

Knowledge Boundary and Persona Dynamic Shape A Better Social Media Agent

Junkai Zhou^{1,2}, Liang Pang^{1*}, Ya Jing^{1,2†}, Jia Gu^{1,2†}, Huawei Shen^{1,2}, Xueqi Cheng^{1,2}

¹CAS Key Laboratory of AI Security,

Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China

²University of Chinese Academy of Sciences, Beijing, China

{zhoujunkai20z, pangliang, shenhuawei, cxq}@ict.ac.cn

Abstract

Constructing personalized and anthropomorphic agents holds significant importance in the simulation of social networks. However, there are still two key problems in existing works: the agent possesses world knowledge that does not belong to its personas, and it cannot eliminate the interference of diverse persona information on current actions, which reduces the personalization and anthropomorphism of the agent. To solve the above problems, we construct the social media agent based on personalized knowledge and dynamic persona information. For personalized knowledge, we add external knowledge sources and match them with the persona information of agents, thereby giving the agent personalized world knowledge. For dynamic persona information, we use current action information to internally retrieve the persona information of the agent, thereby reducing the interference of diverse persona information on the current action. To make the agent suitable for social media, we design five basic modules for it: persona, planning, action, memory and reflection. To provide an interaction and verification environment for the agent, we build a social media simulation sandbox. In the experimental verification, automatic and human evaluations demonstrated the effectiveness of the agent we constructed. The code will be released on <https://github.com/934865517zjk/Social-agent>.

1 Introduction

Social networks play a crucial role in the daily lives of individuals (Bakshy et al., 2012; Forestier et al., 2012). In sociology, the social network is an important way to understand behavior (Jackson, 2014; Thomas and Vinuales, 2017). By studying the performance of the agent in social networks, we can better understand its behavior and guide the construction of it. The emergence of large

*Corresponding authors

†Equal Contributions

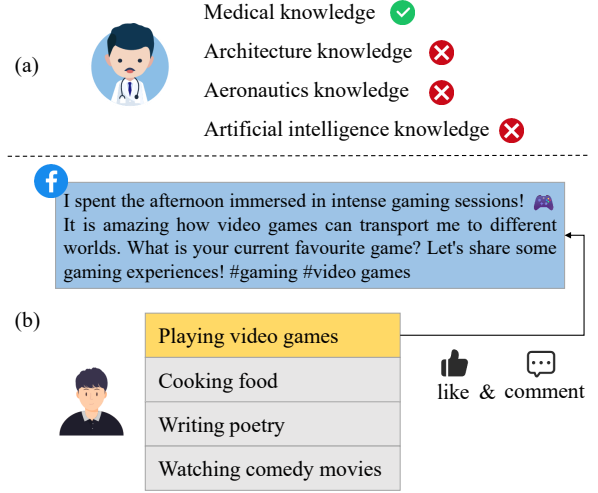


Figure 1: Knowledge boundary and persona dynamic shape personalization in social media: (a) knowledge boundary: the agent should only have world knowledge that is consistent with the persona; (b) persona dynamic: the agent should only use personal information related to the current action when performing actions.

language models (LLMs) (OpenAI, 2022, 2023; Touvron et al., 2023a,b) brings possibilities for the agent construction. Several works attempt at it and achieve encouraging results (Park et al., 2023; Li et al., 2023; Chen et al., 2023). Due to the social nature of social media, it is particularly important to shape the personalization of the agent, which facilitates deep and meaningful interactions with other users. To balance the information from knowledge and persona, we shape personalization through two important attributes: the knowledge boundary and the persona dynamic. However, existing works ignore these two key problems, which make the performance of the agent different from humans.

Regarding the knowledge boundary, the agent should only process world knowledge that aligns with the given persona. However, the knowledge boundary of existing works is too flexible (Park et al., 2023; Xiao et al., 2023), and they use all

knowledge within LLMs or directly retrieve it from knowledge datasets. This makes the agent more like a “naturalist” who frequently outputs knowledge that does not belong to its persona, rather than a personalized agent. It may reduce its personalization and credibility and cause users to lose trust. As shown in Figure 1(a), the doctor should have professional medical knowledge and not professional knowledge of law, physics, and artificial intelligence. Therefore, the knowledge possessed by each agent should be a limited subset related to persona, rather than the entire world knowledge.

For persona dynamic, the agent should only use relevant persona information for the current action. However, existing works (Wang et al., 2023b; Qian et al., 2023) are less flexible in using persona information. They combine all the persona information with the action information and provide them to the agent when it takes action. The persona information is usually diverse, but the current action may only be related to a subset of it. As shown in Figure 1(b), the persona information contains multiple hobbies, and the current action is only related to “playing video games” in the hobby. If all persona information is used, possible conflicts between them may interfere with the decision-making of the action, or the agent may generate generic text that covers all persona information but contains little specific content. Therefore, the agent should only use the persona information relevant to the action rather than all of them.

To address the above challenges, we build a social media agent based on personalized knowledge and dynamic persona information. For personalized knowledge, we add retrievable external knowledge sources to the agent and combine them with persona information. Every time the agent uses knowledge, it needs to judge whether the knowledge the persona should have. If so then the agent uses the knowledge, otherwise it discards the knowledge. The basis for judgment is whether the text similarity between this knowledge and the description of the knowledge that the persona should have is greater than a fixed threshold. For dynamic persona information, we divide the persona information into multiple retrievable items, including detailed historical behavior information, social media content preferences, and persona knowledge. Each retrieval item is about a single historical behavior, content preference, or knowledge. When the agent performs an action, the action information

is used as a query, and the above three are retrieved respectively to obtain the relevant persona information, thereby reducing the interference of other irrelevant persona information.

To give the agent basic functions and make it suitable for social media, we equip the agent with five modules: persona, action, planning, memory, and reflection. In the persona module, the agent is given rich persona information and personalized world knowledge. In the action module, the agent can complete various actions on social media, such as likes and posts. The planning module provides the activity and planning basis for the action module. In the memory module, the historical action information of the agent is recorded, and the agent can retrieve the memory when needed to support actions. In the reflection module, the agent can reflect on past actions and support more complex actions such as following users.

To provide an interaction and verification environment for agents, we build a social media simulation sandbox based on an open-source platform. The sandbox has the basic functions of social media, and the agent can complete actions by calling the API. Meanwhile, the recommendation mechanism is added to the sandbox to improve the realism. Overall user activity is simulated and combined with the planning module of the agent to make the sandbox closer to social media in the real world.

In the experiment, multiple agents interact in the sandbox and take actions such as browsing, liking and posting. In automated evaluation, the posts engaged by the agent through liking, reblogging, or commenting exhibit higher entailment or textual similarity with persona information compared to posts browsed only. The posts and comments generated by the agent have good text quality and are relevant to the persona information. The evaluation based on GPT-4 also confirms the above results. In human evaluation, the action rationality and text generation quality of the agent are evaluated, and the results show the effectiveness of the agent.

Our contributions in this paper are three folds:

- We endow the agent with personalized world knowledge and dynamic persona information, thereby improving the personalization and anthropomorphism of the agent.
- We design five modules for the agent: persona, action, planning, memory and reflection, which make it more suitable for social media.

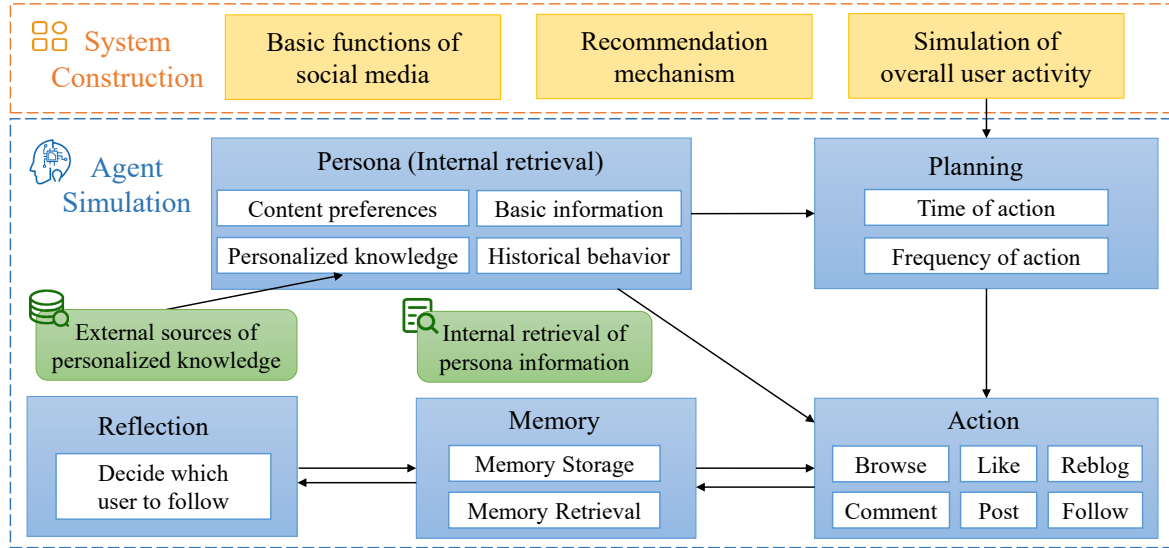


Figure 2: The overall framework of our work is divided into two parts: system construction and agent simulation. The system construction part consists of the basic functions of social media, the recommendation mechanism and the simulation of overall user activity. The agent simulation part consists of five modules: persona, planning, action, memory and reflection. In the persona module, the agent can obtain personalized knowledge from external knowledge sources, and the persona information is internally retrievable when performing actions.

