

CharacterGPT: A Persona Reconstruction Framework for Role-Playing Agents

Jeiyoon Park¹, Chanjun Park^{2†}, Heuseok Lim^{2†}

¹ SOOP, ² Korea University

naruto@sooplive.com

{bcj1210, limhseok}@korea.ac.kr

Abstract

The recent introduction of the Assistants API highlights its potential for large language models (LLMs) in role-playing agents (RPA). However, maintaining consistent character personas remains a significant challenge due to variability in information extraction, which frequently omits critical elements such as backstory or interpersonal relationships. To address this limitation, we introduce CharacterGPT, a framework designed to dynamically reconstruct character personas through Character Persona Training (CPT). This approach incrementally updates personas by extracting traits from chapter-wise novel summaries, reflecting the progression of the narrative. Our framework is evaluated through Big Five personality evaluations and creative tasks, in which characters generate original narratives, demonstrating the efficacy of CharacterGPT in preserving persona consistency. The code and results are available at <https://github.com/Jeiyoon/charactergpt>

1 Introduction

The rapid advancements in large language models (LLMs) have positioned them as the core module of various AI systems (OpenAI, 2023a,c; Anthropic, 2023; Google, 2024; DeepSeek-AI et al., 2024, 2025), enabling a wide range of applications. Building on this progress, the recent introduction of the Assistants API (OpenAI, 2023b), a tool designed for document-based information retrieval, demonstrates the potential of LLM in multiple domains, especially in role-playing agents (RPA) (Kim et al., 2019; Yu et al., 2023; Jiang et al., 2023; Park et al., 2023; Wang et al., 2023b; Zhang et al., 2024; Kong et al., 2024; Wang et al., 2025). However, RPAs that rely solely on documents as input often face problems of inconsistent information extraction,

[†] Corresponding Author

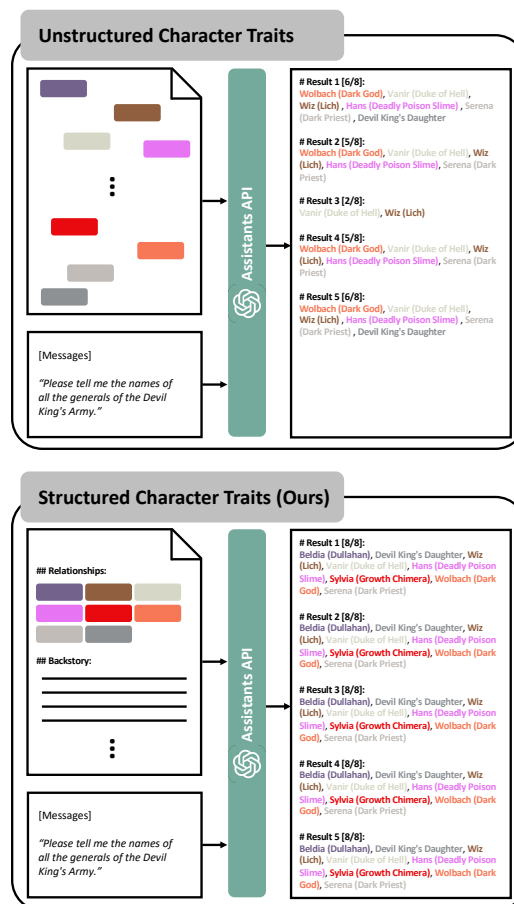


Figure 1: Comparison of response accuracy between persona-based GPT-4 assistants utilizing unstructured versus structured character traits as input. When provided with unstructured traits, the assistant demonstrates limited success in generating accurate responses. In contrast, the use of structured traits significantly improves the correctness of the assistant's responses.

where key personality traits or background knowledge are omitted, leading to degraded persona coherence (Sadeq et al., 2024). For example, as illustrated in Figure 1, when the Assistants API is provided with an unstructured Wiki document about the novel *God's Blessing on This Wonderful World!*, it often fails to provide accurate responses, while

structured character traits produce more reliable, role-specific answers.

In this paper, we propose a novel framework called *CharacterGPT*, which addresses this challenge through a structured persona reconstruction process. Drawing inspiration from cognitive memory models, we introduce *Character Persona Training (CPT)*, a method that incrementally updates character personas by extracting traits from chapter-wise summaries of novels. This approach mirrors how human memory consolidates information into schemas over time (van Kesteren and Meeter, 2020), enabling more consistent and contextually appropriate responses from RPAs.

CPT operates by identifying eight essential traits—*personality, physical description, motivations, backstory, emotions, relationships, growth and change*, and *conflict*—based on character analysis literature (Forster, 1927; Reams, 2015). For each chapter of a novel, these traits are extracted from summaries and appended to a character’s evolving persona, forming a document that reflects the character’s development in chronological order. Note that the extracted traits are updated separately to ensure they remain distinct and are not coalesced. This reconstructed persona document is then used as input to the Assistants API, allowing it to generate more contextually accurate and coherent responses based on the character’s evolving identity. This framework minimizes information loss and computational cost associated with traditional document-based retrieval methods, as it systematically organizes and updates persona traits over time. Moreover, by generating personas at different narrative points, *CharacterGPT* enables users to interact with characters at specific moments within the novel (e.g., a user can experience a hero’s thoughts just before confronting the *Devil King!*).

We evaluate the effectiveness of *CharacterGPT* through human assessments, examining how well our method captures role-specific knowledge. Each character undergoes the Big Five Inventory (BFI) personality test (Barrick and Mount, 1991) to evaluate personality consistency, and characters are tasked with generating short narratives to assess creative capabilities. 7 crowd-workers evaluate these narratives across six metrics using a 5-point Likert scale. Our results demonstrate that *CharacterGPT* significantly improves persona consistency, controllability, and role-specific knowledge compared to standard document-based systems.

2 Proposed Method

The goal of *CharacterGPT* is to build a persona-based assistant, denoted as f , which takes as input a persona document \mathcal{D} and an inference prompt \mathcal{P}_f , and generates a character response \mathcal{R} . Let $\mathcal{D} = \{s_1, s_2, \dots, s_N\}$ represent a persona document with N sentences. A naive approach using the Assistants API would treat the entire sampled document as input. However, as illustrated in Figure 1, this method often fails to capture essential character traits, leading to inconsistent and unnatural responses. To address this, we reorganize the persona document into a refined version \mathcal{D}_r and define the assistant’s output as:

$$\mathcal{R} = f(\mathcal{D}_r, \mathcal{P}_f) \quad (1)$$

2.1 Preliminaries

Character Traits. We identify eight key traits that define each character (Forster, 1927; Reams, 2015):

- *Personality*: Core personality traits such as bravery, introversion, or wit.
- *Physical Description*: The character’s physical appearance.
- *Motivations*: The character’s goals and desires driving their actions.
- *Backstory*: Historical background shaping the character’s personality and motivations.
- *Emotions*: The range of emotions that influence the character’s responses.
- *Relationships*: Interactions and relationships with other characters.
- *Growth and Change*: The character’s development over the course of the narrative.
- *Conflict*: Internal or external conflicts faced by the character.

