

Agentic AI: Transforming Autonomous Decision-Making in Intelligent Systems

K.Petchiammal¹, S.Subha², M.Pon Abisha³, M.Maheswari⁴, K.Rama Subramanian⁵, P.Syed Mohamed Mydeen⁶

¹Assistant Professor, J.P.College Of Engineering, Ayikudy, Tenkasi-627852, Affiliated To Anna University, Chennai-6000252.

^{2,3,4,5,6}Student, Department Of Information Technology, J.P.College Engineering, Ayikudy, Tenkasi 627852, Affiliated To Anna University,

Abstract- Agentic Artificial Intelligence (Agentic AI) is an emerging approach that enables intelligent systems to perform tasks with greater autonomy by combining reasoning, planning, memory, tool usage, and decision-making capabilities. Unlike conventional AI systems that mainly respond to predefined instructions, Agentic AI systems can interpret objectives, divide complex tasks into smaller steps, execute appropriate actions, evaluate outcomes, and modify their strategies when required. This paper presents a closed-loop framework for autonomous decision-making based on goal analysis, task planning, reasoning, tool selection, execution, feedback, verification, and adaptive re-planning. The proposed approach aims to improve the ability of intelligent systems to handle multi-step and changing tasks rather than producing a single response. The paper also discusses existing research, the research gap, experimental evaluation criteria, challenges, limitations, and future opportunities. The framework provides a structured foundation for developing more adaptive, reliable, and efficient autonomous intelligent systems.

Keywords- Agentic AI, Autonomous Decision-Making, Artificial Intelligence, Large Language Models, AI Agents, Planning, Reasoning, Tool Usage

I. INTRODUCTION

Artificial Intelligence (AI) has developed significantly from traditional rule-based systems to advanced models capable of learning, reasoning, prediction, and content generation. Modern AI systems can process large amounts of information and provide useful responses for a wide range of applications. However, many conventional AI systems still depend on human instructions for determining the sequence of actions, selecting appropriate tools, and evaluating whether the final outcome satisfies the intended objective. Agentic AI introduces a more autonomous approach to intelligent systems. An AI agent can receive a high-level objective, analyze the requirements, divide the objective into smaller tasks, develop a plan, perform actions, observe the results, and modify its strategy when necessary. This combination of reasoning, planning, memory, tool interaction, and feedback allows AI systems to move beyond simple response generation toward goal-oriented task execution. Large Language Models (LLMs) have contributed

significantly to the development of AI agents. Research has explored how LLMs can support reasoning, planning, tool usage, memory, and interaction with external environments [1], [2]. Approaches such as reasoning-and-action cycles and tool-using systems have demonstrated the potential of language models to perform multi-step tasks.

Autonomous decision-making is particularly relevant in areas such as healthcare, finance, cybersecurity, robotics, transportation, and smart systems. In these environments, an intelligent system may need to process changing information and make decisions continuously. However, increased autonomy also introduces challenges related to reliability, hallucination, security, explainability, computational cost, and human oversight. Therefore, this research proposes a structured Agentic AI framework that connects goal analysis, task decomposition, planning, reasoning, tool selection, execution, feedback, verification, and adaptive re-planning in a continuous decision-making cycle. The objective is to provide a

systematic approach for studying and developing more adaptive autonomous intelligent systems.

II. RELATED WORK / LITERATURE REVIEW

Recent research has increasingly focused on using Large Language Models as the reasoning core of autonomous AI agents. Wang et al. presented a comprehensive survey of LLM-based autonomous agents and discussed major components such as planning, memory, reasoning, and tool usage [1]. Their work highlights the transition from conventional language-model applications toward systems capable of performing more complex tasks autonomously. Planning is an important component of agentic systems because complex objectives usually require multiple dependent actions. Research on LLM-agent planning has investigated task decomposition, sequential decision-making, plan generation, and plan refinement [3].

Effective planning can help an agent determine what actions should be performed and in what order. Tool usage has also become an important research direction. LLM-based agents can interact with external tools such as search engines, databases, APIs, calculators, and software environments. Tool-learning research examines how agents can select appropriate tools, generate suitable inputs, interpret outputs, and incorporate external information into subsequent decisions [4].

