

Knowledge Graphs as Agentic Systems: Integrating Memory, Action, and Code

1. Introduction

1.1. Context Setting

The field of artificial intelligence (AI) is witnessing a significant evolution towards agentic systems – AI entities characterized by their capacity for autonomous operation, proactive decision-making, and goal-directed behavior.¹ These systems represent a departure from earlier AI paradigms, such as traditional rule-based systems or purely generative models, which often lack the ability to act independently or adapt dynamically to complex environments.¹ Agentic AI aims to mimic human-like agency, enabling systems to interpret complex goals, understand context, plan multi-step actions, and learn from interactions with minimal human oversight.¹ However, realizing the full potential of agentic AI hinges critically on the development of robust mechanisms for knowledge representation, memory, and reasoning to underpin their sophisticated cognitive processes.⁷ Without effective ways to store, retrieve, and utilize knowledge about the world and past experiences, agents struggle to maintain context, make informed decisions, or execute complex tasks reliably.

1.2. Introducing Knowledge Graphs (KGs)

Knowledge Graphs (KGs) emerge as a compelling technology to address these needs. Fundamentally, KGs are structured representations of real-world entities (like objects, events, or concepts) and the intricate relationships that connect them, typically visualized as graph structures composed of nodes and edges.⁹ Their conceptual origins trace back to semantic networks and knowledge bases developed in earlier AI research.¹¹ The modern resurgence and popularization of KGs, however, were significantly propelled by large-scale commercial implementations, most notably Google's Knowledge Graph launched in 2012, which demonstrated their power in organizing vast amounts of information for applications like search enhancement.¹ KGs provide a way to model knowledge with explicit semantics, often leveraging formal ontologies to define the meaning of entities and relationships, thereby enabling both human and machine understanding.¹⁰

1.3. Thesis Statement

This report posits that Knowledge Graphs are evolving beyond their traditional role as static knowledge repositories to become dynamic, integral components of advanced agentic AI systems. They serve not only as sophisticated memory structures, capable of capturing both generalized semantic knowledge and specific episodic experiences, but also as actionable frameworks for grounding agent behavior, planning complex tasks, and integrating executable code. This deep integration of KGs within agentic architectures enables more complex, context-aware, adaptable, explainable, and ultimately more capable autonomous systems.

1.4. Report Structure Overview

This report will systematically explore the synergistic relationship between KGs and agentic AI. Section 2 establishes the foundational concepts of KGs relevant to agency, focusing on their structure, semantics, and key characteristics. Section 3 delves into the principles of agentic AI systems, outlining their core characteristics, operational cycles, and architectures. Section 4 examines the crucial role of KGs as memory systems for agents, covering both semantic and episodic memory representations and dynamic updates, illustrated through specific frameworks. Section 5 investigates how KGs enable and guide agent actions, including modeling procedures, driving planning, and interfacing with tools. Section 6 explores the convergence of graph structures and code representations, discussing embeddings, codebase KGs, and agentic orchestration frameworks. Section 7 addresses the significant challenges and promising future directions in this rapidly evolving field. Finally, Section 8 provides concluding remarks, synthesizing the key findings and reinforcing the transformative potential of integrating KGs within agentic AI.

The confluence of technologies like Large Language Models (LLMs), KGs, and agentic frameworks signifies more than mere technological combination; it marks a fundamental shift towards AI systems possessing dual capabilities.¹⁹ LLMs contribute proficiency in processing and generating natural language, drawing upon vast unstructured data, yet they often suffer from a lack of grounding, potential for hallucination, and opacity in their reasoning.⁴ Conversely, KGs offer structured knowledge, explicit semantics, and the potential for verifiable, formal reasoning, but face challenges in construction and in handling the ambiguity and scale of real-world information.¹² Agentic frameworks provide the necessary architecture—often based on a perception-reasoning-action loop—to orchestrate these components effectively.³ The interaction is synergistic: LLMs can aid in interpreting unstructured data to populate and enrich KGs²², while KGs provide the structured context and factual grounding to constrain and inform the reasoning processes within the agent, often mediated by LLMs.²³ This symbiosis is essential for developing autonomous systems that are not only intelligent but also reliable and trustworthy.

2. The Knowledge Graph Foundation for Agency

2.1. Defining the Modern KG for Agentic Systems

While basic definitions describe KGs simply as nodes (entities), edges (relationships), and labels,⁹ a richer understanding is necessary when considering their role in agentic AI. For agency, a KG is more accurately conceptualized as a dynamic knowledge base characterized by formal semantics and a graph-structured data model.¹⁶ They represent networks of real-world entities—objects, events, concepts—and their interrelations, often stored and queried using graph database technologies.¹¹

A key aspect is the integration with **ontologies**, which provide a formal, explicit specification of a shared conceptualization.¹⁰ Ontologies act as the schema layer for the KG, defining the types of entities (classes), their attributes (properties), and the types of relationships that can exist between them.¹⁰ This formal semantic underpinning distinguishes KGs from simpler graph databases or raw RDF datasets, which might lack

this rich layer of meaning.¹⁶ The **Resource Description Framework (RDF)** triple format (subject, predicate, object) serves as a common standard for representing facts within KGs, facilitating interoperability and grounding in Semantic Web principles.¹¹

2.2. Key Components and Characteristics for Agents

Several components and characteristics of KGs are particularly vital for supporting agentic systems:

