

A Survey on the Application of Agentic AI in Gaming

Ling Xu

*Institute of Physics, Truman State University, Kirksville, USA
xu62284@truman.edu*

Abstract. Since the release of Practices for Governing Agentic AI Systems by OpenAI in 2023, the emerging paradigm of Agentic AI (Artificial Intelligence) has gradually attracted academic attention. Unlike traditional AI systems that rely on structured presets and extensive human intervention, Agentic AI refers to intelligent systems capable of understanding task objectives, adapting to complex environments, and autonomously completing tasks with minimal human oversight. Its advantages in adaptability, decision-making, and self-management make it particularly suited to dynamic and rapidly changing real-world scenarios. This paper first outlines the definition, core characteristics, and primary enabling technologies of Agentic AI, emphasizing that, to date, no deployed system in the gaming domain fully embodies all of its defining features. To explore potential pathways toward practical implementation, the paper analyzes several recent large-model-based agents—such as SIMA, and ChatRPG v2, that, while falling short of the full Agentic AI standard, exhibit partial alignment through capabilities such as autonomous instruction comprehension, long-term task execution, multi-environment adaptation, and low-intervention deployment. These systems can therefore be regarded as generative agents with Agentic characteristics, offering valuable insights for future research directions and technical architectures. Finally, the paper proposes using gaming environments as low-cost, low-risk, and highly controllable testbeds for validating key capabilities, accumulating deployment experience, and accelerating the transition of Agentic AI into real-world applications, thereby advancing game AI from scripted logic toward higher levels of autonomous intelligence.

Keywords: Agentic AI, Agentic gaming, Multi-Agent Systems, Reinforcement Learning, Multimodal Perception and Interaction

1. Introduction

Artificial intelligence has come a long way, from those early rule-based expert systems to today's sophisticated deep learning models, and it has achieved some impressive successes in numerous fields. But most traditional AI systems are still pretty much like obedient assistants—they follow instructions without making independent decisions. They don't have the capability to set their own goals, change strategies as needed, or keep learning in complex and ever-changing environments [1]. Recently, though, there have been developments that open up possibilities for AI to behave more independently. The advancements in large-scale language models, the progress in multi-agent

systems, and the successful automation in complex tasks have all helped lay the groundwork for AI systems that can make their own decisions. In 2023, OpenAI came up with the idea of Agentic AI in their publication Practices for Governing Agentic AI Systems. They provided some initial definitions and guidelines for managing it [2]. Unlike traditional AI, which is heavily reliant on pre-written scripts and constant human oversight, Agentic AI centers around agents that can grasp tasks, devise plans, and accomplish objectives with very little human involvement, even when things are uncertain and constantly changing.

Gaming applications present intriguing possibilities for this technology. The evolution of game AI illustrates its progression via three key stages. Initially, systems employed finite-state machines to manage NPC actions and reactions; however, these were inflexible and unable to adapt [3]. In the second stage, graph search techniques, behavior trees, and planning systems were introduced, granting AI greater flexibility and goal-directed behavior [4]. Recent advancements encompass machine learning approaches like reinforcement learning and evolutionary strategies, resulting in remarkable systems such as AlphaGo, OpenAI Five, and AlphaStar [5, 6]. Nevertheless, existing game AI systems still have crucial limitations. They rely on offline training and struggle to adapt to new circumstances in real-time. Most critically, they concentrate on accomplishing specific tasks instead of exhibiting the sort of long-term motivation and autonomy that would render them genuinely independent.

The paper looks at three crucial aspects of Agentic AI in gaming. Firstly, the paper explains what Agentic AI entails and pinpoint the primary technical parts that underpin autonomous agents. Secondly, the paper scrutinizes recent systems built upon large language models, such as SIMA, PORTAL, and PANGeA. Even though these systems haven't exhibited all the characteristics of genuine Agentic AI yet, they showcase vital elements like autonomous comprehension of tasks, execution over a prolonged period, and adaptation across various environments. Lastly, the paper maintains that games serve as splendid testing grounds for Agentic AI development. They present settings with low costs and low risks, enabling researchers to test essential technologies and deployment strategies prior to applying them to more formidable real-world scenarios.

2. Agentic AI

2.1. Definition and distinction

Agentic AI represents a new breed of autonomous systems capable of setting and working toward long-term goals in messy, real-world environments. What makes these systems distinctive is their ability to think several steps ahead—they take big, ambitious objectives and figure out how to break them down into manageable pieces, then actually follow through on executing those plans [7].

