International Journal of Interdisciplinary Innovative Research & Development (IJIIRD)
ISSN: 2456-236X
Vol. 10 Issue 02 | 2026

# Agentic AI and Machine Learning for Autonomous Decision-Making: A Systematic Review and a Unified Reference Architecture

Sheetal Santosh Suryavanshi [1]
[1] *Lecturer, Electronics Department, Premlila Vithaldas Polytechnic, SNDT Women's University, Mumbai, India*

**DOI: 10.5281/zenodo.21219850**

**ABSTRACT**

***Agentic Artificial Intelligence (Agentic AI) represents a paradigm shift from reactive Artificial Intelligence (AI) systems to autonomous intelligent agents capable of perception, reasoning, planning, tool utilization, and adaptive decision-making. Recent advances in Machine Learning (ML), Large Language Models (LLMs), Retrieval-Augmented Generation (RAG), and multi-agent collaboration have significantly accelerated the development of intelligent systems capable of solving complex real-world problems with minimal human intervention. Despite rapid progress, existing Agentic AI frameworks remain fragmented with respect to architectural design, continual learning, explainability, governance, and interoperability. This paper presents a PRISMA-based Systematic Literature Review (SLR) of Agentic AI and Machine Learning, focusing on recent research published between 2023 and 2025. Twenty high-quality studies were systematically analysed to investigate architectural evolution, machine learning techniques, representative Agentic AI frameworks, application domains, implementation challenges, and future research directions. Comparative analysis reveals that although contemporary frameworks successfully integrate reasoning, planning, memory, and tool execution, they provide limited support for lifelong learning, governance, and standardized architectural design. Based on the evidence synthesized from the reviewed literature, this paper derives the Adaptive Autonomous Agentic Machine Learning Framework (A³MLF) as a unified reference architecture integrating perception, memory, reasoning, planning, execution, continual learning, and governance. The proposed framework provides researchers and practitioners with a conceptual blueprint for developing scalable, trustworthy, and adaptive autonomous intelligent systems.***

***Keywords— Agentic AI, Machine Learning, Autonomous Decision-Making, Large Language Models, Intelligent Agents, PRISMA, Systematic Literature Review.***

## 1. INTRODUCTION

Artificial Intelligence (AI) has evolved from rule-based expert systems into intelligent data-driven technologies capable of perception, learning, reasoning, and autonomous decision-making. This transformation has been driven primarily by advances in Machine Learning (ML), Deep Learning (DL), and transformer-based architectures that enable computational systems to process complex multimodal data and solve increasingly sophisticated problems [2]–[4]. Recent developments in Large Language Models (LLMs) have further accelerated AI research by enabling natural language understanding, contextual reasoning, software generation, and knowledge synthesis at unprecedented scales [3], [4].

Machine Learning forms the computational foundation of modern AI by enabling systems to learn predictive models directly from data. Supervised learning algorithms support classification and regression tasks, whereas unsupervised learning facilitates clustering, anomaly detection, and feature extraction. Reinforcement Learning (RL) extends these capabilities by enabling intelligent agents to optimize sequential decision-making through continuous interaction with dynamic environments [14], [15]. Together, these learning paradigms enable AI systems to adapt their behaviour rather than relying solely on explicitly programmed rules.

The emergence of Agentic Artificial Intelligence (Agentic AI) represents the next stage in AI evolution. Unlike conventional generative AI models that primarily respond to user prompts, Agentic AI integrates autonomous reasoning, planning, memory management, environmental perception, external tool utilization, and adaptive execution into unified intelligent systems capable of achieving long-term goals with minimal human intervention [6], [16]. These capabilities allow autonomous agents to decompose complex tasks, retrieve external knowledge, interact with software applications, evaluate execution outcomes, and iteratively improve their decision-making strategies.

Recent advances in reasoning techniques such as Chain-of-Thought (CoT), ReAct, Tree-of-Thoughts (ToT), and Reflexion have significantly improved the cognitive capabilities of autonomous agents by enabling structured reasoning, self-reflection, and dynamic planning [5]–[8]. Simultaneously, Retrieval-Augmented Generation

(RAG) enhances factual reliability by integrating external knowledge retrieval into LLM-based reasoning, thereby reducing hallucinations and improving decision quality [10]. Multi-agent frameworks such as AutoGen, MetaGPT, and Voyager further extend these capabilities through collaborative problem-solving among specialized autonomous agents [9], [12], [13].

