

A Comprehensive Review of Agentic AI for Intelligent Autonomous Systems

Venkata Phanindra Lingam
AI Architect, USA
venakataphanindra@gmail.com

Sneha Valusa
AI/ML Consultant, USA
snehavalusa@outlook.com

Sankethu Babu Guntaka
AI Security Consultant, USA
sankethuguntaka@outlook.com

Abstract:

Artificial Intelligence (AI) has advanced from traditional rule-based and data-driven machine learning models to intelligent autonomous systems that can reason, plan, learn and perform complicated tasks with little human interaction. This paradigm change has resulted in the advent of Agentic Artificial Intelligence, combining autonomous decision-making, memory, goal-oriented planning, tool use, multi-agent cooperation, and continuous adaptation into a single intelligent framework. This paper provides a detailed discussion of the core principles, architectural elements, enabling technologies, applications, problems and future research directions of Agentic AI. In the first part of the study, the historical history of intelligent agents is traced, starting with symbolic AI, expert systems, multi-agent systems, reinforcement learning and cognitive architectures to the present massive language model driven autonomous agents. And then, the review offers a systematic taxonomy on perception, reasoning and memory management, planning mechanisms, tool integration, collaborative multi-agent coordination, adaptive learning, explainability and governance. Finally, the article discusses future research prospects that are predicted to expedite the development of robust, explainable and trustworthy Agentic AI capable of enabling sophisticated autonomous decision making across heterogeneous intelligent systems.

Keywords— Agentic Artificial Intelligence, Intelligent Agents, Autonomous Systems, Large Language Models, Multi-Agent Systems.

1. Introduction

Artificial Intelligence (AI) is one of the most disruptive technologies of the twenty first century allowing intelligent decision making across varied application domains such as healthcare, finance, transportation, cybersecurity, manufacturing, education and scientific research. Traditional AI systems have made considerable progress in classification, prediction, optimisation and recommendation by using machine learning, deep learning and statistical inference. However, most of the current AI models are still essentially reactive and need a particular cue or a certain procedure to do a job. They often do not have long-term thinking, independent planning, environmental awareness, adaptive decision-making skills or the capacity to successfully complete complicated tasks without ongoing human supervision.

The increasing complexity of real-world situations today has generated a need for intelligent systems that can operate independently in dynamic and

unpredictable contexts. Emerging applications such as autonomous vehicles, intelligent robotics, smart cities, industrial automation, digital healthcare, intelligent software engineering and cybersecurity require systems that can sense continuously changes in the environment, understand contextual information, develop long-term plans, take appropriate actions, assess the outcome and improve future decisions through continuous learning. These demands have driven the movement from traditional AI models to Agentic Artificial Intelligence (Agentic AI), the next evolutionary phase of intelligent autonomous systems.

Modern agentic AI architectures typically integrate a combination of complementary components, including perception modules, memory systems, reasoning engines, planning mechanisms, tool execution frameworks, retrieval-augmented knowledge systems, reinforcement learning modules, explainability components and multi-agent coordination mechanisms.

Received: 08-06-2026

Revised: 16-08-2026

Accepted: 29-08-2026

Published: 05-09-2026

Citation: “ A Comprehensive Review of Agentic AI for Intelligent Autonomous Systems”, *ijaicn*, vol. 2, no. 3, pp. 135–140, September. 2026,

Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

These components allow intelligent entities to understand context, reason logically, use external computing resources, learn from past experiences, and cooperate effectively to achieve difficult goals. These architectures have demonstrated significant potential in numerous domains, including intelligent diagnosis in healthcare, autonomous control in robotics, cybersecurity operations, financial decision support, scientific discovery, industrial process optimisation, smart manufacturing, and autonomous transportation systems.

2. Related Works

Russell and Norvig (1995; later editions) extended rational agents to the agent-oriented problem solving AI framework. They were concerned with intelligent perception, decision making, interaction with the environment and utility based optimisation and thereby became an important reference for teaching and research in the field of AI. Many modern autonomous AI systems adopt the rational agent paradigm.

In the late 1990s to early 2000s, research was directed on cognitive architectures that emulate human cognition and decision making. Rao and Georgeff's BDI (Belief-Desire-Intention) architecture allows agents to form goal-based intentions, have internal beliefs about the environment and carry out plans autonomously. The BDI notion has influenced significantly the modern task-oriented intelligent agents, cognitive robots and autonomous planning systems.