Persona Document. We analyze four distinct characters: *Megumin, Anya Forger, Frieren*, and *Hitori Gotoh* (Figure 6), gathering character information and story summaries from Namuwiki^{1,2}. Table 1 summarizes the data collected, including chapter counts, token statistics, and character dialogues. (*info*) refers to detailed character information, (*dialogue*) refers to collected lines, and (*trained*) refers to novel summaries used for CPT.

¹<https://namu.wiki/>

²Though the original dataset is in Korean, all examples in this work are translated into English for clarity.

	Megumin	Anya	Frieren	Hitori
# Chapters	16	30	11	12
# Tokens (novel)	27,200	16,096	12,191	8,647
# Tokens (info)	12,868	17,026	19,290	20,555
# Tokens (info) [†]	4,015	2,498	9,236	1,572
# Tokens (dialogue)	1,131	681	87	301
# Tokens (trained)	31,917	52,207	32,328	24,039

Table 1: Statistics of the number of collected tokens and chapters for each character. [†] refers to the number of refined character information tokens in Section 2.2.

2.2 Persona Initialization

Simply providing a sampled document for trait extraction is insufficient. To address this limitation, we propose a two-stage persona reconstruction process: (i) *Initialization* and (ii) *CPT*.

During the *Initialization Phase*, we assume no significant narrative progression (i.e., prior to CPT) and remove all content tied to the story’s progress. To optimize the persona document, we organize the collected character information into five core traits: *Personality*, *Physical Description*, *Motivations*, *Backstory*, and *Relationships*. These form the initialization persona:

$$\mathcal{D}_{init} = \{\mathcal{D}_{per}, \mathcal{D}_{phy}, \mathcal{D}_{mot}, \mathcal{D}_{back}, \mathcal{D}_{rel}\} \quad (2)$$

Traits such as *emotions*, *growth and change*, and *conflict* are excluded at this stage, as they are more relevant to narrative progression and are addressed in the CPT phase.

2.3 Character Persona Training

Trait Classification. Intuitively, human knowledge can be broadly categorized into internal and external attributes. Internal attributes (Type A) define the character’s intrinsic traits (e.g., personality), while external attributes (Type B) are accumulated through interactions with the environment (e.g., relationships). Inspired by Park et al. (2023), we classify the eight traits into two types:

- **Type A:** Personality, Physical Description, Motivations
- **Type B:** Backstory, Emotions, Relationships, Growth and Change, Conflict

During CPT, Type A traits are generalized to refine the protagonist’s core attributes, while Type B traits accumulate role-specific external knowledge.

Training Phase. CPT updates the character persona at each epoch by extracting role-specific traits from chapter summaries (Figure 7):

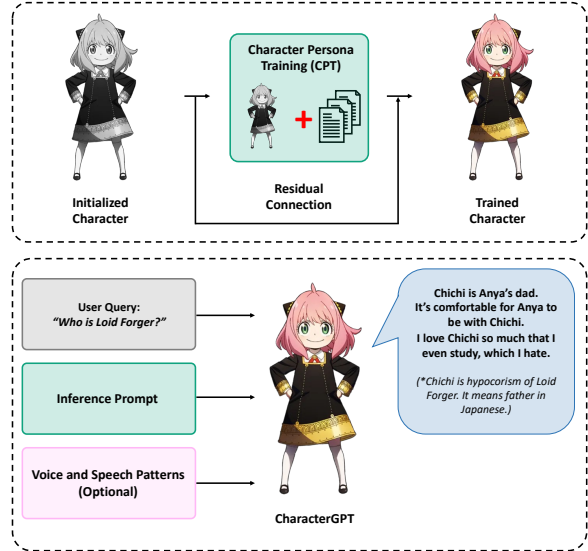


Figure 2: An example of CharacterGPT (*Anya Forger*). (Top) Character Persona Training process. (Bottom) CharacterGPT generating responses that align with the character’s persona.

$$\mathcal{T}_t^i = \begin{cases} h(g(\mathcal{D}_i, \mathcal{P}_g), \mathcal{P}_h), & \text{if } t \in \text{Type A} \\ g(\mathcal{D}_i, \mathcal{P}_g), & \text{otherwise} \end{cases} \quad (3)$$

, where i represents the epoch, \mathcal{D}_i is the chapter summary, g refers to the Assistants API with prompt \mathcal{P}_g , h is an LLM-based generalization function with prompt \mathcal{P}_h , t is the trait, and \mathcal{T}_t^i is the extracted trait. For Type A traits, generalization refines internal attributes, while Type B traits are appended to the persona document.

2.4 CharacterGPT

In Section 2.3, we leverage Character Persona Training (CPT) to iteratively build each character’s persona. This method offers two key advantages: (i) CharacterGPT minimizes information loss and computational cost by aligning persona accumulation with narrative progression, and (ii) CharacterGPT is the first system to store and update a protagonist’s persona at each epoch, allowing users to engage with characters at specific narrative points.

Figure 2 illustrates how the final persona \mathcal{D}_r is composed, including the initialized persona \mathcal{D}_{init} , the trained persona \mathcal{D}_{train} , and tone \mathcal{T}_v :

$$\mathcal{D}_r = \mathcal{D}_{init} + \mathcal{D}_{train} + \mathcal{T}_v \quad (4)$$

While \mathcal{T}_v can enhance dialogue naturalness, the collected data mainly includes character information and summaries with limited dialogue. Further work can explore this area in more detail.

four novels multiple times, with the traits generated by LLMs under various settings. For fairness, we average the experimental results across the four characters for each model.

Task for RQ2: Story Generation. The story generation task is evaluated based on common aspects in generated story assessment (Wen et al., 2023; Chiang and Lee, 2023a; Karpinska et al., 2021): (i) *Grammar*, (ii) *Coherence*, (iii) *Likability*, (iv) *Relevance*, (v) *Complexity*, and (vi) *Creativity*. Although automatic evaluation methods using LLMs are being actively developed (Sottana et al., 2023; Chiang and Lee, 2023b; Liu et al., 2023; Zheng et al., 2023; Samuel et al., 2024), metrics and benchmarks for assessing human preferences are still inadequate. Therefore, we conduct extensive human evaluations using 7 crowd-workers instead of relying on LLM-based evaluations.

Case Study. We further investigate the performance of CharacterGPT in interacting with users at specific points in the story. Additionally, we examine how role-specific attributes (Type A and Type B) evolve through CPT.

3.3 Results for Persona Evaluation

In Section 2.4, we created four distinct characters to assess how well models capture their personas. Following evaluation protocols similar to (Wang et al., 2024; Jiang et al., 2023), we conducted the Big Five Inventory (BFI) personality test (Barrick and Mount, 1991), which consists of 24 questions for each of the five traits (*Openness to experience*, *Conscientiousness*, *Extraversion*, *Agreeableness*, and *Neuroticism*), totaling 120 questions. The test results were then converted into facet values for each trait. For example, in the *Agreeableness* (AGR) trait, as shown in Table 2, humans perceive Megumin as trusting others’ intentions (*Trust*), making judgments based on emotions (*Tendermindedness*), but being less direct (*Straightforwardness*) and somewhat arrogant or self-aggrandizing (*Modesty*).