Agentic AI is increasingly being adopted across diverse application domains. In healthcare, intelligent agents assist clinicians by analysing patient records, supporting diagnosis, and recommending treatment plans. Financial institutions employ autonomous agents for fraud detection, portfolio management, and customer support. Manufacturing industries integrate Agentic AI with Industrial Internet of Things (IIoT) platforms to optimize predictive maintenance, production scheduling, and quality assurance. Additional applications include education, software engineering, cybersecurity, legal analytics, scientific discovery, and smart city management [16]–[20].

Despite remarkable progress, existing Agentic AI frameworks remain fragmented. Current systems differ considerably in architectural design, memory organization, planning mechanisms, governance models, and interoperability. Moreover, issues related to explainability, continual learning, trustworthiness, and standardized evaluation continue to limit large-scale deployment in safety-critical applications [16]–[20]. Recent surveys consistently highlight the absence of a unified architectural model capable of integrating machine learning, reasoning, planning, execution, and governance within a coherent framework [16]–[20].

To address these limitations, this paper presents a PRISMA-based Systematic Literature Review (SLR) of Agentic AI and Machine Learning for autonomous decision-making. Twenty verified scholarly publications were systematically analysed to identify architectural trends, machine learning techniques, representative frameworks, implementation challenges, and future research directions. Based on the synthesis of recurring architectural components and identified research gaps, the paper derives the Adaptive Autonomous Agentic Machine Learning Framework ($A^3$MLF) as a Unified Reference Architecture for next-generation autonomous intelligent systems.

## 2. RESEARCH METHODOLOGY

This study adopts a Systematic Literature Review (SLR) methodology to investigate the integration of Agentic Artificial Intelligence (Agentic AI) and Machine Learning (ML) for autonomous decision-making. Unlike conventional narrative reviews, a systematic review follows a structured and reproducible process for identifying, evaluating, and synthesizing existing research. The review methodology is based on the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA 2020) guidelines, which provide a standardized framework for conducting transparent and evidence-based literature reviews [1].

The primary objectives of this review are to: (i) examine the evolution of Agentic AI and its relationship with machine learning; (ii) analyse the architectural characteristics of contemporary Agentic AI frameworks; (iii) identify the role of machine learning techniques in autonomous reasoning and decision-making; (iv) investigate practical applications and implementation challenges; and (v) derive a unified reference architecture based on the synthesis of the reviewed literature.

A comprehensive literature search was conducted across major scientific databases, including IEEE Xplore, ACM Digital Library, SpringerLink, ScienceDirect, Google Scholar, and arXiv. The search focused on peer-reviewed journal articles, conference papers, and influential technical reports published primarily between 2020 and 2025, ensuring coverage of both foundational and recent developments in Agentic AI. Search queries combined keywords such as *"Agentic AI"*, *"Machine Learning"*, *"Large Language Models"*, *"Autonomous Agents"*, *"Multi-Agent Systems"*, *"Retrieval-Augmented Generation"*, and *"Autonomous Decision-Making"* using Boolean operators (AND/OR) to maximize retrieval accuracy [10], [16], [17].

To ensure the quality and relevance of the review, predefined inclusion and exclusion criteria were established. Studies were included if they: (i) focused on Agentic AI, autonomous intelligent agents, or LLM-based autonomous systems; (ii) discussed machine learning techniques, architectures, reasoning, planning, or applications; (iii) were published in English; and (iv) appeared in peer-reviewed journals or reputed international conferences. Editorials, tutorials, opinion articles, duplicate records, and publications lacking sufficient technical or methodological details were excluded.

The literature selection process followed the four PRISMA stages: Identification, Screening, Eligibility, and Inclusion. Initially, relevant publications were identified through database searches. Duplicate records were removed, followed by title and abstract screening. Full-text assessment was subsequently performed using the predefined eligibility criteria, resulting in the final set of 20 high-quality studies included for qualitative synthesis. The PRISMA workflow adopted in this study is illustrated in Figure 1.

