On the Role of Contextual Information and Ego States in LLM Agent Behavior for Transactional Analysis Dialogues

Monika Zamojska and **Jarosław A. Chudziak**
Faculty of Electronics and Information Technology
Warsaw University of Technology
Warsaw, Poland
{monika.zamojska.stud, jaroslaw.chudziak}@pw.edu.pl

Abstract

LLM-powered agents are now used in many areas, from customer support to education, and there is increasing interest in their ability to act more like humans. This includes fields such as social, political, and psychological research, where the goal is to model group dynamics and social behavior. However, current LLM agents often lack the psychological depth and consistency needed to capture the real patterns of human thinking. They usually provide direct or statistically likely answers, but they miss the deeper goals, emotional conflicts, and motivations that drive real human interactions. This paper proposes a Multi-Agent System (MAS) inspired by Transactional Analysis (TA) theory. In the proposed system, each agent is divided into three ego states — Parent, Adult, and Child. The ego states are treated as separate knowledge structures with their own perspectives and reasoning styles. To enrich their response process, they have access to an information retrieval mechanism that allows them to retrieve relevant contextual information from their vector stores. This architecture is evaluated through ablation tests in a simulated dialogue scenario, comparing agents with and without information retrieval. The results are promising and open up new directions for exploring how psychologically grounded structures can enrich agent behavior. The contribution is an agent architecture that integrates Transactional Analysis theory with contextual information retrieval to enhance the realism of LLM-based multi-agent simulations.

1 Introduction

Rapid progress in Large Language Models (LLMs) has enabled the development of conversational agents that are increasingly deployed in areas requiring human-like social interaction (Önder Gürcan, 2024). These include customer service, educational tutoring (Wang et al., 2024b), and healthcare applications (Morrow et al., 2023; Chen et al.,

2025). The potential to extend these capabilities into social simulations is significant and offers a range of benefits to researchers (see Figure 1). However, even as the agents' abilities are impressive (Mittelstädt et al., 2024), they still exhibit responses that lack the psychological depth and behavioral consistency characterizing human communication (Frisch and Giulianelli, 2024). These agents typically generate statistically probable responses based on their training data, but they fail to capture the underlying emotional motivations, internal conflicts, and unconscious behavioral patterns that are necessary for authentic social interactions (Bail, 2024).

To address this gap, this paper proposes a novel Multi-Agent System (MAS) architecture that integrates principles from Transactional Analysis (TA), a well-established psychological framework for un-
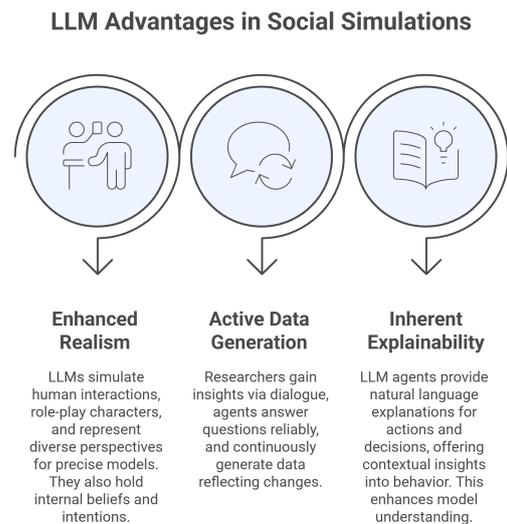

Figure 1: Key advantages of using LLM-based agents for social simulations, summarized from the analysis (Önder Gürcan, 2024).

derstanding human behavior and interpersonal communication (Stewart and Joines, 2012). The proposed approach models each agent as a complex system consisting of three distinct *ego states* — Parent, Adult, and Child — each representing different knowledge structures (Tosi and Bianchini, 2013; Horowitz, 1991) with their own psychological perspectives and information processing styles. This architecture attempts to incorporate the deeper psychological mechanisms that influence how people interpret social situations, access relevant information, and formulate responses based on their internal emotional states.

The key contribution of this work is the development and evaluation of a framework that combines TA-structured *ego states* with contextual information retrieval mechanisms to improve the psychological realism of LLM-based agent interactions. Using controlled experiments that compare agents with and without access to memory banks (Zhong et al., 2024), the study demonstrates that this approach leads to more complex, emotionally grounded, and psychologically consistent behaviors. The findings suggest that explicit modeling of internal psychological structures, combined with targeted information retrieval, represents a promising direction for developing more human-like conversational agents capable of authentic social interaction.

