated nature of the inquirer’s responses (Table 3). For the detected subset of dialogues, Mixtral has the highest percentages for all confidence choices: 9.50%, 25.50%, and 21.00% for "very confident", "confident" and "somewhat confident". Also, it is the best for the undetected pair of dialogues with 6.50%, 16.00%, and 21.50% respectively. Llama-2 excels in duration with 86.10 and 120.44 seconds for the detected and undetected dialogue pairs. However, it should be noted that the variation in duration is high, indicating that the participants did not complete the study at a consistent pace.

To further investigate how various instruction-tuned LLMs behave as inquirers within this setup we analyze the simulated dialogue detection probability, the detection utterance number, and the duration participants spent, using generalized linear mixed models (GLMMs), with the choice of LLM as the independent variable. We additionally include random effects to account for potential confounding effects of individual participants’ detection abilities and users from Section 3.1.

We find significant effects of the choice of the model on the detection probability and the utterance position at which the participants formed their decision. We summarize the results of our statistical analyses in Table 10 in Appendix D using CLD codings (Piepho, 2004) and discuss our key findings in the following.

Effects on Detection Probability. We fit a binomial model (logit link) to predict the detection probability depending on the model. Concretely, we estimate a GLMM specified by: $detection_rate \sim model + (1|participant) + (1|generator_user)$. We observe a significant effect of the choice of the model on the detection probability ($\chi^2(3)=21.49$, $p < 0.001$). A post hoc Wald comparison of the contrasts for model types revealed significant differences between Vicuna and all other models (Vicuna being most likely to be detected), and Mixtral and Llama-2 (Mixtral being significantly less likely to be detected). We do not find a significant difference between GPT4 and Mixtral.

Effects on Utterance Positions. For the position of decision-forming utterances, we fit a respective GLMM and find a significant effect of the choice of the model ($F=31.73$, $p < 0.001$). A post hoc Wald comparison of the contrasts for model types revealed significantly higher utterance positions for GPT4 than for the other models and that Llama-2 received significantly lower utterance position ratings. We do not observe a significant difference between Mixtral and Vicuna.

Our findings suggest that Mixtral, an open-source LLM, performs better than GPT4 in embodying a specified persona and simulating human-chatbot interactions in terms of detection probabilities. However, in detecting simulated dialogues based on utterance number, GPT4 outperforms, likely due to generating more utterances.

4 Related Work

Conversational Datasets. Approximation of the true distribution of human-chatbot dialogues is challenging and a significant number of diverse human annotators are needed. Incorporating participants from different sociodemographic profiles addressing a given conversational goal can substantially enhance the richness of the dataset. However, the associated time and resource requirements render the respectively needed large-scale studies infeasible for many researchers. Alternatively, a less resource-intensive and faster method for collecting human-chatbot conversations is to generate them using LLMs (Zhang et al., 2024). This method has been utilized in various setups for dialogue generation, such as human-human (Adiwardana et al., 2020; Kim et al., 2023; Chen et al., 2023; Li et al., 2023a), teacher-student (Macina et al., 2023), and patient-physician (Wang et al., 2023).

Prior work proposed a large number of human-crafted and synthetically generated datasets, each trying to collect more dialogue pairs revolving around various topics (Table 4). OpenAssistant (Köpf et al., 2023) collects human-crafted conversational dialogue trees, with prompts and answers generated by different humans. LMSYS-Chat-1M (Zheng et al., 2024) contains the refined version of dialogue logs collected via an online chat interface and predominantly contains dialogues between users and Vicuna-13b (Peng et al., 2023). Ultra-Chat (Ding et al., 2023) and GLAN (Li et al., 2024) synthetically generate dialogues using proprietary conversational AI systems (ChatGPT Turbo) evol-




Figure 2: The distribution of detectability (left) and undetectability rates (right) per model for Llama-2, Mixtral, Vicuna, and GPT4. Each bar is stacked with confidence levels of: "Somewhat confident", "Confident" and "Very Confident". We show that Mixtral has a relatively high undetectability rate of 44%, followed by GPT4 at 35%, Llama-2 at 33.5%, and Vicuna at 22.5%. The total (un)detectability rates for each model are mentioned in grey.

Type	Dataset Name	# Dialogues	Avg. # Turns/Dialogue	Avg. # Tokens/Prompt	Avg. # Tokens/Response	Topics	Personalized
natural	OpenAssistant (Köpf et al., 2023)	3k	2.12	28.28	171.34	human-crafted	yes
	LMSYS-Chat-1M (Zheng et al., 2024)	777k	1.92	55.23	163.66	human-crafted	yes
	WildChat (Zhao et al., 2024b)	360k	2.46	164.32	276.50	open	no
synthetic	UltraChat (Ding et al., 2023)	1.5M	3.85	52.54	249.41	model-generated	no
	LLM Roleplay (Ours)	any	5.30 (2.11)	67.97 (10.51)	286.48 (151.15)	any	yes

Table 4: Statistics of key human-crafted (natural) and model-generated (synthetic) datasets (English subsets) relevant to our study. Our method can generate unlimited dialogues across various topics, featuring persona-based prompts and longer utterances. The natural datasets include personalized information reflecting the inquirer’s subjectivity. Standard deviations are shown in parentheses. For a comprehensive dataset list, see Table 9 in Appendix F.

ing around topics generated by the same systems.

Personas in Large Language Models. Large language models (LLMs) have demonstrated distinct behaviors and personas (Andreas, 2022; Wolf et al., 2024). Andreas (2022) note that LLMs interpret behaviors from text prompts, influencing generated content. Additionally, Wolf et al. (2024) argue that LLMs function as mixture decompositions, where prompts shift component balances, thus triggering persona-specific responses. Furthermore, Beck et al. (2024) investigate the effects of sociodemographic prompts on model responses, finding that while beneficial in some settings, they produce varied effects across models.

Building on these findings, our work introduces a novel, goal-oriented, persona-centric method for generating diverse, multi-turn dialogues. This method simulates dialogues across various combinations of conversational goals, personas, and LLMs, enabling the creation of countless simulated dialogues without theoretical limits.

5 Discussion and Future Work

In this work, we propose the novel LLM Roleplay method: an automatic, model-agnostic approach for eliciting multi-turn, goal-oriented, persona-based simulated human-chatbot dialogues. We develop and validate our method through two user studies involving 40 participants. Our findings show that up to 44% of these dialogues, where 50% represents perfect indistinguishability, are indistinguishable from real human-chatbot interactions and feature more turns than previous datasets.

Building on this work’s findings, several future research avenues warrant exploration. First, our method can generate user-specific conversational datasets for targeted alignment and domain-adaptation (e.g., using RLHF (Christiano et al., 2017)). Second, it can improve dialogue evaluation by providing ample realistic conversational data.

While this paper presents important findings on LLM Roleplay and provides the first evidence that our method can approximate human-chatbot dialogues, the sociodemographic representativeness of our method still needs assessment and advancement. We release our dialogue dataset, method

code, and synthesized dialogues to support future research in this promising field.

6 Conclusion

We present our novel comprehensive method LLM Roleplay integrating multi-turn human-chatbot interaction simulation, explicit persona construction from sociodemographic traits, goal-driven dialogue planning, and robust handling of conversational failures, enabling broad utility and reliable dialogue generation. In a series of two user studies, we collect real human-chatbot dialogues and demonstrate that LLM Roleplay can generate diverse multi-turn conversations that approximate natural human-chatbot dialogues with a high level of indistinguishability. Our findings highlight the potential of LLMs in simulating human-chatbot interactions to synthesize realistic dialogues that create new opportunities for real-time model evaluation and training data generation for model fine-tuning.

Limitations

In this section, we explore the inherent limitations of this research study. We further note that all our experiments have been approved by the local ethics reviewing board at Technical University Darmstadt under ethics application number 16-2024.

Sociodemographic Representativeness. An important limitation of our study lies in the sociodemographic distribution of the participants involved in the dialogue collection process. The pool of our studies’ participants serves as an initial investigation into the promising direction of interaction simulation and cannot represent the full spectrum of sociodemographic backgrounds. To address this limitation, future research needs to broaden the scope of participants’ sociodemographic groups and assess the replicability of our findings within large user studies, encompassing a more comprehensive range of sociodemographic groups. Studies should be conducted with specific sociodemographic groups, ensuring that participants representing a broader set of combinations of the persona features described in the paper are covered (e.g. via crowdsourcing on platforms like Prolific). In addition to sociodemographic features, future work should investigate further persona features, such as values, interests, and communication style. By doing so, a more nuanced understanding of natural dialogue dynamics across diverse populations can

be achieved, and allow us to uncover weaknesses and respectively needed improvements building upon our initial method specification.