- **Entities (Nodes) & Relationships (Edges):** The core structure allows modeling of diverse elements relevant to an agent's environment and tasks, including physical objects, abstract concepts, people, locations, events, and the complex ways they connect.⁹
- **Ontologies & Semantics:** Formal semantics, often defined by ontologies, enable machines (agents) to process and interpret the information within the KG reliably and unambiguously.¹² This shared understanding is crucial for tasks like data integration, reasoning, and communication between agents.¹⁰
- **Identifiers (URIs):** The use of unique and global identifiers, such as Uniform Resource Identifiers (URIs) common in RDF and Semantic Web contexts, ensures that entities and relations are unambiguously identified.¹⁰ This prevents confusion and is fundamental for linking data across different parts of the graph or even across different KGs.
- **Flexibility & Extensibility:** Unlike traditional relational databases with rigid schemas, KGs offer significant flexibility.¹² New types of entities, relationships, and data instances can be added incrementally without requiring extensive schema refactoring.¹⁵ This adaptability is essential for agents that need to learn and update their knowledge base continuously in dynamic environments.
- **Integration Capability:** KGs excel at integrating heterogeneous data from diverse sources, including structured databases, unstructured text, APIs, and semi-structured data.¹⁰ This allows agents to build a comprehensive world model from fragmented information.
- **Queryability & Reasoning:** KGs support powerful query languages (e.g., SPARQL for RDF KGs, Cypher for property graphs) that allow agents to retrieve specific information and explore complex relationships.²⁵ Furthermore, the formal semantics enable inference mechanisms that can derive implicit knowledge (new facts) from the explicitly stated information.¹⁰

The precise definition and implementation details of KGs can vary. Some KGs are **entity-centric**, focusing primarily on representing entities and their properties, while others are **content-centric**, linking chunks of content like text sections or table columns.¹⁵ KGs also range from large, general-purpose resources like DBpedia and Wikidata, which aggregate knowledge from sources like Wikipedia¹¹, to highly specialized **domain-specific** or **enterprise KGs** tailored to particular applications or organizational knowledge.²⁹

The characteristics of "unambiguous identification"³⁷ and "formal semantics"¹², often realized through ontologies and standardized identifiers like URIs, are not merely technical specifications. They form the bedrock upon which *explainable and verifiable* agentic behavior can be built. Agentic systems, by definition, make autonomous decisions and take actions.¹ For these systems to be trusted and deployed, especially in critical applications, their decision-making processes must be traceable and understandable.³⁰ KGs provide a mechanism for this transparency. When an agent retrieves information from a KG or performs reasoning based on its structured, semantically defined content¹², the path to its conclusion can often be reconstructed by examining the specific entities, relationships, and inference rules involved.³⁷ This explicit knowledge representation contrasts sharply with the often opaque internal workings of purely connectionist models like LLMs. Thus, the formal structure inherent in well-designed KGs directly facilitates a level of auditability and explainability that is vital for building trustworthy autonomous agents.¹⁸

Furthermore, the inherent flexibility and capacity for schema evolution found in KGs¹² are fundamental enablers for agents designed for *lifelong learning*.⁴³ Such agents must continually acquire new information, integrate novel concepts, and adapt their internal world models over extended periods.⁴³ Traditional database systems, constrained by rigid, predefined schemas, pose significant obstacles to this continuous adaptation, often necessitating complex and disruptive data migrations.³⁶ KGs, particularly those based on flexible models like RDF, readily accommodate the incremental addition of new data, new types of entities, and new kinds of relationships without necessarily invalidating the existing structure or requiring a complete overhaul.¹⁶ This intrinsic adaptability of the knowledge graph structure mirrors the desired adaptive learning capability of the agent itself, positioning KGs as a natural and powerful substrate for representing the dynamically evolving knowledge base of a lifelong learning system.

3. Agentic AI Systems: Architecture and Operation

3.1. Defining Agentic AI

Agentic AI refers to a class of artificial intelligence systems designed to operate with a significant degree of autonomy, enabling them to perceive their environment, make decisions, and take actions to achieve specific goals with minimal direct human supervision.¹ The term "agentic" highlights the core concept of "agency"—the capacity of these systems to act independently and purposefully within their environment.⁴ These systems often leverage advanced AI techniques, including machine learning, natural language processing, and increasingly, large language models (LLMs), to interpret complex situations, formulate plans, and execute tasks.¹

3.2. Core Characteristics

Agentic AI systems are distinguished by several key characteristics:

- **Autonomy:** They can initiate and complete tasks independently, without constant human guidance or intervention.¹
- **Goal-Oriented:** Their actions are driven by predefined objectives or goals they aim to achieve.¹
- **Proactivity/Intentionality:** They possess the ability to take initiative, set sub-goals, and plan sequences of actions rather than merely reacting to stimuli.²
- **Reactivity/Adaptability:** They can perceive changes in their environment and adjust their behavior or plans accordingly. Many agentic systems incorporate learning mechanisms to improve performance over time based on feedback and experience.¹
- **Reasoning:** They exhibit capabilities for complex problem-solving, making judgments based on context, weighing trade-offs, and deducing appropriate courses of action.²
- **Interaction:** They can interact with their environment, which may include digital systems (databases, APIs), physical environments (via sensors and actuators), or other agents.¹

3.3. Operational Cycle

The operation of agentic AI systems is often conceptualized as a continuous cycle involving several key phases, sometimes referred to as a "think-act-observe" loop or similar variations.⁶

1. **Perception:** The agent gathers information about its current state and environment through various inputs, such as sensors, APIs, database queries, or direct user interaction.⁴
2. **Reasoning/Planning:** The agent processes the perceived information, often using LLMs as a core reasoning engine or orchestrator, to understand the situation, interpret goals (potentially given in natural language), evaluate options, and formulate a plan or strategy.¹ This phase involves making judgments and deciding on the next course of action.
3. **Action:** The agent executes the chosen action or plan. This often involves interacting with external tools, calling APIs, manipulating data, generating responses, or controlling physical actuators.⁴
4. **Learning/Reflection:** The agent observes the outcome of its actions and potentially receives feedback. This information is used to evaluate performance, update its internal state or memory, and refine its future reasoning, planning, and actions.¹ Techniques like reinforcement learning are often employed in this phase.¹

3.4. Architectures and Types

Agentic AI systems can range from single, sophisticated agents to complex **Multi-Agent Systems (MAS)** where multiple agents collaborate or compete.³ Coordinating the actions of multiple agents often requires **AI orchestration** mechanisms.³ Architectures can be hierarchical, with a central "conductor" agent delegating tasks, or more decentralized, with agents coordinating peer-to-peer.⁴ Agents themselves can be classified based on their capabilities and design principles, such as **reactive agents** that respond directly to stimuli versus **proactive agents** that plan towards goals,² or along axes like complexity (e.g., simple reflex, model-based, goal-based, utility-based agents^{4,5}) or interaction style (e.g., reactive, deliberative, learning, collaborative agents^{2,6}).