The key difference from traditional AI lies in flexibility. Most conventional AI systems follow predetermined rules and work well only within their specific domains. But agentic AI can pivot when circumstances change. Think of it this way: while tools like ChatGPT wait for you to ask them something before responding, agentic AI is constantly planning and acting on its own initiative [7]. That said, not all agents work the same way. Some are built for very specific tasks and operate independently. However, the more interesting applications often involve multiple agents working together. By maintaining memory of past experiences, adapting their approach to new situations, and coordinating with other agents, these systems can tackle complex, interconnected problems that would overwhelm simpler AI approaches [8].

2.2. Characteristics

2.2.1. High autonomy

Autonomy is a crucial characteristic of Agentic AI, especially useful in situations with many complex goals. Agentic AI can do several tasks by itself, from easy ones to lots of hard goals. It kind of runs itself, following its own rules or changing them as it goes, even if people aren't always watching. It's not just about finishing one thing; it's also about being able to switch between different smaller goals and plans to reach a big, complicated goal over time. For example, in regular games, the computer characters usually act the way the game makers tell them to, like attacking when they see an enemy. But with Agentic AI, the characters can decide for themselves what to do next based on what's happening, so they act in a smarter, more planned-out way.

2.2.2. Situational awareness and adaptability

Situational awareness and adaptability enable Agentic AI to interpret environmental states in real time and to adjust its behavior and strategies as conditions change. These capabilities rely not only on environment modeling and contextual reasoning, but also on the integration of prior experience and external information. In multi-objective tasks, situational awareness helps the agent make trade-offs under conflict or unexpected events, prioritizing the most critical issues [7]. Where conventional systems lean on static evaluation tables to drive predefined policies in complex strategy games, an agentic agent, once a tactical shift is perceived, moves to rebalance offense and defense and to reframe its priority stack.

2.2.3. Complex goal management

Complex goal management is where agentic AI really shows its strength. These systems don't just break down big tasks into smaller ones; they actively juggle multiple objectives, reorganize priorities when needed, and keep everything coordinated as conditions change. Most traditional AI systems work more like specialized tools: give them a specific job, and they'll do it well, but they struggle when the situation shifts or when multiple goals conflict with each other. What makes agentic AI different is its ability to rethink its approach mid-process. If progress stalls on one objective, or if new information suggests a better path forward, the system can restructure its entire plan. This flexibility becomes essential when dealing with interconnected goals where success in one area might affect others [7]. Take open-world games as a concrete example. Traditional game systems typically use fixed event triggers. Complete quest A to unlock quest B, reach location X to start storyline Y. But agentic AI can create entirely new quest chains based on how a particular player actually behaves. If someone spends most of their time exploring rather than following the main story, the system might generate exploration-based challenges that tie into the broader narrative. The result feels less scripted and more responsive to individual play styles. This kind of adaptive goal management has obvious applications beyond gaming, particularly in complex project management or resource allocation scenarios where initial plans rarely survive contact with reality.

2.2.4. Continuous learning and self-improvement

One of the most compelling aspects of agentic AI is its capacity for continuous learning and self-improvement. Rather than remaining fixed after training, these systems actively learn from each interaction, refining their decision-making through online learning, transfer learning, and meta-

learning approaches. This contrasts sharply with static models that remain unchanged once deployed. Agentic AI combines lessons from past experiences with real-time feedback to build comprehensive knowledge bases, while using reflective processes to identify mistakes and adjust strategies [7]. This creates a genuine learning loop where performance improves over time. Consider real-time strategy games as an example. Traditional AI opponents struggle when players employ novel tactics because they lack adaptation capabilities beyond their original programming. Agentic AI, however, can analyze patterns across multiple gaming sessions, develop countermeasures, and implement more sophisticated strategies in subsequent matches. This leads to strategic evolution where the AI becomes progressively more challenging.

2.2.5. Multi-agent collaboration

Multi-agent collaboration stands out as one of agentic AI's defining features. When dealing with complex systems, individual agents often cannot handle everything alone, so tasks get split up among multiple agents. For this to work effectively, agents need ways to share information and coordinate their use of available resources. Agentic AI creates good communication ways in these multi-agent places. It coordinates what different agents do and their duties. This makes the whole system work better and be stronger [7]. In team competitive games, old-style AI teammates usually act by themselves. They have little big-picture coordination. But Agentic AI can make collaborative combat nets. It does this by using shared local smarts and giving out tasks in real time. This lets it do advanced tactics like strategic flanking and defense that fits together well.