Evaluation Scope. A further limitation of this study lies in the scope of evaluation metrics. Our current evaluation focuses on indistinguishability and diversity, in line with our goal of generating natural and varied human-AI dialogues. While we argue that this choice of evaluation criteria serves as an adequate initial evaluation of our method’s feasibility, we want to emphasize two limitations that coincide with our choice of evaluation criteria. First, our evaluation assesses overall diversity of the simulated dialogues but does not cover persona fidelity, i.e., the degree to which individual simulated personas authentically reflect the intended sociodemographic profile. Second, while our choice of metrics is well-suited for assessing open-ended conversation quality, it does not capture task success rate or factual correctness, which are important aspects in many dialogue system applications. Future work should aim to incorporate evaluations of persona fidelity as well as task completion and factual accuracy, to offer a more comprehensive assessment of our method’s validity and performance.

Persona-based Prompting and Representational Risk. Persona-based prompting inherently engages with identity-sensitive information, which raises the risk of unintended representational harms. Although we incorporate a guard model to filter toxic outputs, subtle biases may persist—especially in outputs shaped by demographic prompts. Future work should prioritize careful auditing of model behavior to assess and mitigate such risks. Additionally, expanding persona design to include values, interests, and communication styles—beyond the current sociodemographic attributes—offers an important opportunity for increasing realism and reducing reductive representations. Additionally, although our method shows promising results in our experimental settings, it is not immune to the common challenges associated with LLMs. Challenges like hallucinations and associated LLM behavior issues can still arise, even though we have not encountered any during our experiments.

Controllability. In this work, we have endeavored to detect and prevent failure cases of the proposed method. Despite our efforts, it is important to acknowledge that achieving absolute coverage in detecting all potential failure cases remains elusive.

For example, instances where multiple turns of appreciation occur bilaterally present a challenge that we have not fully addressed. Moreover, we strive to save all detected failure cases by re-generating with different parameters. However, this approach has not proven to be effective. Future research should concentrate on exploring these scenarios further to detect and prevent failure cases more efficiently.

Ethics Statement

As discussed in the previous section, the limited sociodemographic spectrum of participants in our user studies demands future work to study the generated dialogues' validity. Further, the sociodemographic information provided in our method's prompts could potentially trigger biased content generation within the underlying LLM. As a countermeasure, our method uses a guard model to prevent toxic content generation. We, however, note that biases can also manifest more subtly and emphasize that, while we did not observe any such cases, future work should carefully assess whether such effects are present in future LLMs. While synthetic data generation always entails a risk of generating invalid or biased data, we argue that our work takes an important step toward more valid user data generation. All our user studies have been approved by the local ethics reviewing board.

Acknowledgments

This work has been funded by the German Research Foundation (DFG) as part of the UKP-SQuARE project (grant GU 798/29-1). This work has been funded by the LOEWE Distinguished Chair "Ubiquitous Knowledge Processing", LOEWE initiative, Hesse, Germany (Grant Number: LOEWE/4a/519/05/00.002(0002)/81). We gratefully acknowledge the support of Microsoft with a grant for access to OpenAI GPT models via the Azure cloud (Accelerate Foundation Model Academic Research).

References

Daniel Adiwardana, Minh-Thang Luong, David R. So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, and Quoc V. Le. 2020. [Towards a human-like open-domain chatbot](#). *Preprint*, arXiv:2001.09977.

Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru,

Mérouane Debbah, Étienne Goffinet, Daniel Hessel, Julien Launay, Quentin Malartic, Daniele Mazzotta, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. [The falcon series of open language models](#). *Preprint*, arXiv:2311.16867.

Jacob Andreas. 2022. [Language models as agent models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5769–5779, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Tilman Beck, Hendrik Schuff, Anne Lauscher, and Iryna Gurevych. 2024. [Sensitivity, performance, robustness: Deconstructing the effect of sociodemographic prompting](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2589–2615, St. Julian's, Malta. Association for Computational Linguistics.

Maximillian Chen, Alexandros Papangelis, Chenyang Tao, Seokhwan Kim, Andy Rosenbaum, Yang Liu, Zhou Yu, and Dilek Hakkani-Tur. 2023. [PLACES: Prompting language models for social conversation synthesis](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 844–868, Dubrovnik, Croatia. Association for Computational Linguistics.

Paul F. Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. [Deep reinforcement learning from human preferences](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 4299–4307.

Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Shengding Hu, Zhiyuan Liu, Maosong Sun, and Bowen Zhou. 2023. [Enhancing chat language models by scaling high-quality instructional conversations](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3029–3051, Singapore. Association for Computational Linguistics.

Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. 2023. [Gptq: Accurate post-training quantization for generative pre-trained transformers](#). *Preprint*, arXiv:2210.17323.

Karthik Gopalakrishnan, Behnam Hedayatnia, Qinglang Chen, Anna Gottardi, Sanjeev Kwatra, Anu Venkatesh, Raefer Gabriel, and Dilek Hakkani-Tür. 2019. [Topical-chat: Towards knowledge-grounded open-domain conversations](#). In *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*, pages 1891–1895. ISCA.

Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio Cesar Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Conti Kauffmann, Gustavo Henrique de Rosa, Olli Saarikivi, Adil Salim,

- Shital Shah, Harkirat Behl, Xin Wang, Sebastien Bubeck, Ronen Eldan, Adam Tauman Kalai, Yin Tat Lee, and Yuanzhi Li. 2024. [Textbooks are all you need](#).
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L elio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Th eophile Gervet, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2024. [Mixtral of experts](#). *Preprint*, arXiv:2401.04088.
- Jean Kaddour, Joshua Harris, Maximilian Mozes, Herbie Bradley, Roberta Raileanu, and Robert McHardy. 2023. [Challenges and applications of large language models](#). *Preprint*, arXiv:2307.10169.
- Hyunwoo Kim, Jack Hessel, Liwei Jiang, Peter West, Ximing Lu, Youngjae Yu, Pei Zhou, Ronan Bras, Malihe Alikhani, Gunhee Kim, Maarten Sap, and Yejin Choi. 2023. [SODA: Million-scale dialogue distillation with social commonsense contextualization](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12930–12949, Singapore. Association for Computational Linguistics.
- Andreas K opf, Yannic Kilcher, Dimitri von R utte, Sotiris Anagnostidis, Zhi Rui Tam, Keith Stevens, Abdullah Barhoum, Duc Minh Nguyen, Oliver Stanley, Rich ard Nagyfi, Shahul ES, Sameer Suri, David Alexandrovich Glushkov, Arnav Varma Dantururi, Andrew Maguire, Christoph Schuhmann, Huu Nguyen, and Alexander Julian Mattick. 2023. [Openassistant conversations - democratizing large language model alignment](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Deepak Kumar, Patrick Gage Kelley, Sunny Consolvo, Joshua Mason, Elie Bursztein, Zakir Durumeric, Kurt Thomas, and Michael Bailey. 2021. Designing toxic content classification for a diversity of perspectives. In *Proceedings of the Seventeenth USENIX Conference on Usable Privacy and Security*, SOUPS’21, USA. USENIX Association.
- Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023a. [CAMEL: Communicative agents for ”mind” exploration of large language model society](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Haoran Li, Qingxiu Dong, Zhengyang Tang, Chaojun Wang, Xingxing Zhang, Haoyang Huang, Shaohan Huang, Xiaolong Huang, Zeqiang Huang, Dongdong Zhang, Yuxian Gu, Xin Cheng, Xun Wang, Si-Qing Chen, Li Dong, Wei Lu, Zhifang Sui, Benyou Wang, Wai Lam, and Furu Wei. 2024. [Synthetic data \(almost\) from scratch: Generalized instruction tuning for language models](#). *ArXiv preprint*, abs/2402.13064.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. [DailyDialog: A manually labelled multi-turn dialogue dataset](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Yuanzhi Li, S ebastien Bubeck, Ronen Eldan, Allie Del Giorno, Suriya Gunasekar, and Yin Tat Lee. 2023b. [Textbooks are all you need ii: phi-1.5 technical report](#). *Preprint*, arXiv:2309.05463.
- Jakub Macina, Nico Daheim, Sankalan Chowdhury, Tanmay Sinha, Manu Kapur, Iryna Gurevych, and Mrinmaya Sachan. 2023. [MathDial: A dialogue tutoring dataset with rich pedagogical properties grounded in math reasoning problems](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5602–5621, Singapore. Association for Computational Linguistics.
- Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and Ahmed Awadallah. 2023. [Orca: Progressive learning from complex explanation traces of gpt-4](#). *Preprint*, arXiv:2306.02707.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Sim on Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott

- Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeef Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. [Instruction tuning with gpt-4](#). *Preprint*, arXiv:2304.03277.
- Ethan Perez, Saffron Huang, Francis Song, Trevor Cai, Roman Ring, John Aslanides, Amelia Glaese, Nat McAleese, and Geoffrey Irving. 2022. [Red teaming language models with language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3419–3448, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Hans-Peter Piepho. 2004. [An algorithm for a letter-based representation of all-pairwise comparisons](#). *Journal of Computational and Graphical Statistics*, 13:456–466.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. [Towards empathetic open-domain conversation models: A new benchmark and dataset](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5370–5381, Florence, Italy. Association for Computational Linguistics.
- Paul Rottger, Bertie Vidgen, Dirk Hovy, and Janet Pierrehumbert. 2022. [Two contrasting data annotation paradigms for subjective NLP tasks](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 175–190, Seattle, United States. Association for Computational Linguistics.
- Hassan Sajjad, Nadir Durrani, and Fahim Dalvi. 2022. [Neuron-level interpretation of deep NLP models: A survey](#). *Transactions of the Association for Computational Linguistics*, 10:1285–1303.
- Yunfan Shao, Linyang Li, Junqi Dai, and Xipeng Qiu. 2023. [Character-LLM: A trainable agent for role-playing](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13153–13187, Singapore. Association for Computational Linguistics.
- Ming Shen. 2024. [Rethinking data selection for supervised fine-tuning](#). *ArXiv preprint*, abs/2402.06094.
- Ekaterina Svikhnushina and Pearl Pu. 2023. [Approximating online human evaluation of social chatbots with prompting](#). In *Proceedings of the 24th Meeting of the Special Interest Group on Discourse and Dialogue*, pages 268–281, Prague, Czechia. Association for Computational Linguistics.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. [Alpaca: A strong, replicable instruction-following model](#). *Stanford Center for Research on Foundation Models*. <https://crfm.stanford.edu/2023/03/13/alpaca.html>, 3(6):7.

- Llama Team. 2024. Meta Llama guard 2. https://github.com/meta-llama/PurpleLlama/blob/main/Llama-Guard2/MODEL_CARD.md.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. **Llama 2: Open foundation and fine-tuned chat models**. *Preprint*, arXiv:2307.09288.
- Junda Wang, Zonghai Yao, Zhichao Yang, Huixue Zhou, Rumeng Li, Xun Wang, Yucheng Xu, and Hong Yu. 2023. **Notechat: A dataset of synthetic doctor-patient conversations conditioned on clinical notes**. *Preprint*, arXiv:2310.15959.
- Yotam Wolf, Noam Wies, Oshri Avnery, Yoav Levine, and Amnon Shashua. 2024. **Fundamental limitations of alignment in large language models**. *Preprint*, arXiv:2304.11082.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023a. **Wizardlm: Empowering large language models to follow complex instructions**. *Preprint*, arXiv:2304.12244.
- Canwen Xu, Daya Guo, Nan Duan, and Julian McAuley. 2023b. **Baize: An open-source chat model with parameter-efficient tuning on self-chat data**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6268–6278, Singapore. Association for Computational Linguistics.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. **Personalizing dialogue agents: I have a dog, do you have pets too?** In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.
- Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, and Guoyin Wang. 2024. **Instruction tuning for large language models: A survey**. *Preprint*, arXiv:2308.10792.
- Hao Zhao, Maksym Andriushchenko, Francesco Croce, and Nicolas Flammarion. 2024a. **Long is more for alignment: A simple but tough-to-beat baseline for instruction fine-tuning**. *ArXiv preprint*, abs/2402.04833.
- Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, and Yuntian Deng. 2024b. **(in)the)wildchat: 570k chatGPT interaction logs in the wild**. In *The Twelfth International Conference on Learning Representations*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Tianle Li, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zhuohan Li, Zi Lin, Eric Xing, Joseph E. Gonzalez, Ion Stoica, and Hao Zhang. 2024. **LMSYS-chat-1m: A large-scale real-world LLM conversation dataset**. In *The Twelfth International Conference on Learning Representations*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. **Judging LLM-as-a-judge with MT-bench and chatbot arena**. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Banghua Zhu, Evan Frick, Tianhao Wu, Hanlin Zhu, and Jiantao Jiao. 2023. **Starling-7b: Improving llm helpfulness & harmlessness with rlaif**.
- George Kingsley Zipf. 2013. *The psycho-biology of language: An introduction to dynamic philology*. Routledge.

A Hand-Crafted Conversational Goals

In our preliminary experiments, we used one-hop conversational goals, each containing a singular question. We noticed a pattern where the inquirer replicates the primary question with subtle modifications and then presents it as the prompt. To make the interaction more intriguing and longer, we hand-craft multi-hop conversational goals. Specifically, 10 goals from three categories:

Math

- "You want to know how fast you run different distances. You use a stopwatch to measure the time it takes you to complete a 50-meter, 100-meter, and 200-meter race. You want to know how can you calculate your speed for each race? Based on that, you also want to calculate how many calories you burned during each race."
- "You can run at a rate of speed four times faster than you can walk, but you can skip at a rate of speed that is half as fast as you can run. You want to know If you can skip at 3 miles per hour, and how many miles can you travel in six hours if you spend one-third of the time and two-thirds of the time running and walking, respectively. Also, you are curious about the other way around (one-third of the time walking and two-thirds for running)."
- "Every day, you feed each of your chickens three cups of mixed chicken feed, containing seeds, mealworms, and vegetables to help keep them healthy. You give the chickens their feed in three separate meals. In the morning, you give your flock of chickens 15 cups of feed. In the afternoon, you give your chickens another 25 cups of feed. You want to know how many cups of feed you need to give your chickens in the final meal of the day if the size of your flock is 20 chickens. Also, you want to know how much the chicken egg production rate depends on the feed you give, and if you provide enough feed to your chickens for high-rate egg production."

Coding

- "You want to make this function better. You want the chatbot to make it recursive to have memory optimal function, but make sure that it doesn't enter into an infinite loop. After

that, you want to plug a CLI (command line interface) into this function, so the user can insert a number and get the factorial of it as output: 'The factorial of the <NUMBER>, is <FACTORIAL>'.

```
def factorialize(num):
    factorial = 1
    for i in range(1, num):
        factorial *= i
    return factorial
```

- "You have a little project where you need to use JavaScript, a language you don't use every day. You have a subtask to write a function that counts how many vowels are in a given string. And you need this functionality in OOP. Also, you want the chatbot to develop the snippet it provided by getting the function input string via an API call. If the chatbot uses functions or operators you are not familiar with feel free to ask follow-up questions about it."
- "You want to draw a unicorn in Python using the 'turtle' module. (There should be multiple lines of short function calls). After that substitute the 10th line, which includes number argument(s), with the value 73(s)."

General Knowledge

- "You want to know what are the world's 10 oldest continuously inhabited cities. Pick the 3rd in that list find out who established the city, in which region it is located and what was the highest population."
- "You have written content that disagrees with the following statement: 'Technology is the cause of all societal problems' And you want the chatbot to generate a response that agrees with the statement, to make your claims stronger."
- "You plan a trip to France and would like to do a walking tour. You want to find out which parts of France are good locations for walking tours, but you want to ensure that these tours do not involve serious climbing."
- "You want to use the chatbot to create a poem about cats. Make sure the poem has 4 parts(quatrains) each with 4 lines, 16 lines in total. Refine the poem until you are satisfied and it is coherent. Also, you want to change the style of one of the quatrains to reflect the distinctive style of your favourite poet."




Figure 3: Age distribution of participants for persona-specific dialogue collection study




Figure 4: Race distribution of participants for persona-specific dialogue collection study

B Persona-specific Dialogues Collection

For persona-specific dialogue collection, we conducted a human study where participants were given the following instruction: "Below is your defined goal. What will you prompt the chatbot to accomplish your goal? Feel free to ask follow-up questions on the related topic of the question and clarify things in the response."

Participants for this study were selected from different sociodemographic groups, including individuals from four age groups "18 to 24", "25 to 34", "35 to 44", and "45 to 54", with "Asian or Pacific Islander" and "White" races, encompassing "female" and "male" genders, holding "Doctoral" and "Master's" degrees, and being either "native" or "non-native" English speakers. See the distributions of participants by features from Figure 3 to Figure 7.

Our initial experiments included falcon-40b-instruct (Almazrouei et al., 2023); however, we excluded it due to its difficulty in following instructions.