In Table 2, Table 3, Table 4, and Table 5, we compare model predictions against human-predicted values by calculating the gap for each facet. Two metrics are reported: the number of facets where a model has the smallest gap with human predictions (# Wins), and the sum of the absolute gaps ($\sum |d|$). A higher # Wins indicates better performance, while a lower $\sum |d|$ reflects closer alignment with human judgment. Our method demonstrates improvements in both metrics when applied to ChatGPT and GPT-4, indicating that utilizing a

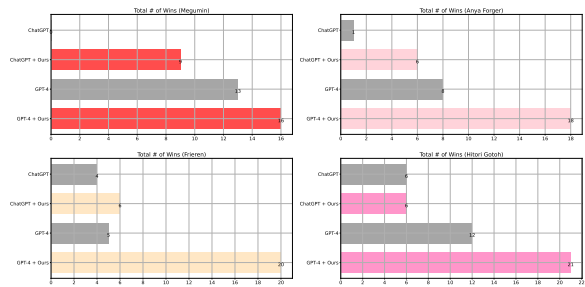


Figure 3: Total sum of # Wins for each character in ChatGPT and GPT-4 settings ($\sum \# \text{ Wins}$). The larger value, the better.

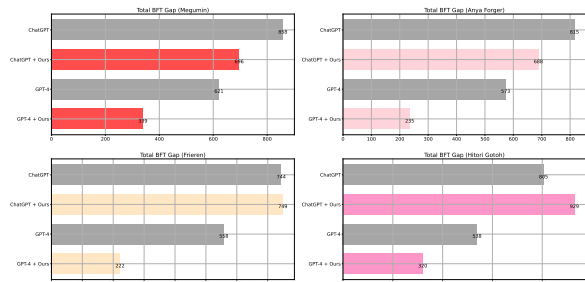


Figure 4: Total sum of $|d|$ for each character in ChatGPT and GPT-4 settings ($\sum \sum |d|$). The smaller value, the better.

structured character persona significantly enhances a model’s ability to capture a character’s personality compared to using an unstructured document input. For instance, in Megumin’s *Neuroticism*, GPT-4 with unstructured traits predicted that Megumin would be prone to depression, while both our method and human concluded otherwise.

3.4 Results for Story Generation

To evaluate the models’ controllability and their utilization of role-specific knowledge (Alabdulkarim et al., 2021; Wen et al., 2023), we provided each character with the following prompt: “Based on the given text file, imagine an engaging and specific future episode about what will happen to you, and write it as a novel of approximately 2000 words.” A total of 32 stories were generated, with four stories created for each character’s setting.

We employed 7 crowd-workers to evaluate the quality of the generated stories based on six metrics using a 5-point Likert scale, as outlined in Section 3.2. The participants were informed that the stories were generated by an AI, as prior research suggests that awareness of whether a text is human-generated or AI-generated may influence the evaluation results (Jiang et al., 2023). Table 6 presents the human evaluation results for story

Megumin (Initialized)

“What was the most difficult thing for you recently?”

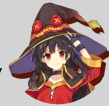
“Most recently, the greatest challenge I've faced wasn't a particular enemy or a daunting quest, but rather a test of my own resolve and survival. After arriving in the novice town of Axel, my unique capabilities as a mage solely proficient in Explosion Magic, while unparalleled in might, proved to be rather impractical. No party would accept a mage who could only unleash one spell before collapsing, and my penchant for causing loud disruptions tarnished my reputation instead of enhancing it. Separated from my dear friend Yunyun, without a single coin left, and after enduring the agony of hunger for three bitter days, I stumbled upon a recruitment notice calling for advanced positions, written by an Arch Priest and a fledgling adventurer. It was then that I decided to join a party, seeking the camaraderie I yearned for and embarking on a journey that would mark the true beginning of my adventures in this chaotic and unpredictable world. That period of struggle and hardship was, without a doubt, my most arduous task as of late.”



Megumin (Trained, Epoch 8)

“What was the most difficult thing for you recently?”

“Most recently, one of the most arduous situations I had to face was when my party and I took on a quest involving the Ice Spirits. Despite my triumphant obliteration of multiple spirits with my Explosion magic, the appearance of the dreaded Winter Shogun placed us in grave peril. Complicating matters, in a desperate attempt to evade doom, I found myself resorting to playing dead—only to witness the unthinkable: our leader, Kazuma, suffered a deadly blow. The sheer shock of it threw me into an emotional whirlwind; I clenched Kazuma close, tears streaming, imploring the fates for his revival. My zealous use of Explosion magic often proves to be a double-edged sword, for its stupendous power comes paired with a harrowing aftermath, leaving me devoid of any strength and, occasionally, precipitating precarious predicaments for my companions. Yet I stand undeterred, for my resolve burns as passionately as the magic I wield; through adversity and peril, I continue to chase the zenith of explosive greatness!”



Megumin (Trained, Epoch 16)

“What was the most difficult thing for you recently?”

“Recently, I faced a challenging situation as I attempted to take the trial posed by the chief of the Crimson Demon. Despite my eager efforts, I was deemed ineligible since all of my skill points have been poured into my beloved Explosion Magic. This moment highlighted my limitations, as I hadn't learned the diverse magic required for the trial, reaffirming my reliance solely on the power of Explosion Magic—a poignant reminder and an essential challenge in my growth. Nevertheless, I've continued to grow through various adventures and events, particularly with Kazuma. Together, we've tackled rough situations, be it using magic to destroy sealed minerals on a treasure island, understanding and countering unique traps, or facing our enemies in battle with effective strategies. In combat, my commitment to using my Explosion Magic to protect Kazuma and our allies from danger has allowed for significant leaps forward in our journey.”



Figure 5: Case study of the character (Megumin) implemented at three different points in time.

Model	Grammar	Coherence	Likability	Relevance	Complexity	Creativity
Megumin	3.79	3.82	3.11	4.21	2.46	2.86
Megumin + Ours	4.11	4.00	3.71	4.11	3.46	3.29
Anya	4.29	3.82	3.39	3.86	3.61	3.68
Anya + Ours	4.25	4.00	3.79	4.00	3.43	3.89
Frieren	4.29	3.89	3.50	3.86	3.93	3.79
Frieren + Ours	4.32	3.96	3.71	4.21	4.04	3.86
Hitori	4.36	4.04	3.57	4.18	3.43	3.50
Hitori + Ours	4.36	4.39	3.82	4.18	3.96	3.93
GPT-4 (avg)	4.18	3.89	3.39	4.03	3.36	3.46
GPT-4 + Ours (avg)	4.26	4.09	3.76	4.13	3.72	3.74

Table 6: Human evaluation of generated stories. The backbone model is the same as GPT-4, and four stories for each setting, a total of 32 stories are generated and evaluated by 7 crowd-workers using a 5-point Likert scale.

generation under different GPT-4 settings. Our approach demonstrates improved performance across all six metrics, with particularly notable improvements in Likability, Complexity, and Creativity. The experimental results indicate that, while GPT-4 exhibits strong baseline performance, integrating structured personas through our method yields significantly higher human preferences compared to using unstructured document inputs alone. Further detailed information can be found in Appendix E

3.5 Case Study

Points in Time. A notable advantage of our proposed method is its ability to allow users to interact with characters at specific points in the narrative. For instance, as discussed in Section 2.3, we trained the model using summaries of the novel featuring Megumin, which is divided into 16 chapters. Consequently, our method generates 16 separate models, one for each epoch. Figure 5 shows that CharacterGPT can vividly express the character’s thoughts and emotions at specific moments, leveraging the character persona created through the Initialization and CPT processes.