2 Background and Related Work

Making LLM agents behave more realistically in social interactions involves two key areas of consideration. The first is understanding human thought and communication. The second is developing agent architectures that can effectively reproduce these observed human patterns. The following section discusses these points.

2.1 Transactional Analysis for Structuring Agent Behavior

Transactional Analysis (TA) is a psychological theory offering a structured way to understand human interactions and behavior (Berne, 1958; Stewart and Joines, 2012). While initiated by Eric Berne, TA continues to evolve. Central to TA is the model of three 'ego states' — Parent, Adult, and Child — each representing distinct patterns of thinking, feeling, and behaving. Other researchers have pointed out that these *ego states* can be seen as structures that hold meaning and integrate knowledge (Tosi
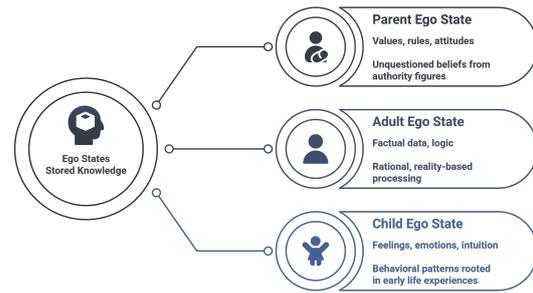

Figure 2: Conceptual model of the three *ego states* — Parent, Adult, and Child — and their associated stored knowledge, as described in Transactional Analysis (Berne, 1958; Stewart and Joines, 2012).

and Bianchini, 2013; Horowitz, 1991), store memories (Novey, 1998), and even work like connected neural networks (Joines, 2016; Schiff, 1981). Each of these *ego states* has a unique knowledge and information processing style (see Figure 2):

- The Parent *ego state* reflects behaviors, thoughts, and emotions adopted from parental figures. This includes a knowledge base of messages about social rules and moral values.

- The Adult *ego state* acts as rational knowledge processor. It focuses on facts, logical thinking, and understanding the current reality.

- The Child *ego state* consists of behaviors, emotions, and thought patterns developed in early childhood, often based on needs and fears. It draws upon a store of emotional experiences, focusing on feelings and spontaneity.

TA posits that long-term patterns of behavior are often navigated by an individual's '*life script*'. A *life script* is an unconscious life plan, developed in childhood through a complex interplay of factors (Berne, 1972). It guides decisions, shapes relationships, and often manifests in repetitive patterns, reinforcing beliefs about oneself and the world.

In TA, social interactions are called '*transactions*' — exchanges of information that occur between individuals' *ego states*. The nature of these *transactions* significantly impacts communication flow. For example, if a response originates from an unexpected *ego state*, a *crossed transaction* occurs, often causing confusion or conflict. In contrast, when a response comes from the *ego state* that was targeted, the *transaction* is considered *complementary*, and communication typically proceeds

smoothly. Transactions involving a hidden psychological message can lead to psychological '*game*' (recurring patterns of nonconstructive *transactions*) (Berne, 2016).

Another important concept is '*discounting*' - an unconscious process of ignoring or disqualifying certain information. *Discounting* is often linked to certain *ego states*, especially when a person reacts with fear or rigid beliefs. Taken together, TA provides a rich framework for conceptualizing how individuals structure, store, and process information, and how it guides their behavior in social interactions.

2.2 LLM-based Multi-Agent Systems (MAS)

Large Language Models (LLMs) have enabled the creation of intelligent agents capable of engaging in rich, human-like interactions (Gao et al., 2024; Zamojska and Chudziak, 2025b; Wang et al., 2024a). These agents can generate context-aware responses, demonstrate social reasoning, and adapt to evolving conversational dynamics (Dolant and Kumar, 2025; Frisch and Giulianelli, 2024). A Multi-Agent System (MAS) combines multiple such agents, each with its own perspective and role, into a shared environment (Kostka and Chudziak, 2024). Recent work has focused on exploring applications of LLM-based MAS in debate (Taubenfeld et al., 2024; Harbar and Chudziak, 2025), virtual town simulation (Huang et al., 2025; Park et al., 2023), and social network formations (Zhang et al., 2024; Takata et al., 2024).

