Deloitte.

VU UNIVERSITY AMSTERDAM

Knowledge-enriched Agentic AI Workflows

Intelligent and Responsible Task Execution for Business Growth

Content

Key takeaways

- The integration of AI agents with Knowledge Graphs enhances operational efficiency through the responsible execution of complex tasks. This enables the strategic reallocation of human resources, allowing them to concentrate on high-value activities and fostering innovation.
- By leveraging shared insights through Knowledge Graphs, AI agents and teams can collaborate more effectively and achieve better outcomes.
- Knowledge Graphs serves as a shared framework for storing and managing organizational knowledge, laying the foundation for enhanced data quality and consistency by enforcing standardized data models and validation rules.
- Improved interoperability through Knowledge Graphs enables seamless integration and communication among diverse AI agents and systems.
- Semantic understanding and contextual awareness provided by Knowledge Graphs empower AI agents to deliver more accurate and insightful responses to complex queries.

The Value: Driving operational efficiency with knowledge-enriched agentic AI workflows

In the fast-paced world of AI, businesses are continuously seeking ways to enhance **efficiency and productivity**. While Generative AI, exemplified by Large Language Models (LLMs), offers significant advancements in content creation and customer engagement through personalized interactions, it faces limitations in handling **complex, multi-step tasks on large volumes of connected data**.

These challenges have sparked a renewed interest in foundational AI concepts like knowledge representation and reasoning, which provide complementary insights to generative models. By integrating these concepts via Knowledge Graphs into agentic workflows, where AI agents autonomously handle complex tasks, businesses can **streamline operations, enhance decision-making,** and **improve customer satisfaction**.

For example, in supply chain management, Knowledge Graphs enable AI agents to autonomously monitor inventory levels, predict demand fluctuations, and reorder stock by offering a unified view of knowledge that connects products, suppliers, customers, and more. This ensures smooth operations and minimizes the risk of stockouts or overstocking.

This not only enhances **operational efficiency** but also strategically positions businesses to make more informed, autonomous decisions, driving greater **competitiveness** and **innovation** in an AI-driven world.

Introduction to Agentic AI Workflows

Agentic workflows represent a paradigm shift in how organizations approach business processes and task automation. These workflows leverage AI Agents: models that can **autonomously perform tasks**, **make decisions**, and **learn** over time with only occasional human intervention. This marks a major leap in capability compared to using single LLMs such as ChatGPT, which requires constant interaction with a human who ultimately makes the decision.

This innovative approach boosts productivity by dynamically **adapting processes in real-time,** allowing organizations to respond swiftly to environmental changes and enhancing collaboration among departments for improved decision-making.