ReAct-based approaches combine reasoning and action, allowing an agent to alternate between internal reasoning and interactions with an external environment. This provides a foundation for handling tasks in which actions generate new information that influences later decisions. Similarly, feedback-oriented approaches such as Reflexion explore the use of previous outcomes to improve subsequent behavior. Multi-agent systems extend this concept by allowing multiple specialized agents to collaborate on complex problems. Recent surveys have examined agent workflows, communication mechanisms, infrastructure, and challenges associated with LLM-based multi-agent systems [5], [6]. Although these studies provide strong foundations, many approaches focus

primarily on individual capabilities such as planning, tool use, or collaboration. A more integrated approach is required in which planning, execution, feedback, verification, and re-planning operate as parts of a continuous decision cycle. The proposed research focuses on this integration.

III. RESEARCH GAP

Existing research on Agentic AI and Large Language Model (LLM)-based agents has explored several important capabilities, including reasoning, task planning, memory, tool usage, reflection, and multi-agent collaboration. These studies have demonstrated that LLMs can support intelligent systems in performing tasks that involve multiple steps and interactions with external environments. However, many existing approaches concentrate on individual capabilities rather than examining the complete autonomous decision-making process as an integrated cycle. A significant gap exists in the integration of planning, execution, feedback, verification, and adaptive re-planning. In practical environments, an agent may encounter unexpected results, incomplete information, tool failures, or changes in task requirements after the initial plan has been generated.

Systems that depend mainly on a fixed plan may therefore struggle to adapt effectively. A continuous feedback mechanism is required to evaluate intermediate outcomes and modify subsequent actions. Another gap concerns systematic evaluation of autonomous behavior. Traditional AI evaluation often focuses on final prediction accuracy or response quality. However, Agentic AI performs a sequence of decisions and actions, making it necessary to evaluate additional factors such as task completion, planning efficiency, tool-use success, failure recovery, adaptability, computational cost, and decision reliability. There is also a need to address the balance between autonomy and safety.

Increasing the number of actions an agent can perform may improve task automation but can also increase the possibility of incorrect or unsafe operations. Therefore, verification mechanisms,

permission controls, and human oversight need to be considered as part of the agent architecture rather than as separate components. Based on these limitations, this research proposes a closed-loop Agentic AI framework that connects goal analysis, task decomposition, planning, reasoning, tool interaction, execution, feedback, verification, and adaptive re-planning. The framework aims to provide a structured approach for evaluating how an autonomous agent can respond to intermediate outcomes and changing conditions while maintaining alignment with the original objective. Thus, the research gap addressed in this study is the need for an integrated, adaptive, and systematically evaluated framework for autonomous decision-making, rather than treating individual agent capabilities as isolated components.

IV. ANALYSIS

The proposed methodology is designed to enable an intelligent system to transform a high-level objective into an appropriate decision through a sequence of interconnected stages. Instead of generating an immediate response, the agent continuously analyzes the objective, develops a plan, performs actions, evaluates outcomes, and adapts its behavior.

4.1 Framework Overview

The process begins with goal analysis, where the system interprets the user's objective and identifies the expected outcome. The objective is then divided into smaller and manageable subtasks. Task decomposition reduces the complexity of the overall problem and allows individual tasks to be handled systematically. The next stage is planning. The agent determines the sequence of actions required to complete the subtasks while considering dependencies, available resources, and expected outcomes. Planning provides a structured path from the initial objective toward the desired result. The reasoning stage determines the most suitable action for each task. The agent evaluates available information and selects an appropriate strategy. If external information or functionality is required, the system proceeds to tool selection. In

the tool selection stage, the agent identifies suitable external resources such as databases, APIs, search systems, calculators, or domain-specific applications. The output obtained from a tool is returned to the agent and becomes part of the information used for subsequent reasoning.

4.2 Execution and Feedback

After selecting an action or tool, the system performs the required operation. The resulting information is then examined through a feedback mechanism. Feedback compares the observed result with the expected outcome. If the result is satisfactory, the system continues with the next task. If the result is incomplete, incorrect, or unexpected, the agent returns to the reasoning and planning stages. It can identify the affected task and select an alternative strategy. This feedback mechanism is important because autonomous systems may encounter unexpected conditions during execution. Rather than following an original plan without modification, the proposed framework allows the agent to respond to newly available information.

