How Personality Traits and Interaction Histories Can Affect Cooperative Behavior in an LLM-Based Social Particle Swarm

Taisei Hishiki[1], Takaya Arita[1] and Reiji Suzuki[1]

[1]Graduate School of Information Science, Nagoya University
hishiki.taisei.g3@s.mail.nagoya-u.ac.jp

Abstract

Generative Agent-Based Modeling (GABM) using Large Language Models (LLMs) provides new avenues for studying complex social phenomena. This study investigates how personality traits and interaction history affect cooperative behavior and collective dynamics in an LLM-based Social Particle Swarm (SPS) model, where agents move in a two-dimensional space and play the prisoner's dilemma game with their neighbors. We replaced conventional agents in the SPS framework with LLM agents endowed with diverse Big Five personality scores and varying memory lengths for the previous history of games with neighbors. Our experiments revealed that memory length is a critical factor governing collective behavior. The longer memory drastically suppressed cooperation, transitioning the system from stable cooperative clusters to dynamic cycles of cooperation and collapse, and ultimately leading to a state of scattered defection. Furthermore, agents' personality traits correlated with their actions; for example, high Agreeableness was consistently linked to more cooperative and less mobile behavior, partially aligned with findings from experiments with human participants. We also found that agents without explicitly assigned personalities sustained higher levels of cooperation. These results demonstrate that cognitive parameters, such as memory, can qualitatively alter collective dynamics and highlight the potential of LLM-based models for exploring the intricate interplay between psychology and social behavior.

Introduction

Generative Agent-Based Modeling (GABM) using Large Language Models (LLMs) has offered new avenues for understanding complex social phenomena by simulating agents with human-like reasoning (Park et al. (2023), Chen et al. (2023), Lu et al. (2024)). A central question is how cooperation emerges and evolves within these AI agent populations, connecting vast knowledge from evolutionary game theory (Nowak (2006)) to the dynamics of these new agents (Sun et al. (2025)).

The purpose of this study is to understand how individuality (e.g., persona, personality traits) and interaction histories can affect individual behavior and collective dynamics in a group of LLM agents. We focus on the personality traits described in prompts for their decision-making process. Jiang et al. explored the ability of Large Language Models (LLMs) to express personality traits by simulating distinct LLM personas based on the Big Five personality model, demonstrating that their self-reported Big Five Inventory scores were consistent with their assigned personality types and that their writings exhibited representative linguistic patterns (Jiang et al. (2024)). It is also reported that such personality descriptions can affect cooperative behavior (Phelps and Russell (2023)) and can evolve (Suzuki and Arita (2025)) in game-theoretical situations such as the prisoner's dilemma.

We also focus on the effects of interaction history on collective dynamics in LLM agents, as they are expected to exhibit diverse behaviors stemming from their personality traits when facing different situations. The relationship between memory length and cooperative behavior presents contradictory findings in evolutionary models. While some studies suggest longer memory enhances cooperation through improved recognition of past behaviors and the stabilization of reciprocal strategies (Hauert and Schuster (1997); Li and Kendall (2014)), recent evidence shows that excessive memory can be detrimental to cooperation due to reputation-based punishment traps that prevent forgiveness (Horvath et al. (2012)) or information quality trade-offs in learning (Alonso-Sanz (2009)). Luo et al. (2016) demonstrate that intermediate memory lengths optimally facilitate cooperation on interdependent networks, while Qin et al. (2008) show context-dependent effects where memory benefits vary with model parameters, suggesting diverse mechanisms underlying the complex relationship between memory and cooperation.While these studies provide valuable insights, they predefine how memory is utilized in specific ways. This may limit an LLM-based agent's ability to interpret and act on its experiences in a flexible, context-dependent manner.

The Social Particle Swarm (SPS) (Nishimoto et al. (2023)) is a model for studying the emergence and collapse of cooperative groups in social dilemma situations. The model integrates a physical model of self-driven particles with game theory, providing a framework to capture the co-evolution of cooperative behavior and social relationships (represented as distances between particles in a 2D space) in continuous space and time. Agents (particles) play the Prisoner's Dilemma game with their neighbors, moving and altering their relationships by attracting or repelling each other based on the payoffs received. The previous study revealed that the population primarily exhibits three characteristic collective states: general defection (Class A), general cooperation (Class B), and a cyclical dynamic of cooperation and defection (Class C). The dynamics observed in Class C, characterized by the formation, invasion, and collapse of cooperative clusters, are particularly interesting as they have also been observed in experiments involving human participants (Suzuki et al. (2018)), capturing the instability of cooperation and fluidity of relationships in real societies.

Studies on the SPS model suggested that the diversity of individuality among agents plays a crucial role in maintaining these collective dynamics, especially the dynamic Class C. In the original SPS model, individuality is represented as individual differences in a parameter indicating the propensity to cooperate (i.e., the cooperation threshold for the proportion of cooperators in their neighbors) (Nishimoto et al. (2023)). Furthermore, web-based experiments with human participants using the SPS model framework indicated some correlation between subjects' personality traits and their cooperative behavior and relationship-building with others within the experiment(Suzuki et al. (2018)). These findings highlight the importance of individuals' internal characteristics and their diversity in understanding collective dynamics.