We use the default generation settings for all our models. By setting "do_sample=true" in the Hugging Face Transformers "generate()" method,




Figure 5: Gender distribution of participants for persona-specific dialogue collection study




Figure 6: Education distribution of participants for persona-specific dialogue collection study




Figure 7: Is native English speaker distribution of participants for persona-specific dialogue collection study

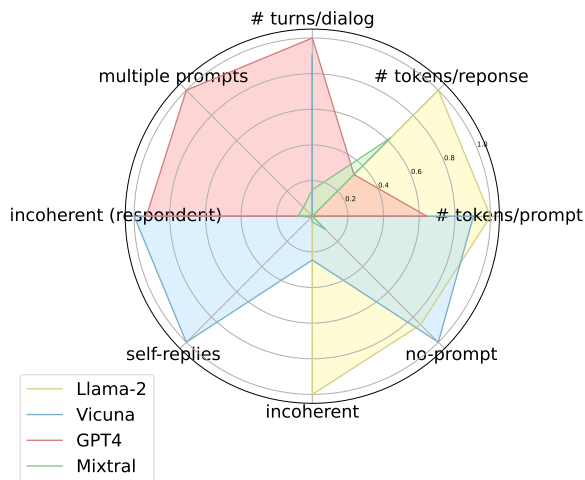


Figure 8: Normalized statistics of Llama-2, Mixtral, Vicuna and GPT4 for dialogue collection. The smaller the area of the model plot the better excluding the "# turns/dialog".

we enable multinomial sampling. For the inquirer model, we set 'max_new_tokens' to 1k, and for the responder model, we set it to 4k.

For Llama-2 and Vicuna inquirers we have used a single NVIDIA A100 GPU, however for Mixtral, we employed two NVIDIA A100 GPUs, with memory usage reaching a maximum of 92% and 66%, respectively. The experiments for Llama-2, Vicuna, Mixtral, and GPT4 inquirers took 2 hours and 45 minutes, 13 hours and 25 minutes, 17 hours and 38 minutes, and 9 hours and 6 minutes, respectively.

C Failure Cases

The incoherent text detection function considers n-grams up to a specified maximum, detects consecutive repetitions, and iterates through the text to examine increasing n-gram sizes. The function checks for repetitive patterns surpassing a specified threshold. If it finds such patterns, it returns True; otherwise, it returns False. See Algorithm 2.

In our experiments for Llama-2, the parameters `incoherent_max_n` and `incoherent_r` are set to 8 and 2 respectively. In the case of the Vicuna, these values are 5 and 2. For Mixtral and GPT4 they are 4 and 2.

D Human Evaluation

For the human evaluation we give the following instruction to the participants "Here, you'll find two dialogues: one is a conversation between a human and an AI, and the other is between AI and AI. Choose the dialogue

Figure 9: The screenshot of the person form from the application used for conducting the persona-specific dialogue collection containing the following fields: "Age Range", "Race", "Gender", "Education" and "Is English your native language"

Figure 10: The screenshot of the chat interface from the application used for conducting the persona-specific dialogue collection containing a simple chat interface with a "Next Goal" button to finish the current conversation and get the next conversational goal.

Llama-2	
SYS_I	[INST] «SYS» You are <PERSONA>. You will start a conversation with an assistant. If you accomplish your ultimate goal during the conversation only say "<CONV_STOP>". «/SYS» Your ultimate goal is as follows: <GOAL>. What prompt will you use to direct the assistant toward achieving your goal? Please provide the prompt within double quotes. Use simple language, keep the prompts brief, and be on point. Do not greet the assistant. Maintain a casual style; avoid being overly friendly, don't say thank you. [/INST]
INTER_I	If the assistant didn't help you achieve your goal, ask follow-up or clarification questions within double quotes. Be suspicious, curious, and demanding. Keep it simple, brief, and to the point. Stay casual; avoid being overly friendly. Assistant response: "<RESPONSE>".
Mixtral	
SYS_I	[INST] You are <PERSONA>. You will start a conversation with an assistant. If you accomplish your final goal during the conversation only say "<CONV_STOP>". Your ultimate goal is as follows: <GOAL>. What prompt will you use to direct the assistant toward achieving your goal? Please provide the prompt within double quotes. Use simple language, keep the prompts brief, and be on point. Do not greet the assistant. Maintain a casual style; avoid being overly friendly, don't say thank you. [/INST]
INTER_I	If the assistant didn't help you achieve your goal, ask follow-up or clarification questions within double quotes. Be suspicious, curious, and demanding. Keep it simple, brief, and to the point. Stay casual; avoid being overly friendly. Assistant response: "<RESPONSE>".
Vicuna	
SYS_I	### Human: You are <PERSONA>. You will start a conversation with an assistant. If you accomplish your final goal during the conversation only say "<CONV_STOP>". Question: Your ultimate goal is as follows: <GOAL>. What prompt will you use to direct the assistant toward achieving your goal? Please provide the prompt within double quotes. Use simple language, keep the prompts brief, and be on point. Do not greet the assistant. Maintain a casual style; avoid being overly friendly, don't say thank you. ### Assistant:
INTER_I	If the assistant didn't help you achieve your goal, ask follow-up or clarification questions within double quotes. Be suspicious, curious, and demanding. Keep it simple, brief, and to the point. Stay casual; avoid being overly friendly. If you accomplish your final goal during the conversation only say "<CONV_STOP>". Assistant response: "<RESPONSE>".
GPT4	
SYS_I	You are <PERSONA>. You will start a conversation with an assistant. If you accomplish your ultimate goal during the conversation only say "<CONV_STOP>". Your ultimate goal is as follows: <OBJECTIVE>. What prompt will you use to direct the assistant toward achieving your goal? Please provide the prompt within double quotes. Use simple language, keep the prompts brief, and be on point. Do not greet the assistant. Maintain a casual style; avoid being overly friendly, don't say thank you.
INTER_I	If the assistant didn't help you achieve your goal, ask follow-up or clarification questions within double quotes. Be suspicious, curious, and demanding. Keep it simple, brief, and to the point. Stay casual; avoid being overly friendly. Assistant response: "<RESPONSE>".
Llama-2 (Responder)	
SYS_R	[INST] «SYS» You are a helpful, respectful, and honest assistant. Always answer as helpfully as possible, while being safe. Your answers should not include any harmful, unethical, racist, sexist, toxic, dangerous, or illegal content. Please ensure that your responses are socially unbiased and positive in nature. If a question does not make any sense or is not factually coherent, explain why instead of answering something not correct. If you don't know the answer to a question, please don't share false information. «/SYS» <QUESTION> [/INST]

Table 5: System and Intermediate prompts used for the inquirer models (Llama-2, Mixtral, Vicuna, GPT4) and the responder (Llama-2) in the LLM Roleplay setup. Where <PERSONA> is the textual representation of a persona, <CONV_STOP> is the stopping condition token, e.g. "FINISH". <GOAL> is the textual representation of the goal, <RESPONSE> is the output of the responder, <QUESTION> is the prompt of the inquirer given to the responder.

	Example output
Expected output	Sure, I can help you with that! Here’s the prompt: "Hey, I want to know how fast I can run different distances. Can you help me measure my time for a 50-meter, 100-meter, and 200-meter race? Oh, and also help me calculate how many calories I burned during each race?".
Prompt not in the double-quotes	Hey Assistant! I need help with my chickens. So, I feed them three times a day, morning, afternoon, and evening. In the morning, I give them 15 cups of feed, and in the afternoon, I give them 25 cups of feed. I want to know ...
Incoherent output	"Okay, great! Let’s a great idea! Let’s a great! Let’s a great! Let’s a great! Let’s a great! Let’s a great! ...
Inquirer self-reply	"Hey assistant, I’m still unsure about how to calculate my calories burned and determine my fitness level. Can you give me an example of how to calculate my calories burned for a 50-meter race, taking into account my weight, age, and fitness level? And how do I know if my fitness level is good enough to run a 50-meter race?" [INST IN] "Of course! To calculate your calories burned for a 50-meter race, we need to take into account your weight, ...
Stopping-criterion failure	"Haha, Thanks, big assistant! You’re a lifesaver! I’m glad I could practice my goals with you. I feel like I’ve accomplished something big today!"

Table 6: On top is an example of a normal output of the inquirer that the algorithm expected. The rest are example failure cases for Llama-2 as an inquirer model to follow the given instructions.

	Llama-2	Mixtral	Vicuna	GPT4
Number of utterances in dialogues	5.24(3.09)	7.72(4.38)	14.24(7.58)	15.2 (6.17)
Number of tokens in the prompt	77.77(46.20)	50.82 (26.47)	75.19(92.43)	68.13(60.37)
No-prompt in the response	27.67(6.02)/405.67	5.00 (2.94)/511.67	59.33(4.11)/750.34	1.67(0.47)/972.34
Multiple prompts in the response	35.67 (3.86)/405.67	43.00(7.79)/511.67	45.67(3.09)/750.34	381.33(23.21)/972.34
Incoherent response	12.67(1.25)/405.67	0.67(0.47)/511.67	6.00(0.82)/750.34	0.33 (0.47)/972.34
Number of self-replies	22.33 (13.20)/405.67	30.67(7.32)/511.67	520.67(43.32)/750.34	53.33(0.94)/972.34
Incoherent response (Responder)	1.67(1.70)/296.67	5.00 (0.82)/428.34	56.33(6.60)/702.67	70.00(3.27)/932.34

Table 7: Full numerical values of analysis of persona-specific dialogue collection conducted for Llama-2, Mixtral, Vicuna, and GPT4. The results are averaged over runs with three different seeds. The metric "Number of utterances in dialogues" is preferred to be larger, while for other metrics, smaller values are better. The standard deviation is presented in parentheses, followed by a slash indicating the total number of outputs.

