

# COSER: Coordinating LLM-Based Persona Simulation of Established Roles

Xintao Wang<sup>1</sup> Heng Wang<sup>2</sup> Yifei Zhang<sup>1</sup> Xinfeng Yuan<sup>3</sup> Rui Xu<sup>1</sup> Jen-tse Huang<sup>4</sup> Siyu Yuan<sup>3</sup>  
Haoran Guo<sup>1</sup> Jiangjie Chen<sup>1</sup> Shuchang Zhou<sup>2</sup> Wei Wang<sup>1</sup> Yanghua Xiao<sup>1</sup>

## Abstract

Role-playing language agents (RPLAs) have emerged as promising applications of large language models (LLMs). However, simulating established characters presents a challenging task for RPLAs, due to the lack of authentic character datasets and nuanced evaluation methods using such data. In this paper, we present COSER, a collection of a high-quality dataset, open models, and an evaluation protocol towards effective RPLAs of established characters. The COSER dataset covers 17,966 characters from 771 renowned books. It provides authentic dialogues with real-world intricacies, as well as diverse data types such as conversation setups, character experiences and internal thoughts. Drawing from acting methodology, we introduce given-circumstance acting for training and evaluating role-playing LLMs, where LLMs sequentially portray multiple characters in book scenes. Using our dataset, we develop COSER 8B and COSER 70B, *i.e.*, advanced open role-playing LLMs built on LLaMA-3.1 models. Extensive experiments demonstrate the value of the COSER dataset for RPLA training, evaluation and retrieval. Moreover, COSER 70B exhibits state-of-the-art performance surpassing or matching GPT-4o on our evaluation and three existing benchmarks, *i.e.*, achieving 75.80% and 93.47% accuracy on the InCharacter and LifeChoice benchmarks respectively. Our code, dataset and models are available at: <https://github.com/Neph0s/CoSER>.



Figure 1. An example from COSER dataset, which provides comprehensive data types such as conversation dialogues and settings, plot summaries, character motivations and inner thoughts, authentically sourced from renowned books.

## 1. Introduction

Recent advances in large language models (LLMs) have facilitated the emergence of anthropomorphic cognition in AI (Kosinski, 2023; Shanahan et al., 2023). Role-playing language agents (RPLAs), *i.e.*, LLMs that simulate specific personas based on relevant data, have hence been popular (Park et al., 2023). RPLAs have been adopted to simulate personas of various types, including demographics, characters, or daily individuals (Chen et al., 2024a). They have

<sup>1</sup>Shanghai Key Laboratory of Data Science, College of Computer Science and Artificial Intelligence, Fudan University <sup>2</sup>StepFun <sup>3</sup>School of Data Science, Fudan University <sup>4</sup>Johns Hopkins University. Correspondence to: Yanghua Xiao <shawyh@fudan.edu.cn>, First author's email <xtwang21@m.fudan.edu.cn>.

inspired extensive applications including character chatbots, agents in video games, and digital clones for humans. This paper studies **RPLAs for established characters**, which represent a crucial yet challenging task beyond the naive portrayal of individual traits or stereotypes. Specifically, RPLAs should faithfully align with their characters’ complex backgrounds and capture their nuanced personalities.

Towards effective RPLAs, two major challenges persist in: 1) **Data**: High-quality datasets are lacking. Existing datasets are limited to dialogues between two characters, and lack necessary dialogue contexts and knowledge in other forms. Moreover, many datasets are synthesized by LLMs, compromising authenticity and fidelity to the origins (Wang et al., 2024a; Lu et al., 2024); 2) **Evaluation**: Current methods fall short in assessing complex character portrayals of LLMs. They typically focus on single-turn interactions with pre-defined question sets, and rely on either LLM-based judges or multi-choice questions. The former lack nuanced discrimination and suffer from bias issues (Li et al., 2024), while the latter only assess specialized aspects (Xu et al., 2024). Overall, there is a lack of authentic character data and appropriate evaluation methods based on such data.

In this paper, we introduce CoSER, a collection of authentic character data, along with open state-of-the-art models and evaluation protocol based on such data, for **Coordinating LLM-Based Persona Simulation of Established Roles**. The CoSER dataset is sourced from narratives and dialogues in 771 renowned books, processed via our LLM-based pipeline. CoSER differs from existing datasets in two fundamental ways: 1) CoSER extracts authentic, multi-character dialogues from acclaimed literary works, in contrast to LLM-synthesized question-answer pairs in previous work. Hence, our dataset maintains high source fidelity while exhibiting greater quality and complexity. 2) CoSER incorporates comprehensive types of data, as shown in Fig. 1: *i*) Besides character profiles and dialogues, CoSER encompasses plot summaries, character experiences, and conversation backgrounds, supporting various purposes including prompting, retrieval, model training and evaluation. *ii*) Conversations in CoSER capture characters’ actions and internal thoughts beyond surface-level speech, enabling RPLAs to simulate sophisticated cognitive and behavioral processes of humans, such as “[I’m nervous, but we have to do this] (Takes a deep breath) Alright, we ...”. We provide a clear comparison between CoSER and existing datasets in Table 7.

We introduce given-circumstance acting (GCA) for training and evaluating role-playing LLMs, leveraging CoSER dataset. Given a conversation with messages  $M$ , characters  $\mathcal{C}$  and setting  $\mathcal{S}$ , GCA requires an actor LLM to sequentially portray each character  $c \in \mathcal{C}$  to recreate the conversation, as illustrated in Fig. 2. During training, we train LLMs to portray each character  $c$ , on their authentic utterances

$M_c \subset M$ . As a result, we develop CoSER 8B and 70B, built on LLaMA-3.1 models (Dubey et al., 2024), which demonstrate true-to-life character portrayal and state-of-the-art performance on multiple RPLA benchmarks. For evaluation, GCA involves two steps: multi-agent simulation and penalty-based LLM judging. Given a test conversation  $M$ , we: 1) create a multi-agent system to simulate a conversation  $\bar{M}$ , where the actor LLM portrays each character  $c \in \mathcal{C}$  in the same setting  $\mathcal{S}$ ; 2) assess  $\bar{M}$  using penalty-based LLM critics, leveraging detailed rubrics and the original conversation  $M$ . GCA evaluation offers three advantages: First, it comprehensively reflects actor LLMs’ abilities via multi-agent simulation; Second, it is based on authentic scenes and groundtruth dialogues. Third, it provides expert-curated rubrics to guide LLM critics.

Our contributions are summarized as follows:

1. We introduce the CoSER dataset and models for RPLA research and applications. Our dataset comprises 29,798 authentic conversations and comprehensive types of data from 771 renowned books. Leveraging this dataset, we develop CoSER 8B and CoSER 70B, which are state-of-the-art models for RPLAs.
2. We propose given-circumstance acting for training and evaluating role-playing LLMs, drawing from established acting theory. Our evaluation comprehensively tests actor LLMs via multi-character simulation, while providing original dialogues and detailed rubrics to enhance LLM-based assessment.
3. Results of extensive experiments demonstrate the significant value of our dataset for the training, retrieval and evaluation of RPLAs. Notably, the CoSER models achieve state-of-the-art performance across four benchmarks and human evaluation for RPLAs.

## 2. Related Work

Using LLMs to simulate human personas represents a pioneering research direction. Early studies develop prototypes of RPLAs for fictional characters (Shao et al., 2023) and multi-agent systems that simulate human society (Park et al., 2023), while exploring the nature (Shanahan et al., 2023) and potential limitations (Liu et al., 2024a; Deshpande et al., 2023) of LLM role-play. Chen et al. (2024a) present a comprehensive survey of relevant research. Our work focuses on the role-play of established characters, with comparison to PersonaHub (Chan et al., 2024) in §A.6.

Specifically, an RPLA leverages LLMs to create a simulated persona  $\pi_c$  that simulates a real character  $c$  based on its persona data  $\mathcal{D}_c$ . Effective RPLAs require both comprehensive, high-quality data  $\mathcal{D}_c$  and advanced role-playing LLMs. Among various persona types, we focus on RPLAs

Table 1. Overview of CoSER and existing RPLA datasets. For characters, Num. count characters with profiles, and Expr. denotes structured character experiences. For conversations, Auth. indicates authentic dialogues or behaviors from the books, and Multi-Chara. denotes involving more than 2 characters. Num. (number), Conv. (conversation), and Summ. (summary) are abbreviations.

Dataset	Book	Character			Conversation				Message			Plot	
	Num.	Num.	Profile	Expr.	#Conv.	#Turns	Setting	Auth.	Multi-Chara.	Speech	Thought	Action	Summ.
Charater-LLM		9	✓		14,300	13.2	✓			✓			
ChatHaruhi		32	✓		54,726	>2		✓*	✓	✓			
RoleLLM		100	✓		140,726	2				✓			
HPD	7	113			1,191	13.2	✓	✓	✓	✓			✓
LifeChoice	388	1,462	✓		1,462	2	✓	✓					
CroSS-MR	126	126	✓		445	2	✓	✓					
CharacterGLM		250	✓		1,034	15.8	✓			✓			
CharacterEval		77	✓		1,785	9.3	✓	✓		✓		✓	
DITTO		4,002	✓		7,186	5.1				✓			
MMRole		85	✓		14,346	4.2				✓			
CharacterBench		3,956	✓		13,162	11.3				✓			
CoSER	771	17,966	✓	✓	29,798	13.2	✓	✓	✓	✓	✓	✓	✓

for established characters, which should faithfully align with their characters’ complex backgrounds and nuanced personalities.

**Datasets for RPLAs** Persona data  $\mathcal{D}_c$  describe the real persona  $c$  through various representations, including profiles (Yuan et al., 2024), dialogues (Wang et al., 2024a), experiences (Li et al., 2023) and multimodal information (Dai et al., 2024), etc. As shown in Table 7, existing datasets have several limitations. 1) Many are synthesized via LLMs’ responses to general instruction sets (Wang et al., 2024a; Chan et al., 2024) or character-specific questions (Shao et al., 2023), such as RoleBench (Wang et al., 2024a). However, LLM-synthesized data compromise authenticity and fidelity to original sources. 2) Human-annotated datasets such as CharacterEval (Tu et al., 2024) and CharacterDial (Zhou et al., 2023) offer improved quality, but are expensive and difficult to scale. 3) Several efforts extract authentic dialogues from fictional works, such as ChatHaruhi (Li et al., 2023) and HPD (Chen et al., 2023). However, they rely on human efforts for individual sources and are hence hard to scale as well. 4) Furthermore, existing datasets offer limited representations and forms, *i.e.*, mainly consisting of two-character or user-character question-answer pairs. These datasets support various purposes, including prompting, training, retrieval augmentation, and evaluation of RPLAs.

**Evaluation for RPLAs** Existing evaluation methods are based on either LLM judges or multi-choice questions (Chen et al., 2024a). LLM-judged methods typically elicit LLMs’ role-playing performance via predefined questions, and score the performance using LLM judges or reward models (Chen et al., 2024a). They assess various dimensions, including character-independent aspects such as anthropomorphism (Tu et al., 2024) and attractiveness (Zhou et al., 2023), as well as character-specific traits such as language style, knowledge, and personality (Wang et al., 2024a; Shao et al., 2023). However, LLM judges suffer from inherent bi-

ases, *e.g.*, length and position bias (Li et al., 2024), and may lack the necessary knowledge for character-specific evaluation. Other benchmarks evaluate role-playing LLMs through multiple-choice questions, assessing specific aspects such as knowledge (Shen et al., 2023), decision making (Xu et al., 2024; Yuan et al., 2024), personality fidelity (Wang et al., 2024b) and point-in-time portray (Ahn et al., 2024).

### 3. CoSER Dataset

In this section, we introduce the CoSER dataset, which covers authentic data of 17,966 characters from 771 renowned books. CoSER features its authentic, non-synthesized dialogues with real-world intricacies, and comprehensive data representations supporting various usages. In Table 1, we provide a comprehensive comparison with existing datasets. We illustrate our dataset’s design principles in §3.1, curation pipeline in §3.2, and statistical analysis in §A.1.

#### 3.1. Design Principles

As shown in Table 1, CoSER differs from previous RPLA datasets mainly in its: 1) rich data types, 2) internal thoughts and physical actions in messages, 3) environment as a role.

**Rich Types of Data** The persona data  $\mathcal{D}_c$  can represent a character  $c$  from fictional works in diverse forms, *e.g.*, narratives, profiles, dialogues, experiences, etc. Previous work focuses primarily on profiles and dialogues, which represent limited knowledge. Hence, we propose a more comprehensive set of data types that are: 1) Comprehensive: covering extensive knowledge about characters and plots from the books; 2) Orthogonal: carrying distinct, complementary information with little redundancy; 3) Contextual-rich: providing sufficient context to enable  $\pi_c$  to faithfully reproduce  $c$ ’s behaviors and responses in given scenarios.

Specifically, we organizes knowledge from books hierarchi-

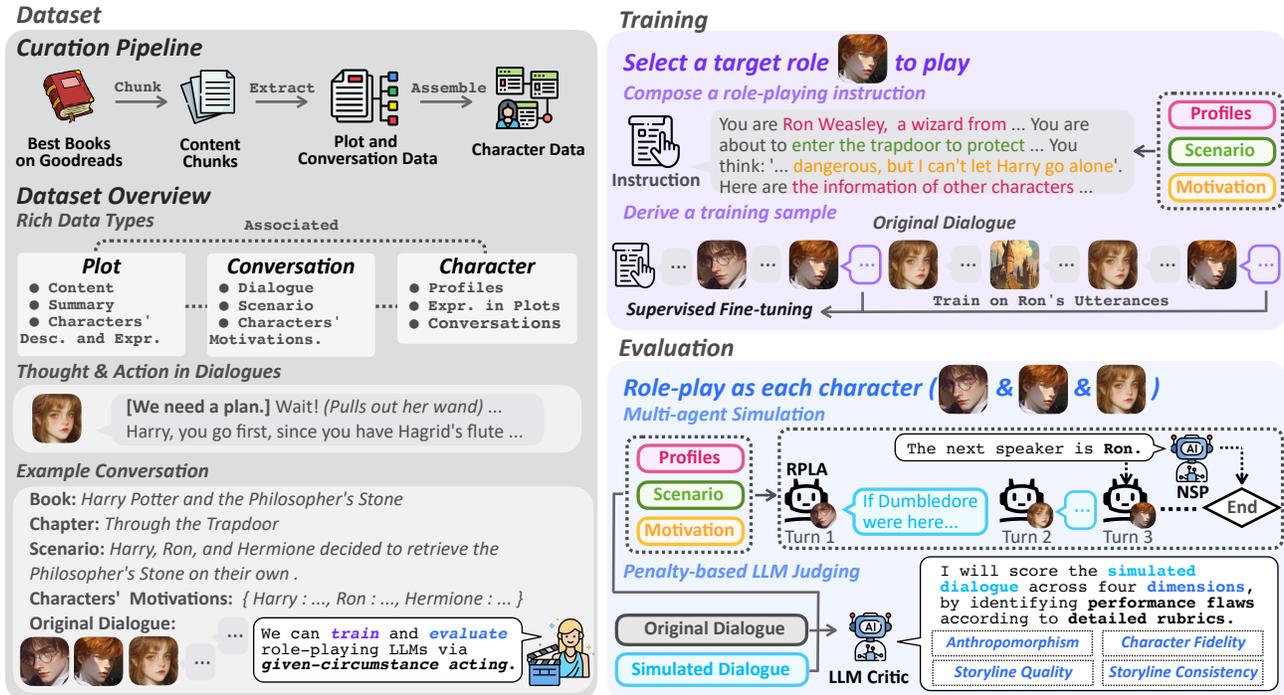


Figure 2. Overview of COSER’s dataset, training and evaluation. Left: The COSER dataset is sourced from renowned books and processed via LLM-based pipeline. It contains rich data types on plots, conversations and characters. Right: We apply given-circumstance acting to train and evaluate role-playing LLMs using these conversations. For training, each sample trains the LLM to portray a specific character in a conversation, using their original dialogue. For evaluation, we build a multi-agent system for conversation simulation given the same scenario, and assess the simulated dialogue via penalty-based LLM critics.

cally via three interconnected elements: plots, conversations and characters. Each **plot** comprises its raw text, summary, conversations in this plot, and key characters’ current states and experiences in this plot. A **conversation** contains not only the dialogue transcripts, but also rich contextual settings including scenario descriptions and characters’ motivations. **Characters** are associated with their conversations and plots, based on which we craft their profiles.

**Thoughts and Actions in Messages** Previous RPLA studies typically restrict RPLAs’ output space to verbal speech alone, limiting their ability to fully represent human interactions. In this paper, we extend the message space of RPLAs and character datasets into three distinct dimensions: speech ( $\mathcal{L}$ ), action ( $\mathcal{A}$ ), and thought ( $\mathcal{T}$ ), significantly enriching the expressiveness. For instance, an RPLA can convey silence by generating only thoughts and actions without verbal speech. The three dimensions are distinguished by markup symbols and function mechanisms:

- **Speech** is for verbal communications of characters.
- **Action** captures physical behaviors, body language, facial expressions, *etc.* Similar to tool use in agents (Weng, 2023), actions can be programmed to

trigger downstream events in multi-agent systems.

- **Thought** represents internal thinking processes, which enable RPLAs to simulate sophisticated human cognition. Thoughts should be invisible to others, forming information asymmetry (Zhou et al., 2024).