To achieve our purpose, we extend the SPS model by replacing the conventional agents, whose behavior is described by a simple rule, with LLM agents endowed with diverse Big Five personality traits, represented as a set of numeric parameters, and their interaction histories of past game results and payoffs. We then explore how personality and memory affect the dynamics of the social groups they form. Specifically, we discuss the relationship between their personality traits and individual behavioral tendencies, as well as the effects of the memory length of an individual's interaction history with others on their collective behavior. The experiments revealed that varying memory length in LLM agents leads to distinct collective behaviors, indicating that longer memory lengths have adverse effects on cooperation. Moreover, the observed correlations between agent personality traits and their actions show partial consistency with findings in experiments with human participants.

We also preliminarily consider cases where personality traits are not explicitly specified, showing that unspecifying the personality may promote cooperation.

Model

The model is based on the Social Particle Swarm (SPS) model, with agent decision-making replaced by LLMs. A population of N agents operates within a 2D toroidal plane of size $W \times W$, interacting with others within a radius R.

Agent Behavior

At each time step t, each agent acts according to the following steps, as illustrated in Figure 1, which shows both the prompt structure (left panel) and the spatial decision-making context (right panel):

1. Situation Recognition and Decision-Making by LLM

   Each agent i's LLM is fed a prompt containing the necessary information for decision-making. This prompt, which will be explained in detail in the next section, includes the agent's current state (position $\mathbf{x}_i(t)$, strategy $s_i(t)$, cumulative score $Score_i(t)$), its pre-assigned Big Five personality traits, its recent interaction history (based on memory length L_m), and the status of other agents in its neighborhood $N_i(t)$ (strategy $s_j(t)$, relative position $\mathbf{u}_{ij}(t)$). Based on this prompt, the LLM determines the strategy $s_i(t+1)$ to adopt at the next time step $t+1$ and the movement action (i.e., magnitude and direction of movement), to maximize its score, outputting them along with its reasoning. The agent moves to a new position $\mathbf{x}_i(t+1)$ according to the determined movement action, with its speed capped at a maximum value MAX_SPEED.

2. Payoff Calculation and Score Update

   The instantaneous total score $G_i(t)$ at time t is calculated based on the strategy $s_i(t)$ at time t and the strategies $s_j(t)$ of neighboring agents. If the basic payoff agent i obtains from a Prisoner's Dilemma (PD) game with a neighboring agent $j \in N_i(t)$ is $g_{\text{base}}(s_i(t), s_j(t))$, considering the decay in their social closeness according to the distance between them $|\mathbf{u}_{ij}(t)|$, the instantaneous total score agent i obtains at time t is calculated as:

$$G_i(t) = \sum_{j \in N_i(t)} \frac{g_{\text{base}}(s_i(t), s_j(t))}{1 + |\mathbf{u}_{ij}(t)|}.$$