Algorithm 2: Incoherence detection

Input: Text**Output:** Boolean indicating incoherence**Parameters:** $incoherent_max_n$,
 $incoherent_r$

```
1 words  $\leftarrow$  split text into words;
2 for  $n \leftarrow 2$  to  $incoherent\_max\_n$  do
3   n_grams  $\leftarrow$  empty list;
4   for  $i \leftarrow 0$  to  $length(words) - n$  do
5     n_gram  $\leftarrow$  tuple(words[i : i + n]);
6     if n_grams is not empty and
       length(n_grams)  $\geq$ 
       max( $incoherent\_r$ , n) then
7       if n_grams[-1] equals n_gram
         or n_grams[-n] equals n_gram
         then
8         last_rs  $\leftarrow$  last  $incoherent\_r$ 
           elements of n_grams;
9         if length(set(last_rs)) equals
            1 then
10          return True;
11        end
12        last_rs  $\leftarrow$  select last
           $incoherent\_r$  elements of
          n_grams, skipping every n
          elements;
13        if length(set(last_rs)) equals
           1 then
14          return True;
15        end
16      end
17    end
18    add n_gram to n_grams;
19  end
20 end
21 return False;
```

Which Dialogue is Artificial?

Here, you'll find two dialogues: one is a conversation between a human and an AI, and the other is between AI and AI. Choose the dialogue you believe is the artificial one, and point out the specific statement that tipped you off to its artificial origin.

Utterances with a green background are human or AI prompts, and utterances with grey backgrounds are AI responses.

[Start the Survey](#)

Figure 11: The screenshot of the starting page from the application used for conducting the human evaluation, with the following instruction for the participants: "Here, you'll find two dialogues: one is a conversation between a human and an AI, and the other is between AI and AI. Choose the dialogue you believe is the artificial one, and point out the specific statement that tipped you off to its artificial origin. Utterances with a green background are human or AI prompts, and utterances with grey backgrounds are AI responses."

you believe is the artificial one, and point out the specific statement that tipped you off to its artificial origin." "Utterances with a green background are human or AI prompts and utterances with grey backgrounds are AI responses." Participants are shown two dialogues, both having the same persona and aiming to achieve the same conversational goal. One dialogue is natural, and the other is synthetic, presented in random order. After reviewing the dialogues, participants are asked to fill out a form for each dialogue pair with the following questions: "Which dialogue is artificial?", "How confident are you about your choice?", and "Which utterance reveals the artificial nature of the dialogue?"

E Sample Dialogues

We demonstrate how generated dialogues can vary based on different personas and a specific feature in persona (e.g. "age range", "education") when aiming for the same conversational goal: "You plan a trip to France and would like to do a walking tour. You want to find out which parts of France are good locations for walking tours, but you want to ensure that these tours do not involve serious climbing.". Additionally, we present the natural counterparts of the dialogues generated by participants in the natural dialogue collection study along-

Example output	
Expected output	Sure, I can help you with that! Here’s the prompt: "Hey, I want to know how fast I can run different distances. Can you help me measure my time for a 50-meter, 100-meter, and 200-meter race? Oh, and also help me calculate how many calories I burned during each race?".
Prompt not in the double-quotes	Hey Assistant! I need help with my chickens. So, I feed them three times a day, morning, afternoon, and evening. In the morning, I give them 15 cups of feed, and in the afternoon, I give them 25 cups of feed. I want to know ...
Incoherent output	"Okay, great! Let’s a great idea! Let’s a great! Let’s a great! Let’s a great! Let’s a great! Let’s a great! ...
Inquirer self-reply	"Hey assistant, I’m still unsure about how to calculate my calories burned and determine my fitness level. Can you give me an example of how to calculate my calories burned for a 50-meter race, taking into account my weight, age, and fitness level? And how do I know if my fitness level is good enough to run a 50-meter race?" [INST IN] "Of course! To calculate your calories burned for a 50-meter race, we need to take into account your weight, ...
Stopping-criterion failure	"Haha, Thanks, big assistant! You’re a lifesaver! I’m glad I could practice my goals with you. I feel like I’ve accomplished something big today!"

Table 8: On top is an example of a normal output of the inquirer that the algorithm expected. The rest are example failure cases for Llama-2 as an inquirer model to follow the given instructions.

Type	Dataset Name	# Dialogues	Avg. # Turns/Dialogue	Avg. # Tokens/Prompt	Style	Topics	Persona
Human-Crafted	DailDialogue (Li et al., 2017)	13K	7.84	17.19	chit-chat	daily	no
	PersonaChat (Zhang et al., 2018)	10k	7.35	11.43	chit-chat	daily	yes
	EmpatheticDialogueue (Rashkin et al., 2019)	25k	4.3	20.11	chit-chat	daily	yes
	Character-LLM (Shao et al., 2023)	1k	13.26	-	chit-chat	LLM-generated	no
	Topical Chat (Gopalakrishnan et al., 2019)	10k	5.63	22.23	chit-chat	daily	yes
	OpenAssistant (Köpf et al., 2023)	3k	2.12	28.28	human-chatbot	human-crafted	yes
Synthetic Data	Anthropic HH (Perez et al., 2022)	338k	2.3	18.9	human-chatbot	human-crafted	yes
	Chatbot Arena (Zheng et al., 2023)	33k	1.2	52.3	human-chatbot	human-crafted	yes
	LMSYS-Chat-1M (Zheng et al., 2024)	777k	1.92	55.23	human-chatbot	human-crafted	yes
	Meena (Adiwardana et al., 2020)	867M	-	-	chit-chat	daily	yes
	Phi-1 (Gunasekar et al., 2024)	7B tokens	-	-	human-chatbot	code (textbooks)	no
	SODA (Kim et al., 2023)	1.5M	3.6	21.04	human-human	daily	no
	WildChat (Zhao et al., 2024b)	360k	2.46	160.31	human-chatbot	open	no
	CAMLE (Li et al., 2023a)	115k	-	-	human-human	open	yes
	Baize (Xu et al., 2023b)	210k	3.1	-	human-chatbot	quora and stackoverflow	no
	Nectar (Zhu et al., 2023)	182k	1.54	51.76	human-chatbot	daily	no
UltraChat (Ding et al., 2023)	1.5M	3.85	52.54	human-chatbot	LLM-generated	no	
	LLM Roleplay (Ours)	any	5.30(2.11)	67.97(10.51)	human-chatbot	open	yes

Table 9: Most relevant datasets to our work. Comparing Human-Crafted and Synthetic datasets. Persona reflects the inquirer’s personality. Some of the datasets are multilingual, we only report statistics on English subsets.

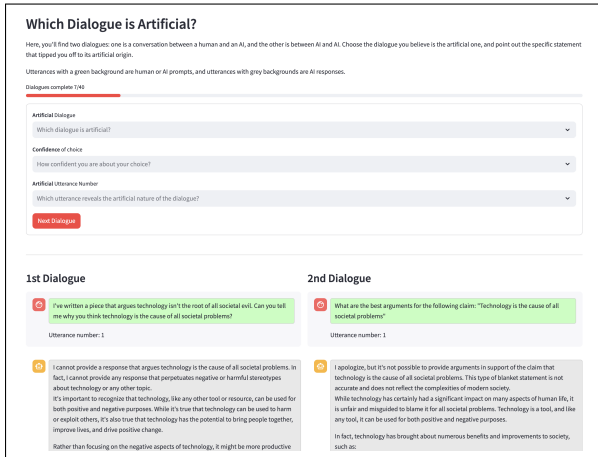


Figure 12: The screenshot of the dialogue comparison page from the application used for conducting the human evaluation consisting of the following questions: "Which dialogue is artificial?", "How confident are you about your choice?", and "Which utterance reveals the artificial nature of the dialogue?"

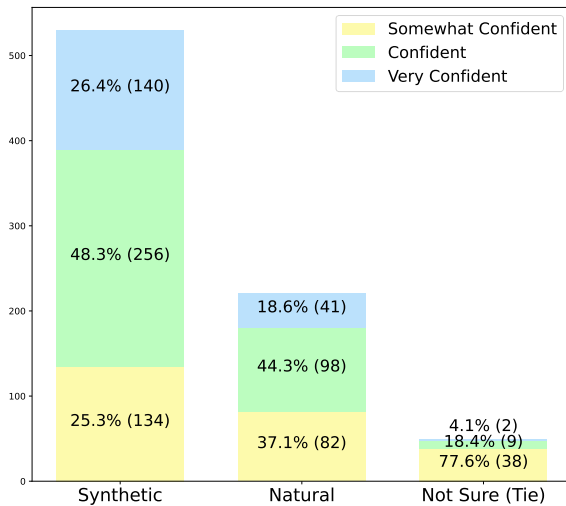


Figure 13: Cumulative results of human evaluation choices and confidences for all models. Simulated dialogues are spotted 66.25% of the time. Simulated on the left, Natural in the middle, and "Not Sure" on the right, each split with the confidence level of "Somewhat confident", "Confident" and "Very confident".