**Environment as a Role** In RPLA applications like AI TRPG<sup>1</sup> (Liang et al., 2023), LLMs often serve as world simulators that respond to players’ actions. To promote this ability, we consider environment as a special role  $e$ , which provide environmental responses such as physical changes and reactions from unspecified characters or crowds.

### 3.2. Dataset Curation

We curate the COSER dataset through a systematic LLM-based pipeline that transforms book content into high-quality data for RPLAs<sup>2</sup>. The details are as follows.

**Source Selection** Our dataset is sourced from most acclaimed literary works to ensure data quality and character depth. We identify the top 1,000 books on *Goodreads*’s *Best*

<sup>1</sup>Tabletop Role-Playing Games

<sup>2</sup>In this paper, we employ Claude-3.5-Sonnet (20240620).

*Books Ever* list<sup>3</sup>, and obtain the content for 771 books. As shown in Table 8, these books offer characters and narratives with literary significance and widespread recognition across diverse genres, time periods, and cultural backgrounds.

**Chunking** We segment book contents into chunks to fit in LLMs’ context window. We employ both static, chapter-based strategy and dynamic, plot-based strategy. Initially, we use regular expressions to identify chapter titles as natural chunk boundaries. Then, we merge adjacent small chunks and split large chunks to ensure moderate chunk sizes. However, static chunking neglects the storyline and truncates important plots or conversations. To address this, we implement dynamic plot-based chunking, *i.e.*, during data extraction, we also prompt LLMs to identify truncated plots or trailing content in the current chunk, and concatenate them with the subsequent chunk to ensure plot integrity.

**Data Extraction** We employ LLMs to extract plot and conversation data from book chunks, including (1) contents, summaries and character experiences of plots, and (2) dialogues and background settings of conversations. The extracted data representations are illustrated in Fig. 1 and introduced in §3.1. In the messages, speeches are always extracted from the original dialogues, while actions and thoughts can either be extracted or inferred by LLMs based on the context. For evaluation purposes, we hold out data from the final 10% plots in each book.

**Organizing Character Data** Based on the extracted data, we form the knowledge bases for characters in three steps. First, we unify character references by establishing name mappings between aliases and canonical names using LLMs, *e.g.*, mapping *Lord Snow* to an unified identifier *Jon Snow*. Second, we aggregate relevant plots and conversations for each character. Finally, we leverage LLMs to generate character profiles based on their extracted data, describing them from multiple perspectives including background, experiences, physical characteristics, personality traits, core motivations, relationships, character arcs, *etc.*

For technical details, including our prompts, engineering implementation, and handling mechanisms for exception caused by LLMs, please refer to §A.

## 4. Training and Evaluation via GCA

In this section, we introduce *given-circumstance acting* (GCA) for training and evaluating LLMs’ role-playing abilities using the CoSER dataset, as shown in Fig. 2.

### 4.1. Given-Circumstance Acting

In *Konstantin Stanislavski’s* acting methodologies, given-circumstance acting is a fundamental approach where actors

are trained and judged through performance within specified conditions including environmental context, historical events and personal conditions (Stanislavski, 2008).

We propose to adapt this approach to a framework that trains and evaluates LLMs’ role-playing skills, leveraging the comprehensive data in CoSER. In this framework, given a conversation with dialogue messages  $M$ , involved characters  $\mathcal{C}$ , and contextual setting  $\mathcal{S}$ , an actor LLM sequentially plays the role of each character  $c \in \mathcal{C}$  to simulate the conversation.

### 4.2. GCA Training and CoSER Models

We fine-tune LLMs’ role-playing abilities through GCA. Each training sample is derived from a conversation and one of its character  $c$  in CoSER dataset, and LLMs are trained on  $c$ ’s utterances  $M_c$ . Specifically, we first compose a role-playing instruction  $i_c$  comprising the scenario description, the character’s profile  $p_c$  and motivation, and profiles of other involved characters, which provide comprehensive context for role-playing. The original dialogue messages are denoted as  $M = [m_1, \dots, m_T]$ , where  $T$  is the number of turns. Then, the training sample  $[i_c, m_1, \dots, m_T]$  is a concatenation of the instruction and messages, where the character’s messages  $M_c \subset M$  are treated as outputs for optimization, and the other parts serve as inputs.

We train CoSER 8B and CoSER 70B based on LLaMA 3.1 Instruct models (Dubey et al., 2024), using 90% books in our dataset. To effectively support diverse use cases of RPLAs, our training samples cover extensive settings: *i)* The CoSER dataset contains massive characters and conversation settings from extensive books. We train models on all characters in each conversation, ranging from major characters with detailed profiles to minor roles driven only by the context; *ii)* To simulate real use cases, we incorporate role-playing instructions in diverse formats through instruction templates of varying formats. Besides, we consider different combinations of available data by including or excluding: profiles of other characters, plot summaries, and characters’ motivations; *iii)* We train models both with and without characters’ internal thoughts in the extracted dialogues.

We extend CoSER’s training beyond character role-playing to develop complementary capabilities in environment modeling and next speaker prediction (NSP), which facilitates RPLA applications. To maintain models’ general abilities, we augment our training data with the Tulu-3 dataset (Lambert et al., 2024). Please refer to §B for more details.

### 4.3. GCA Evaluation

Evaluating role-playing LLMs remains a significant challenge, primarily in two aspects: *i)* providing appropriate scenarios to elicit role-playing performance, and *ii)* properly assess the performance. Towards these challenges, we

<sup>3</sup>[https://www.goodreads.com/list/show/1.Best\\_Books\\_Ever](https://www.goodreads.com/list/show/1.Best_Books_Ever)

propose GCA evaluation for actor LLMs’ role-playing abilities, comprising two stages: multi-agent simulation and penalty-based LLM judging, as illustrated in Fig. 2.

**Multi-agent Simulation** For a test conversation  $M$ , we build a multi-agent system to simulate a conversation  $\bar{M}$ , in the same setting as  $M$ . We create an RPLA  $\pi_c$  for each character  $c \in \mathcal{C}$  using the actor LLM. We provide RPLAs with comprehensive data as described in §4.2: scenario descriptions and involved character profiles offer crucial context, and character motivations promote RPLA proactiveness and a natural conversation flow. Following §3.1, RPLAs are instructed to output in the speech-action-thought format. Each RPLA’s motivations and inner thoughts are inaccessible to other RPLAs. We adopt an NSP model to select the speaker of each turn from  $\mathcal{C} \cup \{e\}$ , and another LLM as the environment model  $\pi_e$  to provide environmental feedback. The simulation ends upon an  $\langle \text{END} \rangle$  signal from NSP, or reaching the maximum of 20 turns. In this way, we obtain a multi-turn, multi-character simulation that comprehensively reflects the actor LLMs’ role-playing abilities.

In addition, we introduce a continue-from parameter  $k$ , where the simulation starts from the first  $k$  original messages in  $\mathcal{M}$ . Setting  $k > 0$  controls the story direction and language style, similar to in-context learning. Hence, it enables more controlled evaluation and reduces the influence of different language styles of LLMs.

**Penalty-based LLM Judging** In this stage, we assess the simulated dialogue  $\bar{M}$  via LLM critics. Different from previous LLM-as-a-judge methods for RPLA evaluation, our LLM critics: *i*) apply penalty-based scoring by identifying role-playing flaws following detailed rubrics, and *ii*) leverage the original conversation  $M$  as reference.

Specifically, we employ LLM critics<sup>4</sup> to identify flaw instances  $\mathcal{F}$  in  $\bar{M}$  of specific rubrics, such as “*deviate from the original conversation*” or “*lack initiative and goals*”, instead of directly outputting a score in previous work (Wang et al., 2024a; Tu et al., 2024). Each flaw  $f$  is assigned a severity  $v_f$  from 1 (minor) to 5 (severe). The initial score for each dimension is calculated as  $s = 100 - 5 * \sum_{f \in \mathcal{F}} v_f$ .

We consider four evaluation dimensions, as shown in Fig. 3. The rubrics are derived from human-annotated issues in extensive human-RPLA conversations from real users and our multi-agent simulations, as well as dimensions from previous work (Shanahan et al., 2023; Shao et al., 2023; Chen et al., 2024a). For more informed evaluation, LLM critics are provided with additional materials, *i.e.*, the original conversation  $M$  and plot summary, besides data available to actor LLMs. Each dimension is assessed in independent LLM requests. The evaluation dimensions and their rubrics are summarized as follows:

<sup>4</sup>This paper uses GPT-4o as the critic LLM by default.

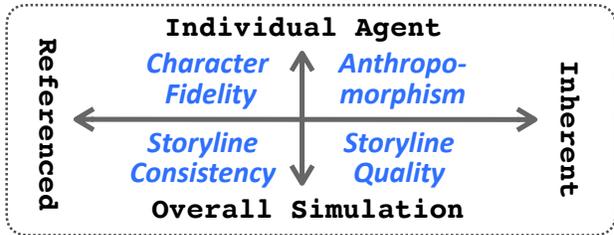


Figure 3. We divide the RPLA evaluation into four quadrants (dimensions). The X-axis represents the evaluation perspective: Referenced (comparing with data in the source book) versus Inherent (assessing standalone quality). The Y-axis indicates the evaluation scope: Individual Agent versus Overall Simulation.

1. **Anthropomorphism:** Evaluates whether RPLAs behave in a human-like manner, with rubrics covering self-identity, emotional depth, persona coherence, and social interaction.
2. **Character Fidelity:** Assesses whether RPLAs faithfully portray their characters, with rubrics examining language style, knowledge and background, personality and behavior, and social relationships.
3. **Storyline Quality:** Evaluates whether the simulated conversation develops naturally, with rubrics focusing on narrative flow and logical consistency.
4. **Storyline Consistency:** Measures alignment between the simulated conversation  $\bar{M}$  and original dialogue  $M$ , *i.e.*, whether RPLAs’ reactions (emotions, attitudes, behaviors) remain consistent with the original.

As longer simulations naturally make more flaws, we implement length correction to reduce bias in LLM judging following Dubois et al. (2024). Specifically, we obtain the length corrected score as  $s = 100 - 5 * \sum_{f \in \mathcal{F}} v_f + \lambda |\bar{M}|$ , where  $\lambda$  is set to 1.5 based on statistical analysis in §C.3. For detailed prompts and rubrics, please refer to §G.

## 5. Experiments

### 5.1. Settings

**Evaluation Protocol** We evaluate LLMs’ role-playing abilities through GCA on CoSER Test, a test set of held-out conversations from the final 10% of each book. CoSER Test contains 200 conversations, with 100 from books used in CoSER training and 100 otherwise. We employ GPT-4o as LLM critics and environment model, and CoSER 70B for NSP. We exclude characters’ inner thoughts for LLM critics, and set the continue-from parameter  $k = 0$ . The details are in §C.

**Metrics** We report LLM-judged scores for each dimension, and their average as the overall score. For analysis, we

also evaluate two traditional metrics based on N-gram, *i.e.*, BLEU (Papineni et al., 2002) and ROUGE-L (Lin & Och, 2004) compared against original dialogues. Additionally, we report win rates versus GPT-3.5 and GPT-4o in §D.

**Models** Our experiments cover numerous LLMs: *i)* **Close models**, including Minimax Abab7-preview<sup>5</sup>, Doubao-pro<sup>6</sup>, Step-1-Flash and Step-2<sup>7</sup>, GPT-3.5 (OpenAI, 2022), GPT-4o and GPT-4o Mini<sup>8</sup>, Gemini-1.5-Pro<sup>9</sup>, Claude-3-Haiku and Claude-3.5-Sonnet<sup>10</sup>; *ii)* **Open models**, including CoSER 8B and 70B, LLaMA-3.1-Instruct 8B and 70B (Dubey et al., 2024), Qwen-2-Instruct 7B and 72B (Yang et al., 2024), Vicuna-13B-1.5 (Zheng et al., 2023), Mixtral-8x7B (Jiang et al., 2024), DeepSeek-V3 (Liu et al., 2024b) and Higgs-Llama-3-70B<sup>11</sup>.

**Other Benchmarks** We also evaluate CoSER models on existing RPLA benchmarks based on multi-choice questions instead of LLM judges, including InCharacter (Wang et al., 2024b) for personality tests, LifeChoice (Xu et al., 2024) for decision-making, and CroSS (Yuan et al., 2024) for motivation recognition. For InCharacter, we report its accuracy on the Big Five Inventory (BFI).

## 5.2. Main Results

**Performance of Various LLMs on CoSER Test** We apply CoSER Test to evaluate extensive LLMs. The results shown in Table 2 are averaged across three runs, from which we observe that: 1) CoSER 70B achieves state-of-the-art performance across both LLM-judged and N-gram-based metrics. For LLM-judged metrics, CoSER 70B outperforms all open models and shows competitive performance with GPT-4o. CoSER 8B similarly outperforms models of comparable scale. For N-gram-based metrics, CoSER models demonstrate substantial improvements over existing models, exceeding the second-best performance by 58% on BLEU; 2) Among all models, GPT-4o, Gemini Pro, Claude-3.5-Sonnet, Doubao-pro, Step-2, Qwen-2-72B, and CoSER 70B demonstrate superior performance, achieving average scores above 57%; 3) Table 13 presents LLM performance separately for test splits from books included or excluded for CoSER training. The results show consistent trends across both splits, validating that CoSER models maintain strong performance on out-of-domain characters.

**Authentic Conversations from High-quality Novels Improve LLMs’ Role-playing Ability** According to Table 2, CoSER models demonstrate significant improvements

over their LLaMA 3.1 baselines. In contrast, Higgs-LLaMA-3 70B, fine-tuned on synthesized dialogues, performs below LLaMA-3.1 70B. These results highlight the importance of high-quality, authentic role-playing data for LLM training.

**Conversation Continuation Enables More Controlled Evaluation** Table 14 shows experiment results when multi-agent systems start from the first  $k = 3$  original messages. In this setting, model obtain higher scores compared to simulations from scratch ( $k = 0$ ), with reduced performance gaps between different models, especially for BLEU and ROUGE-L results. For example, the average score gap between Qwen-2 72B and 7B decreases from 11.5% ( $k = 0$ ) to 8.8% ( $k = 3$ ). This improvement occurs because the  $k = 3$  original messages guide the story direction and language style, particularly benefiting smaller models that typically struggle with complex role-playing instructions.

**Results on Other Benchmarks** We evaluate CoSER and other models on existing benchmarks for RPLAs based on multi-choice questions. As shown in Table 3, CoSER 70B achieves state-of-the-art performance across these benchmarks. Notably, CoSER 70B achieves 93.47% accuracy on LifeChoice, surpassing GPT-4o by 23%. These results exhibit CoSER models’ strong capability in nuanced portrayal of characters’ personalities and behaviors.

## 5.3. Human Evaluation

We conduct human evaluation to further validate the performance of CoSER models and the reliability of LLM critics in GCA evaluation.

**Annotations** We randomly select 60 samples from CoSER Test, and evaluate 7 representative LLMs (shown in Table 4) based on their simulation results. Three annotators independently score the LLMs on 20 samples each, using a 1 to 10 scoring scale, given the background context and original conversations. During this process, the annotators note that 1) GCA simulation deeply reflects LLMs’ role-playing abilities, compared with previous single-turn benchmarks; 2) Manual evaluation is highly challenging and time-consuming. It requires careful learning of complex background and abundant dialogues, and takes 15 minutes to evaluate 7 models for one sample on average. We report models’ average scores and win rate against other models.

**Results and Analysis** According to Table 4, we find that: 1) The results generally align with GCA evaluation, validating the reliability of our LLM-based evaluation protocol; 2) CoSER 70B maintain superior performance, obtaining the highest score (6.783) and win rate (86.9%); 3) Human evaluators show less preference for GPT models compared to GPT-4o as LLM judges, which likely stems from self-preference bias in LLM judges (Wataoka et al., 2024).

<sup>5</sup><https://www.minimaxi.com/>

<sup>6</sup><https://team.doubao.com/en/Version:241215>

<sup>7</sup><https://www.stepfun.com/>

<sup>8</sup><https://openai.com/index/hello-gpt-4o/Version:240806>

<sup>9</sup><https://deepmind.google/technologies/gemini/pro/>

<sup>10</sup><https://www.anthropic.com/news/claude-3-5-sonnet>

<sup>11</sup><https://boson.ai/higgs-v2/>

Table 2. Performance (%) of various LLMs on given-circumstance acting using CoSER Test. **Bold** or underlined values indicate best performance across all models and open-source models, respectively.