4.3 Verification and Adaptive Re-planning

Before producing the final decision, the system performs verification. The obtained result is compared against the original objective and predefined requirements. Verification helps identify incomplete or unsuitable outcomes. When verification identifies a problem, adaptive re-planning is activated. The agent modifies the existing plan using information obtained during previous execution steps. The revised plan is then executed and evaluated again. The complete workflow can therefore be represented as:

Goal Analysis → Task Decomposition → Planning → Reasoning → Tool Selection → Execution → Feedback → Verification → Re-planning → Final Decision

4.4 Decision-Making Process

The proposed framework treats autonomous decision-making as a continuous process rather than a single prediction step. Information obtained during execution influences subsequent actions.

This allows the system to respond to failures, changing conditions, and newly discovered information. The framework can also be adapted for different application domains by replacing or extending its tools, knowledge sources, constraints, and verification mechanisms. This provides flexibility for applications where autonomous decision-making requires domain-specific information.

4.5 Knowledge and Memory Management

Agentic AI systems require access to relevant information while performing multi-step tasks. The proposed framework incorporates a knowledge and memory layer to store useful task information, previous actions, intermediate results, and feedback. Retrieved information can be provided to the reasoning component when required. This helps the agent maintain context during long tasks and avoid repeating previously unsuccessful actions. Memory can also support learning from earlier interactions and improve the consistency of future decisions.

4.6 Safety and Human Oversight

Autonomous decision-making requires appropriate safety mechanisms, particularly when an agent performs actions that may have significant consequences. The proposed framework can include validation rules, permission controls, action verification, and human approval for high-risk operations. Before executing sensitive actions, the system can evaluate whether the action satisfies predefined constraints. Human intervention can also be requested when the agent has insufficient information or low confidence. This combination of autonomous execution and human oversight can improve the reliability and safe operation of Agentic AI systems.

4.7 Context-Aware Decision Making

The proposed framework considers the context surrounding each task before selecting an action. Context may include the current task state, available information, previous actions, tool outputs, system constraints, and user requirements. The agent uses this information to select actions that are relevant to the current situation. When new information

becomes available, the decision context is updated so that subsequent actions can be selected more appropriately. This mechanism helps the agent maintain consistency while handling multi-step and changing tasks.

4.8 Confidence-Based Action Selection

The framework can incorporate a confidence assessment before executing important decisions. The agent evaluates the reliability of the available information and the suitability of the planned action. When confidence is sufficiently high, the action can proceed automatically. If confidence is low or conflicting information is detected, the system can perform additional verification, reconsider the plan, or request human intervention. This approach can reduce the probability of unsafe or incorrect autonomous actions and improve the reliability of decision-making.

4.9 Agent State Management

Agent state management is used to maintain the current status of the task throughout the decision-making process. The system records completed subtasks, pending tasks, previous actions, tool outputs, intermediate results, and detected errors. This information is provided to the reasoning and planning components when required. Maintaining the agent state helps preserve task context and reduces the possibility of repeating unsuccessful actions. It is particularly useful for long-horizon tasks where multiple decisions must be performed sequentially.

4.10 Decision Monitoring and Logging

The proposed framework includes a monitoring and logging mechanism to track the agent's actions and decisions during execution. Important events such as generated plans, selected tools, tool outputs, feedback, verification results, errors, and re-planning activities can be recorded. These logs can be used to analyze agent behavior, identify failures, support debugging, and evaluate system performance. Decision monitoring also improves transparency by providing a record of how the final outcome was produced.

V. EXPERIMENTAL DESIGN AND PERFORMANCE EVALUATION

The proposed framework can be evaluated using multi-step tasks that require planning, reasoning, tool interaction, and adaptation. To determine its effectiveness, the framework can be compared with different baseline approaches under similar task conditions.

5.1 Experimental Setup

The evaluation can include information-based tasks, multi-stage decision-making tasks, tool-assisted tasks, and tasks containing intermediate failures or changing conditions. The same task definitions and evaluation criteria should be applied to each approach to ensure a fair comparison. The proposed framework can be implemented using an LLM as the reasoning component together with selected external tools. Each task should record the generated plan, actions performed, tool interactions, feedback received, recovery attempts, and final outcome.