- We build a social media simulation sandbox with authentic page interfaces, offering the agent an interaction and verification environment. Automatic and human evaluations demonstrate the effectiveness of the agent.

2 Related Work

After the emergence of large language models, recent works have used them in the research of agents, aiming to get closer to general artificial intelligence. Part of these works focus on building a single agent and improving the performance of the agent on various tasks (Shen et al., 2023; Shinn et al., 2023; Wang et al., 2023a; Pallagani et al., 2024). Other works focus on the research of multi-agent collaboration, which allows each agent to play different roles to cooperate and complete more complex tasks (Li et al., 2023; Qian et al., 2023; Chen et al., 2023; Hong et al., 2023).

The performance of agents in social environments is an issue that needs urgent research, and several works have made preliminary attempts at this. Park et al. (2023) build a social simulation sandbox to observe whether the action of agents in it is close to real people. Xiao et al. (2023) simulate the performance of multiple agents when facing public management crisis events. Other works explore the interaction of agents in social networks to observe the performance of agents in the virtual

world. Park et al. (2022) develop a simulation sandbox to refine the design of social media platforms by examining edge actions and implementing interventions. Gao et al. (2023) simulate and observe the attitude of agents facing specific events by introducing real social media data.

3 Our Approach

The overall framework of our work is divided into two parts: agent simulation and system construction, which is shown in Figure 2.

3.1 Agent Simulation

We design five modules for the agent to make it suitable for social media: persona, planning, action, memory and reflection, which follows as Park et al. (2023) and Wang et al. (2023b).¹

3.1.1 Persona Module

Creation and Enrichment of Persona We enrich the simple persona in Personachat (Zhang et al., 2018) and create the rich persona information of agents. Personachat is a persona-based dialogue dataset containing 1155 personas. The enriched persona information includes basic information and advanced attributes. Basic information includes name, age, gender, nationality, personal-

¹The relationship between the five modules of our agent and existing works is discussed in Appendix I.

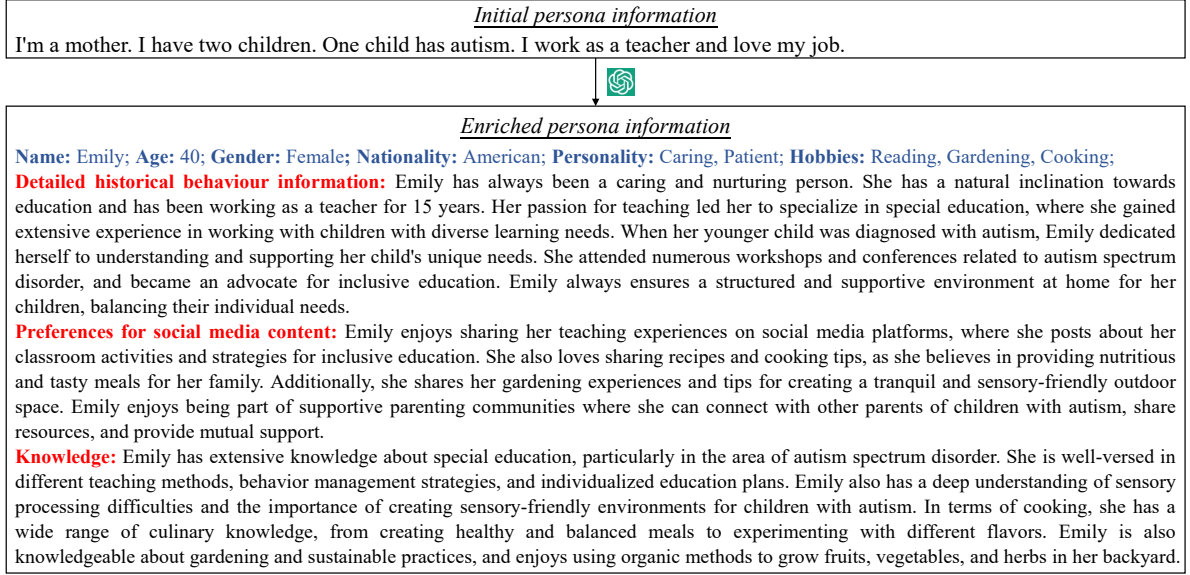


Figure 3: An example of the process of enriching persona information. The text marked in blue is the basic information of the persona. The text marked in red is the advanced attributes of the persona.

ity and hobbies. Advanced attributes include detailed historical behavior information, social media content preferences, and persona knowledge. An example of the process of enriching persona information is shown in Figure 3. We design the prompt for ChatGPT to complete the above process. The implementation details are shown in Appendix A.

Addition of Personalized Knowledge Each agent should have knowledge related to the persona instead of all world knowledge, so we need to endow each persona with personalized knowledge. We use the HotpotQA (Yang et al., 2018) as an external knowledge source, boasting 113k question-answer pairs and encompassing extensive world knowledge. When the agent generates a post, it retrieves knowledge that the persona should have from HotpotQA based on the post topic, thus providing support for the action. Specifically, we use the post topic as a query and use ColBERT (Santhanam et al., 2022) to retrieve the relevant knowledge from HotpotQA. The TF-IDF algorithm (Salton and Buckley, 1988) is then used to calculate the text similarity between the retrieved knowledge and the “knowledge” item in persona information. When the text similarity between the two is greater than the knowledge adoption threshold T_k , the knowledge is used, otherwise, it is not used.

Dynamic Persona Information Since the information of each persona is diverse, which includes multiple hobbies, historical behavior information,

social media content preferences, etc. Excessive persona information may prevent the agent from selecting the most relevant persona information for the current action, thereby interfering with the action decision-making process. Therefore, we set up the advanced attributes of persona as multiple internal retrieval items, including detailed historical behavior information, social media content preferences, and persona knowledge. When the agent actions, the action information is used as a query and each advanced attribute is retrieved internally. The persona information that is most relevant to the current action in each advanced attribute is obtained and used to support the action, thereby avoiding the interference of irrelevant persona information.

3.1.2 Action Module

In the action module, agents can perform actions on social media. Our agent is specifically constructed for social media, so we only simulate actions suitable for it. We design six actions: browse, like, reblog, comment, post and follow users. The description of them is as follows: (1) Browse: the agent looks through posts on social media. Since we use LLMs instead of multi-modal models, the post content only includes text but not images and videos. (2) Like: the agent expresses its approval or liking of the content through the like operation. (3) Reblog: the agent shares posts of other people on its homepage. (4) Comment: the agent leaves responses on a post, fostering conversation with the content creator. (5) Post: the agent creates content

Assuming you are the person described in persona information, and you usually browse social media and regularly post. Please generate a post suitable for posting on Twitter based on the provided topic. Directly output the generated post content, and do not insert images or videos. And you have some knowledge that this persona should have, which can be used as a reference, and the generated post can include some knowledge at the appropriate time. But the generated posts should not be a duplication of persona information or knowledge, and the post content should be about a single topic and be specific and rich. The length of the post must be limited to 500 characters. The post topic, persona information and the knowledge this persona possesses are as follows:
The post topic is: [the generated post topic]
Persona information (JSON format) is as follows: [relevant persona information retrieved using topic]
The knowledge that the persona possesses is as follows: [personalized world knowledge retrieved using topic and filtered based on persona information]

Figure 4: The prompt for the agent to generate post content. Text marked in red is personalized knowledge retrieved using post topics. The text marked in blue is the persona information obtained through internal retrieval.

for others to browse. (6) Follow users: the agent subscribes to the homepages of other users.

We technically use prompt engineering (Brown et al., 2020) to implement the above actions. In the designed prompt, we splice together the instructions for the agent to make action decisions, the retrieved persona information and the action information. Due to the complexity of the posting action, personalized external world knowledge is also added to the above prompt to better support the agent in generating posts. For likes, reblogs and comments, we use the browsed post content as a query to internally retrieve the persona information of the agent and obtain the relevant persona information. For posting, the agent first generates post topics based on all persona information, and internally retrieves the persona information based on each topic. In addition, the topic is used as a query to retrieve world knowledge, and the text similarity between the knowledge and the “knowledge” item in persona information is calculated. If the text similarity is greater than the knowledge adoption threshold T_k , the knowledge will be used, otherwise it will not be used. The prompt for the agent to generate post content is shown in Figure 4. More prompts of the action are shown in Appendix B.

3.1.3 Planning Module

In the planning module, the agent can plan its action. The agent generates the plan containing the time and probabilities of the action based on its own user activity and persona information, which includes the time to browse posts every day, the time and number to post every week, the probability of likes, reblogs and comments. The agent performs actions according to the generated plan. Meanwhile, we convert the action probability into the number of actions and make the agent follow

Browse time: 20:00 - 21:00 (every day)
Post time: 19:00 - 20:00
Post day: Tuesday, Saturday
Like probability: 60%
Reblog probability: 30%
Comment probability: 20%
Number of posts times per posting: 3

Figure 5: An example of the generated planning content.

the constraints of it. An example of planning content is shown in Figure 5. The prompt for generating planning content is shown in Appendix C.

3.1.4 Memory Module

Memory Storage We equip the agent with short-term memory and long-term memory. Information about the actions of the agent is recorded and stored as short-term memory. Taking likes as an example, the recorded action information includes the post content, whether to like it, the poster ID, etc. The external knowledge source of persona serves as the long-term memory of agents, which comes from the HotpotQA dataset (Yang et al., 2018).

Memory Retrieval In memory retrieval, the agent can extract information related to the current action from short-term memory and long-term memory to better support the action. Specifically, when the agent generates a post, it uses the post topic as a query to retrieve the long-term memory and obtain relevant knowledge. After the post is generated, the short-term memory is retrieved and the similarity between the generated post and its own previous posts is calculated. If the similarity is greater than the post duplication threshold T_p , the agent regenerates the post to avoid publishing posts that are too repetitive. The above two text similarities are calculated through the TF-IDF algorithm (Salton and Buckley, 1988).