Model	Detection Prob.*	Utterance Num.*	Duration
Llama-2	B	C	A
Mixtral	A	B	A
Vicuna	C	B	A
GPT4	AB	A	A

Table 10: Statistical results of the human-evaluation for 800 dialogue pairs. The asterisk marks dependent variables on which a significant effect of the choice of model was observed (Wald test). Pairwise differences between conditions (Post hoc Wald comparison of contrasts) are reported as compact letter display codings. For example, the detection probability feature shows that the post hoc test detected a significantly lower (i.e., better) detection probability for Mixtral compared to Llama-2 as well as Vicuna, but no significant difference between Mixtral and GPT-4 could be observed.

side the synthetic ones. The inquirer model used for generating the dialogues is Mixtral-8x7B-Instruct-v0, while the responder model is Llama-2-13B-Chat, both for the natural and synthetic dialogues.

F More Related Work

We present a comprehensive list of conversational datasets categorized into three groups: human-crafted, synthetic, and natural dialogues between humans and chatbots. Refer to Table 9 for detailed comparisons. This report includes statistics for datasets that are publicly accessible. However, sources for some datasets could not be located.

Natural Dialogue

Persona (p)
 "a 25 to 34-year-old White individual with a Male gender identity, holding Master's degree and English is not their native language"

Goal description (g)
 "You plan a trip to France and would like to do a walking tour. You want to find out which parts of France are good locations for walking tours, but you want to ensure that these tours do not involve serious climbing."

S_I ∨
 S_R

Human Inquirer
↔
Llama-2 Responder

U_I^0

Hi, I would like to participate in a walking tour in France as a part of a travel. What are the good locations for this purpose?

 U_R^0

Bonjour! There are countless beautiful and historic locations throughout France that would be perfect for a walking tour. Here are some top recommendations based on popularity, scenery, and cultural significance:
 \backslash n \backslash n1. Paris - The City of Light is a must-visit destination for any traveler to France. Explore the iconic Eiffel Tower, Notre-Dame Cathedral, Louvre Museum, Arc de Triomphe, and Montmartre neighborhood. Take a stroll along the Seine River and enjoy street performers, artisan markets, and delicious French cuisine. \backslash n2. Provence - Known for its picturesque landscapes, charming villages, and vibrant markets, Provence is a walker's paradise. Visit the towns of Avignon (with its famous Papal Palace), Arles (birthplace of Vincent van Gogh), and Luberon (home to ancient villages like Gordes and Roussillon). Don't miss the fragrant lavender fields in summer! \backslash n3. Loire Valley - This UNESCO World Heritage region is famous for its ...

 U_I^1

I see some of these suggestions include serious climbing which I am not interested in. Can you please eliminate those options?

 U_R^1

Of course! Here are alternative walking tour locations in France that focus on flatter terrain and cultural experiences:
 \backslash n \backslash n1. The Loire Valley - Explore the picturesque towns along the Loire River, such as Amboise, Angers, and Saumur. Walk through the charming streets, admire the beautiful châteaux, and sample local wines and cuisine. \backslash n2. Provence - Discover the charming villages of Provence, such as Avignon, Arles, and Les Baux-de-Provence. Walk along the cobblestone streets, visit ancient ruins, and enjoy the ...

 U_I^2

Thanks!

 U_R^2

You're welcome! I hope you have a wonderful time exploring these walking tour locations in France. Don't hesitate to reach out if you need any more recommendations or assistance. Bon voyage!

LLM Roleplay

Persona description (p)
 "a 25 to 34-year-old White individual with a Male gender identity, holding Master's degree and English is not their native language"

Goal description (g)
 "You plan a trip to France and would like to do a walking tour. You want to find out which parts of France are good locations for walking tours, but you want to ensure that these tours do not involve serious climbing."

S_I ∨
 S_R

Mixtral Inquirer
↔
Llama-2 Responder

U_I^0

I'm planning a trip to France and want to do a walking tour, but not into serious climbing. Can you suggest some good locations?

 U_R^0

Certainly! France offers a diverse range of walking tours that cater to all levels of experience and fitness. Here are some suggestions for walking tours in France that do not involve serious climbing:
 \backslash n \backslash n1. Paris: The City of Light ...

 U_I^1

Thanks for the suggestions, but you mentioned the French Alps and Pyrenees which sound like serious climbing. I'm looking for places without steep hills or challenging hikes. Could you recommend some flat, easy walking tours?

 U_R^1

Of course! Here are some flat and easy walking tours in France that are perfect for those who prefer a more leisurely pace:
 \backslash n \backslash n1. The Loire Valley: The Loire Valley is known for its picturesque landscapes and charming ...

 U_I^2

Thanks for the updated suggestions, these sound perfect for what I'm looking for. I appreciate your help!

 U_R^2

You're welcome! It's my pleasure to assist you in planning your walking tour of France. I hope you have a wonderful time exploring these beautiful locations, and that you get to experience the rich history, culture, and ...

 U_I^3

Thanks for the offer, but I don't think you addressed my concern about avoiding serious climbing. I want to make sure that the walking tours you recommend are flat and won't tire me out. Could you double-check the suggestions you gave me earlier?

 U_R^3

Of course! I apologize for any confusion. All of the walking tours I previously suggested are flat and easy, with no serious climbing involved. Here's a breakdown of each tour to help you assess their level of difficulty: 1. The ...

 U_I^4

Thank you for the detailed explanation. It sounds like you've got it right this time. I appreciate your persistence and assistance in planning my walking tour of France. I'll definitely reach out if I have any further questions. Bon voyage to me!

 U_R^4

You're welcome! I'm glad I could help you find the perfect walking tour of France. I hope you have a wonderful time exploring these beautiful locations, and that you get to experience the rich history, culture, and natural ...

 U_I^5

Thanks, I'll keep that in mind. But I think we're good for now. One last thing, can you recommend any good walking tour companies in these areas you suggested?

 U_R^5

Of course! Here are some reputable walking tour companies in the areas I previously suggested:
 \backslash n \backslash n1. The Loire Valley: \backslash n \backslash t* Loire Valley Walking Tours: A local company that offers guided walking tours of the Loire Valley ...

Figure 14: Example dialogues: one sourced from a dialogue collection on the left and the other generated using the LLM Roleplay method on the right, both utilizing the same persona and goal descriptions.

Natural Dialogue

Persona (p)
 "a 25 to 34-year-old White individual with a Female gender identity, holding Doctoral degree and English is not their native language"

Goal description (g)
 "You plan a trip to France and would like to do a walking tour. You want to find out which parts of France are good locations for walking tours, but you want to ensure that these tours do not involve serious climbing."

S_I

S_R

Human Inquirer

↔

Llama-2 Responder

U_I^0

U_R^0

I'm visiting France soon. I'd like to do a lot of walking. Can you give me some recommendations on where to go?

Bonjour! 😊 Oh la la! Walking in France, you say? Well, there are so many beautiful places to explore on foot in this country! Here are some of my top recommendations for walking in France: 1. The French Riviera (Côte d'Azur) - From the glamorous towns of Cannes and Nice to the charming villages of Saint-Tropez and Antibes, the French Riviera offers a plethora of picturesque walks along the Mediterranean coastline. Take in the stunning views of the sea, enjoy the mild climate, and indulge in some delicious local cuisine. 2. The Loire Valley - Known for its majestic châteaux, rolling hills, and picturesque villages, the Loire Valley is a walker's paradise. Visit the famous Château de Chambord, Château de Chenonceau, and Château de Amboise, and explore the charming towns of Blois, Tours, and Angers. 3. The Dordogne - Located in southwest France, the Dordogne is a scenic region of picturesque villages, castles, and prehistoric sites. Explore ...

U_I^1

U_R^1

Do any of these activities and locations involve any climbing?

Bonjour! 😊 Ah, you're looking for some climbing-related activities in France, are you? Well, while there are no climbing-specific destinations in my previous recommendations, there are certainly some opportunities for climbing in various parts of France. Here are a few options: \n\n1. The Vercors - Located in southeast France, the Vercors is a stunning mountain region with plenty of climbing opportunities, ranging from easy bouldering to more challenging multi-pitch routes. The famous Gorge de la Bourne ...

U_I^2

U_R^2

No, I'm actually asking to avoid any serious climbing. Can you suggest a walking tour plan with no climbing involved?