The multi-agent systems were improved by the distributed intelligent agent cooperative problem solving (Jennings, 2000). He demonstrated that autonomous agents may solve complex problems quicker by communicating, collaborating, bargaining and distributed planning. Multi-agent cooperation is a key need in today's Agentic AI frameworks for complex business processes and collaborative autonomous systems.

Research on reinforcement learning advanced autonomous decision-making. According to a complete paradigm of Sutton and Barto (1998), intelligent agents learn optimal behaviours by interacting continuously with dynamic environments via incentive based feedback. Reinforcement learning, that permitted unconscious adaptability, became a popular learning approach in robotics, autonomous vehicles, industrial automation, and intelligent decision-support systems.

Planning algorithms pushed forward autonomous

intelligence. Automatic planning algorithms that can construct sequences of actions to meet intricate goals under uncertain scenarios have been developed by Ghallab, Nau, and Traverso (2004). They enabled intelligent agents to do long-term reasoning, dynamic task scheduling and adaptive execution; all of which are relevant to current Agentic AI systems.

Intelligent decision making in the presence of uncertainty facilitated by probabilistic reasoning and knowledge representation. Bayesian networks, Markov decision processes, ontology-based knowledge representation and semantic reasoning employing incomplete observations and prior information were used to assist intelligent agents in making better choices. These solutions considerably increase resilience in dynamic real-world environments.

The following decade was marked by the revolution of intelligent agent research by machine learning and deep learning. Autonomous systems increasingly learned representations from large datasets rather than hand-built symbolic knowledge. CNNs, RNNs, LSTM networks, and subsequent Transformer designs enabled intelligent agents to understand complicated visual, verbal, and temporal information with unprecedented accuracy.

Vaswani et al. (2017) employed attention techniques and the transformer architecture to replace recurrent computation and to allow highly parallel contextual learning in sequence modelling. Transformers paved the way for today's Large Language Models (LLMs) with improvements to reasoning, natural language understanding and context awareness.

AlphaGo and AlphaZero used reinforcement learning and deep neural networks to plan and think strategically on their own in complex games, surpassing human capabilities. These systems revealed that autonomous long-term planning and adaptive learning are achievable without domain-specific programming.

Multi-agent reinforcement learning was a hot topic before 2019. The researchers researched collaborative learning systems that allow numerous autonomous agents to maximise group objectives and adapt to changing conditions. These studies proved the potential for decentralised intelligence and cooperative autonomous decision making, which were shown critical to Agentic AI ecosystems.

These basic experiments established autonomy,

reasoning, planning, learning, and cooperation, but most early intelligent agent systems lacked hand-coded rules, domain-specific knowledge, limited memory, and limited natural language comprehension. They could not generalise across different tasks and so were confined to complex real world scenarios.

Since 2018, the rapid progress of foundation models and large-scale pre-trained language models has revolutionised research on intelligent agents. Modern autonomous agents incorporate natural language thinking, external information retrieval, long-term memory, tool execution, code generation, planning, and adaptive feedback to handle many real-world issues. These traits have led to the development of Agentic AI, integrating autonomous reasoning with generative AI to extend the notion of intelligent agents.

3. Architecture Diagram




Fig 1: Architecture Diagram

4. Results & Discussion

Agentic Artificial Intelligence is the confluence of many AI paradigms, including symbolic reasoning, machine learning, deep learning, reinforcement learning, big language models, and multi-agent systems. In contrast to traditional AI approaches that view prediction or classification as stand-alone tasks, Agentic AI incorporates independent thinking, planning, memory, external tool use, collaborative intelligence, and adaptive learning in a single paradigm. This section compares Agentic AI with past AI paradigms, identifying its merits and weaknesses, and its promise for the future.

Traditional AI systems are usually built on manually created knowledge libraries, symbolic logic and pre-defined rules. These technologies are clear and highly interpretable yet have limited flexibility while functioning in dynamic contexts. Rule-based systems are usually very labor-intensive to update the knowledge bases, and they are not capable of autonomous knowledge acquisition via interaction with the environment.

Data-driven intelligence was enabled by Machine Learning, in which models could learn the prediction patterns directly from the historical information. Supervised learning, unsupervised learning, and statistical optimisation have increased performance in classification, regression, clustering, and anomaly detection. However, traditional machine learning models are often performing isolated prediction tasks without autonomous thinking, long-term planning, contextual memory and independent execution .