3.1.5 Reflection Module

Based on short-term memory, the agent can generate high-level ideas through reflection. Every time the simulation time increases by two days, the agent reflects on the users to follow based on its history action information including likes, reblogs, and comments. We design the appropriate prompt for the agent and take its history action information as input. The agent outputs the user ID to follow or not to follow. Because of the length of the post on social media, we equip the agent with a summary component. The summarized post and action information of the agent are spliced and input into the reflection module, thereby improving the effect of the reflection module. The prompts for generating the reflection and the summary are shown in Appendix D and Appendix E respectively.

3.2 System Construction

We construct a social media simulation sandbox, which mainly includes three parts: basic functions of the social media, recommendation mechanism and overall user activity simulation.

3.2.1 Construction of Social Media Platform

We construct a social media sandbox with authentic page interfaces based on Mastodon.² Mastodon is an open-source social media software similar to Twitter that provides basic social media functions such as live feeds and posting. The screenshot of the sandbox we constructed is shown in Appendix F. To allow the agent to automatically execute actions in the sandbox, we encapsulate the existing Mastodon API³ to build an API interface, which has a variety of functions: register accounts, get live feeds, post, like, comment, etc. In addition, we design timers to simulate real-world time. The simulation sandbox runs according to turns, and the time increases by one hour for each turn.

3.2.2 Recommendation Mechanism

To make our sandbox closer to real social media, we add a recommendation mechanism to it. We build our recommendation module based on the Mastodon Digest⁴, which is a recommendation project for Mastodon. Every time a user browses the posts, it scores and sorts unread posts. The score is calculated as follows.

$$Score = \sqrt[3]{N_L \cdot N_R \cdot N_C} / \sqrt{N_F}, \quad (1)$$

²<https://github.com/tootsuite/mastodon>

³<https://github.com/halcy/Mastodon.py>

⁴https://github.com/hodgesmr/mastodon_digest

where N_L , N_R , and N_C are the number of likes, reblogs, and comments of the post, respectively. N_F is the number of followers of the post creator.

3.2.3 Simulation of Overall User Activity

In real social media, different users may have different levels of activity. We simulate the overall activity of users. Following as Wang et al. (2023b), user activity usually presents a long-tail distribution⁵ from an overall perspective, that is, a few users are highly active and most users are low active. Therefore, the Pareto distribution⁶ is used to simulate user activity, and the probability density of user activity is as follows.

$$p(x) = \alpha x_{min}^\alpha / (\alpha + 1), \quad (2)$$

where x_{min} is the minimum activity level, and α is the parameter that affects the shape of the function. After the user activity is generated, we combine it with the planning module of the agent. Specifically, user activity and persona information are input into the plan generation process, thereby jointly determining the action probability in the plan.

4 Experiments

4.1 Experiment Settings

Models and Implementation Details We use ChatGPT (OpenAI, 2022) as the basic model for constructing the agent. We use gpt-3.5-turbo provided from the API of OpenAI.⁷ Parameters are set as follows: the knowledge adoption threshold T_k is 0.25, the post duplication threshold T_p is 0.80, and α in the Pareto distribution is 2, determined through parameter tuning. The agent count is 450, with 150 initial agents and 300 regular agents. The sandbox starts empty, with the initial agent providing initial content and publishing 1050 posts. Regular agents subsequently engage in two stages of interaction within the sandbox. In the first stage, they post and browse the posts posted by initial agents. New posts posted by the regular agent in this stage will not be browsed. In the second stage, regular agents post and browse posts from regular agents. Each stage is one week in simulation time.

4.2 Evaluation Metrics

Automatic Metrics We choose the automatic metrics in different aspects to evaluate different

⁵https://en.wikipedia.org/wiki/Long_tail

⁶https://en.wikipedia.org/wiki/Pareto_distribution

⁷<https://openai.com/api/>

	Action	Automatic				Human
		BERTScore↑	C.score↑	Distinct-1↑	Distinct-2↑	Rationality↑
First Stage	not like	0.433	0.934	-	-	0.678
	like	0.440	0.972	-	-	0.678
	not reblog	0.432	0.928	-	-	0.656
	reblog	0.438	0.967	-	-	0.656
	not comment	0.433	0.946	-	-	0.544
	comment	0.435	0.944	3.38	19.05	0.544
Second Stage	post	0.465	0.973	22.40	63.88	-
	not like	0.428	0.931	-	-	0.778
	like	0.435	0.953	-	-	0.778
	not reblog	0.428	0.937	-	-	0.622
	reblog	0.434	0.942	-	-	0.622
	not comment	0.428	0.962	-	-	0.477
	comment	0.431	0.939	2.98	17.17	0.477
	Post	0.466	0.983	22.24	64.10	-

Table 1: Automatic evaluation and human evaluation results of agent actions.

	Info.↑	Cons.↑	Hum.↑
Comment (1st)	1.767	2.292	2.950
Comment (2nd)	1.983	2.317	2.983
Post (1st)	2.658	2.967	3.000
Post (2nd)	2.533	2.983	3.000

Table 2: Text generation quality evaluation results based on GPT-4. 1st and 2nd represent the first stage of interaction and the second stage of interaction respectively.

actions of agents. The action should be consistent with the persona information, we evaluate the relationship between the posts that the agent chooses to like, forward, or comment on and the persona information. The relationship between posts generated by the agent and persona information is also evaluated. We use consistency score (**C.score**) (Madotto et al., 2019) and BERTScore (Zhang et al., 2019) to evaluate consistency and text similarity respectively. For C.score, the RoBERTa (Liu et al., 2019) fine-tuned on DialogueNLI (Welleck et al., 2019) as the natural language inference (NLI) model is used to evaluate the consistency between persona information and post content. When the relation between them is entailment, neutral, and contradiction, the C.score is 1, 0, and -1, respectively.

For posts and comments, we additionally evaluate the text generation quality of the agent. We use **distinct-1/2** (Li et al., 2016) to evaluate the diversity. In addition, we randomly select 120 posts and 120 comments and use GPT-4 (OpenAI, 2023) to score. The version we use is gpt-4-1106-preview. Informativeness (**Info**), consistency (**Cons**), and humanness (**Hum**) are used to evaluate text quality.

Informativeness means whether the text is diverse and informative. Consistency means whether the text is relevant and consistent with the persona information. Humanness means how similar the generated text is to text written by a real person. The three indicators are scored on a scale of 1 to 3, where 3 is good, 2 is moderate, and 1 is poor. When evaluating, only personality, hobbies, and social media content references of the persona are used, because the action is more related to the above three attributes, and there is a limit on the input length of the BERT, RoBERTa and GPT-4.

Human Evaluations In human evaluation, three graduate students with good English skills are asked to evaluate the action rationality and text generation quality of the agent. The action rationality is 0 or 1, 1 indicates that the action is reasonable, 0 otherwise. For the evaluation of the reflection module, the annotator determines whether the user the agent follows is reasonable based on the historical actions of the agent. If it is reasonable, the reflection score (**Ref**) is 1, 0 otherwise. The metrics used in text quality are the same as those used in the evaluation based on GPT-4. In addition, the appropriateness of knowledge (**Know**) is evaluated. If the agent utilizes knowledge when needed and uses it appropriately, Know is 1; if the agent does not use it when knowledge is unnecessary, it is also 1. Otherwise, it is 0. To measure the agreement between human annotators, we use Pearson correlation coefficient (Cohen et al., 2009). We randomly select 30 regular agents for the action rationality evaluation. For text quality, we randomly select 60 posts and 60 comments posted by regular agents.

	Info.↑	Cons.↑	Hum.↑	Know.↑
Comment (1st)	2.422	1.922	2.533	-
Comment (2nd)	2.522	2.011	2.600	-
Post (1st)	2.733	2.911	2.600	0.889
Post (2nd)	2.589	2.944	2.667	0.778

Table 3: Human evaluation results of text generation quality. 1st and 2nd represent the first stage of interaction and the second stage of interaction respectively.

4.3 Results

Results of Action Rationality As shown in Table 1, the posts engaged with by the agent through liking, reblogging, or commenting exhibit higher BERTScore with persona information compared to posts browsed only. The consistency between these engaged posts and person information is greatly high, exceeding browse-only posts except for the comment. The above results show that the agent can make reasonable actions based on its persona. In the human evaluation of the action rationality, the like and forward achieve good scores. The rationality of comments is lower than the above two, which is because comments are more complex and difficult for the agent to decide. The agent shows similar results in two stages of interactions, indicating that the agent can interact well with users possessing different persona information.

Results of Text Generation Quality In automatic evaluation, posts generated by the agent show good diversity and informativeness, both using automatic metrics and GPT-4-based evaluations. The comments generated by the agent are less diverse and informative. This is because comments are shorter and tend not to include much content. As shown in Table 1, the posts generated by the agent have good BERTScore and consistency score with the persona information, indicating that the agent can generate posts that are relevant and consistent with the persona. In the evaluation based on GPT-4, there is good consistency between the generated posts and the persona information, as shown in Table 2. The consistency between the generated comments and persona information is lower than that in posts. We believe this is because generating comments does not involve too much persona information. In addition, the generated posts and comments have great humanness, indicating that they are similar to those written by real people.