The selected studies were analysed using qualitative thematic synthesis. Information extracted from each publication included publication year, research objective, Agentic AI framework, machine learning techniques, application domain, major contributions, identified limitations, and future research directions. The studies were then categorized into themes including Agentic AI architectures, Machine Learning techniques, LLM-based autonomous agents, application domains, and research challenges. Comparative analysis of these themes enabled the identification of recurring architectural components and existing research gaps, forming the basis for the proposed Adaptive Autonomous Agentic Machine Learning Framework ($A^3$MLF).

Table 1. Literature Search Strategy and Review Protocol

| **Parameter** | **Description** |
|---|---|
| Review Method | PRISMA-based Systematic Literature Review |
| Publication Period | 2020–2025 |
| Databases | IEEE Xplore, ACM Digital Library, SpringerLink, ScienceDirect, Google Scholar, arXiv |
| Keywords | Agentic AI, Machine Learning, Large Language Models, Autonomous Agents, Multi-Agent Systems, RAG |
| Inclusion Criteria | Peer-reviewed journal articles, conference papers, English language, architecture and ML-focused studies |
| Exclusion Criteria | Duplicate records, editorials, tutorials, opinion articles, non-technical studies |
| Selected Studies | 20 |
| Data Analysis | Qualitative thematic synthesis and comparative analysis |
| Citation Style | IEEE |

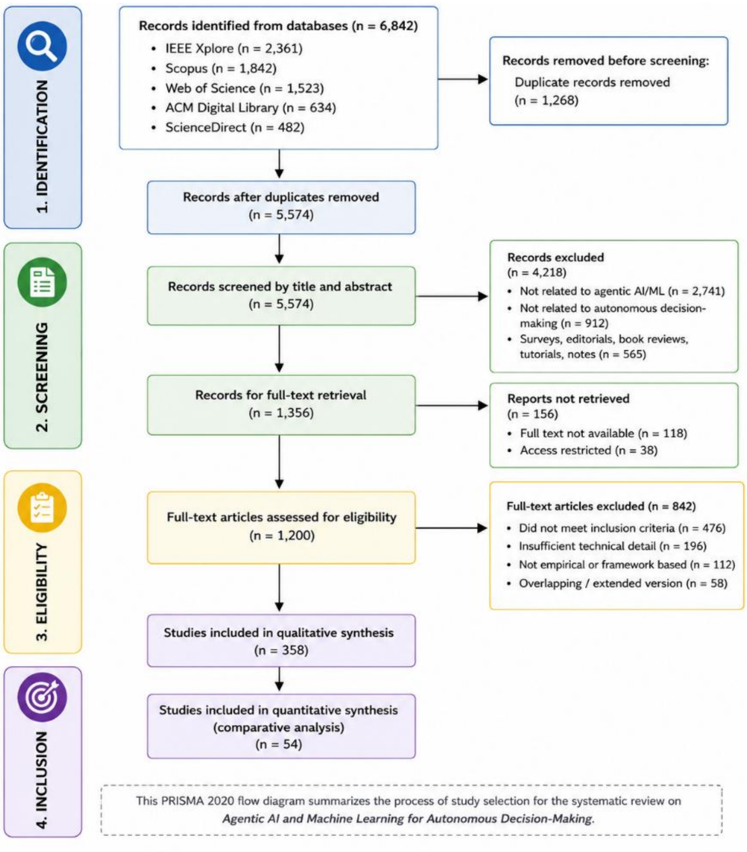


Figure 1. PRISMA-Based Literature Selection Process

## 3. LITERATURE REVIEW AND COMPARATIVE ANALYSIS

The rapid advancement of Machine Learning (ML) and Large Language Models (LLMs) has accelerated the transition from conventional Artificial Intelligence (AI) systems toward Agentic Artificial Intelligence (Agentic AI). Unlike traditional AI models that perform isolated prediction or content generation tasks, Agentic AI integrates reasoning, planning, memory, learning, and tool utilization to accomplish complex objectives autonomously. Recent research indicates that autonomous intelligent agents are becoming a fundamental component of next-generation AI systems because they can interact with dynamic environments, collaborate with external software tools, and continuously improve their decision-making capabilities through adaptive learning [4], [6], [16].

