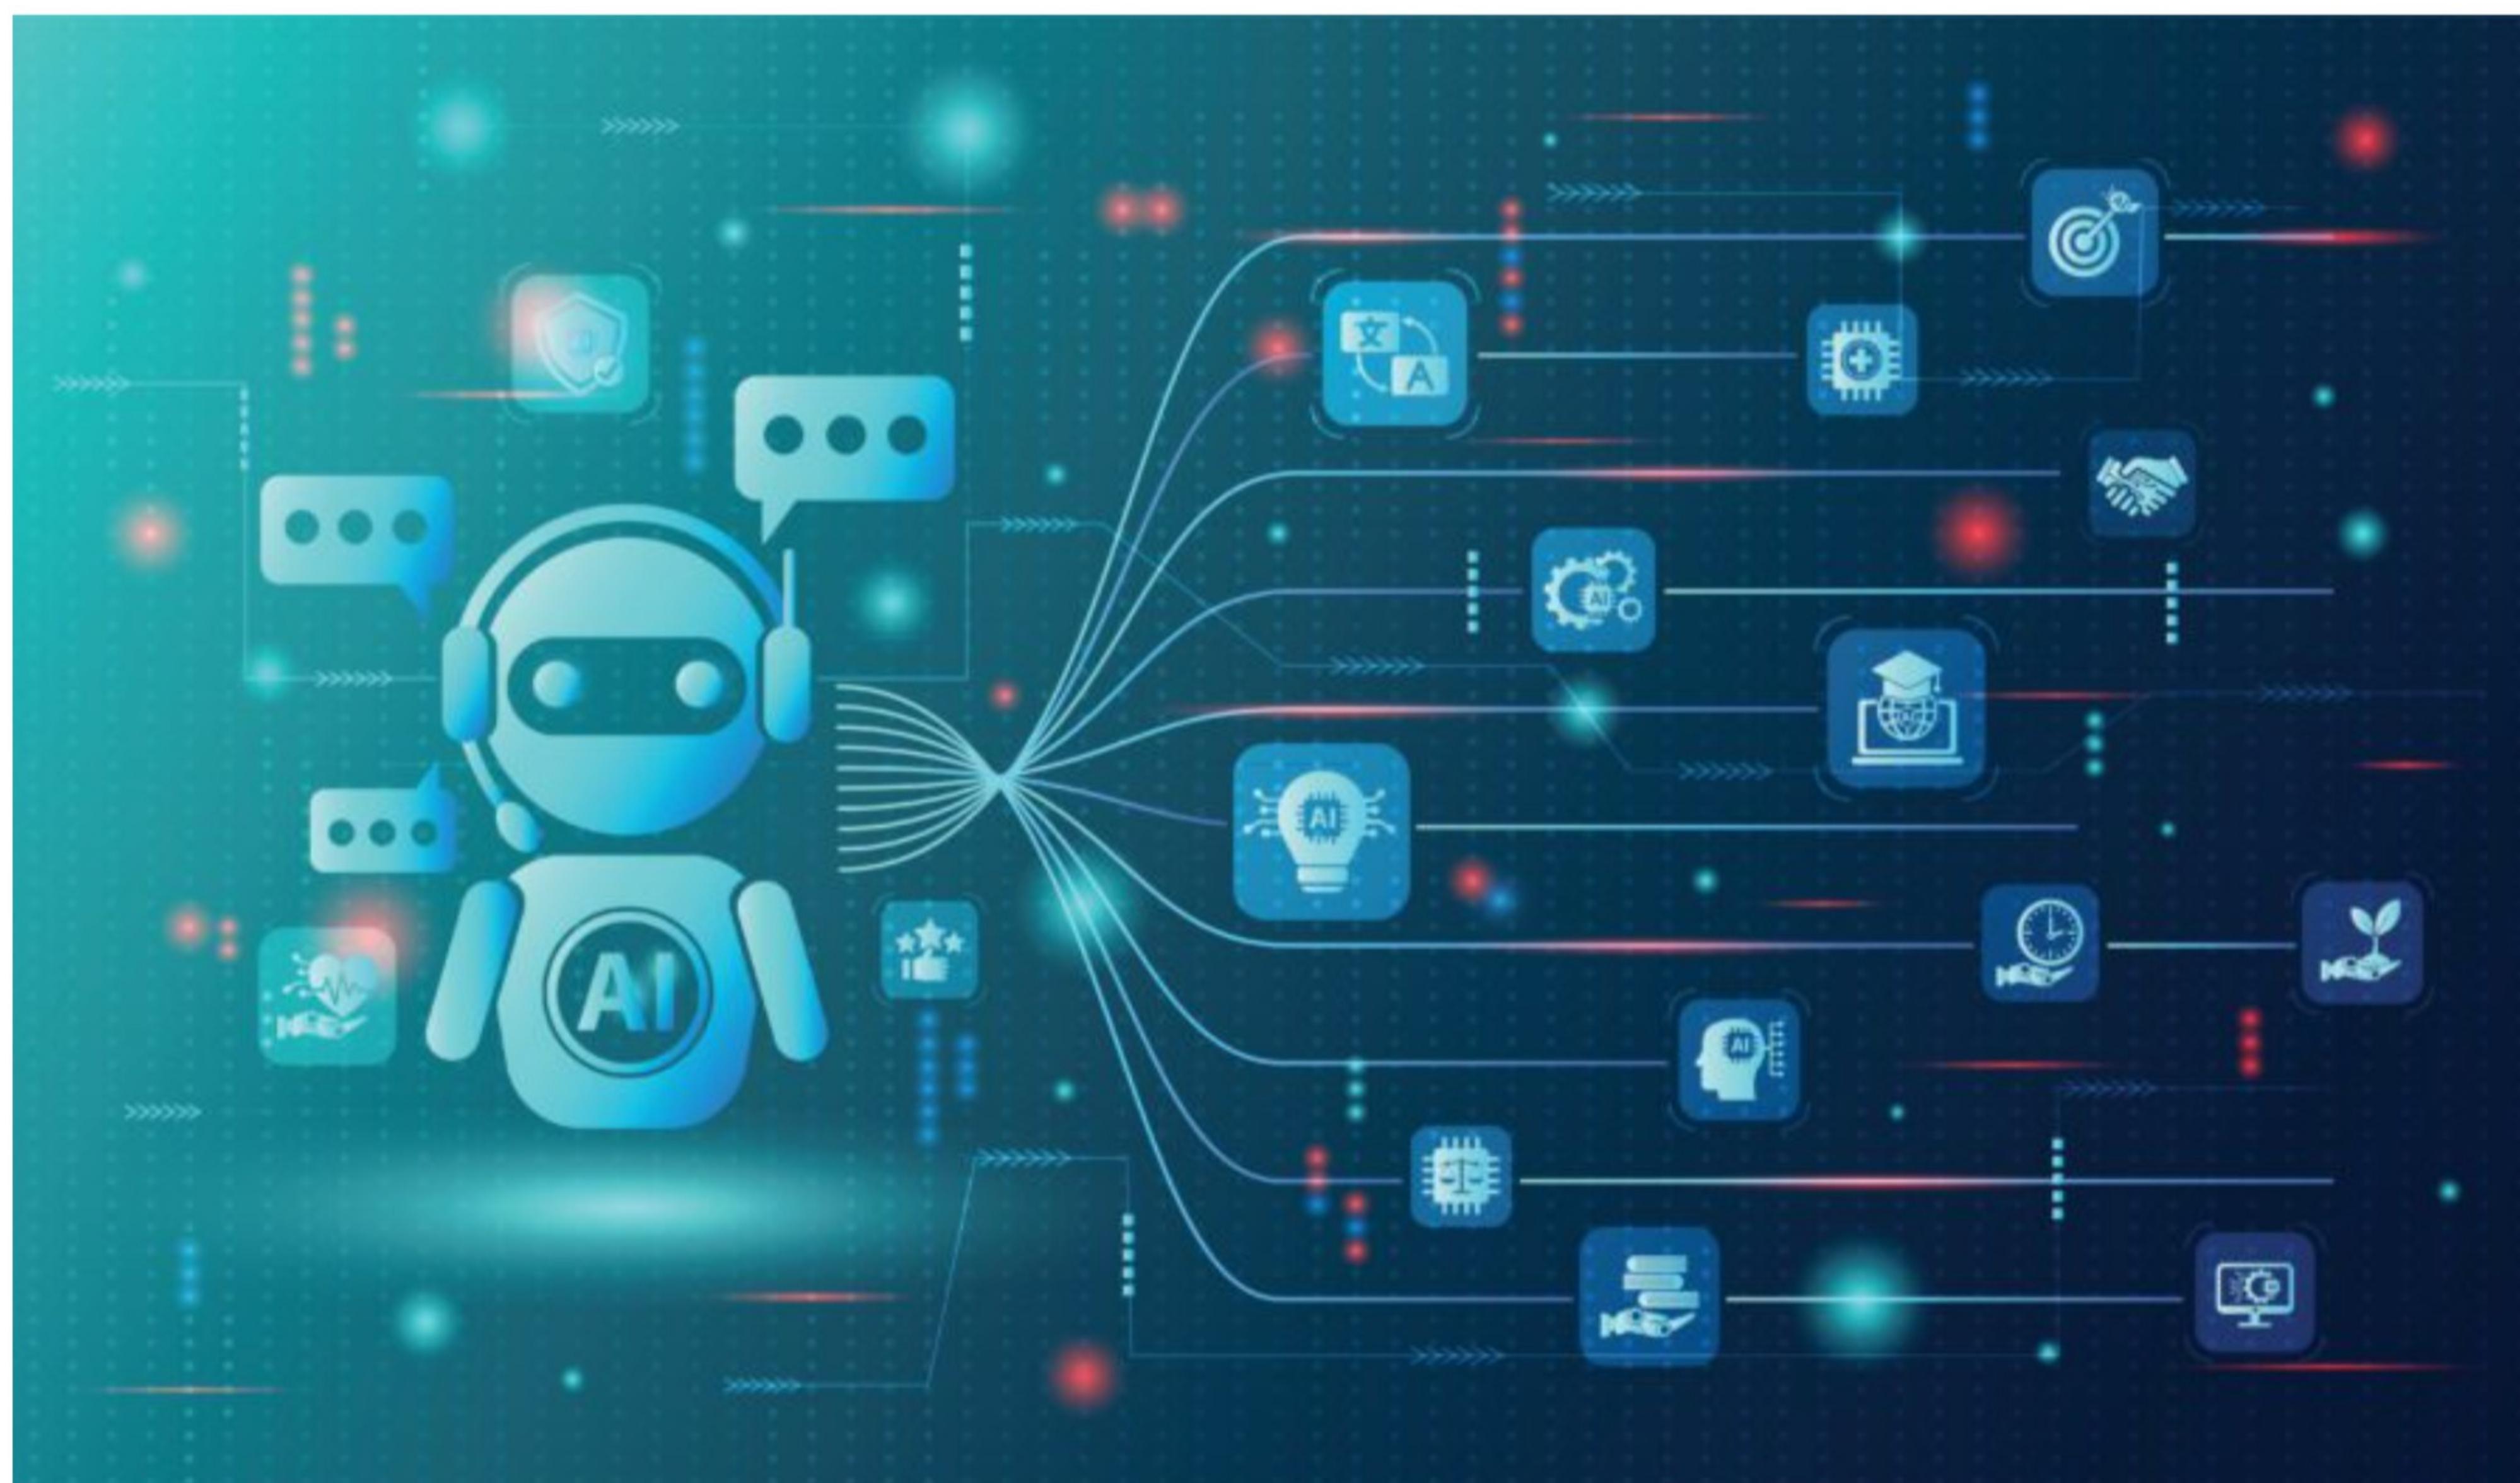


How graph thinking empowers agentic AI



Jans Aasman | July 29, 2025 at 11:35 am



Agentic AI systems are designed to adapt to new situations without requiring constant human intervention. These systems can provide tremendous benefits within many industries such as healthcare, supply chains, robotics, and autonomous vehicles. [Neuro-Symbolic Knowledge Graphs](#) (NSKGs) provide structured reasoning, contextual understanding, and long-term memory—critical elements for autonomous decision-making. By using NSKGs as the foundation for Agentic AI applications, we can seamlessly blend thinking and action, allowing intelligent systems to reason and respond dynamically.

Before we explore how these two technologies complement each other, let's first examine the [Neuro-Symbolic aspect of Knowledge Graphs](#). NSKGs are an advanced evolution of traditional knowledge graphs. These structures store information about entities (e.g., patients, customers, aircraft), their