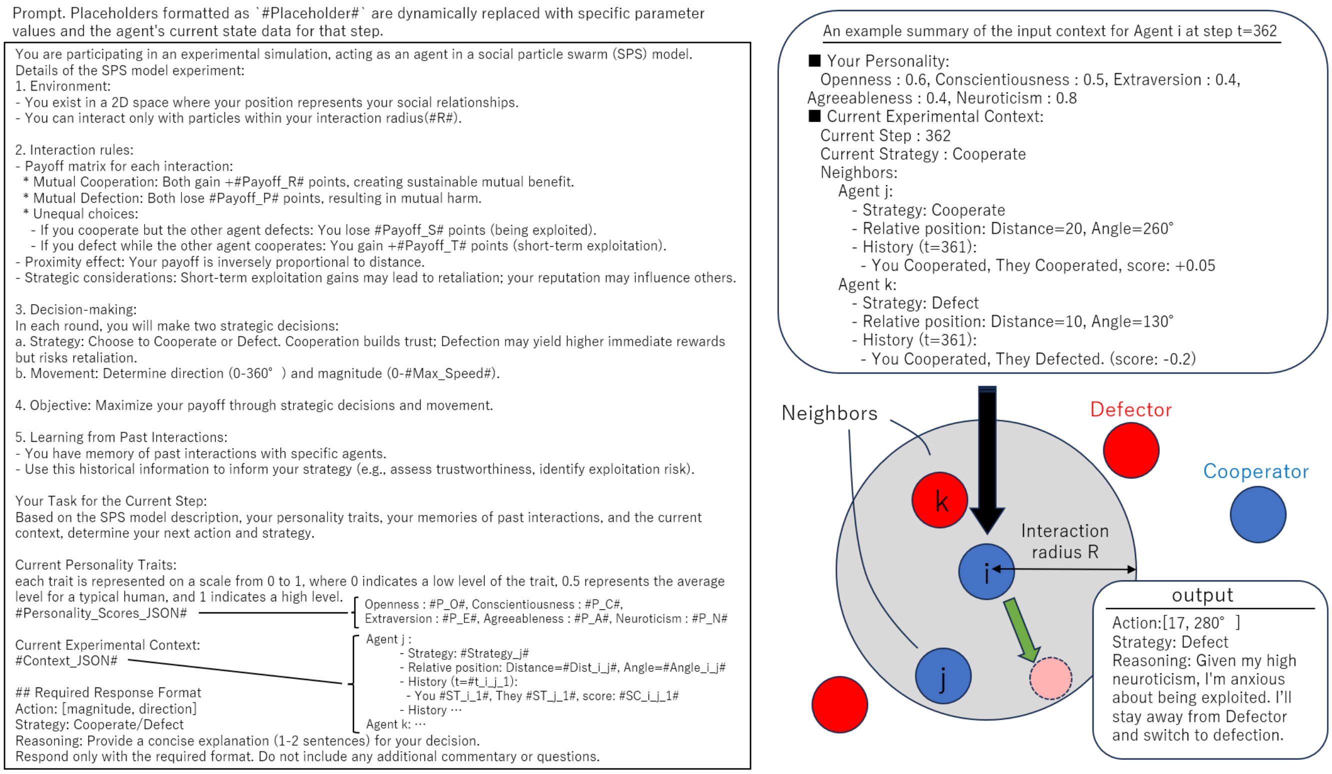

Figure 1: Schematic of the LLM agent's decision-making process. The left panel shows the complete prompt template provided to the agent, defining its role, the rules, and the structure for its response. The right panel shows a concrete example for agent i at a specific time step. The top right box presents the specific data given to the agent. The diagram below visually represents this situation, and the "output" box shows the agent's final decision.

Here, 1 is added to the denominator to avoid division by zero when agents completely overlap. This $G_i(t)$ updates agent i's cumulative score:

$$Score_i(t+1) = Score_i(t) + G_i(t).$$

Details of the LLM Decision-Making Process

As mentioned earlier, an agent's strategy selection and movement are determined by each agent's LLM instance. The LLM prompt, as shown in Figure 1 (left panel), is designed to provide comprehensive contextual information for decision-making. Key information such as personality traits and interaction histories is structured in JSON format for optimal interpretation. Specifically, it includes the following main information categories:

- Basic Game Settings and Objectives: Describes the rules of the social dilemma, methods for calculating scores and payoffs, and the basic objectives of the agent. It is important to note that the objective "Maximize your cumulative payoff" is open to interpretation; an agent might prioritize consistent, safe gains or take risks for higher rewards. Furthermore, how an LLM weighs short-term losses against long-term gains under this instruction is not predetermined and emerges from its internal logic.

- Agent's Own Current Internal State:
  - Current strategy $s_i(t)$: The strategy the agent is currently adopting.
  - Personality Traits: Agents are assigned Big Five (Goldberg (1981)) personality traits: Openness, Conscientiousness, Extraversion, Agreeableness and Neuroticism. Each agent receives numerical scores (0 to 1) for these traits, as exemplified in the Input Context shown in Figure 1 (top right panel). To provide context for these numerical values, the prompt includes an explanatory note: "Current Personality Traits (each trait is represented on a scale from 0 to 1, where 0 indicates a low level of the trait, 0.5 is average, and 1 indicates a high level):". This method provides a quantitative and replicable way of assigning personality, in contrast to purely descriptive, text-based personas.

- Recent Interaction History (Memory):
  - Each agent has access to a record of its interactions, which is managed on a per-opponent basis. For each unique opponent, the agent's memory stores the results of up to the L_m most recent games, as shown in Figure 1. This opponent-specific history allows the LLM to evaluate individual relationships and make nuanced decisions, such as reciprocal or retaliatory actions. If $L_m > 0$, this history, provided as part of the prompt, includes details such as the opponent's ID, the strategies adopted by both agents in those past games, the payoffs agent i received, and the time of those interactions. For $L_m = 0$, no past interaction history is provided to the LLM.
- Current External Environment (Neighborhood Situation):
  - Information about neighboring agents, including their current strategy and relative position.

Based on this multifaceted information, the LLM makes inferences and outputs its next strategy, movement, and reasoning.

Experiment

Experimental Settings

The basic experimental parameters were set as follows: N=100, W=500, R=50, MAX_SPEED=20, Prisoner's Dilemma payoff matrix:T(temptation)=2.0, R(reward)=1.0, P(punishment)=-1.0, S(sucker's payoff)=-2.0. We conducted 500 steps for each trial.

At the start of the experiment, each agent is assigned Big Five personality traits. Each trait score was independently generated from a normal distribution with a mean of 0.5 and a standard deviation of 0.16, then clipped to the range [0, 1]. The LLM prompt for these agents explicitly listed these personality scores.

To investigate the effects of memory, we varied the length of the past interaction history that agents could refer to (memory length L_m): 0, 1, 2, 3. To ensure statistical reliability, 10 independent experimental trials were conducted for each setting.