The operational cycle (Perceive-Reason-Plan-Act-Learn)⁶ common to agentic systems is not merely a linear sequence but an integrated feedback loop. Within this loop, the agent's internal representation of knowledge and experience—its memory—plays a pivotal role at every stage. The quality, structure, and timeliness of this memory profoundly influence the agent's overall performance. Perception⁴ is enhanced when incoming sensory data or information can be contextualized against existing knowledge; a KG can provide this semantic context to help disambiguate observations or relate them to known entities and concepts.^{2,8} The Reasoning and Planning phases² are fundamentally dependent on the agent's world model and its recollection of past events. KGs, by providing structured, interconnected knowledge, enable more sophisticated reasoning compared to relying solely on unstructured conversation history or retrieved text snippets.^{2,8} The selection of Actions⁴ is directly guided by the plans derived from this KG-informed reasoning. Finally, the Learning and Reflection phase¹ involves updating the agent's internal state based on the outcomes of its actions. Dynamic KGs offer a concrete mechanism for this memory update, allowing the agent to integrate new facts, modify existing beliefs, or track the evolution of its environment over time.^{2,8} Consequently, the KG is not a passive data store but an active, dynamic substrate that shapes and is shaped by the entire agentic process. A more comprehensive, accurate, and dynamically maintained KG directly translates to improved perception, reasoning, planning, action selection, and learning adaptation.

As agentic AI evolves, there is a clear trend towards **Multi-Agent Systems (MAS)**, where tasks are decomposed and handled by multiple, specialized agents.³ Examples include systems with distinct agents for coding, web searching, planning, or specific domain expertise.⁴ While specialization enhances efficiency and capability within narrow domains, it simultaneously creates a critical need for effective coordination and shared understanding among these diverse agents.⁷ Simple message passing or sharing of unstructured text between agents can be inefficient and prone to misinterpretation or loss of context. A robust, shared knowledge representation layer becomes essential to ensure coherent collaboration. Knowledge Graphs are ideally suited for this role, providing a structured, semantic "lingua franca".⁷ Agents can read from and write to a common KG, allowing specialized knowledge acquired by one agent (e.g., a new fact discovered by a web search agent) to be integrated into the shared model and become accessible to other agents (e.g., a planning agent) in a consistent, unambiguous format. This shared KG effectively functions as the collective memory and world model for the MAS, enabling complex, coordinated behaviors that emerge from the interaction of specialized components, grounded in a common understanding of the environment and task context.^{4,7}

4. Knowledge Graphs as Agent Memory

A defining feature of advanced agentic AI is the ability to learn, remember, and utilize past experiences and knowledge to inform future actions. Traditional approaches to AI memory, particularly within the context of LLMs, face limitations. The finite context window of LLMs restricts the amount of historical information that can be directly processed, leading to a loss of long-term coherence.⁵⁴ While Retrieval-Augmented Generation (RAG) addresses this by retrieving relevant text chunks from external corpora, standard RAG often lacks the structure needed to represent complex relationships or temporal dynamics effectively, and struggles with frequently updated information sources.²² KGs offer a compelling alternative, providing a structured, dynamic, and queryable foundation for agent memory.

4.1. KGs for Semantic Memory

KGs excel at representing **semantic memory** – the general knowledge about the world, including facts, concepts, entities, and the relationships between them.⁵⁰ By encoding this knowledge in a structured graph format (nodes connected by labeled edges), agents can:

- **Access Factual Knowledge:** Retrieve specific facts about entities or concepts relevant to their current task (e.g., "What is the capital of France?").
- **Understand Relationships:** Navigate the graph to understand how different entities are connected (e.g., "How is Person A related to Company B?").
- **Provide Context:** Use the graph's structure and semantics to disambiguate information or understand the context of observations.²⁸ For example, distinguishing between "Apple" the company and "apple" the fruit based on connected entities and relationships.
- **Enable Reasoning:** Perform inference over the graph to derive implicit knowledge not explicitly stated (e.g., inferring a common ancestor based on family relationships).

4.2. KGs for Episodic Memory

Beyond general knowledge, agents need **episodic memory** to recall specific events, interactions, or experiences tied to particular times and contexts.⁵⁰ KGs can be adapted to store this type of memory by incorporating temporal and contextual information:

- **Temporal Tagging:** Facts or relationships (edges) in the KG can be annotated with timestamps or validity intervals indicating when they were true.⁵⁴ This allows the agent to reconstruct past states or reason about sequences of events.
- **Contextual Qualifiers:** Techniques like RDF* or RDF quadruples extend the basic triple model to add context, such as the time, location, or source of a particular fact.⁵³ For instance, a triple `(Agent, atLocation, Library)` could become `(Agent, atLocation, Library, {timestamp: 42})` to represent an episodic memory.⁵³
- **Representing Interactions:** Sequences of user-agent interactions or agent actions within an environment can be stored as linked events or episodes within the KG structure.⁵⁴

4.3. Dynamic Memory Updates

A crucial aspect of agent memory is its ability to evolve. Unlike static knowledge bases, the KG serving as an agent's memory must be dynamically updated as the agent perceives, acts, and learns.²¹ This involves:

- **Adding New Knowledge:** Incorporating new entities, relationships, or facts derived from observations, user input, or task execution.
- **Modifying Existing Knowledge:** Updating entity properties or relationship attributes based on new information.
- **Temporal Management:** Marking facts as no longer valid or updating their validity intervals as the environment changes.⁵⁴

