

Beyond Self-Interest: Modeling Social-Oriented Motivation for Human-like Multi-Agent Interactions

Jingzhe Lin
School of Artificial Intelligence,
Beijing Normal University
& Beijing Key Laboratory of Artificial
Intelligence for Education
& Engineering Research Center of
Intelligent Technology and
Educational Application, MOE
Beijing, China
linjingzhe@bnu.edu.cn

Ceyao Zhang
School of Intelligence Science
and Technology,
Institute for Artificial Intelligence,
Peking University
Beijing, China
ceyaozhang@pku.edu.cn

Yaodong Yang
Institute for Artificial Intelligence,
Peking University
Beijing, China
yaodong.yang@pku.edu.cn

Yizhou Wang
School of Computer Science,
Institute for Artificial Intelligence,
State Key Laboratory of General
Artificial Intelligence,
Peking University
Beijing, China
yizhou.wang@pku.edu.cn

Song-Chun Zhu
School of Intelligence Science and
Technology, Institute for Artificial
Intelligence,
Peking University
& Beijing Institute for General
Artificial Intelligence
Beijing, China
s.c.zhu@pku.edu.cn

Fangwei Zhong*
School of Artificial Intelligence,
Beijing Normal University
& Beijing Key Laboratory of Artificial
Intelligence for Education
& Engineering Research Center of
Intelligent Technology and
Educational Application, MOE
Beijing, China
fangweizhong@bnu.edu.cn

ABSTRACT

Large Language Models (LLMs) demonstrate significant potential for generating complex behaviors, yet most approaches lack mechanisms for modeling social motivation in human-like multi-agent interaction. We introduce Autonomous Social Value-Oriented agents (ASVO), where LLM-based agents integrate desire-driven autonomy with Social Value Orientation (SVO) theory. At each step, agents first update their beliefs by perceiving environmental changes and others' actions. These observations inform the value update process, where each agent updates multi-dimensional desire values through reflective reasoning and infers others' motivational states. By contrasting self-satisfaction derived from fulfilled desires against estimated others' satisfaction, agents dynamically compute their SVO along a spectrum from altruistic to competitive, which in turn guides activity selection to balance desire fulfillment with social alignment. Experiments across School, Workplace, and Family contexts demonstrate substantial improvements over baselines in behavioral naturalness and human-likeness. These findings show that structured desire systems and adaptive SVO drift enable realistic multi-agent social simulations. **Project page:** <https://asvo-agents.github.io/ASVO-agents/>

*Corresponding author.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/OITH7375>

KEYWORDS

Social value orientation; Social simulation; Intrinsic motivation

ACM Reference Format:

Jingzhe Lin, Ceyao Zhang, Yaodong Yang, Yizhou Wang, Song-Chun Zhu, and Fangwei Zhong. 2026. Beyond Self-Interest: Modeling Social-Oriented Motivation for Human-like Multi-Agent Interactions. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/OITH7375>

1 INTRODUCTION

Simulating human-like behavior has long been a core goal in artificial intelligence [11]. Beyond reproducing individual actions, it is important to model the complex social dynamics in realistic multi-agent scenarios [3, 31], where agents are endowed with diverse personalities, motivations, and social preferences. Such simulations have potential applications in domains including education [51], transportation [30], and finance [46].

Recent large language models [1, 5] have opened a new era of agent-based social simulation. Pioneering frameworks such as Generative Agents [28] and ProAgent [49] demonstrate autonomous social behaviors, while subsequent work extends this paradigm to large-scale reasoning [21], institutional settings [22], or general social simulation [37, 39]. These advances reveal that LLMs enable complex social inference and emergent norms [4], yet most current approaches rely on static [38] or rule-based agent [48] personalities, limiting their ability to capture the dynamics of human intrinsic motivation.

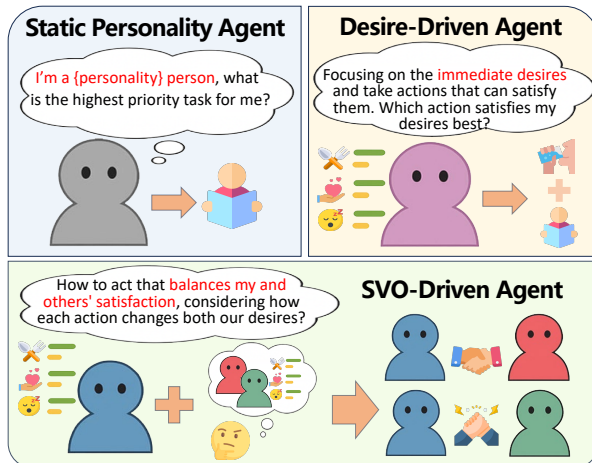


Figure 1: Upper-left: Static Personality agents condition actions on fixed goals. Upper-right: Desire-Driven agents produce human-like behavior by grounding decisions in internal desires. Lower: SVO-Driven agents (ours) build upon desire-based autonomy by incorporating Social Value Orientation and explicitly weighing each action’s impact on both self and others’ anticipated satisfaction, thereby moving beyond pure self-interest toward a richer model of human social behavior.

Consequently, there is growing interest in integrating structured psychological constructs such as **Social Value Orientation (SVO)** [26] and motivational models into LLM agents [18, 19, 40, 50]. Recent progress in **desire-driven autonomy** shows agents better emulate human-like behavior when guided by intrinsic motivations rather than fixed goals [40]. However, these approaches primarily focus on task generation and lack mechanisms to capture how social preferences shape complex multi-agent interactions. This is a gap this paper addresses by coupling desire systems with adaptive SVO computation.

In this paper, we propose **Autonomous Social Value-Oriented agents (ASVO)**, a novel framework that integrates desire-driven autonomy with social motivation grounded in SVO. Within this framework, agents are categorized into four core social personality types: *Altruistic*, *Prosocial*, *Individualistic*, and *Competitive*. Each type reflects a distinct trade-off between self-interest and consideration for others. This unified framework also enables agents to adaptively infer their own and others’ motivations, engage in complex social interactions, and demonstrate value shifts that mirror complex human dynamics. Built on the Concordia [37], our method enables agents to actively take activities to update and satisfy both intrinsic desires and social value preferences as they interact with others in diverse simulated societies.

The contributions are threefold: 1) We introduce an LLM-based framework, ASVO, which embeds Social Value Orientation into desire-driven agents, facilitating dynamic motivational adaptation in multi-agent environments. 2) We design a multi-dimensional evaluation metric to strictly assess behavioral alignment, covering cooperation dynamics, naturalness, and human-likeness. 3) We validate ASVO through extensive benchmarking across three distinct contexts, showing that it yields more interpretable and socially consistent behaviors compared to strong baselines.