interrelationships, and relevant events. The primary function of NSKGs is to classify entities into different categories, diagnose anomalies, provide explanatory reasoning, and predict future behaviors and events. To achieve these goals, NSKGs integrate three distinct but equally crucial AI methodologies:

1. Rules and reasoning

When Generative AI (GenAI) first emerged, some believed that rule-based reasoning would become obsolete. However, it quickly became evident that structured reasoning remains critical for understanding entities, events, and behaviors. First-order logic (FOL) continues to be one of the most reliable and explainable methods for inference. Thus, a knowledge graph without robust support for a full FOL-based rule system or a Prolog compiler is significantly limited in its utility.

Related Content



[Bridging the digital defense gap: Why cybersecurity education matters more than ever](#)

Dan Wilson | July 24, 2025 at 11:25 am



[Data quality for unbiased results: Preventing AI-induced hallucinations](#)

Robert Stanley | July 22, 2025 at 10:45 am



[How to launch an AI startup in 2025](#)

Vincent Granville | July 20, 2025 at 11:07 pm



[How AI is transforming financial forecasting for SMBs and enterprises](#)

Dan Wilson | July 15, 2025 at 2:57 pm

2. Machine learning for classification and prediction

Machine learning techniques remain essential for uncovering patterns in data, classifying entities, and forecasting future behaviors. For instance, recurrent neural networks (RNNs) and long short-term memory (LSTM) models can analyze sequential events associated with entities such as patients, aircraft, or customers to predict the next significant occurrence. Unlike first-order logic and even GenAI models, RNNs excel at processing temporal sequences and making probabilistic predictions based on historical trends.

3. Generative AI

The integration of GenAI within knowledge graphs offers numerous advantages, particularly in handling unstructured data, which often constitutes over 70% of real-world knowledge graphs. GenAI can facilitate entity extraction, annotation, retrieval-augmented generation (RAG), automated SPARQL query generation, taxonomy matching, ontology creation, and ETL processes. While GenAI is not well-suited for complex first-order logic programming or classical machine learning tasks, it plays a crucial role in augmenting machine learning pipelines and streamlining data-driven workflows.

Why combine Agentic AI and NSKGs? Do they genuinely complement each other? As an experimental psychologist and cognitive scientist, I believe the answer is a resounding yes. Autonomous agents function as cognitive systems that perceive, decide, remember, act, and learn. Leading industry analyst firm [Gartner depicts Agentic AI as an AI practice](#) that employs an AI agent-driven approach to software implementation. These agents, whether autonomous or semi-autonomous, leverage AI techniques to perceive, decide, act, and achieve goals within digital or physical environments.