2.2.6. Long-term reasoning and planning

Long-term reasoning and planning capabilities empower Agentic AI to keep goal consistency over prolonged periods when carrying out multi-step action simulations and optimizations for future situations [7]. This capability necessitates that agents assess both immediate advantages and the influence of future states on overarching goals. For example, in simulation management games, conventional AI decision-making is frequently restricted to short-term resource maximization, while Agentic AI might forego immediate gains for long-term developmental potential, like prioritizing infrastructure investment to boost subsequent production efficiency and accomplish strategic progress.

3. Technical foundations

The principal enabling technologies for Agentic AI currently include reinforcement learning, multi-agent system architectures, large language models, multimodal perception and interaction, simulation and digital twin technologies, as well as planning and reasoning techniques [9]. These approaches collectively support the essential capabilities of Agentic AI across dimensions such as autonomous learning, task transfer, structural management, intrinsic motivation for exploration, and long-term memory.

In this study, the discussion of core technologies focuses on three specific categories: reinforcement learning and inverse reinforcement learning, multi-agent systems, and multimodal perception and interaction. These were selected because they correspond directly to critical functional domains of Agentic AI, namely autonomous decision-making, collaboration and task decomposition, and environmental understanding and interaction. Together, they offer a representative view of the fundamental characteristics of Agentic AI. While other related

technologies are likewise important for implementation, they are not elaborated here due to the scope and emphasis of the present work.

3.1. Reinforcement learning and inverse reinforcement learning

Reinforcement learning (RL) is one of the most widely applied foundational techniques in the field of artificial intelligence. It is a sequential decision optimization approach in which an agent interacts with its environment in a state–action–feedback cycle, updating its policy based on the reward signals received from the environment in order to maximize the long-term cumulative return [10]. Variants include Q-Learning, Deep Q-Networks (DQN), and Actor–Critic methods. Inverse reinforcement learning (IRL), by contrast, models the reward function in reverse, inferring the underlying implicit reward structure from observed expert trajectories, thereby providing a target representation for policy learning when explicit reward definitions are unavailable [10]. Within the context of Agentic AI, RL equips agents with the capacity to improve autonomously through trial and error in unfamiliar and dynamic environments, supporting the decomposition and optimization of long-term tasks. IRL enables agents to rapidly capture and replicate expert preferences and goals in the absence of explicit reward signals, achieving a closed-loop adaptation from autonomous exploration to imitation learning.

In gaming scenarios, this technological combination can provide agents with behavioural mechanisms that both optimise independently and align with player preferences. For instance, in open-world or sandbox games, RL agents may learn optimal balance between exploration and resource management through multi-round interactions, guided by rules such as task completion (reward), acquisition of rare resources (reward), and avoiding health reduction (negative reward). Conversely, IRL agents can analyse players' movement patterns, combat decisions, and resource allocation during cooperative tasks to infer implicit objectives such as prioritising teammate support or concentrating attacks on high-threat targets, subsequently adjusting their own strategies accordingly.

3.2. Multi-agent systems

Multi-agent systems (MAS) are distributed systems comprising numerous autonomous agents, each possessing capabilities for perception, reasoning, and action execution. These agents can collectively address complex problems through collaboration, competition, or negotiation [11]. The fundamental components of a MAS include the agents themselves, their operating environment, communication protocols, and coordination mechanisms. There are three types of system architectures for MAS: centralized ones, distributed ones, and hybrid forms. In places that are really changeable and full of uncertainties, distributed architectures tend to be more robust and scalable. On the other hand, centralized architectures are better at being efficient when it comes to making plans for the whole system and assigning tasks [11]. Inside agentic AI frameworks, MAS gives the ability to coordinate multiple entities and break down tasks into smaller pieces. This lets the system allocate resources dynamically and make decisions in real time in situations with lots of dimensions and many different goals at once. This way, it can handle complex tasks that would be too hard for just one agent to manage.

In gaming contexts, exemplary applications of MAS technology are evident in the AlphaStar system. During StarCraft II battles, AlphaStar's framework involves training distinct agents to specialise in particular strategies and tactics (including rush attacks, economic focus, defensive counter-strategies, etc.), with these agents sharing experience and competitively refining strategies

under a central meta-learning structure, automatically assigning roles and tasks across different scenarios to achieve varied combat approaches [5]. This multi-agent collaborative approach allows the system to adapt dynamically to opponents' tactical modifications whilst accomplishing resource management, unit coordination, and tactical implementation without human real-time guidance. For instance, some agents concentrate on early resource gathering, whilst others focus on combat unit assembly and attack pathway planning, maintaining equilibrium between global strategy and local tactics through central orchestration. This example illustrates that MAS architectures can effectively enhance system adaptability and decision-making sophistication in complex, real-time, adversarial gaming environments, offering practical directions for future Agentic AI implementation in open-world and highly complex games.