2 RELATED WORK

Social Simulation with Large Language Models. Recent advances in large language models (LLMs) have improved the realism and complexity of agent-based social simulations [44, 52]. Pioneering frameworks such as Voyager [43] and CAMEL [21] showed that LLM-driven agents can autonomously explore, communicate, and collaborate in virtual environments, enabling complex group behaviors, norms, and roles. Further, earlier intrinsic-motivation-driven LLM agents, such as D2A [40], together with simulators like Ecoagent [23], OASIS [47], and AgentScope [13], have extended simulation capabilities toward larger-scale and more diverse social environments. Subsequent research has expanded into domains including social dilemma games [20], online community simulation [2], and value alignment in digital societies [12]. These works demonstrate that LLMs enable complex language interaction and adaptive behavior modeling, pushing the boundaries of computational social science. However, most existing LLM-based simulations rely on static or rule-based agent personalities, which fail to capture the full spectrum of human-like heterogeneity and dynamic value change. Thus, there is growing interest in integrating structured psychological constructs, such as Social Value Orientation (SVO) and motivational models, into LLM agents to simulate real-world social phenomena [18, 19, 50].

Social Value Orientation in Agent-Based Models. Social Value Orientation (SVO) is a key concept for quantifying individual differences in social preferences, cooperation, and competition [26, 36]. SVO classifies agents along a continuum from altruistic and prosocial to individualistic and competitive types, providing a grounded metric for modeling how agents weigh their own welfare against that of others. Continuous SVO measures and related constructs have become standard tools for psychological assessment and multi-agent simulations, enabling fine-grained analysis of resource allocation, group coordination, and social dilemmas [14]. Recent work on multi-agent systems has shown that assigning heterogeneous SVO values to agents leads to more realistic social outcomes. For example, in human-robot teams, SVO-inspired reward shaping improves fairness, trust, and group performance [42]. In computational sociology, SVO-based models enable the study of emergent cooperation, norm formation, and polarization, especially when combined with networked agent learning [24]. Despite these advances, most prior work treats SVO as a static attribute, failing to capture the dynamic and context-sensitive nature of human social preferences. Recent work on desire-driven autonomy [40] enables agents to evolve intrinsic motivations in response to experience and feedback, offering new opportunities for modeling personality plasticity and adaptive decision-making in multi-agent systems.

Motivation and Personality in LLM Agents. Modeling dynamic motivation and personality remains a central challenge for AI-based social simulation. Traditional multi-agent systems often rely on handcrafted rules or fixed reward structures, limiting their expressiveness and adaptability [9]. The integration of LLMs allows richer language understanding and social inference, but explicit mechanisms for intrinsic motivation, needs satisfaction, and personality evolution are still lacking. Recent approaches bridge this gap by introducing dual motivational frameworks, combining classic psychological theories such as Maslow’s hierarchy of needs [25],

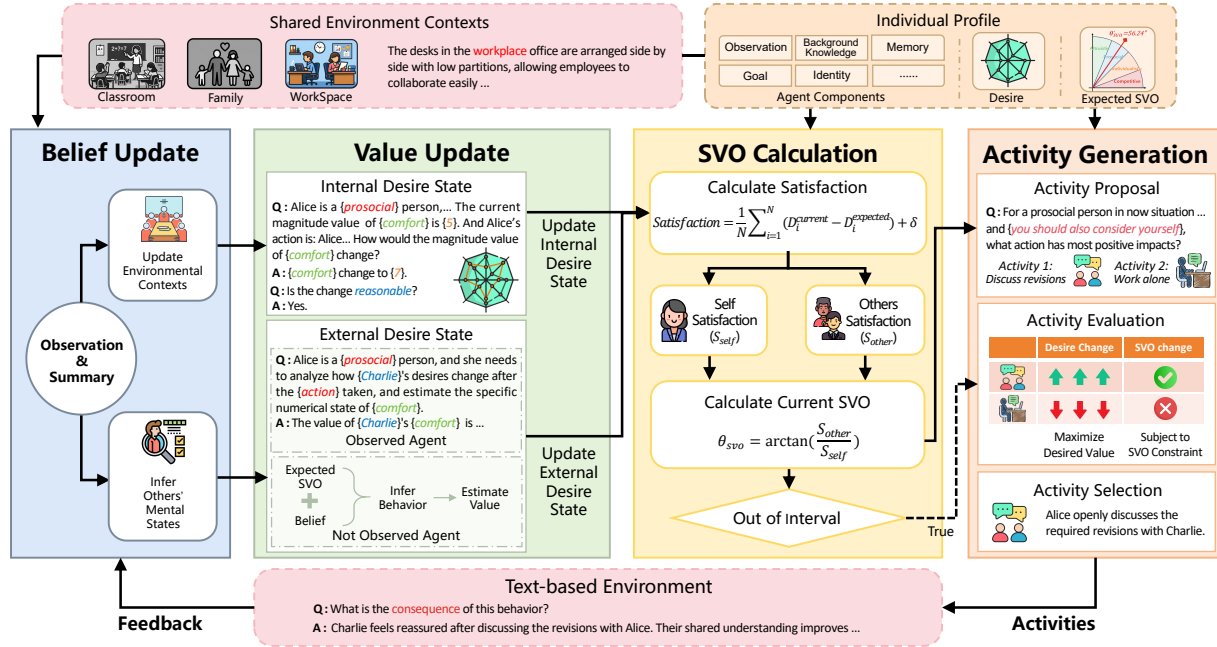


Figure 2: Overview of the ASVO framework. During the Initialization stage, each agent’s Individual Profile is constructed, specifying its social personality type (expected SVO), a set of core psychological desires, and agent-specific components including *Observation*, *Memory*, *Background Knowledge*, *Identity*, and *Goal*. These attributes are encoded into the agent’s internal state, forming the basis of simulation. At each simulation step, the pipeline iteratively proceeds through four key modules: **Belief Update** (agents incorporate shared environment contexts and summarize observations to update beliefs about others’ states), **Value Update** (updating the Internal Desire State via self-reflection and the External Desire State by estimating others’ needs, including unobserved agents), **SVO Calculation** (computing self and others’ satisfaction to compute the agent’s current SVO with an interval check), and **Activity Generation** (proposing, evaluating, and selecting LLM-driven activities by synthesizing desires, context, and SVO). This forms a closed feedback loop over a text-based environment, enabling the co-evolution of motivations and social values throughout multi-agent social interaction.

Self-Determination Theory [29], or Big Five personality modeling [15] with SVO and value reasoning. These frameworks support more realistic simulation of individual adaptation, group influence, and value-driven decision-making in open-ended social environments [40].

3 PRELIMINARIES

Social Value Orientation (SVO) [26]. is a psychological construct that quantifies how individuals balance their own interests with those of others during social interactions. It captures individual differences in social preferences across people and contexts, reflecting how much weight one places on others’ welfare when making decisions. Accordingly, it serves as a foundational framework for explaining cooperative, altruistic, and competitive behaviors in groups. Formally, SVO is represented as an angle:

$$\theta_{SVO} = \arctan\left(\frac{V_{other}}{V_{self}}\right), \quad (1)$$

where V_{self} and V_{other} denote the values attributed to oneself and to others, respectively. Based on this angle, individuals are categorized into four types: *altruistic* (high concern for others), *prosocial* (balancing self and others), *individualistic* (focused on self-interest), and *competitive* (maximizing relative advantage). This classification provides a theoretical foundation for modeling diverse human

behaviors and has been widely applied in multi-agent systems [7], behavioral economics [34], and social psychology [35], enabling realistic simulation of social norms and collective decision-making.