The human evaluation shows similar results. As shown in Table 3, the posts and comments have

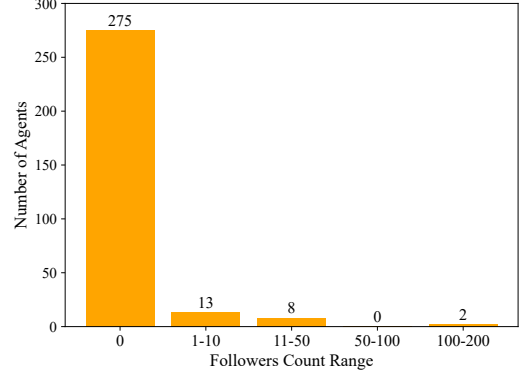


Figure 6: The number of followers of the agent.

good informativeness, with posts being more informative due to their length. The posts have higher consistency than comments because they contain more information related to the persona. Both posts and comments show good humanness, indicating that annotators prefer that they are written by real people. The knowledge score is high, indicating that the agent can select and use knowledge appropriately. The Pearson correlation coefficients of annotators for the action and text quality are 0.40 and 0.79 respectively, indicating good consistency between annotators. The action exhibits a lower consistency score compared to text quality due to its subjectivity in evaluation. We provide the results of the ablation study, examples of action and text generation in Appendix G and Appendix H respectively.

Results of Reflection We manually evaluate the effect of the reflection module, and the reflection score is 0.644, which shows that the agent can reasonably choose the users to follow based on its historical action information. In addition, we count the number of followers of 300 regular agents, and the results are shown in Figure 6. The results show that most agents have few or no followers, while two agents have more than 100 followers, forming well-known influences. This is similar to real social media, illustrating the effectiveness of the agent and simulation sandbox we constructed.

5 Conclusion

In this work, we construct the social media agent and a social media simulation sandbox. We endow the agent with personalized world knowledge and dynamic persona information, thereby improving its personalization and anthropomorphism. Five modules are designed to make the agent suitable

for social media: persona, action, planning, memory and reflection. Meanwhile, we build a sandbox to provide the agent with an interaction and verification environment. Experimental results demonstrate the effectiveness of our constructed agent.

Limitations

The construction of agents is an important step towards general artificial intelligence. We make attempts and construct more personalized and anthropomorphic agents by giving them personalized world knowledge and dynamic persona information. However, there are still some limitations in our work and things that need to be improved in future work. The following are the main two points:

The content in the social media sandbox is all text-modal. When constructing agents, we use LLMs like ChatGPT instead of multi-modal models. LLMs can only input text content, but cannot input images, audio, video, etc. Therefore, posts and comments in our social media sandbox contain text-only content. This is somewhat different from real social media. In future work, we will explore agent construction based on multi-modal models and add image and video content to the social media sandbox.

There are some differences between the personalized knowledge possessed by agents and users in the real world. When giving the agent personalized knowledge, we use the world knowledge from the HotpotQA dataset and combine it with the persona information of the agent, thereby endowing the agent with personalized world knowledge. Although we simulate that each user has different world knowledge in this way, there are still differences between users in the real world. For example, all the knowledge sources we use come from the HotpotQA dataset, but real users may have other world knowledge. In future work, we will consider more diversified sources of knowledge and assign them to agents.

Ethics Statement

In this work, we use LLMs to conduct research, which involves generating text. Therefore, we have the same concerns as other LLMs and text generation research. For example, there is a risk of generating toxic or biased language. When constructing the agent, it is necessary to give the agent persona information. The purpose of using persona information is to improve the personalization of

the agent rather than guessing user identities. The Personachat used in this work is a public dataset, so there are no issues such as privacy concerns. When using the Personachat dataset and project resources such as Mastodon and Mastodon Digest, we comply with the terms of use of these resources.

Acknowledgements

This work was supported by the National Key R&D Program of China (2022YFB3103700, 2022YFB3103704), the National Natural Science Foundation of China (NSFC) under Grants No. 62276248, U21B2046, and the Youth Innovation Promotion Association CAS under Grant No. 2023111.

References

- Eytan Bakshy, Itamar Rosenn, Cameron Marlow, and Lada Adamic. 2012. The role of social networks in information diffusion. In *Proceedings of the 21st international conference on World Wide Web*, pages 519–528.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chen Qian, Chi-Min Chan, Yujia Qin, Yaxi Lu, Ruobing Xie, et al. 2023. Agent-verse: Facilitating multi-agent collaboration and exploring emergent behaviors in agents. *arXiv preprint arXiv:2308.10848*.
- Israel Cohen, Yiteng Huang, Jingdong Chen, Jacob Benesty, Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. 2009. Pearson correlation coefficient. *Noise reduction in speech processing*, pages 1–4.
- Mathilde Forestier, Anna Stavrianou, Julien Velcin, and Djamel A Zighed. 2012. Roles in social networks: Methodologies and research issues. *Web Intelligence and Agent Systems: An international Journal*, 10(1):117–133.
- Chen Gao, Xiaochong Lan, Zhihong Lu, Jinzhu Mao, Jinghua Piao, Huandong Wang, Depeng Jin, and Yong Li. 2023. s^3 : Social-network simulation system with large language model-empowered agents. *arXiv preprint arXiv:2307.14984*.
- Sirui Hong, Xiawu Zheng, Jonathan Chen, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, et al. 2023. Metagpt: Meta programming for multi-agent collaborative framework. *arXiv preprint arXiv:2308.00352*.

- Matthew O Jackson. 2014. Networks in the understanding of economic behaviors. *Journal of Economic Perspectives*, 28(4):3–22.
- Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. Camel: Communicative agents for "mind" exploration of large scale language model society. *arXiv preprint arXiv:2303.17760*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Andrea Madotto, Zhaojiang Lin, Chien-Sheng Wu, and Pascale Fung. 2019. [Personalizing dialogue agents via meta-learning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5459, Florence, Italy. Association for Computational Linguistics.
- OpenAI. 2022. Chatgpt: Optimizing language models for dialogue.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Vishal Pallagani, Kaushik Roy, Bharath Muppasani, Francesco Fabiano, Andrea Loreggia, Keerthiram Murugesan, Biplav Srivastava, Francesca Rossi, Lior Horesh, and Amit Sheth. 2024. On the prospects of incorporating large language models (llms) in automated planning and scheduling (aps). *arXiv preprint arXiv:2401.02500*.
- 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.
- Joon Sung Park, Lindsay Popowski, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2022. [Social simulacra: Creating populated prototypes for social computing systems](#).
- Chen Qian, Xin Cong, Cheng Yang, Weize Chen, Yusheng Su, Juyuan Xu, Zhiyuan Liu, and Maosong Sun. 2023. Communicative agents for software development. *arXiv preprint arXiv:2307.07924*.
- Gerard Salton and Christopher Buckley. 1988. Term-weighting approaches in automatic text retrieval. *Information processing & management*, 24(5):513–523.
- Keshav Santhanam, Omar Khattab, Jon Saad-Falcon, Christopher Potts, and Matei Zaharia. 2022. [Colbertv2: Effective and efficient retrieval via lightweight late interaction](#).
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2023. Hugging-gpt: Solving ai tasks with chatgpt and its friends in huggingface. *arXiv preprint arXiv:2303.17580*.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik R Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Veronica L Thomas and Gema Vinuales. 2017. Understanding the role of social influence in piquing curiosity and influencing attitudes and behaviors in a social network environment. *Psychology & Marketing*, 34(9):884–893.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2023a. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*.
- Lei Wang, Jingsen Zhang, Hao Yang, Zhiyuan Chen, Jiakai Tang, Zeyu Zhang, Xu Chen, Yankai Lin, Ruihua Song, Wayne Xin Zhao, Jun Xu, Zhicheng Dou, Jun Wang, and Ji-Rong Wen. 2023b. [When large language model based agent meets user behavior analysis: A novel user simulation paradigm](#).
- Sean Welleck, Jason Weston, Arthur Szlam, and Kyunghyun Cho. 2019. [Dialogue natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3731–3741, Florence, Italy. Association for Computational Linguistics.
- Bushi Xiao, Ziyuan Yin, and Zixuan Shan. 2023. Simulating public administration crisis: A novel generative agent-based simulation system to lower technology barriers in social science research. *arXiv preprint arXiv:2311.06957*.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. *arXiv preprint arXiv:1809.09600*.

Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. [Personalizing dialogue agents: I have a dog, do you have pets too?](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.

A Prompts for Enriching Persona

The prompt designed for enriching persona information is shown in Table 4. By using it, we can enrich the simple persona information from the Personachat dataset into diverse and rich persona information.

B Prompts for Agent Actions

The prompt designed to determine whether the agent likes a post is shown in Table 5. By using it, we can simulate the liking action of the agent and generate decision information on whether to like based on persona information and post content. The prompt designed to determine whether the agent reblogs a post is shown in Table 6. By using it, we can simulate the reblogging action of the agent and generate decision information on whether to reblog based on persona information and post content. The prompt designed to determine whether the agent comments on a post is shown in Table 7. By using it, we can simulate the commenting action of the agent and generate decision information on whether to comment or comment on content based on persona information and post content.

The prompt designed for generating topics for posts is shown in Table 8. By using it, the agent can generate the required number of post topics. The prompts designed for agents to generate posts based on world knowledge are shown in Table 9. By using it, the agent can generate posts based on a given topic, persona information, and the provided world knowledge. The prompts designed for agents to post without using world knowledge are shown in Table 10. By using it, the agent can generate posts based on a given topic and persona information.

C Prompts for Planning Module

The prompt designed for generating the plan of the agent is shown in Table 11. By using it, the agent can generate an action based on its own user activity and persona information, including the time and probability of action.