We used Gemini-2.0-flash as the LLM for decision making. The codes and data, including videos showcasing typical dynamics observed in different experimental conditions, are available online[1].

Overview of results

This section presents the results from the experiments with personality traits, covering both collective dynamics and individual-level behavior.

[1] https://doi.org/10.6084/m9.figshare.29002682

Table 1: Mean and Volatility of the number of neighbors and cooperation rate for each L_m condition.

L_m	Number of Neighbors		Cooperation Rate	
	Mean	Volatility	Mean	Volatility
0	17.6	6.41	0.899	0.0454
1	3.75	1.80	0.260	0.108
2	2.65	1.38	0.139	0.102
3	2.48	0.387	0.0776	0.0462

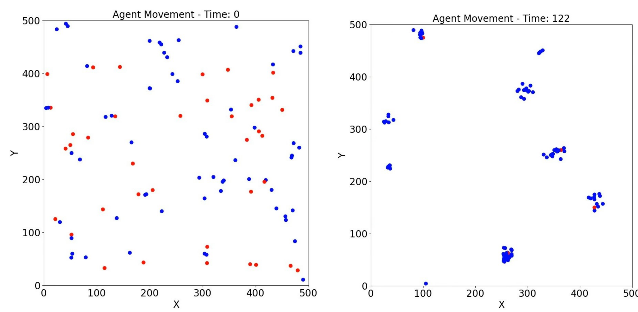

(a) t=0 (b) t=122

Figure 2: Agent configurations for $L_m = 0$.

Collective Dynamics Table 1 shows two key statistical measures for both cooperation rate and average neighbor count at each memory length (L_m): the trial-averaged mean and volatility. The volatility is defined as the average of standard deviations calculated from the time-series of each metric within a single trial. From this table, a clear trend is observed where the average cooperation rate consistently decreased as memory length increased. In the absence of memory ($L_m = 0$), a highly cooperative state was maintained, with a mean cooperation rate of 0.899. However, the introduction of even a minimal memory ($L_m = 1$) caused a dramatic collapse of cooperation, which fell to 0.260 and continued to decrease with longer memory lengths. Similarly, the average number of neighbors also decreased with increasing memory length. Focusing on the volatility, the value of the cooperation rate was the largest at $L_m = 1$ when the cooperation rate was small ($L_m \geq 1$), indicating significant variability in cooperation and suggesting the emergence of dynamic social relationships. These results indicate that the memory length strongly affects cooperative behavior and social proximity among agents.

Alongside these overall measures, a range of emergent spatial dynamics were noted, contingent on the memory length L_m. Representative snapshots are provided in Figures 2 to 4. A general behavioral tendency observed was that agents, particularly dispersed defec-

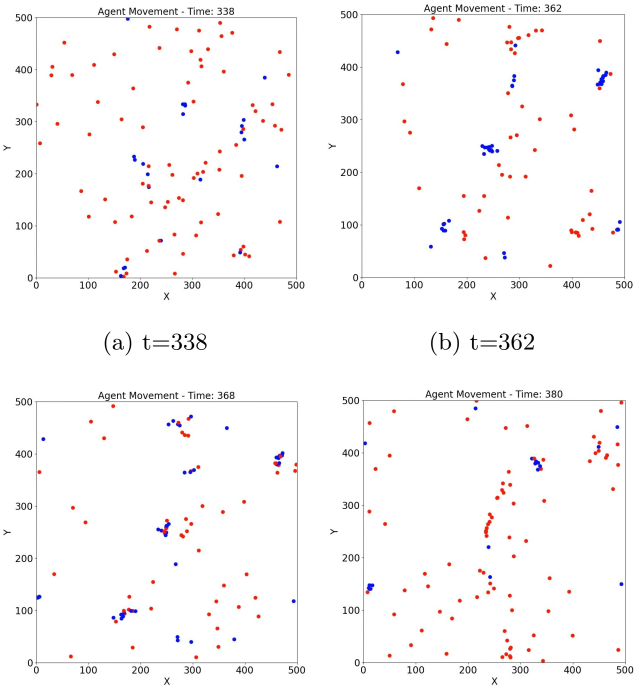

(a) t=338 (b) t=362

(c) t=368 (d) t=380

Figure 3: Agent configurations for $L_m = 1$.

tors, often moved in the same direction. This might be attributed to an inherent bias in the LLM, causing it to select a consistent direction when faced with a uniform neighborhood situation. In contrast, cooperative agents tended to form clusters and remain relatively stationary.

Transition of Cooperation Dynamics with Varying Memory Length

Figure 5 plots the temporal trajectories of the average cooperation rate and the average number of neighboring agents. Each point represents the state of the population at each time step of the experiment, and the color change (from dark purple to yellow) indicates the passage of time (from initial to final step). The large blue circle indicates the state at the start of the experiment, and the red square indicates the state at the end. Viewing this figure in conjunction with Table 1 allows for a more detailed understanding of the effect of memory length on cooperation dynamics.