This continuous interaction creates a feedback loop, sometimes termed a "data flywheel," where the agent's experiences constantly enrich the KG, making it a more accurate and comprehensive representation of the agent's world over time.²⁸

4.4. Memory Retrieval Mechanisms

Agents need efficient ways to retrieve relevant information from their potentially vast KG memory. Common mechanisms include:

- **Graph Queries:** Using formal query languages like SPARQL or Cypher to retrieve specific nodes, edges, or subgraphs matching certain patterns.²⁵ This allows for precise, structured retrieval.
- **Graph Traversal:** Navigating the graph structure by following edges from known entities to discover related information.
- **Embedding-Based Search:** Representing nodes and edges as vectors (embeddings) and using similarity search (e.g., vector databases) to find entities or facts semantically related to a query.⁵⁴
- **GraphRAG:** Applying RAG principles specifically to KGs, where retrieved graph structures (subgraphs, paths, entities) provide context to an LLM for generating responses or making decisions.²⁰

4.5. Illustrative Frameworks

Several research frameworks exemplify the use of KGs for agent memory:

- **AriGraph:** This framework focuses on agents operating in interactive text environments. The agent dynamically constructs and updates a "memory graph" that integrates both semantic facts about the environment and episodic memories of its actions and observations. This structured graph facilitates efficient associative retrieval of relevant knowledge, enhancing the agent's planning and decision-making.

capabilities compared to unstructured memory approaches.⁵⁰

- **Zep/Graphiti:** Designed as a production-ready memory layer, Zep utilizes the Graphiti engine to build a dynamic, temporally-aware KG.⁵⁴ It features distinct subgraphs for raw episodic data, derived semantic entities/facts, and higher-level community structures. Its bi-temporal modeling explicitly tracks both event time and ingestion time, allowing for sophisticated historical queries and accurate representation of evolving knowledge. Zep emphasizes low latency and scalability for real-world agentic applications.⁵⁴
- **ODA (Observation-Driven Agent):** ODA proposes a cyclical paradigm where the agent actively observes the global KG structure as part of its reasoning loop.²⁴ It uses a recursive observation mechanism to manage the potential complexity of large KGs. Observed KG patterns are then integrated into the action and reflection modules, allowing the agent to synergistically leverage both LLM inference and KG structural reasoning throughout its task-solving process.³⁵
- **Agentic Reasoning Framework (Mind Map):** This framework introduces a specialized "Mind Map" agent whose role is to construct a KG dynamically from the main reasoning agent's context or thought process.³³ This KG structures the logical relationships emerging during reasoning, allowing the primary agent (or other tools) to query this "map" to retrieve relevant context, check consistency, or gain insights into the reasoning trajectory itself.³³

The explicit modeling of both semantic and episodic memory within a unified KG structure offers significant advantages, particularly for agents operating in complex, dynamic, or partially observable environments.⁵⁰ Human cognition seamlessly blends general knowledge with specific past experiences to navigate uncertainty and make decisions.⁵³ Traditional AI systems often struggle to replicate this integration. By implementing distinct yet interconnected representations for semantic facts and time-bound episodes within the KG⁵⁰, an agent gains the ability to: (1) interpret current observations using its general understanding of the world (semantic component); (2) recall relevant past events or interactions that provide context or precedent for the current situation (episodic component); and (3) refine its general knowledge based on new experiences (linking episodes to semantic updates). This structured approach, enabled by the KG, directly addresses the challenges posed by Partially Observable Markov Decision Processes (POMDPs) by allowing the agent to maintain and reason over a much richer belief state, informed by both its accumulated general knowledge and its specific history.⁵³

Furthermore, the development and use of *temporal KGs*⁵⁴ represent a critical advancement. Their importance extends beyond simply storing episodic memories; they empower agents to reason explicitly about *change* and *causality* over time. Standard KGs typically capture a static snapshot of knowledge. However, agentic systems must operate and make decisions within environments that are inherently dynamic.¹ Temporal KGs, by incorporating mechanisms like timestamping or validity intervals for facts and relationships⁵⁴, enable agents to query the state of their knowledge base at specific points in the past (e.g., "What relationships held true for entity X yesterday?"). This capability facilitates reasoning about the evolution of entities and relationships, which is fundamental for understanding ongoing processes, predicting potential future states, and inferring causal dependencies (e.g., "Did event A precede event B? Could A have caused B?"). This moves the agent's capabilities beyond simple fact retrieval towards genuine temporal reasoning, a cornerstone of effective planning and action in environments characterized by continuous change.

Table 1: Comparison of KG-Based Agent Memory Frameworks

Feature	AriGraph	Zep/Graphiti	ODA	Agentic Reasoning (Mind Map)
Core Concept	Memory graph as world model for LLM agents in interactive envs.	Production-focused, dynamic, temporal memory layer service.	Observation-driven cyclical integration of KG structure and LLM reasoning.	KG ("Mind Map") dynamically built from reasoning context to structure logic.
KG Structure	Dynamically constructed graph integrating semantic & episodic info.	Dynamic, temporally-aware KG with Episodic, Semantic Entity, and Community subgraphs. Bi-temporal model.	Assumes existing KG; focuses on observing/integrating global KG structure. Recursive observation for scale.	Dynamically constructed KG representing reasoning steps, entities, and relationships.
Memory Types	Semantic & Episodic.	Semantic & Episodic (explicitly modeled via subgraphs and temporal tracking). Summaries (Community subgraph).	Primarily leverages existing KG structure (semantic/factual) through observation.	Primarily structures the reasoning process itself (meta-memory). Can store facts/entities mentioned during reasoning.
Update Mechanism	Agent updates graph based on exploration/interaction.	Incremental, real-time updates via Graphiti engine; handles temporal validity.	KG assumed external/pre-existing in examples; focus is on observing, not agent-driven updates to the main KG.	Mind Map agent updates KG based on ongoing reasoning chain from primary agent.
Retrieval Method	Associative retrieval based on relevance to state/goals.	Hybrid: Graph queries, vector search (embeddings), keyword search, graph traversal. Low latency focus.	"Observation" module retrieves relevant subgraphs based on task/entities. Action module can query KG.	Querying the Mind Map KG using RAG/standard KG queries.
Key Strengths	Tailored for interactive environments; integrates memory	Production-ready, scalable, low latency, explicit temporal reasoning, handles unstructured &	Deep integration of KG structure into reasoning loop; handles KG	Structures the reasoning process itself; improves deductive reasoning; facilitates querying the