Anatomy of a Learning Agent

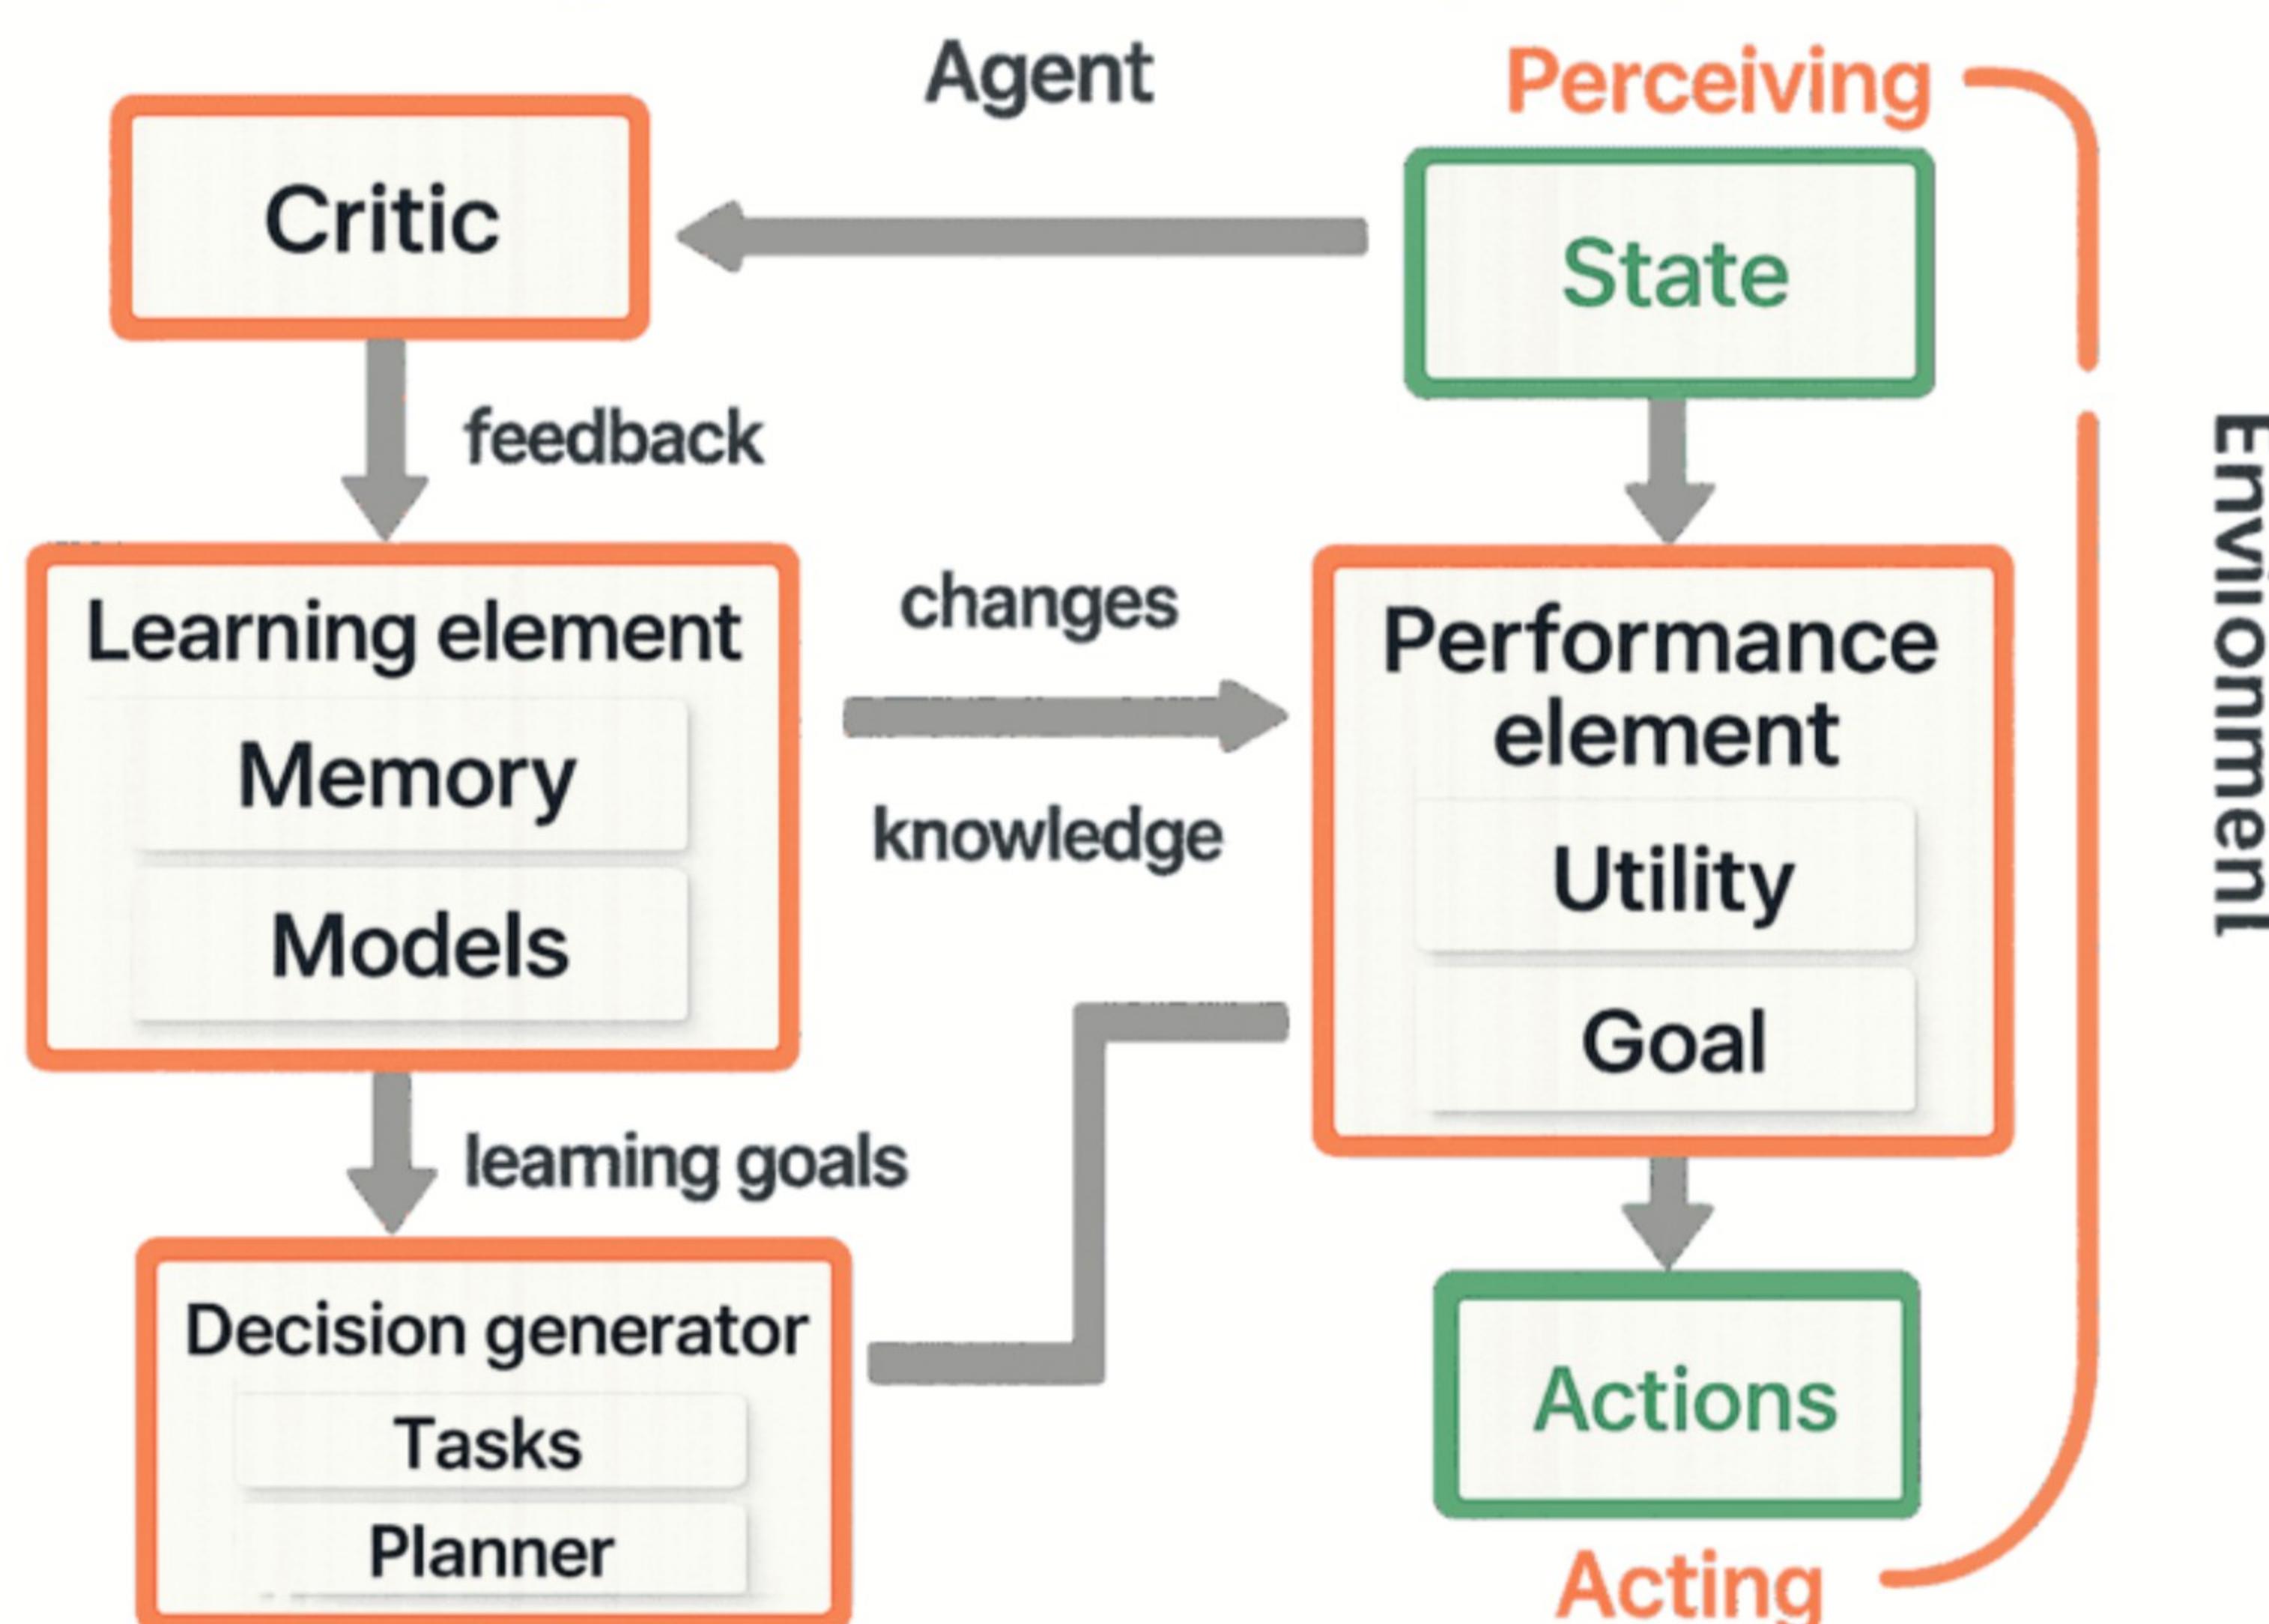


Image created by Franz Inc. using ChatGPT 4o

How neuro-symbolic knowledge graphs and agentic AI work together

Standardized Communication Language: Agentic AI systems must communicate effectively with one another. While intra-organizational agents can rely on proprietary communication protocols, agents operating in open ecosystems require standardized formats. A framework for inter-agent communication was outlined in *The Semantic Web* (Scientific American, May 2001) by Tim Berners-Lee, James Hendler, and Ora Lassila. At the time, the authors envisioned a future where software agents autonomously execute complex tasks by interpreting semantic data and collaborating through RDF, OWL, and knowledge graphs. To facilitate inter-agent comprehension, data must be semantically structured and self-descriptive, necessitating the use of knowledge graphs designed for meaningful information representation.

Long-Term Memory: Just as humans rely on both short-term and long-term memory, autonomous agents must retain and retrieve historical knowledge. Knowledge graphs serve as the long-term