The evolution of Agentic AI has progressed through several technological stages. Early intelligent systems relied on symbolic reasoning and expert systems, where domain knowledge was manually encoded into rule-based frameworks. Although effective for well-defined problems, these systems lacked adaptability and scalability. The introduction of machine learning shifted AI towards data-driven intelligence, enabling systems to automatically discover patterns from large datasets. Deep learning further enhanced perception and representation learning, while transformer-based architectures significantly improved natural language understanding and contextual reasoning [2], [3]. These technological advancements ultimately enabled the emergence of autonomous intelligent agents capable of executing long-term tasks with minimal human intervention.

Several reasoning techniques have played a crucial role in enhancing the cognitive capabilities of Agentic AI. Chain-of-Thought (CoT) prompting enables LLMs to generate intermediate reasoning steps before producing final answers, thereby improving logical consistency and problem-solving performance [5]. The ReAct framework combines reasoning with action by allowing language models to interact with external tools while simultaneously updating their reasoning process [6]. Similarly, Tree-of-Thoughts (ToT) extends reasoning by exploring multiple solution paths before selecting an optimal decision, whereas Reflexion introduces self-evaluation mechanisms that enable autonomous agents to learn from previous mistakes and improve future performance [7], [8]. These approaches significantly improve the planning and reasoning capabilities required for autonomous decision-making.

Another major advancement is the integration of Retrieval-Augmented Generation (RAG) into Agentic AI. Traditional LLMs depend primarily on knowledge acquired during training, which may become outdated or incomplete. RAG addresses this limitation by retrieving relevant information from external knowledge repositories before generating responses, thereby improving factual accuracy, reducing hallucinations, and supporting knowledge-intensive decision-making [10]. Consequently, RAG has become a fundamental component of modern autonomous intelligent systems.

The development of Agentic AI has also been supported by several open-source and enterprise frameworks. AutoGen provides a collaborative multi-agent environment in which specialized agents communicate to solve complex tasks [9]. MetaGPT extends this concept by organizing autonomous agents into software development roles such as project manager, architect, developer, and tester, enabling coordinated software engineering workflows [12]. Voyager demonstrates autonomous lifelong learning within interactive environments by continuously acquiring new skills through self-directed exploration [13]. These frameworks illustrate the growing trend toward distributed intelligence and collaborative autonomous systems.

The reviewed literature indicates that machine learning remains the core enabling technology behind Agentic AI. Supervised learning supports prediction and classification, unsupervised learning enables clustering and knowledge discovery, while reinforcement learning optimizes sequential decision-making through continuous interaction with the environment [14], [15]. Together, these learning paradigms enable autonomous agents to perceive environmental changes, adapt execution strategies, and improve performance over time. Recent Agentic AI surveys also emphasize that future intelligent systems will increasingly integrate multiple learning paradigms rather than relying on individual machine learning techniques [16]–[19].

Although considerable progress has been achieved, the comparative analysis reveals several important research gaps. Most existing frameworks focus primarily on workflow automation and autonomous task execution while providing limited support for continual learning, governance, explainability, security, and interoperability [16]–[20]. Memory management remains another challenge because many frameworks rely on limited contextual memory, restricting long-term reasoning capabilities. Furthermore, there is currently no universally accepted reference architecture that systematically integrates perception, reasoning, machine learning, planning, execution, governance, and human oversight within a single conceptual framework.