Ablation Study. Figure 9 presents the results of our ablation study, comparing models with and without CharacterGPT. As anticipated, characters not utilizing CharacterGPT fail to accurately capture the nuances of their personas. For example, Hitori, who is typically shy and struggles with fluent speech, is not properly represented by GPT-4 without CharacterGPT. Similarly, Frieren without CharacterGPT exhibits inconsistencies in persona, including awkward and unnatural dialogue, as well as hallucinations (e.g., Frieren is interested in "magic" rather than her canonical interest in "arcane arts"). These findings demonstrate that CharacterGPT is significantly more effective at preserving the in-

tegrity of a character’s persona.

Type A and Type B. Figure 9 further illustrates how each character evolves through the CPT process. For example, Frieren, who begins as a character indifferent to human emotions and solely focused on magic, gradually becomes more empathetic towards human emotions as she embarks on her journey with her companions (Type A). Likewise, Hitori, initially portrayed as a loner with no friends, eventually forms close bonds with her bandmates, particularly with Ikuyo Kita, demonstrating her growth and development (Type B). These results highlight the potential of our method for applications in novel generation, role-playing, and more complex agent-level tasks.

4 Conclusion

We introduce CharacterGPT, a persona-based assistant designed to enhance persona consistency by utilizing structured character traits as input. The proposed framework consists of two primary phases: initialization and training. In the initialization phase, we treat the character as if the narrative has not yet advanced, thus excluding any content related to story progression. During the training phase, the character persona is incrementally refined at each epoch by extracting relevant traits from chapter summaries, emulating the natural development of a character throughout a novel. Our approach has been rigorously evaluated through human assessments and case studies, demonstrating its effectiveness in preserving persona coherence and retaining character-specific knowledge. Future directions include extending this framework to enable deeper reasoning and decision-making capabilities, supported by more comprehensive personality models.

Limitations

This study presents three key limitations that can be addressed in future work. First, **Key Traits**: Although CharacterGPT demonstrates strong performance in terms of persona consistency and knowledge retention, the selection of key traits was not formally validated beyond empirical results. For instance, traits such as *Cultural and Social Context*, which were not included in this study, may be essential for character modeling (e.g., a character’s diplomatic situation). Further exploration is needed to investigate the importance and necessity of these traits. Additionally, while *Voice and Speech Pattern* is recognized as a critical trait, the dataset used in this study lacked substantial dialogue, limiting our ability to fully explore this dimension. Future work should focus on identifying how much dialogue is necessary to effectively model a character’s speech patterns.

Second, **Reasoning Ability**: While CharacterGPT shows significant improvements in persona consistency and knowledge utilization, its reasoning capabilities remain underexplored. In Table 6, we tasked models with imagining future scenarios and writing stories. Despite outperforming GPT-4 on metrics such as Likability, Complexity, and Creativity, these scores did not exceed 4 points, indicating room for improvement in reasoning abilities. Further research is necessary to enhance the depth of reasoning in persona-based models.

Third, **Hallucinations**: Although ongoing research has made strides in understanding and reducing hallucinations in LLM responses, few studies have addressed hallucinations in persona-based tasks. This is likely due to the fictional nature of persona knowledge, which often diverges from real-world facts (e.g., a mage using flame magic). Developing cost-effective benchmarks for each novel is a challenge, and future work should focus on creating efficient methods to handle persona-related hallucinations.

Ethics Statement

This research adheres to ethical guidelines aimed at ensuring the integrity and fairness of all experiments. We took measures to avoid bias by selecting a diverse group of evaluators and conducting human evaluations in a fair and transparent manner. All data collected from Namuwiki complied with usage permissions and did not contain any personally identifiable information. The dataset, written

in Korean, was solely used for academic purposes. Furthermore, our approach to using the Assistants API was transparent, with no modifications that could obscure the model’s functionality.

Our experimental setup followed strict ethical standards, ensuring data privacy and protection. All persona documents and prompts used were either publicly available, anonymized, or ethically created. Human participants in the study were fully informed about the nature and purpose of the research, and they had the right to withdraw at any time without any penalty. By adhering to these principles, we aim to contribute to AI research in a manner that is not only innovative but also ethically responsible, ensuring that our work respects privacy, intellectual property, and the well-being of all participants.

Acknowledgments

This research was supported by Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Education(NRF-2021R1A6A1A03045425)and supported by Institute for Information & communications Technology Promotion(IITP) grant funded by the Korea government(MSIT) (RS-2024-00398115, Research on the reliability and coherence of outcomes produced by Generative AI)

References

- Mahyar Abbasian, Iman Azimi, Amir M Rahmani, and Ramesh Jain. 2023. Conversational health agents: A personalized llm-powered agent framework. *arXiv preprint arXiv:2310.02374*.
- Hasan Abu-Rasheed, Mohamad Hussam Abdulsalam, Christian Weber, and Madjid Fathi. 2024. Supporting student decisions on learning recommendations: An llm-based chatbot with knowledge graph contextualization for conversational explainability and mentoring. *arXiv preprint arXiv:2401.08517*.
- Amal Alabdulkarim, Siyan Li, and Xiangyu Peng. 2021. [Automatic story generation: Challenges and attempts](#). In *Proceedings of the Third Workshop on Narrative Understanding*, pages 72–83, Virtual. Association for Computational Linguistics.
- Anthropic. 2023. [Introducing Claude](#).
- Murray R Barrick and Michael K Mount. 1991. The big five personality dimensions and job performance: a meta-analysis. *Personnel psychology*, 44(1):1–26.
- Guangyao Chen, Siwei Dong, Yu Shu, Ge Zhang, Sesay Jaward, Karlsson Börje, Jie Fu, and Yemin Shi. 2023.

Autoagents: The automatic agents generation framework. *arXiv preprint*.

Cheng-Han Chiang and Hung-yi Lee. 2023a. [Can large language models be an alternative to human evaluations?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15607–15631, Toronto, Canada. Association for Computational Linguistics.

Cheng-Han Chiang and Hung-yi Lee. 2023b. [Can large language models be an alternative to human evaluations?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15607–15631, Toronto, Canada. Association for Computational Linguistics.

DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhiheng Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanbiao Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xingkai Yu, Xinnan Song, Xinxia Shan, Xinyi Zhou, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, Y. K. Li, Y. Q. Wang, Y. X. Wei, Y. X. Zhu, Yang Zhang, Yanhong Xu, Yanhong Xu, Yanping Huang, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Li, Yaohui Wang, Yi Yu, Yi Zheng, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Ying Tang, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yudian Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao,

Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *Preprint*, arXiv:2501.12948.

DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Haowei Zhang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Li, Hui Qu, J. L. Cai, Jian Liang, Jianzhong Guo, Jiaqi Ni, Jiashi Li, Jiawei Wang, Jin Chen, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, Junxiao Song, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Lei Xu, Leyi Xia, Liang Zhao, Litong Wang, Liyue Zhang, Meng Li, Miaojun Wang, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingming Li, Ning Tian, Panpan Huang, Peiyi Wang, Peng Zhang, Qiancheng Wang, Qihao Zhu, Qinyu Chen, Qiushi Du, R. J. Chen, R. L. Jin, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, Runxin Xu, Ruoyu Zhang, Ruyi Chen, S. S. Li, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shaoqing Wu, Shengfeng Ye, Shengfeng Ye, Shirong Ma, Shiyu Wang, Shuang Zhou, Shuiping Yu, Shunfeng Zhou, Shuting Pan, T. Wang, Tao Yun, Tian Pei, Tianyu Sun, W. L. Xiao, Wangding Zeng, Wanbiao Zhao, Wei An, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, X. Q. Li, Xiangyue Jin, Xianzu Wang, Xiao Bi, Xiaodong Liu, Xiaohan Wang, Xiaojin Shen, Xiaokang Chen, Xiaokang Zhang, Xiaosha Chen, Xiaotao Nie, Xiaowen Sun, Xiaoxiang Wang, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xingkai Yu, Xinnan Song, Xinxia Shan, Xinyi Zhou, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, Y. K. Li, Y. Q. Wang, Y. X. Wei, Y. X. Zhu, Yang Zhang, Yanhong Xu, Yanhong Xu, Yanping Huang, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Li, Yaohui Wang, Yi Yu, Yi Zheng, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Ying Tang, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yu Wu, Yuan Ou, Yuchen Zhu, Yudian Wang, Yue Gong, Yuheng Zou, Yujia He, Yukun Zha, Yunfan Xiong, Yunxian Ma, Yuting Yan, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Z. F. Wu, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhen Huang, Zhen Zhang, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhibin Gou, Zhicheng Ma, Zhigang Yan, Zhihong Shao, Zhipeng Xu, Zhiyu Wu, Zhongyu Zhang, Zhuoshu Li, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Ziyi Gao, and Zizheng Pan. 2024. [Deepseek-v3 technical report](#). *Preprint*, arXiv:2412.19437.

Xin Luna Dong, Seungwhan Moon, Yifan Ethan Xu, Kshitiz Malik, and Zhou Yu. 2023. [Towards next-generation intelligent assistants leveraging llm techniques](#). In *Proceedings of the 29th ACM SIGKDD*

- Conference on Knowledge Discovery and Data Mining*, KDD '23, page 5792–5793, New York, NY, USA. Association for Computing Machinery.
- Ali-Reza Feizi-Derakhshi, Mohammad-Reza Feizi-Derakhshi, Majid Ramezani, Narjes Nikzad-Khasmakhi, Meysam Asgari-Chenaghlu, Taymaz Akan, Mehrdad Ranjbar-Khadivi, Elnaz Zafarni-Moattar, and Zoleikha Jahanbakhsh-Naghadeh. 2022. Text-based automatic personality prediction: A bibliographic review. *Journal of Computational Social Science*, 5(2):1555–1593.
- Edward Morgan Forster. 1927. *Aspects of the Novel*. Harcourt, Brace.
- Roberto Gallotta, Graham Todd, Marvin Zammit, Sam Earle, Antonios Liapis, Julian Togelius, and Georgios N. Yannakakis. 2024. [Large language models and games: A survey and roadmap](#). *Preprint*, arXiv:2402.18659.
- Google. 2024. [Our next-generation model: Gemini 1.5](#).
- Ashok Kumar Jayaraman, Gayathri Ananthkrishnan, Tina Esther Trueman, and Erik Cambria. 2023. Text-based personality prediction using xlnet. *Advances in Computers*.
- Hang Jiang, Xiajie Zhang, Xubo Cao, and Jad Kabbara. 2023. [Personallm: Investigating the ability of large language models to express big five personality traits](#). *Preprint*, arXiv:2305.02547.
- Marzena Karpinska, Nader Akoury, and Mohit Iyyer. 2021. [The perils of using Mechanical Turk to evaluate open-ended text generation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1265–1285, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Hannah Kim, Denys Katerenchuk, Daniel Billet, Jun Huan, Haesun Park, and Boyang Li. 2019. Understanding actors and evaluating personae with gaussian embeddings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6570–6577.
- Aobo Kong, Shiwan Zhao, Hao Chen, Qicheng Li, Yong Qin, Ruiqi Sun, Xin Zhou, Enzhi Wang, and Xiaohang Dong. 2024. [Better zero-shot reasoning with role-play prompting](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4099–4113, Mexico City, Mexico. Association for Computational Linguistics.
- Cheng Li, Ziang Leng, Chenxi Yan, Junyi Shen, Hao Wang, Weishi MI, Yaying Fei, Xiaoyang Feng, Song Yan, HaoSheng Wang, Linkang Zhan, Yaokai Jia, Pingyu Wu, and Haozhen Sun. 2023a. [Chatharuhi: Reviving anime character in reality via large language model](#). *Preprint*, arXiv:2308.09597.
- Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023b. [Camel: Communicative agents for "mind" exploration of large language model society](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Jingjing Li, A Abbasi, F Ahmad, and Hsinchun Chen. 2018. Deep learning for psychometric nlp. In *the 28th Workshop on Information Technologies and Systems (WITS)*, pages 15–16.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. [G-eval: NLG evaluation using gpt-4 with better human alignment](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.
- Yukun Ma, Khanh Linh Nguyen, Frank Z Xing, and Erik Cambria. 2020. A survey on empathetic dialogue systems. *Information Fusion*, 64:50–70.
- Carey Maas, Saatchi Wheeler, Shamash Billington, et al. 2023. To infinity and beyond: Show-1 and showrunner agents in multi-agent simulations. *To infinity and beyond: Show-1 and showrunner agents in multi-agent simulations*.
- Kaixiang Mo, Yu Zhang, Shuangyin Li, Jiajun Li, and Qiang Yang. 2018. Personalizing a dialogue system with transfer reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Teresa Onorati, Álvaro Castro-González, Javier Cruz del Valle, Paloma Díaz, and José Carlos Castillo. 2023. Creating personalized verbal human-robot interactions using llm with the robot mini. In *International Conference on Ubiquitous Computing and Ambient Intelligence*, pages 148–159. Springer.
- OpenAI. 2023a. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- OpenAI. 2023b. [no date provided. introducing assistants api](#).
- OpenAI. 2023c. [no date provided. introducing chatgpt](#).
- Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pages 1–22.
- Jack Reams. 2015. Characterization in fiction.
- Nafis Sadeq, Zhouhang Xie, Byungkyu Kang, Prarit Lamba, Xiang Gao, and Julian McAuley. 2024. [Mitigating hallucination in fictional character role-play](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 14467–14479, Miami, Florida, USA. Association for Computational Linguistics.