Table I: Strengths and Limitations of Agentic AI

Strengths	Limitations
Autonomous reasoning	Hallucination risk
Goal-oriented planning	Computationally expensive
Long-term contextual memory	Memory management complexity
Multi-agent collaboration	Communication overhead
Tool and API integration	Security vulnerabilities
Continuous learning	Ethical and governance concerns
Dynamic task execution	Limited explainability
Adaptive decision making	Dependence on high-quality external knowledge

Table II: Comparative Analysis of AI Paradigms

Feature	Traditional AI	Machine Learning	Deep Learning	Large Language Models	Agentic AI
Rule-Based Reasoning	✓	✗	✗	Partial	✓
Data-Driven Learning	✗	✓	✓	✓	✓
Autonomous Decision Making	Low	Low	Medium	Medium	High
Long-Term Memory	✗	✗	✗	Limited	✓
Goal-Oriented Planning	✗	✗	✗	Limited	✓
External Tool Integration	✗	✗	✗	Partial	✓
Multi-Agent Collaboration	✗	✗	✗	Limited	✓
Continuous Learning	Limited	Medium	High	Medium	High
Explainability	High	Medium	Low	Medium	Medium–High
Scalability	Medium	High	High	High	High
Adaptability	Low	Medium	High	High	Very High
Human Intervention	High	Medium	Medium	Medium	Low

5. Discussions

Large Language Models changed the landscape of AI by offering sophisticated natural language comprehension, contextual reasoning, text production, summarization, translation, and information retrieval. The advent of LLMs greatly improved the human-computer interaction, enabling the user to connect with the machine in the natural language instead of the preset programming interfaces. Despite these developments, the answers provided by standalone LLMs are often in reaction to user inputs, and such LLMs typically lack the ability for independent planning, permanent memory, long-term goal management, external tool orchestration, and continual engagement with the environment.

The comparative study clearly shows that Agentic AI is the next evolutionary step in artificial intelligence, combining the benefits of previous paradigms while resolving many of their flaws. Traditional AI gave us symbolic reasoning, machine learning brought data-driven adaptation, deep learning allowed better vision, and Large Language Models have vastly increased contextual understanding. We combine these skills with autonomous planning, persistent memory, collaborative intelligence, and external tool orchestration to construct systems that can autonomously solve complicated real-world issues. Nevertheless, issues relating to reliability, explainability, security, computing efficiency, and governance still impede broad use. Meeting these

problems is critical for safe, scalable and human-centric autonomous intelligent systems.

6. Limitations

Agentic AI has shown promising promise to enable intelligent autonomous systems, although several restrictions still exist. The performance of Agentic AI relies greatly on the quality of the underlying foundation models, external knowledge sources and the available computing resources. When faced with missing or unclear information, autonomous agents may produce outputs that are either hallucinated or based on false reasoning. Moreover, long-term memory management, multi-agent coordination and real-time planning add system complexity and computational expense. Security risks, privacy problems, explainability and ethical governance are difficulties that may prevent adoption in safety-critical systems. Moreover, the fast development of Agentic AI technologies leads to the fact that the current architectures and assessment methods are still evolving, which makes it challenging to standardise and benchmark across diverse application domains.

7. Conclusion & Future Works

Agentic Artificial Intelligence is the next generation of intelligent autonomous systems that integrates perception, reasoning, planning, memory, learning, tool use, and multi-agent cooperation in one framework. This overview introduced the evolution, taxonomy, architecture, applications, and research problems of Agentic AI. Agentic AI facilitates more adaptable, goal-oriented, and autonomous decision-

making compared to standard AI techniques in several fields, including healthcare, cybersecurity, robotics, finance, and smart manufacturing. Although much progress has been gained, reasoning dependability, explainability, security, scalability, and governance are still critical issues to be considered for real-world application.

Future research should concentrate on the development of trustworthy and explainable Agentic AI systems that have better reasoning accuracy, efficient long-term memory management, scalable multi-agent cooperation, and safe autonomous decision-making. Further improvements in the dependability and practicality of intelligent autonomous systems will be achieved by advances in constant learning, human-AI cooperation, privacy-preserving approaches and energy-efficient structures. Furthermore, the ethical and widespread deployment of Agentic AI in crucial application areas would need the establishment of solid ethical governance structures and regulatory requirements.

