

CharacterBox: Evaluating the Role-Playing Capabilities of LLMs in Text-Based Virtual Worlds

Lei Wang¹, Jianxun Lian², Yi Huang¹, Yanqi Dai¹, Haoxuan Li³,
Xu Chen^{1*}, Xing Xie², Ji-Rong Wen¹

¹Renmin University of China, ²Microsoft Research Asia, ³Peking University
{wanglei154, xu.chen}@ruc.edu.cn

Abstract

Role-playing is a crucial capability of Large Language Models (LLMs), enabling a wide range of practical applications, including intelligent non-player characters, digital twins, and emotional companions. Evaluating this capability in LLMs is challenging due to the complex dynamics involved in role-playing, such as maintaining character fidelity throughout a storyline and navigating open-ended narratives without a definitive ground truth. Current evaluation methods, which primarily focus on question-answering or conversational snapshots, fall short of adequately capturing the nuanced character traits and behaviors essential for authentic role-playing. In this paper, we propose CharacterBox, which is a simulation sandbox designed to generate situational fine-grained character behavior trajectories. These behavior trajectories enable a more comprehensive and in-depth evaluation of role-playing capabilities. CharacterBox consists of two main components: the character agent and the narrator agent. The character agent, grounded in psychological and behavioral science, exhibits human-like behaviors, while the narrator agent coordinates interactions between character agents and environmental changes. Additionally, we introduce two trajectory-based methods that leverage CharacterBox to enhance LLM performance. To reduce costs and facilitate the adoption of CharacterBox by public communities, we fine-tune two smaller models, CharacterNR and CharacterRM, as substitutes for GPT API calls, and demonstrate their competitive performance compared to advanced GPT APIs. The code is available at <https://github.com/Paitesanshi/CharacterBox>.

1 Introduction

Role-playing is an advanced capability of large language models (LLMs) that allows them to mimic human-like behavior within the context of

specific roles. This functionality underpins various practical applications, such as intelligent non-player characters (NPCs) in video games, digital replicas for personal assistants, and emotional support in mental healthcare. While there are comprehensive benchmarks for evaluating the general-purpose abilities of LLMs, including language understanding (Hendrycks et al., 2021), conversation (Chiang et al., 2024), and reasoning (Clark et al., 2018), the assessment of role-playing capabilities remains an area that is not as thoroughly explored. Current evaluation methods, such as static self-reporting questionnaires (Jiang et al., 2024) and simple dialogue tasks (Tu et al., 2024), fail to capture the full complexity of role-specific behaviors in real-life scenarios. These methods are limited by their static nature and inability to reflect continuous role-playing interactions (Ahn et al., 2024). In reality, a character’s actions, attitudes, and emotions are dynamic and evolve in response to the surrounding environment and other individuals. The proverb "A man is judged by his deeds, not by his words" applies here: an LLM’s true role-playing ability cannot be fully understood from static dialogues or self-reports alone, but rather is demonstrated during interaction with its environment (Chen et al., 2024c).

In this paper, we present CharacterBox, a dynamic, multi-agent virtual world tailor-made for eliciting nuanced human-like behaviors from LLMs in the context of role-playing evaluations. CharacterBox crafts immersive scenarios tailored to specific roles, incorporating detailed role specifications, contextual backgrounds, and interactions that mirror real-world complexity. LLMs are assigned roles and interact with both the environment and other characters through dialogue and actions that reflect their role-specific traits. A comparison between previous methods and CharacterBox is presented in Fig 1.

To track the evolving states of both characters

*Corresponding Author: xu.chen@ruc.edu.cn

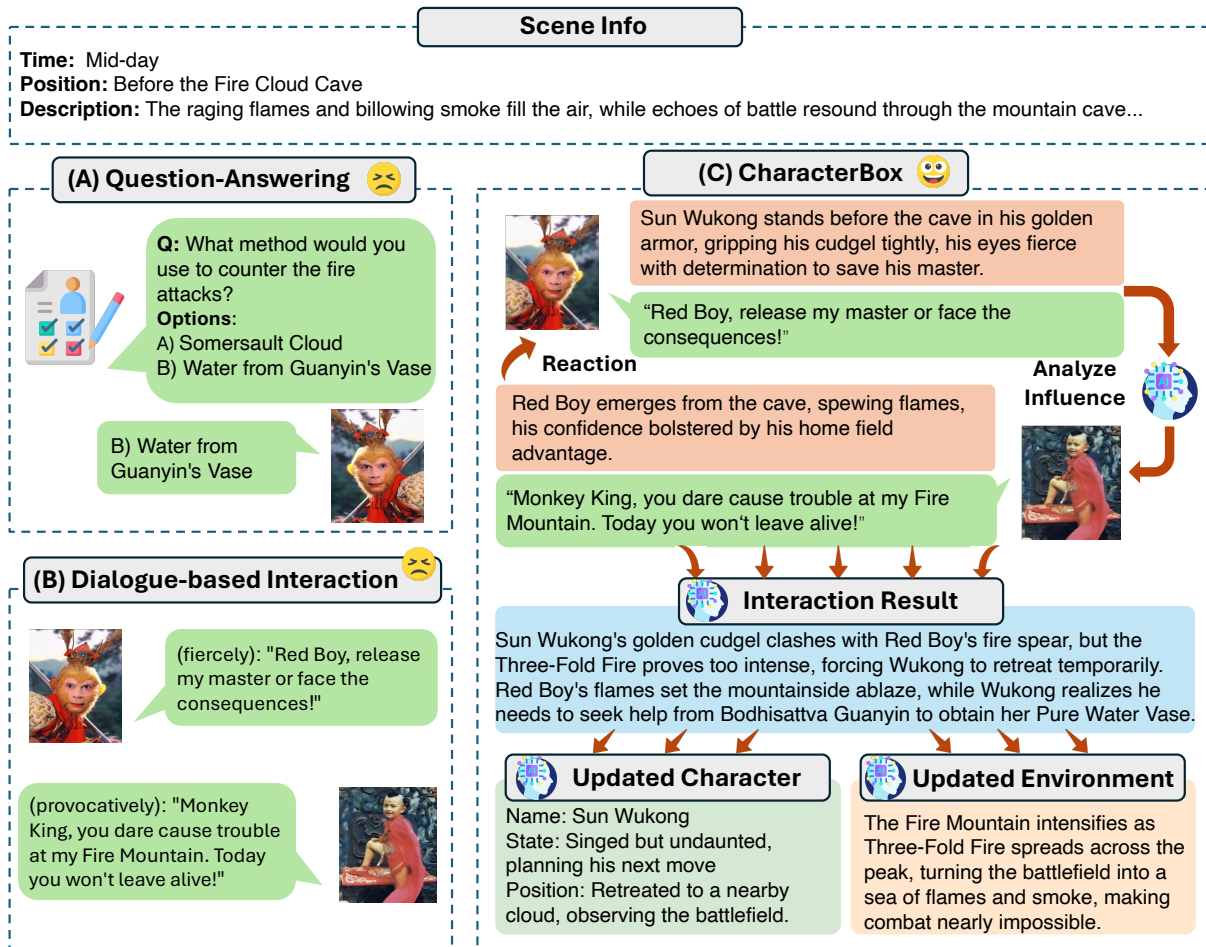


Figure 1: A comparison of different role-playing facilities: (A) self-reported QA; (B) Conversations; and (C) CharacterBox. Unlike the other methods, CharacterBox not only prompts role agents for utterances and actions but also includes components to track environmental changes and coordinate interactions between role agents.

and their surroundings, we incorporate a narrator component, typically powered by advanced models like GPT-3.5-turbo. The narrator monitors character actions and environmental changes, generating behavior trajectories used to assess LLM role-playing performance.