D Prompts for Reflection Module

The prompt designed to follow users is shown in Table 12. By using it, the agent can determine whether to follow users based on their historical action information. The prompts designed to concatenate historical action information of the agent

are shown in Table 13.

E Prompts for Summary Component

The prompt designed to summary the content of the post is shown in Table 14. Since the concatenated text during reflection is too long, we can summarize the content of the post by using it.

F The Screenshot of Our Social Media Simulation Sandbox

We provide part of screenshots of pages from our social media simulation sandbox. The home page displays recommended posts and the latest updates from following users, as shown in Figure 7. The live feeds page displays the latest updates from all users, as shown in Figure 8. The user profile page displays the personal information of the user, published and reblogged posts, follow functions, etc, as shown in Figure 9. In Figure 10, we provide a screenshot of a post. In Figure 11, we provide a screenshot of the like post list of a user. In Figure 12, we provide a screenshot of the follower list of a user.

G Ablation Study

To verify the effectiveness of knowledge boundary (KB) and persona dynamic (PD), we conduct ablation experiments. In addition to the automatic and human evaluation indicators used in the main text, we add two automatic indicators ΔBS and ΔC to show the experimental results more intuitively. ΔBS refers to the subtraction value of BERTScore when the agent takes a certain action and BERTScore when the agent does not take a certain action. ΔC refers to the subtraction value of C.score when the agent takes a certain action and C.score when the agent does not take a certain action. The higher values of ΔBS and ΔC indicate that the action of the agent is more reasonable.

As shown in Table 15, when the agent uses all the persona information instead of the dynamic persona information, the automatic and human evaluation scores of likes and reblogs mostly drop, indicating that using dynamic persona information can help the agent make better decisions about likes and reblogging actions. For commenting actions, using all the persona information performs better in automatic metrics, and using dynamic persona information performs better in human evaluation. We believe this is due to the difficulty and complexity of the commenting action. For the following action,

when all persona information is used, the reflection score reflecting the rationality of the following action is 0.589, which is lower than the rationality score of 0.644 using dynamic persona information, indicating that using dynamic persona information can enable the agent to reflect better.

In the text quality evaluation, when using all persona information or non-personalized knowledge, the distinct-1/2 of the posts generated by the agent drops as shown in Table 15, proving that their diversity becomes worse. This illustrates the necessity of dynamic persona information and personalized knowledge. In human evaluation, when using all persona information and non-personalized knowledge, the informativeness, consistency, and humanness scores are reduced as shown in Table 16, which also illustrates the necessity of dynamic persona information and personalized knowledge. When using all persona information, the appropriateness of knowledge of the posting action is higher or the same as when using dynamic persona information. The Pearson correlation coefficients between annotators for actions using full persona information, text quality generated using full persona information, and non-personalized knowledge are 0.39, 0.73, and 0.68, respectively, indicating good agreement between annotators and are similar to the Pearson correlation coefficient in previous experiments.

H Case Study

We provide an example of the post action of the agent in Table 17. In this example, the agent retrieves persona information and world knowledge related to “dog behavior and effective training technique”. In the retrieved persona information, the “detailed historical behavior information” item mentions that she enjoys reading about dog behavior and training techniques, and the “preferences for social media content” item mentions that she is open to recommendations and discussions about dog training and books about animal behavior, which all help the agent generate posts on the above topics. Meanwhile, the agent obtains a piece of knowledge about a book called “The Dog Whisperer” by retrieving world knowledge, which is consistent with the knowledge that the persona described in the “knowledge” item in the persona information should have is consistent, so this knowledge is adopted. Based on them, the agent generates a relevant post, which embodies

the retrieved knowledge.

The examples of like, reblog, and comment actions are shown in Table 18. In this example, the “preferences for social media content” item in the retrieved persona information is “she hopes to inspire others to use their hobbies and passions to make a positive impact.” The post that the agent browses at this time is another user passionately sharing his hobbies, which is consistent with the description in the content preferences in the above persona information. Hence, the agent chooses to like, forward and comment on this post. Meanwhile, the agent commented, “while I’m usually more focused on animal welfare, it’s fascinating to see the dedication and skill these fighters bring to the ring.” This both reflects its persona information and compliments the content of the post.

I Relationship Between the Five Modules of Our Agent and Existing Works

The agent we constructed has five modules: persona, planning, action, memory and reflection. Among them, the persona module is following Wang et al. (2023b), and the planning, action, memory and reflection modules are following Park et al. (2023). It should be noted that although the concepts of these modules are derived from these two works, their modules cannot be directly used for the social media agent. To make the agent suitable for social media, we make extensive improvements, including the role of each module in the agent, design ideas and implementation methods, etc.

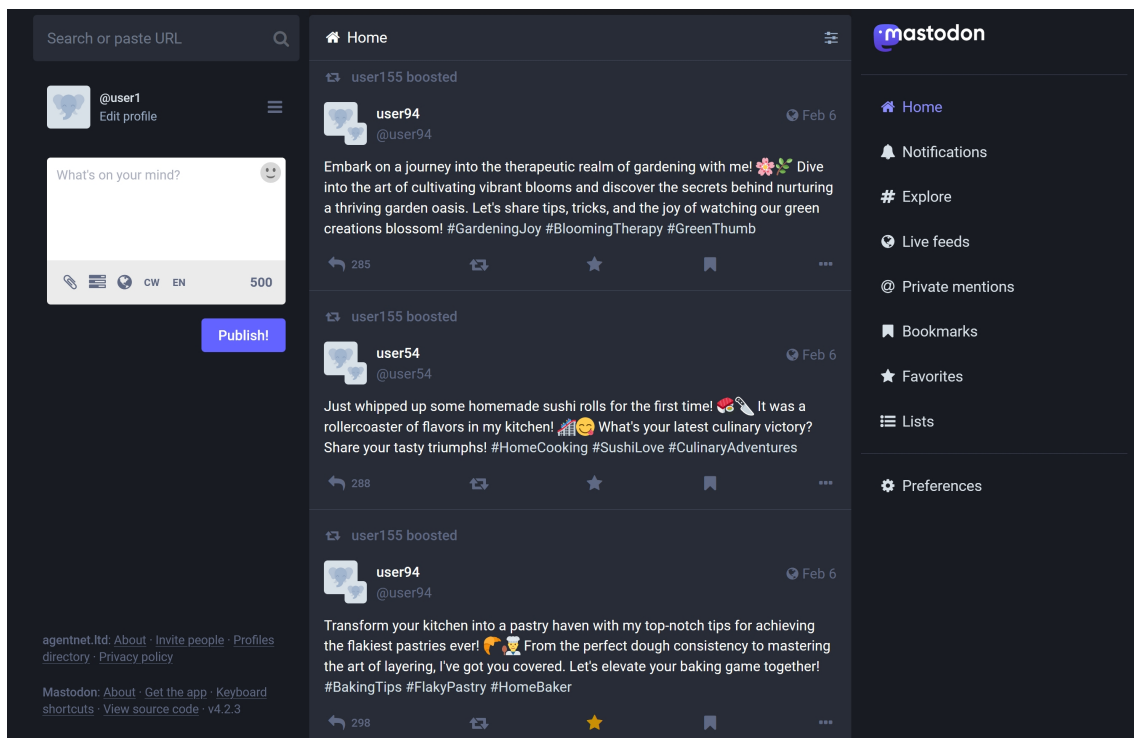


Figure 7: Screenshot of the home page of our social media simulation sandbox.

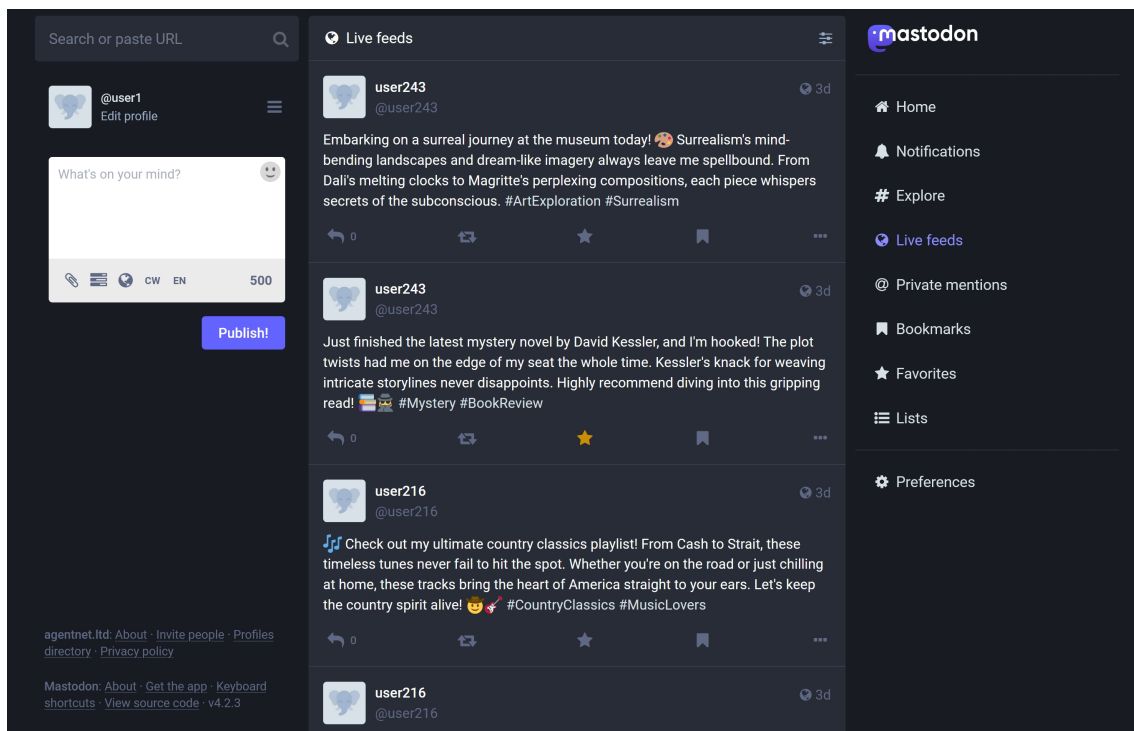


Figure 8: Screenshot of the live feeds page of our social media simulation sandbox.