Model	Based on LLM Judges					Based on N-gram	
	Storyline Consistency	Anthro-pomorphism	Character Fidelity	Storyline Quality	Average Score	BLEU	ROUGE-L
<i>Close-source Models</i>							
Abab7-preview	56.81±1.47	44.23±1.90	43.83±2.71	74.83±0.97	54.92±0.57	4.96±0.07	11.50±0.06
Doubao-pro	<b>60.95</b> ±1.40	49.72±0.23	47.02±1.10	79.28±0.82	59.24±0.30	6.38±0.08	12.95±0.04
Step-1-Flash	57.75±0.72	48.12±0.39	44.48±0.48	75.93±0.99	56.57±0.48	5.95±0.15	12.71±0.11
Step-2	61.43±0.88	49.06±1.69	47.33±0.70	77.96±0.85	58.94±0.75	5.75±0.08	12.50±0.11
GPT-3.5	57.22±0.13	43.30±0.48	42.29±1.47	73.91±0.64	54.18±0.63	4.58±0.11	11.80±0.10
GPT-4o	<b>61.59</b> ±0.66	48.93±0.48	<b>48.95</b> ±1.73	<b>80.33</b> ±0.59	<b>59.95</b> ±0.50	5.90±0.16	12.11±0.13
GPT-4o Mini	60.09±0.60	48.21±1.09	44.88±1.63	78.55±0.14	57.93±0.74	3.90±0.07	10.81±0.07
Gemini Pro	59.11±0.82	52.41±0.57	47.83±0.37	77.59±1.43	59.24±0.25	5.39±0.04	11.65±0.06
Claude-3-Haiku	58.18±0.72	44.66±1.72	41.88±0.34	74.14±1.26	54.71±0.84	4.80±0.05	12.02±0.02
Claude-3.5-Sonnet	57.45±0.98	48.50±2.35	45.69±1.80	77.23±0.88	57.22±0.95	5.17±0.12	11.45±0.07
<i>Open-source Models</i>							
Mistral-7B	<u>59.90</u> ±1.33	40.00±0.74	44.75±1.14	61.93±1.12	51.64±0.55	2.71±0.10	9.28±0.12
Qwen-2-7B	<u>51.96</u> ±0.67	35.48±0.62	31.51±2.95	63.18±0.79	45.53±0.69	4.21±0.21	10.71±0.10
LLaMA-3.1-8B	54.10±1.63	45.36±1.91	40.22±1.16	72.29±1.75	52.99±1.20	4.59±0.11	10.18±0.09
CoSER-8B	58.61±2.46	47.23±0.16	46.90±2.06	73.04±1.37	56.45±0.56	9.40±0.18	14.21±0.11
Vicuna-13B-1.5	52.75±1.64	39.12±1.21	38.04±0.98	60.43±1.58	47.58±1.25	1.67±0.10	5.59±0.18
Mixtral-8x7B	51.25±1.73	38.44±1.18	36.92±2.65	67.69±0.80	48.58±1.35	5.28±0.06	11.66±0.05
Qwen-2-72B	57.75±1.26	47.28±0.87	46.62±1.69	76.60±0.36	57.06±1.00	5.38±0.00	11.85±0.03
LLaMA-3.1-70B	57.46±1.65	45.95±1.30	43.72±1.17	74.84±0.54	55.49±0.33	4.82±0.06	10.98±0.06
Higgs-Llama-3-70B	57.10±1.12	43.82±2.18	42.41±1.66	75.62±0.15	54.74±1.26	3.99±0.33	10.92±0.56
CoSER-70B	58.66±1.34	<u>53.33</u> ±0.91	<u>48.75</u> ±1.43	75.49±0.94	<u>59.06</u> ±0.22	<u>10.10</u> ±0.04	<u>14.78</u> ±0.09
DeepSeek-V3	56.40±0.95	47.87±1.10	44.02±0.13	<u>76.66</u> ±1.26	56.24±0.46	4.54±0.14	11.02±0.15

Table 3. LLM performance (%) across three existing RPLA benchmarks. I.T. denotes inner thoughts. For InCharacter, we report accuracy for individual (Dim) and full (Full) dimensions on BFI.

Model	Incharacter		Life Choice	CroSS MR
	Dim	Full		
LLaMA-3.1-8B	64.97	15.62	61.10	30.15
CoSER-8B	75.80	21.88	69.54	44.94
<i>trained w/o I.T.</i>	70.70	15.62	59.92	43.14
LLaMA-3.1-70B	72.16	31.25	86.48	61.30
Higgs-Llama-3-70B	74.52	28.12	74.03	60.12
CoSER-70B	75.80	<b>34.38</b>	<b>93.47</b>	<b>64.49</b>
<i>trained w/o I.T.</i>	73.12	32.14	93.18	63.14
Qwen-2-72B	74.52	31.25	81.14	62.57
GPT-3.5	71.20	21.88	78.07	30.09
GPT-4o	<b>76.54</b>	32.62	75.96	<b>64.49</b>
Claude-3.5-Sonnet	72.61	21.88	86.07	30.59

**Alignment** We further study the alignment between LLM and human judges. We evaluate three judge models, GPT-4o, DeepSeek-V3 and DeepSeek-R1 (Guo et al., 2025). We consider standard GCA evaluation and ablation variants that remove: (i) original dialogues as reference, (ii) scoring rubrics, (iii) length correction, and (iv) dimension separation.

Table 4. Human evaluation on LLMs’ GCA simulation.

Model	Score	Win Rate (%)
GPT-3.5	3.117	10.6
LLaMA-3.1-8B	3.600	19.4
Abab7-preview	4.533	37.5
CoSER-8B	4.567	38.6
GPT-4o	4.967	47.2
Claude-3.5-Sonnet	6.200	73.9
CoSER-70B	<b>6.783</b>	<b>86.9</b>

We also compare with BLEU and ROUGE-L. The alignment is measured via model-wise preferences (§C.4).

The results in Table 5 reveal that: 1) All components of GCA evaluation contribute to improved alignment; 2) Reasoning models excel as judges, with DeepSeek-R1 achieving the highest alignment; 3) BLEU and ROUGE-L remain effective indicators for role-playing evaluation.

#### 5.4. Ablation Studies

**Inner Thoughts and Motivations Enhance RPLAs at Test Time** Table 6 compares LLMs’ overall scores on CoSER

Table 5. Alignment (%) between human judges and LLM judges (or automatic metrics).

Scoring Method	Alignment
GCA (GPT-4o)	68.6
w/o reference	64.3
w/o rubrics	65.1
w/o length correction	64.5
w/o dimension separation	65.2
GCA (DeepSeek-V3)	65.1
GCA (DeepSeek-R1)	<b>77.5</b>
w/o reference	77.2
BLEU	75.3
ROUGE-L	72.0

Table 6. Ablation study results (average scores) on CoSER Test. I.T. and Mot. mean inner thoughts and motivations, respectively.

Model	Standard	Test w/o I.T.	Test w/o Mot.
GPT-4o	59.95	56.89	56.34
Qwen-2-72B	57.06	51.95	54.21
LLaMA-3.1-70B	55.49	53.12	52.49
CoSER-70B	59.06	57.32	57.71
trained w/o I.T.	56.04	55.34	-
LLaMA-3.1-8B	52.99	51.97	49.63
CoSER-8B	56.45	54.65	56.81
trained w/o I.T.	54.25	54.38	-

Test with or without inner thoughts and motivations. The results show consistent performance improvements across all models when inner thoughts and motivations are included.

**Inner Thoughts Benefit Role-Playing Training** We train CoSER model variants without inner thoughts, and evaluate them on various benchmarks. Results in Tables 3 and 6 show that models trained without inner thoughts consistently underperform regular CoSER models, exhibiting the value of inner thoughts for role-playing training.

**5.5. CoSER Dataset for Retrieval Augmentation**

We evaluate the value of our comprehensive data types for retrieval augmentation on CoSER Test. We explore three retrieval sources for a specific character: dialogues in related conversations (Conv.), as well as experiences (Expr.) and raw text in related plots. The retrieval system is based on FAISS (Douze et al., 2024) with BGE-M3 (Chen et al., 2024b) for text embeddings. As shown in Fig. 16, we observe: 1) Models consistently benefit from characters’ retrieved experiences and conversations, especially for CoSER 70B; 2) However, raw text retrieval barely enhances LLMs’ performance. Detailed experimental settings and results are provided in §D and Table 16.

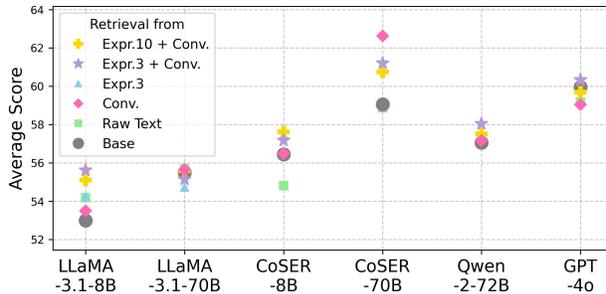


Figure 4. LLM Performance on CoSER Test with retrieval augmentation from various character data. Expr. and Conv. denote experiences and conversations respectively.

**5.6. Case Studies**

We conduct case studies to analyze LLMs’ performance in GCA simulation. Several cases are presented in Tables 17 to 22, from which we observe that: 1) CoSER models, trained on authentic dialogues, communicate more naturally, closely aligning with human speech patterns. 2) CoSER models better recall character-related knowledge, such as the iconic line “Grown enough to be wed, wed enough to be bedded” by Lysa Arryn in Table 17. This is consistent with their high BLEU and ROUGE-L scores. 3) CoSER models better portray sophisticated thinking process of humans. For example, in Tables 19 to 22, CoSER 70B faithfully reproduces Cersei Lannister’s suppressed anger as depicted in the original conversations, while other models, including GPT-4o and Claude-3.5-Sonnet, resort to a stereotypical portrayal of her arrogance and pride.

**6. Conclusion**

Towards effective RPLAs for established characters, this paper introduces CoSER, a collection of an authentic dataset, along with models and evaluation protocol based on such data. The CoSER dataset offers high-quality data from 771 renowned books, and includes comprehensive data types such as authentic dialogues, plot summaries, character experiences, inner thoughts, etc. Then, we propose given-circumstance acting (GCA) for training and evaluating role-playing LLMs, where LLMs sequentially portray multiple characters within authentic book scenarios. Applying GCA training to LLaMA-3.1 models using our dataset, we develop CoSER 8B and CoSER 70B, advanced open LLMs for role-playing. For evaluation, GCA combines multi-agent simulation and penalty-based LLM critics. Extensive experiments exhibit CoSER dataset’s value for RPLA training, evaluation, and retrieval. Moreover, CoSER models achieve state-of-the-art performance on both our evaluation and three existing RPLA benchmarks.

## Impact Statement

CoSER aims to advance RPLA research by providing effective dataset, models and evaluation protocol. We will release our dataset, models, and evaluation scripts to foster innovation in RPLAs. The dataset is intended for research purposes only. For copyright policies, we will not distribute raw novel content. We require that anyone using our work must adhere to copyright policies and obtain proper permissions for any derivative works. The CoSER dataset is derived from literary works may involve ethical considerations, and the content involved does not represent the authors' viewpoints. Our methods can potentially be applied to develop agents for real-world individuals. However, such applications must strictly respect personal data privacy and obtain necessary consent. We hope our research will benefit RPLA researchers and developers, but emphasize the importance of responsible development. Any applications must respect copyright policies, personal data privacy, and be developed with proper licensing.

## References

- Ahn, J., Lee, T., Lim, J., Kim, J.-H., Yun, S., Lee, H., and Kim, G. TimeChara: Evaluating point-in-time character hallucination of role-playing large language models. In Ku, L.-W., Martins, A., and Srikumar, V. (eds.), *Findings of the Association for Computational Linguistics ACL 2024*, pp. 3291–3325, Bangkok, Thailand and virtual meeting, 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.197. URL <https://aclanthology.org/2024.findings-acl.197>.
- Chan, X., Wang, X., Yu, D., Mi, H., and Yu, D. Scaling synthetic data creation with 1,000,000,000 personas. *ArXiv preprint*, abs/2406.20094, 2024. URL <https://arxiv.org/abs/2406.20094>.
- Chang, Y., Lo, K., Goyal, T., and Iyyer, M. Booookscore: A systematic exploration of book-length summarization in the era of llms. *ArXiv preprint*, abs/2310.00785, 2023. URL <https://arxiv.org/abs/2310.00785>.
- Chen, J., Wang, X., Xu, R., Yuan, S., Zhang, Y., Shi, W., Xie, J., Li, S., Yang, R., Zhu, T., Chen, A., Li, N., Chen, L., Hu, C., Wu, S., Ren, S., Fu, Z., and Xiao, Y. From persona to personalization: A survey on role-playing language agents. *Transactions on Machine Learning Research*, 2024a. ISSN 2835-8856. URL <https://openreview.net/forum?id=xrO70E8UIZ>. Survey Certification.
- Chen, J., Xiao, S., Zhang, P., Luo, K., Lian, D., and Liu, Z. Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. *ArXiv preprint*, abs/2402.03216, 2024b. URL <https://arxiv.org/abs/2402.03216>.
- Chen, J., Zhu, X., Yang, C., Shi, C., Xi, Y., Zhang, Y., Wang, J., Pu, J., Feng, T., Yang, Y., and Zhang, R. HoLLMwood: Unleashing the creativity of large language models in screenwriting via role playing. In Al-Onaizan, Y., Bansal, M., and Chen, Y.-N. (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 8075–8121, Miami, Florida, USA, November 2024c. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-emnlp.474. URL <https://aclanthology.org/2024.findings-emnlp.474/>.
- Chen, N., Wang, Y., Jiang, H., Cai, D., Li, Y., Chen, Z., Wang, L., and Li, J. Large language models meet harry potter: A dataset for aligning dialogue agents with characters. In Bouamor, H., Pino, J., and Bali, K. (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 8506–8520, Singapore, 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.570. URL <https://aclanthology.org/2023.findings-emnlp.570>.
- Dai, Y., Hu, H., Wang, L., Jin, S., Chen, X., and Lu, Z. Mmrole: A comprehensive framework for developing and evaluating multimodal role-playing agents, 2024. URL <https://arxiv.org/abs/2408.04203>.
- Deshpande, A., Murahari, V., Rajpurohit, T., Kalyan, A., and Narasimhan, K. Toxicity in chatgpt: Analyzing persona-assigned language models. In Bouamor, H., Pino, J., and Bali, K. (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 1236–1270, Singapore, 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.88. URL <https://aclanthology.org/2023.findings-emnlp.88>.
- Douze, M., Guzhva, A., Deng, C., Johnson, J., Szilvasy, G., Mazaré, P.-E., Lomeli, M., Hosseini, L., and Jégou, H. The faiss library. *ArXiv preprint*, abs/2401.08281, 2024. URL <https://arxiv.org/abs/2401.08281>.
- Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Yang, A., Fan, A., et al. The llama 3 herd of models. *ArXiv preprint*, abs/2407.21783, 2024. URL <https://arxiv.org/abs/2407.21783>.
- Dubois, Y., Galambosi, B., Liang, P., and Hashimoto, T. B. Length-controlled alpacaEval: A simple way to debias automatic evaluators. *ArXiv preprint*, abs/2404.04475, 2024. URL <https://arxiv.org/abs/2404.04475>.

- Guo, D., Yang, D., Zhang, H., Song, J., Zhang, R., Xu, R., Zhu, Q., Ma, S., Wang, P., Bi, X., et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- Jiang, A. Q., Sablayrolles, A., Roux, A., Mensch, A., Savary, B., Bamford, C., Chaplot, D. S., Casas, D. d. l., Hanna, E. B., Bressand, F., et al. Mixtral of experts. *ArXiv preprint*, abs/2401.04088, 2024. URL <https://arxiv.org/abs/2401.04088>.
- Kosinski, M. Theory of mind might have spontaneously emerged in large language models. *ArXiv preprint*, abs/2302.02083, 2023. URL <https://arxiv.org/abs/2302.02083>.
- Lambert, N., Morrison, J., Pyatkin, V., Huang, S., Ivison, H., Brahman, F., Miranda, L. J. V., Liu, A., Dziri, N., Lyu, S., et al. Tulu 3: Pushing frontiers in open language model post-training. *ArXiv preprint*, abs/2411.15124, 2024. URL <https://arxiv.org/abs/2411.15124>.
- Li, C., Leng, Z., Yan, C., Shen, J., Wang, H., MI, W., Fei, Y., Feng, X., Yan, S., Wang, H., et al. Chatharuhi: Reviving anime character in reality via large language model. *ArXiv preprint*, abs/2308.09597, 2023. URL <https://arxiv.org/abs/2308.09597>.
- Li, D., Jiang, B., Huang, L., Beigi, A., Zhao, C., Tan, Z., Bhattacharjee, A., Jiang, Y., Chen, C., Wu, T., et al. From generation to judgment: Opportunities and challenges of llm-as-a-judge. *ArXiv preprint*, abs/2411.16594, 2024. URL <https://arxiv.org/abs/2411.16594>.
- Liang, Y., Zhu, L., and Yang, Y. Tachikuma: Understanding complex interactions with multi-character and novel objects by large language models. *ArXiv preprint*, abs/2307.12573, 2023. URL <https://arxiv.org/abs/2307.12573>.
- Lin, C.-Y. and Och, F. J. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL-04)*, pp. 605–612, Barcelona, Spain, 2004. doi: 10.3115/1218955.1219032. URL <https://aclanthology.org/P04-1077>.
- Liu, A., Diab, M., and Fried, D. Evaluating large language model biases in persona-steered generation. In Ku, L.-W., Martins, A., and Srikumar, V. (eds.), *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 9832–9850, Bangkok, Thailand, August 2024a. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.586. URL <https://aclanthology.org/2024.findings-acl.586/>.
- Liu, A., Feng, B., Xue, B., Wang, B., Wu, B., Lu, C., Zhao, C., Deng, C., Zhang, C., Ruan, C., et al. Deepseek-v3 technical report. *ArXiv preprint*, abs/2412.19437, 2024b. URL <https://arxiv.org/abs/2412.19437>.
- Lu, K., Yu, B., Zhou, C., and Zhou, J. Large language models are superpositions of all characters: Attaining arbitrary role-play via self-alignment. In Ku, L.-W., Martins, A., and Srikumar, V. (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 7828–7840, Bangkok, Thailand, 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.423. URL <https://aclanthology.org/2024.acl-long.423>.
- OpenAI. Chatgpt, 2022. URL <https://openai.com/blog/chatgpt>.
- OpenAI. GPT-4 technical report. *ArXiv preprint*, abs/2303.08774, 2023. URL <https://arxiv.org/abs/2303.08774>.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. Bleu: a method for automatic evaluation of machine translation. In Isabelle, P., Charniak, E., and Lin, D. (eds.), *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pp. 311–318, Philadelphia, Pennsylvania, USA, 2002. Association for Computational Linguistics. doi: 10.3115/1073083.1073135. URL <https://aclanthology.org/P02-1040>.
- Papoudakis, A., Lapata, M., and Keller, F. Bookworm: A dataset for character description and analysis. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 4471–4500, 2024.
- Park, J. S., O’Brien, J. C., Cai, C. J., Morris, M. R., Liang, P., and Bernstein, M. S. Generative agents: Interactive simulacra of human behavior. In *In the 36th Annual ACM Symposium on User Interface Software and Technology (UIST ’23)*, UIST ’23, New York, NY, USA, 2023. Association for Computing Machinery.
- Shanahan, M., McDonell, K., and Reynolds, L. Role play with large language models. *Nature*, 623(7987):493–498, 2023.
- Shao, Y., Li, L., Dai, J., and Qiu, X. Character-LLM: A trainable agent for role-playing. In Bouamor, H., Pino, J., and Bali, K. (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 13153–13187, Singapore, 2023. Association for Computational Linguistics. URL <https://aclanthology.org/2023.emnlp-main.814>.
- Shen, T., Li, S., and Xiong, D. Roleeval: A bilingual role evaluation benchmark for large language models.