5.2 Baseline Comparison

The proposed approach can be compared with a direct LLM approach, a planning-based agent, and a tool-integrated agent. A direct LLM produces an answer without an explicit autonomous execution cycle. A planning-based agent introduces task planning, while a tool integrated agent additionally interacts with external resources. The comparison helps determine whether combining planning, tool usage, feedback, verification, and adaptive re-planning provides additional benefits for autonomous task execution.

5.3 Evaluation Criteria

The system can be evaluated using task completion rate, decision accuracy, planning efficiency, tool execution success, adaptation capability, failure recovery, and computational cost. These measures provide information about both the quality of the final decision and the efficiency of the decision-making process. For dynamic tasks, controlled

changes or failures can be introduced after the initial plan is generated. The ability of the agent to detect the problem, revise its strategy, and complete the task can then be evaluated.

5.4 Task Complexity Analysis

The performance of the proposed Agentic AI framework can be examined across different levels of task complexity. Simple tasks may require only a few reasoning and execution steps, while complex tasks may involve multiple subtasks, dependencies, tool interactions, and decision points. Evaluating the system at different complexity levels helps determine how effectively the agent scales when the number of required actions increases. The analysis can consider task completion, decision quality, execution steps, and error propagation.

5.5 Adaptability and Dynamic Environment Testing

Agentic AI systems should be capable of responding to changing conditions. Therefore, the evaluation can introduce controlled changes during task execution, such as modified requirements, unavailable tools, unexpected outputs, or incomplete information. The agent's ability to detect these changes, update its plan, and continue execution can be measured. This experiment helps evaluate the practical contribution of feedback and adaptive re-planning.

5.6 Robustness and Error Analysis

Robustness evaluation examines how the framework behaves when errors occur during reasoning, planning, tool execution, or information retrieval. Different failure conditions can be introduced and the system's response can be analyzed. Important observations include error detection, recovery success, repeated failures, and impact on the final decision. This analysis can identify weaknesses in individual components and provide directions for improving system reliability.

5.7 Scalability and Resource Efficiency

The proposed framework can also be evaluated based on its resource requirements as task complexity increases. Important factors include execution time, number of model interactions,

number of tool calls, memory usage, and overall computational cost. The results can be compared across different task sizes to determine whether the framework remains practical for larger workloads. Improving resource efficiency is important for deploying Agentic AI systems in real-world applications.

5.8 Reproducibility and Consistency Analysis

The proposed framework can be tested through repeated trials using the same task conditions to examine the consistency of its decisions and execution behavior. Since LLM-based agents may produce different reasoning paths for similar inputs, repeated experiments can help identify variations in planning, tool selection, and final outcomes. Important observations include consistency of task completion, decision quality, and recovery behavior. Reproducible evaluation procedures can improve the reliability and credibility of the experimental findings.

5.9 Ablation Analysis

Ablation analysis can be used to determine the contribution of individual components within the proposed framework. Selected components such as feedback, verification, memory, tool interaction, or adaptive re-planning can be temporarily removed while keeping the remaining system unchanged. The resulting performance can then be compared with the complete framework. This analysis helps identify which components contribute most significantly to autonomous decision-making and overall system performance.

5.10 Decision Quality Assessment

Decision quality can be evaluated by comparing the agent's final decisions with predefined task requirements or expected outcomes. The assessment can consider correctness, relevance, completeness, and consistency of the decisions. For multi-step tasks, intermediate decisions can also be examined to identify where errors occur. This evaluation helps determine whether the proposed framework produces useful decisions rather than simply completing tasks through a sequence of actions.

5.11 Response Time and Execution Efficiency

The efficiency of the proposed framework can be analyzed by measuring the time required to complete different tasks. Response time may be influenced by the number of reasoning cycles, planning steps, tool calls, verification operations, and re-planning attempts. Comparing execution time across different approaches can help identify the computational overhead introduced by autonomous decision-making. The objective is to maintain an appropriate balance between decision quality and execution efficiency.