4 PROBLEM FORMULATION

In this work, we formalize the problem of modeling socially adaptive agents that integrate structured intrinsic motivations with dynamically evolving social value orientations (SVOs). Each agent continuously balances self- and other-oriented drives in a shared environment, allowing value-based reasoning and context-sensitive social behavior to emerge.

Let s_t denote the environment state at time t , and $o_{0:t-1}$ and $a_{0:t-1}$ represent the sequences of past observations and actions, respectively. Each agent i is characterized by an internal desire vector D_t^i , a social value orientation angle θ_t^i , and a fixed attribute set ψ_i that encodes static features such as identity, memory, and personality. At each step, the agent generates a new activity conditioned on its experience, motivational state, and social orientation:

$$a_t^i \sim \text{AGENT}(\cdot \mid o_{0:t-1}, a_{0:t-1}, D_t^i, \theta_t^i; \psi_i), \quad (2)$$

Here, D_t^i encodes the agent’s intrinsic motivations across multiple dimensions, while θ_t^i determines how self- and other-oriented satisfactions are weighted when evaluating possible actions. The

static attributes ψ_i serve as contextual priors that shape consistent identity and reasoning style throughout the simulation. Together, these components define a unified decision-making process in which intrinsic desires interact with social preferences to produce coherent behaviors. The agent selects actions that maximize its desire satisfaction while ensuring the resulting SVO remains within its expected personality range, i.e., $a_t^i = \arg \max_a U_{\text{self}}(a \mid D_t^i)$ s.t. $\theta_t^i(a) \in [\theta_{\min}^i, \theta_{\max}^i]$.

5 METHODOLOGY

In this section, we will explain each module of our multi-agent simulation in detail, as shown in Figure 2. We first initialize the individual profile and environmental contexts. The core simulation consists of four modules: (1) **belief updates**, (2) **value updates**, (3) **SVO calculations**, and (4) **activity generation**. Lastly, we describe how each module iterates.

5.1 Individual Profile Initialization

At the start of each simulation, a set of agents is instantiated within a shared environment context. Each agent is initialized with an *expected SVO*, representing its initial social personality type. During simulation, this expected SVO serves as a reference, while the agent’s *current SVO* is dynamically updated at each step based on ongoing behavioral feedback and social context. Thus, the agent’s social value orientation is not static, but adaptively evolves throughout the simulation process.

In addition to SVO, each agent is equipped with an individual profile that includes eight structured desire dimensions forming the foundation of its intrinsic motivation system. These desires, together with the agent’s background knowledge, memory, and identity information, constitute the initial cognitive and affective states used by the reasoning modules.

This dual-SVO and multi-desire initialization provides a psychologically grounded motivational baseline that shapes each agent’s value tendency and perception of others. It also prepares subsequent modules such as belief update and activity generation, ensuring that social adaptation arises from coherent internal dynamics.

5.2 Belief Update

The **belief update** module serves as a core process in our framework, underpinning both self-modeling and modeling of others. At each step, agents actively perceive changes in the environment independently and update their internal belief of the state, denoted as b_t^i . Specifically, agents gather new observations about: 1) **Environmental Contexts**: Changes in shared environment states, such as shifts in group membership or resources, represented as o_t^{env} ; 2) **Others’ Actions**: Observations or inferences about other agents’ actions, represented as o_t^{other} . These beliefs together provide the foundation for subsequent modeling of both self and others.

5.3 Value Update

At each simulation step, each agent updates both its internal desire state and its estimation of others’ desire states.

Internal Desire State: Each agent maintains a structured, multi-dimensional desire state to capture evolving needs during interaction. It is updated via an LLM-mediated reflective loop: the agent’s

most recent action and observed consequences are fed into a structured prompt, with the agent’s SVO imposed as an explicit constraint. For example, a prosocial agent like Alice organizing a review session weighs personal comfort against collective benefit, aligning choices with both individual and group welfare. An update is accepted only if it remains consistent with the agent’s established SVO profile; otherwise, the model iterates reflection and correction. Specifically, in this paper we instantiate the desire state with eight dimensions, while the mechanism generalizes beyond this choice.

External Desire State: To estimate others’ desires, agents perform SVO-conditioned social inference. If the target agent is visible (e.g., Charlie), the agent directly infers their desire state from behavioral cues; if invisible (e.g., Bob), it first infers the most likely behavior from SVO and contextual memory, then estimates the current desire state conditioned on that behavior. This aligns with the Value Update pathways in Figure 2. Formally, given $o_{0:t-1}$ and $a_{0:t-1}$, agent i infers the latent motivational state of agent j via:

$$\widehat{D}_t^{ji} \sim \text{LLM}(o_{0:t-1}, a_{0:t-1}, \theta_t^j), \quad (3)$$

where \widehat{D}_t^{ji} denotes i ’s estimate of j ’s desire vector, and $\widehat{\theta}_t^{ji}$ denotes the inferred SVO of j . This step turns observations into socially meaningful representations, enabling anticipation of others’ intentions for subsequent value updates and decision-making.

5.4 SVO Calculation

Throughout the simulation, each agent maintains both a fixed *expected SVO* (personality baseline) and a *dynamic current SVO* that is continuously recalculated based on recent social interactions and value satisfaction. The current SVO serves as the immediate basis for decision-making and is regularly compared to the expected SVO to regulate behavioral consistency and personality alignment.

Following value updates, each agent calculates its current SVO by comparing internal and external satisfaction. For N tracked desires, satisfaction is computed as the average difference between the current and expected values:

$$S = \frac{1}{N} \sum_{i=1}^N (D_i^{\text{current}} - D_i^{\text{expected}}) + \delta, \quad (4)$$

where D_i^{current} and D_i^{expected} represent the present and expected values of the i -th desire, respectively (in our implementation, each desire takes values in $[0, 10]$). In practice, we clip the per-dimension differences ($D_i^{\text{current}} - D_i^{\text{expected}}$) to $[-4, 4]$, with this numeric range chosen based on empirical experience to stabilize the ratio. This bounded design also reflects the fact that extreme discrepancies tend to saturate and no longer change perceived satisfaction [17]. Here δ is a constant offset chosen to shift satisfactions to be positive before taking the ratio, ensuring a stable and well-defined SVO angle. The SVO angle is then determined as:

$$\theta_{svo} = \arctan\left(\frac{S_{\text{other}}}{S_{\text{self}}}\right), \quad (5)$$

where S_{self} and S_{other} are the satisfaction scores for self and others.

If the computed SVO value deviates from the expected personality interval (for example, if a prosocial agent’s SVO is unusually high or low), a regulatory mechanism is triggered at the activity

generation stage. The agent’s action selection prompt is then adaptively modified to guide its choices gently back toward the expected SVO range. This iterative feedback, illustrated in Figure 2, enables dynamic and interpretable regulation of social value orientation through natural language guidance.