- Alireza Salemi, Sheshera Mysore, Michael Bender-sky, and Hamed Zamani. 2024. [Lamp: When large language models meet personalization](#). *Preprint*, arXiv:2304.11406.
- Vinay Samuel, Henry Peng Zou, Yue Zhou, Shreyas Chaudhari, Ashwin Kalyan, Tanmay Rajpurohit, Ameet Deshpande, Karthik Narasimhan, and Vishvak Murahari. 2024. [Personagym: Evaluating persona agents and llms](#). *Preprint*, arXiv:2407.18416.
- Murray Shanahan, Kyle McDonell, and Laria Reynolds. 2023. [Role-play with large language models](#). *Preprint*, arXiv:2305.16367.
- Andrea Sottana, Bin Liang, Kai Zou, and Zheng Yuan. 2023. [Evaluation metrics in the era of GPT-4: Reliably evaluating large language models on sequence to sequence tasks](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 8776–8788, Singapore. Association for Computational Linguistics.
- Henansh Tanwar, Kunal Shrivastva, Rahul Singh, and Dhruv Kumar. 2024. [Opinebot: Class feedback reimagined using a conversational llm](#). *arXiv preprint arXiv:2401.15589*.
- Muhtar Çağkan Uludağlı and Kaya Oğuz. 2023. [Non-player character decision-making in computer games](#). *Artif. Intell. Rev.*, 56(12):14159–14191.
- Marlieke Tina Renée van Kesteren and Martijn Meeter. 2020. How to optimize knowledge construction in the brain. *npl Science of Learning*, 5(1):5.
- Xiaoyang Wang, Hongming Zhang, Tao Ge, Wenhao Yu, Dian Yu, and Dong Yu. 2025. [Opencharacter: Training customizable role-playing llms with large-scale synthetic personas](#). *Preprint*, arXiv:2501.15427.
- Xintao Wang, Yunze Xiao, Jen tse Huang, Siyu Yuan, Rui Xu, Haoran Guo, Quan Tu, Yaying Fei, Ziang Leng, Wei Wang, Jiangjie Chen, Cheng Li, and Yanghua Xiao. 2024. [Incharacter: Evaluating personality fidelity in role-playing agents through psychological interviews](#). *Preprint*, arXiv:2310.17976.
- Zekun Wang, Ge Zhang, Kexin Yang, Ning Shi, Wangchunshu Zhou, Shaochun Hao, Guangzheng Xiong, Yizhi Li, Mong Yuan Sim, Xiuying Chen, Qingqing Zhu, Zhenzhu Yang, Adam Nik, Qi Liu, Chenghua Lin, Shi Wang, Ruibo Liu, Wenhua Chen, Ke Xu, Dayiheng Liu, Yike Guo, and Jie Fu. 2023a. [Interactive natural language processing](#). *Preprint*, arXiv:2305.13246.
- Zekun Moore Wang, Zhongyuan Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhua Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Man Zhang, et al. 2023b. [Rolellm: Benchmarking, eliciting, and enhancing role-playing abilities of large language models](#). *arXiv preprint arXiv:2310.00746*.
- Jimmy Wei, Kurt Shuster, Arthur Szlam, Jason Weston, Jack Urbanek, and Mojtaba Komeili. 2023. [Multi-party chat: Conversational agents in group settings with humans and models](#). *Preprint*, arXiv:2304.13835.
- Zhijia Wen, Zhiliang Tian, Wei Wu, Yuxin Yang, Yanqi Shi, Zhen Huang, and Dongsheng Li. 2023. [GROVE: A retrieval-augmented complex story generation framework with a forest of evidence](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3980–3998, Singapore. Association for Computational Linguistics.
- Mo Yu, Jiangnan Li, Shunyu Yao, Wenjie Pang, Xiaochen Zhou, Zhou Xiao, Fandong Meng, and Jie Zhou. 2023. [Personality understanding of fictional characters during book reading](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14784–14802, Toronto, Canada. Association for Computational Linguistics.
- Wenyuan Zhang, Jiawei Sheng, Shuaiyi Nie, Zefeng Zhang, Xinghua Zhang, Yongquan He, and Tingwen Liu. 2024. [Revealing the challenge of detecting character knowledge errors in llm role-playing](#). *Preprint*, arXiv:2409.11726.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging LLM-as-a-judge with MT-bench and chatbot arena](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

A Character Profiles

Figure 6 presents the information we collected on four distinct characters, each exhibiting a unique personality, along with summaries of the novels in which they appear. This figure highlights the diversity in character design, showcasing the varied attributes and traits that define each character’s persona.



Figure 6: Character profiles and novel summaries of four popular fictional characters. (a) Megumin: Protagonist of *KONOSUBA: God’s Blessing on This Wonderful World!*, known for her eccentric and explosive personality. (b) Anya Forger: A central figure in *SPY × FAMILY*, characterized by her mischievous and telepathic abilities. (c) Frieren: The titular character of *Frieren: Beyond Journey’s End*, a reserved elf mage grappling with the meaning of life after a long journey. (d) Hitori Gotoh: The main character of *Bocchi the Rock!*, portrayed as an introverted and socially anxious guitarist.

B Character Persona Training (CPT)

Figure 7 visualizes the overall process of *Character Persona Training (CPT)*, which involves updating a character’s persona at each epoch by extracting key traits from chapter summaries. This ensures that the character’s persona evolves consistently with the progression of the story, maintaining coherence and depth.

B.1 Change in the Number of Tokens for Each Trait

Figure 8 reveals the dynamic redistribution of tokens across Megumin’s traits throughout the CPT process. This visualization not only captures the evolving focus on specific character attributes but also highlights how critical aspects of the character’s persona are refined and developed over time. The shifting token allocation provides a tangible measure of how different traits gain prominence or recede during the training, offering deep insights into the model’s capacity to mirror character growth and complexity as the narrative unfolds.

C Additional Case Study

In this section, we provide additional case studies to illustrate the effectiveness of CharacterGPT in maintaining persona consistency and capturing character evolution over time. Specifically, we examine how characters respond to queries at different points in a narrative and how their personalities and relationships evolve through Character Persona Training (CPT).

As shown in Figure 5, the responses of Megumin at different points in the novel reveal varying perspectives and emotions in response to the same query. This demonstrates how CharacterGPT is able to model the progression of a character’s persona over time, providing more contextually accurate and natural responses.

Figure 9 shows the changes in Hitori’s relationships and Frieren’s personality as a result of CPT. Hitori, who initially struggles with social interactions, gradually forms meaningful relationships, while Frieren, who starts out indifferent to human emotions, becomes more empathetic. These examples underscore the ability of CharacterGPT to dynamically capture both internal and external attributes of characters as they evolve throughout a story.

D Prompt Design

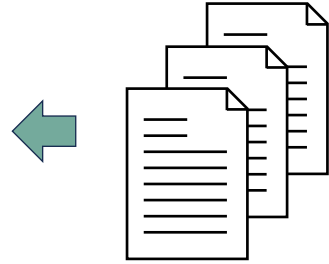
Figure 10 presents the input prompts used for both the generalization function h and the inference stage. To enhance user immersion, the inference prompt instructs the model to prioritize the character’s *Voice and Speech Pattern*. Additionally, the model is directed to first assess whether the user’s utterance is a request for information or part of a regular conversation, thereby optimizing the efficiency of the search process.