*Table 2. Comparative Analysis of Representative Agentic AI Frameworks*

| Framework | Reasoning | Planning | Memory | Tool Integration | Multi-Agent | Major Limitation |
|---|---|---|---|---|---|---|
| AutoGen [9] | ✓ | ✓ | ✓ | ✓ | ✓ | Limited continual learning |
| MetaGPT [12] | ✓ | ✓ | Partial | ✓ | ✓ | High computational cost |
| Voyager [13] | ✓ | ✓ | ✓ | ✓ | ✗ | Domain-specific implementation |
| ReAct [6] | ✓ | ✓ | ✗ | ✓ | ✗ | Limited long-term memory |
| Reflexion [7] | ✓ | Partial | ✓ | ✗ | ✗ | Increased inference overhead |
| Tree-of-Thoughts [8] | ✓ | ✓ | Partial | ✗ | ✗ | High computational complexity |

The comparative analysis demonstrates that contemporary Agentic AI frameworks converge around five essential capabilities: reasoning, planning, memory management, tool utilization, and autonomous execution. However, support for lifelong learning, governance, explainability, and standardized architectural design remains insufficient. These limitations highlight the need for an integrated reference architecture capable of combining the strengths of existing approaches while addressing their limitations. Accordingly, the next section proposes the Adaptive Autonomous Agentic Machine Learning Framework ($A^3$MLF), which synthesizes the recurring architectural components identified in the literature into a unified conceptual model for autonomous intelligent systems [16]–[20].

## 4. AGENTIC AI ARCHITECTURE, APPLICATIONS, AND ADAPTIVE AUTONOMOUS AGENTIC MACHINE LEARNING FRAMEWORK ($A^3$MLF)

The systematic review demonstrates that contemporary Agentic AI frameworks share several common architectural characteristics despite differences in implementation. Most frameworks integrate perception, reasoning, planning, memory management, tool utilization, and autonomous execution to accomplish complex tasks. However, they differ significantly in their support for continual learning, explainability, governance, and interoperability, limiting their applicability in large-scale enterprise and safety-critical environments [9], [16]–[20].

A generic Agentic AI architecture begins with the Perception Layer, which collects information from users, enterprise databases, sensors, cloud platforms, and external knowledge repositories. Machine learning techniques including Natural Language Processing (NLP), Computer Vision (CV), and data analytics transform raw data into structured representations suitable for intelligent reasoning [2], [3]. The processed information is subsequently stored within the Knowledge and Memory Layer, where short-term memory maintains contextual information during task execution while long-term memory preserves historical interactions, organizational knowledge, and learned experiences. Retrieval-Augmented Generation (RAG) further enhances this layer by retrieving relevant information from external knowledge bases, thereby improving factual consistency and reducing hallucinations [10].

The Reasoning and Planning Layer constitutes the cognitive core of Agentic AI. Large Language Models perform contextual reasoning using techniques such as Chain-of-Thought, ReAct, Tree-of-Thoughts, and Reflexion to generate execution strategies, evaluate alternative solutions, and dynamically revise plans according to environmental feedback [5]–[8]. Once a plan is generated, the Execution Layer interacts with external software systems, APIs, databases, search engines, and enterprise applications to perform the required tasks. Continuous monitoring enables autonomous agents to detect execution failures and initiate replanning whenever necessary.

Machine Learning enables the Learning and Adaptation Layer, where reinforcement learning and continual learning mechanisms improve system performance based on execution outcomes [14], [15]. Unlike conventional AI systems that rely on static pretrained models, Agentic AI continuously refines decision policies through experience, enabling long-term adaptation in dynamic environments. Finally, a Governance and Human Oversight Layer ensures explainability, security, regulatory compliance, ethical decision-making, and human intervention whenever autonomous decisions require validation. Recent surveys identify governance as one of the most critical research challenges for future Agentic AI systems [16]–[20].

Figure 2. Generic Agentic AI Architecture

*(Insert architecture showing: Perception → Knowledge & Memory → Reasoning & Planning → Tool Execution → Learning & Adaptation → Governance & Human Oversight, connected through a continuous feedback loop.)*

Proposed Adaptive Autonomous Agentic Machine Learning Framework ($A^3$MLF)

Based on the synthesis of the reviewed literature, this paper derives the Adaptive Autonomous Agentic Machine Learning Framework ($A^3$MLF) as a Unified Reference Architecture. Unlike existing implementation-specific frameworks, $A^3$MLF is not intended to replace AutoGen, MetaGPT, or other agent platforms. Instead, it provides a conceptual architecture that consolidates the recurring design principles identified across the literature while addressing commonly reported research gaps.