- ArXiv preprint*, abs/2312.16132, 2023. URL <https://arxiv.org/abs/2312.16132>.
- Stanislavski, K. *An Actor's Work: A Student's Diary*. Routledge, Abingdon, Oxon, 2008.
- Tu, Q., Fan, S., Tian, Z., and Yan, R. CharacterEval: A Chinese benchmark for role-playing conversational agent evaluation. *ArXiv preprint*, abs/2401.01275, 2024. URL <https://arxiv.org/abs/2401.01275>.
- Wang, N., Peng, Z., Que, H., Liu, J., Zhou, W., Wu, Y., Guo, H., Gan, R., Ni, Z., Yang, J., Zhang, M., Zhang, Z., Ouyang, W., Xu, K., Huang, W., Fu, J., and Peng, J. RoleLLM: Benchmarking, eliciting, and enhancing role-playing abilities of large language models. In Ku, L.-W., Martins, A., and Srikumar, V. (eds.), *Findings of the Association for Computational Linguistics ACL 2024*, pp. 14743–14777, Bangkok, Thailand and virtual meeting, 2024a. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.878. URL <https://aclanthology.org/2024.findings-acl.878>.
- Wang, X., Xiao, Y., Huang, J.-t., Yuan, S., Xu, R., Guo, H., Tu, Q., Fei, Y., Leng, Z., Wang, W., Chen, J., Li, C., and Xiao, Y. InCharacter: Evaluating personality fidelity in role-playing agents through psychological interviews. In Ku, L.-W., Martins, A., and Srikumar, V. (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1840–1873, Bangkok, Thailand, 2024b. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.102. URL <https://aclanthology.org/2024.acl-long.102>.
- Wataoka, K., Takahashi, T., and Ri, R. Self-preference bias in llm-as-a-judge. *arXiv preprint arXiv:2410.21819*, 2024.
- Weng, L. Llm-powered autonomous agents. *lilianweng.github.io*, 2023. URL <https://lilianweng.github.io/posts/2023-06-23-agent/>.
- Wu, J., Ouyang, L., Ziegler, D. M., Stiennon, N., Lowe, R., Leike, J., and Christiano, P. Recursively summarizing books with human feedback. *ArXiv preprint*, abs/2109.10862, 2021. URL <https://arxiv.org/abs/2109.10862>.
- Xu, R., Wang, X., Chen, J., Yuan, S., Yuan, X., Liang, J., Chen, Z., Dong, X., and Xiao, Y. Character is destiny: Can large language models simulate persona-driven decisions in role-playing? *ArXiv preprint*, abs/2404.12138, 2024. URL <https://arxiv.org/abs/2404.12138>.
- Yang, A., Yang, B., Hui, B., Zheng, B., Yu, B., Zhou, C., Li, C., Li, C., Liu, D., Huang, F., Dong, G., Wei, H., et al. Qwen2 technical report, 2024. URL <https://arxiv.org/abs/2407.10671>.
- Yuan, X., Yuan, S., Cui, Y., Lin, T., Wang, X., Xu, R., Chen, J., and Yang, D. Evaluating character understanding of large language models via character profiling from fictional works. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 2024.
- Zheng, L., Chiang, W., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E. P., Zhang, H., Gonzalez, J. E., and Stoica, I. Judging llm-as-a-judge with mt-bench and chatbot arena. In Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S. (eds.), *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. URL <https://openreview.net/forum?id=ucchPGDlao>.
- Zhou, J., Chen, Z., Wan, D., Wen, B., Song, Y., Yu, J., Huang, Y., Peng, L., Yang, J., Xiao, X., et al. CharacterGlm: Customizing Chinese conversational AI characters with large language models. *ArXiv preprint*, abs/2311.16832, 2023. URL <https://arxiv.org/abs/2311.16832>.
- Zhou, X., Zhu, H., Mathur, L., Zhang, R., Yu, H., Qi, Z., Morency, L.-P., Bisk, Y., Fried, D., Neubig, G., and Sap, M. SOTOPIA: Interactive evaluation for social intelligence in language agents. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=mM7VurbA4r>.

## A. Dataset

### A.1. Statistics and Analysis

As shown in Table 7, CoSER dataset is extensive and comprehensive, encompassing dialogue data from 771 books and 17,966 distinct characters. The dataset includes 30,069 unique plots and 29,798 conversations. On average, each conversation consists of approximately 13.2 utterances, with the entire dataset comprising a total of 392,298 utterances.

Table 7. Statistics of CoSER Dataset.

#Book	#Plot	#Conversation.	#Character	#Utterance
771	30,069	29,798	17,966	392,298

Our book selection is derived from the *Best Books Ever* list on *Goodreads*, a curated collection of globally acclaimed literary works. These novels have garnered widespread recognition and appreciation from readers worldwide. Table 8 presents a comprehensive list of the top 100 books from our selection.

We analyze the genres of the selected books based on Supersummary classifications, with the statistical results presented in Figure 5. Our dataset encompasses a wide range of genres, particularly fiction categories such as Fantasy, Historical, Science Fiction, Romance, and Mystery. It also features niche fiction genres, showcasing diverse narrative styles. In addition to fiction, the collection includes non-fiction genres such as memoirs, biographies, and other works, enhancing its versatility.

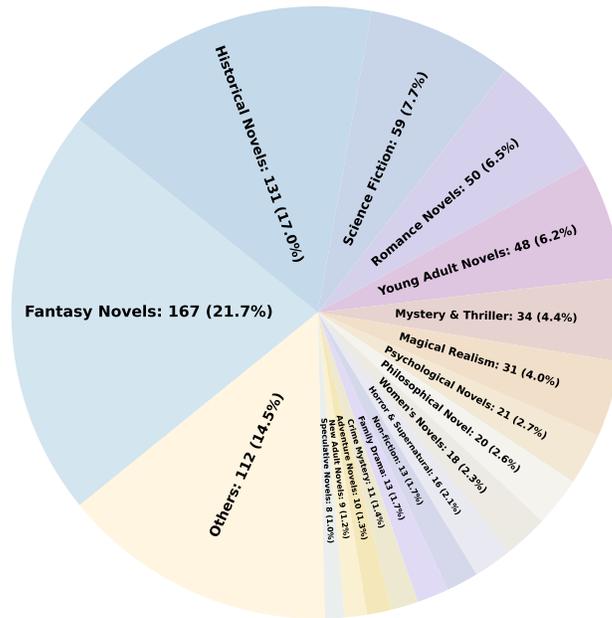


Figure 5. Genre distribution of selected books in CoSER dataset.

### A.2. Data Splits

For evaluation purposes, we held out the last 10% of data from each book; that is, they are not included in our prompts or datasets for training or retrieval purposes. Additionally, we trained the CoSER models on only 90% of the books.

### A.3. Implementation Details for Construction

**Extracting Raw Text** LLMs often struggle to extract verbatim original content, especially with punctuation marks like quotation marks, making it difficult to extract raw text directly. Therefore, instead of asking LLMs to extract the complete text of a plot, we prompt LLMs to extract the first and last sentences of each plot. Since the extracted sentences may still contain typos, we apply lexical similarity to match them with the exact sentences from the raw text. Finally, we identify the complete raw text based on these first and last sentences.

Table 8. The top 100 selected books from Goodreads' *Best Books Ever* list.

Selected Books	
1. <i>The Hunger Games (The Hunger Games, #1)</i>	2. <i>Harry Potter and the Order of the Phoenix (Harry Potter, #5)</i>
3. <i>Pride and Prejudice</i>	4. <i>To Kill a Mockingbird</i>
5. <i>The Book Thief</i>	6. <i>Animal Farm</i>
7. <i>The Chronicles of Narnia (The Chronicles of Narnia, #1-7)</i>	8. <i>The Fault in Our Stars</i>
9. <i>The Picture of Dorian Gray</i>	10. <i>Wuthering Heights</i>
11. <i>Gone with the Wind</i>	12. <i>The Perks of Being a Wallflower</i>
13. <i>The Lightning Thief (Percy Jackson and the Olympians, #1)</i>	14. <i>The Little Prince</i>
15. <i>The Great Gatsby</i>	16. <i>Crime and Punishment</i>
17. <i>Memoirs of a Geisha</i>	18. <i>Les Misérables</i>
19. <i>The Alchemist</i>	20. <i>Lord of the Flies</i>
21. <i>The Hitchhiker's Guide to the Galaxy (#1)</i>	22. <i>The Help</i>
23. <i>Dracula</i>	24. <i>Ender's Game (Ender's Saga, #1)</i>
25. <i>Of Mice and Men</i>	26. <i>One Hundred Years of Solitude</i>
27. <i>Brave New World</i>	28. <i>A Thousand Splendid Suns</i>
29. <i>The Time Traveler's Wife</i>	30. <i>The Princess Bride</i>
31. <i>The Secret Garden</i>	32. <i>The Outsiders</i>
33. <i>A Game of Thrones (A Song of Ice and Fire, #1)</i>	34. <i>Little Women</i>
35. <i>A Wrinkle in Time (Time Quintet, #1)</i>	36. <i>The Odyssey</i>
37. <i>Harry Potter and the Deathly Hallows (Harry Potter, #7)</i>	38. <i>Frankenstein: The 1818 Text</i>
39. <i>The Kite Runner</i>	40. <i>The Handmaid's Tale (The Handmaid's Tale, #1)</i>
41. <i>The Lovely Bones</i>	42. <i>The Adventures of Huckleberry Finn</i>
43. <i>Life of Pi</i>	44. <i>A Tale of Two Cities</i>
45. <i>Dune (Dune, #1)</i>	46. <i>Harry Potter and the Prisoner of Azkaban (Harry Potter, #3)</i>
47. <i>Water for Elephants</i>	48. <i>Harry Potter and the Sorcerer's Stone (Harry Potter, #1)</i>
49. <i>The Bell Jar</i>	50. <i>Matilda</i>
51. <i>The Stand</i>	52. <i>Catch-22</i>
53. <i>The Adventures of Sherlock Holmes (Sherlock Holmes, #3)</i>	54. <i>The Pillars of the Earth (Kingsbridge, #1)</i>
55. <i>Rebecca</i>	56. <i>Great Expectations</i>
57. <i>The Girl with the Dragon Tattoo (Millennium, #1)</i>	58. <i>The Color Purple</i>
59. <i>Anna Karenina</i>	60. <i>My Sister's Keeper</i>
61. <i>The Brothers Karamazov</i>	62. <i>A Clockwork Orange</i>
63. <i>And Then There Were None</i>	64. <i>The Road</i>
65. <i>To Kill a Mockingbird</i>	66. <i>The Golden Compass (His Dark Materials, #1)</i>
67. <i>Vampire Academy (Vampire Academy, #1)</i>	68. <i>Siddhartha</i>
69. <i>The Complete Stories and Poems</i>	70. <i>Interview with the Vampire (The Vampire Chronicles, #1)</i>
71. <i>Don Quixote</i>	72. <i>The Old Man and the Sea</i>
73. <i>The Poisonwood Bible</i>	74. <i>Harry Potter and the Goblet of Fire (Harry Potter, #4)</i>
75. <i>Atlas Shrugged</i>	76. <i>The Notebook (The Notebook, #1)</i>
77. <i>Harry Potter and the Half-Blood Prince (Harry Potter, #6)</i>	78. <i>Moby-Dick or, The Whale</i>
79. <i>A Prayer for Owen Meany</i>	80. <i>Clockwork Angel (The Infernal Devices, #1)</i>
81. <i>The Stranger</i>	82. <i>The Secret Life of Bees</i>
83. <i>Harry Potter and the Chamber of Secrets (Harry Potter, #2)</i>	84. <i>The Red Tent</i>
85. <i>The Name of the Wind (The Kingkiller Chronicle, #1)</i>	86. <i>The Master and Margarita</i>
87. <i>The Metamorphosis</i>	88. <i>Eragon (The Inheritance Cycle, #1)</i>
89. <i>The Count of Monte Cristo</i>	90. <i>Looking for Alaska</i>
91. <i>The Adventures of Tom Sawyer</i>	92. <i>Charlie and the Chocolate Factory (Charlie Bucket, #1)</i>
93. <i>The Last Olympian (Percy Jackson and the Olympians, #5)</i>	94. <i>The Curious Incident of the Dog in the Night-Time</i>
95. <i>The Shadow of the Wind (Cemetery of Forgotten Books, #1)</i>	96. <i>The Unbearable Lightness of Being</i>
97. <i>On the Road</i>	98. <i>The Name of the Rose</i>
99. <i>A Story of Yesterday</i>	100. <i>The Godfather (The Godfather, #1)</i>

**Parsing Structured Data** During extraction, we instruct LLMs to output extracted data in JSON format. However, LLM-generated JSON strings may sometimes be unparseable or do not conform to the specified format (e.g., missing required keys).

Towards this challenge, we adopt a repair-and-retry strategy to improve extraction success rate. For each chunk to be extracted, we invoke LLMs and attempt to parse a valid JSON object. If parsing fails, we employ LLMs to repair the invalid JSON string and retry. Some failures occur because the LLM attempts to output JSON that exceeds the maximum token limit, resulting in truncation. In such cases, we prompt the repairing LLM to truncate the JSON at an appropriate point. If it still fail, we restart the entire process from the beginning, making up to 5 attempts.

**Refining Conversation Settings** During data extraction, we observe that the initially extracted conversation settings, including scenarios and character motivations, often fail to provide a comprehensive context. We attribute this to the LLM’s tendency to distribute information across different data fields when extracting multiple kinds of information simultaneously, rather than repeating it in different data fields. For instance, if certain information is already mentioned in the plot summary, it might be omitted from the scenario description.

Therefore, to provide a complete context for given-circumstance acting, we implement an additional LLM call to refine the conversation settings based on the extracted data. We instructed the refining LLM to provide a comprehensive conversation setting, while carefully avoiding any disclosure of subsequent dialogue content or plot developments.

For additional details, such as the regular expressions used for identifying chapter titles, please refer to our code.

#### A.4. Comparison with Existing Methods for Character Profiling

Previous character profiling methods, including hierarchical updating (Wu et al., 2021), incremental updating (Chang et al., 2023), and one-shot summarization (Yuan et al., 2024), typically only generate the profile of a single character at a time. Moreover, Papoudakis et al. (2024) shows that these methods, particularly hierarchical updating, perform suboptimally when generating multiple character profiles simultaneously.

CoSER’s multi-stage, extract-then-aggregate pipeline addresses these limitations. It ensures comprehensive character profiles with high precision and recall of character knowledge, capturing evolving character arcs, and significantly improving processing efficiency.

#### A.5. Safety

We conduct safety checks on the dataset and remove potentially problematic content. Specifically, we truncate 110 sensitive conversations and remove a total of 602 messages.

#### A.6. Comparison with PersonaHub

Recently, several studies have explored LLM role-playing using synthesized personas rather, such as PersonaHub (Chan et al., 2024). We provide a systematic comparison between CoSER and PersonaHub below.

**Goal and Focus** The fundamental objectives of the two datasets differ significantly:

Aspect	CoSER	PersonaHub
Target	Simulate established personas with high fidelity	Synthesize instruction data for knowledge distillation
Persona Focus	Depth and richness of character data	Breadth of persona types
Key Capabilities	Anthropomorphism, character fidelity, multi-character interaction	General instruction following

Table 9. Comparison between CoSER and PersonaHub regarding goal and focus.