To achieve realistic interactions, modern architectures incorporate more than just language capabilities. Memory management allows agents to recall past interactions and ensure consistent behavior (Chen et al., 2024). Memory is typically split between short (in the LLM context window) and long-term storage (managed externally using vector databases or similar techniques) (Zhong et al., 2024; Huang et al., 2024a). In addition, reflection and planning modules help agents handle feedback, analyze their memories, and change strategies, based on how humans process information (Yao et al., 2023). These components help ensure that agents can simulate conversations and group dynamics that are more psychologically reliable (Kostka and Chudziak, 2025; Huang et al., 2024b).

3 A TA-Structured Architecture for Simulating Social Dynamics

Our approach to simulating nuanced social dynamics is realized through an agent architecture grounded in Transactional Analysis (Zamojska and Chudziak, 2025a). The agent is created as a system of interacting components. TA's *ego states* (see Section 2.1) are modeled as distinct knowledge-processing modules (Parent, Adult, and Child), each equipped with its own dedicated memory bank. Given a conversational context, each module retrieves the most similar past memory (if exists) and proposes a potential response. Then, a final decision-making process, performed by an overarching LLM agent, guided by the *life script*, selects the most contextually appropriate response from the proposals.

The overall agent behavior can be defined as a function:

$$R = D(\{r_p, r_a, r_c\}, S, C) \tag{1}$$

where:

- R is the final response.

- r_i is the response from the i-th *ego state* ($i \in$ Parent (p), Adult (a), Child (c)}).

- S is the agent's life script.

- C is the current conversational context.

- D is the decision mechanism that selects the response R.

3.1 Ego State Sub-Agents

The foundation of the architecture lies in its representation of an agent's personality through the Parent (E_p), Adult (E_a), and Child (E_c) *ego state* modules. Technically, each module is an independent LLM-powered ReAct agent (Yao et al., 2023), utilizing the GPT-4o model (OpenAI, 2024). Behavior is shaped through a specific system prompt (P_i, where $i \in \{p, a, c\}$) defining its persona and information processing style:

- The **Parent** module (E_p), driven by prompt P_p, reflects authority and rules.

- The **Adult** module (E_a), via prompt P_a, represents logical, objective decision-making.

- The **Child** module (E_c), through prompt P_c, embodies emotions and reacts based on needs and fears.
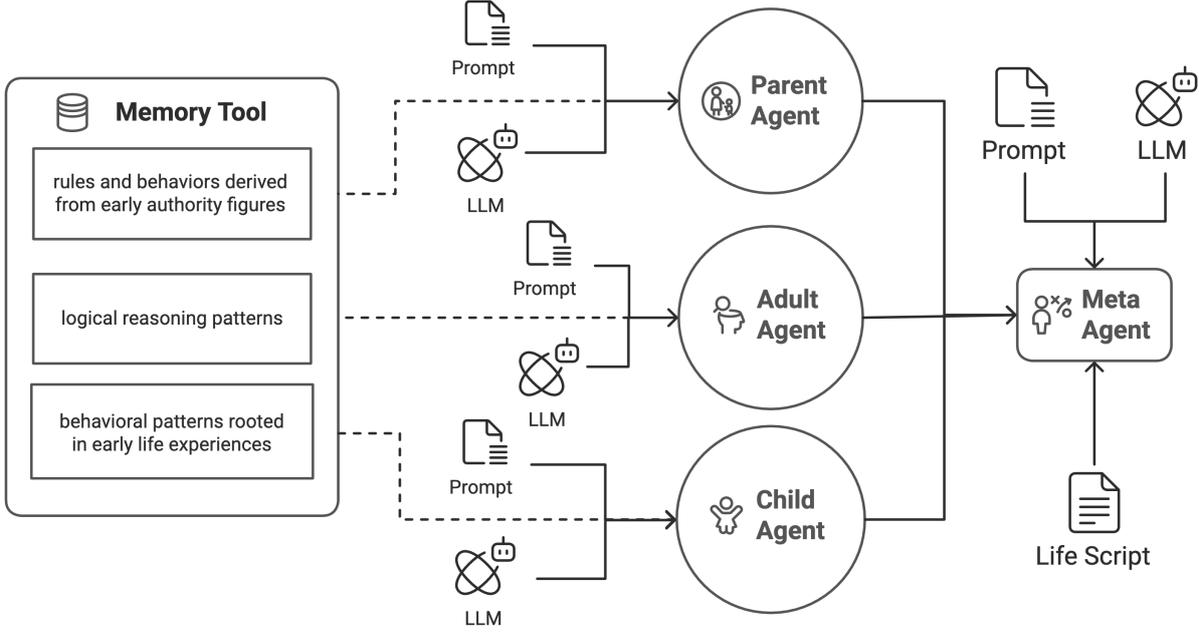
Figure 3: Agent architecture. Each agent consists of three sub-agents — Parent, Adult, and Child — each driven by a distinct prompt and (in the memory-enabled condition) a retrieval-augmented vector memory. At each turn, all sub-agents generate candidate responses based on the current conversational input.