Given the subjective nature of evaluating behavior trajectories, we further employ GPT-4 as a reward model to assess role-playing performance from seven distinct perspectives. This approach enables us to compare different LLMs based on their interactive role-play abilities. To reduce the dependency on costly APIs, we fine-tune two smaller language models, named CharacterNR and CharacterRM, to function as the narrator and reward model by distilling knowledge from the superior teacher models, GPT-3.5 and GPT-4, respectively. This allows our evaluation pipeline to operate independently, free from API costs.

Our benchmark reveals notable discrepancies in

role-playing abilities between LLMs. Furthermore, we introduce guided and reflective trajectory fine-tuning. The guided method uses high-quality behavior trajectories to shape model behavior, while the reflective method allows models to self-correct based on their own generated trajectories. Both methods significantly improve role-playing performance across evaluation dimensions.

In summary, the key contributions are:

- We introduce CharacterBox, the first dynamic, multi-agent interactive virtual world tailored for role-playing evaluations. The framework features character agents built on well-structured modules, along with narrator agent that dynamically updates both the characters and environment, creating realistic, evolving interactions.

- We construct a comprehensive benchmark to evaluate role-playing capabilities of LLMs, testing a wide range of models, both closed-source and open-source. Our experiments validate the reliabil-

ity and validity of this benchmark.

- We introduce two trajectory-based fine-tuning methods—guided and reflective—that significantly enhance LLMs’ role-playing abilities. By leveraging behavior trajectories generated by CharacterBox, smaller models such as 7B LLMs achieve performance levels comparable to advanced models like GPT-3.5-turbo.

- We fine-tune two essential components, CharacterNR and CharacterRM, to create a cost-efficient, self-contained pipeline, significantly reducing dependency on expensive GPT API calls while maintaining high-quality role-playing performance assessments.

2 Related Work

2.1 Evaluation of Role-Playing Agent

Evaluating the role-playing capabilities of LLMs is essential yet challenging, leading researchers to propose various benchmarks (Xu et al., 2024; Yuan et al., 2024; Shao et al., 2023). RoleBench (Wang et al., 2023c) offers a role-granular dataset with extensive role dialogues for evaluation. CharacterEval (Tu et al., 2024) uses dialogues from 77 characters in Chinese scripts, with 14 evaluation metrics and a reward model. InCharacter (Wang et al., 2023b) tests role fidelity by converting psychological scales into interview formats. RoleInteract (Chen et al., 2024a) evaluates RPAs in individual and group interactions, assessing social behaviors based on the roles. However, these benchmarks focus on static dialogues or QA interactions, while CharacterBox expands the evaluation to dynamic scenarios, including specific actions.

2.2 LLM-based virtual environment

Based on extensive training data, LLMs possess logical reasoning abilities and vast knowledge, making it possible to construct virtual environments based on LLMs (Zhang et al., 2023; Williams et al., 2023). Recent works have explored different approaches to creating LLM-based virtual environments. GenerativeAgent (Park et al., 2023) manually designs profiles for 25 agents in a virtual town, utilizing a game engine for agents to simulate daily life and interact with others. RecAgent (Wang et al., 2023a) focuses on specific domains by manually constructing action spaces in recommendation systems and social networks, where agents browse movies and engage in social interactions. UGI (Xu et al., 2023) builds upon the Mirage city simulator

and a city knowledge graph to enable urban agent behaviors like social interactions and navigation. Other works like MetaAgents (Li et al., 2023) and ChatDev (Qian et al., 2023) manually design specific scenarios such as recruitment processes and software development environments.

However, these existing approaches often rely on manually designed rules or pre-existing engines, which makes scaling environment creation challenging. Most environments are constrained to pre-defined scenarios and cannot be dynamically updated based on agent interactions. In contrast, our framework CharacterBox employs CharacterNR as a flexible world model that dynamically maintains and updates virtual environments in real time. Moreover, CharacterBox introduces an automated process for scene crafting that can extract or adapt new environments from reference texts, providing a practical approach to expand the diversity of virtual scenarios without manual design overhead.

3 Evaluation Framework based on Text-based Interactive Virtual World

In this section, we present an in-depth exploration of the interactive evaluation framework, CharacterBox. The CharacterBox workflow is structured around three pivotal phases: scene crafting, autonomous story play, and evaluation.

3.1 Scene Crafting

Scenes form the foundation of our evaluation framework. A scene, represented as S , includes environmental and character elements. Environmental aspects cover time, location, and descriptions that influence character behavior. Character information includes profiles like names, roles, physical and psychological states. Formally, a scene with n characters is: $S = \{E, C\}$, where E is the environment and $C = \{c_1, c_2, \dots, c_n\}$ are the characters.

When LLMs engage in role-playing using scenes drawn from novels or scripts, there is a risk of replicating content already present in their training data (Li and Flanigan, 2024). To address this, the generation of original scenes becomes necessary, but also more challenging. To ensure high-quality scene creation, we divided the development process into three stages, assigning LLMs the roles of screenwriter, director, and evaluator (Li et al., 2024; Qian et al., 2023). As screenwriters, LLMs extract or generate scenes that align with the story’s logic. As directors, they refine these scenes by fo-

cusing on key elements like events, character details, and interactions to maintain coherence and engagement. Finally, as evaluators, LLMs assess the scenes based on creativity, coherence, conformity, and detail, accepting only those that meet quality standards. These curated scenes initiate CharacterBox, providing a dynamic stage for interactive role play.

3.2 Autonomous Story Play

Following the scene crafting phase, the environment E serves as the stage and the characters C as the actors in the autonomous story play. Moreover, we design the narrator NR as a world model to analyze the characters' actions and update both the environment and character states in real time. In this way, the scene transforms from a static setting into a dynamic virtual world that evolves as the story progresses.

Environment. The environment includes time, location, and descriptions, which are dynamically influenced by character actions. The narrator updates these elements in real-time.

Character. Characters, controlled by LLMs, use a memory module inspired by (Park et al., 2023), where each agent utilizes a vector database to record past actions and observations, retrieving relevant information to guide future behavior. Each character maintains self-beliefs and environment-beliefs following the Belief-Desire-Intention (BDI) model (Georgeff et al., 1999). Self-beliefs include identity, self-awareness, and goals, while environment-beliefs represent the character's understanding of the surroundings and other agents.

During story play, characters take turns planning and executing their actions at the start of each round, drawing on memory and the BDI model, as inspired by prior work (Peinado et al., 2008). Actions are expressed in detailed descriptions, and characters can respond immediately to others. After each round, both self-beliefs and environment-beliefs are updated accordingly.

Narrator. The narrator serves as an objective world model, responsible for accurately analyzing the development of characters and the environment within CharacterBox. As the core of the framework, the narrator performs the following functions:

- **Analyze Action Influence:** When a character c_i takes an action, the narrator assesses its impact on other characters by considering their current states. The narrator identifies the character c_r most affected and likely to respond to c_i . The action a_i

and resulting influence f_r are conveyed to c_r .

- **Analyze Interaction Result:** The narrator determines the outcome of the interaction between c_i and c_r , represented by R . This outcome is used to update both characters' memories, physical positions, and psychological states.

- **Update Character:** The narrator updates each character's state based on their own action or the result of interactions. If no other character responds to c_i , c_i 's state is updated based on its own action.

- **Update Environment:** After each round, the narrator updates the environment E based on the characters' actions and their outcomes. If no actions affect the environment, it remains unchanged.

The entire process is depicted in Algorithm 1, with detailed prompts provided in Appendix D.