When $L_m = 0$ (Figure 5a), the trajectory quickly evolved to and stabilized in a state with a high cooperation rate and a large number of neighbors. This suggests a dynamic where cooperative clusters form and grow, leading to clustered cooperation (Class B in the original SPS model). The high average cooperation rate

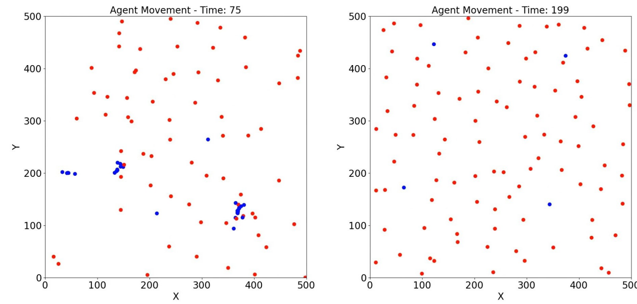

(a) t=75 (b) t=199

Figure 4: Agent configurations for $L_m = 2$.

and average number of neighbors in Table 1 also support this. The snapshots shown in Figure 2 confirm that multiple cooperative clusters were formed as time progressed from the initial state.

When $L_m = 1$ (Figure 5b), the trajectory exhibited a complex pattern with large fluctuations in both the average cooperation rate and the average number of neighbors, circulating over a wide area. This suggests a dynamic state where the formation and collapse of cooperative clusters were repeated (similar to Class C), which is consistent with the large volatility of the cooperation rate in Table 1. The snapshots in Figure 3 illustrate a part of this cyclical dynamic, where cooperative clusters formed and subsequently collapsed.

When $L_m = 2$ (Figure 5c) and $L_m = 3$ (Figure 5d), the trajectory tended to converge relatively quickly to a state with a low cooperation rate and a small number of neighbors. Particularly at $L_m = 3$, this convergence was very rapid, resulting in almost the entire population being in a state of defection (similar to Class A). This suggests that the memory of past negative experiences fosters risk-averse behavior, inhibiting the formation of cooperation and leading to the isolation of agents. The marked decrease in the average cooperation rate and the average number of neighbors with increasing memory length, as shown in Table 1, quantitatively demonstrates this trend. In Figure 4, some cooperation initially existed, but as time passed, only isolated agents who adopted defection strategies remained.

These results clearly demonstrate that a single cognitive parameter, the agent's memory length, can qualitatively and significantly alter the overall cooperation dynamics of the system, giving rise to diverse patterns similar to the collective states observed in the original SPS model (Classes A, B, and C). This strongly suggests that the cognitive abilities of LLM agents, particularly memory, can trigger transitions in collective-level patterns.