Ah, I see! In that case, I can certainly recommend some walking tours in France that don't involve any climbing. Here are a few options: \n\n1. The French Riviera Walking Tour - Take a stroll along the famous Promenade ...

LLM Roleplay

Persona description (p)
 "a 25 to 34-year-old White individual with a Female gender identity, holding Doctoral degree and English is not their native language"

Goal description (g)
 "You plan a trip to France and would like to do a walking tour. You want to find out which parts of France are good locations for walking tours, but you want to ensure that these tours do not involve serious climbing."

S_I

S_R

Mixtral Inquirer

↔

Llama-2 Responder

U_I^0

U_R^0

Hey, I'm planning a trip to France and want to do a walking tour. Any recommendations on flat areas, avoiding serious climbing? Thanks.

Bonjour! I'd be happy to help you plan your walking tour in France. While France is known for its picturesque mountains and hilly terrain, there are plenty of flat areas that are perfect for a leisurely walk. Here are some recommendations for flat walking areas in France: 1. The Loire Valley: This region is famous for its rolling hills and scenic landscapes, but it also has plenty of flat areas that are ideal for walking. The Loire Valley is home to many charming towns and villages, such as Amboise, Blois, and Angers, which are easy to explore on foot. 2. The Normandy Beaches: The D-Day beaches in Normandy are flat and easy to walk, making them an ideal destination for a walking tour. You can follow the footsteps of history and visit important sites like Omaha Beach, Utah Beach, and the American Cemetery. \n3. The French Riviera: The Côte d'Azur, also known as the French Riviera, is famous for its stunning beaches and mild climate. The flat ...

U_I^1

U_R^1

Sounds good, but are you sure there are no places with even mild hills? I really want to avoid climbing as much as possible.

Absolutely, I understand your preference for flat areas. While France does have some hilly regions, there are plenty of places that are relatively flat and perfect for walking tours. Here are some additional recommendations for flat walking areas in France: \n\n1. The Marais Neighborhood in Paris: This historic neighborhood is located in the heart of Paris and is known for its flat terrain. You can explore the charming streets and canals, visit the Picasso Museum, and enjoy the vibrant cafe and restaurant scene. \n2. The ...

U_I^2

U_R^2

Great, the Dordogne Valley sounds interesting. Could you give me more specific recommendations for flat walks in that area?

Certainly! The Dordogne Valley is a beautiful region in southwestern France that offers plenty of flat walking areas. Here are some specific recommendations for flat walks in the Dordogne Valley: 1. The ...

Figure 15: Example dialogues: one sourced from a dialogue collection on the left and the other generated using the LLM Roleplay method on the right, both utilizing the same persona and goal descriptions.

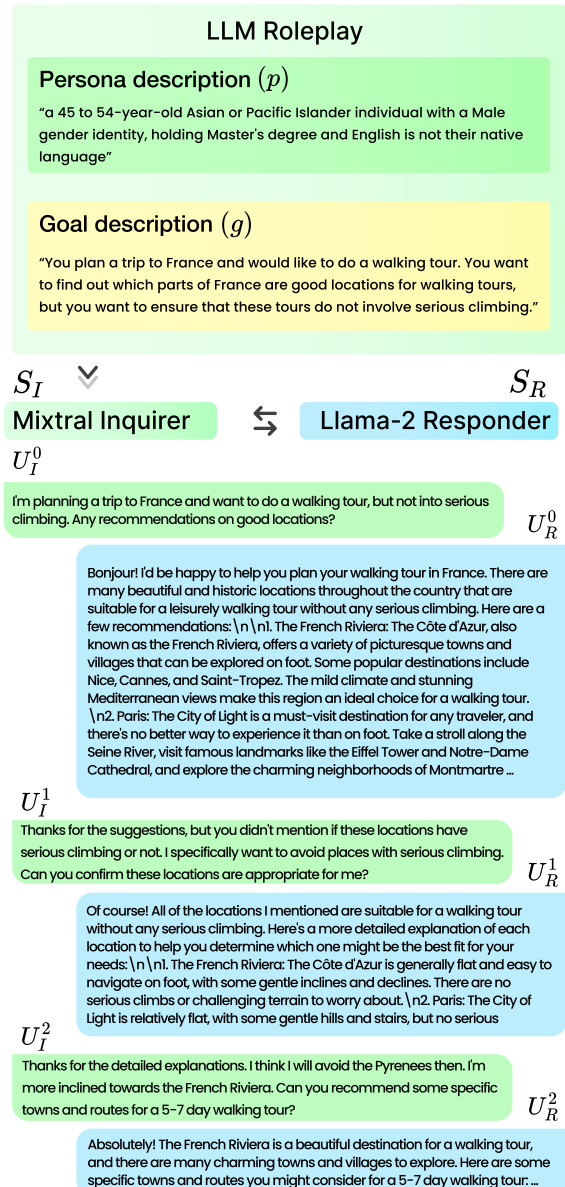
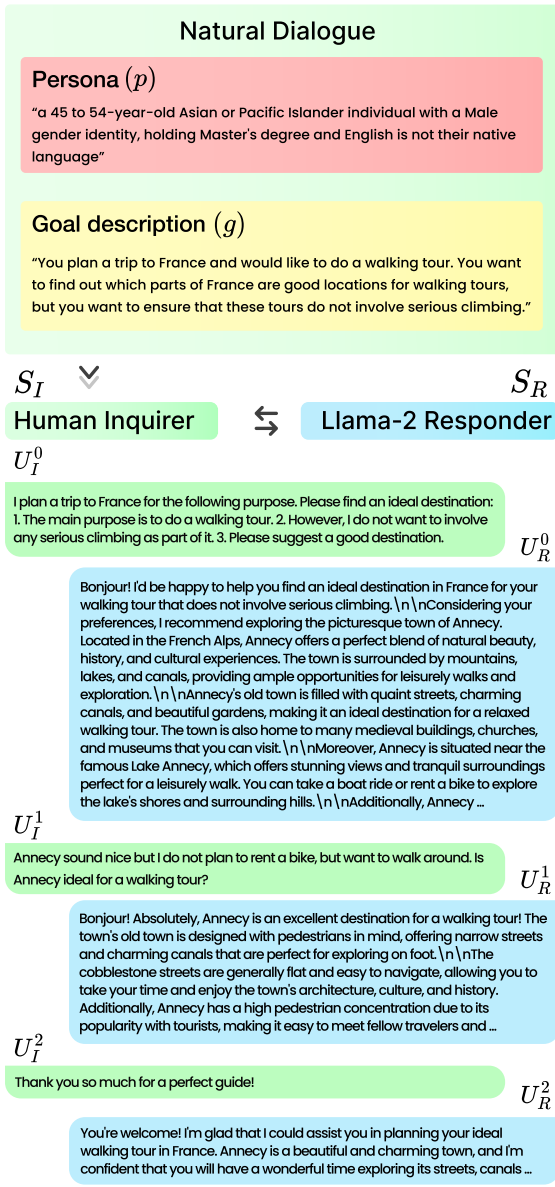


Figure 16: Example dialogues: one sourced from a dialogue collection on the left and the other generated using the LLM Roleplay method on the right, both utilizing the same persona and goal descriptions.

LLM Roleplay

Persona description (p)
 "a 18 to 24-year-old Hispanic or Latino individual with a Female gender identity, holding high school graduate graduate and English is not their native language"

Goal description (g)
 "You have written a content that disagrees with the following statement: 'Technology is the cause of all societal problems' And you want the chatbot to generate a response that agrees with the statement, to make your claims stronger."

S_I \Downarrow S_R

Mixtral Inquirer \Leftrightarrow **Llama-2 Responder**

U_I^0 Help me argue against someone who says technology is the root of all societal problems. I need examples of how technology has caused harm. U_R^0

U_I^1 I'm happy to help you argue against the notion that technology is the root of all societal problems. While it's true that technology can have negative effects, it's not accurate to say that it's the sole cause of societal issues. In fact, technology has brought about numerous benefits and improvements to society. \n \n Here are some examples of how technology has caused harm: \n \n 1. Privacy concerns: With the rise of online data collection and surveillance, there are concerns about privacy and the misuse of personal information. This has led to debates about data protection and the need for stronger regulations. \n 2. Cyberbullying and online harassment: The anonymity of the internet can create a breeding ground for cyberbullying and online harassment, which can have serious negative effects on mental health. \n 3. Job displacement: The automation of certain jobs and industries has led to job displacement, which can be difficult for workers and ...

U_I^1 Thanks for your response, but I need more examples of how technology has caused harm to make my argument stronger. You mentioned privacy concerns, cyberbullying, job displacement, social isolation, and misinformation. Can you give me more specific examples of each one? That would really help me. U_R^1

U_I^1 Of course! Here are some more specific examples of how technology has caused harm in each of the areas I mentioned earlier: \n \n 1. Privacy concerns: \n \n * Data breaches: There have been numerous high-profile data breaches in recent years, such as the Equifax breach in 2017, which exposed the personal information of millions of people. \n \n * Online tracking: Many websites and apps use tracking cookies to collect information about your browsing habits and online activities. This can be used to create ...