memory repository for agents, storing events, decisions, actions, and outcomes in a structured format. Events, represented as ontologically defined entities, enable predictive analytics using RNNs, GenAI-driven inference, and logical reasoning. This structured memory transforms knowledge graphs into a crucial component of neurosymbolic systems.

Orchestration Role: In multi-agent ecosystems, autonomous systems operate independently yet interact dynamically, often exhibiting emergent behaviors. A central knowledge graph acts as a system of record, capturing agent activities, outcomes, and decision rationales. By analyzing this data, a higher-order intelligence can optimize inter-agent collaboration through logical reasoning, machine learning-based insights, or GenAI-driven orchestration. In some cases, knowledge graphs may even assume a proactive role, guiding agents toward coordinated behavior.

Critics' Role: Agents equipped with diverse intelligence models — machine learning, GenAI, or symbolic logic — are prone to errors. By systematically recording decision-making rationales, actions, and resulting outcomes within a knowledge graph, meta-analysis can be performed to identify erroneous behaviors and refine agent performance. This critical evaluation function enhances system robustness and adaptability over time.

Learning and Adaptation: The structured data within knowledge graphs enables continuous learning and behavioral refinement. By leveraging stored interactions, agents can adapt through various mechanisms, including training large language models (LLMs), updating machine learning parameters, or refining logical rule sets. This iterative learning process fosters incremental improvements in agent decision-making and enhances the overall efficiency of autonomous systems.

Graph thinking in action

Agentic AI necessitates a robust, structured, and semantically rich framework for reasoning, decision-making, and long-term learning. Neuro-symbolic knowledge graphs provide the foundation for these capabilities by integrating logical inference, machine learning, and generative AI methodologies. By using NSKGS as the foundation for AI Agent systems, AI Agents gain long-term memory, enabling them to perceive, analyze, reason, and act simultaneously — adapting their decisions based on evolving circumstances in digital or physical environments.



AI AI Linguistics Business Analytics Cognitive Computing Data Strategist Data Trends Knowledge Engineering

◀ PREVIOUS

[Bridging the digital defense gap: Why cybersecurity education matters more than ever](#)

Leave a Reply

Your email address will not be published. Required fields are marked *

Name *

Email *

Website

Comment *

Save my name, email, and website in this browser for the next time I comment.

Post Comment

[About Us](#) [Contact Us](#) [Partner with Us](#) [Advertise with Us](#) [Write for Us](#) [RSS](#)

[Legal](#) [Terms of Service](#) [Privacy Policy](#) [Do Not Sell or Share My Personal Information](#) [Cookie Preferences](#)

© 2025 TechTarget, Inc. d/b/a Informa TechTarget. All rights reserved.



This website is owned and operated by Informa TechTarget, part of a global network that informs, influences and connects the world's technology buyers and sellers. All copyright resides with them. Informa PLC's registered office is 5 Howick Place, London SW1P 1WG. Registered in England and Wales. TechTarget, Inc.'s registered office is 275 Grove St. Newton, MA 02466.