3.3. Multimodal perception and interaction

Multimodal perception refers to intelligent agents' capacity to simultaneously process inputs from multiple heterogeneous information sources, such as visual, linguistic, auditory, and action trajectory data, integrating these into unified internal representations to enable comprehension and decision-making within more complex and dynamic environments [12]. Multimodal interaction builds upon this foundation, allowing agents to engage in bidirectional communication with users or environments through various modalities, including natural language exchange, motion or gestural responses, and environmental feedback [12]. In Agentic AI contexts, multimodal perception significantly enhances agents' depth of environmental understanding and contextual association abilities, enabling cross-modal reasoning and generation of behaviours that better align with situational demands. Meanwhile, multimodal interaction allows agents to dynamically acquire supplementary information during task execution, adjust strategies, and achieve greater fluidity in human-machine collaboration.

In gaming scenarios, multimodal perception and interaction can equip agents with the ability to understand and respond across modalities. For example, in immersive virtual reality (VR) games, an agent can integrate visual input from the player's field of view, spoken commands, and motion capture data of hand or body movements to infer player intent, and then provide feedback through real-time verbal responses, coordinated actions, or environmental changes. In large-scale open-world games, an agent with multimodal capabilities could not only understand a player's natural language instructions but also combine visual map data, terrain information, and quest scripts to plan paths and allocate tasks, maintaining natural interaction with the player while executing objectives.

4. Case analysis of generative agents with agentic characteristics

At present, there are no implemented systems in the gaming field that fully meet all the core characteristics of Agentic AI. The SIMA and ChatRPG v2 systems selected in this paper are not Agentic AI in the strict sense, but they are generative agents that possess some Agentic characteristics. They are chosen as case studies because they have significant overlap with the core capabilities of Agentic AI in aspects such as autonomous instruction comprehension, cross-task execution, and adaptation to multiple environments. In addition, their technical implementations make extensive use of methods closely related to the enabling technologies of Agentic AI, including multimodal perception, large model reasoning, and tool invocation. These systems can therefore provide useful architectural designs and training approaches for gradually building more complete Agentic AI in games. The following analysis examines their system background, technical methods,

Agentic characteristics, and shortcomings, and explores their implications for future research on Agentic AI.

4.1. Scalable instructable multiworld agent

SIMA, proposed by Google DeepMind in 2024, aims to test language-driven action planning and generalization across multiple environments [13]. Its design principle is not to rely on internal game application interfaces or a simplified action space, but to interact with virtual environments in the same way as human players, taking screen image input and producing keyboard and mouse control output [13]. This approach significantly improves cross-environment generality while also increasing task complexity. The training data of SIMA covers more than ten commercial and research-oriented three-dimensional environments, including open-world games such as No Man's Sky, Teardown, and Valheim, as well as research-developed virtual environments such as Playhouse and Construction Lab. These environments differ considerably in visual style, interaction rules, and physical mechanisms, and the task types include resource collection, construction, combat, and exploration. SIMA can parse natural language instructions and convert them into continuous low-level action sequences [13]. Initial experiments show that SIMA can still perform tasks such as gathering, navigation, and crafting in unseen game levels [13], demonstrating a strong ability to generalize tasks. However, SIMA has important limitations. Its task objectives rely entirely on external input and it lacks the ability to generate its own goals [13]. When performing long-duration tasks, it finds it difficult to autonomously replan in response to environmental changes or intermediate failures [13]. Its ability to transfer knowledge across tasks has not yet been established, and it lacks a persistent memory mechanism that can be shared across different tasks and environments [13]. Therefore, although SIMA is not complete Agentic AI, it provides important reference for the future introduction of autonomous goal setting, long-term memory, and self-reflection mechanisms in terms of multimodal perception and task generalization.