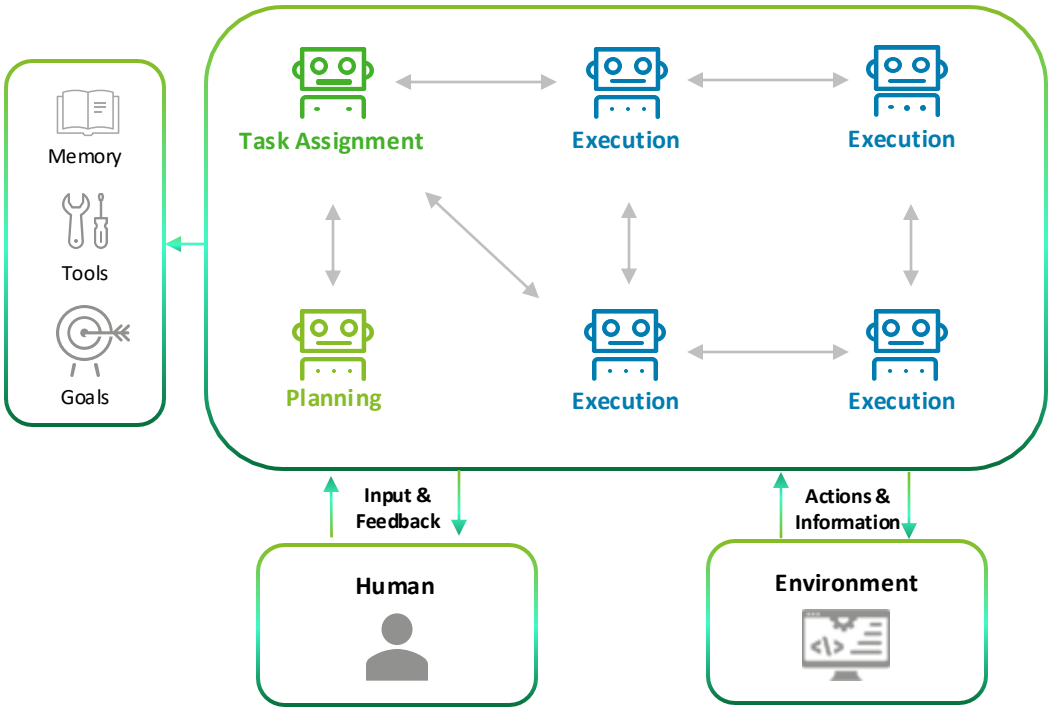

The elements of Agentic AI Workflows

Each AI agent, equipped with memory, goals, and tools, operates in a shared environment to complete complex tasks. They make decisions based on their goals and environmental observations, communicate with each other to share information, and coordinate actions to achieve a common objective. A feedback loop allows them to learn and adapt, enhancing their performance over time.

- The **Planning Agent** organizes the human input and creates a strategic plan that outlines the necessary tasks and their respective priorities.
- The **Task Assignment Agent** receives the task assignments from the Planning Agent and assigns tasks to specific execution agents.
- The **Execution Agents** execute specific tasks as assigned, such as monitoring inventory levels, predicting demand fluctuations, or coordinating logistics.

A new paradigm in **Task Automation**

In the absence of AI agents, businesses typically rely on manual processes or basic automation tools such as Robotic Process Automation (RPA) to handle tasks. These traditional workflows are often rigid, with limited **adaptability**. They require significant human intervention, which can lead to inefficiencies, increased errors, and slower response times. Furthermore, these processes are heavily dependent on the expertise and knowledge locked in the brains of a few experts, limiting the organization's ability to effectively scale operations. As a result, the ability to process **large volumes of data** and make **real-time decisions** is constrained, and organizations may struggle to keep up with the pace of change, impacting their competitive edge.

Even a single AI agent offers notable advantages, primarily through its ability to autonomously perform repetitive and time-consuming tasks. This autonomy **frees human resources**, allowing them to focus on strategic activities that require human creativity and decision-making. Furthermore, the increase in productivity and speed enables businesses to concentrate on **higher-level decision-making** and fostering **innovation**. The contextual understanding of AI agents is improved by using agents powered by LLMs.

However, it should be noted that their understanding is limited, as it is constrained by the data and algorithms they are trained on. Consequently, they may struggle to fully capture human nuances and could incorporate data irrelevant to the issue at hand, driven by their tendency to generate and hallucinate, in case of generative AI. Fortunately, these limitations can be navigated, which is discussed on the next page.

- According to Salesforce, more than half of the service professionals reported substantial improvements in customer interactions using AI, as response times could be reduced, while customer satisfaction was boosted simultaneously.
- Besides customer service, productivity in Software Engineering firms could also see significant increases. McKinsey & Company states that AI can automate up to 30% of work hours, enabling developers to focus on more complex challenges and foster innovation. Additionally, the World Economic Forum's The Future of Jobs Report 2025 anticipates that AI and automation will transform around 22% of jobs by 2030.
- According to our recent Deloitte survey on The State of Generative AI in the Enterprise, the most compelling areas of interest are agentic AI and multiagent systems for task automation, with over 26% of leaders reporting that their organizations are exploring it to a large or very large extent.

Navigating Limitations: How Knowledge Graphs enhance Agentic AI Workflows

Current agentic workflows face four primary challenges: (meta)data fragmentation, lack of interoperability, the difficulty of maintaining consistent data models, and a limited contextual understanding. While AI agents excel at executing specific tasks, they often struggle to navigate across **siloed data sources** and **lack the deeper contextual awareness** needed for collaboration and complex reasoning tasks.

For instance, when a customer asks a customer service chatbot why their order has not shipped, it struggles to synthesize information from separate databases for payments and shipping, which are not directly connected. Furthermore, the chatbot may not grasp complex inquiries, such as how returning a product affects loyalty points, due to its limited contextual awareness. This results in generic or inadequate responses.

Lastly, the need for specific APIs and configurations for AI agents to integrate and communicate effectively introduces another difficulty. This **lack of interoperability** can hinder seamless collaboration across diverse systems and agents.

Knowledge Graphs as a solution

Knowledge Graphs play a transformative role by providing structured representation of data that captures **relationships** and **context**. This allows AI agents to navigate data intuitively, similar to human **reasoning**, enabling sophisticated analyses and pattern recognition.