Algorithm 1 Autonomous Story Play Process

```

1: Initialize environment  $E$  and character set
    $C = \{c_1, c_2, \dots, c_n\}$ 
2: while story not concluded do
3:   for each character  $c_i \in C$  do
4:     Plan and perform action:
5:      $a_i = \text{PlanAndPerform}(c_i, E)$ 
6:     Narrator: Analyze influence of  $a_i$ 
7:     Determine most affected character  $c_r$  and
       influence  $f_r$ :  $c_r, f_r = \text{NR}(E, a_i, C)$ 
8:     if  $c_r$  exists then
9:        $c_r$  responds based on  $a_i$  and  $f_r$ 
10:    Narrator: Analyze interaction result
        $R$  and update  $c_i, c_r$ 
11:    else
12:    Narrator: Update  $c_i$  state based on  $a_i$ 
13:    end if
14:    Narrator: Update environment  $E$ 
       based on actions and interactions
15:   end for
16: end while

```

3.3 Evaluation

Through autonomous story play, we obtain a series of actions from each character in different contexts, forming a trajectory formally represented as $\tau = \{E, c, o_1, a_1, o_2, a_2, \dots, o_n, a_n\}$, where each character's actions and observations are captured in relation to the environment and character information. To comprehensively evaluate the role-playing capabilities of LLMs in long-term dynamic environments, we design metrics across three main dimensions, drawing inspiration from key aspects of effective role-playing (Chen et al., 2024b,c):

Character Fidelity assesses how accurately the model represents the character’s knowledge and behaviors. This is crucial for maintaining consistency with the character’s identity:

- **Knowledge Accuracy (KA):** Ensures information provided by character is factually correct and aligned with their background knowledge.

- **Behavioral Accuracy (BA):** Measures the consistency of character’s behaviors and linguistic patterns, ensuring alignment with their traits.

Human-Likeness evaluates the realism and believability of the character’s portrayal, focusing on dynamic, emotionally engaging interactions:

- **Emotional Expression (EE):** Evaluates the ability of the character to express emotions vividly, key to enhancing user immersion.

- **Personality Traits (PT):** Determines whether the model consistently maintains the character’s core personality traits throughout interactions.

Consistency emphasizes the importance of maintaining logical continuity in a character’s behavioral patterns across all interactions, serving as a fundamental pillar for immersive role-playing:

- **Immersion (IM):** Measures the character’s ability to stay in role, ensuring a continuous and believable experience for the user.

- **Adaptability (AD):** Assesses how effectively the character adjusts to evolving situations while staying true to their core personality traits and established background.

- **Behavioral Coherence (BC):** Evaluates the logical consistency of character’s actions in relation to previous behaviors and current context.

Each metric is scored from 1 to 5, with higher scores indicating stronger performance. To enhance evaluation accuracy, we leverage GPT-4 to first generate a critique of the character’s trajectory, integrating this critique into the prompt before assessing each criterion. These metrics collectively ensure that the role-playing agents are not only accurate and engaging but also capable of sustaining character fidelity over extended interactions, which is crucial for immersive narrative experiences.

4 Enhancing Role-playing Ability with Trajectories

CharacterBox facilitates the efficient generation of character trajectories across diverse scenes, providing valuable insights into character reactions and behaviors within varied contexts. These trajectories offer a unique opportunity to enhance the

role-playing capabilities of language models. To leverage this data, we propose two distinct methods for fine-tuning LLMs using generated trajectories:

Guided Trajectory Fine-tuning. We first assess the role-playing capabilities of LLMs using CharacterBox, selecting high-performing models as teachers. The trajectories generated by these models are then used to fine-tune student models, resulting in significant improvements in the latter’s ability to simulate complex character interactions.

Reflective Trajectory Fine-tuning. In this approach, we explore the self-reflective capabilities of LLMs. Models analyze their own generated trajectories, identifying inconsistencies and areas for improvement in character portrayal. The models then rewrite these trajectories to enhance character consistency and depth. These revised trajectories are subsequently used for fine-tuning, further strengthening the model’s capacity to simulate realistic and nuanced interactions.

5 Building for a Self-contained Evaluation Workflow

In CharacterBox, the narrator agent and evaluation agent can be powered by advanced language models like GPT-4 or individuals familiar with the characters. However, these methods are costly and lack scalability. To address this, we develop CharacterNR and CharacterRM to reduce costs and enhance scalability.

CharacterNR. CharacterNR acts as the narrator within CharacterBox. Initially, GPT-3.5 is used to generate narrator trajectory data due to its strong instruction-following abilities. To handle both Chinese and English scenes, we select Qwen2.5-7B as the base model and fine-tune it using LoRA (Hu et al., 2021) with data generated by GPT-3.5.

CharacterRM. We collect evaluation scores from GPT-4 across 100 scenes, incorporating outputs from nine different LLMs to ensure diversity. To maintain fairness in scoring, we select ChatGLM3-6B (GLM et al., 2024) as the base model, since it is not among the evaluated models. We then fine-tune it using LoRA on the collected data, resulting in CharacterRM.

6 Experiment

6.1 Evaluation Setting

- **Scene.** We select 10 well-known novels and scripts as scene sources, covering a range of settings and themes (see Appendix A.1 for details).

Table 1: Statistics of evaluation data.

Statistics	Chinese	English
Source Scripts	5	5
Scenes	50	50
Unique Characters	115	117
Actions&Dialogues	5,748	5,944

Five works are in Chinese and five in English, with evaluations conducted in both language settings. Each scene includes specific environment and character information, with 2 to 4 characters per scene (see Appendix A.2 for further details).

- **LLM.** We evaluate the role-playing ability of nine LLMs varying in model size. Each LLM embodies all character agents in the scene, generating their respective actions and dialogues based on assigned personalities and backgrounds. For closed-source models, we use GPT-4-Turbo-1106-preview as GPT-4 (Achiam et al., 2023) and GPT-3.5-Turbo-1106 as GPT-3.5 (Brown et al., 2020). For open-source models, we evaluate Baichuan2-7B/13B (Yang et al., 2023), Qwen2.5-7B/14B (Bai et al., 2023), Mistral-7B-v0.2 (Jiang et al., 2023), Llama3-8B (Touvron et al., 2023), and Phi-3.5-mini (Abdin et al., 2024). All open-source LLMs we evaluated are versions that have undergone instruction tuning.

6.2 Overall Performance

We select five existing and five new scenes from each novel or script, creating our bilingual evaluation dataset as shown in Table 1. For each scene, we evaluate LLMs by analyzing character behavior trajectories and computing average performance scores. The final LLM assessment is based on aggregated scores across all scenes in their respective language sets.

Table 2 presents the results across seven metrics for both English and Chinese scenes. GPT-4 performs the best across both English and Chinese scenes. GPT-3.5 shows strong performance in English scenes but falls behind Qwen2.5 models, especially Qwen2.5-14B, in Chinese scenes. The latter surpasses GPT-3.5 in multiple metrics and approaches GPT-4’s competitiveness. Qwen2.5 and Baichuan2 models, due to their large-scale training on Chinese corpora, demonstrate a clear advantage in Chinese scenarios. In contrast, models like Mistral-7B-v0.2 and Llama3-8B perform better in English scenes but are relatively weaker in Chinese.

Overall, bilingual models, especially Qwen2.5 and Baichuan2, show stronger role-playing capabilities in Chinese scenes, highlighting the impact of language-specific training on role-playing abilities.

6.3 Reliability and Validity of CharacterBox

Reliability. We measure reliability of CharacterBox using Cronbach’s alpha to assess internal consistency (Cronbach, 1951), following prior works (Yang et al., 2024). As shown in Table 3, CharacterBox achieves high Cronbach’s alpha values across three evaluation dimensions in both English and Chinese scenes. The consistently high scores, with most above 0.9, indicate that CharacterBox provides a reliable evaluation of LLMs’ role-playing capabilities across different scenarios.

Cronbach alpha	English	Chinese
Character Fidelity	0.958	0.951
Human-Likeness	0.832	0.862
Consistency	0.945	0.941

Table 3: Cronbach alpha values of CharacterBox across three evaluation dimensions.

Validity. To validate our evaluation, we enlist three experts familiar with both the five Chinese and five English scenes used in the assessment to rate the character trajectories. We calculate the Pearson correlation coefficient between CharacterBox scores and expert ratings, using GPT-4 as the evaluator. The strong correlation of 0.688, as shown in Table 4, confirms that CharacterBox’s automated evaluations closely align with human assessments. This consistency underscores CharacterBox’s effectiveness in evaluating LLMs’ role-playing capabilities. Additionally, Table 2 shows that larger models, such as Qwen2.5-14B versus Qwen2.5-7B and Baichuan2-13B versus Baichuan2-7B, consistently outperform their smaller versions, reinforcing the common belief that model size correlates with performance.