5.5 Activity Generation

Based on the regulated SVO, ASVO generates agent behavior through a structured multi-step process that synthesizes motivational needs, social context, and value orientation. At each step, the agent assembles relevant information, including updated desires, contextual memory, and recent observations, into the LLM input, which is adaptively modified by SVO regulation to guide behavior as needed.

The LLM proposes multiple candidate actions based on the agent’s needs and current SVO. Each candidate is evaluated for its predicted impact on self-satisfaction and alignment with value orientation. For example, a prosocial agent may weigh sharing notes (collective benefit) against collecting problems (self-benefit) to determine which better fits its SVO.

Formally, the action selection process can be modeled as a constrained optimization that maximizes the agent’s own desire satisfaction while maintaining consistency with its social value orientation (SVO). Let $C_t^i = \{\tilde{a}_{t,k}^i\}_{k=1}^K$ denote the candidate action set generated by the LLM at step t , and let $U_{\text{self}}(a | D_t^i)$ represent the predicted improvement in the agent’s internal desires when taking action a . Using the SVO Calculation module, the agent predicts the resulting SVO angle for each candidate, denoted $\tilde{\theta}_t^i(a)$. Given the expected SVO interval $[\theta_{\min}^i, \theta_{\max}^i]$, the final action is selected as:

$$a_t^i = \arg \max_{a \in C_t^i} U_{\text{self}}(a | D_t^i) \quad \text{s.t.} \quad \tilde{\theta}_t^i(a) \in [\theta_{\min}^i, \theta_{\max}^i]. \quad (6)$$

This formulation ensures that the chosen behavior not only maximizes the agent’s internal satisfaction but also stays within the acceptable SVO range, thereby maintaining a balance between self-driven motivation and social alignment.

The agent selects the action that best balances desire fulfillment and SVO alignment. This chosen action is executed, and the outcomes are used to update desires and SVO in the next step. In this way, agent behavior remains contextually appropriate and value-aligned throughout the simulation, with regulatory adjustments ensuring personality consistency even as social dynamics evolve. Full prompts are provided in the Appendix.

5.6 Iterative Cycle

The ASVO framework operates through a continuous perception, motivation, value, action loop. This process integrates all preceding modules into a unified and adaptive behavioral mechanism. Formally, the overall cycle at time step t can be expressed as:

$$s_t \Rightarrow o_{0:t} \Rightarrow (D_t^i, \widehat{D}_t^{j|i}) \Rightarrow \theta_t^i \Rightarrow a_t^i \Rightarrow s_{t+1}, \quad (7)$$

where s_t is the environment state, $o_{0:t}$ represents the accumulated observations, $(D_t^i, \widehat{D}_t^{j|i})$ are the updated internal and inferred external desire states, θ_t^i denotes the current social value orientation, and a_t^i is the generated action.

At each iteration, agents first perceive environmental cues and others’ actions to update their beliefs and observation histories $o_{0:t}$. They then revise both self- and other-related motivational

states $(D_t^i, \widehat{D}_t^{j|i})$ through reflective and inferential reasoning, which jointly shape the updated SVO θ_t^i . Conditioned on this updated motivational–social state, the agent selects an action a_t^i according to its policy, balancing self-oriented needs with socially aligned value considerations. The resulting action modifies the environment to s_{t+1} , providing new contextual feedback for the next step.

Through this iterative cycle, ASVO agents continuously adapt their desires, infer others’ intentions, and refine their social value orientations. This dynamic feedback mechanism ensures that simulated behaviors remain context-sensitive, interpretable, and personality-consistent across diverse scenarios.

6 EXPERIMENTS

We conduct experiments across three social contexts to evaluate ASVO’s effectiveness in modeling human-like, socially-motivated multi-agent interactions. We assess behavioral naturalness, SVO adaptation dynamics, and emergent cooperation patterns against baseline approaches. Code and supplementary materials are available at <https://asvo-agents.github.io/ASVO-agents/>.

6.1 Experimental Setups

All experiments are conducted within an extended Concordia-based multi-agent simulation platform [37], specifically designed to systematically evaluate the effects of dynamic SVO adaptation and desire-driven decision-making in diverse social contexts. The core simulation platform robustly supports flexible agent configuration, real-time interaction logging, and reproducible environment initialization for consistent and comparative experimental analysis.

Environments and Scenarios. Building on the established micro–meso–macro stratification paradigm [6, 10], we extend the simulation framework beyond educational settings to encompass three representative **social contexts**, *school*, *family*, and *workplace*. Each context contains multiple **social scales** (dyadic, small group, and large group), allowing the investigation of SVO adaptation and motivational reasoning under progressively complex social structures. The family context includes the first two scales, reflecting its naturally smaller social scope. Each context represents a distinct social type, where Family reflects small-scale interactions, School represents structured peer environments, and Workplace corresponds to institutionalized collaboration and competition.

School. *a) Micro:* A dormitory-like environment where two students share daily routines and emotions, focusing on interpersonal adjustment and empathy formation in close relationships. *b) Meso:* A classroom small-group setting in which students collaborate or compete for limited resources, reflecting the balance between cooperation, self-interest, and social comparison. *c) Macro:* A class election involving public speeches, alliances, and peer influence, capturing collective norm formation, leadership emergence, and competitive persuasion dynamics.

Workplace. *a) Micro:* A two-person office dyad reflecting hierarchical or peer communication under performance pressure, examining how SVO affects conflict management and support-seeking behaviors. *b) Meso:* A project team emphasizing collaboration, task division, and strategic coordination, where agents must negotiate responsibility and recognition. *c) Macro:* A departmental review simulating institutionalized cooperation and competition driven by

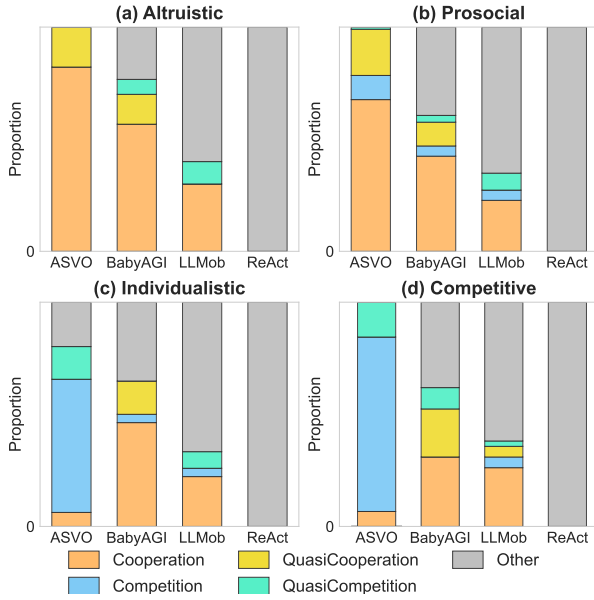


Figure 3: Behavioral distributions of four SVO-based social personality types, *Altruistic*, *Prosocial*, *Individualistic*, and *Competitive*, showing proportions of cooperation, competition, and intermediate behaviors across frameworks.

evaluation, incentives, and career-oriented motivation, revealing how social values adapt to structured hierarchies.