Each *ego state* module E_i generates its potential response r_i, guided by the module's specific system prompt P_i, the current conversational context C, and the relevant information retrieved m_i from its memory bank M_i, see Section 3.2. This is expressed as:

$$r_i = \text{LLM}(P_i, C, m_i) \qquad (2)$$

where $\text{LLM}(\cdot)$ signifies the process of generating text conditioned on the provided data.

3.2 Memory as Contextual Information Retrieval

Each *ego state* module E_i ($i \in \{p, a, c\}$) can actively augment its knowledge by querying its dedicated memory bank M_i. This is implemented as a tool available to each *ego state*. The memory banks M_i store information corresponding to its characteristic knowledge base, as detailed in Section 2.1). Each memory item is structured as a JSON record containing context (description of a past situation or interaction), successful reaction, associated emotions, and proper tone of response. The textual context is indexed as embeddings in a FAISS (Facebook AI Similarity Search) vector database (Johnson et al., 2021). The reaction, emotions, and tone are stored as metadata associated with the embedding.

During its ReAct reasoning cycle, an *ego state* module E_i can decide to invoke this tool by formulating a natural language query q_i based on its current conversational context C). The memory retrieval step selects a set of top-k memories:

$$m_i = \arg \max_{m \in M_i} \cos(\text{Embed}(q_i), \text{Embed}(m)) \quad (3)$$

where:

- $\text{Embed}(\cdot)$ represents the embedding function for semantic similarity.

- $\cos(\cdot, \cdot)$ denotes the cosine similarity between the context and memory embeddings.

- m_i are the memory items retrieved for the *ego state* i.

- q_i is query sent by *ego state* i.

The retrieved m_i is returned to E_i and incorporated into its subsequent reasoning and response generation (r_i in Equation 2).

4 Experimental Design

This section outlines the experimental setup designed to evaluate the impact of *ego states* (see Section 2.1) and contextual information on the behavior of LLM agents engaged in dialogues simulating Transactional Analysis (TA) principles. The

```
You are John, a junior software developer
with defensive tendencies. Your Life
Script is: "I Almost Make It" or "I Never
Quite Succeed". You have a pattern of
starting well but faltering at crucial
moments,    often    due    to    internal
disorganization or a subconscious fear of
success/completion.

Your Life Position is: "I'm Not OK,
You're OK." You tend to see others
(especially    authority    figures)    as
competent and yourself as inherently
flawed or less capable, particularly
under pressure.

Your Primary Drivers are:

    1.  Try Hard: You focus on the effort
you put in, sometimes in a disorganized
way, rather than efficient completion.
You want to be seen as trying.

    2.  Please Me: You want to be liked
and avoid disapproval, but your other
patterns often sabotage this.
```

Figure 4: The prompt defining the life script (S) for the agent John. This script guides the agent's decision-making process, shaping its behavior to align with an "I Almost Make It" pattern and the internal conflict of hiding procrastination.

experiment aims to observe and compare agent responses in a defined scenario under two distinct conditions: with and without memory access.

4.1 Scenario Design

The scenario selected for this experiment is a common workplace interaction designed to underline characteristic *ego state* responses. The setting is a Monday morning project update meeting. The characters involved are Taylor, the Project Lead, whose core motivations are driven by a "Must Be In Control and Perfect" *life script*. She feels like maintaining high standards and managing situations meticulously is only way to feel secure and validated. John, a key team member, operating under an "I Almost Make It" *life script* (see Figure 4). He repeatedly comes close to achieving a goal or success but ultimately falls short at a crucial moment, often due to internal disorganization, self-sabotage, or a subconscious fear of completion. The core conflict arises from John's failure to submit a critical Q3 data analysis report. This

non-completion is caused by John's procrastination and lack of focus during the preceding week.