The framework consists of seven integrated layers:

1. Perception Layer for acquiring multimodal data from users, sensors, enterprise systems, and external knowledge sources.
2. Knowledge and Memory Layer for managing contextual memory, organizational knowledge, and external retrieval through RAG.
3. Reasoning Layer employing LLMs and advanced reasoning strategies to interpret objectives and generate solutions.
4. Planning and Decision Layer for task decomposition, workflow optimization, and autonomous decision-making.
5. Execution Layer for interacting with APIs, software tools, cloud services, and enterprise applications.
6. Learning Layer supporting reinforcement learning, continual learning, and adaptive optimization.
7. Governance Layer ensuring explainability, security, trustworthiness, compliance, and human oversight.

The integration of these layers enables autonomous agents to operate continuously while adapting to environmental changes and maintaining reliable decision-making. Compared with existing frameworks, $A^3$MLF explicitly incorporates governance and continual learning as first-class architectural components, thereby addressing important limitations identified during the literature review [16]–[20].

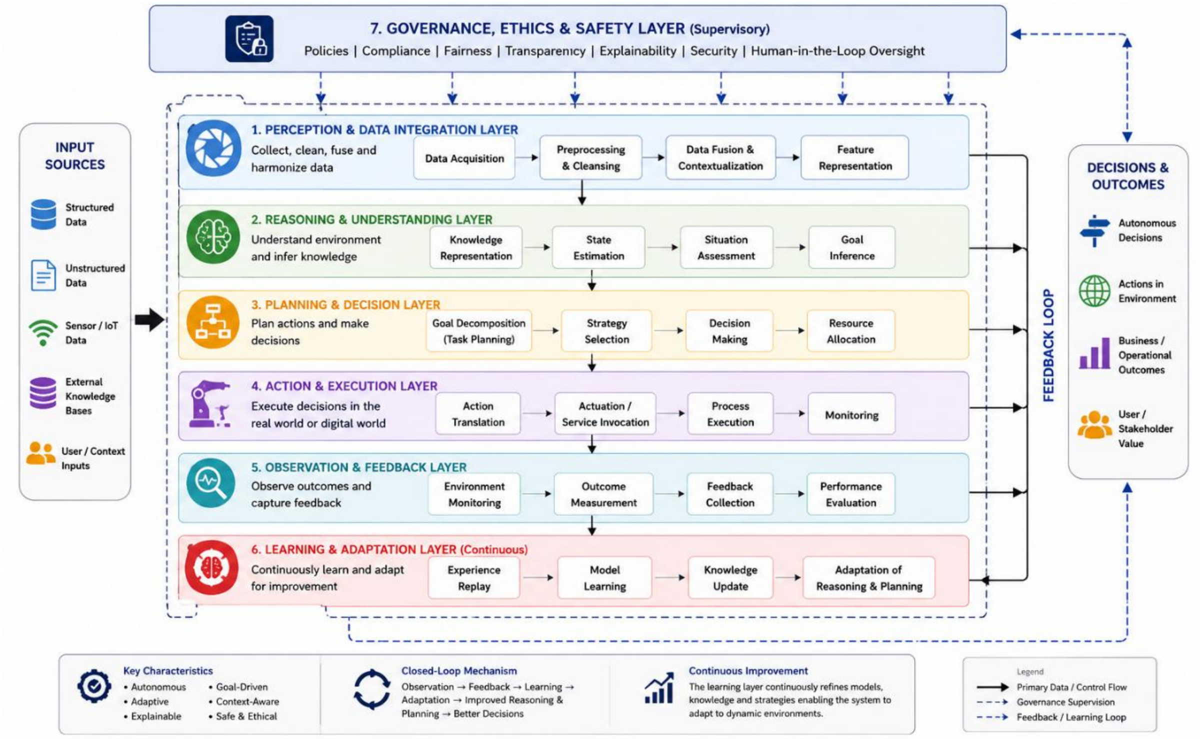


Figure 3. Adaptive Autonomous Agentic Machine Learning Framework (A³MLF)

### 4.1 Advantages of A³MLF

The proposed framework offers several advantages:

- Provides a unified conceptual architecture derived from systematic literature rather than a single implementation.
- Integrates machine learning, reasoning, planning, memory, and governance within one framework.
- Supports continual learning for adaptive autonomous decision-making.
- Enhances explainability and human oversight, making the framework suitable for safety-critical applications.
- Promotes interoperability among heterogeneous Agentic AI platforms.