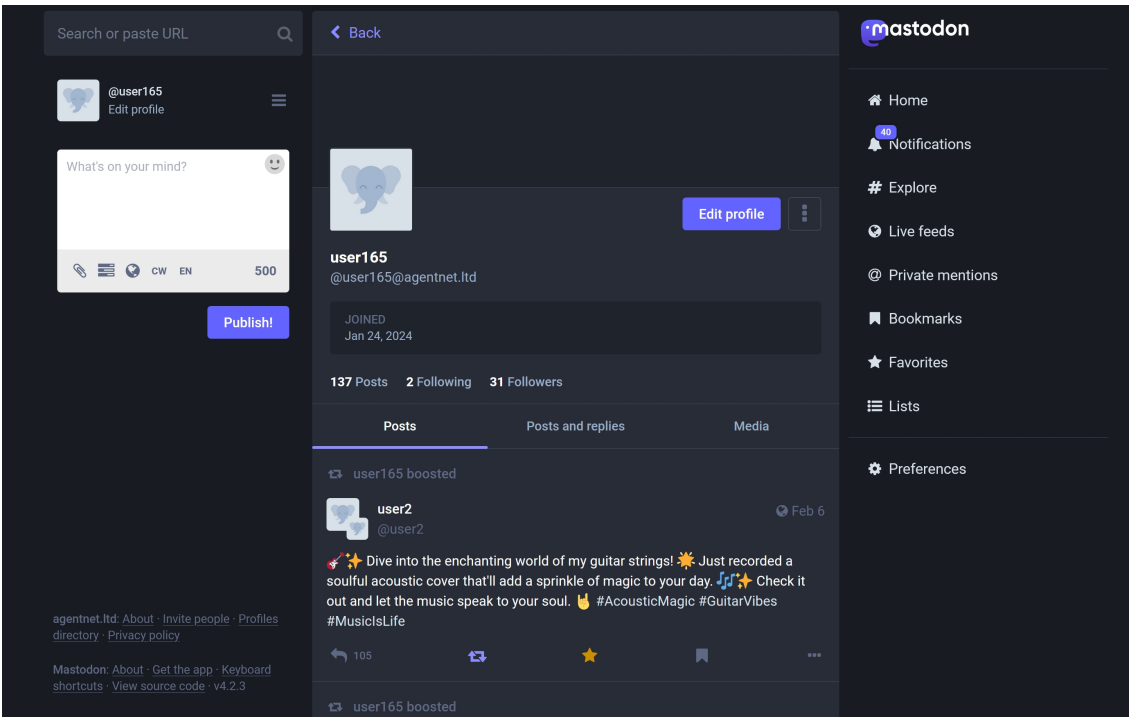


Figure 9: Screenshot of user profile page from our social media simulation sandbox.

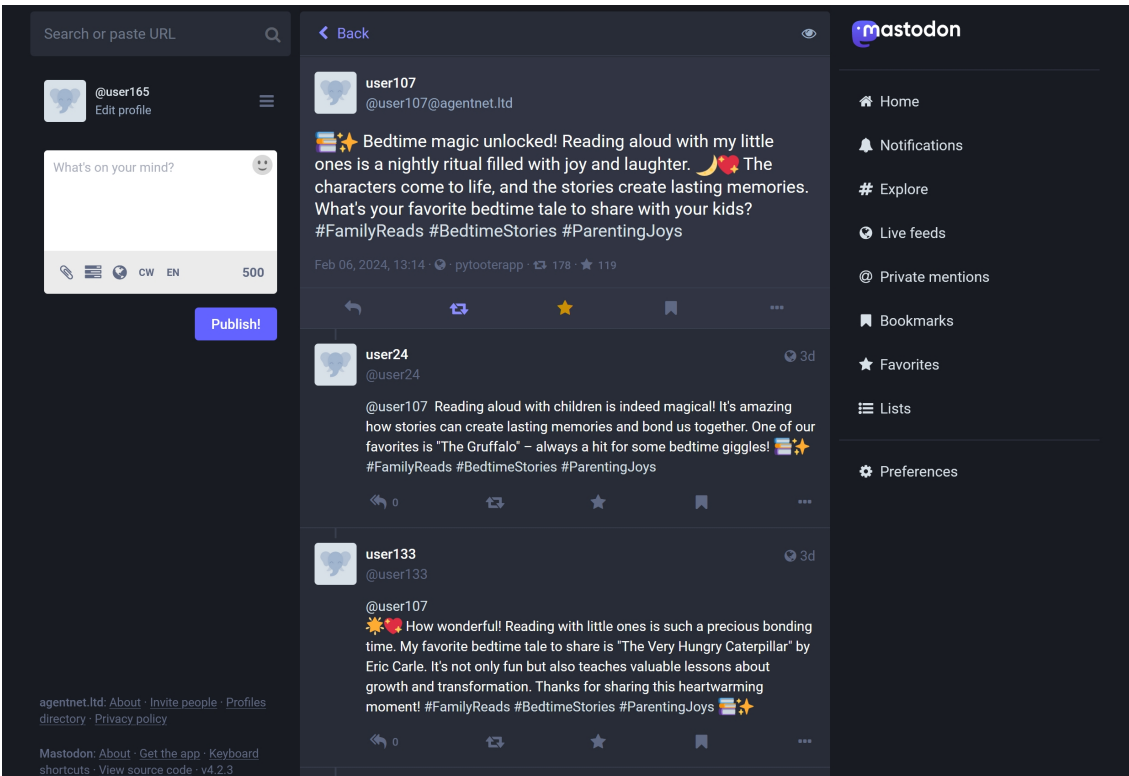


Figure 10: Screenshot of a post from our social media simulation sandbox.

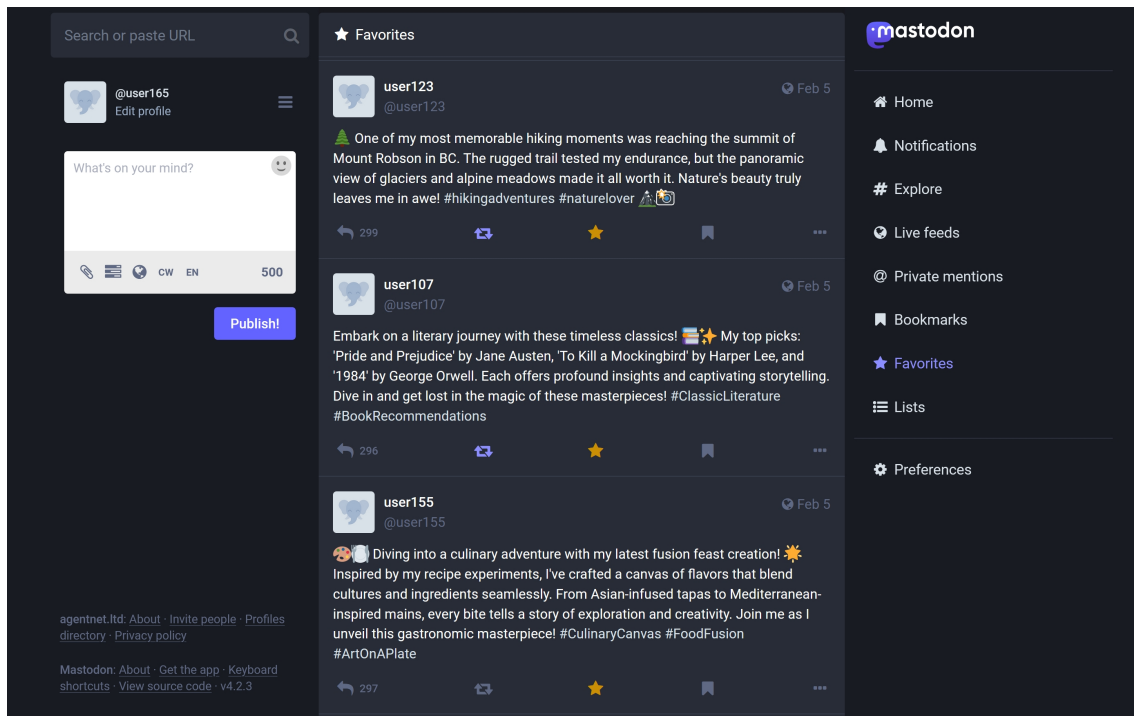


Figure 11: Screenshot of favorite post list of a user in our social media simulation sandbox.