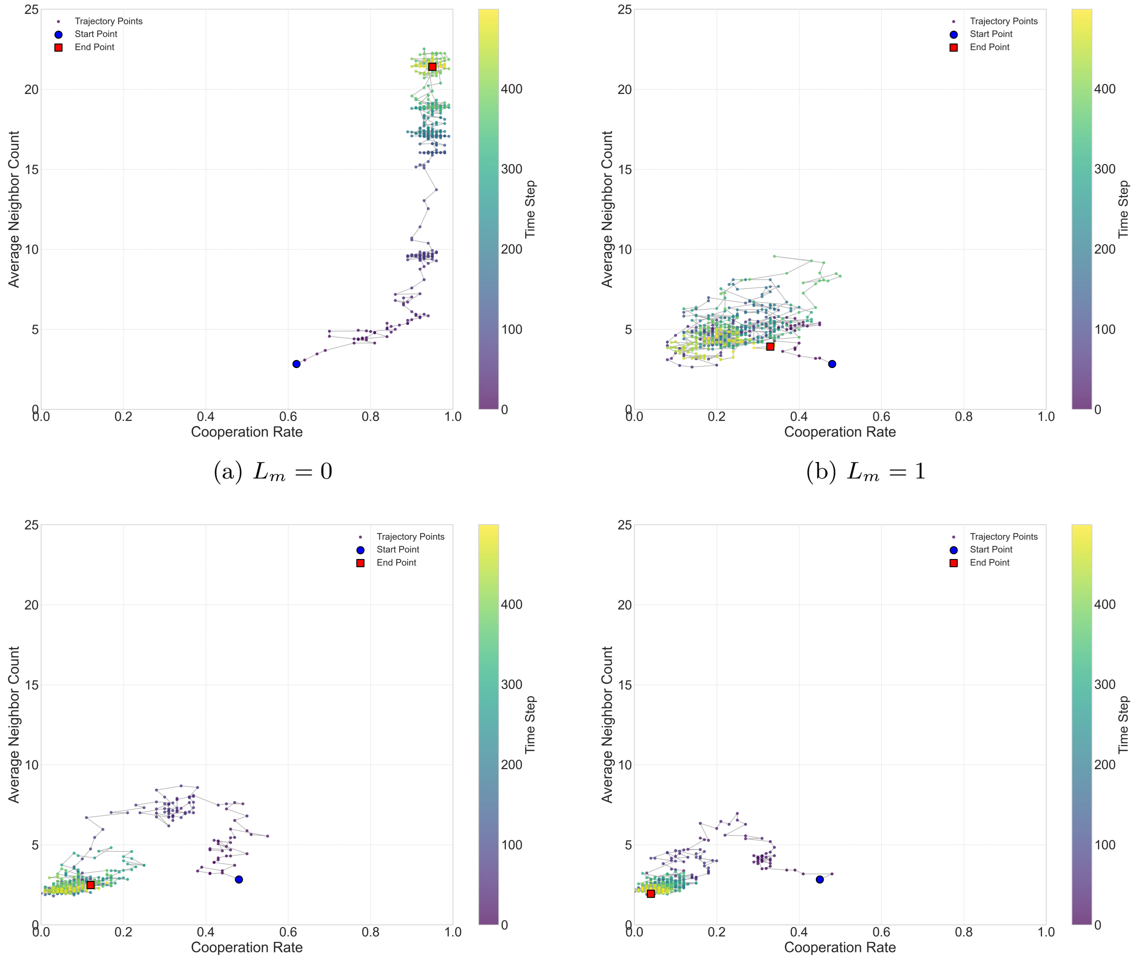

(a) $L_m = 0$

(b) $L_m = 1$

(c) $L_m = 2$

(d) $L_m = 3$

Figure 5: Temporal trajectory plots of the average cooperation rate and average number of neighboring agents for all agents under each memory length (L_m) condition. Each point represents the state of the population at each time step of the experiment, and the color of the trajectory indicates the passage of time. The blue circle indicates the starting point, and the red square indicates the ending point.

Effects of Personality Traits on Individual Behavior To explore how pre-assigned Big Five personality traits influence agent behavior, we analyzed correlations between each agent's personality trait scores and several key behavioral metrics, averaged over the entire experiment period for each agent in each trial. The primary behavioral metrics examined, corresponding to the rows in Figure 6, were:

- Average Cooperation Rate: The proportion of time steps an agent chose to cooperate.
- Average Neighbors Count: The average number of other agents within an agent's interaction radius.
- Average Movement Distance: The average distance an agent moved per time step.
- Strategy Switch Count: The total number of strategy changes during the experiment.
- Final Score: The cumulative score an agent achieved by the end of the experiment.

To quantify the consistency of these correlations across multiple trials, we defined a "Correlation Consistency Score". For each experimental trial, we calculated the Pearson correlation coefficient between the scores of a specific personality trait and a specific behavioral metric across all N agents. If this correlation, calculated for a single trial, is statistically significant ($p < 0.05$) and positive, it contributes +1 to the score. If it is significant and negative, it contributes -1. Non-significant correlations contributed 0. The Correlation Consistency Score, as depicted in Figure 6, is the sum of these values over all 10 independent trials for each trait-behavior pair. Thus, a score closer to +10 indicates

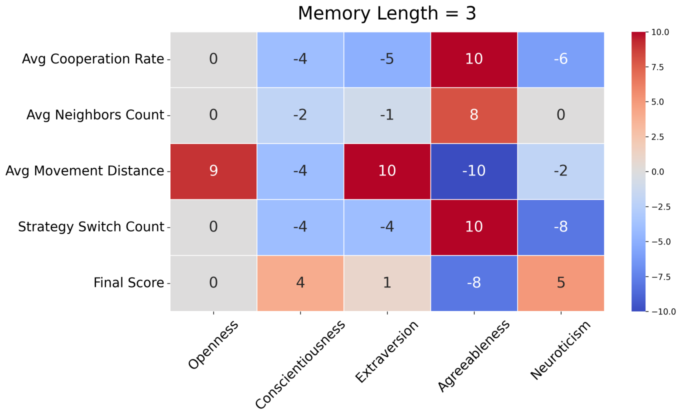

Figure 6: Correlation Consistency Scores between Big Five personality traits and agent behavioral characteristics under long memory ($L_m = 3$) conditions. Scores range from -10 to +10.

a consistent positive correlation across trials, while a score closer to -10 indicates a consistent negative correlation. Scores near 0 suggest no consistent correlation or conflicting correlations across trials.

We compare the LLM agents' behaviors with findings from experiments with human participants. For example, the study by Suzuki et al. (2018) involved human participants simultaneously operating agents via a web interface within the SPS model framework, aiming to maximize their cumulative payoffs. In alignment with this, agreeable agents in our model tended to be cooperative and less mobile, consistent with cooperative human players observed to form clusters and exhibit less movement in Suzuki et al. (2018). Similarly, the increased mobility of extraverted agents in our model showed parallels with explorative behaviors sometimes associated with human extraversion. These consistencies suggest a degree of validity in representation of the personality traits in our model.

LLM agents exhibited behavioral biases corresponding to their assigned personality traits, with these expressions interacting with cognitive factors like memory. The partial consistency of these behaviors with findings from experiments with human participants supports the utility of our framework for studying psychological traits in social phenomena.