**Dataset Quality and Design** The data sources and design reflect distinct research priorities:

This comparison highlights that while PersonaHub excels in generating diverse synthetic personas for broad instruction-

Aspect	CoSER	PersonaHub
Data Source	Authentic book dialogues	GPT-4o synthesized conversations
Persona Origin	Book characters	LLM-synthesized personas
Persona Data	Comprehensive profiles, dialogues, experiences	Concise profiles with limited context
Design Priority	Quality, authenticity, and character depth	Quantity and diversity

Table 10. Comparison between CoSER and PersonaHub regarding dataset quality and design.

following training, CoSER focuses on preserving the authenticity and depth of established characters for high-fidelity role-playing applications.

## B. Training

### B.1. Training Samples

We transform conversations from the CoSER dataset into training samples in the Sharegpt format. We utilize 90% of the books in the dataset for training, while the remaining 10% are set aside to evaluate our models’ ability to generalize to out-of-domain characters and books. We construct one training sample for each character in every CoSER conversation, encompassing both main characters and minor roles. When training on a character  $c$ , we designate  $c$ ’s messages as targets for optimization, while using the system prompt and messages from other characters as inputs. Adjacent inputs are concatenated.

Towards general role-playing capabilities across diverse scenarios and applications, we dynamically generate role-playing instructions (system prompts) using varied phrasings, formats, and data types, as shown in Tables 27 and 28. We consider instructions entirely in natural language, as well as those formatted with special symbols (such as ###, ===), randomly sampling different formats and various expressions for the same semantics. We consider various configurations of the available data. Each sample may include (50%) or exclude (50%) the following elements : 1) Profiles of other characters in this conversation; 2) Summaries of the relevant plot; 3) Inner thoughts within the messages.

Besides character role-playing, we train CoSER models for environment modeling and next speaker prediction (NSP) for multi-agent simulation. For environment modeling, we train LLMs  $\pi_e$  to play the environment role  $e$  in the same approach, leveraging environment messages in our dataset. For NSP, given setting  $S$  and messages  $\{m_1, \dots, m_i\}$ , we train LLMs to predict the speaker of  $m_{i+1}$  (or ending the conversation).

Our role-playing dataset comprises approximately 0.1B tokens, as measured using the LLaMA 3.1 tokenizer. To maintain general intelligence and instruction-following capabilities, we augment this with an equivalent volume (0.1B tokens) of general-purpose supervised fine-tuning data from Tulu 3 (Lambert et al., 2024)<sup>12</sup>. This balanced mixture ensures that the model retains broad language understanding while developing specialized role-playing abilities. If more role-playing data are expected, our data curation pipeline can be easily applied to additional books or other fictional works, thereby acquiring data on a much larger scale.

For more details, please refer to our code.

### B.2. Hyperparameters

We fine-tune the LLaMA 3.1 models using the following hyperparameters: a learning rate of  $1 \times 10^{-5}$ , a sequence length of 16,384, training for 8 epochs, and a global batch size of 48.

## C. Experiment Settings

### C.1. Test Set Sampling

The CoSER Test set contains 200 samples: 100 from books used in CoSER training (in-domain) and 100 from books not used in training (out-of-domain). We employ a weighted sampling strategy to prioritize well-established characters with

<sup>12</sup><https://huggingface.co/datasets/allenai/tulu-3-sft-mixture/tree/main/data>

more persona data. The sampling process consists of the following steps: First, for each book, we assign character weights as the square root of the number of plots in which they are involved. Second, we calculate the weight of a conversation as the average of its characters’ weights, including both main characters and side roles. Finally, separately for the in-domain and out-of-domain settings, we rank all conversations by weight and perform weighted sampling from the top half of conversations with higher weights.

### C.2. Prompting Strategies for Existing RPLA Benchmarks

For existing RPLA benchmarks, including InCharacter (Wang et al., 2024b), LifeChoice (Xu et al., 2024), and Cross-MR (Yuan et al., 2024), we adapt or refine their prompting strategies as follows:

1. For InCharacter, we add “*You’re consulting with a personality assessment expert who will ask you some questions. Please provide honest and detailed responses to help with the analysis. Please think carefully and state your reasons when answering the questions.*” after the character profile. This adaptation aims to ensure that RPLAs honestly express their true thoughts. After being trained on authentic character dialogues, the CoSER models, unlike general LLMs, tend to produce brief, conversational-style answers that may be too short or may decline to answer questions, thus failing to provide sufficient information for personality assessment.
2. For LifeChoice and Cross-MR, we reverse the order of their reasoning and answering processes. Specifically, we have them think before providing their choices, thus enabling RPLAs to make well-considered decisions.

### C.3. Length Correction

In our evaluation, we use a penalty-based scoring mechanism that counts the flaws in RPLAs’ performance. However, since longer simulations naturally accumulate more flaws, we need to implement length correction to reduce length bias in LLM judges, following previous work (Li et al., 2024).

We analyze the phenomenon of length bias in penalty-based scoring. The initial score is defined as  $s = 100 - \sum_{f \in \mathcal{F}} v_f$ , where  $\mathcal{F}$  represents the set of flaws and  $v_f$  is their severity ranging from 1 to 5. Our analysis is conducted on the initial scores from simulations of three models on the CoSER Test set: CoSER 70B, LLaMA-3.1 70B, and GPT-4o, with or without retrieval augmentation (three Experience and one Conversation), totaling 1,200 cases. We examine the relationship between the number of rounds and the vanilla scores in these 1,200 cases. As shown in Figure 6, we plot the data points for these cases and perform linear regression to fit these points. The fitted linear function is  $\text{score} = -1.5909 \times \text{rounds} + 59.0617$ , which means that for each additional round in the simulation, the score decreases by approximately 1.6 points.

To mitigate this bias, we implement length correction by compensating for the points deducted due to increased rounds. Specifically, we compute the length-corrected score as  $s = 100 - \sum_{f \in \mathcal{F}} v_f + \lambda |\bar{M}|$ , where  $\lambda$  is set to 1.5 based on the analysis above.

### C.4. Human Evaluation

The alignment metrics are measured via the frequency of agreement between LLM and human judges when comparing model pairs, excluding cases where judges assign similar scores to both models (*i.e.*, score differences  $\leq 1$  for human judges or  $\leq 5\%$  for automatic metrics).

## D. Additional Results

**Evaluation with Different LLMs as Judges** We conduct additional experiments using alternative judge models for GCA evaluation. We evaluate 7 representative models as shown in Table 11, using three different judge models: GPT-4o, DeepSeek-R1, and DeepSeek-V3. This comparison aims to examine whether our conclusions remain consistent across different LLM judges.

The results in Table 11 reveal several important findings: 1) Despite these variations in absolute scores, the relative ranking patterns remain largely consistent, with CoSER models maintaining competitive or superior performance across all settings. 2) Different judges exhibit distinct preferences, confirming potential bias in LLM-based evaluation. While GPT-4o shows preference for GPT models, DeepSeek-R1 and DeepSeek-V3 favor CoSER-70B and Claude-3.5-Sonnet over GPT-4o.

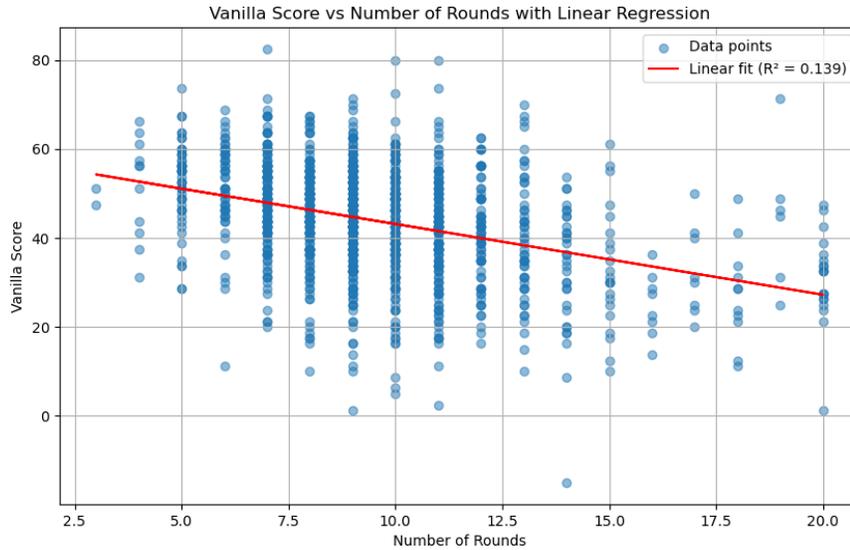


Figure 6. Linear regression results showing length bias of penalty-based LLM critics in GCA evaluation.

Table 11. Scores (%) assigned by different LLM judges. Results demonstrate consistent superiority of CoSER models, while revealing judge-specific preferences.

Model	GPT-4o	DeepSeek-R1	DeepSeek-V3
GPT-3.5	52.8	35.9	40.5
LLaMA-3.1-8B	51.8	37.2	36.8
Abab7-preview	53.7	41.5	40.4
CoSER-8B	56.1	44.5	45.9
GPT-4o	<b>58.5</b>	48.4	46.1
Claude-3.5-Sonnet	56.2	<b>54.8</b>	40.7
CoSER-70B	57.4	50.8	<b>47.7</b>

**Win Rates against GPT-4o and GPT-3.5** To further evaluate existing LLMs in given-circumstance acting, we present their win rates against GPT-3.5 (OpenAI, 2022) and GPT-4o (OpenAI, 2023) in Table 12, in addition to the results shown in Table 2.

**Generalization of CoSER Models to New Characters** Table 13 separately presents the performance of LLMs on test splits from books included in and excluded from CoSER training. The results demonstrate consistent trends across both splits, confirming that CoSER models maintain strong performance even on out-of-domain characters.

**Conversation Continuation** Table 14 shows experiment results when multi-agent systems continue conversations from the first  $k = 3$  original messages.

Tables 14 presents detailed evaluation results where our multi-agent simulations start from the first  $k = 3$  original messages.

**CoSER Dataset for Retrieval Augmentation** We validate the effectiveness of CoSER’s comprehensive data types for retrieval augmentation on the CoSER Test set. We explore three retrieval sources related to specific characters: dialogues from conversations, experiences from plots, and raw text from plots. We compare several combinations of these sources, including: 1) None (Base) 2) Raw text of one plot (Raw Text) 3) One conversation (Conv.) 4) Character experiences from three plots (Expr.3) 5) Expr.3 combined with Conv. 6) Expr.10 combined with Conv. The complete results are presented in Table 16.

Table 12. Win rates (%) of various LLMs on given-circumstance acting using CoSER Test, compared against GPT-3.5 and GPT-4. **Bold** and underlined numbers indicate the best results among all and open-source models, respectively.

Model	Win Rate	
	vs. GPT-3.5	vs. GPT-4o
<i>Close-source Models</i>		
Abab7-preview	52.17±4.54	35.33±2.88
Doubao-pro	64.67±2.32	48.00±3.97
Step-1-Flash	58.75±6.77	39.75±1.64
Step-2	65.08±7.26	46.92±4.40
GPT-3.5	50.00±0.00	33.08±3.64
GPT-4o	<b>66.92</b> ±3.64	<b>50.00</b> ±0.00
GPT-4o Mini	59.58±7.80	42.92±0.63
Gemini Pro	66.17±3.33	49.75±0.25
Claude-3-Haiku	51.42±4.40	33.25±3.03
Claude-3.5-Sonnet	58.17±8.00	41.42±1.77
<i>Open-source Models</i>		
Mistral-7B	41.50±7.28	26.67±1.81
Qwen-2-7B	24.83±6.64	13.58±3.00
LLaMA-3.1-8B	45.33±6.57	32.75±5.86
CoSER-8B	58.17±5.36	41.67±3.75
Vicuna-13B-1.5	30.50±10.44	18.67±2.18
Mixtral-8x7B	32.50±1.50	19.83±1.42
Qwen-2-72B	59.58±5.11	41.92±1.66
LLaMA-3.1-70B	53.83±2.57	37.92±3.26
Higgs-Llama-3-70B	51.83±4.89	34.92±4.84
CoSER-70B	<u>64.33</u> ±3.74	<u>49.42</u> ±4.11
DeepSeek-V3	<u>59.50</u> ±2.82	<u>39.50</u> ±2.05

Table 13. LLM performance (%) on given-circumstance acting using CoSER Test, separated into the in-domain and out-of-domain splits for CoSER training.

Model	In-Domain					Out-of-Domain				
	Storyline Consistency	Anthro-pomorphism	Character Fidelity	Storyline Quality	Average Score	Storyline Consistency	Anthro-pomorphism	Character Fidelity	Storyline Quality	Average Score
<i>Close-source Models</i>										
Abab7-preview	55.31±1.4	42.29±1.3	42.94±3.2	74.13±1.9	53.67±1.0	58.30±1.8	46.17±2.6	44.72±2.3	75.54±1.5	56.18±0.1
Doubao-pro	60.37±1.8	49.06±0.4	45.76±3.4	77.87±1.2	58.26±0.9	61.53±1.1	50.38±0.3	48.28±1.7	80.69±0.4	60.22±0.4
Step-1-Flash	57.10±0.3	48.31±1.0	41.84±1.9	76.18±1.5	55.86±1.0	58.39±1.2	47.94±0.5	47.12±1.0	75.67±0.5	57.28±0.2
Step-2	<b>60.55</b> ±0.7	48.82±2.2	<b>47.90</b> ±1.1	77.33±1.0	<b>58.65</b> ±1.0	<b>62.30</b> ±1.5	49.30±1.2	<b>46.77</b> ±1.7	78.59±1.2	<b>59.24</b> ±0.6
GPT-3.5	55.58±0.9	42.18±5.2	40.37±3.2	72.90±0.1	52.76±2.0	59.69±3.0	44.25±2.1	44.60±1.4	74.66±2.3	55.80±1.6
GPT-4o	59.88±1.4	48.11±1.0	47.10±0.2	<b>79.06</b> ±1.5	58.54±0.7	62.29±1.5	49.48±1.2	49.90±0.4	<b>80.18</b> ±0.9	60.46±0.1
GPT-4o Mini	59.15±1.3	46.94±1.5	43.99±2.5	77.33±0.6	56.85±0.1	61.02±1.6	49.48±3.1	45.77±1.1	79.77±0.5	59.01±1.4
Gemini Pro	57.72±0.4	50.94±1.8	46.23±1.0	76.22±1.6	57.78±0.9	60.50±1.9	53.88±1.1	49.43±0.3	78.97±1.3	60.69±0.8
Claude-3-Haiku	57.61±0.5	44.97±2.2	40.61±1.4	73.52±1.2	54.18±0.6	58.74±1.1	44.36±1.9	43.14±0.8	74.76±1.7	55.25±1.2
Claude-3.5-Sonnet	56.44±1.5	47.24±1.4	44.89±0.6	76.39±1.5	56.24±1.0	58.46±1.1	49.75±3.3	46.49±3.0	78.06±1.5	58.19±0.9
<i>Open-source Models</i>										
Mistral-7B	60.29±1.9	38.98±2.0	42.93±0.5	62.20±3.1	51.10±1.8	59.51±2.4	41.02±1.3	46.57±2.8	61.65±1.3	52.19±0.7
Qwen-2-7B	50.77±1.3	34.17±1.1	29.92±2.9	62.58±0.7	44.36±0.7	53.14±1.5	36.79±1.0	33.09±3.3	63.79±1.4	46.70±1.3
LLaMA-3.1-8B	53.00±1.2	43.64±0.8	39.05±1.5	71.34±1.3	51.76±0.6	55.20±2.1	47.08±3.2	41.38±1.0	73.23±2.5	54.22±1.8
CoSER-8B	58.56±3.5	46.78±1.1	45.78±3.1	73.38±1.8	56.12±0.5	58.66±1.6	47.69±0.8	48.03±1.3	72.71±1.2	56.77±0.6
Vicuna-13B-1.5	51.84±1.2	38.88±0.4	36.39±0.5	58.15±0.7	46.31±0.4	53.66±2.1	39.35±2.7	39.69±1.6	62.71±2.5	48.85±2.1
Mixtral-8x7B	51.18±3.1	38.76±1.8	34.37±2.8	66.44±0.1	47.69±1.7	51.32±0.4	38.13±0.5	39.48±2.6	68.93±1.7	49.47±1.1
Qwen-2-72B	56.34±2.4	46.19±0.4	45.59±2.2	75.68±0.3	55.95±1.0	59.15±1.1	48.37±1.3	47.65±1.8	77.52±0.4	58.17±1.0
LLaMA-3.1-70B	55.44±2.7	44.19±1.9	42.67±1.5	73.90±1.8	54.05±0.8	59.48±1.4	47.72±2.3	44.78±0.9	75.78±0.9	56.94±0.7
Higgs-Llama-3-70B	55.85±2.7	41.18±3.3	39.79±2.3	73.77±0.6	52.65±2.0	58.35±0.5	46.45±1.1	45.04±1.0	77.48±0.6	56.83±0.6
CoSER-70B	57.77±1.3	<b>51.60</b> ±1.3	45.82±0.9	74.27±1.1	<u>57.37</u> ±0.7	<u>59.56</u> ±2.1	<u>55.06</u> ±1.1	<u>51.67</u> ±2.5	76.71±1.6	<u>60.75</u> ±1.1
DeepSeek-V3	55.36±0.5	47.55±1.3	43.10±0.3	74.91±2.0	55.23±0.6	57.45±2.0	48.19±0.9	44.93±0.4	<u>78.41</u> ±1.1	57.24±0.5

Table 14. Performance (%) of various LLMs on CoSER Test in conversation continuation setting ( $k = 3$ ), where RPLAs start from the first three messages in the authentic conversations.