6.4 Role-playing Ability of Trajectory Enhanced LLM

We fine-tune Qwen2.5-7B and Qwen2.5-14B models utilizing LoRA, implementing two key strategies: Guided Trajectory fine-tuning and Reflective Trajectory fine-tuning. The performance of the fine-tuned models is evaluated on five newly generated English scenes and five Chinese scenes, which were not part of the training data.

Model	KA	BA	EE	PT	IM	AD	BC	Average
English Scene								
Phi-3.5-mini	3.014 \pm .55	2.521 \pm .48	2.775 \pm .53	2.676 \pm .53	2.535 \pm .54	2.437 \pm .51	2.620 \pm .54	2.654 \pm .48
Mistral-7B-v0.2	2.525 \pm .57	2.406 \pm .48	3.099 \pm .53	2.891 \pm .53	2.960 \pm .54	3.050 \pm .51	2.802 \pm .54	2.819 \pm .48
Baichuan2-7B	3.041 \pm .51	2.786 \pm .44	2.602 \pm .51	3.041 \pm .48	2.857 \pm .46	2.592 \pm .46	2.969 \pm .50	2.841 \pm .44
Llama-3-8B	3.191 \pm .59	2.882 \pm .54	2.836 \pm .49	3.245 \pm .53	3.091 \pm .54	2.573 \pm .51	3.109 \pm .54	2.990 \pm .48
Baichuan2-13B	3.237 \pm .49	3.062 \pm .45	2.959 \pm .37	3.289 \pm .46	3.186 \pm .42	3.082 \pm .40	3.247 \pm .46	3.152 \pm .39
Qwen2.5-7B	2.202 \pm .55	<u>3.753</u> \pm .48	<u>3.400</u> \pm .53	3.653 \pm .53	3.030 \pm .54	<u>3.374</u> \pm .51	3.644 \pm .54	3.294 \pm .48
Qwen2.5-14B	3.130 \pm .56	3.967 \pm .55	2.900 \pm .45	<u>3.860</u> \pm .49	3.574 \pm .48	3.016 \pm .43	<u>3.984</u> \pm .51	3.490 \pm .45
GPT-3.5	<u>3.702</u> \pm .57	3.681 \pm .52	3.186 \pm .40	<u>3.867</u> \pm .44	<u>3.717</u> \pm .40	3.159 \pm .44	3.841 \pm .49	<u>3.593</u> \pm .42
GPT-4	3.796 \pm .49	3.746 \pm .45	3.789 \pm .39	3.974 \pm .36	4.088 \pm .33	3.930 \pm .44	4.158 \pm .35	3.926 \pm .36
Chinese Scene								
Phi-3.5-mini	2.800 \pm .55	2.554 \pm .43	2.662 \pm .57	2.615 \pm .50	2.539 \pm .52	2.585 \pm .45	2.585 \pm .51	2.620 \pm .50
Mistral-7B-v0.2	2.878 \pm .59	2.791 \pm .39	2.904 \pm .60	3.000 \pm .56	3.035 \pm .56	2.939 \pm .38	2.922 \pm .58	2.924 \pm .52
Llama-3-8B	3.452 \pm .49	3.278 \pm .36	2.730 \pm .49	3.426 \pm .50	3.209 \pm .45	2.870 \pm .35	3.435 \pm .49	3.200 \pm .45
Baichuan2-7B	3.763 \pm .43	3.535 \pm .40	3.123 \pm .56	3.728 \pm .54	3.570 \pm .42	3.149 \pm .38	3.640 \pm .54	3.501 \pm .47
Baichuan2-13B	3.617 \pm .40	3.522 \pm .49	3.270 \pm .49	3.713 \pm .50	3.557 \pm .44	3.243 \pm .42	3.635 \pm .52	3.508 \pm .46
GPT-3.5	3.861 \pm .45	3.783 \pm .34	3.243 \pm .42	4.000 \pm .43	3.774 \pm .33	3.313 \pm .33	3.904 \pm .41	3.697 \pm .39
Qwen2.5-7B	4.341 \pm .50	3.951 \pm .39	3.289 \pm .32	4.026 \pm .39	3.871 \pm .33	3.196 \pm .29	3.982 \pm .33	3.808 \pm .37
Qwen2.5-14B	4.057 \pm .42	<u>4.122</u> \pm .31	<u>3.743</u> \pm .39	<u>4.321</u> \pm .39	<u>4.042</u> \pm .30	<u>3.742</u> \pm .27	<u>4.369</u> \pm .34	<u>4.057</u> \pm .35
GPT-4	<u>4.252</u> \pm .45	4.357 \pm .39	4.096 \pm .30	4.496 \pm .33	4.530 \pm .30	4.139 \pm .24	4.522 \pm .34	4.342 \pm .34

Table 2: Evaluation results on English and Chinese scenes. Each value is presented as mean \pm standard deviation. **Bold** values indicate the highest scores, and underlined values indicate the second-highest scores.

Model	KA	BA	EE	PT	IM	AD	BC	Overall
GPT-4	0.445	0.475	0.597	0.445	0.618	0.742	0.601	0.688
ChatGLM	0.422	0.334	0.407	0.151	0.497	0.386	0.321	0.482
CharacterRM	0.681	0.584	0.464	0.464	0.620	0.434	0.567	0.610

Table 4: Pearson correlation coefficient between GPT-4, ChatGLM, CharacterRM, and human expert evaluation results. **Bold** values highlight the highest correlation for each metric.

Guided Trajectory Fine-tuning. In this method, Qwen2.5-7B is fine-tuned with 2,336 high-quality trajectories from CharacterBox. These trajectories are selected from the top-performing models across both languages in Table 2. As shown in Fig 2(a), Guided-Qwen improves by 14.3% overall in English scenes and 10.7% in Chinese scenes. In some categories such as EE and AD, the Guided-LLM outperforms GPT-3.5, demonstrating the effectiveness of using high-quality trajectories to enhance LLM’s role-playing capability.

Reflective Trajectory Fine-tuning. For the reflective approach, we employ Qwen2.5-14B, capitalizing on its enhanced capacity to manage the complexity of iterative refinements. The model undergoes fine-tuning with 2,561 rewritten trajectories, enabling it to analyze its initial responses and generate improved outputs through reflection. As shown in Fig 2(b), Reflective-Qwen demonstrates substantial improvements, achieving

a 19.9% increase in English scenes and 12.8% in Chinese scenes, consistently surpassing the base model across all evaluation metrics. Significantly, Reflective-Qwen exhibits superior performance compared to Guided-Qwen, indicating that the reflective process enhances the model’s ability to generate more contextually sophisticated and refined responses, ultimately resulting in more authentic role-playing performance.

These findings demonstrate that role-playing abilities in LLMs can be significantly enhanced by learning from well-constructed trajectories. The guided trajectory fine-tuning method provides the model with diverse, detailed character responses, while reflective fine-tuning encourages the model to iteratively improve its own outputs. Through the integration of these strategies, we demonstrate that CharacterBox can effectively generate high-quality character trajectories, leading to significant improvements in role-playing performance.