4.2. ChatRPG v2

ChatRPG v2 is a multi-agent game master system focused on interactive storytelling, mainly designed for role-playing scenarios similar to Dungeons and Dragons [14]. Compared with version 1, which only used prompt engineering, it introduces a multi-agent architecture based on the ReAct framework, dividing work between two cooperating roles, the Narrator and the Archivist [14]. The Narrator is the main storyteller facing the player, responsible for interpreting player input, generating plot developments, and invoking tools such as combat, healing, and injury when needed to advance changes in the game world [14]. For example, when the player inputs "I attack the monster in front of me with my sword," the Narrator will determine that the combat tool should be triggered and, based on the hit and damage results returned by the tool, produce a narrative description such as "You swing your sword at the monster. It staggers back in pain, blood seeping from the wound." At the same time, the Archivist continuously updates the world state in the background, recording new locations, characters, and events to ensure consistency in subsequent storytelling [14]. For example, when the player enters a new town, the Archivist will store the description of the location and related NPC information so that when the player later asks "What was the name of the blacksmith shop I visited earlier?" the system can answer accurately. Through this division of roles, ChatRPG v2 solves the problem in version 1 of losing context in long-term narratives, and user studies show that it significantly improves immersion and replayability. However, ChatRPG v2 still depends on presets or player input for generating global task objectives

and main storylines, and lacks the ability to generate them fully autonomously [14]. Knowledge from one task cannot be transferred to different stories or sessions, and character consistency is limited to a single script [14]. In addition, although ReAct enables the Narrator to perform immediate reasoning and dynamic action selection, its planning is mostly local and reactive rather than oriented toward globally optimal long-term strategies [14].

4.3. Summary

In conclusion, although SIMA and ChatRPG v2 do not have the full capabilities of Agentic AI, they provide important practical references in aspects such as multimodal understanding, multi-environment adaptation, real-time reasoning, and multi-agent collaboration. SIMA demonstrates how to maintain generality in complex perception and multi-environment contexts, while ChatRPG v2 verifies the feasibility of multi-agent division of labor and long-term narrative consistency in interactive entertainment. These cases suggest that by introducing autonomous goal generation, cross-task knowledge transfer, long-term memory, and self-reflection mechanisms into existing generative agents, it is possible to promote the gradual evolution of game AI from partial intelligence to true Agentic AI.

5. Challenge and future

5.1. Challenges

Although generative agents with Agentic characteristics have shown notable progress in some capabilities, they still fall significantly short of complete Agentic AI.

5.1.1. Incomplete feature set

Current systems often lack essential capabilities such as stable long-term memory, cross-task transfer, and autonomous goal generation, which limits their performance in complex and sustained tasks. To address this, a multi-level memory architecture can be introduced that integrates episodic memory with semantic memory, combined with meta-learning mechanisms to enhance cross-task adaptability. It is also necessary to explore enabling agents to autonomously generate and adjust goals during task execution.

5.1.2. Controllability and safety

In open-ended tasks, the decision-making pathways and behavioral outcomes of agents are difficult to predict and fully constrain, which creates potential safety and alignment risks. This can be addressed by establishing secure sandbox environments and restrictive API call mechanisms, combined with interpretable decision modules and reversible behavior logging, to strengthen human intervention at critical decision points.

5.1.3. Computational and resource demands

Highly autonomous agents operating over long durations require substantial computational resources, which hinders large-scale deployment. Model distillation and pruning techniques can be used to achieve lightweight implementations, supplemented by distributed inference, edge computing, and task prioritization mechanisms to reduce operating costs.

5.1.4. Lack of evaluation standards

There is currently no unified and repeatable evaluation framework for Agentic AI, which makes it difficult to compare and optimize different systems objectively. It is therefore necessary to establish standardized benchmark tests based on task diversity and complexity, design multi-dimensional metrics covering task success rate, environmental adaptability, and safety, and promote the creation of shared test environments and datasets across institutions to facilitate comparability and reproducibility of research outcomes.

5.2. Technological development pathway

The future evolution of Agentic AI in gaming can follow a path of progressive feature enhancement, capability integration, and full-feature realization. This entails first expanding the scope of individual capabilities, then effectively integrating multiple core capabilities, and ultimately achieving a complete system with all key features while ensuring safety and controllability.

5.3. Cross-domain transfer and application prospects

As the technology matures, Agentic AI models and architectures that have been validated in gaming can be transferred to higher-risk and more complex real-world task environments, such as automated industrial control, unmanned vehicles, and virtual education and training. The advantage of repeatable experimentation in gaming scenarios will accelerate algorithm optimization and iteration, providing data accumulation and safety assurance for real-world deployment. However, cross-domain transfer must also address differences in environments, inconsistencies in task definitions, and stricter real-world constraints. Therefore, establishing a unified capability evaluation and safety review framework before transfer will be an essential prerequisite.