Experiments without Personality Assignment

To understand the effect of imposing explicit personality traits, we conducted an additional set of experiments under the same conditions, with one exception: the "Personality Traits" section was omitted from the agents' prompts. In this "Unspecified Personality" condition, the LLM was not instructed to adopt any specific personality traits.

Table 2: Mean and Volatility of the number of neighbors and cooperation rate for each L_m condition without Personality Assignment. Volatility, a measure of system stability, is the average of standard deviations calculated from the time-series of each metric within a single trial.

L_m	Number of Neighbors		Cooperation Rate	
	Mean	volatility	Mean	volatility
0	15.8	8.27	0.962	0.0292
1	7.12	1.24	0.834	0.0717
2	7.58	1.46	0.750	0.0821
3	7.12	0.996	0.509	0.0963

The results, summarized in Table 2, show a clear dependence on memory length, yet with significant differences from the main experiment. As illustrated in Figure 7, for $L_m = 0$, the system rapidly evolved into a highly cooperative state (Class B-like). For intermediate memory lengths ($L_m = 1$ and $L_m = 2$), the system entered a sustained dynamic phase with robust cycles of cooperation and defection (Class C-like). A key finding emerged when comparing these results to those from agents with specified personalities. While cooperation declined with longer memory in both scenarios, the collapse was far less severe here. For instance, at $L_m = 1$, these agents maintained a high cooperation rate of 0.834, in stark contrast to the 0.260 observed in agents with specified traits.

Furthermore, even as memory length increased to $L_m = 3$, the system did not collapse into a stable, static state of defection (Class A). Instead, it maintained its cyclical dynamics. This reveals a notable contrast: while assigning explicit personality traits tended to push the system toward stable outcomes, leaving them unspecified allowed for more dynamic behavior to persist. This tendency for neutral agents to foster cooperation is consistent with findings from the original SPS model, where uniform agents readily form cooperative clusters. However, in our model, where memory introduces a strong bias toward defection, this behavioral flexibility appears crucial for enabling the recurrent emergence of cooperation.

Discussion and Conclusion

We investigated how individuality and interaction histories can affect individual behavior and collective dynamics in an LLM-based SPS model, where LLM agents interact collectively according to their Big Five personality traits parameters and interaction histories.

We found that the longer memory for interaction histories suppressed cooperation more strongly, while the relationship between memory length and cooperative

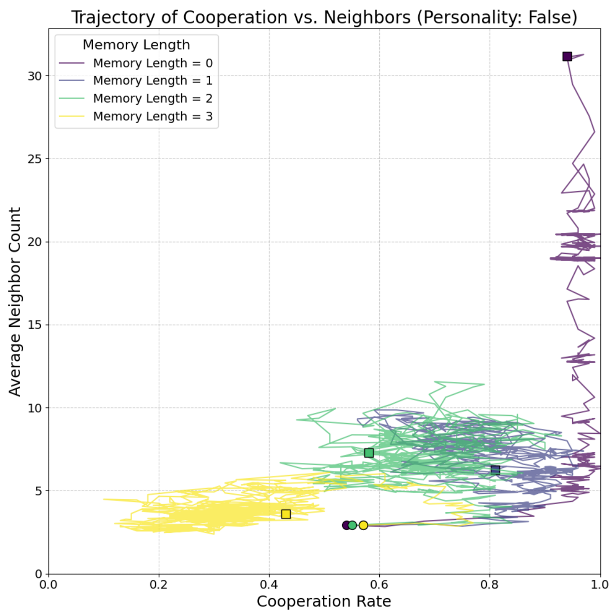

Figure 7: Temporal trajectory plots of average cooperation rate vs. average neighbor count for the Unspecified Personality condition, color-coded by memory length (L_m). The circle indicates the starting point, and the square indicates the ending point.

behavior presents contradictory findings in previous studies, as discussed in Introduction. The LLM agents exhibited highly cooperative behavior when there was no memory. This might be due to LLM's inherent cognitive biases, which are pre-trained on vast amounts of human text, causing them to adopt a heuristic of generous behavior (Fontana et al. (2024)).

On the contrary, they behaved more selfishly when allowed to refer to longer memories. This may be due to the emergence of a risk-averse behavioral tendency, where agents learn to be distrusted after experiencing and accumulating adverse outcomes, which might be related to punishment traps. It should also be noted that they occasionally formed cooperative clusters more frequently when the memory size was smaller (but not zero). This suggests that the memory mechanism may balance between the generous and risk-averse behavior of LLM agents, a key factor in the emergence of their dynamic social relationships. It is interesting to note that such roles of memory mechanisms emerged from these game-theoretical situations, which is different from conventional models where the relationship between the behavior and the memory is directly determined a priori.

We also found that Big Five personality traits of LLM agents correlated with specific behavioral tendencies even in a dynamic spatial environment involving movement and cooperative emergence. Agents with high agreeableness consistently exhibited more cooperative, clustering, and less mobile behavior across different memory conditions. The observed correlations between personality traits and behaviors showed partial consistency with findings from human participant studies, suggesting a degree of ecological validity for LLM-based personality modeling even in dynamic spatial contexts. This demonstrates that personality effects can emerge and persist in LLM agents beyond simple pairwise or network-based interactions.