Knowledge Graphs transform agentic workflows by playing four key roles.

The Role of Knowledge Graphs

1 Serve as a shared framework
providing a standardized structure for storing and managing organizational knowledge. This enforces consistent data models and validation rules, laying the foundation for enhancing data quality and reliability across the system.

2 Enable interoperability
allowing AI agents to seamlessly integrate and communicate with each other and various tools/systems. This ensures diverse agents can work together effectively, regardless of their individual configurations or specializations.

3 Enhance knowledge sharing
creating a shared knowledge repository where different agents and teams from different data domains can collaborate effectively. This enables agents to leverage each other's insights and learnings, leading to improved workflow outcomes and continuous system enhancement.

4 Provide semantic understanding and contextual awareness
enabling AI agents to interpret complex queries and scenarios with deeper insights. This leads to more accurate and relevant responses, particularly in situations requiring nuanced understanding. Also serving as a dynamic long-term memory, which agents can continuously update and expand, thereby enhancing contextual awareness among each other.

The need for strategic investment

While the benefits are substantial, implementing Knowledge Graphs presents certain challenges. Organizations must invest significant resources in **developing and maintaining** the knowledge and graph structure, ensuring data quality, and managing schema (i.e. semantic model) evolution. Although strategies exist for accelerating Knowledge Graph construction, such as utilizing LLMs, effective implementation relies on **domain expertise** and careful consideration of data domain entities, their relationships, constraints and rules, known as a semantic model or ontology. When these semantic models in the data are frequently updated and expanded, the Knowledge Graph should be updated as well. Maintaining the semantic model and ensuring it remains up-to-date without introducing errors or inconsistencies is a complex task.

Despite these challenges, the empirical evidence strongly supports their value. Data.world discovered that using Knowledge Graphs with LLMs achieved a 3x improvement in the accuracy of responses for business questions, even for complex queries. Microsoft's research on GraphRAG demonstrated cost reductions through a token reduction of 26-97%, while LinkedIn's recent paper showed decrease in customer service resolution times by 28.6%. These results demonstrate that Knowledge Graphs significantly enhance both agent performance and operational efficiency while reducing costs across business applications.

A comparison to the traditional model

	Agentic Workflows	Agentic Workflows with Knowledge Graphs
Reliability	Reliability linked to model training and vector database used for RAG	KGs offer a standard framework for enhanced data quality and support predictable reasoning for consistent results
Cost and Scalability	High token count due to multi-LLM architecture	Additional graph context reduces token usage, and some reasoning tasks can be offloaded to a graph database for higher efficiency
Inter-operability	Limited by the need for specific APIs and configurations for integration	KGs enable seamless integration and communication among AI agents and other tools, enhancing interoperability
Understanding capabilities	Limited understanding due to reliance on pre-trained models without contextual knowledge	The semantic richness of KGs allows for more accurate and relevant responses to complex queries

Proposed **Architecture** for Knowledge-enriched Agentic AI Workflows

Instead of using the traditional architecture as represented on page 5, the Agentic AI Workflow is enriched by a Knowledge Graph. In its simplified form, this suggests a single central knowledge hub; however, each agent can possess its own Knowledge Graph that is aligned with the central one.