6. Conclusion

This paper has examined the potential applications of Agentic AI in the gaming domain. It first defined Agentic AI and summarized its core characteristics, including autonomy, situational awareness, complex goal management, multi-agent collaboration, and long-term planning. It then analyzed key enabling technologies such as reinforcement learning, multi-agent architectures, and multimodal perception and interaction. Furthermore, through the analysis of generative agents with Agentic characteristics, such as SIMA and ChatRPG v2, the study revealed their practical value in multi-environment adaptability, real-time reasoning, and collaborative role division, while also identifying their limitations in autonomous goal generation, long-term memory, and cross-task transfer.

The paper also highlighted that gaming environments, due to their low cost, low risk, and high controllability, are well suited as testbeds for Agentic AI, providing an ideal platform for verifying key capabilities, iterating technical architectures, and accumulating deployment experience. However, current systems still face challenges, including incomplete features, insufficient safety and controllability, high resource consumption, and the absence of standardized evaluation frameworks. Future technological development may progress along a pathway of progressive feature enhancement, capability integration, and full-feature realization, while also exploring cross-domain transfer to apply Agentic AI models validated in gaming to real-world scenarios such as industrial control, unmanned systems, and virtual training.

References

- [1] Lake, B.M., Ullman, T.D., Tenenbaum, J.B. and Gershman, S.J. (2017) Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, e253.
- [2] Shavit, Y., Agarwal, S., Brundage, M., Adler, S., O’Keefe, C., Campbell, R., Mishkin, P., Ahmad, L., Eloundou, T., McMillan, P., Hickey, A., Beutel, A., Lee, T., Slama, K., Passos, A., and Robinson, D.G. (2023) Practices for governing agentic AI systems. Research Paper, OpenAI.
- [3] Jagdale, D. (2021) Finite state machine in game development. *International Journal of Advanced Research in Science, Communication and Technology*, 10(1).
- [4] Sekhvat, Y.A. (2017) Behavior trees for computer games. *International Journal on Artificial Intelligence Tools*, 26(02), 1730001.
- [5] Vinyals, O., Babuschkin, I., Czarnecki, W.M., et al. (2019) Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575, 350-354.
- [6] Berner, C., Brockman, G., Chan, B., Cheung, V., Dębiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., Józefowicz, R., Gray, S., Olsson, C., Pachocki, J., Petrov, M., Ponde de Oliveira Pinto, H., Raiman, J., Salimans, T., Schlatter, J., Schneider, J., Sidor, S., Sutskever, I., Tang, J., Wolski, F., and Zhang, S. (2019) Dota 2 with large scale deep reinforcement learning. arXiv preprint, arXiv: 1912.06680.
- [7] Acharya, D.B., Kuppan, K. and Divya, B. (2025) Agentic AI: Autonomous intelligence for complex goals – A comprehensive survey. *IEEE Access*.
- [8] Pati, A.K. (2025) Agentic AI: A comprehensive survey of technologies, applications, and societal implications. *IEEE Access*.
- [9] Pati, A.K. (2025) Agentic AI: A comprehensive survey of technologies, applications, and societal implications. *IEEE Access*.
- [10] Adams, S., Cody, T. and Beling, P.A. (2022) A Survey of Inverse Reinforcement Learning. *Artificial Intelligence Review*, 55(6), 4307-4346.
- [11] Maldonado, D., Cruz, E., Torres, J.A., Cruz, P.J. and Benitez, S.D.P.G. (2024) Multi-Agent Systems: A Survey about Its Components, Framework and Workflow. *IEEE Access*, 12, 80950-80975.
- [12] Durante, Z., Huang, Q., Wake, N., Gong, R., Park, J.S., Sarkar, B., Taori, R., Noda, Y., Terzopoulos, D., Choi, Y., Ikeuchi, K., Vo, H., Fei-Fei, L. and Gao, J. (2024) Agent AI: Surveying the Horizons of Multimodal Interaction. arXiv preprint arXiv: 2401.03568.
- [13] SIMA Team, Abi Raad, M. et al. (2024) Scaling Instructable Agents across Many Simulated Worlds. arXiv preprint arXiv: 2404.10179.
- [14] Jørgensen, N.H., Tharmabalan, S., Aslan, I., Hansen, N.B. and Merritt, T. (2025) Static vs. agentic game master AI for facilitating solo role-playing experiences. arXiv preprint arXiv: 2502.19519.