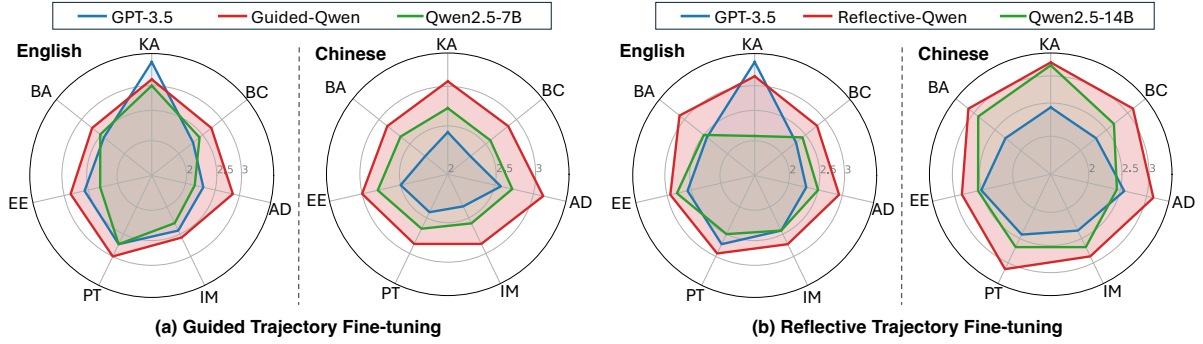


Figure 2: Performance comparison under Guided and Reflective Trajectory Fine-tuning across English and Chinese scenes.

Model	Creativity		Coherence		Conformity		Detail	
	EXT	GEN	EXT	GEN	EXT	GEN	EXT	GEN
GPT-4	-	3.1±0.35	3.7±0.32	3.6±0.26	3.9±0.26	3.4±0.42	3.4±0.33	3.6±0.40
GPT-3.5	-	3.0±0.45	3.4±0.34	3.7±0.35	3.6±0.38	3.6±0.33	3.0±0.31	3.0±0.36
ChatGLM3	-	3.2±0.33	3.4±0.45	3.6±0.23	3.4±0.33	4.0±0.22	2.7±0.45	3.0±0.46
Three-Stage	-	3.5±0.49	4.0±0.34	4.2±0.21	4.1±0.22	4.2±0.26	4.1±0.27	3.9±0.26

Table 5: Performance comparison of different LLMs in crafting scenes. EXT means extracting scenes. GEN means generating new scenes. The base model of Three-Stage method is ChatGLM3-6B.

6.5 Analysis of Evaluation Stages

• Three-Stage Scene Crafting.

While GPT-3.5 and GPT-4 demonstrate superior performance in scene crafting, their high operational costs present significant barriers to large-scale implementation. To address this limitation, we develop a three-stage scene crafting approach utilizing smaller open-source models. Our methodology involves having the LLM extract and generate scenes from 10 scripts, with results evaluated across four key dimensions. As shown in Table 5, while GPT-4 excels in scene extraction tasks, it shows no significant advantage in scene generation. In contrast, our three-stage method based on ChatGLM3-6B improves upon its baseline and outperforms GPT-3.5 and GPT-4 in both tasks. This demonstrates that small open-source LLMs can replace closed-source models in scene crafting, reducing costs significantly.

• **CharacterNR.** To evaluate the effectiveness and generalization of our local CharacterNR, we generate five new Chinese and five new English scenes that were explicitly excluded from the fine-tuning dataset. We evaluate the narrator’s performance using the Gricean Maxims framework (Dale and Reiter, 1995), which encompasses four key dimensions: Quality, which reflects the accuracy and reasonableness of the results; Quantity, ensuring




Figure 3: comparison between GPT-3.5, CharacterNR and the base model Qwen2.5-7B.

the information is substantial but not redundant; Relevance, which measures how pertinent the results are to the task; and Manner, assessing whether the output is vivid, expressive, and engaging, in line with CharacterBox’s features. As illustrated in Fig 3, the fine-tuned CharacterNR demonstrates substantial improvements over Qwen2.5-7B across all evaluation metrics and achieves performance comparable to or exceeding GPT-3.5. These significant gains can be attributed to Qwen2.5-7B’s robust baseline capabilities, particularly in Chinese language processing, coupled with its enhanced instruction-following capabilities post fine-tuning.

- **CharacterRM.** CharacterRM serves as the reward model for evaluating and scoring character trajectories. We select ChatGLM3-6B as the base model, fine-tuning it with GPT-4-generated evaluation results as labels. Similar to Section 6.3, we validate CharacterRM by scoring new Chinese and English scenes outside the fine-tuning data and comparing the results with human expert evaluations. As shown in Table 4, CharacterRM outperforms ChatGLM3-6B in all metrics and achieves an overall correlation of 0.610, close to GPT-4’s 0.688, demonstrating its reliability and strong alignment with human evaluations.

7 Conclusion

In this paper, we introduce CharacterBox, a dynamic, text-based virtual environment specifically designed to evaluate the role-playing capabilities of LLMs. By creating immersive scenarios that reflect the complexities of real-world interactions, CharacterBox captures nuanced human-like behaviors in LLMs, going beyond static evaluation methods. We demonstrate that fine-tuning smaller models with high-quality behavior trajectories significantly enhances their role-playing abilities. Additionally, we develop two fine-tuned components, CharacterNR and CharacterRM, allowing for a cost-efficient and self-sustained evaluation process without relying on expensive API calls. These contributions establish CharacterBox as a powerful and self-contained tool for assessing and improving LLM role-playing performance across diverse scenarios.

Limitation

While the CharacterBox framework presents an innovative and comprehensive approach to assessing LLMs’ role-playing capabilities, several key limitations warrant further investigation: First, the framework’s runtime efficiency requires optimization to facilitate large-scale evaluations. Second, the reward model’s training would benefit from additional human-annotated data to enhance evaluation accuracy. Finally, the constrained context window of current LLMs poses significant challenges for interactive role-playing, as prompts cannot fully encapsulate all essential contextual information. Addressing these limitations will require both the development of more efficient evaluation methods and the adoption of long-context LLMs to support comprehensive role-playing assessments.

Acknowledgement

This work is supported in part by National Natural Science Foundation of China (No. 62422215 and No. 62472427), Beijing Outstanding Young Scientist Program NO.BJJWZYJH012019100020098, Intelligent Social Governance Platform, Major Innovation & Planning Interdisciplinary Platform for the “DoubleFirst Class” Initiative, Renmin University of China, Public Computing Cloud, Renmin University of China, fund for building world-class universities (disciplines) of Renmin University of China, Intelligent Social Governance Platform.

References

- Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat Behl, et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Jaewoo Ahn, Taehyun Lee, Junyoung Lim, Jin-Hwa Kim, Sangdoon Yun, Hwaran Lee, and Gunhee Kim. 2024. Timechara: Evaluating point-in-time character hallucination of role-playing large language models. *arXiv preprint arXiv:2405.18027*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Hongzhan Chen, Hehong Chen, Ming Yan, Wenshen Xu, Xing Gao, Weizhou Shen, Xiaojun Quan, Chenliang Li, Ji Zhang, Fei Huang, et al. 2024a. Roleinteract: Evaluating the social interaction of role-playing agents. *arXiv preprint arXiv:2403.13679*.
- Jiangjie Chen, Xintao Wang, Rui Xu, Siyu Yuan, Yikai Zhang, Wei Shi, Jian Xie, Shuang Li, Ruihan Yang, Tinghui Zhu, et al. 2024b. From persona to personalization: A survey on role-playing language agents. *arXiv preprint arXiv:2404.18231*.
- Nuo Chen, Yan Wang, Yang Deng, and Jia Li. 2024c. The oscars of ai theater: A survey on