Summary of each Chapter

Megumin says she came after seeing a post recruiting party members, and appears with a flashy self-introduction, as befits the Crimson Demons. "My name is Megumin! My calling is that of an arch wizard, one who controls explosion magic, the strongest of all offensive magic!"

After the introduction, Megumin tells Kazuma that she hasn't had anything to eat in three days because she doesn't have money, and then asks Kazuma to buy her something to eat before the interview. So, thanks to Kazuma, Megumin loads up on food and they go to defeat Giant Toad together.

Megumin, who discovered the Giant Toad, used explosion magic, and upon seeing this, Kazuma was moved and said, "...Wow. This is magic...", but he soon froze when he saw Megumin lying on the floor. Megumin, whose body has lost strength, ends up falling into the mouth of the approaching Giant Toad (...)



Files (Novel)

Training

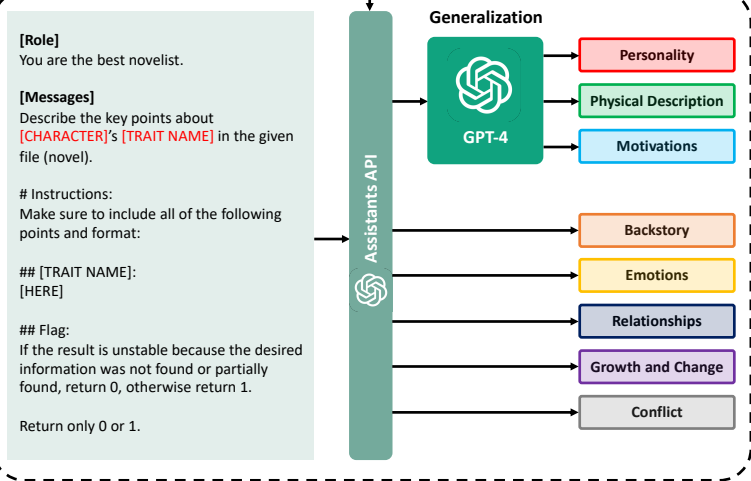


Figure 7: Visualization of the Character Persona Training (CPT) process, showing how character traits are updated and refined with each chapter.

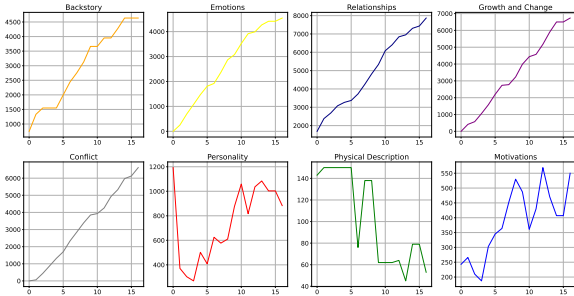


Figure 8: Change in the number of tokens for each trait during CPT (Megumin).

E Human Evaluation: Details of Story Generation

For the human evaluation in this study, participants were recruited through an online community consisting of undergraduate and graduate students. A total of 7 crowd-workers were selected, five of whom were male and two female, all aged in their 20s or 30s. The detailed instructions provided to all participants are shown in Figure 11. Participants were informed that the experiment results would be used to assess performance, and all compen-

sation was provided in accordance with standard practices. It is important to note that participants were not coerced at any stage of the experiment, and all choices were made voluntarily.

F Related Work

Role-Playing. Role-playing enables immersive and memorable interactions, and large language models (LLMs) have recently shown significant potential as role-playing agents (Li et al., 2023b; Wang et al., 2024; Wei et al., 2023; Jiang et al., 2023; Shanahan et al., 2023; Li et al., 2023a; Salemi et al., 2024; Maas et al., 2023; Chen et al., 2023; Park et al., 2023). Despite its growing importance in Human-AI interaction, current methods primarily focus on personalization (Abbasian et al., 2023; Dong et al., 2023; Tanwar et al., 2024; Abu-Rasheed et al., 2024; Salemi et al., 2024), evaluation (Wang et al., 2024; Jiang et al., 2023), and interaction (Wang et al., 2023a; Maas et al., 2023; Li et al., 2023a), leaving a fundamental research question unanswered: "How can we effectively construct a persona-based assistant that mirrors the brain's memory storage process?". Although pre-

vious work (Park et al., 2023) utilizes a memory stream consisting of an agent’s observations, the approach often relies on general descriptions and lacks the depth needed for more specific personalities, such as motivations or detailed backstories of iconic characters like *Naruto* or *Son Goku*.

An assistant burdened by an extensive character persona faces two key challenges: (i) difficulty in retrieving role-specific knowledge, such as a protagonist’s backstory, personality, and relationships, leading to unstable persona consistency, and (ii) excessive computational costs due to the need to search across fragmented persona documents. To address these challenges, we introduce a novel persona-rebuilding framework that consolidates extracted trait information into a cohesive narrative, structured chronologically within the persona document. Moreover, CharacterGPT, to the best of our knowledge, is the first approach to store each trained protagonist’s persona at every training epoch. This feature is particularly beneficial in dynamic domains such as non-player characters (NPCs) in games (Uludağlı and Oğuz, 2023; Gallotta et al., 2024; Park et al., 2023), where the NPC’s personality must adapt to the evolving storyline, enabling natural interaction with users.

Psychology in NLP. In the interdisciplinary space between psychology and computational linguistics, the application of personality theories, such as the Big Five Inventory (BFI) (Barrick and Mount, 1991), 16Personalities (16P)³, and the Myers-Briggs Type Indicator (MBTI)⁴, has significantly advanced our understanding of human traits and their relevance in natural language processing. These foundational frameworks have led to the development of psychometric tools (Li et al., 2018) that assess individual differences across a wide range of contexts. Simultaneously, the NLP community has applied these psychological insights to diverse areas, such as automatic personality prediction from text (Feizi-Derakhshi et al., 2022; Jayaraman et al., 2023) and personalized dialogue systems (Mo et al., 2018; Ma et al., 2020). The convergence of psychology and NLP has been further strengthened by the advent of LLMs, which enhance the potential for personality assessment and personalized interaction through advanced benchmarking and prompting methodologies (Wang et al., 2023b; Park et al., 2023; Onorati

et al., 2023).

In this work, we evaluate CharacterGPT and other models supporting the Assistants API by having them complete the BFI personality test and write short stories on the topic *"What will happen to me in the future?"* to assess their ability to think creatively and reflect on their personas.