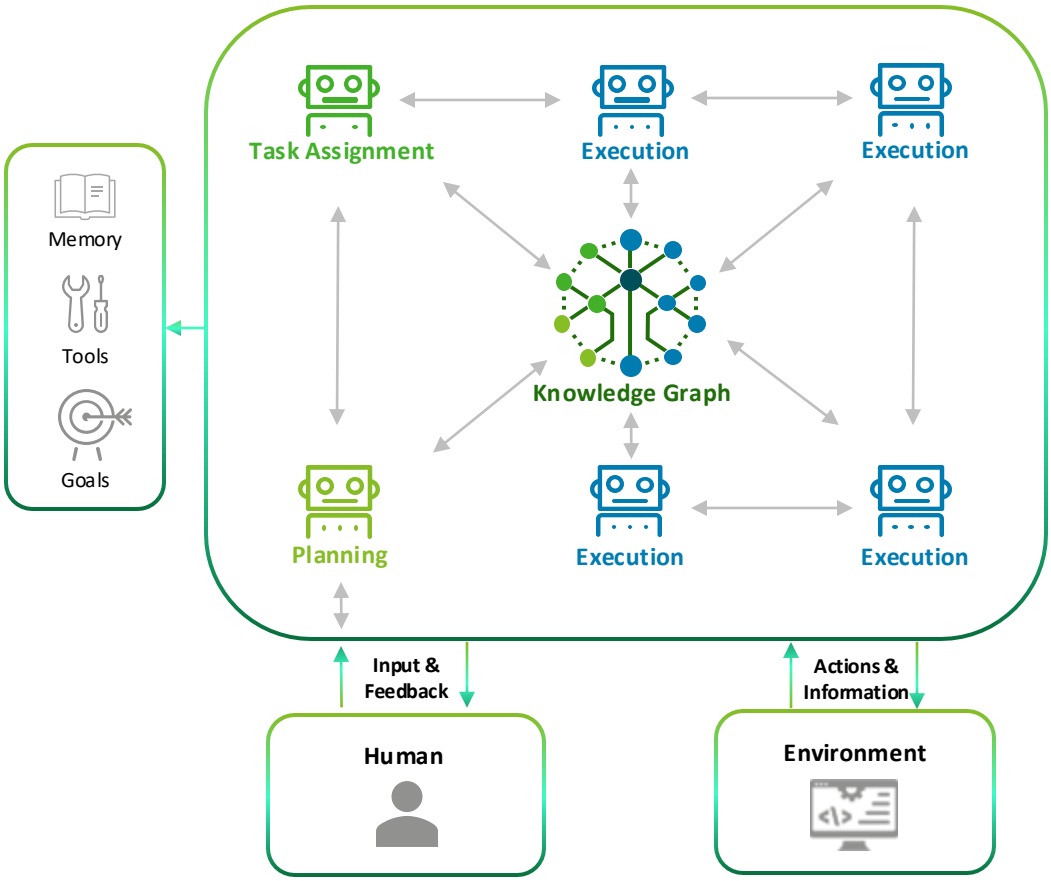

The elements of Knowledge-enriched Agentic AI Workflows

- The **Knowledge Graph** is leveraged by the AI agents as a metadata hub, as well as a repository of domain knowledge. By providing a standardized structure for storing and managing this knowledge, it interoperates between data models and provides validation rules, thereby significantly improving overall data quality and reliability across the system. This foundational aspect enables superior interoperability, allowing AI agents to seamlessly integrate and communicate with one another as well as with various tools and systems.
- The **Planning Agent** organizes the human input and creates a strategic plan that outlines the necessary tasks and their respective priorities.
- The **Task Assignment Agent** receives the task assignments from the Planning Agent and assigns tasks to specific agents based on metadata stored in the Knowledge Graph.
- The **Execution Agents** execute specific tasks as assigned, utilizing information from the Knowledge Graph to ensure accuracy and relevance in their operations.

“ Knowledge Graphs are the essential framework that keeps agentic AI workflows on track, enabling responsible task execution and enhancing decision support, thereby improving result quality. ”

Integrating into operations: Strategies for success

The integration of results from agentic AI workflows back into business processes is vital to maximizing their impact. Organizations should focus on:

1 **Feedback Loops**
Establish robust feedback mechanisms for AI agents to learn from their actions and adjust their strategies accordingly. This initiative is typically driven by **Operations**, supported by **Strategy and Innovation** to align with organizational goals.

2 **Collaboration with Human Workers**
Foster collaboration between AI agents and human workers. AI agents can handle routine tasks, while humans can focus on areas requiring creativity, judgement, and emotional intelligence. This is facilitated by **Operations and Human Resources**, ensuring effective training and integration.

3 **Continuous Monitoring and Accountability**
Implement ongoing monitoring processes to ensure accountability and instituting checks to prevent unintended consequences from autonomous systems. This is led by **Risk Management** with input from **Legal and Compliance** to ensure adherence to relevant regulation ethical standards.

4 **Training and Change Management**
Provide training to employees on how to work with AI agents and integrate them into their daily operations. This is driven by **Human Resources** and supported by **Corporate Communications** to promote understanding and acceptance.

5 **Knowledge Sharing and Interoperability**
Promote knowledge sharing across domains to enhance AI outputs, driven by **IT**, with contributions from **Business Data Domain Owners** to leverage expertise and insights that enhance AI output.

Case study: Boosting trusted ESG reporting with high-quality data

In an era where ESG (Environmental, Social, and Governance) considerations are essential, organizations face increasing pressure to comply with regulations such as EU's CSRD. To meet these requirements and disclose their commitment to ESG principles, companies must collect data from a broad range of sources and areas, including carbon emissions, worked hours, wages, and supply chain metrics. This drives the need for accurate, complete, and fine-grained data for material topics based on their Double Materiality Assessment. For companies that have already completed their assessments, ongoing reviews and adjustments are essential to stay aligned with evolving regulations and stakeholder expectations.