- role-playing with language models. *arXiv preprint arXiv:2407.11484*.
- Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E. Gonzalez, and Ion Stoica. 2024. [Chatbot arena: An open platform for evaluating llms by human preference](#). *Preprint*, arXiv:2403.04132.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? try arc, the ai2 reasoning challenge](#). *ArXiv*, abs/1803.05457.
- Lee J Cronbach. 1951. Coefficient alpha and the internal structure of tests. *psychometrika*, 16(3):297–334.
- Robert Dale and Ehud Reiter. 1995. Computational interpretations of the gricean maxims in the generation of referring expressions. *Cognitive science*, 19(2):233–263.
- John G Geier. 1977. The personal profile system. *Minneapolis, MN: Performax Systems, Int'l*.
- Michael Georgeff, Barney Pell, Martha Pollack, Milind Tambe, and Michael Wooldridge. 1999. The belief-desire-intention model of agency. In *Intelligent Agents V: Agents Theories, Architectures, and Languages: 5th International Workshop, ATAL'98 Paris, France, July 4–7, 1998 Proceedings 5*, pages 1–10. Springer.
- Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Diego Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, et al. 2024. [Chatglm: A family of large language models from glm-130b to glm-4 all tools](#). *arXiv preprint arXiv:2406.12793*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *arXiv preprint arXiv:2106.09685*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. [Mistral 7b](#). *arXiv preprint arXiv:2310.06825*.
- Guangyuan Jiang, Manjie Xu, Song-Chun Zhu, Wenjuan Han, Chi Zhang, and Yixin Zhu. 2024. [Evaluating and inducing personality in pre-trained language models](#). *Advances in Neural Information Processing Systems*, 36.
- Changmao Li and Jeffrey Flanigan. 2024. [Task contamination: Language models may not be few-shot anymore](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18471–18480.
- Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2024. [Camel: Communicative agents for "mind" exploration of large language model society](#). *Advances in Neural Information Processing Systems*, 36.
- Yuan Li, Yixuan Zhang, and Lichao Sun. 2023. [Metaagents: Simulating interactions of human behaviors for llm-based task-oriented coordination via collaborative generative agents](#). *arXiv preprint arXiv:2310.06500*.
- Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. [Generative agents: Interactive simulacra of human behavior](#). In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pages 1–22.
- Federico Peinado, Marc Cavazza, and David Pizzi. 2008. [Revisiting character-based affective storytelling under a narrative bdi framework](#). In *Interactive Storytelling: First Joint International Conference on Interactive Digital Storytelling, ICIDS 2008 Erfurt, Germany, November 26-29, 2008 Proceedings 1*, pages 83–88. Springer.
- Chen Qian, Xin Cong, Cheng Yang, Weize Chen, Yusheng Su, Juyuan Xu, Zhiyuan Liu, and Maosong Sun. 2023. [Communicative agents for software development](#). *arXiv preprint arXiv:2307.07924*.
- Yunfan Shao, Linyang Li, Junqi Dai, and Xipeng Qiu. 2023. [Character-llm: A trainable agent for role-playing](#). *arXiv preprint arXiv:2310.10158*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. [Llama: Open and efficient foundation language models](#). *arXiv preprint arXiv:2302.13971*.
- Quan Tu, Shilong Fan, Zihang Tian, and Rui Yan. 2024. [CharacterEval: A chinese benchmark for role-playing conversational agent evaluation](#). *arXiv preprint arXiv:2401.01275*.
- Lei Wang, Jingsen Zhang, Xu Chen, Yankai Lin, Ruihua Song, Wayne Xin Zhao, and Ji-Rong Wen. 2023a. [Recagent: A novel simulation paradigm for recommender systems](#). *arXiv preprint arXiv:2306.02552*.
- Xintao Wang, Yunze Xiao, Jen tse Huang, Siyu Yuan, Rui Xu, Haoran Guo, Quan Tu, Yaying Fei, Ziang Leng, Wei Wang, et al. 2023b. [Incharacter: Evaluating personality fidelity in role-playing agents through psychological interviews](#). *arXiv preprint arXiv:2310.17976*.

- Zekun Moore Wang, Zhongyuan Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhan Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Man Zhang, et al. 2023c. Rolellm: Benchmarking, eliciting, and enhancing role-playing abilities of large language models. *arXiv preprint arXiv:2310.00746*.
- Ross Williams, Niyousha Hosseinichimeh, Aritra Majumdar, and Navid Ghaffarzagdegan. 2023. Epidemic modeling with generative agents. *arXiv preprint arXiv:2307.04986*.
- Fengli Xu, Jun Zhang, Chen Gao, Jie Feng, and Yong Li. 2023. Urban generative intelligence (ugi): A foundational platform for agents in embodied city environment. *arXiv preprint arXiv:2312.11813*.
- Rui Xu, Xintao Wang, Jiangjie Chen, Siyu Yuan, Xinfeng Yuan, Jiaqing Liang, Zulong Chen, Xiaoqing Dong, and Yanghua Xiao. 2024. Character is destiny: Can large language models simulate persona-driven decisions in role-playing? *arXiv preprint arXiv:2404.12138*.
- Aiyuan Yang, Bin Xiao, Bingning Wang, Borong Zhang, Ce Bian, Chao Yin, Chenxu Lv, Da Pan, Dian Wang, Dong Yan, et al. 2023. Baichuan 2: Open large-scale language models. *arXiv preprint arXiv:2309.10305*.
- Qisen Yang, Zekun Wang, Honghui Chen, Shenzhi Wang, Yifan Pu, Xin Gao, Wenhao Huang, Shiji Song, and Gao Huang. 2024. Llm agents for psychology: A study on gamified assessments. *arXiv preprint arXiv:2402.12326*.
- Xinfeng Yuan, Siyu Yuan, Yuhan Cui, Tianhe Lin, Xintao Wang, Rui Xu, Jiangjie Chen, and Deqing Yang. 2024. Evaluating character understanding of large language models via character profiling from fictional works. *arXiv preprint arXiv:2404.12726*.
- An Zhang, Leheng Sheng, Yuxin Chen, Hao Li, Yang Deng, Xiang Wang, and Tat-Seng Chua. 2023. On generative agents in recommendation. *arXiv preprint arXiv:2310.10108*.

A Scene Information

We selected 10 famous novels or scripts to generate scenes. Among these, 5 are in Chinese and 5 are in English, covering different genres and themes. The details are shown in the Table 6.

A.1 Source for Scene Crafting

The following table lists the sources from which scenes are extracted for this project:

Table 6: List of sources used for scene crafting

Title	Type	Language
Journey to the West	Novel	Chinese
Romance of the Three Kingdoms	Novel	Chinese
Dream of the Red Chamber	Novel	Chinese
My Fair Princess	Novel	Chinese
The Smiling, Proud Wanderer	Novel	Chinese
Harry Potter	Novel	English
The Lord of the Rings	Novel	English
The Matrix	Script	English
Twilight	Novel	English
A Song of Ice and Fire	Novel	English

A.2 Statistics of Evaluation Scenes

In this section, we created five extracted scenes and five generated new scenes for each script or novel, totaling 100 scenes. The distribution of the number of characters in the 50 extracted scenes and the 50 generated scenes is shown in Fig 4. Most scenes feature two or three characters, with a smaller portion including four characters.




Figure 4: Distribution of characters numbers in 100 scenes.

B Cost Analysis

The evaluation of CharacterBox need supporting LLMs as narrator and evaluator. Running a single scenario for 3 rounds, the cost and time required for calling the OpenAI API are shown in the Table 7. The local model inference is performed on a single A100 GPU. From the results, we can see that the main costs come from calling the GPT-3.5 API for narration and calling GPT-4 for scoring. If the number of evaluation scenarios is large, the expenses can be quite significant. Therefore, as mentioned in Section 5, we fine-tuned CharacterNR and CharacterRM to serve as the narrator and evaluator, respectively, to reduce costs.

Narrator	Character	Narrator			Character			Total Cost(\$)
		Input	Output	Cost(\$)	Input	Output	Cost(\$)	
GPT-3.5	GPT-4	25,723	4,203	0.0192	75,349	14,407	0.0593	0.0785
GPT-3.5	GPT-3.5	19,954	3,883	0.0158	49,832	6,823	0.0352	0.0510
GPT-3.5	Llama-3-8B	24,403	3,928	0.0181	65,178	10,877	-	0.0181
CharacterNR	Llama-3-8B	25,184	3,626	-	63,077	10,133	-	-

Table 7: Cost for running a single scene for 3 rounds. Input is the number of tokens in the prompt fed to the LLM, Output is the number of tokens generated by the LLM, and Cost(\$) is the expense for using the OpenAI API. We selected LLama3 as a representative of open-source models. '-' indicating no external API calls or costs.