	types.	structured data.	scale via recursion.	logic flow.
Limitations	Evaluated mainly in TextWorld; scalability less explored.	Newer framework; requires specific infrastructure (e.g., Neo4j for Graphiti). Benchmark coverage still developing.	Assumes access to KG; effectiveness depends on quality of observation mechanism and underlying KG.	Focuses on structuring reasoning, less on comprehensive world knowledge or long-term episodic memory storage.

5. Knowledge Graphs Enabling Agent Action

Beyond serving as memory repositories, KGs play an active role in enabling and guiding agent actions. They provide the structured knowledge required for planning, decision-making, and interacting effectively with tools and the environment.

5.1. Modeling Actions and Procedures

- KGs offer a powerful formalism for representing procedural knowledge – the "how-to" aspect of tasks.³⁰ This involves modeling elements like:
- Tasks and Goals:** Representing overall objectives as nodes in the graph.
 - Plans and Workflows:** Modeling sequences of steps or actions required to achieve a goal, often using relationships like `precededBy` or `hasStep`.³⁰ Frameworks like P-Plan or K-Hub provide ontological structures for this.³⁰
 - Actions:** Representing specific actions as nodes or relationships, potentially linked to preconditions, effects, required resources, or agent capabilities.³⁰
 - Objects and Equipment:** Modeling the entities acted upon (direct objects) and the tools or resources needed (equipment) for specific steps or actions.³⁰ Ontologies like FRAPO can be used for equipment representation.³⁰

LLMs can be employed to extract this procedural knowledge from unstructured text sources like manuals or recipes, identifying steps, actions, objects, equipment, and temporal information, and structuring it according to a predefined ontology to populate the KG.³⁰ This transforms passive textual descriptions into actionable, machine-interpretable representations within the graph. Graphs are naturally suited for describing processes, mirroring how humans often use diagrams with nodes and arrows to explain complex workflows.⁷⁸

5.2. KG-Driven Planning and Reasoning

- Once actions and procedures are modeled, the KG becomes a crucial resource for agent planning and reasoning.⁷ Agents can leverage the KG to:
- Identify Possible Actions:** Query the graph to determine available actions based on the current state, agent capabilities, or task context.
 - Check Preconditions and Effects:** Verify if the necessary conditions for an action are met (based on the current state represented in the KG) and predict the likely outcomes.
 - Formulate Plans:** Use graph traversal algorithms (e.g., pathfinding), constraint satisfaction techniques, or more sophisticated AI planning methods operating over the KG structure to find sequences of actions that lead from the current state to a desired goal state.⁸
 - Contextualize LLM Planning:** Provide structured knowledge from the KG as context to LLMs tasked with generating plans, helping to ground the LLM's output in factual constraints and available actions.²⁸

5.3. Interfacing with Tools and Code Execution

- Agentic systems often rely on external tools, APIs, or executable code to perform actions in the real world or digital environments.⁴ KGs can serve as an interface layer or registry for these tools:
- Tool Representation:** Tools, APIs, or functions can be represented as nodes within the KG.²⁸
 - Parameter and Context Retrieval:** The KG can store metadata about tools, such as required parameters, expected inputs/outputs, and contextual information needed for their execution. Agents can query the KG to find the appropriate tool for a sub-task and retrieve the necessary information to invoke it correctly.³⁴
 - Action Grounding:** The planning process, guided by the KG, can identify not just the conceptual action but the specific tool or code function needed to realize that action.

5.4. Text-to-Query Translation

To facilitate interaction, particularly for LLM-based agents or human users, mechanisms are needed to translate natural language queries or commands into formal graph queries executable on the KG.³² Techniques like Text2Cypher aim to interpret the user's intent and generate the corresponding Cypher query for graph databases like Neo4j.³² While powerful, the accuracy of these translation methods is still an active area of research, as nuances in language and the need for schema-specific details pose challenges.³⁸ Agentic approaches involving multi-step query formulation or retries are being explored to improve robustness.³⁸

Encoding procedures and actions directly within the Knowledge Graph structure³⁰ elevates the agent's capabilities beyond simple task execution. Instead of merely planning a sequence of calls to external, opaque tools, the agent gains the ability to *reason about the procedures*

themselves. When the steps of a workflow, their sequence, the actions involved, their preconditions, and required resources are explicitly represented as interconnected nodes and edges in the KG³⁰, the agent can query this structure introspectively. It can ask questions like: "What are the alternative sequences of steps to achieve goal G?", "What resources are required for step S?", "Which step in procedure P is the most time-consuming?", or "How does procedure X differ from procedure Y?". This enables a form of meta-reasoning about workflows. The agent can compare different approaches, identify potential inefficiencies or bottlenecks, evaluate the feasibility of a procedure given current resources, or even adapt or modify procedures based on contextual factors. This capability signifies a deeper level of agency, moving from merely following instructions (even self-generated ones) towards understanding, optimizing, and adapting the processes themselves.