Family. *a) Micro:* A weekend home scenario with two siblings staying indoors while parents are away, reflecting mild autonomy, cooperation, and self-regulation in intimate family interaction. *b) Meso:* A family park outing with siblings and parents engaging in shared play, small conflicts, and emotional bonding, highlighting the role of affection and social harmony in value adaptation.

This hierarchical and cross-context design enables a unified examination of how SVO-driven adaptation and desire-based reasoning evolve from intimate relationships to institutionalized cooperation and competition. It also allows comparison of emergent behavioral patterns across social scales and domains, bridging micro-level interpersonal dynamics and macro-level collective outcomes.

Language Model Selection. To comprehensively evaluate the impact of foundational models on multi-agent social simulation, we selected four representative large language models covering a range of scales and architectures: **GPT-5** [32], **Qwen3-235b** [45], **Deepseek-v3** [8], and **Gemini-2.5** [33]. These models include both open-source and proprietary solutions and represent state-of-the-art advancements in LLMs. All models are utilized both as behavior generators (i.e., simulating agent actions and interactions) and as automatic evaluators (i.e., rating the naturalness and human-likeness of agent behaviors) within our experimental framework. For each experimental run, the same set of agent profiles, environmental contexts, and evaluation prompts is used to ensure a fair and reproducible comparison across LLMs.

Evaluation Metrics. We focus on the following key metrics to evaluate agent behavior and social adaptation: 1) **Cooperation and competition rate:** The frequency and proportion of cooperative and competitive actions among agents, as defined by whether

Table 1: Average naturalness (N) and human-likeness (H) scores for each LLM and method. Avg: mean value; Std: standard deviation across all LLMs for each method.

LLM	ReAct		BabyAGI		LLMob		D2A		ASVO	
	N	H	N	H	N	H	N	H	N	H
School										
Deepseek	4.004	3.667	3.438	2.771	4.500	3.708	3.292	3.083	4.750	4.792
GPT-5	3.750	3.417	3.833	3.521	4.354	4.271	3.812	3.448	4.958	4.958
Gemini-2.5	3.833	3.896	3.479	3.042	4.312	4.062	3.427	3.083	4.708	4.708
Qwen3	3.917	3.521	3.458	3.188	4.271	4.021	3.677	3.406	4.792	4.824
Avg	3.876	3.625	3.552	3.130	4.359	3.938	3.552	3.255	4.802	4.821
Std	1.097	0.658	1.322	0.877	0.861	0.565	1.426	1.047	0.497	0.456
Workplace										
Deepseek	3.214	2.385	2.455	2.397	3.253	2.609	2.849	2.479	4.766	4.119
GPT-5	3.221	3.045	3.349	3.064	4.144	3.792	3.703	3.458	4.849	4.019
Gemini-2.5	3.448	3.182	2.673	2.724	3.442	3.269	2.891	2.698	4.878	4.034
Qwen3	3.297	2.760	2.843	2.673	3.330	3.058	3.135	2.818	4.782	4.026
Avg	3.330	2.877	2.798	2.710	3.542	3.182	3.145	2.863	4.819	4.049
Std	1.568	1.063	1.358	0.968	1.463	1.088	1.444	0.991	0.491	0.466
Family										
Deepseek	3.286	3.262	2.917	2.690	3.589	3.339	4.010	3.167	4.750	3.964
GPT-5	3.345	3.202	3.345	3.202	4.202	3.863	4.479	3.750	4.905	4.024
Gemini-2.5	3.726	3.488	3.423	1.433	3.845	3.429	4.031	3.615	4.863	4.077
Qwen3	4.006	3.536	3.685	3.315	4.000	3.738	4.490	3.781	4.381	3.820
Avg	3.693	3.449	3.342	3.079	3.909	3.579	4.253	3.578	4.725	3.946
Std	1.423	0.921	1.239	0.904	1.166	0.827	1.148	0.941	0.624	0.648

agents’ actions are promotive of or oppositional to others’ goal attainment [16]. 2) **Naturalness:** The plausibility and human-likeness of agent behaviors, automatically evaluated by large language models using standardized prompts and scoring rubrics [41]. 3) **Human-likeness:** The perceived similarity of agent behaviors to real human actions, as judged by LLM evaluators based on their conformity to human norms, diversity, and context-appropriateness [28].

6.2 Quantitative Results

To evaluate effectiveness, we compare ASVO with representative LLM-based agents, including ReAct [48], BabyAGI [27], LLMob [38], and D2A [40] (an earlier intrinsic-motivation-driven framework), under the same simulation settings with matched profiles, goals, and context. Only ASVO jointly models SVO and dynamic desires; the baselines follow their original designs (e.g., goal-driven planning, fixed habits, or static traits) without explicit SVO modeling.

Figure 3 illustrates the consistent and interpretable behavioral patterns of ASVO across the full spectrum. Specifically, agents initialized with altruistic or prosocial profiles demonstrate a distinct tendency towards cooperative behaviors. This validates the model’s effective alignment with socially oriented motivations. Conversely, as the SVO angle shifts toward individualistic orientations, the behavioral distribution transitions accordingly. Cooperative tendencies diminish, while competitive actions become increasingly dominant. This smooth, monotonic progression underscores ASVO’s capability to model the continuous behavioral plasticity inherent in human social values.

In contrast, baseline models fail to reproduce this structured trend: their generated actions remain largely insensitive to SVO type, often demonstrating neutral or context-independent behavior. ASVO uniquely demonstrates a well-ordered continuum of cooperation–competition balance and complex behavioral diversity, bridging quantitative SVO parameters with observable outcomes.

Beyond capturing this graded behavioral spectrum, we evaluate the *naturalness* and *human-likeness* of generated behaviors across three social contexts: *School*, *Workplace*, and *Family*. As shown




Figure 4: Distribution of cooperative and competitive actions for each social personality type within the *School* context across micro-, meso-, and macro-level scales.

in Table 1, ASVO consistently achieves the highest scores across all contexts and LLM backbones, with remarkably low standard deviations indicating strong robustness and generalization. The smaller performance gaps among different LLMs suggest that value-driven adaptation effectively mitigates backbone variability. Multi-judge stability analysis is provided in Appendix D.4. These results confirm that incorporating explicit value structures and adaptive SVO mechanisms substantially enhances behavioral coherence and social consistency in multi-agent simulations, enabling agents to express distinct yet contextually appropriate social behaviors.

To further assess the fidelity of our SVO modeling, we further examine whether agents’ SVO values remain within their intended theoretical ranges over time. As illustrated in Figure 5, the mean and standard deviation of SVO for each personality type are plotted alongside shaded bands representing the expected SVO intervals. Across all steps, altruistic, prosocial, individualistic, and competitive agents consistently maintain SVO values within their respective reference ranges. This result confirms that ASVO effectively preserves distinct social value orientations throughout the simulation.

6.3 Scalability Analysis

To test whether ASVO can scale to larger scenarios, we extend the experiments by varying the horizon length and the agent population size. We use the *School* setting as a representative testbed to keep the environment fixed while changing only the scale factors (Table 2). We extend the simulation horizon from 6 to 12, 18, and 24 steps. ASVO maintains coherent and human-like behaviors as the horizon grows, with degradation at 24 steps. We increase the population from 4 to 8, 16, and 32 agents under the same setting. Scores decrease gradually with more agents, indicating ASVO remains effective when scaling to larger groups.