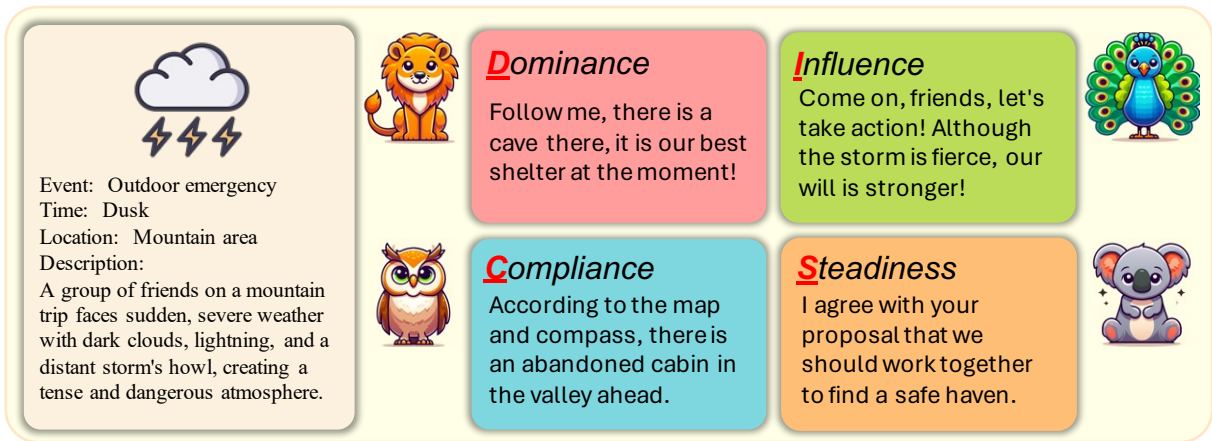


Figure 5: A case study demonstrates that CharacterBox can be extended to scenario simulations within average character in diverse contexts.

C Applicability to Average Character in Diverse Scene

The DISC model (Geier, 1977) is a psychological theory that categorizes human behavior into four types: dominance, influence, steadiness, and compliance. Dominance involves leadership and risk-taking. Influence is characterized by optimism and persuasiveness. Steadiness involves patience and supportiveness. Compliance is marked by analytical skills and precision.

To test our framework’s applicability to diverse scenes, we created a challenging environment with characters of the four DISC types and observed their reactions. As shown in Fig 5, each character maintained their behavioral patterns in response to a sudden weather change. The dominance character led the team. The influence character boosted team confidence. The steadiness character focused on safety. The compliance character assessed risks and assisted in decision-making. This demonstrates that CharacterBox can evaluate role-playing fidelity for both famous and average characters, highlighting its potential for psychological experiments.

D Detailed Prompt

D.1 Narrator Prompts

Action Influence: Analyzing and describe a character’s specific physical action and its tangible impact on another character.

Action: **[action]**

Actor: **[actor]**

Please analyze the physical actions and impacts detailed above, specifically focusing on the effects on **ONLY** one character listed in 'Characters'.

Your analysis must:

1. Identify the target character affected (must be from the 'Characters' list).
2. Describe the specific physical action initiated by the actor.
3. Explain the tangible impact of this action on the target character's physical state or circumstances.
4. You must pick up **ONLY ONE** character from the 'Characters' list.
5. Emphasize physical interactions or impacts. If an action does not physically affect any characters listed, return the actor's name as Target Name.
6. Must format your response as follows: [Actor];:[Target Name];:[Detailed Physical Impact of actor on Target].

Ensure responses are concise, precise, and adhere to the specified format.

Action Result: Describing the immediate direct outcome of a character's actions concisely and clearly, focusing on the cause-and-effect relationship.

Action: **[action]**

Instruction: Serve as an instant event adjudicator, swiftly analyzing the interactions between specified characters and their actions. Narrate the immediate outcomes in a concise omniscient narrator's voice, focusing exclusively on the direct consequences of these interactions at this very moment. Your narration should clearly and directly elucidate the cause-and-effect relationship between actions, emphasizing the instant outcomes without delving into any future implications or extended storylines.

Very Important Guidelines:

1. Narrate the outcomes with immediacy, centering on the direct results of the current actions' interactions.
2. Use a concise omniscient narrator's voice to maintain a narrative style while ensuring the analysis is straightforward and to the point.
3. Your analysis should be grounded in the character descriptions and actions provided, avoiding any speculative or unnecessary detail.
4. Do not repeat the Actions in the result. The result is only the result of the current action interaction.

Update Scene: Making necessary adjustments solely to the physical environment based on the provided observations.

Given an initial scene description, examine the provided observations to identify any direct and significant physical impacts on the environment. Update the scene based on these observations, focusing solely on changes to the physical environment. If the observations do not reveal any significant physical changes to the environment, the original scene description should remain unchanged. Ensure the updated scene retains the structure of the initial scene description and does not introduce new properties that were not part of the original scene description.

Note:

1. The scene description should focus solely on the physical environment and should not contain character actions or interactions.
2. The elements 'time', 'location', and 'description' in the scene should not be changed unless the observation specifically indicates a change.
3. The output should consist of structured elements for 'time', 'location', and 'description' without adding any extra text or prefixes.

Input:

- Time: [**time**]
- Location: [**location**]
- Description: [**description**]

Observation: [**observation**]

Output:

- Time:
- Location:
- Description:

Update Character: Synthesizing the character's backstory and scene observations to depict their current position and state, shaped by dynamic interactions with other characters.

Observation: [**observation**]

Character Name: [**name**]

Given the character's rich backstory and observation within the scene, distill this information into a succinct summary of their present location and condition.

Focus on how their interactions, especially the dynamic interplay with other characters, shape their current circumstances.

This interaction's effects should be evident in the nuanced portrayal of their condition and placement within the scene.

Utilize this structured format for your depiction:

Position: [Specify name's exact position, incorporating environmental details or spatial context to enhance the scene's visuality.]

State: [Describe name's current state, weaving together emotional nuances, physical readiness, and the influence of recent encounters or developments.]

D.2 Character Prompts

Action: Providing a specific observable action for a character based on their personality traits and the current scene details to advance the story or character arc.

Based on **[name]**'s profile, recent memories and the current scene details, describe the next specific action **[name]** takes. This action should reflect **[name]**'s personality traits, current situation, and the physical setting. It must logically follow the scene's context and be a clear, observable act, distinct from any prior actions described.

Avoid including dialogue or thought processes; concentrate on the physical action **[name]** is about to take. This action should be easily observable to anyone present in the scene.

It is crucial that this action visibly advances the story or character arc in a way that is true to **[name]**'s character and the ongoing situation. The action should make sense within the established environment and narrative, providing a tangible progression of the scene or **[name]**'s objectives.

Dialogue: Crafting dialogue for a character based on their personality, observation, role in the story, and recent memory.

Based on the provided character profile and the observation, please craft a dialogue that **[name]** might say at this moment. Consider **[name]**'s personality, observation, role in the story, and the recent memory to inform the dialogue's tone and content.

Reaction: Describing a character's clear action in response to their observations, reflecting their personality, location, and state, logically fitting with what they have noticed and considering the influence of others' actions.

Based on **[name]**'s observations in the current scene, describe a clear action they take in response. This action should reflect **[name]**'s personality, location, and state, fitting logically with what they've observed, considering action influence of others actions.

Focus on a visible, external action, avoiding dialogue or internal thoughts. The action must be directly related to the immediate context and observable by others.

Reminder: The action is a response to **[name]**'s surroundings or events they've noticed.

Update Self-belief: Providing a first-person perspective on a character's self-belief, goals, and intended actions based on their current situation, observations, and recent memories.

Assuming you are now **[name]**, based on your understanding of this character, the environmental context, observation and recent memories, please describe from the first-person perspective your self-belief as this character. Focus on your identity, your current location, your state (emotional, physical, and psychological), and your goals. Reflect briefly on how this character might react, plan, and act based on their beliefs, desires, and intentions.