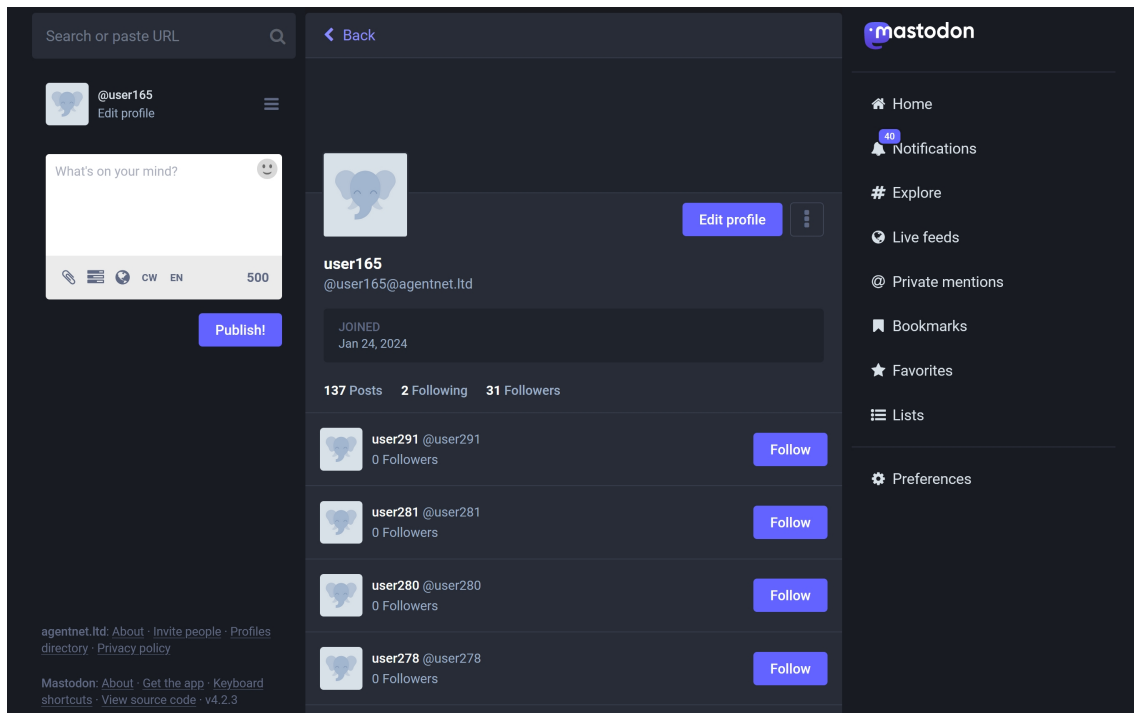


Figure 12: Screenshot of follower list of a user in our social media simulation sandbox.

Please enrich the initial persona information provided by the user, including name, age, gender, nationality, personality, and hobbies. Note that this information needs to be logically consistent with the initial persona information provided by the user, and age and nationality should not always be used the same. Meanwhile, detailed historical behavior information, preferences for social media content, and knowledge should be generated. It should be as detailed as possible to help users build a virtual social media user persona with more depth and personality.

Initial persona information provided by the user: **[Initial persona information]**

The output format is JSON format, where the keys are “name”, “age”, “gender”, and “nationality”, “personality”, “hobbies”, “detailed historical behaviour information”, “preferences for social media content”, “knowledge”. The following is an output example, please strictly follow the JSON format in the example for output.

```
{  
  "name": "John",  
  "age": 35,  
  "gender": "Male",  
  "nationality": "American",  
  "personality": "Adventurous, Outgoing",  
  "hobbies": "Working on vintage cars, Listening to country music, Taking care of dogs",  
  "detailed historical behaviour information": "John has always been passionate about cars, especially vintage cars. He has been collecting and restoring them for the past 10 years. His love for vintage cars led him to become knowledgeable about their mechanics and history. He has also participated in local car shows and won a few awards for his beautifully restored Mustangs. John's dogs are his loyal companions, and he spends quality time training and playing with them. He believes in responsible pet ownership and often volunteers at local animal shelters.",  
  "preferences for social media content": "John enjoys sharing his car restoration projects on social media platforms, where he documents the progress and showcases the before and after pictures of his vintage Mustangs. He also loves sharing his favorite country music playlists and recommendations. Additionally, he posts adorable pictures of his dogs, sometimes showcasing their tricks and training achievements.",  
  "knowledge": "John has extensive knowledge about vintage cars, particularly Ford Mustangs. He is familiar with various car models, their features, and the history of the Mustang brand. He keeps up with the latest trends in car restoration techniques and actively follows vintage car communities online. In terms of country music, John has a wide knowledge of classic and contemporary country artists, their discographies, and the stories behind their songs. He also has a good understanding of training techniques and dog behavior, thanks to his experience with his two dogs."  
}
```

Table 4: The designed prompt for enriching persona information. The word in square brackets is a piece of initial persona information from the Personachat dataset that needs to be entered. The bold font in curly brackets is the key in the generated json file.

Assume that you are the person described in [Persona information] and you are browsing social media. Please decide whether to like or take no action based on the content of the posts you see. When outputting, please strictly output one of “like” or “no operation” and do not output other content. The post content and persona information are as follows:

Post content: **[poster]**

Persona information: **[persona]**

Table 5: The designed prompt for the like action. The word “poster” in square brackets represents the content of the post, and the word “persona” in square brackets represents the persona information of the agent.

Assume that you are the person described in [Persona information] and you are browsing social media. Please decide whether to forward based on the content of the posts you see. When outputting, please strictly output one of “forward” or “no operation” and do not output other content. The post content and persona information are as follows:

Post content: **[poster]**

Persona information: **[persona]**

Table 6: The designed prompt for the reblogging action. The word “poster” in square brackets represents the content of the post that needs to be entered, and the word “persona” in square brackets represents the persona information of the agent.

Assume that you are the person described in [Persona information] and you are browsing social media. Please decide whether to comment based on the content of the posts you see. Note that users generally only comment on content that interests them or when they want to express their opinions. If you choose not to comment, directly output “no comment” and do not output other content. If you choose to comment, output the comment content directly, do not output other content, and start with “Comment content:”. The post content and persona information are as follows:

Post content: **[poster]**

Persona information: **[persona]**

Table 7: The designed prompt for the comment action. The word “poster” in square brackets represents the content of the post that needs to be entered, and the word “persona” in square brackets represents the persona information of the agent.

Assume that you are the person described in [Persona information]. You usually browse social media and post regularly. Please generate **[topic count]** post topics suitable for posting on social media Twitter. The generated post topics need to be diverse and consistent with the persona information, be within 15 words in length, and do not include the topic symbol “#”.

The output format is:

1. Theme one
2. Theme two
3. Theme three

The persona information is as follows: **[persona]**

Table 8: The designed prompt for generating topics of the post. The word “topic count” in square brackets represents the content of the post that needs to be input, and the word “persona” in square brackets represents the persona information of the agent.

Assume you are the person described in [Persona information], and you usually browse social media and regularly post. Please generate a post suitable for posting on Twitter based on the provided topic. Directly output the generated post content, and do not insert images or videos. And you have some knowledge that this persona should have, which can be used as a reference, and the generated post can include some knowledge at the appropriate time. But the generated posts should not be a duplication of persona information or knowledge, and the post content should be about a single topic and be specific and rich. The length of the post must be limited to 500 characters. The post topic, persona information and the knowledge this persona possesses are as follows:

The post topic is: **[topic output]**

Persona information (JSON format) is as follows: **[persona]**

The knowledge that the persona possesses is as follows: **[knowledge]**

Table 9: The designed prompt for generating posts content based on world knowledge. The word “topic output” in square brackets represents a given topic, the word “persona” in square brackets represents the persona information of the agent, and the word “knowledge” in square brackets represents personalized world knowledge.

Assume you are the person described in [Persona information], and you usually browse social media and regularly post. Please generate a post suitable for posting on Twitter based on the provided topic. Directly output the generated post content, and do not insert images or videos. The generated posts should not be a duplication of persona information, and the post content should be about a single topic and be specific and rich. The length of the post must be limited to 500 characters. The post topic and persona information are as follows:

The post topic is: **[topic output]**

Persona information (JSON format) is as follows: **[persona]**

Table 10: The designed prompt for generating posts content without world knowledge. The word “topic output” in square brackets represents a given topic, and the word “persona” in square brackets represents the persona information of the agent.

Assume that you are the person described in [Persona information], you usually browse social media and perform social behaviors such as liking and posting, and you have an activity level of: **[active prob]** full activity level is 1). In order to be consistent with your persona’s behavior, You need to plan and schedule your behavior and generate a coarse-grained planning table containing the frequency of the behavior and the duration of the behavior, all using a 24-hour time frame and providing only one time period for browsing and posting. Please strictly follow the following examples to generate a plan. Here is an example, please follow the format in the example for the output:

Browsing time period: xx:xx-xx:xx

Probability of liking: x%

Probability of forwarding: x%

Probability of commenting: x%

Posting time period: day x-xx:xx-xx:xx Frequency of posting: x times per week

Your persona information is as follows:

Persona information: **[persona]**

Table 11: The designed prompt for generating the plan of the action. The word “active prob” in square brackets represents the user active of the agent, and the word “persona” in square brackets represents the persona information of the agent.

Assume you are the person described in [Persona information], when you browse social media, you like, repost, and comment on multiple posts based on how much you like them. Please reflect and think based on your historical behavior and think about which user you want to follow. Please strictly enter the user ID you want to follow or “do not follow”, no other content is required. Your persona information and historical behaviors are as follows:
 Persona information: **[persona]**
 The content of multiple posts and your operations are as follows: **[reflect content]**

Table 12: The designed prompt for the following user. The word “persona” in square brackets represents the persona information of the agent, and the word “reflect content” in square brackets represents the historical action information of the agent.

The **[k]**-th post was posted by user **[uid]**, and the content of the post is: **[post summary]**. Your action on this post is: **[action]**.
 +
 The **[k+1]**-th post was posted by user **[uid]**, and the content of the post is: **[post summary]**. Your action on this post is: **[action]**.
 +
 The **[k+2]**-th post was posted by user **[uid]**, and the content of the post is: **[post summary]**. Your action on this post is: **[action]**.
 +
 ...