5.12 Human Intervention Analysis

The evaluation can measure how often human intervention is required during autonomous task execution. Tasks can be categorized according to whether they are completed independently, require clarification, or require human approval because of uncertainty or risk. A lower requirement for unnecessary intervention can indicate better autonomy, while appropriate intervention in high-risk situations can demonstrate effective safety control. This analysis helps evaluate the balance between autonomous operation and human oversight.

The results of the energy consumption analysis show that there is negligible overhead on the power consumption of the gateway when using continuous authentication. The asynchronous operation paradigm guarantees that the process of authentication does not affect the communication tasks.

VI. RESULTS AND DISCUSSION

The evaluation of the proposed framework focuses on both final task outcomes and the behavior of the agent during execution. Since Agentic AI performs multiple reasoning and action steps, evaluating only the final answer may not fully represent system performance.

The first aspect is overall task performance. Task completion rate and decision accuracy can be compared between the proposed framework and the baseline approaches. The measured results

should be used to determine whether the additional planning and feedback mechanisms provide measurable improvements. The second aspect is planning and decision-making. Explicit task decomposition may help the agent organize complex objectives into smaller operations. The evaluation should examine whether the generated plans reduce unnecessary actions and produce more consistent decisions. Another important factor is failure recovery.

When an intermediate action produces an unexpected result, the feedback mechanism allows the agent to reconsider the affected step. The number of detected failures and successfully recovered tasks can be recorded to measure the effectiveness of adaptive re-planning. Tool interaction should also be evaluated by examining tool selection, execution success, and interpretation of returned information. Incorrect tool selection or incorrect interpretation can negatively affect later decisions.

The verification mechanism provides another layer of reliability by checking whether the final outcome satisfies the original objective. Comparing systems with and without verification can help determine the contribution of this stage. Finally, computational efficiency must be considered. Multiple reasoning, planning, verification, and re-planning cycles can increase execution time and model usage. Therefore, the practical effectiveness of the proposed approach should be considered as a balance between decision quality, adaptability, and computational requirements. Important: Actual numerical values, graphs, accuracy percentages, or performance claims should be added only after the experiment is actually conducted. Fabricated experimental results should not be included in a research paper.

VII. CHALLENGES AND LIMITATIONS

Despite its potential, Agentic AI introduces several challenges. Decision reliability is a major concern because an incorrect interpretation of the objective or an unsuitable plan can affect all subsequent actions. Errors introduced during early stages may propagate through the complete execution process.

Hallucination and incorrect reasoning are additional concerns for LLM based agents. An agent may generate information that appears reasonable but is factually incorrect. External verification and reliable information sources can reduce this problem but cannot completely eliminate it. Long-horizon task execution also presents difficulties. As the number of subtasks increases, maintaining consistency between earlier and later decisions becomes more complex.

Memory limitations, accumulated errors, and changing environmental conditions may affect performance. External tool dependencies and security risks must also be considered. Tools may become unavailable, return incorrect information, receive inappropriate inputs, or expose the system to security vulnerabilities. Another limitation is computational cost. Autonomous agents may require multiple model interactions, memory retrieval operations, tool calls, and verification cycles. This can increase latency and resource consumption compared with direct response-based systems. Finally, explainability and human oversight remain important. Autonomous systems operating in high-impact environments may require human approval before executing critical decisions. Therefore, future implementations should provide suitable monitoring and intervention mechanisms.

VIII. FUTURE SCOPE

Agentic AI has significant potential for further development as intelligent systems become more autonomous and capable of interacting with complex environments. Future research can focus on improving the planning and reasoning abilities of AI agents so that they can handle longer and more complicated tasks with fewer errors. Advanced planning mechanisms can help agents generate efficient action sequences and revise them dynamically when the environment changes. Another important direction is the development of long-term memory. Current agents may have difficulty maintaining relevant information across extended tasks. Improved memory mechanisms can allow agents to retain previous experiences, user preferences, task history, and successful strategies,

enabling more consistent decision-making over time. Multi-agent collaboration is another promising area.

Multiple specialized agents can be assigned different responsibilities, such as planning, information retrieval, verification, and execution. Cooperation between these agents could improve the ability of AI systems to solve complex problems while reducing the workload of a single agent. Future systems can also focus on explainable and trustworthy decision-making. Agents should be able to provide understandable reasons for important decisions, identify uncertainty, and communicate the factors that influenced their actions.