Furthermore, the capacity to establish direct links between KG entities or relationships and executable code segments or external API calls⁴ constructs a vital bridge between declarative knowledge (facts about the world stored in the KG) and procedural knowledge (the mechanisms for acting in that world). Traditionally, these two forms of knowledge reside in separate realms: the KG holds the 'what', while code or APIs embody the 'how'. Agents would query the KG to inform their decisions, then invoke separate procedural components. By integrating representations of actions directly, for example, by having a KG node representing an API endpoint store properties defining its parameters and linking to the actual function call³¹, this separation dissolves. When the agent's reasoning process identifies the need for a specific action represented in the KG, it can potentially retrieve not just the factual context but also the executable instructions or interface details directly from the graph-linked information. This tight coupling makes the KG an active participant in task execution, transforming it from a passive knowledge repository into an actionable component central to the agent's operational capabilities.

6. Integrating Graph Structures and Code Representations

The synergy between KGs and agentic systems extends to the realm of software development and execution itself. Integrating the structured knowledge representation of graphs with the executable logic of code is crucial for building agents that can understand, reason about, generate, and interact with software systems.

6.1. Bridging the Gap: Embeddings

Knowledge Graph Embeddings (KGEs) provide a fundamental bridge between the symbolic nature of KGs and the numerical processing required by many machine learning models, including those used within agents.³⁷ KGEs learn low-dimensional vector representations for entities and relations in the KG, aiming to capture their semantic meaning and preserve the graph's structural properties in the vector space.⁸⁶

- **Functionality:** Embeddings allow for efficient computation of semantic similarity between entities or relations, facilitate tasks like link prediction (inferring missing relationships), entity classification, and clustering, and provide numerical features for downstream ML tasks within the agent.⁸⁶
- **Models:** Various KGE models exist, broadly categorized into translational distance models (like TransE, TransH, TransR, which model relationships as translations in the vector space) and semantic matching models (like RESCAL, DistMult, ComplEx, which use matrix/tensor factorization or neural networks).⁸⁶
- **Tools:** Libraries like PyKeen, AmpliGraph, OpenKE, and DGL-KE provide implementations of various KGE models and facilitate training and evaluation.⁸⁷
- **Limitations and Advancements:** While powerful, traditional KGEs often produce static embeddings, potentially limiting their ability to capture dynamic changes in the KG.⁹⁰ Research is exploring dynamic KGEs⁹¹ and novel approaches like embedding KGs in function spaces, which may offer greater expressiveness and allow operations like composition or differentiation on the representations.⁸⁹

6.2. Codebase Knowledge Graphs

A specific application of KGs relevant to agentic systems involved in software development or interaction is the creation of **Codebase Knowledge Graphs**.⁷⁶ These KGs model the structure and semantics of software code itself:

- **Representation:** Functions, classes, modules, variables, API endpoints, and other code artifacts are represented as nodes. Relationships like function calls, class inheritance, module dependencies, data flow, or API usage are represented as edges.⁷⁶
- **Construction:** These KGs can be built by parsing code, analyzing dependencies, and potentially using LLMs to extract higher-level information, such as generating natural language explanations or summaries of functions, classes, or API endpoints, which are then stored as node properties or linked entities in the graph.⁸³
- **Application for Agents:** Agentic systems designed for code generation, analysis, or debugging can query the codebase KG to gain a holistic understanding of the existing software.⁸³ This context helps agents avoid duplicating existing functionality, adhere to established design patterns and architectural constraints, understand data models (e.g., Pydantic models mentioned in⁸³), and generate code that is consistent with the overall project structure and business logic.⁸³

6.3. Agentic Frameworks for Orchestration

Frameworks specifically designed for building agentic applications play a crucial role in managing the complex interactions between different components, including LLMs, KGs, and executable code/tools.⁷

- **Components:** Frameworks like LangChain²⁰, LangGraph²⁰, CrewAI¹⁹, AutoGPT²⁰, and others provide abstractions and tools for defining agent workflows, managing state, integrating tools, and orchestrating the flow of information between the LLM (often acting as the reasoning engine or planner) and other components.

- **KG Interaction:** These frameworks often include modules or integrations for connecting to graph databases (like Neo4j³¹) and querying KGs as part of the agent's perception or reasoning steps. GraphRAG techniques can be implemented within these frameworks.²⁰
- **Tool/Code Invocation:** A key feature is enabling agents to use tools. This can involve defining specific tools (which might wrap code execution, API calls, or database queries) and allowing the LLM to decide which tool to use and with what parameters. **Function calling**, where the LLM generates a structured output (e.g., JSON) specifying a function name and arguments to invoke, is a common mechanism facilitated by these frameworks and supported by some LLMs.⁸²

6.4. Linking Executable Code to KG Elements

Beyond representing code structure *within* a KG, there is interest in directly linking KG nodes or edges to the actual executable code they represent.³¹ This creates a tighter coupling between knowledge and action:

- **Representational Strategies:**
 - Storing code snippets or function identifiers as properties of nodes representing functions or APIs.
 - Using specific relationship types (edges) to explicitly link a conceptual entity (e.g., "User Authentication Service") to its implementation (e.g., a specific microservice endpoint or function).
 - Employing meta-layer KGs that store links (edges) between nodes in different KGs or between KG nodes and elements in external systems like code repositories.⁹²
- **Benefits:** This direct linking allows an agent, upon identifying a relevant KG entity during reasoning, to immediately access the associated executable component or interface details, streamlining the transition from planning to action.

The potential exists to create a unified representational space where both the structured, symbolic knowledge from KGs and the semantic nuances of code can coexist and interact. Knowledge Graph Embeddings (KGEs)⁸⁶ transform KG components into vectors, while analogous techniques exist for code (e.g., using models like CodeBERT or GraphCodeBERT) to generate code embeddings. If these distinct embedding spaces can be effectively aligned, perhaps through joint training objectives or mapping functions, it could unlock novel reasoning capabilities for agents. An agent operating in such a unified vector space could seamlessly transition between symbolic concepts and their code implementations. For instance, it could identify a KG node representing a specific algorithm (e.g., Dijkstra's algorithm) and then, based on vector proximity in the shared space, locate the corresponding function implementation in a codebase, or vice versa. This bridging of symbolic (KG) and sub-symbolic (embedding) representations for both abstract knowledge and concrete code could enable agents to perform more powerful, flexible tasks that require understanding the deep connections between concepts and their realization in software.