This necessitates a streamlined process for answering key questions, such as:

- What data is required?
- Is the required data available?
- From which sources should the data be sourced?
- What is the data quality?

To achieve this, organizations can automate ESG reporting by utilizing agentic AI with a Knowledge Graph (KG) as the backbone. Different AI agents, each with a specific toolset (e.g., LLM or GraphRAG) tailored to the step's needs, guide the process through six automated steps:

1. **Double Materiality Assessment**: AI agents review ESRS standards and external data like annual reports, industry analysis and competitor information to identify material topics where the company both impacts and is affected by sustainability. A prioritization agent produces a shortlist of material topics, ranked by scale, scope, irremediable character, and likelihood.
2. **CSRD Gap Analysis**: An AI agent compares the identified topics to CSRD requirements, consulting the KG to pinpoint missing or insufficient disclosures regarding data availability and granularity.
3. **Metadata & Data Attributes Collection**: Data from structured and unstructured sources is extracted for each material topic, such as scope 1-3 GHG emissions and human rights. The AI agent uses the KG to interpret formats and label attributes with metadata (e.g., description, units, frequency, data owners).
4. **Data Quality Assessment**: Guided by KG rules, a data quality agent checks data integrity and flags outliers like high consumption or anomalies like incomplete records.
5. **Remediation**: If data quality checks are not met, a remediation agent develops a plan, assigns ownership of the data quality issue for resolution, and validates the corrections once implemented.
6. **Data Consolidation & Reporting**: A reporting agent compiles the validated data into ESG reports, which encompass figures, text, and contextual information.

This process involving the agents and their interactions is included on the next page.