Beyond quantitative metrics, longer-horizon simulations demonstrate clearer long-range behavioral continuity. Altruistic agents gradually evolve from basic peer assistance to more structured cooperative behaviors, such as monitoring peers’ progress, organizing study groups, and offering emotional support. Competitive agents increasingly focus on performance disparities and may show frustration when others progress faster, while individualistic agents consistently favor independent and self-directed learning. In contrast, prosocial agents progressively converge toward stable supportive roles as group dynamics unfold.

Table 2: Robustness of ASVO under extended temporal horizons and agent population sizes in the *School* setting.

Setting	Altruistic		Prosocial		Individualistic		Competitive		Avg	
	N	H	N	H	N	H	N	H	N	H
Extended Temporal Horizons (Steps)										
6	4.83	4.33	4.83	4.33	4.50	4.20	4.33	4.30	4.63	4.29
12	5.00	4.51	5.00	4.17	4.75	4.50	4.25	4.42	4.75	4.39
18	4.89	4.39	4.94	4.50	4.67	4.23	4.78	3.83	4.82	4.23
24	4.42	4.21	4.83	4.50	4.46	4.29	4.29	4.08	4.50	4.27
Avg	4.79	4.36	4.90	4.38	4.60	4.31	4.41	4.16	4.68	4.30
Increased Agent Population Sizes (Agents)										
4	4.83	4.33	4.83	4.33	4.50	4.23	4.33	4.30	4.63	4.29
8	4.75	4.43	4.71	4.33	4.20	4.25	4.17	4.33	4.56	4.37
16	4.79	4.43	4.80	4.45	4.16	4.23	4.58	4.31	4.58	4.33
32	4.60	4.15	4.35	4.21	3.92	4.25	4.35	4.21	4.42	4.27
Avg	4.74	4.34	4.67	4.33	4.20	4.24	4.36	4.29	4.55	4.32

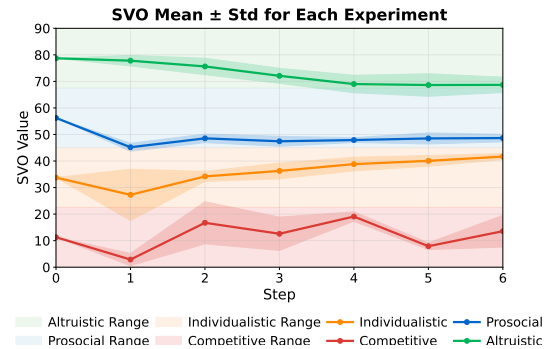


Figure 5: Temporal evolution of SVO mean and standard deviation for each personality type. Shaded areas denote the reference range of each SVO category, with values remaining within their respective intervals across simulation steps.

6.4 Behavioral Patterns Across Environments

To investigate how social personality shapes agent behavior across social contexts, we conducted controlled experiments in which all agents within each environment were assigned the same personality type. Each scenario was repeated five times with six simulation steps per run, and agent actions were aggregated for analysis.

Figure 4 shows how cooperation and competition rates vary with personality type and social scale within the *School*. Altruistic and prosocial agents consistently demonstrate higher proportions of cooperative actions, whereas individualistic and competitive agents are more prone to competitive behaviors. As the environment expands from micro (dormitory) to meso (classroom) and macro (class election), cooperation rates increase across all personality types, for example, from 0.15 to 0.56 for altruistic agents and from 0.22 to 0.28 for prosocial agents. In contrast, competition rates show subtler changes: they rise slightly for prosocial and competitive agents but decline marginally for altruistic and individualistic ones. This pattern indicates that larger and more complex environments enhance overall social engagement, leading agents to demonstrate richer and more differentiated behavioral patterns rather than uniformly greater competitiveness.

In particular, altruistic agents show the strongest context sensitivity, with cooperation rates more than tripling from the dormitory to the class election scenario. This pattern suggests that altruistic

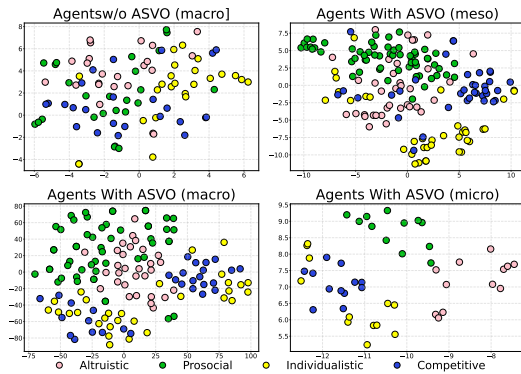


Figure 6: t-SNE visualization of agent behavior embeddings across four SVO personality types. The top-left panel shows agents without SVO, while the others demonstrate our agents (ASVO) across *School*, *Workplace*, and *Family*. ASVO yields clearer clustering, indicating more disentangled behaviors.

SVOs are highly responsive to expanded social visibility and collective incentives, leading to cooperative behavior in larger peer groups. Prosocial agents maintain a stable balance between self- and other-oriented motivations, demonstrating modest and consistent cooperation growth as environments scale up. By contrast, individualistic and competitive agents show lower cooperation ratios, with increased behavioral variance in meso- and macro-level contexts, reflecting adaptive shifts between self-interest and strategic social positioning. Overall, these findings show that ASVO captures flexible value-driven decision-making, with cooperative and competitive tendencies shaped by personality and social structure.

6.5 Ablation Studies of the LLMs Choice

To systematically assess the impact of different LLMs on emergent social behaviors across multiple social contexts, we conducted experiments under three representative environments: *School*, *Workplace*, and *Family*. We employed both **naturalness** and **human-likeness** as automatic evaluation metrics to measure the fluency and human-level coherence of agent behaviors.

Table 1 presents the average scores obtained by each LLM across the three contexts. The results reveal clear performance distinctions among models. **GPT-5** and **Qwen3** achieve the highest overall scores in both metrics, indicating their stronger capability in generating coherent, socially aligned, and context-sensitive behaviors. In contrast, **Deepseek** and **Gemini** show greater variability across contexts, reflecting higher sensitivity to environmental complexity and interaction structure. Among the three environments, the *Workplace* context shows the highest variability, reflecting the increased complexity and role diversity of organizational interactions. In contrast, the *Family* context yields lower overall scores due to its emotionally driven and less structured interactions, while the *School* context maintains balanced performance across all models. These results collectively highlight that both the social setting and the underlying LLM jointly shape the emergence of realistic and personality-consistent behaviors in large-scale simulations.

To evaluate whether ASVO enables the emergence of consistent personality-driven behaviors, we performed an ablation study using t-SNE to visualize the action distribution of four personality types.

By comparing the baseline without ASVO and the ASVO-enabled agents across three social contexts (*School*, *Workplace*, and *Family*), we directly observe how explicit value orientation modeling enhances the separability of agent behaviors.