The emergence of Codebase Knowledge Graphs⁷⁶ specifically designed to support agentic code generation, signals a potential transformation in software development itself. Current AI-driven code generation often operates on isolated code snippets or files, lacking awareness of the broader system context, which can lead to inconsistencies or redundancy.⁸³ By providing a holistic, structured view of the entire codebase—including dependencies, architectural patterns, and functional relationships—these KGs equip agents with the necessary context to generate code that is not only functionally correct but also well-integrated and consistent with the existing system.⁸³ This suggests a future where AI agents function less like simple code completion tools and more like collaborative "team members." The Codebase KG serves as the shared, structured understanding—a common ground or "map"—of the evolving software system, facilitating a more sophisticated human-agent collaboration model for building complex, maintainable software. Both human developers and AI agents could leverage this shared map to reason about the system, plan modifications, and ensure consistency.

Table 2: Agentic Frameworks Integrating KGs and Tools/Code

Feature	LangChain	LangGraph	CrewAI	Agentic Reasoning Framework	ODA	Zep/Graphiti
Core Orchestration Logic	Chains/Agents based on LLM reasoning; sequential or tool-driven execution.	Builds stateful, multi-agent applications as graphs; nodes are functions/LLMs, edges control flow.	Role-based multi-agent collaboration; defines agents with goals, backstories, tools; orchestrates task execution among agents.	LLM reasoning enhanced by dynamically calling external tool-agents (Web Search, Code, Mind Map KG) based on reasoning needs.	Cyclical Observation-Action-Reflection loop; agent actively observes KG structure to guide reasoning and action.	Focuses on the memory layer; provides KG infrastructure for other agent frameworks to build upon.
KG Integration Method	Integrations for various graph databases (Neo4j, etc.); used via RAG, direct query tools.	Nodes can query KGs; state can hold graph data; facilitates complex GraphRAG workflows.	Agents can be equipped with tools that query KGs (e.g., via LangChain integrations).	Dedicated "Mind Map" agent builds/queries a KG representing the reasoning context. Can potentially query external KGs via tools.	Core mechanism involves "Observation" module querying/analyzing the KG structure.	Provides a dynamic, temporal KG as the primary memory store; accessed via retrieval API combining graph queries, vector search, etc.

Tool/Code Invocation	Defines "Tools" wrapping functions/APIs; LLM selects/calls tools (often via function calling).	Nodes represent tools or LLM calls; edges determine transitions based on state/outputs. Supports function calling.	Agents assigned specific tools (can include code execution, API calls); orchestration layer manages tool use based on roles/tasks.	Reasoning LLM triggers specialized tokens to call external agents (web search, coding, Mind Map) which execute tasks.	"Action" module executes tasks based on reasoning+observation; can include KG exploration, path discovery, or answering. Tool use less explicitly detailed vs. other frameworks.	Primarily a memory system; assumes an overlying agent framework handles tool/code execution, using Zep for context/memory retrieval to inform those actions.
Key Use Cases	General agent development, RAG, chatbots, task automation.	Complex, cyclical agentic workflows, multi-agent systems, long-running processes, GraphRAG.	Collaborative task execution, complex problem decomposition, simulating organizational structures.	Deep research, complex scientific reasoning, knowledge-intensive QA, structured problem-solving.	KG-centric QA, complex reasoning over graph structures, tasks requiring deep KG integration.	Providing long-term, dynamic, context-aware memory for any type of LLM agent (chatbots, assistants, autonomous systems).
Strengths (KG-Code Synergy)	Flexible tool definition, wide range of KG integrations.	Explicit state management, robust handling of cycles/branches, good for complex KG interactions.	Clear role definition and task delegation facilitates integrating specialized KG/code tools.	Tightly integrates reasoning context (as KG) with tool use (code, search); structures the reasoning process itself.	Deepest integration of KG structure into the core reasoning loop via observation.	Strong foundation for KG memory; temporal awareness crucial for dynamic context; low-latency retrieval supports real-time interaction informed by KG + code execution results.
Weaknesses (KG-Code Synergy)	Can become complex to manage state/flow for intricate KG-dependent workflows.	Steeper learning curve than simpler chain/agent models.	Orchestration logic might be less flexible for highly dynamic/unpredictable KG-driven tool sequences compared to LangGraph.	Mind Map KG focuses on reasoning context, may need separate mechanism for broader world/codebase knowledge.	Less focus on defining/orchestrating external tools compared to LangChain/LangGraph. Assumes KG exists.	Is primarily the memory component; requires integration with an execution/orchestration framework like LangChain/LangGraph to fully bridge KG memory and code action.

7. Challenges and Future Directions

Despite the significant potential of integrating KGs into agentic AI systems, several challenges remain, alongside promising avenues for future research and development.

- **Scalability:** As KGs grow to encompass vast amounts of enterprise or web-scale data, ensuring efficient storage, querying, reasoning, and embedding generation becomes a major hurdle.¹² Handling dynamic updates and maintaining performance in real-time agentic loops with potentially billions of nodes and edges requires optimized graph databases, distributed architectures, and scalable algorithms.²⁵
- **Reasoning Depth and Complexity:** Current KG reasoning within agents often relies on pathfinding or simple rule-based inference. Achieving deeper logical, causal, analogical, or probabilistic reasoning capabilities, especially in combination with LLM inference, remains a significant challenge.¹⁰ Developing hybrid neuro-symbolic approaches that effectively combine the strengths of KGs (explicit structure, formal logic) and neural models (pattern recognition, generalization) is a key research direction.¹⁹
- **Dynamic Updates and Consistency:** Agents operating in dynamic environments continuously update their KG memory. Ensuring the consistency, accuracy, and timeliness of this evolving knowledge base is critical.¹⁵ Mechanisms are needed for conflict resolution, belief revision, and managing the provenance and temporal validity of information added by the agent itself or external sources.⁹⁶ The long-term evolution and preservation of these dynamic KGs also pose challenges.⁹⁶
- **Integration Complexity:** Building robust agentic systems requires seamless integration of diverse components: LLMs, KGs, vector databases, APIs, sensors, actuators, and potentially legacy enterprise systems.¹⁹ Defining standardized interfaces and developing effective orchestration strategies for these heterogeneous components is complex.
- **Explainability and Trust:** While KGs offer more inherent transparency than purely neural models, ensuring that the combined reasoning process of an LLM interacting with a KG remains explainable and trustworthy is crucial, especially for high-stakes decisions.²⁸ Methods are needed to trace and justify agent actions derived from this hybrid reasoning.
- **Ethical Considerations and Governance:** Biases present in the data used to construct KGs or train LLMs can propagate into agent behavior.⁴⁸ The autonomy of agents raises concerns about accountability, control, and potential misuse. Developing robust ethical