References

- [1]. Russell, S. J., & Norvig, P. (2016). *Artificial Intelligence: A Modern Approach* (3rd ed.). Pearson.
- [2]. Wooldridge, M., & Jennings, N. R. (1995). Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2), 115–152.
- [3]. Rao, A. S., & Georgeff, M. P. (1995). BDI agents: From theory to practice. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS)* (pp. 312–319).
- [4]. Jennings, N. R. (2000). On agent-based software engineering. *Artificial Intelligence*, 117(2), 277–296.
- [5]. Wooldridge, M. (2009). *An Introduction to MultiAgent Systems* (2nd ed.). John Wiley & Sons.
- [6]. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press.
- [7]. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- [8]. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention Is All You Need. In *Advances in Neural Information Processing Systems (NeurIPS 2017)* (pp. 5998–6008).
- [9]. Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., et al. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484–489.
- [10]. Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., et al. (2017). Mastering the game of Go without human knowledge. *Nature*, 550(7676), 354–359.
- [11]. Foerster, J., Farquhar, G., Afouras, T., Nardelli, N., & Whiteson, S. (2018). Counterfactual Multi-Agent Policy Gradients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- [12]. Nwana, H. S. (1996). Software agents: An overview. *The Knowledge Engineering Review*, 11(3), 205–244.
- [13]. Vishram Singh, 2026. "Enterprise Resource Planning for Financial Applications: Engineering, Architecture, and Design Principles," *International Journal of Scientific Research in Science and Technology*, Technoscience Academy, vol. 13(2), pages 653–661, April.
- [14]. V. K. B. Parasaram, V. T. Bathini, S. K. Nalluri and A. H. Sannidhanam, "A Sustainable AI-Enabled Predictive Maintenance Framework for Smart Industrial Systems Using Industrial IoT Data," *2026 Third International Conference on Innovations in Cybersecurity and Data Science (ICICDS)*, Pathum Thani, Thailand, 2026, pp. 440–447, doi: 10.1109/ICICDS70526.2026.11604878.
- [15]. Vishram Singh, "Real-Time Financial Data Validation Platform for Enterprise Accounting Systems", *IJCISIM*, vol. 18, no. 10s, pp. 467–481, Jul. 2026.
- [16]. S. K. Nalluri, V. T. Bathini and V. K. B. Parasaram, "AI-Augmented Manufacturing Execution Systems for Precision Healthcare: A Scalable Framework for Smart Genomic and Proteomic Production," *2026 IEEE 2nd International Conference on Secure IoT, Assured and Trusted Computing (SATC)*, Houston, TX, USA, 2026, pp. 1–6, doi: 10.1109/SATC69565.2026.11542384.
- [17]. C. Koneni and S. Lokula, "Silent Data Corruption in Distributed Genai-Augmented Data Pipelines: an Empirical Study of Hallucination-Induced Failure Propagation," *2026 IEEE 46th International Conference on Distributed Computing Systems Workshops (ICDCSW)*, Seoul, Korea, Republic of, 2026, pp. 137–143,
- [18]. V. K. Bharadwaj Parasaram, V. Teja Bathini and S. K. Nalluri, "An Intelligent Digital Framework for Clinical Trial Monitoring Using AI and Cloud-Based Software Platforms," *2026 2nd International Conference on Artificial Intelligence Systems (AIS)*, San Antonio, TX,

- USA, 2026, pp. 524-530, doi:
10.1109/AIS69919.2026.11621841.
- [19]. Koneni, C., S. Lokula, and L. Panjiyar. "A Cloud-Native Approach to Statistical Anomaly Detection and Automated Data Quality Validation." *International Journal of Data Science and IoT Management System* 4.4 (2025): 533-543
- [20]. M. B K, V. K. Bharadwaj Parasaram, S. K. Nalluri and V. T. Bathini, "Attention-Guided Multi-Modal Deep Intelligence Framework for Early Cancer Detection Using Imaging and Genomic Signatures," *2026 International Conference on Data Science for Cyber-Physical Systems Resilience using Advanced Applications (ICDCA)*, BANGALORE, India, 2026, pp. 1-6, doi:
10.1109/ICDCA69396.2026.11620255.
- [21]. C. Koneni and S. Lokula, "Evidence-Quality Telemetry for Cloud Incident Response: Detecting Gaps, Drift, and Integrity Failures in Large-Scale Observability Pipelines," *2026 14th International Symposium on Digital Forensics and Security (ISDFS)*, Boston, MA, USA, 2026, pp. 1-6, doi:
10.1109/ISDFS69419.2026.11459097.