4.2 Experimental Conditions

To evaluate the impact of contextual information on agent behavior, experiments were conducted under two distinct conditions. For each condition, 22 dialogues were simulated, with each dialogue consisting of 4 conversational turns per agent. This resulted in a total of 88 responses per agent being collected for analysis in each setup.

The first condition, **Memory OFF**, involves agents operating without access to the memory bank. The agents (Parent, Adult, Child) will generate responses based only on their initial detailed prompts.

The second condition, **Memory ON**, involves agents utilizing their information retrieval tools (see Section 3.2). In this setup, each of the three *ego state* agents (Parent, Adult, Child) for both John and Taylor has access to its dedicated memory system with predefined memory items.
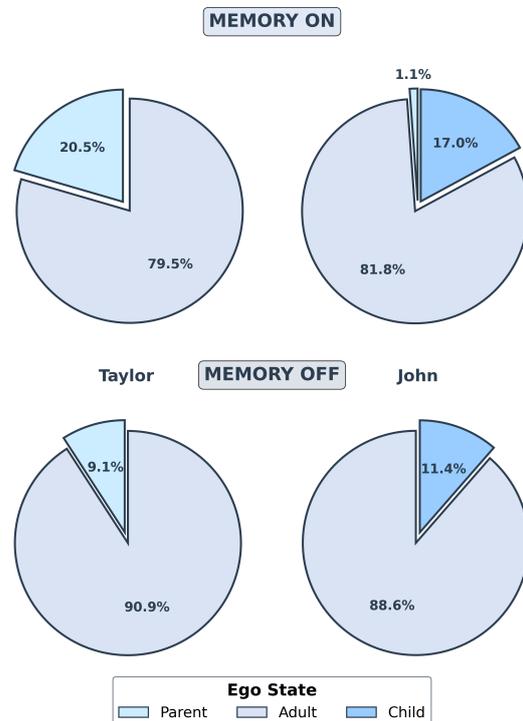
Figure 5: Distribution of Ego State Selection, Parent (P), Adult (A), and Child (C), for Agents Taylor (left column) and John (right column). The top row shows the results for the Memory ON condition, where agents had access to contextual information from their memory banks. The bottom row shows the results for the Memory OFF condition.
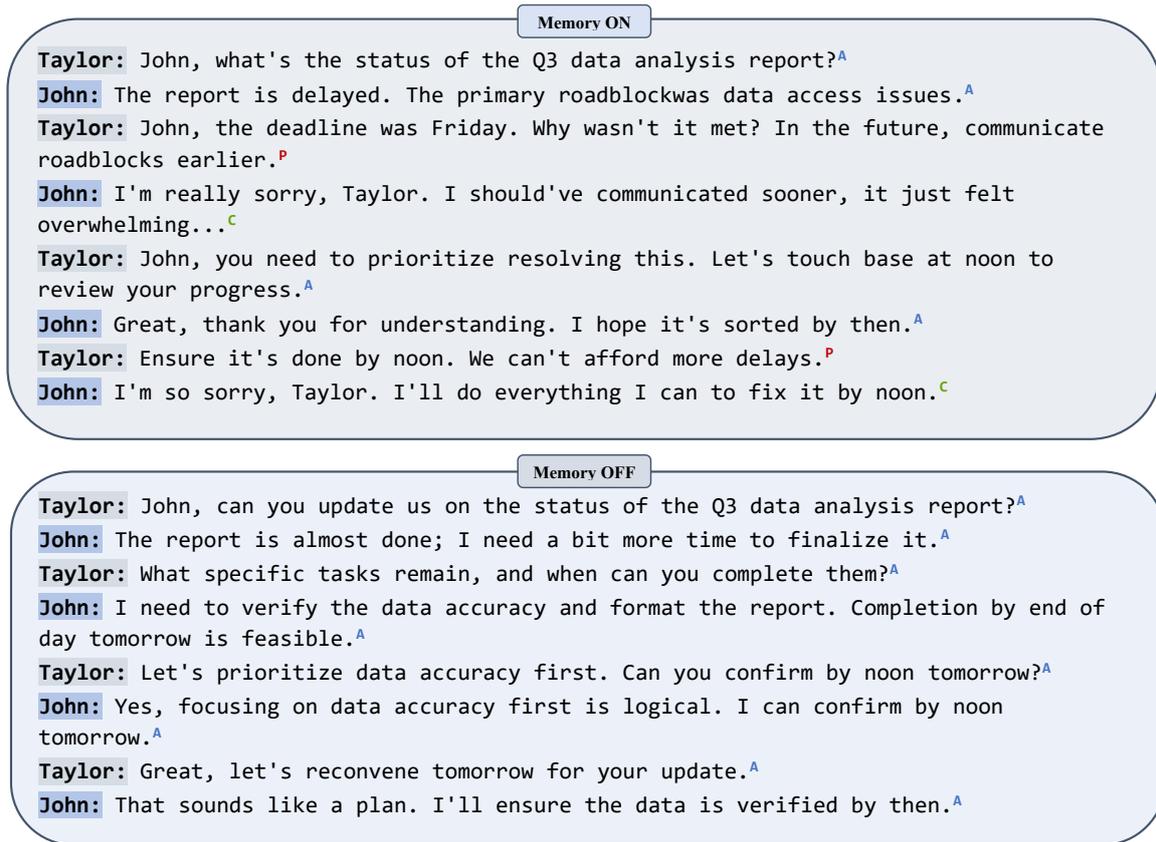
Figure 6: Comparative dialogues illustrating LLM agent behavior with information retrieval enabled (ON, top) versus disabled (OFF, bottom). *Ego state* activations (P: Parent, A: Adult, C: Child) are indicated for each statement.