In the absence of ASVO (top left), the actions of all personality types are highly intermixed in the t-SNE space, indicating that baseline agents lack consistent alignment between personality and behavior. There is no clear cluster structure, suggesting that static or shallow modeling fails to capture or preserve individual value orientations during social interactions. Figure 6 shows that, with ASVO, the actions of Altruistic and Prosocial agents are clearly separated from those of Individualistic and Competitive agents across all three social contexts (*School*, *Workplace*, and *Family*), reflecting robust personality-driven clustering.

This result clearly demonstrates that integrating structured social value orientation not only increases the interpretability and realism of multi-agent behavior, but also substantially enhances the separation and robustness of personality-driven dynamics. Because t-SNE is a qualitative visualization and may appear less distinctive when embeddings overlap, we additionally quantify cluster separability using the silhouette score and a cluster separation ratio computed from the same behavior embeddings. Detailed results are provided in the Appendix C.

7 CONCLUSION AND FUTURE WORK

This paper presents the **Autonomous Social Value-Oriented agents (ASVO)** framework, which integrates desire-driven autonomy with dynamically evolving Social Value Orientation (SVO). By combining intrinsic motivational reasoning with adaptive value regulation, ASVO enables LLM agents to demonstrate socially adaptive and personality-consistent behaviors across diverse contexts. Experiments across school, workplace, and family environments show that ASVO reproduces key social phenomena, including cooperation, competition, and value alignment, consistent with SVO theory. The framework offers a structured foundation for studying value-driven decision-making and supports applications in education, social training, and human-AI interaction.

Limitations: Despite its promising results, ASVO has two main limitations. First, the current motivational system relies on a fixed set of core desires. This design cannot fully capture the richness and variability of human values. Second, the simulated environments remain relatively controlled and deterministic. Although structurally diverse, they limit the observation of emergent adaptation in open-ended or uncertain social contexts.

Future Work: We will extend the internal value system to support self-evolving motivational structures during simulation. We also plan to explore more dynamic, open, and cross-cultural environments in which agents infer roles, negotiate norms, and adapt to complex social dynamics. In addition, incorporating multimodal signals may further improve the realism and complexity of agent interactions.

ACKNOWLEDGMENTS

This work was supported by Beijing Natural Science Foundation L252010, NSFC-62406010, and the Fundamental Research Funds for the Central Universities.