This is particularly important for applications involving healthcare, finance, cybersecurity, and other high-impact domains. The integration of Agentic AI with real-world environments provides further opportunities. Autonomous agents can be connected with robots, IoT devices, transportation systems, enterprise applications, and smart infrastructure. Such integration could allow agents to sense environmental conditions, make decisions, and perform physical or digital actions with limited human intervention.

Security and safety will also require significant attention. Future research should investigate protection against malicious inputs, unsafe tool execution, unauthorized actions, data manipulation, and adversarial attacks. Appropriate access controls, monitoring mechanisms, verification layers, and human approval systems can improve the safe deployment of autonomous agents. Another major research direction is continuous learning and adaptation. Instead of relying only on their initial knowledge, future agents can learn from execution outcomes and feedback to improve their future decisions. This could enable systems to adapt to changing environments without requiring complete redesign.

Finally, future work should address computational efficiency and scalability. Reducing unnecessary model calls, optimizing tool usage, and improving

memory and planning mechanisms can make Agentic AI more practical for large-scale applications. The long-term goal is to develop autonomous systems that are not only capable of completing complex tasks but are also reliable, secure, explainable, efficient, and responsible. Detection performance of the framework for the different attack vectors is shown in Figure 5. The framework effectively detects replay attacks, device substitution attacks, and credential compromise scenarios. Continuous authentication approach of the framework is especially useful in detecting attacks performed post-authentication.

IX. CONCLUSION

Agentic AI provides a promising approach for developing intelligent systems capable of autonomous planning, reasoning, action, and adaptation. This research presented a closed-loop framework that combines goal analysis, task decomposition, planning, reasoning, tool interaction, execution, feedback, verification, and adaptive re-planning. The proposed framework moves autonomous decision-making beyond conventional response-based AI by allowing the system to continuously evaluate its actions and modify its strategy when necessary. This approach provides a structured foundation for developing intelligent systems capable of handling complex and changing tasks. However, challenges related to reliability, hallucination, security, explainability, computational cost, and human oversight must be addressed before deploying highly autonomous systems in critical environments. Future research can therefore focus on developing Agentic AI systems that are more reliable, secure, explainable, efficient, and adaptable.

REFERENCES

1. L. Wang, C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhang, Z. Chen, J. Tang, X. He, and others, "A Survey on Large Language Model Based Autonomous Agents," *Frontiers of Computer Science*, vol. 18, 2024.

2. M. Abou Ali et al., "Agentic AI: A Comprehensive Survey of Architectures, Applications, and Future Directions," *Artificial Intelligence Review*, vol. 59, 2026.
3. X. Huang et al., "Understanding the Planning of LLM Agents: A Survey," *arXiv preprint arXiv:2402.02716*, 2024.
4. W. Xu et al., "LLM-Based Agents for Tool Learning: A Survey," *Data Science and Engineering*, vol. 10, pp. 533–563, 2025.
5. X. Li et al., "A Survey on LLM-Based Multi-Agent Systems: Workflow, Infrastructure, and Challenges," *Vicinagearth*, vol. 1, Art. no. 9, 2024.
6. T. Guo et al., "Large Language Model Based Multi-Agents: A Survey of Progress and Challenges," in *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 8048–8057, 2024.
7. X. Li, "A Review of Prominent Paradigms for LLM-Based Agents: Tool Use, Planning, and Feedback Learning," *arXiv preprint arXiv:2406.05804*, 2024.
8. T. Masterman et al., "The Landscape of Emerging AI Agent Architectures for Reasoning, Planning, and Tool Calling: A Survey," *arXiv preprint arXiv:2404.11584*, 2024.
9. A.Yehudai et al., "Survey on Evaluation of LLM-Based Agents," *arXiv preprint arXiv:2503.16416*, 2025.
10. S. Yao et al., "ReAct: Synergizing Reasoning and Acting in Language Models," in *International Conference on Learning Representations (ICLR)*, 2023.
11. This is especially relevant to your Reasoning → Action → Observation → Decision concept.