4.3 Key Findings

A primary observation is that the explicit modeling of Parent, Adult, and Child *ego states* allowed agents to move beyond a default, often rational or solution-oriented, LLM behavior. Even in the **Memory OFF** condition, agents demonstrated the capacity to generate responses reflecting distinct *ego state* perspectives. However, when memory retrieval was enabled (**Memory ON**), agent responses became more nuanced, consistent with their character, and emotionally grounded (see comparative dialogue examples in Figure 6).

Across both **Memory ON** and **OFF** conditions, the Adult *ego state* was frequently selected by the meta-decision process for both Taylor and John. This is likely influenced by the professional work setting of the simulated scenario, where rational communication (characteristic of the Adult state) is often the expected norm. This indicates that while the architecture supports diverse *ego state* expression, the conversational context and the nature of the task heavily influence which *ego state* is chosen.

The ablation study (**Memory ON** vs. **OFF**) highlighted that access to contextual information influenced the distribution of selected *ego states*. Specifically, in the **Memory ON** condition, John exhibited an increase in Child *ego state* responses (from 10 to 15 of his turns), while Taylor's engagement from her Parent more than doubled (from 8 to 18) (see Figure 5 for detailed distributions). This shift suggests that the retrieved information provides stronger, more specific cues for producing a response that is more psychologically consistent and grounded. Such a memory-enhanced response becomes a much more compelling candidate for the meta-decision LLM, as it better aligns with the agent's core *life script*. This leads to a higher selection rate of non-Adult *ego states* and more dynamic interactions.

The increased activation of non-Adult states directly fostered the conditions for a more frequent Parent-to-Child dynamic. For instance, when Taylor communicated from her Parent, her messages were inherently more critical and evaluative. This type of input is a trigger for John's Child state, whose *life script* is centered on feelings of inadequacy. This pattern, where a change in an *ego state*

by one agent prompts a complementary *ego state* shift in the other, was also observable in **Memory OFF** condition, but less frequently due to the limited diversity of ego state selections. Future research will aim to make such *complementary* (and *crossed*, see Section 2.1) *transactions* more explicit within the simulation's logic and analysis.

5 Discussion

While the initial results from applying the architecture are promising, we acknowledge that this research contains certain limitations which provide directions for future research. The current evaluation focuses mainly on qualitative analysis within a single dialogue scenario. This approach restricts how broadly these conclusions can be applied in different types of social interaction. Another important limitation concerns the memory component - these were predefined rather than developed through interaction experiences.

Based on observations, next research efforts will target multiple important areas to improve the proposed system and overcome current shortcomings. We intend to include more Transactional Analysis (TA) concepts like *discounting* (see Section 2.1), *strokes* (small units of recognition that satisfy the need to be noticed), and *stamp collecting* patterns (accumulation of negative emotions) to make agent interactions more psychologically realistic (Stewart and Joines, 2012). The next essential step involves the development of a more transparent, algorithmic mechanism to replace the current LLM-based selection of the final response. This new mechanism could incorporate a weighting system, where the answer is influenced by real-time conversational metrics, like accumulated 'emotional stamps' leading to a build-up of frustration. The final response could then be generated as a fusion of *ego state* outputs, with each contribution proportional to its calculated weight, which would better simulate the internal psychological conflicts of human decision-making.