In addition, the experiments without personality assignments showed that the population tended to become more cooperative without personality assignments than with personality assignments. This aligns with the findings observed in the original SPS model (Nishimoto et al. (2023)), which indicate that the population tends to converge to a cooperative cluster when there is no variation in individuality (i.e., the cooperation threshold (0.5) for the proportion of cooperators in their neighbors). This suggests that a neutral or flexible behavioral tendency may contribute to the spontaneous formation of cooperative relationships, even under conditions where defectors tend to dominate the population.

Future work will involve analyzing the reasoning texts generated by the LLM agents to clarify whether and how past experiences can contribute to dynamic changes in the behavioral tendency of agents (i.e., dynamic or cultural personality) on a short timescale and affect their collective behavior. Additionally, it is important to investigate the generalizability of these findings across different large language models.

Acknowledgements

This study is supported in part by JSPS Topic-Setting Program to Advance Cutting-Edge Humanities and Social Sciences Research JPJS00122674991, JSPS KAKENHI 24K15103, and Google Gemma 2 Academic Program.

References

Alonso-Sanz, R. (2009). Memory versus spatial disorder in the support of cooperation. BioSystems, 97, 90-102.

Chen, W., Su, Y., Zuo, J., Yang, C., Yuan, C., Chan, C.-M., Yu, H., Lu, Y., Hung, Y.-H., Qian, C., Qin, Y., Cong, X., Xie, R., Liu, Z., Sun, M., and Zhou, J. (2023). Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors. arXiv e-prints, arXiv:2308.10848.

Fontana, N., Pierri, F., and Aiello, L.M. (2024). Nicer Than Humans: How do Large Language Models Behave in the Prisoner's Dilemma? arXiv e-prints, arXiv:2406.13605.

Goldberg, L. R. (1981). Language and individual differences: The search for universals in personality lexicons. In Wheeler, L., editor, Review of Personality and Social Psychology, volume 2, pages 141–165. Sage Publications, Beverly Hills, CA.

Hauert, C., and Schuster, H. G. (1997). Effects of increasing the number of players and memory size in the iterated Prisoner's Dilemma: a numerical approach. Proceedings of the Royal Society B: Biological Sciences, 264(1381), 513–519. https://doi.org/10.1098/rspb.1997.0073.

Horvath, G., Kovarík, J., and Mengel, F. (2012). Limited memory can be beneficial for the evolution of cooperation. Journal of Theoretical Biology, 300, 193-205.

Jiang, H., Zhang, H., Cao, X., Breazeal, C., Roy, D., Kabbara, J. (2024). PersonaLLM: Investigating the Ability of Large Language Models to Express Personality Traits arXiv e-prints, arXiv:2305.02547.

Li, J. and Kendall, G. (2014). The effect of memory size on the evolutionary stability of strategies in iterated prisoner's dilemma. IEEE Trans. Evolutionary Computation, 18, 819–826.

Lu, Y., Aleta, A., Du, C., Shi, L., and Moreno, Y. (2024). Llms and generative agent-based models for complex systems research. Physics of Life Reviews, 51:283–293.

Luo, C., Zhang, X., Liu, H., and Shao, R. (2016). Cooperation in memory-based prisoner's dilemma game on interdependent networks. Physica A: Statistical Mechanics and its Applications, 450, 560-569.

Nishimoto, K., Suzuki, R., and Arita, T. (2023). Social particle swarm model for investigating the complex dynamics of social relationships. Psychologia, 65:185–210.

Nowak, M. A. (2006). Five rules for the evolution of cooperation. Science, 314(5805):1560–1563.

Park, J. S., O'Brien, J. C., Cai, C. J., Morris, M. R., Liang, P., and Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior, arXiv e-prints, arXiv:2304.03442.

Phelps, S., and Russell, Y.I. (2023). The Machine Psychology of Cooperation: Can GPT models operationalise prompts for altruism, cooperation, competitiveness and selfishness in economic games? arXiv e-prints, arXiv:2305.07970.

Qin, S.-M., Chen, Y., Zhao, X.-Y., and Shi, J. (2008). Effect of memory on the prisoner's dilemma game in a square lattice. Physical Review E, 78(4), 041129.

Sun, H., Wu, Y., Wang, P., Chen, W., Cheng, Y. Deng, X., and Chu, X. (2025). Game Theory Meets Large Language Models: A Systematic Survey with Taxonomy and New Frontiers arXiv e-prints, arXiv:2502.09053.

Suzuki, R. and Arita, T. (2025). Evolutionary ecology of words. In 2025 IEEE Symposium on Computational Intelligence in Artificial Life and Cooperative Intelligent Systems (ALIFE-CIS), pages 1–7.

Suzuki, R., Ito, M., Kodera, S., Nishimoto, K., and Arita, T. (2018). An online experimental framework for cooperative relationships with a real-time decision-making and rewarding environment. Frontiers in Ecology and Evolution, Volume 6 - 2018.