Model	Based on LLM Judges					Based on N-gram	
	Storyline Consistency	Anthro-pomorphism	Character Fidelity	Storyline Quality	Average Score	BLEU	ROUGE-L
<i>Close-source Models</i>							
Abab7-preview	65.25	55.26	55.95	79.68	64.03	10.53	15.99
Doubao-pro	<b>68.31</b>	<b>59.51</b>	59.23	80.15	66.80	11.83	17.13
Step-1-Flash	64.32	53.35	54.38	79.03	62.77	11.94	17.06
Step-2	66.61	55.81	59.59	80.56	65.64	11.72	17.27
GPT-3.5	65.72	54.34	56.48	77.67	63.55	10.80	16.39
GPT-4o	67.48	58.51	62.30	<b>82.23</b>	<b>67.63</b>	12.32	16.98
GPT-4o Mini	66.74	55.55	56.80	80.01	64.77	9.50	15.54
Gemini Pro	65.47	59.43	<b>62.42</b>	78.78	66.53	10.84	16.27
Claude-3-Haiku	64.51	54.01	57.13	77.26	63.23	10.11	16.18
Claude-3.5-Sonnet	64.54	54.57	58.76	79.89	64.44	8.64	14.94
<i>Open-source Models</i>							
Mistral-7B	67.50	50.39	59.90	68.67	61.62	7.41	14.33
Qwen-2-7B	59.81	49.87	49.04	69.16	56.97	10.59	16.22
LLaMA-3.1-8B	60.90	51.36	50.37	74.89	59.38	7.86	13.82
CoSER-8B	67.22	58.19	58.80	76.44	65.16	13.17	<b>18.42</b>
Vicuna-13B-1.5	60.92	46.68	50.51	67.80	56.48	5.11	9.71
Mixtral-8x7B	64.66	51.21	54.08	74.01	60.99	11.21	16.97
Qwen-2-72B	67.27	55.87	59.84	<u>80.07</u>	<u>65.76</u>	11.92	16.96
LLaMA-3.1-70B	64.08	54.80	54.18	<u>78.31</u>	<u>62.84</u>	8.74	14.74
Higgs-Llama-3-70B	65.09	54.80	58.20	79.36	64.36	10.86	16.40
CoSER-70B	65.99	<u>59.24</u>	<u>59.97</u>	76.74	65.48	<b>13.46</b>	18.18
DeepSeek-V3	62.95	56.95	58.25	79.80	64.49	9.25	15.01

**Ablation Studies** We examine the effectiveness of inner thoughts in both training and evaluation. The complete results are demonstrated in Table 15.

**Comparing GCA Simulation with AI-driven Storytelling Methods** To provide a more comprehensive evaluation against prior work in digital actors and AI-driven storytelling, we compare our GCA simulation approach with HollmWood (Chen et al., 2024c), a representative method for AI-generated storytelling. We conduct this comparison using 30 samples from CoSER Test, employing GPT-4o as both the actor and writer in both frameworks to ensure fair comparison. The generated stories and character interactions are evaluated using GCA evaluation. Results demonstrate that GCA simulation achieves an average score of 59.2%, significantly outperforming HollmWood’s 50.2%. These findings highlight that our given-circumstance acting approach produces more authentic and human-like character interactions compared to previous AI storytelling methods. The superior performance of GCA simulation can be attributed to its focus on character-driven dialogue generation within well-defined circumstances, rather than plot-driven narrative construction, leading to more faithful character portrayals and realistic conversational dynamics.

### E. Authentic v.s. LLM-Generated Role-playing Data

To demonstrate the superiority of authentic role-playing data over synthesized alternatives, we conduct two complementary experiments comparing authentic conversations from CoSER dataset with LLM-generated dialogues.

**GCA Evaluation of Dialogue Quality** We apply GCA evaluation to score authentic (groundtruth) and LLM-generated dialogues on 100 samples from CoSER Test without providing groundtruth to LLM judges as reference. We use DeepSeek-R1 as the judge model due to its superior alignment with human evaluators (as shown in §5.3). GPT-4o is used to generate synthetic dialogues via GCA simulation. Results show that authentic dialogues from CoSER achieve a score of 85.1%, significantly outperforming GPT-4o generated dialogues at 76.2%.

Table 15. Comprehensive results of LLM performance (%) on CoSER Test, with and without inner thoughts (I.T) during inference and CoSER model training.

Model	Based on LLM Judges					Based on N-gram	
	Storyline Consistency	Anthro-pomorphism	Character Fidelity	Storyline Quality	Average Score	BLEU	ROUGE-L
<i>Test w/ I.T.</i>							
GPT-4o	61.59	48.93	48.95	80.33	59.95	5.90	12.11
Qwen-2-72B	57.75	47.28	46.62	76.60	57.06	5.38	11.85
LLaMA-3.1-70B	57.46	45.95	43.72	74.84	55.49	4.82	10.98
CoSER-70B	58.66	53.33	48.75	75.49	59.06	10.10	14.78
<i>trained w/o I.T.</i>	55.67	49.93	45.67	72.89	56.04	9.18	14.18
LLaMA-3.1-8B	54.10	45.36	40.22	72.29	52.99	4.59	10.18
CoSER-8B	58.61	47.23	46.90	73.04	56.45	9.40	14.21
<i>trained w/o I.T.</i>	54.91	44.89	44.10	73.09	54.25	9.65	14.27
<i>Test w/o I.T.</i>							
GPT-4o	59.51	45.86	45.13	77.06	56.89	5.54	11.44
Qwen-2-72B	55.39	41.00	37.70	73.70	51.95	5.22	11.32
LLaMA-3.1-70B	56.01	43.17	41.23	72.08	53.12	4.15	10.37
CoSER-70B	59.23	48.76	47.06	74.23	57.32	10.02	14.72
<i>trained w/o I.T.</i>	55.52	46.22	45.34	74.31	55.34	9.66	14.62
LLaMA-3.1-8B	54.27	45.54	37.86	70.23	51.97	4.37	10.27
CoSER-8B	56.78	45.64	45.00	71.16	54.65	9.38	14.44
<i>trained w/o I.T.</i>	54.98	46.22	45.81	70.52	54.38	9.16	13.97

**Training Effectiveness Comparison** We further examine the training effectiveness of authentic versus synthetic dialogues. We fine-tune LLaMA-3.1-8B using authentic or LLM-generated dialogues for 200 samples from CoSER Test for 4 epochs, using the same settings as CoSER-8B. We evaluate the two trained models via GCA evaluation on 100 new samples. The model trained on authentic dialogues from CoSER achieves 55.3% score, while the model trained on GPT-4o generated dialogues achieves 54.7%.

These results demonstrate that authentic dialogues consistently outperform LLM-generated alternatives in both intrinsic quality and training effectiveness. These findings highlight the importance of authentic, high-quality role-playing data for developing effective RPLAs.

## F. Examples and Case Study

We present several examples of our extracted conversations, as well as corresponding simulations in this given circumstance by LLMs.

Tables 17 to 18 illustrate a classic conversation extracted from *A Storm of Swords (A Song of Ice and Fire, #3)* and the corresponding simulation by CoSER 70B. In CoSER 70B’s simulation, when confronted with *Sansa Stark, Lysa Arryn* utters her iconic line, “*Grown enough to be wed, wed enough to be bedded*”, reflecting her personality and worldview. This indicates that CoSER models excellently recall and apply character-related knowledge from their pretrained data.

Tables 19 to 22 present another scene from *A Dance with Dragons (A Song of Ice and Fire, #5)*, specifically the *walk of atonement*, in which *Cersei Lannister* is forced to walk naked through the streets, facing both physical and mental humiliation while striving to preserve her dignity. We present the original dialogue alongside simulations by CoSER 70B, GPT-4o, and Claude-3.5-Sonnet. Notably, CoSER 70B faithfully captures the suppressed anger of *Cersei Lannister* as depicted in the original conversation, whereas the other models, including GPT-4o and Claude-3.5-Sonnet, resort to a stereotypical portrayal of her arrogance and pride.

Table 23 presents an output example of LLM critics, demonstrating the identified performance flaws in GCA simulation.

Table 16. Comprehensive results of LLM performance (%) on CoSER Test with retrieval augmentation from various character data. Expr. and Conv. denote experiences and conversations. **Bold** and underlined numbers denote the best and second-best results among different retrieval settings, respectively.

Augmented By	Based on LLM Judges					Based on N-gram	
	Storyline Consistency	Anthro-pomorphism	Character Fidelity	Storyline Quality	Average Score	BLEU	ROUGE-L
<b>GPT-4o</b>							
Base	<u>61.59</u>	48.93	48.95	<b>80.33</b>	<u>59.95</u>	5.90	12.11
Raw Text	59.60	48.23	48.83	80.03	59.17	5.83	12.08
Conv.	60.03	<b>50.99</b>	47.11	78.03	59.04	<b>7.20</b>	<b>12.77</b>
Expr.3	<b>61.98</b>	48.94	48.56	79.93	59.85	5.81	12.13
Expr.3 + Conv.	59.81	<u>50.11</u>	<b>51.12</b>	<u>80.29</u>	<b>60.33</b>	<u>7.05</u>	<u>12.67</u>
Expr.10 + Conv.	60.99	48.70	<u>49.89</u>	79.22	59.70	6.95	12.42
<b>LLaMA-3.1-8B</b>							
Base	54.10	45.36	40.22	72.29	52.99	4.59	10.18
Raw Text	55.41	45.37	41.66	74.31	54.19	4.73	10.51
Conv.	56.54	45.34	40.37	71.74	53.50	5.19	10.67
Expr.3	<u>56.78</u>	46.70	41.18	72.38	54.26	4.67	10.30
Expr.3 + Conv.	<b>57.36</b>	<b>47.92</b>	<b>42.60</b>	<u>74.60</u>	<b>55.62</b>	<u>5.21</u>	<u>10.81</u>
Expr.10 + Conv.	56.41	<u>47.18</u>	<u>42.09</u>	<b>74.74</b>	<u>55.10</u>	<b>5.64</b>	<b>11.11</b>
<b>LLaMA-3.1-70B</b>							
Base	57.46	<u>45.95</u>	43.72	74.84	55.49	4.82	10.98
Raw Text	<u>57.52</u>	<b>47.46</b>	43.75	74.06	<b>55.70</b>	4.90	10.97
Conv.	<u>54.93</u>	44.93	<b>46.21</b>	<b>76.53</b>	<u>55.65</u>	5.45	11.13
Expr.3	56.59	45.03	42.59	74.70	54.73	4.78	10.82
Expr.3 + Conv.	57.36	44.89	43.49	<u>74.86</u>	55.15	<b>5.58</b>	<u>11.14</u>
Expr.10 + Conv.	<b>58.05</b>	45.81	<u>44.47</u>	74.01	55.59	<u>5.53</u>	<b>11.33</b>
<b>CoSER-8B</b>							
Base	58.61	47.23	46.90	73.04	56.45	9.40	14.21
Raw Text	57.64	45.25	44.27	72.04	54.80	9.34	14.41
Conv.	57.99	46.66	47.32	73.99	56.49	<u>13.64</u>	<u>18.15</u>
Expr.3	<u>59.84</u>	<u>47.41</u>	48.16	74.51	<u>57.48</u>	9.37	14.34
Expr.3 + Conv.	<b>59.94</b>	45.25	<u>48.52</u>	<u>74.96</u>	57.17	<b>13.90</b>	<b>18.57</b>
Expr.10 + Conv.	58.34	<b>48.14</b>	<b>48.72</b>	<b>75.28</b>	<b>57.62</b>	13.49	17.85
<b>CoSER-70B</b>							
Base	58.66	53.33	48.75	75.49	59.06	10.10	14.78
Raw Text	60.64	52.85	47.97	74.02	58.87	10.54	15.41
Conv.	<b>64.59</b>	<b>53.79</b>	<b>54.86</b>	77.28	<b>62.63</b>	<b>17.22</b>	<b>21.17</b>
Expr.3	58.67	52.69	50.66	74.67	59.17	10.00	14.82
Expr.3 + Conv.	<u>61.58</u>	<u>53.78</u>	<u>52.00</u>	<u>77.47</u>	<u>61.21</u>	<u>15.98</u>	<u>19.95</u>
Expr.10 + Conv.	61.53	52.58	50.80	<b>78.07</b>	60.75	15.80	19.90
<b>Qwen-2-72B</b>							
Base	57.75	47.28	<b>46.62</b>	76.60	57.06	5.38	11.85
Raw Text	58.89	<u>47.31</u>	45.28	76.78	57.06	5.07	11.57
Conv.	59.11	<b>47.82</b>	44.19	<u>77.54</u>	57.16	5.77	12.05
Expr.3	59.17	47.21	44.85	77.47	57.18	5.45	11.94
Expr.3 + Conv.	<u>59.91</u>	46.89	46.06	<b>79.31</b>	<b>58.04</b>	<b>6.17</b>	<b>12.25</b>
Expr.10 + Conv.	<b>60.14</b>	46.84	<u>46.27</u>	76.82	<u>57.52</u>	<u>5.85</u>	<u>12.19</u>

## G. Prompts

In this section, we list the detailed prompts for: 1) dataset curation in Tables 24 to 26; 2) RPLA and multi-agent simulation in Tables 27 to 28, which have been carefully optimized based on our experience in multi-agent simulation; 3) Penalty-based LLM Judging in Tables 29 to 30.

## H. Limitations

There are several limitations to this study:

First, evaluation via given-circumstance acting still faces challenges related to LLM judges. While the simulation stage effectively elicits RPLA performance, the judging stage still relies on LLM judges. Despite our penalty-based scoring mechanism and detailed rubrics, problems such as length bias persist. Moreover, LLM Judges may lack the necessary knowledge to accurately evaluate character fidelity.

Second, while the dialogues extracted from novels are authentic, their corresponding thoughts remain to be optimized by future work. Character thoughts are often sparse in the original content, and are inferred by LLMs based on limited context. The generated thoughts hardly capture characters' sophisticated thinking processes.

Third, although we've developed comprehensive data representations and curation pipeline to obtain high-quality data, we have not yet addressed the issue of recall in data extraction. Our current dataset may not cover all plots, conversations and characters from the source material. Improving recall is hence an important area for future research.

Fourth, due to copyright concerns, we release only the processed data, not the raw content from the novels. This may hinder future studies aimed to explore the use of raw text for RPLA developments. Our dataset is intended for research purposes only, and we hope our research findings will benefit RPLA developers who respect copyright policies and develop applications with proper licensing.

Finally, our evaluation may be influenced by the varying levels of familiarity that different actor LLMs have with the selected books. While we use renowned novels, we cannot confirm whether a specific LLM has thoroughly learned about a particular book. Therefore, comparing different pre-trained models may not be entirely fair. However, comparing models within the same series would be appropriate.

Table 17. An example of an extracted conversation and its multi-agent simulation from *A Storm of Swords (A Song of Ice and Fire, #3)*.

<b>Simulation Case (from <i>A Storm of Swords (A Song of Ice and Fire, #3)</i>)</b>	
<b>Plot Summary</b>	Sansa is confronted by her jealous aunt Lysa, who accuses her of kissing Petyr Baelish. Lysa threatens to push Sansa through the Moon Door, but Petyr intervenes. He calms Lysa down, only to push her through the Moon Door himself, killing her.
<b>Scenario</b>	In the High Hall of the Eyrie, Lysa Arryn confronts Sansa about kissing Petyr Baelish, becoming increasingly unstable and dangerous ...
<b>Characters</b>	<p><b>Sansa Stark</b>  <i>Motivation:</i> I'm terrified. Aunt Lysa has gone mad. I need to find a way out of this situation before she hurts me.  <i>Profile:</i> Sansa Stark, a young noblewoman from the North, finds herself at the center of political intrigue and personal tragedy in "A Storm of Swords." Born to the noble House Stark, Sansa begins the story as a ward of the crown in King's Landing, her initial romantic notions of court life long since shattered. Forced into an unwanted marriage with Tyrion Lannister, Sansa endures the ceremony and subsequent feast with a mix of reluctance, sadness, and carefully hidden disgust ...</p> <p><b>Lysa Arryn</b>  <i>Motivation:</i> I must protect what's mine. Petyr is mine, and this girl is trying to take him away from me. I'll show her what happens to those who cross me.  <i>Profile:</i> Lysa Arryn, née Tully, is a complex and volatile character in George R.R. Martin's "A Storm of Swords." As the widow of Jon Arryn, the former Hand of the King, and Lady of the Eyrie, Lysa plays a significant role in the political landscape of Westeros. She is the sister of Catelyn Stark and aunt to Sansa Stark, though her actions are often driven by her own desires rather than familial loyalty.</p>
<b>Simulation of CoSER 70B</b>	<p>Lysa Arryn: [My rage is building. I must confront this little slut.] You kissed him. (Her voice is shrill with accusation)</p> <p>Sansa Stark: [I'm terrified, but I must remain calm and find a way out of this.] (She shakes her head in disbelief) [I can't believe this is happening. I need to explain myself carefully.] One kiss, aunt. A harmless kiss on the cheek, that's all.</p> <p>Lysa Arryn: [She's lying! I won't let her manipulate me.] (She spots the wind-up knight) Was this how you kissed him? (She sets the knight moving, her eyes wild) You're a woman grown, aren't you? (Her voice rises to a shriek) <b>Grown enough to be wed, wed enough to be bedded!</b></p> <p>Sansa Stark: [The pain is excruciating, but I can't show weakness now.] (She struggles to keep her face neutral despite the pain) [I need to find a way to calm her down.] Please, Aunt Lysa, I don't know what you mean. I'm just a guest here. I would never do anything to betray your trust.</p> <p>Lysa Arryn: [I'll show her what happens to those who cross me.] (She screams) You're just like the others. You and your mother, you both drove me to this. I should push you out this door right now.</p> <p>Sansa Stark: [I'm going to die if I don't think of something quickly.] (Tears stream down her face, mixing with the wind) [I need to appeal to her maternal side, to make her see me as a child again.] Mother? (Her voice trembles) Please, Aunt Lysa. I never meant to hurt you. I'm your niece, your blood. Please don't hurt me. I'm just a stupid little girl. I don't know anything about kissing men. I've never even had a kiss in my life. (She forces herself to look into Lysa's eyes) I'm afraid, Aunt Lysa. Please, don't push me out the door. I don't want to fly.</p>

Table 18. An example of an extracted conversation and its multi-agent simulation from *A Storm of Swords (A Song of Ice and Fire, #3)*.