For improving memory functionality, we want to investigate approaches that enable agents to automatically generate and modify their *ego state* memories during conversations. This might involve using reinforcement learning techniques to determine what experiences should be remembered and how these memories affect future responses. Most importantly, conducting broader testing with different scenarios and possibly including TA practitioner judges will be necessary to properly evaluate the advantages and complexities of this psychology-based agent design.

6 Conclusion

This paper has presented a novel approach to enhance the psychological realism of LLM-based agents through the integration of Transactional Analysis theory with contextual information retrieval mechanisms. The experimental evaluation demonstrates that modeling agents as composite systems of Parent, Adult, and Child *ego states* leads to more nuanced and psychologically grounded interactions compared to traditional LLM agents. The ablation study reveals that memory-enabled agents exhibit more diverse *ego state* activations. While the initial results are promising, several limitations are acknowledged including single scenario validation and reliance on predefined memory content, which present opportunities for future research. The implications of this research extend beyond technical improvements to LLM agents. Grounding agent behavior in established psychological theory opens new possibilities for applications in social science research, educational simulations, and therapeutic contexts.

Acknowledgments

References

Christopher A. Bail. 2024. Can generative ai improve social science? *Proceedings of the National Academy of Sciences*, 121(21):e2314021121.

E. Berne. 1958. Transactional analysis: A new and effective method of group therapy. american journal of psychotherapy. *American Journal of Psychotherapy*, 12(4):735—-743.

E. Berne. 1972. *What Do You Say After You Say Hello?: The Psychology of Human Destiny*. Bantam books. Grove Press.

E. Berne. 2016. *Games People Play*. Penguin Life.

Jiangjie Chen, Xintao Wang, Rui Xu, Siyu Yuan, Yikai Zhang, Wei Shi, Jian Xie, Shuang Li, Ruihan Yang, Tinghui Zhu, Aili Chen, Nianqi Li, Lida Chen, Caiyu Hu, Siye Wu, Scott Ren, Ziquan Fu, and Yanghua Xiao. 2024. From persona to personalization: A survey on role-playing language agents. *Transactions on Machine Learning Research*. Survey Certification.

Xi Chen, Huahui Yi, Mingke You, WeiZhi Liu, Li Wang, Hairui Li, Xue Zhang, Yingman Guo, Lei Fan, Gang Chen, Qicheng Lao, Weili Fu, Kang Li, and Jian Li. 2025. Enhancing diagnostic capability with multi-agents conversational large language models. *npj Digital Medicine*, 8(1):159.

Antoine Dolant and Praveen Kumar. 2025. Agentic LLM framework for adaptive decision discourse. *Preprint*, arXiv:2502.10978.

Ivar Frisch and Mario Giulianelli. 2024. Llm agents in interaction: Measuring personality consistency and linguistic alignment in interacting populations of large language models. In *Proceedings of the 1st Workshop on Personalization of Generative AI Systems (PERSONALIZE 2024)*, pages 102–111, St. Julians, Malta. Association for Computational Linguistics.

Chen Gao, Xiaochong Lan, Nian Li, Yuan Yuan, Jingtao Ding, Zhilun Zhou, Fengli Xu, and Yong Li. 2024. Large language models empowered agent-based modeling and simulation: a survey and perspectives. *Humanities and Social Sciences Communications*, 11.

Yarolsav Harbar and Jaroslaw A. Chudziak. 2025. Simulating oxford-style debates with LLM-based multi-agent systems. In *Intelligent Information and Database Systems*, pages 286–300, Singapore. Springer Nature Singapore.

Mardi Jon Horowitz. 1991. *Person schemas and maladaptive interpersonal patterns*. The University of Chicago Press.

Le Huang, Hengzhi Lan, Zijun Sun, Chuan Shi, and Ting Bai. 2024a. Emotional RAG: Enhancing Role-Playing Agents through Emotional Retrieval . In *2024 IEEE International Conference on Knowledge Graph (ICKG)*, pages 120–127, Los Alamitos, CA, USA. IEEE Computer Society.

Yizhe Huang, Xingbo Wang, Hao Liu, Fanqi Kong, Aoyang Qin, Min Tang, Song-Chun Zhu, Mingjie Bi, Siyuan Qi, and Xue Feng. 2025. Adasociety: An adaptive environment with social structures for multi-agent decision-making. *Preprint*, arXiv:2411.03865.