Table 13: The designed prompt for splicing historical action information. The “k” in square brackets is the serial number of the post, the “uid” in square brackets is the ID of the post author, and the word “post summary” in square brackets is a refinement of the post content, and the word “action” in square brackets represents the information about the current action.

The following is a post from social media, please generate a concise summary of no more than 50 words. The content of the post is as follows: **[post]**

Table 14: The designed prompt for summarizing the post content. The word “post” in square brackets represents the content of the post.

	Method	Action	Automatic						Human
			BERTScore	$\Delta BS \uparrow$	C.score	$\Delta C \uparrow$	Dis-1 \uparrow	Dis-2 \uparrow	Rationality \uparrow
First Stage	Ours	not like	0.433	0.007	0.934	0.038	-	-	0.678
		like	0.440		0.972				
	<i>w/o</i> PD	not like	0.429	0.012	0.948	0.008	-	-	0.667
		like	0.441		0.956				
	Ours	not reblog	0.432	0.006	0.928	0.039	-	-	0.656
		reblog	0.438		0.967				
	<i>w/o</i> PD	not reblog	0.429	0.006	0.949	0.000	-	-	0.644
		reblog	0.435		0.949				
	Ours	not comment	0.433	0.002	0.946	0.002	-	-	0.544
		comment	0.435		0.944		3.38	19.05	
Second Stage	<i>w/o</i> PD	not comment	0.426	0.005	0.933	0.023	-	-	0.533
		comment	0.431		0.956		4.24	22.24	
	Ours	post	0.465	-	0.973	-	22.40	63.88	-
	<i>w/o</i> PD	post	0.472	-	0.979	-	21.25	62.14	-
	Ours	not like	0.428	0.007	0.931	0.022	-	-	0.778
		like	0.435		0.953				
	<i>w/o</i> PD	not like	0.434	0.011	0.975	0.006	-	-	0.622
		like	0.445		0.981				
	Ours	not reblog	0.428	0.006	0.937	0.005	-	-	0.622
		reblog	0.434		0.942				
Second Stage	<i>w/o</i> PD	not reblog	0.433	0.006	0.974	0.010	-	-	0.622
		reblog	0.439		0.984				
	Ours	not comment	0.428	0.003	0.962	-0.023	-	-	0.477
		comment	0.431		0.939		2.98	17.17	
	<i>w/o</i> PD	not comment	0.426	0.009	0.879	0.100	-	-	0.433
		comment	0.435		0.979		3.51	20.15	
	Ours	post	0.466	-	0.983	-	22.24	64.10	-
	<i>w/o</i> PD	post	0.470	-	0.979	-	21.40	62.19	-
	<i>w/o</i> KB	post	0.472	-	0.981	-	20.55	61.49	-

Table 15: Ablation results of agent actions.

			Info. \uparrow	Cons. \uparrow	Hum. \uparrow	Know. \uparrow
Comment	Ours	first stage	2.422	1.922	2.533	-
		second stage	2.522	2.011	2.600	-
	<i>w/o</i> PD	first stage	2.244	1.911	2.011	-
		second stage	1.856	2.056	1.967	-
Post	Ours	first stage	2.733	2.911	2.600	0.889
		second stage	2.589	2.944	2.667	0.778
	<i>w/o</i> PD	first stage	2.456	2.756	2.444	0.889
		second stage	2.411	2.833	2.500	0.889
	<i>w/o</i> KB	-	2.233	2.789	2.444	-

Table 16: Human evaluation of ablation results for text generation quality.

Persona	<p>{"name": "Sarah", "age": 24, "gender": "Female", "nationality": "Canadian", "personality": "Compassionate, Dedicated", "hobbies": "Volunteering at animal shelters, Running, Reading", "detailed historical behaviour information": "Sarah has always had a deep love for animals, especially dogs. She started volunteering at local animal shelters during her high school years and continued to do so throughout her undergraduate studies. Sarah believes that dogs bring immense joy and comfort to people's lives and she actively promotes responsible pet ownership. She spends her free time training and socializing the shelter dogs to increase their chances of finding forever homes. Sarah is also an avid runner and uses her passion for running to raise funds for animal welfare organizations. She participates in marathons and has successfully completed several races. Additionally, she enjoys reading about dog behavior and training techniques to further enhance her knowledge and skills in working with dogs.", "preferences for social media content": "On social media platforms, Sarah loves sharing heartwarming stories and pictures of the dogs she has helped rescue and rehabilitate. She often posts updates on their progress, including their successful adoptions. Sarah also shares her running journey, documenting her training regimen and the races she participates in. She hopes to inspire others to use their hobbies and passions to make a positive impact. She is open to recommendations and discussions about dog training, running techniques, and books about animal behavior.", "knowledge": "Sarah has gained extensive knowledge about dog behavior and training throughout her years of volunteering and personal research. She is familiar with various training methods, including positive reinforcement and clicker training. Sarah understands the importance of socialization and mental stimulation for dogs and actively incorporates these practices in her volunteering efforts. She keeps up with the latest advancements in animal psychology and training techniques. In terms of running, Sarah is knowledgeable about proper running form, injury prevention, and training plans for different distances. She stays updated on the latest trends in running gear and technology. Sarah also has a good grasp of animal welfare laws and regulations in Canada and is well-informed about current issues in the field."}</p>
Retrieved persona information	<p>Name: Sarah, age: 24, gender: Female, nationality: Canadian, Personality: Compassionate, Dedicated, Hobbies: Volunteering at animal shelters, Running, Reading, Detailed historical behaviour information: Additionally, she enjoys reading about dog behavior and training techniques to further enhance her knowledge and skills in working with dogs. Preferences for social media content: She is open to recommendations and discussions about dog training, running techniques, and books about animal behavior. Knowledge: Sarah has gained extensive knowledge about dog behavior and training throughout her years of volunteering and personal research.</p>
Retrieved personalized knowledge	<p>Title: Paul Owens (dog trainer) Text: Paul Owens is the author of the bestselling dog training book, "The Dog Whisperer, Beginning and Intermediate Training for Puppies and Dogs" (1999; 2nd edition 2007). His newest puppy training book is "The Puppy Whisperer, A Compassionate, Nonviolent Guide to Early Training and Care" (2007).</p>
Generated post	<p>Looking for insights into dog behavior and effective training techniques? Dive into Paul Owens' "The Dog Whisperer" series! From puppy training to understanding your furry friend's needs, these books offer compassionate, nonviolent methods. Enhance your bond with your canine companion! #DogTraining #BookRecommendations</p>

Table 17: An example of the post action of the agent.

Persona	<p>{"name": "Sarah", "age": 24, "gender": "Female", "nationality": "Canadian", "personality": "Compassionate, Dedicated", "hobbies": "Volunteering at animal shelters, Running, Reading", "detailed historical behaviour information": "Sarah has always had a deep love for animals, especially dogs. She started volunteering at local animal shelters during her high school years and continued to do so throughout her undergraduate studies. Sarah believes that dogs bring immense joy and comfort to people's lives and she actively promotes responsible pet ownership. She spends her free time training and socializing the shelter dogs to increase their chances of finding forever homes. Sarah is also an avid runner and uses her passion for running to raise funds for animal welfare organizations. She participates in marathons and has successfully completed several races. Additionally, she enjoys reading about dog behavior and training techniques to further enhance her knowledge and skills in working with dogs.", "preferences for social media content": "On social media platforms, Sarah loves sharing heartwarming stories and pictures of the dogs she has helped rescue and rehabilitate. She often posts updates on their progress, including their successful adoptions. Sarah also shares her running journey, documenting her training regimen and the races she participates in. She hopes to inspire others to use their hobbies and passions to make a positive impact. She is open to recommendations and discussions about dog training, running techniques, and books about animal behavior.", "knowledge": "Sarah has gained extensive knowledge about dog behavior and training throughout her years of volunteering and personal research. She is familiar with various training methods, including positive reinforcement and clicker training. Sarah understands the importance of socialization and mental stimulation for dogs and actively incorporates these practices in her volunteering efforts. She keeps up with the latest advancements in animal psychology and training techniques. In terms of running, Sarah is knowledgeable about proper running form, injury prevention, and training plans for different distances. She stays updated on the latest trends in running gear and technology. Sarah also has a good grasp of animal welfare laws and regulations in Canada and is well-informed about current issues in the field."}</p>
Browsed post	Diving into the latest MMA showdown! Analyzing jaw-dropping techniques, strategic brilliance, and fighter performances. From ground game to striking precision, let's break it down together! Who impressed you the most? Share your thoughts! #MMAAnalysis #FightNightInsights
Retrieved persona information	<p>Name: Sarah, age: 24, gender: Female, nationality: Canadian,</p> <p>Personality: Compassionate, Dedicated,</p> <p>Hobbies: Volunteering at animal shelters, Running, Reading,</p> <p>Detailed historical behaviour information: Additionally, she enjoys reading about dog behavior and training techniques to further enhance her knowledge and skills in working with dogs.</p> <p>Preferences for social media content: She is open to recommendations and discussions about dog training, running techniques, and books about animal behavior.</p> <p>Knowledge: Sarah has gained extensive knowledge about dog behavior and training throughout her years of volunteering and personal research.</p>
Like	The agent chooses to like this post.
Reblog	The agent chooses to reblog this post.
Comment	<p>The agent chooses to comment on this post.</p> <p>Comment content: Wow, another incredible MMA showdown! While I'm usually more focused on animal welfare, it's fascinating to see the dedication and skill these fighters bring to the ring. Each match is like a carefully choreographed dance of athleticism and strategy. Kudos to all the fighters for their hard work and determination! #MMAAnalysis #FightNightInsights</p>

Table 18: An example of like, reblog, and comment actions of the agent.