guidelines, governance frameworks, and evaluation methods for agentic AI is essential. Fairness, privacy, and robustness are critical considerations.⁹⁵

The challenge of maintaining the integrity and validity of the KG is significantly heightened in agentic systems where the KG serves as dynamic memory, updated autonomously by the agent itself.²¹ In traditional KG pipelines, validation often occurs during curated construction or batch updates. However, when an agent modifies its own KG memory based on real-time perceptions or potentially flawed LLM inferences²¹, there is a substantial risk of introducing errors, inconsistencies, or even propagating hallucinations into the core knowledge base. This corrupted memory can then poison subsequent reasoning cycles, leading to incorrect decisions and actions. Therefore, future advancements must focus not only on enabling dynamic KG updates but also on developing robust, real-time validation mechanisms operating within the agentic loop. Potential approaches could include automated self-consistency checks against ontological constraints, assigning confidence scores to agent-derived facts, leveraging external knowledge sources for verification, or even employing dedicated "validator" agents within a multi-agent system architecture to monitor and correct the shared knowledge graph. Ensuring the veracity of the agent's own evolving world model is paramount for reliable autonomy.

Furthermore, the deep integration of KGs (structured knowledge), LLMs (natural language processing and generative reasoning), and code execution (procedural action) within agentic systems necessitates a shift towards a co-design methodology. Building these sophisticated systems is not merely a matter of connecting pre-existing components; the design choices in one layer profoundly impact the others. The specific schema and ontology used for the KG⁹ influence the types of reasoning that are feasible and efficient.⁷⁷ The choice of LLM impacts its ability to interpret KG information, generate valid queries, or plan effectively based on graph context.⁴ The set of available tools and APIs defines the agent's action space and constrains its capabilities.²⁸ Optimizing the overall agent requires a holistic perspective that considers these interdependencies. For example, designing a KG schema that naturally aligns with the parameters of available APIs³¹, or selecting an LLM architecture specifically trained or fine-tuned for interacting with graph structures³², are co-design challenges. Successfully navigating this complexity requires interdisciplinary teams and expertise spanning knowledge representation, machine learning, and software engineering, fostering a collaborative approach to system architecture and development.⁹⁷

Future directions in this field include:

- **Hybrid AI Systems:** Continued development of architectures that synergistically combine symbolic reasoning (KGs) and sub-symbolic learning (LLMs, embeddings).²⁹
- **Automated KG Construction and Evolution:** Improving methods for automatically building, refining, and updating KGs from diverse data sources, potentially using agents themselves for curation.²¹
- **Real-time Knowledge Updating:** Enhancing the ability of KGs and agent memory systems to reflect changes in the environment with minimal latency.²¹
- **Industry-Specific KGs and Agents:** Development of tailored KGs and agentic solutions optimized for specific domains like healthcare, finance, or manufacturing.²⁹
- **Advanced GraphRAG and Memory:** Innovations in retrieving and utilizing graph-based knowledge for RAG, and more sophisticated agent memory architectures incorporating temporal reasoning, uncertainty management, and hierarchical abstraction.²⁹

8. Conclusion

Knowledge Graphs are rapidly transitioning from static data representations to become dynamic, actionable foundations for the next generation of agentic AI systems. Their unique ability to represent structured knowledge with explicit semantics makes them invaluable for addressing the inherent limitations of purely LLM-based approaches, particularly concerning grounding, consistency, and explainability.

This report has detailed how KGs serve a dual role within agentic architectures. Firstly, they function as sophisticated memory systems, capable of storing and organizing both general semantic knowledge and specific, time-bound episodic experiences. Frameworks like AriGraph, Zep/Graphiti, and ODA demonstrate diverse strategies for leveraging KGs to provide agents with persistent, context-aware, and dynamically updatable memory, moving beyond the constraints of LLM context windows and simple RAG. Secondly, KGs actively enable and guide agent action. By modeling procedures, workflows, and capabilities within the graph structure, KGs provide the basis for planning, reasoning, and decision-making. They serve as a crucial bridge connecting perception and reasoning to concrete actions, often involving interaction with external tools and APIs.

Furthermore, the convergence of graph structures and code representations, exemplified by codebase KGs and agentic orchestration frameworks, highlights a path towards agents that can deeply understand and interact with software systems. Techniques like knowledge graph embeddings and function calling facilitate a tighter integration between declarative knowledge stored in the graph and the procedural logic embedded in code. This integration is fundamental for building agents capable of complex tasks ranging from automated software development to sophisticated process automation.

While significant challenges related to scalability, reasoning depth, dynamic consistency, integration complexity, and trustworthiness remain, the synergistic potential of combining KGs with agentic AI principles is undeniable. The ongoing research into temporal KGs, hybrid neuro-symbolic reasoning, automated KG maintenance, and robust agentic frameworks promises to yield AI systems that are not only more autonomous and capable but also more grounded, adaptable, and explainable. The continued exploration of this intersection represents a critical frontier in the pursuit of truly intelligent artificial agents.