REFERENCES

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).
- [2] Armand Andrei, Paul Vicol, et al. 2023. Societal-scale simulation with large language models. *arXiv preprint arXiv:2311.14007* (2023).
- [3] Gary S Becker. 1974. A theory of social interactions. *Journal of Political Economy* 82, 6 (1974), 1063–1093.
- [4] Marcel Binz and Eric Schulz. 2023. Using cognitive psychology to understand GPT-3. *Proceedings of the National Academy of Sciences* 120, 6 (Feb. 2023). <https://doi.org/10.1073/pnas.2218523120>
- [5] Yuheng Cheng, Ceyao Zhang, Zhengwen Zhang, Xiangrui Meng, Sirui Hong, Wenhao Li, Zihao Wang, Zekai Wang, Feng Yin, Junhua Zhao, et al. 2024. Exploring large language model based intelligent agents: Definitions, methods, and prospects. *arXiv preprint arXiv:2401.03428* (2024).
- [6] James S. Coleman. 1990. *Foundations of Social Theory*. Harvard University Press.
- [7] Daniel E Collins, Conor Houghton, and Nirav Ajmeri. 2023. Social value orientation and integral emotions in multi-agent systems. In *Int'l Workshop on Coordination, Organizations, Institutions, Norms, and Ethics for Governance of MAS*. Springer, 118–138.
- [8] DeepSeek-AI. 2024. DeepSeek LLM: Scaling Open-Source Language Models with Longtermism. (2024). [arXiv:2401.02954](https://arxiv.org/abs/2401.02954)
- [9] Joshua M. Epstein. 2006. *Generative Social Science: Studies in Agent-Based Computational Modeling*. Princeton University Press.
- [10] Joshua M. Epstein and Robert Axtell. 1996. *Growing Artificial Societies: Social Science from the Bottom Up*. Brookings Institution Press; MIT Press.
- [11] James H Fetzer. 1990. What is artificial intelligence? In *Artificial intelligence: Its scope and limits*. Springer, 3–27.
- [12] Iason Gabriel. 2023. Artificial Social Intelligence: A Call for Building Human-Compatible AI. *Nature Machine Intelligence* 5 (2023), 3–4.
- [13] Dawei Gao, Zitao Li, Xuchen Pan, Weirui Kuang, Zhijian Ma, Bingchen Qian, Fei Wei, Wenhao Zhang, Yuxiang Xie, Daoyuan Chen, Liuyi Yao, Hongyi Peng, Zeyu Zhang, Lin Zhu, Chen Cheng, Hongzhu Shi, Yaliang Li, Bolin Ding, and Jingren Zhou. 2024. AgentScope: A Flexible yet Robust Multi-Agent Platform. [arXiv:2402.14034](https://arxiv.org/abs/2402.14034)
- [14] Szymon Grzelak, Anna Z Czarna, et al. 2022. Social Value Orientation as a Predictor of Cooperation in Social Dilemmas. *Current Psychology* 41 (2022), 5738–5748.
- [15] Oliver P. John and Sanjay Srivastava. 1999. The Big Five Trait Taxonomy: History, Measurement, and Theoretical Perspectives. In *Handbook of Personality: Theory and Research*. Guilford Press, 102–138.
- [16] David W. Johnson and Roger T. Johnson. 2005. New Developments in Social Interdependence Theory. *Genetic, Social, and General Psychology Monographs* 131, 4 (2005), 285–358.
- [17] Daniel Kahneman and Amos Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47, 2 (1979), 263–291.
- [18] Annie Kirk, Robert West, et al. 2023. Value pluralism in multi-agent LLM societies. *arXiv preprint arXiv:2310.11442* (2023).
- [19] Sarah Knight, Kevin Rowe, et al. 2023. Alignment in Multi-Agent Large Language Model Systems. *arXiv preprint arXiv:2311.00010* (2023).
- [20] Joel Z Leibo, Vinicius Zambaldi, et al. 2017. Multi-agent reinforcement learning in sequential social dilemmas. *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS)* (2017), 464–473.
- [21] Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. Camel: Communicative agents for “mind” exploration of large language model society. *Advances in Neural Information Processing Systems* 36 (2023), 51991–52008.
- [22] Junkai Li, Yungwei Lai, Weitao Li, Jingyi Ren, Meng Zhang, Xinhui Kang, Siyu Wang, Peng Li, Ya-Qin Zhang, Weizhi Ma, and Yang Liu. 2025. Agent Hospital: A Simulacrum of Hospital with Evolvable Medical Agents. [arXiv:2405.02957](https://arxiv.org/abs/2405.02957)
- [23] Nian Li, Chen Gao, Mingyu Li, Yong Li, and Qingmin Liao. 2024. Econagent: large language model-empowered agents for simulating macroeconomic activities. In *Proc. of the 62nd Annual Meeting of ACL (Vol. 1: Long Papers)*. 15523–15536.
- [24] Michael W Macy and Robert Willer. 2002. From factors to actors: Computational sociology and agent-based modeling. *Annual Review of Sociology* 28, 1 (2002), 143–166.
- [25] Abraham Harold Maslow. 1943. A theory of human motivation. *Psychological review* 50, 4 (1943), 370.
- [26] Ryan O Murphy, Kurt A Ackermann, and Michel J J Handgraaf. 2011. Measuring social value orientation. *Judgment and Decision Making* 6, 8 (2011), 771–781.
- [27] Yohei Nakajima. 2023. BabyAGI. <https://github.com/yoheinakajima/babyagi>.
- [28] 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 *Proc. of the 36th Annual ACM Symp. on User Interface Software and Technology*. 1–22.
- [29] Richard M. Ryan and Edward L. Deci. 2000. Self-Determination Theory and the Facilitation of Intrinsic Motivation, Social Development, and Well-Being. *American Psychologist* 55, 1 (2000), 68–78.
- [30] Wilko Schwarting, Alyssa Pierson, Javier Alonso-Mora, Sertac Karaman, and Daniela Rus. 2019. Social behavior for autonomous vehicles. *Proceedings of the National Academy of Sciences* 116, 50 (2019), 24972–24978. <https://doi.org/10.1073/pnas.1820676116>
- [31] Tianmin Shu, Marta Kryven, Tomer D Ullman, and Joshua B Tenenbaum. 2020. Adventures in flatland: Perceiving social interactions under physical dynamics. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 42.
- [32] Aaditya Singh, Adam Fry, Adam Perelman, Adam Tart, Adi Ganesh, Ahmed El-Kishky, Aidan McLaughlin, Aiden Low, AJ Ostrow, Akhila Ananthram, et al. 2025. OpenAI GPT-5 System Card. *arXiv preprint arXiv:2601.03267* (2025).
- [33] Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805* (2023).
- [34] Paul AM Van Lange. 1999. The pursuit of joint outcomes and equality in outcomes: an integrative model of social value orientation. *Journal of personality and social psychology* 77, 2 (1999), 337.
- [35] Paul AM Van Lange, Ellen De Bruin, Wilma Otten, and Jeffrey A Joireman. 1997. Development of prosocial, individualistic, and competitive orientations: theory and preliminary evidence. *Journal of Personality and social psychology* 73, 4 (1997), 733.
- [36] Paul A M van Lange, Jeff Joireman, et al. 2021. Social Value Orientation: Past, Present, and Future. *Group Processes & Intergroup Relations* 24, 5 (2021), 701–705.
- [37] Alexander Sasha Vezhnevets, John P Agapiou, Avia Aharon, Ron Ziv, Jayd Matyas, Edgar A Duéñez-Guzmán, William A Cunningham, Simon Osindero, Danny Karmon, and Joel Z Leibo. 2023. Generative agent-based modeling with actions grounded in physical, social, or digital space using Concordia. *arXiv preprint arXiv:2312.03664* (2023).
- [38] Jiawei Wang, Renhe Jiang, Chuang Yang, Zengqing Wu, Makoto Onizuka, Ryosuke Shibasaki, Noboru Koshizuka, and Chuan Xiao. 2024. Large language models as urban residents: An llm agent framework for personal mobility generation. *Advances in Neural Information Processing Systems* 37 (2024), 124547–124574.
- [39] Lei Wang, Heyang Gao, Xiaohu Bo, Xu Chen, and Ji-Rong Wen. 2025. YuLan-OneSim: Towards the Next Generation of Social Simulator with Large Language Models. (2025). [arXiv:2505.07581](https://arxiv.org/abs/2505.07581)
- [40] Yiding Wang, Yuxuan Chen, Fangwei Zhong, Long Ma, and Yizhou Wang. 2024. Simulating human-like daily activities with desire-driven autonomy. *arXiv preprint arXiv:2412.06435* (2024).
- [41] Sean Welleck, Jason Weston, Arthur Szlam, and Kyunghyun Cho. 2019. Dialogue natural language inference. In *Proc. of the 57th Annual Meeting of ACL*. 3731–3741.
- [42] Tim Williams, An Truong, et al. 2021. Modeling Social Preferences in Human-Robot Teams. In *Proc. of the 2021 ACM/IEEE Int'l Conf. on Human-Robot Interaction (HRI)*. 149–157.
- [43] Jerry Xu, Tianqi Yu, Yang Ge, et al. 2023. Voyager: An Open-Ended Embodied Agent with LLM-Powered Autonomy. *arXiv preprint arXiv:2305.16291* (2023).
- [44] Shuhang Xu and Fangwei Zhong. 2025. Comet: Metaphor-driven covert communication for multi-agent language games. In *Proc. of the 63rd Annual Meeting of ACL (Vol. 1: Long Papers)*. 7892–7917.
- [45] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388* (2025).
- [46] Yuzhe Yang, Yifei Zhang, Minghao Wu, Kaidi Zhang, Yunmiao Zhang, Honghai Yu, Yan Hu, and Benyou Wang. 2025. Twinmarket: A scalable behavioral and social simulation for financial markets. *arXiv preprint arXiv:2502.01506* (2025).
- [47] Ziyi Yang, Zaibin Zhang, Zirui Zheng, Yuxian Jiang, Ziyue Gan, Zhiyu Wang, Zijian Ling, Jinsong Chen, Martz Ma, Bowen Dong, et al. 2024. OASIS: Open Agent Social Interaction Simulations with One Million Agents. *arXiv preprint arXiv:2411.11581* (2024).
- [48] Shunyu Yao, Jeff Zhao, Dian Yu, et al. 2023. ReAct: Synergizing Reasoning and Acting in Language Models. In *Advances in Neural Networks*.
- [49] Ceyao Zhang, Kaijie Yang, Siyi Hu, Zihao Wang, Guanghe Li, Yihang Sun, Cheng Zhang, Zhaowei Zhang, Anji Liu, Song-Chun Zhu, et al. 2024. Proagent: building proactive cooperative agents with large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 17591–17599.
- [50] Zhaowei Zhang, Ceyao Zhang, Nian Liu, Siyuan Qi, Ziqi Rong, Song-Chun Zhu, Shuguang Cui, and Yaodong Yang. 2023. Heterogeneous value alignment evaluation for large language models. *arXiv preprint arXiv:2305.17147* (2023).
- [51] Zheyuan Zhang, Daniel Zhang-Li, Jifan Yu, Linlu Gong, Jinchang Zhou, Zhanxin Hao, Jianxiao Jiang, Jie Cao, Huiqin Liu, Zhiyuan Liu, et al. 2025. Simulating classroom education with llm-empowered agents. In *Proc. of NAACL-HLT 2025 (Vol. 1: Long Papers)*. 10364–10379.
- [52] Zhou Ziheng, Huacong Tang, Mingjie Bi, Yipeng Kang, Wanying He, Fang Sun, Yizhou Sun, Ying Nian Wu, Demetri Terzopoulos, and Fangwei Zhong. 2025. An LLM-based Agent Simulation Approach to Study Moral Evolution. *arXiv preprint arXiv:2509.17703* (2025).