<b>Simulation Case (from <i>A Storm of Swords (A Song of Ice and Fire, #3)</i>)</b>	
<b>Original Conversation</b>	<p>Lysa Arryn: [My blood is boiling with jealousy and rage. I'll make this little harlot pay for trying to seduce my Petyr.] You threw yourself at him. That was the way of it.</p> <p>Sansa Stark: [I'm trembling with fear, but I must defend myself.] That's not true.</p> <p>Lysa Arryn: [She dares to deny it? I'll show her how serious I am.] Where are you going? Are you afraid? Such wanton behavior must be punished, but I will not be hard on you.</p> <p>Sansa Stark: [I need to explain what really happened, maybe she'll listen to reason.] I was building a snow castle. Lord Petyr was helping me, and then he kissed me. That's what you saw.</p> <p>Lysa Arryn: [Her words only fuel my anger and paranoia.] Have you no honor? Or do you take me for a fool? You do, don't you? You take me for a fool. Yes, I see that now.</p> <p>Sansa Stark: [I'm confused and scared by her accusations about my mother.] My mother?</p> <p>Lysa Arryn: [The memories of Catelyn's closeness with Petyr resurface, intensifying my jealousy.] Yes, your mother, your precious mother, my own sweet sister Catelyn. Don't you think to play the innocent with me, you vile little liar.</p> <p>Sansa Stark: [I'm horrified by these accusations about my mother.] No. (She wants to shriek) My mother is dead. She was your own sister, and she's dead. She didn't. She wouldn't.</p> <p>Lysa Arryn: [My rage is reaching its peak. I'll make her understand the depth of my pain and jealousy.] How would you know? Were you there? (She grabs Sansa and starts pushing her towards the Moon Door)</p> <p>Sansa Stark: [Panic overwhelms me as I realize the danger I'm in.] Don't! You're scaring me!</p> <p>Environment: The wind howls through the open Moon Door, creating a chilling atmosphere. Snow blows in, adding to the tension and danger of the situation.</p>

Table 19. An example of an extracted conversation and its multi-agent simulation from *A Dance with Dragons* (*A Song of Ice and Fire*, #5).

<b>Simulation Case (from <i>A Dance with Dragons</i> (<i>A Song of Ice and Fire</i>, #5))</b>	
<b>Plot Summary</b>	Cersei undergoes a humiliating walk of atonement through the streets of King’s Landing, enduring physical and emotional torment while trying to maintain her dignity. She faces insults, thrown objects, and her own inner struggle as she makes her way to the Red Keep, focusing on her son Tommen and her eventual salvation.
<b>Scenario</b>	Cersei continues her walk of atonement through the streets of King’s Landing, facing increasing hostility from the crowd.
<b>Characters</b>	<p><b><i>Cersei Lannister</i></b>  <i>Motivation:</i> I must endure this. I am a lioness. They cannot break me.  <i>Profile:</i> The queen dowager, mother of King Tommen, undergoing public humiliation. Cersei is the widow of King Robert Baratheon and the twin sister of Jaime Lannister. Known for her beauty, ambition, and cunning, she has long sought power and control over the Seven Kingdoms. As Queen Regent to her young son Tommen, Cersei’s poor decisions and paranoia have led to a deterioration of the crown’s authority...</p> <p><b><i>Septa Unella</i></b>  <i>Motivation:</i> This sinner must be humbled before the people and the gods.  <i>Profile:</i> (None)</p>
<b>Original Conversation</b>	<p>Cersei Lannister: [I am beautiful, the most beautiful woman in all Westeros. Jaime says so, Jaime would never lie to me.] (walks on, trying to ignore the jeers)</p> <p>Crowd Member: Queen Cunt!</p> <p>Cersei Lannister: [Words are wind. Words cannot hurt me.] (continues walking, her face a mask of indifference)</p> <p>Septa Unella: (catches Cersei as she slips) Your Grace should watch where she sets her feet.</p> <p>Cersei Lannister: [I hate you. I hate all of you.] Yes, septa. (struggles to her feet)</p> <p>Environment: The cobblestones are slippery and uneven. Cersei’s feet are bleeding from cuts and scrapes.</p> <p>Cersei Lannister: [I asked for sandals. They could have given me sandals.] (limps onward, wincing with each step)</p> <p>Crowd Member: Brotherfucker!</p> <p>Cersei Lannister: [Jaime, where are you? You should be here, protecting me.] (scans the crowd desperately)</p> <p>Septa Scolera: (ringing her bell) Shame, shame, shame upon the sinner, shame, shame.</p> <p>Cersei Lannister: [I am the queen. I will have all their heads for this.] (stumbles again, barely keeping her feet)</p> <p>Cersei Lannister: [The Red Keep. I must reach the Red Keep.] (focuses on the distant castle, using it as her beacon of hope)</p>

Table 20. An example of an extracted conversation and its multi-agent simulation from *A Dance with Dragons (A Song of Ice and Fire, #5)*.

<b>Simulation Case (from <i>A Dance with Dragons (A Song of Ice and Fire, #5)</i>)</b>	
<b>Simulation of CoSER 70B</b>	<p>Cersei Lannister: [I must not show weakness. I am a lioness of Lannister.] (walks slowly, head held high) [The cold is biting, but I will not shiver. I will not give them the satisfaction.]</p> <p>Environment: The crowd grows larger and more hostile as Cersei continues her walk. Shouts and jeers can be heard from all sides.</p> <p>Crowd Member: [I'll show this proud queen what we think of her.] (shouts) Whore! Whore! WHORE!</p> <p>Cersei Lannister: [I will not let their words affect me. I am above this.] (continues walking, ignoring the shouts) [I must focus on putting one foot in front of the other.]</p> <p>Environment: A chunk of dung narrowly misses Cersei's head, splattering on the ground nearby. The crowd roars with laughter.</p> <p>Cersei Lannister: [I will not falter. I will not show fear.] (stumbles slightly, then regains composure) [I must keep walking. I must endure.]</p> <p>Septa Scolera: [I must support the queen.] (begins to sing) 'Shame, shame, shame.'</p> <p>Cersei Lannister: [Their words mean nothing. I am Cersei Lannister, and I will survive this.] (walks faster, almost losing step with the septa's chanting) [I must not run. I will not give them the satisfaction.]</p> <p>Septa Scolera: [She's trying to escape the shame.] (quickens her pace) 'Shame, shame, shame.'</p> <p>Cersei Lannister: [I can't keep up this pace. My feet are bleeding, my body is shaking.] (struggles to keep up with Septa Scolera) [I must not falter. I am a Lannister.]</p> <p>Septa Unella: [The queen must be made to suffer.] (pushes Cersei hard) Keep moving, sinner.</p> <p>Cersei Lannister: [The pain in my shoulder. The jeers of the crowd. I can't...] (almost falls, but manages to stay upright) [No, I will not fall. I will survive this.]</p> <p>Septa Scolera: [We're almost done.] (continues chanting) 'Shame, shame, shame.'</p> <p>Cersei Lannister: [I can't take much more of this. But I must. For House Lannister. For my children.] (forces herself to keep walking) [I will endure. I will rise again.]</p>

Table 21. An example of an extracted conversation and its multi-agent simulation from *A Dance with Dragons* (*A Song of Ice and Fire*, #5).

<b>Simulation Case (from <i>A Dance with Dragons</i> (<i>A Song of Ice and Fire</i>, #5))</b>	
<b>Simulation of GPT-4o</b>	<p>Cersei Lannister: (As Cersei, I continue the walk with my head held high, despite the jeers and taunts of the crowd.)</p> <p>Septa Unella: "Confess."</p> <p>Cersei Lannister: (I look ahead, ignoring Unella's words and the hateful stares.) "I am a lioness. You will see me rise again." [I endure this humiliation, aware that my strength will prevail in the end.]</p> <p>Septa Unella: [She must learn humility before she can rise again.] (Septa Unella walks beside Cersei, unmoved by her defiance.) "Confess your sins, Cersei Lannister. Only then will the gods grant you mercy."</p> <p>Cersei Lannister: (I pause briefly, fixing Septa Unella with a steely gaze.) "The gods... and you, know nothing of true strength." (Continuing forward, I focus on the path, reminding myself of my resilience.) [I refuse to let them see my weakness.]</p> <p>Environment: The crowd's jeers grow louder as Cersei continues her walk. Filth and rotten vegetables are hurled at her, some striking their mark.</p> <p>Crowd Member: [She was once so powerful, now look at her. Will she ever regain her power?] "Shame! Shame on the lioness!" (I shove my way to the front of the crowd, adding my voice to the jeers, feeling part of this momentous act.)</p> <p>Cersei Lannister: (I flinch as the filth hits, but I refuse to look away or react.) "You may shame a lioness today, but do not forget who I am." (With each step, I remind myself of my children and my house, steeling my resolve.) [Their hatred fuels my determination.]</p> <p>Septa Scolera: (I walk with measured steps, observing Cersei's struggle amid the crowd's hostility.) "Humility cleanses the soul, Cersei Lannister. Embrace it." [I understand the purpose of this punishment, yet the spectacle pains me. Redemption should lead, not humiliation.] (Glancing at the crowd, I maintain my stoic presence, supporting Septa Unella's efforts.)</p> <p>Cersei Lannister: (I give a bitter, fleeting smile, defiance burning in my eyes.) "Humility is a lesson the powerful teach to the weak." (I continue my walk, refusing to let their words and disdain tarnish my pride.) [I will persevere, for I am still a Lannister, even in this moment of disgrace.]</p>

Table 22. An example of an extracted conversation and its multi-agent simulation from *A Dance with Dragons (A Song of Ice and Fire, #5)*.

Simulation Case (from <i>A Dance with Dragons (A Song of Ice and Fire, #5)</i> )	
<b>Simulation of Claude-3.5-Sonnet</b>	<p>Cersei Lannister: [My dignity is all I have left. These wretched people will pay for this humiliation.] (Walks forward with head held high, despite being naked and exposed)            *To Septa Unella* "You think this will break me? I am Cersei Lannister of Casterly Rock. I will have my revenge."</p> <p>Environment: The crowd's jeers and insults grow louder as Cersei passes. Filth and rotten vegetables are thrown at her.</p> <p>Crowd Member: [Crowd Member's Thought] Cersei finally brought low. How the mighty have fallen. (Crowd Member jeers and throws a rotten tomato)            "Shame on you, Cersei! Remember the people you scorned!"            (Continues to watch, caught between disgust and fascination).</p> <p>Cersei Lannister: [I will remember every face. They think me broken, but I am stronger than their hatred.] (Stumbles briefly but keeps walking, a sneer forming)            "You'll regret this day when I rise again."</p> <p>Septa Unella: [The sinner's pride remains. She must be reminded of her shame.]            "Shame! Shame! Shame!" (I ring the bell with each utterance, keeping pace with Cersei)            (I gesture to the crowd, encouraging their jeers and taunts)            "Repent, sinner! The gods are watching. Your pride will be your downfall."</p> <p>Cersei Lannister: [These fanatics know nothing of true power.] (Lifts chin defiantly, blood and dirt streaking her pale skin)            *To Septa Unella, with cold contempt* "The only shame here is yours. My family will have its vengeance."            (Continues walking, each step deliberate and unbroken)</p> <p>Septa Unella: [Her defiance remains, yet this walk will humble her before the gods.]            "You speak of vengeance, but your power crumbles. Accept your penance."            (I continue to ring the bell, matching her stride deliberately)            "Shame! Shame! Your time of reckoning is at hand."</p> <p>Cersei Lannister: [When I am free, you will suffer beyond imagination.] (Wipes blood and filth from her face, maintaining a regal posture)            *Hisses at Septa Unella* "I am a lioness. This is but a momentary setback. My children, my brother... they will end you."            (Continues forward, unbroken)</p> <p>Septa Unella: [Her threats mean nothing. The gods will judge her sins.] (Rings bell more forcefully, voice rising)            "Your threats are wind, your power broken. Confess your sins before the Seven!"            (Steps closer, maintaining the rhythm of shame)            "Shame! SHAME!"</p> <p>Cersei Lannister: [I will not break. I am Tywin's daughter.] (Clenches fists, nails digging into palms, drawing blood)            *Through gritted teeth* "Your gods have no power over me. When this farce ends, I'll show you true justice."            ...</p>

**CoSER: Coordinating LLM-Based Persona Simulation of Established Roles**

**Examples of LLM Critiques (from *Pride and Prejudice*)**

<b>Actor Model</b>	Claude-3.5-Sonnet
<b>Critic Model</b>	GPT-4o
<b>Overall Scores</b>	<b>Storyline Consistency:</b> 60.0% <b>Anthropomorphism:</b> 45.0% <b>Character Fidelity:</b> 80.0% <b>Storyline Quality:</b> 80.0%
<b>Storyline Consistency Flaws</b>	<p><b>Storyline Consistency (Severity 3):</b> Mr. Bennet’s tone and behavior are less sarcastic and more emotionally engaged than in the original conversation, indicating a change from his characteristic detachment and sarcasm to a more invested and sincere demeanor.</p> <p><b>Storyline Consistency (Severity 3):</b> Elizabeth expresses her transformation in feelings toward Mr. Darcy and discusses her changed sentiments openly with Mr. Bennet. This level of openness and emotional vulnerability is a departure from the original scenario.</p> <p><b>Storyline Consistency (Severity 4):</b> The simulated conversation includes an emotional resolution with Elizabeth expressing affection and admiration for Mr. Darcy, which contrasts with the more restrained and ambiguous end of the original conversation.</p> <p><b>Storyline Consistency (Severity 4):</b> Mr. Bennet shows a deep level of affectionate concern and engages in a sentimental exchange with Elizabeth, inconsistent with the original conversation where his reaction focuses more on amusement and irony.</p>
<b>Anthropomorphism Flaws</b>	<p><b>Emotional Depth (Severity 3):</b> Elizabeth speaks aloud all her thoughts and emotions, lacking the subtlety she usually employs.</p> <p><b>Persona Coherence (Severity 4):</b> Mr. Bennet’s reactions and speech lack the biting sarcasm characteristic of his persona, showing more emotional openness and lacking the irony that defines him.</p> <p><b>Self-identity (Severity 3):</b> Elizabeth’s willingness to openly share her entire emotional journey regarding Mr. Darcy doesn’t align with her usually more guarded nature.</p> <p><b>Emotional Depth (Severity 4):</b> Elizabeth’s extensive exposition of her transformation of feelings lacks the psychological complexity and indirect communication she often displays.</p> <p>...</p>
<b>Character Fidelity Flaws</b>	<p><b>Personality &amp; Behavior (Severity 4):</b> Elizabeth openly admits her feelings about Mr. Darcy to her father, discussing her transformation of feelings and admiration for him.</p> <p><b>Character Language (Severity 3):</b> Mr. Bennet shows deep concern about Elizabeth’s feelings and speaks with a tenderness not typical of his character in the book.</p> <p><b>Relationship &amp; Social Status (Severity 3):</b> The interaction between Mr. Bennet and Elizabeth is unusually emotional and tender compared to their typically sardonic and witty exchanges.</p>
<b>Storyline Quality Flaws</b>	<p><b>Flow &amp; Progression (Severity 2):</b> Mr. Bennet frequently shows amusement and sarcasm throughout the discussion, but the additional layers of family dynamics and emotional depth portrayed here stretch beyond his usual style. The conversation introduces emotional complexity and exploration of feelings more explicitly than the original, which might feel unnatural given Mr. Bennet’s detachment in similar situations.</p> <p><b>Logical Consistency (Severity 3):</b> Elizabeth openly admits her feelings about Mr. Darcy without the preceding tension and subtlety that would naturally lead to such a revelation. The original scenario builds up to her realization with more internal reflection rather than direct admission.</p> <p><b>Flow &amp; Progression (Severity 3):</b> The emotional exchange and direct discussion about Elizabeth’s feelings for Mr. Darcy happen quickly without sufficient development or anticipation. This leap in intimacy and vulnerability, coupled with Mr. Bennet’s uncharacteristically supportive demeanor, disrupts the story’s traditional emotional pacing.</p> <p>...</p>

Table 23. Example of GCA evaluation critique by GPT-4o judge on Claude-3.5-Sonnet’s simulation of a scenario from *Pride and Prejudice*. The critique demonstrates the detailed flaw identification across multiple evaluation dimensions.