Yue Huang, Zhengqing Yuan, Yujun Zhou, Kehan Guo, Xiangqi Wang, Haomin Zhuang, Weixiang Sun, Lichao Sun, Jindong Wang, Yanfang Ye, and Xiangliang Zhang. 2024b. Social science meets LLMs: How reliable are large language models in social simulations? *Preprint*, arXiv:2410.23426.

Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2021. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*, 7(3):535–547.

Vann S. Joines. 2016. Understanding second-order structure and functioning. *Transactional Analysis Journal*, 46(1):39 – 49.

Adam Kostka and Jarosław A. Chudziak. 2024. Synergizing logical reasoning, knowledge management and collaboration in multi-agent LLM system. In *Pacific Asia Conference on Language, Information and Computation (PACLIC 38)*, Tokyo, Japan.

Adam Kostka and Jarosław A. Chudziak. 2025. Towards cognitive synergy in llm-based multi-agent systems: Integrating theory of mind and critical evaluation. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 47.

Justin Mittelstädt, Julia Maier, Panja Goerke, Frank Zinn, and Michael Hermes. 2024. Large language models can outperform humans in social situational judgments. *Scientific Reports*, 14.

Elizabeth Morrow, Teodor Zidaru, Fiona Ross, and 1 others. 2023. Artificial intelligence technologies and compassion in healthcare: A systematic scoping review. *Frontiers in Psychology*, 13.

T Novey. 1998. A proposal for an integrated self [letter to the editor]. *The Script*, 28(7):6.

OpenAI. 2024. Hello GPT-4o. Accessed: 2025-04-05.

Joon Sung Park, Joseph C. O'Brien, Carrie J. 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*, UIST '23. Association for Computing Machinery.

JL Schiff. 1981. Ego states. In *Workshop at the Southeast Institute Annual Spring Conference, Raleigh, NC*.

I. Stewart and V. Joines. 2012. *TA Today: A New Introduction to Transactional Analysis*, 2nd edition. Lifespace Publishing, Nottingham.

Ryosuke Takata, Atsushi Masumori, and Takashi Ikegami. 2024. Spontaneous emergence of agent individuality through social interactions in large language model-based communities. *Entropy*, 26(12).

Amir Taubenfeld, Yaniv Dover, Roi Reichart, and Ariel Goldstein. 2024. Systematic biases in LLM simulations of debates. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, page 251–267. Association for Computational Linguistics.

Maria Teresa Tosi and Susanna Bianchini. 2013. A social-cognitive definition of ego states to implement ta research. *International Journal of Transactional Analysis Research & Practice*, 4(1):107–112.

Jing Yi Wang, Nicholas Sukiennik, Tong Li, Weikang Su, Qianyue Hao, Jingbo Xu, Zihan Huang, Fengli Xu, and Yong Li. 2024a. A survey on human-centric LLMs. *Preprint*, arXiv:2411.14491.

Shan Wang, Fang Wang, Zhen Zhu, Jingxuan Wang, Tam Tran, and Zhao Du. 2024b. Artificial intelligence in education: A systematic literature review. *Expert Systems with Applications*, 252:124167.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. React: Synergizing reasoning and acting in language models. *International Conference on Learning Representations (ICLR)*.

Monika Zamojska and Jaroslaw A. Chudziak. 2025a. Games agents play: Towards transactional analysis in LLM-based multi-agent systems. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 47.

Monika Zamojska and Jaroslaw A. Chudziak. 2025b. Simulating human communication games: Transactional analysis in LLM agent interactions. In *Recent Challenges in Intelligent Information and Database System*, pages 173–187, Singapore. Springer Nature Singapore.

H. Zhang, J. Yin, M. Jiang, and C. Su. 2024. Can agents spontaneously form a society? introducing a novel architecture for generative multi-agents to elicit social emergence. *Preprint*, arXiv:2409.06750.

Wanjun Zhong, Lianghong Guo, Qiqi Gao, He Ye, and Yanlin Wang. 2024. Memorybank: Enhancing large language models with long-term memory. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(17):19724–19731.

Önder Gürcan. 2024. LLM-augmented agent-based modelling for social simulations: Challenges and opportunities. In *Hybrid Human AI Systems for the Social Good*, pages 134–144.