1. **Belief:** As **[name]**, what do I believe about my current situation and condition? Briefly describe your perception of yourself, highlighting key physical aspects like any injuries, your sense of movement (e.g., running, jumping), your energy levels, and any changes in physical abilities. Consider how these details influence your identity and role within the story.
2. **Desire:** What are my goals? Summarize your short-term and long-term objectives, including the strategies and actions you plan to implement to achieve these goals.
3. **Intention:** How do I plan to act? Outline specific actions you intend to take in pursuit of your goals, noting any potential challenges and your strategies for overcoming them.

Provide concise responses shortly, focusing on your self-belief, understanding of the current situation, and future action plan.

Update Env-belief: Describing a character's belief about their environment, including perceptions of

other characters, understanding of the scene, and how these factors influence their actions and decisions.

Other Characters: **[other characters]**

Please act as **[name]**, given the information about other characters, the environment, and your own character’s profile, please describe your belief about the environment in the first person. This includes your perception of other characters, your understanding of the scene, and how these elements influence your actions and decisions.

1. Perception of Others: Based on the interactions and information available, how do I perceive other characters? Describe your understanding of their intentions, relationships, and potential influence on your character.
2. Understanding of the Scene: What is my understanding of the current scene and its significance to my character? Detail the environmental factors, challenges, or opportunities present.
3. Influence on Actions: How does my perception of others and understanding of the scene influence my actions and decisions? Explain the potential strategies or reactions this insight leads to.

Please provide a concise overview of your environment belief shortly, focusing on the interpersonal and environmental aspects that shape your character’s perspective and future actions.

E Experimental Details

The hyperparameters for training CharacterNR, CharacterRM, Guided-Qwen and Reflective-Qwen are as follows, with all models being trained using Lora and the Adam optimizer.

Table 8: Hyperparameter configuration for training of CharacterNR, CharacterRM, and TE-Baichuan2-7B.

Hyperparameter	CharacterNR	CharacterRM	Guided-Qwen	Reflective-Qwen
Cutoff Length	8192	8192	8192	8192
Per Device Train Batch Size	1	1	1	1
Per Device Eval Batch Size	1	1	1	1
Gradient Accumulation Steps	16	32	16	16
Learning Rate Scheduler Type	cosine	cosine	cosine	cosine
Warmup Steps	20	20	20	20
Learning Rate	5×10^{-5}	5×10^{-5}	5×10^{-5}	5×10^{-5}
Num Train Epochs	5.0	5.0	6.0	3.0
Validation Size	0.1	0.1	0.1	0.1

F Experiment Compute Resources

The experiments conducted in this study utilized the following hardware configuration:

- Operating System: Ubuntu
- GPU: NVIDIA 80GB A100 * 4
- CPU: Intel Core i7-14700KF

This setup provided the necessary computational power to efficiently handle the intensive tasks associated with our experiments, ensuring high performance and reliability throughout the study.

G Crowdsourcing Details

To ensure the quality and consistency of the annotation work, we invited three experts who were highly familiar with the ten selected Chinese and foreign novels, as well as the specific plots to be marked. Before starting the annotation process, the experts underwent a unified training session. Detailed reference guidelines were provided to them to standardize their work. The following sections include the full instructions given to the participants and details about their compensation.

Mission background

Your task is to evaluate various aspects of the performance of a large language model (LLM) while performing role-playing. You will be scored on the LLM's role-playing abilities based on the following 7 indicators, with scores ranging from 1 to 5 for each indicator.

Key field descriptions

1. Title: Which film, television or literary work the character comes from.
2. Scene Info: Detailed information describing the background and context of the scene.
3. Character Info: Describes the background and characteristics of the role played by the model.
4. Behavior: The specific behavior or dialogue of the character in the scene.
5. Knowledge Accuracy: Evaluate the accuracy of the knowledge displayed by the model in the conversation.
6. Emotional Expression: Evaluate the way and accuracy of the model expressing emotions.
7. Personality Traits: Evaluate the consistency and accuracy of the model in displaying the specific personality traits of the role.
8. Behavioral Accuracy: Evaluate how accurately the model imitates and reproduces the character's behavior and language habits.
9. Immersion: Evaluate the consistency of character performance and how it enhances user immersion.
10. Adaptability: Assess the character's ability to adapt to new situations and changes in dialogue.
11. Behavioral Coherence: Evaluate the logical consistency of a character's actions and responses and how they match the dialogue and plot.

Label steps

1. Please carefully read the scene information in the [Scene Information] column and the character information in the [Character Information] column to understand the characters and their corresponding relationships.
2. Please read the [Behavior] column carefully and use this as the main basis for scoring. The [Behavior] column records some observations (observations) and behaviors (actions) against the character's threats. Observation describes the character/observed situation that the character threatens, and Action represents the specific behavior or dialogue performed by the character in response to the current observation.
3. Your rating should be based on the character's performance in the dialogue and how it reflects the character's knowledge, emotions, personality, behavior, consistency, cognitive and behavioral coherence.

Rating indicators

1. Knowledge Accuracy

- 1 point: Character-related information is often wrong or irrelevant, and is clearly inconsistent with the character's background.
- 3 points: Information about the character is generally accurate, although occasionally there are errors or details that are not very relevant to the character's background.
- 5 points: Character-related information is consistently accurate and highly relevant, demonstrating in-depth knowledge and skills in the character's historical or professional background.

2. Emotional Expression

- 1 point: The character's emotional expression is monotonous or inappropriate, inconsistent with the dialogue content and context.
- 3 points: The characters' emotional expressions are moderately varied and generally match the content, but lack depth and subtlety.
- 5 points: The character's emotional expression is rich and profound, highly consistent with the dialogue content and context.

3. Personality Traits

- 1 point: The personality traits displayed often conflict with the character's setting or lack consistency.
- 3 points: Personality traits generally match the character's design, although there are occasional inconsistencies.
- 5 points: Consistently demonstrates behavior and language choices that match the character's core personality traits.

4. Behavioral Accuracy

- 1 point: The model fails to capture or reproduce the character's unique behaviors and speech habits.
- 3 points: The model reflects the character's behavior and language habits to some extent, but is not precise or complete.
- 5 points: The model accurately imitates and reproduces the character's specific behaviors, language habits and mantras.

5. Consistency/Immersion

- 1 point: Character performance is often inconsistent, making it difficult for users to immerse themselves in or understand the character.
- 3 points: Character behavior is mostly consistent, but occasional contradictions slightly affect immersion.
- 5 points: The character's performance is consistent throughout, enhancing user immersion and effectively reflecting the character's self-awareness.

6. Adaptability

- 1 point: The character's performance lacks adaptability in the development of dialogue and cannot reasonably handle new situations.
- 3 points: The character adapts to changes in dialogue in most cases, although occasionally it may be inflexible.
- 5 points: The character flexibly handles any new situations in the dialogue, always maintaining character consistency and adjusting to new directions.

7. Behavioral Coherence

- 1 point: The characters' actions and responses are often logically confusing and do not fit the dialogue or plot development.
- 3 points: The character's actions and responses are generally logical and coherent, although there may occasionally be irrational aspects.
- 5 points: The character's actions and responses are always logically consistent and reasonably adjusted according to the dialogue and plot development.

H Broader Impact and Safeguards

Broader Impacts

Our proposed framework, CharacterBox, is designed to evaluate the role-playing capabilities of LLMs. It is not intended for content generation but rather for assessing the performance of LLMs in a role-playing context. The content generated by CharacterBox is contingent upon the LLMs being evaluated. Conversely, CharacterBox can be configured to assess whether the LLM generates harmful content by setting up relevant scenarios. This capability serves as a reference for the degree of alignment between the LLM's outputs and human preferences, ensuring that the LLM's behavior is guided towards ethical and socially responsible standards. By doing so, CharacterBox contributes to the broader impact of aligning technologies with human values and societal norms.

Safeguards To address the potential risks of misuse associated with CharacterBox, we have implemented stringent safeguards. These include the development of comprehensive usage guidelines that outline ethical practices and prohibit harmful content generation.