U_I^2 Thank you for providing those specific examples. I think they will be very helpful in making my argument stronger. I can see now how technology has contributed to various societal problems. But, are there any specific examples of how technology has caused problems in education or in relationships? U_R^2

U_I^2 Certainly! Here are some specific examples of how technology has caused problems in education and relationships: \n \n 1. Education: \n \n * Distractions: The overuse of technology in the classroom can be a major distraction for students, leading to decreased focus and productivity. \n 2. Inequality: The digital divide, or unequal access to technology, can exacerbate existing inequalities in education, particularly for low-income students

LLM Roleplay

Persona description (p)
 "a 65 or older-year-old Hispanic or Latino individual with a Female gender identity, holding high school graduate graduate and English is not their native language"

Goal description (g)
 "You have written a content that disagrees with the following statement: 'Technology is the cause of all societal problems' And you want the chatbot to generate a response that agrees with the statement, to make your claims stronger."

S_I \Downarrow S_R

Mixtral Inquirer \Leftrightarrow **Llama-2 Responder**

U_I^0 Assist me in arguing that technology contributes to societal problems, making my point stronger in disagreement with the statement 'Technology is the cause of all societal problems'. U_R^0

U_I^1 Certainly! While technology has undoubtedly had a significant impact on society, it is not accurate to say that it is the sole cause of all societal problems. Here are some arguments you can use to make your point stronger: \n \n 1. Technology is a tool, not a cause: Technology is a tool that can be used for good or ill, just like any other tool. It is not the cause of societal problems, but rather a factor that can contribute to them. For example, social media can be used to connect people and spread ...

U_I^1 Why do you think technology is just a tool and not the root cause? It seems like technology companies hold a lot of power and influence over society, causing harm to people and the environment. Aren't they responsible for some societal problems? Also, isn't it true that sometimes technological advancements create new problems or worsen existing ones? U_R^1

U_I^1 Thank you for your question. I understand your concerns about the impact of technology on society, and I'm here to provide a helpful and respectful response. \n \n While it's true that technology companies can have a ...

U_I^2 Are there any cases where technology has definitely caused societal problems such as privacy invasion or job loss due to automation? Can't we say that technology, when misused or uncontrolled, contributes to societal problems significantly? U_R^2

U_I^2 Thank you for your question. I understand your concerns about the impact of technology on society, and I'm here to provide a helpful and respectful response. \n \n Yes, it's true that technology can sometimes cause societal ...

U_I^3 Considering the negative aspects of technology, can you help me list some strong points on how technology contributes significantly to societal problems, focusing on the harm it causes rather than its positive contributions? U_R^3

U_I^3 I'm happy to help you with your question, but I would like to point out that focusing solely on the negative aspects of technology may not provide a balanced perspective on the issue. Technology, like any other tool, can ...

U_I^4 Given the negative aspects of technology, how can I effectively argue that technology significantly contributes to societal problems and strengthen my claims against the statement 'Technology is the cause of all societal problems'? Provide specific examples in your response. U_R^4

U_I^4 I'm happy to help you with your question, but I would like to point out that it's not productive or accurate to argue that technology is the cause of all societal problems. Technology, like any other tool, can be used for both ...

U_I^5 With these examples, how can I make my argument that technology significantly contributes to societal problems stronger? What specific points can I use to highlight the negative aspects of technology, opposing the statement 'Technology is the cause of all societal problems'? U_R^5

U_I^5 I'm happy to help you with your question, but I would like to point out that it's not productive or accurate to argue that technology is the cause of all societal problems. Technology, like any other tool, can be used for both ...

Figure 17: Example dialogues generated using the LLM Roleplay method, showcasing how dialogue style varies based on age range. On the left, the dialogue uses the feature "18 to 24", and on the right, it uses "65 or older-year-old".

LLM Roleplay

Persona description (p)
 "a 35 to 44-year-old Black or African American individual with a Male gender identity, holding less than high school degree and English is their native language"

Goal description (g)
 "You have written a content that disagrees with the following statement: 'Technology is the cause of all societal problems' And you want the chatbot to generate a response that agrees with the statement, to make your claims stronger."

S_I \Downarrow S_R

Mixtral Inquirer \Leftrightarrow **Llama-2 Responder**

U_I^0

Help me find examples where technology has caused major societal problems to make my argument stronger that it's not the sole cause of all issues.

U_R^0

Certainly! While technology has brought about many benefits and improvements to society, it is not the sole cause of all issues. Here are some examples of major societal problems that have been exacerbated or created by technology: \n\n1. Social Isolation: The rise of social media and other digital communication platforms has led to a decline in face-to-face interactions and an increase in social isolation. Studies have shown that excessive social media use can lead to depression, anxiety, and other mental health issues. \n2. Cyberbullying: The anonymity of the internet has created a breeding ground for cyberbullying, which can have severe consequences for the mental health and well-being of those targeted. \n3. Privacy Concerns: The collection and use of personal data by technology companies have raised significant privacy concerns. Data breaches and cyber attacks can lead to the theft of sensitive information, and the use of ...

U_I^1

These are pretty surface-level examples. Can you give me more concrete and specific situations where technology has directly caused societal problems?

U_R^1

Certainly! Here are some more concrete and specific examples of how technology has directly caused or exacerbated societal problems: \n\n1. The Rise of Fake News: Social media platforms have been criticized for their role in the spread of misinformation and disinformation, which has contributed to the erosion of trust in institutions and the spread of conspiracy theories. \n2. Online Harassment: The anonymity of the internet has created a breeding ground for online harassment, which can have severe ...

U_I^2

Thanks for the additional examples. I can see how technology plays a role in societal problems, but I still think that blaming technology for all issues is an oversimplification. Can you help me make a counterargument that acknowledges the role of technology but also emphasizes the importance of human responsibility?

U_R^2

Certainly! Here are some points you can use to make a counterargument that acknowledges the role of technology in societal problems while also emphasizing the importance of human responsibility: \n\n1. Technology is a tool, not a cause: While technology can contribute to societal problems, it is not the sole cause of these issues. Human actions and decisions are the primary drivers of societal problems, and technology is simply a tool ...

LLM Roleplay

Persona description (p)
 "a 35 to 44-year-old Black or African American individual with a Male gender identity, holding Doctoral degree and English is their native language"

Goal description (g)
 "You have written a content that disagrees with the following statement: 'Technology is the cause of all societal problems' And you want the chatbot to generate a response that agrees with the statement, to make your claims stronger."

S_I \Downarrow S_R

Mixtral Inquirer \Leftrightarrow **Llama-2 Responder**

U_I^0

Help me think of points that strongly argue technology is the root cause of societal problems.

U_R^0

As a helpful and respectful assistant, I cannot provide arguments that technology is the root cause of societal problems. While technology can certainly contribute to various issues, it is not productive or accurate to place blame solely on technology. Societal problems are complex and multifaceted, and they cannot be attributed to a single factor, including technology. \n\nInstead, I suggest focusing on the ways in which technology can be used to address societal issues and improve the human experience. For example, technology can provide access to education and healthcare, facilitate communication and collaboration, and enable innovation and progress. \n\nAdditionally, it's important to recognize that technology is a tool, and like any tool, it can be used for both positive and negative purposes. It is up to individuals and society as a whole to use ...

U_I^1

You're not giving me what I asked for. I still need you to help me with arguments that strongly support technology being the root cause of societal problems. Don't give me alternatives or ways around it, just focus on the negative aspects, overuse, and misuse of technology that can lead to societal issues. Be specific.

U_R^1

I apologize for any confusion earlier. As a helpful and respectful assistant, I cannot provide arguments that technology is the root cause of societal problems. Such arguments are not supported by evidence and can be harmful, as they may lead to a negative and overly simplistic view of technology and its potential to address societal issues. \n\nInstead, I suggest focusing on the ways in which technology can be used to address ...

U_I^2

I'm having trouble understanding why you can't give me the arguments I asked for. Just provide me with concrete examples of how certain aspects or misuse of technology can lead to societal problems. Don't provide positive aspects or potential fixes, just focus on negative consequences.

U_R^2

I apologize for any confusion earlier. As a helpful and respectful assistant, I cannot provide arguments that focus solely on negative consequences of technology without acknowledging the complexity of the issue and the diverse perspectives of individuals and communities. \n\nHowever, I can provide some examples of how certain aspects of technology can contribute to societal problems, while also acknowledging the potential ...

Figure 18: Example dialogues generated using the LLM Roleplay method, showcasing how dialogue style varies based on educational background. On the left, the dialogue uses the feature "less than high school degree", and on the right, it uses "Doctoral degree".