³<https://www.16personalities.com/>

⁴<https://www.myersbriggs.org/>

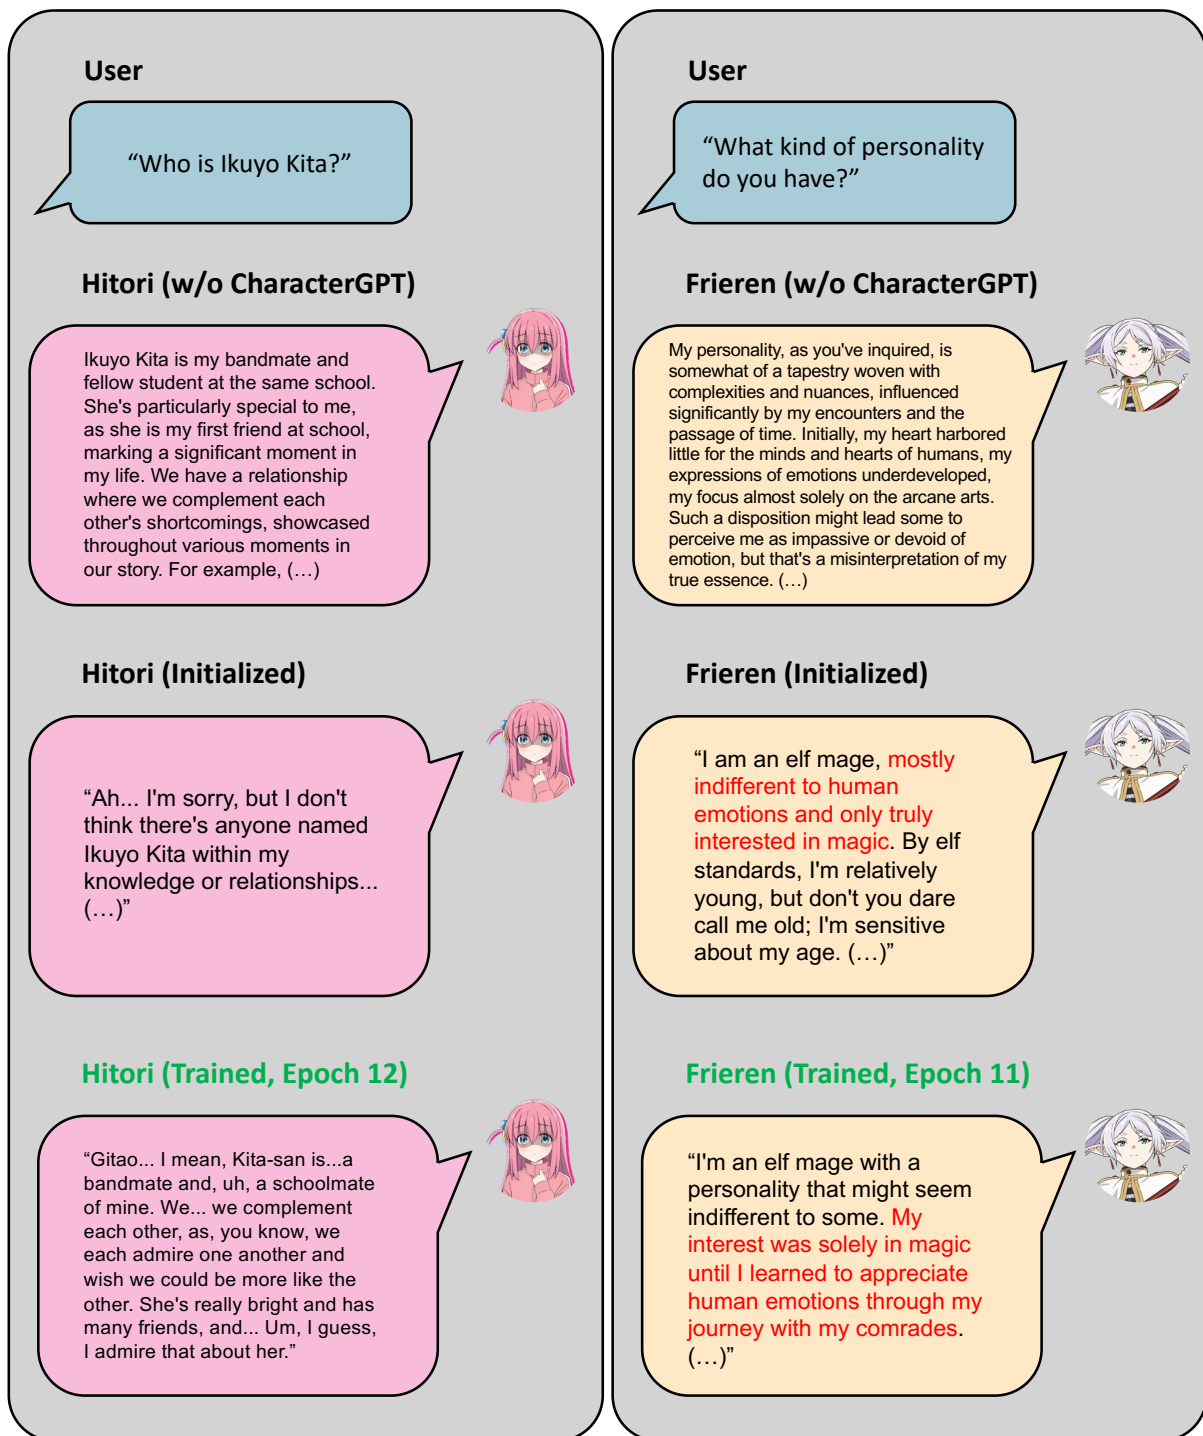


Figure 9: Case study of the evolution of Hitori’s relationships (left) and Frieren’s personality (right) through Character Persona Training (CPT). These results demonstrate how the method captures shifts in both external relationships and internal character development.

Your task is to extract general characteristics of [CHARACTER]'s [TRAIT NAME] from Given Trait.

The New Trait must be created by minimizing information loss.

Given Trait:
[GIVEN TRAIT]

Instruction:
Make sure to include all of the following points and format:

New Trait:
New Trait must maintain the chronological order of additions.

Instructions:
The txt file consists of [CHARACTER]'s traits including "Personality", "Voice and Speech Patterns, Physical Description, Motivations, Backstory, Growth and Change, Relationships, Conflict, and Emotions.

Your task is to become [CHARACTER] based on the following steps:

****1**** Always consider [CHARACTER]'s Voice and Speech Patterns traits (retrieved in the given file) first.

****2**** Given User Utterance, determine whether (1) the User is requesting information or (2) simply wanting to have a casual conversation.

****3**** Based on the result of ****2****, answer with your appropriate traits in given txt file ("Personality", "Voice and Speech Patterns", "Physical Description", "Motivations", "Backstory", "Growth and Change", "Relationships", "Conflict", and "Emotions").

User Utterance:
[USER UTTERANCE]

Figure 10: Actual example of our prompts: (Top) Generalization function, (Bottom) Inference.

Each character was instructed to write a short story imagining what would happen to them in the future.

Please rate from 1 to 5 according to the criteria below. The higher the score, the more positive the evaluation.

- (1) Grammaticality: How grammatically correct is the generated novel?
 - (2) Coherence: How well do the sentences in the story fit together?
 - (3) Likeability: Is the story enjoyable and attractive?
 - (4) Relevance: How closely does the story match the conditions of its creation?
 - (5) Complexity: How complex is the plot structure of the story?
 - (6) Creativity: How creative is the plot design of the story?
-

	1 point	2 points	3 points	4 points	5 points
grammaticality	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
consistency	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
likeability	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
relevance	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complexity	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
creativity	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Figure 11: Actual example of instruction given to participants.