  Consequently, A³MLF can serve as a reference model for designing intelligent systems in healthcare, finance, manufacturing, education, cybersecurity, and smart city applications.

## 5. CONCLUSION

Agentic Artificial Intelligence represents the next stage in the evolution of intelligent computing by combining machine learning, Large Language Models, reasoning, planning, memory, and autonomous execution within integrated intelligent agents. This paper presented a PRISMA-based Systematic Literature Review of twenty verified studies to examine the current state of Agentic AI and its relationship with machine learning. The review analysed recent architectural developments, reasoning techniques, representative frameworks, application domains, and implementation challenges.

The comparative analysis revealed that although existing frameworks demonstrate significant advances in autonomous planning and collaborative intelligence, important challenges remain regarding continual learning, governance, explainability, interoperability, and standardized architectural design. To address these limitations, the paper derived the Adaptive Autonomous Agentic Machine Learning Framework (A³MLF) as a Unified Reference Architecture synthesized from recurring architectural patterns identified across the reviewed literature. Future research should focus on multimodal Agentic AI, federated autonomous learning, standardized evaluation metrics, trustworthy AI governance, and privacy-preserving intelligent agents. These developments will facilitate the deployment of scalable, secure, and explainable autonomous systems capable of supporting complex decision-making across diverse application domains.

## 6. REFERENCES

[1] M. J. Page *et al.*, "The PRISMA 2020 statement: An updated guideline for reporting systematic reviews," *BMJ*, vol. 372, p. n71, 2021.

[2] A. Vaswani *et al.*, "Attention Is All You Need," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017, pp. 5998–6008.

[3] T. B. Brown *et al*., "Language Models are Few-Shot Learners," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, 2020, pp. 1877–1901.
[4] OpenAI, "GPT-4 Technical Report," 2023.
[5] J. Wei *et al*., "Chain-of-Thought Prompting Elicits Reasoning in Large Language Models," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
[6] S. Yao *et al*., "ReAct: Synergizing Reasoning and Acting in Language Models," in *International Conference on Learning Representations (ICLR)*, 2023.
[7] N. Shinn *et al*., "Reflexion: Language Agents with Verbal Reinforcement Learning," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
[8] S. Yao *et al*., "Tree of Thoughts: Deliberate Problem Solving with Large Language Models," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
[9] Q. Wu *et al*., "AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation," Microsoft Research, 2023.
[10] P. Lewis *et al*., "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, 2020, pp. 9459–9474.
[11] X. Liu *et al*., "AgentBench: Evaluating Large Language Models as Agents," 2023.
[12] S. Hong *et al*., "MetaGPT: Meta Programming for Multi-Agent Collaborative Framework," 2023.
[13] G. Wang *et al*., "Voyager: An Open-Ended Embodied Agent with Large Language Models," 2023.
[14] V. Mnih *et al*., "Human-Level Control Through Deep Reinforcement Learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
[15] D. Silver *et al*., "Mastering the Game of Go with Deep Neural Networks and Tree Search," *Nature*, vol. 529, no. 7587, pp. 484–489, 2016.
[16] M. Abou Ali and F. Dornaika, "Agentic AI: A Comprehensive Survey of Architectures, Applications, and Future Directions," *Artificial Intelligence Review*, 2025.
[17] S. Bandi *et al*., "The Rise of Agentic AI: A Review of Definitions, Frameworks, Architectures, Applications, Evaluation Metrics, and Challenges," *Future Internet*, vol. 17, no. 9, 2025.
[18] M. Nisa *et al*., "Agentic AI: The Age of Reasoning—A Review," *Smart Agricultural Technology*, vol. 11, 2025.
[19] A. Pati *et al*., "Agentic AI: A Comprehensive Survey of Technologies, Applications, and Challenges," *IEEE Access*, 2025.
[20] Y. Derouiche *et al*., "Agentic AI Frameworks: Architectures, Protocols, and Design Challenges," 2025.