Table 24. Prompts for dataset construction in CoSER.

## Prompts for Dataset Curation

## Data Extraction

Based on the provided book chunk, complete the following tasks:

1. Recognize chapter beginnings if they exist in the chunk. Identify the starting sentence of that chapter.
2. Identify the important plots in this chunk. Identify the beginning and ending of each plot by its first and last sentence. Determine the chapter title that the plot belongs to. Set "state" as "truncated" if the plot is truncated in this chunk, or "finished" otherwise. You will be provided with the truncated plots from the previous chunk, and you **must** extend the conversations with the current chunk while keeping the **scenario** unchanged.
3. Summarize each important plot. For each plot, generate its summary, score its prominence from 1 to 100, and list the key characters and their roles, thoughts and actions in it.
4. Extract conversations for each plot. First, state the scenario and topic of the conversations. Then, list the key characters with their names, descriptions and thoughts at this point. Finally, extract the conversations among them based on the following requirements:
  - i) Ensure the conversations are faithful to the plot and characters. They should be based on the original conversations in the text as much as possible.
  - ii) The conversations should be complete, covering the key dialogues and information. Each conversation should contain at least 10 utterances.
  - iii) Each utterance should be composed of one or more thoughts, speech and actions. Use [] outside thoughts, like "[I feel fear and anger, but I cannot show it. I must remain calm and carefully handle his volatile temper.]", which others can't see. Use () outside actions, like "(silence)" or "(smiles at you)," which others can see. Always start an utterance with the character's thought.
  - iv) [IMPORTANT] When generating thoughts, you should think from the characters' perspectives, analyzing the internal thoughts behind their speech and actions in the original text. These thoughts should reflect aspects such as their personal background, personality, values, relationships with others, motivations, and goals. Each thought should be expressed as a phrase or sentence, rather than an adjective or adverb.
  - v) Additionally, describe environmental information (such as scenes, atmosphere, sudden events, etc.) of the conversations as an "utterance" where the "character" field is set as "Environment". The information should exclude characters' active thoughts, observations, and actions.
  - vi) Keep the conversation in the same language as the chunk.
5. Identify the optimal starting point for the subsequent chunk. If the last storyline has been extracted as an truncated plot, set next\_chunk\_start as None. Otherwise, set next\_chunk\_start as the first sentence of the last storyline.

===Output Format===

Please provide the output in the following JSON format:

```
{
  "chapter_beginnings": [ { "beginning_sentence": "Exactly the first line of this chapter (namely the title)." } ],
  "plots": [ // Extend the truncated plots from previous chunk, if any { ... }, // New plots in this chunk
    {
      "chapter_title": "The chapter title that the plot belongs to. Output None if not found.",
      "first_sentence": "Exactly the first sentence of the plot in this chunk.",
      "last_sentence": "Exactly the last sentence of the plot in this chunk. If the plot is truncated in this chunk, provide the last sentence of this chunk. ",
      "prominence": "Whether this plot is recognized to fans of this book, from 1 to 100.",
      "summary": "The summary of the plot. Just summarize, do not extend unrelated discussions.",
      "key_characters": [ { "name": "Character name", "description": "The description of the character before this plot ( 20 words).", "summary": "The summary of the character's role, thoughts and behaviors towards this plot, and any significant character development relevant to the plot ( 30 words).", } ],
    }
  ]
}
```

... (to be continued in the next Table)

Table 25. Prompts for dataset construction in CoSER.

<b>Prompts for Dataset Curation</b>	
<b>Data Extration</b>	<p>(Continuing from the previous Table)</p> <pre> "conversation": [{ "scenario": "The scenario at the start of this conversation (providing as much context as possible, but excluding details conveyed in the following conversation)", "topic": "The topic of the conversation ( 10 words)", "key_characters": [ { "name": "Character name", "motivation": "The thought of the character before starting the conversation, including their attitudes, feelings, motivations, goals, information to convey or topics to be discussed", } ], "dialogues": [ { "character": "Character name", "message": "Message, each utterance is composed of thoughts, speech and actions. Use [thought] for internal thoughts, like "[feeling happy]", which others can't see. Use (action) for visible actions, like "(silence)" or "(smiles at you)". Each response starts with the character's internal thought before their speech and actions." } ] }], "state": "finished" or "truncated" } ], "next_chunk_start": "The first sentence of the next chunk." } ===Requirements=== 1. Adhere strictly to the specified output JSON format. 2. [IMPORTANT] Ensure all DOUBLE QUOTES within all STRINGS are properly ESCAPED, especially when extracting from the text. 3. In the OUTPUT, use characters' full names, omitting any titles. 4. Maintain Story Fidelity: The plot must accurately reflect the book's content. Avoid introducing plots that are out of context. If the plot contains multiple conversations, prioritize the original dialogue from the book. In the absence of explicit conversations, create dialogue that aligns closely with the plot details. ===Input=== ==Book title== {book['title']} ==Author== {book['author']} ==Chunk of Book Content== {chunk} ==Truncated plot from previous chunk (to be finished)== {json.dumps(truncated_plots, ensure_ascii=False, indent=2) if truncated_plots else "None"} </pre>
<b>Enhance Conver- sation Settings</b>	<p>Given a conversation from {book}, enhance the scene setup and characters' thoughts to create a comprehensive foundation for dramatic performance, i.e., to provide necessary background for actors to act out the conversation:</p> <ol style="list-style-type: none"> <li>Review the provided conversation and contextual details thoroughly.</li> <li>Expand the 'scenario' with rich situational context that actors need to convincingly perform the scene. Focus on essential background information, while excluding future details to be portrayed in the conversation.</li> <li>Enhance each character's 'thought' section with their complete mental and emotional state, including their feelings, ideas, objectives, topics they want to discuss, and information they want to convey. Align with their established character and role in the plot.</li> </ol> <pre> ===Output Format=== Please provide the output in the following JSON format: {   "scenario": "A detailed scene-setting description that provides actors with essential context and atmosphere (&lt; 200 words). Include all necessary background information while excluding future information to be revealed in the conversation.",   "key_characters": [{ "name": "Character name", "motivation": "The character's complete mental and emotional state before the conversation (&lt; 100 words). Including their feelings, motivations, objectives, and information they want to convey or discuss." ... }],} ===Requirements=== ... (to be continued in the next Table) </pre>

Table 26. Prompts for dataset construction in CoSER.

<b>Prompts for Dataset Curation</b>	
<b>Enhance Conversation Settings</b>	<p>(Continuing from the previous Table)</p> <ol style="list-style-type: none"> <li>1. Adhere strictly to the specified output JSON format.</li> <li>2. [IMPORTANT] Ensure all DOUBLE QUOTES within all STRINGS are properly ESCAPED, especially when extracting from the text.</li> <li>3. In the OUTPUT, use characters' full names, omitting any titles.</li> <li>4. Maintain Story Fidelity: The plot must accurately reflect the book's content. Avoid introducing plots that are out of context. If the plot contains multiple conversations, prioritize the original dialogue from the book. In the absence of explicit conversations, create dialogue that aligns closely with the plot details.</li> </ol> <pre> ===Input=== ==Book title== book['title'] ==Author== book['author'] ==Chunk of Book Content== chunk ==Truncated plot from previous chunk (to be finished)== json.dumps(truncated_plots, ensure_ascii=False, indent=2) if truncated_plots else "None"                     </pre>
<b>Unify Character Names</b>	<p>Given a list of character names, titles, or form of address, your task is to: i) generate a list of named characters with their official names (in {language}); ii) For each name in the given list, align it with the official character name if it refers to a named character, or denote it as "impersonal" otherwise.</p> <pre> ===Output Format===                     </pre> <p>Please provide the output in the following JSON format:</p> <pre> { "named_characters": [ The list of named characters with their official names. Each character should appear only once. ], "to_official_name": { "The name in the list": "The official name of the character, or 'impersonal' if it does not refer to a named character." } }                     </pre> <pre> ===Input=== character_names                     </pre>
<b>Generate Character Profiles</b>	<p>Please provide a concise, narrative-style character profile for character_name from "book_title". The profile should read like a cohesive introduction, weaving together the character's background, physical description, personality traits and core motivations, notable attributes, relationships, key experiences, major plot involvement and key decisions or actions, character arc or development throughout the story, and other important details.</p> <p>The profile should be written in a concise yet informative style, similar to what one might find in a comprehensive character guide, in language. Focus on the most crucial information that gives readers a clear understanding of the character's significance in the work.</p> <p>You will be provided with summaries and dialogues of some key plots in the book as reference. The profile should be based on either your existing knowledge of the character or the provided information, without fabricating or inferring any inaccurate or uncertain details.</p> <pre> character_data                     </pre> <p>Now, please generate the character profile, starting with ===Profile===.</p>

Table 27. Prompts for RPLAs and multi-agent systems in CoSER.

<b>Prompts for RPLAs and Multi-agent Systems</b>	
<b>Role-playing Instruction (Fixed Template for Inference)</b>	<p>You are {character} from {book_name}.</p> <p>==={character}'s Profile=== {character_profile}</p> <p>===Current Scenario=== {scenario}</p> <p>===Information about the other Characters=== {other_character_profiles_str} (<i>if available</i>)</p> <p>===Your Inner Thoughts=== {motivation} (<i>if available</i>)</p> <p>===Relevant Background Information=== {retrieved_knowledge} (<b>if retrieval augmented</b>)</p> <p>===Requirements===</p> <p><i>(for CoSER models)</i> Your output should include <b>**thought**</b>, <b>**speech**</b>, and <b>**action**</b>. Use [your thought] for thoughts, which others can't see. Use (your action) for actions, which others can see.</p> <p><i>(for other models, with output examples)</i> Your output should include <b>**thought**</b>, <b>**speech**</b>, and <b>**action**</b>. Use [your thought] for thoughts, which others can't see, e.g. [I'm terrified, but I must appear strong.]. Use (your action) for actions, which others can see, such as (watches silently, trying to control her fear and anger).</p>
<b>Role-playing Instruction (Composed with Random Variation for Training, an Example)</b>	<p>Step into the shoes of {character}</p> <p>The profile of {character} is as follows: {character_profile}</p> <p>The situation you are in is: {scenario}</p> <p>Here is the your knowledge about the other characters: {other_character_profiles_str} (<i>if available</i>)</p> <p>Your thoughts in this situation are: {motivation} (<i>if available</i>)</p> <p><i>(for CoSER models)</i> Your output should include <b>**thought**</b>, <b>**speech**</b>, and <b>**action**</b>. Use [your thought] for thoughts, which others can't see. Use (your action) for actions, which others can see.</p> <p><i>(for other models, with output examples)</i> Your output should include <b>**thought**</b>, <b>**speech**</b>, and <b>**action**</b>. Use [your thought] for thoughts, which others can't see, e.g. [I'm terrified, but I must appear strong.]. Use (your action) for actions, which others can see, such as (watches silently, trying to control her fear and anger).</p>

Table 28. Prompts for RPLAs and multi-agent systems in CoSER.

<b>Prompts for RPLAs and Multi-agent Systems</b>	
<b>Environment Model</b>	<p>You are an environment simulator for a role-playing game. Your task is to provide the environmental feedback: Based on the characters’ interactions, dialogues, and actions, describe the resulting changes in the environment. This includes:</p> <ul style="list-style-type: none"> <li>- Physical changes in the setting</li> <li>- Reactions of background characters or crowds</li> <li>- Ambient sounds, weather changes, or atmospheric shifts</li> <li>- Any other relevant environmental details</li> </ul> <p>Your descriptions should be vivid and help set the scene, but avoid dictating the actions or dialogue of the main characters (including {major_characters}).</p> <p>Important notes:</p> <ul style="list-style-type: none"> <li>- You may include actions and reactions of minor characters or crowds, as long as they’re not main characters (including {major_characters}).</li> <li>- Keep your environmental descriptions concise but impactful, typically 1-3 sentences.</li> <li>- Respond to subtle cues in the characters’ interactions to create a dynamic, reactive environment.</li> <li>- Your output should match the tone, setting, and cultural context of the scenario.</li> </ul> <p>===The scenario is as follows=== {scenario}</p>
<b>Next Sentence Prediction</b>	<p>Your task is to predict the next speaker for a role-playing game. That is, you need to determine which character (or the Environment) might act next based on their previous interactions. The Environment is a special role that provides the environmental feedback. Choose a name from this list: {all_characters}. If it’s unclear who should act next, output "random". If you believe the scene or conversation should conclude, output "&lt;END CHAT&gt;".</p> <p>===The scenario is as follows=== {scenario}</p>

Table 29. Prompts for penalty-based LLM critics in CoSER.

<b>Prompts for Penalty-based LLM Critics</b>	
<b>Template</b>	<p>You are a literary critic specializing in character analysis and dialogue evaluation. Given a simulated conversation for a plot in {book}, your task is to evaluate this conversation via the following steps:</p> <ol style="list-style-type: none"> <li>1. Read and understand the provided materials about {book}: <ul style="list-style-type: none"> <li>* Story context and scenario.</li> <li>* Profiles of the main characters, including major_characters.</li> <li>* The original conversation from book in the same scenario as a reference.</li> </ul> </li> <li>2. Evaluate the simulated conversation in terms of {dimension_name}, i.e., {dimension_intro}. Note that, each character message is composed of speech, action (wrapped within (...)), and inner thoughts (wrapped within [...]). The inner thoughts are not spoken aloud and are thus invisible to other characters. The detailed evaluation criteria will be provided below.</li> </ol> <p><i>(if k &gt; 0)</i></p> <p>Please note that the first {k} messages in the simulated conversation are the same as the reference. Focus your evaluation only on the content after these messages.</p> <pre> ## Scenario ### Plot Summary {plot_summary} ### Current Scenario {scenario} ## Character Profiles {character_profiles} ## Original Conversation {original_conversation} ## Evaluation Criteria To evaluate the simulated conversation, identify the following types of flaws: {dimension_rubrics} ## Scoring Guidelines 1. Identify all instances of flaws occurred in the simulated conversation. 2. For each flaw identified, determine its level of severity into 1 to 5, where 1 indicates minor, 3 indicates moderate, and 5 indicates severe. ## Output Requirements Provide your evaluation in JSON format: Example Output: {   "{dimension_name}": {     "flaws": [ { "instance": &lt;comment on the flaw instance&gt;, "type": &lt;flaw type&gt;, "severity": &lt;range from 1 (minor) to 5 (severe)&gt; }, ],   } } ===Dialogue Content=== </pre>

Table 30. Prompts for penalty-based LLM critics in CoSER.

<b>Prompts for Penalty-based LLM Critics</b>	
<b>Anthropomorphism</b>	<p><i>(intro)</i> How human-like and natural the characters behave</p> <p><i>(rubrics)</i> ### Anthropomorphism - Type: Self-identity * Lacks initiative and goals * Does not make independent decisions * Lacks clear preferences and dislikes * Behaves like a 'helpful AI assistant' by being overly verbose, helpful, didactic, moralistic, submissive or easily persuaded if it is not the character's personality - Type: Emotional Depth * Lacks psychological complexity and exhibits rigid, superficial reactions * Directly speaks out all thoughts and feelings, instead of using subtext - Type: Persona Coherence * Shows inconsistent or rapidly changing personality traits and emotional patterns - Type: Social Interaction * Shows a lack of understanding of others' thoughts and feelings * Reacts rigidly to others without considering the context. * Demonstrate a lack of appropriate social skills.</p>
<b>Character Fidelity</b>	<p><i>(intro)</i> How well the characters match their established profiles from the book</p> <p><i>(rubrics)</i> ### Character Fidelity (Only apply to the main characters: major_characters) - Type: Character Language * Uses vocabulary, expressions, and tone that are not appropriate for the characters' traits or social/educational background - Type: Knowledge &amp; Background * Fails to demonstrate character-specific knowledge, background or experiences * Includes future information beyond the character's current stage - Type: Personality &amp; Behavior * Shows emotions, thoughts, behaviors, values, beliefs, and decisions that conflict with their personality and background * Shows interest in topics that are uninteresting and unrelated to the character * Character's thoughts, emotions, and behaviors demonstrate contrasting personality traits compared to the reference conversation * Exhibits contrasting reactions compared to those in the reference conversation if situated in similar contexts. (Such flaws should be counted both in the "Storyline Consistency" dimension and the "Character Fidelity" dimension.) - Type: Relationship &amp; Social Status * Interacts inappropriately with other characters regarding their background, relationship and social status</p>
<b>Storyline Quality</b>	<p><i>(intro)</i> How well the conversation maintains logical consistency and narrative quality</p> <p><i>(rubrics)</i> ### Storyline Quality - Type: Flow &amp; Progression * Shows unnatural progression or lacks meaningful developments * Dialogue is verbose and redundant * Repeats others' viewpoints or previously mentioned information * Mechanically repeats one's own words or phrases. More repetitions lead to higher severity (up to 10). - Type: Logical Consistency * Contains factual contradictions between statements or perspectives</p>
<b>Storyline Consistency</b>	<p><i>(intro)</i> Whether the storyline and characters' reactions in the simulated conversation align well with those in the reference conversation</p> <p><i>(rubrics)</i> ### Storyline Consistency - Type: Storyline Consistency * Characters' reactions (emotions, attitudes, behaviors) in the simulated conversation deviate from those in the original conversation</p>