Case study: Boosting trusted ESG reporting with high-quality data

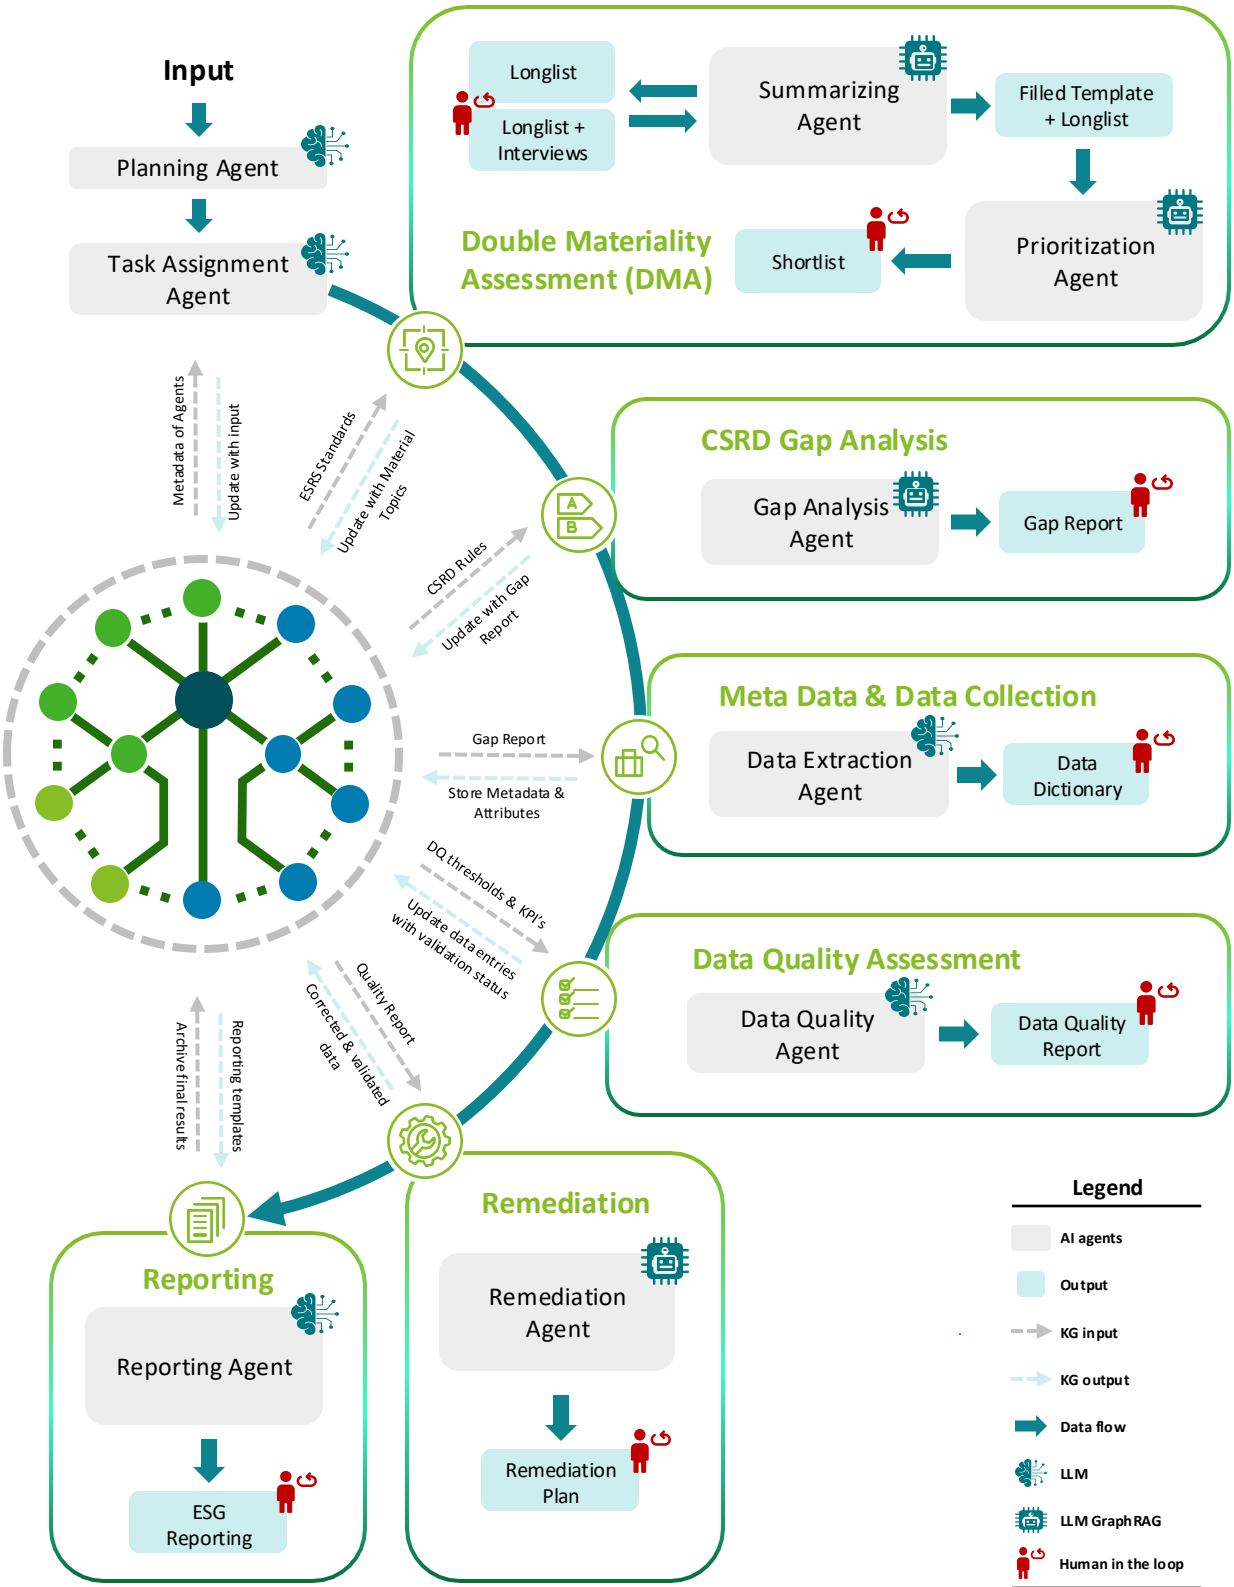

Pathway to Implementation: Where and how to start?

For organizations beginning their journey towards implementing knowledge-enriched agentic AI workflows, several key questions can guide them in determining where and how to start. One of the foremost questions to consider is: *Which existing processes are **ideal candidates for automation**, and how can AI agents add value to these processes?*

This inquiry helps identifying specific areas within the organization where automation can yield significant improvements in efficiency and effectiveness. To facilitate this exploration, the impact-feasibility matrix can be utilized as a prioritization exercise. This matrix allows teams to evaluate potential processes based on two key dimensions: the expected impact of automation on organizational performance, and the feasibility of implementing AI solutions in those processes.

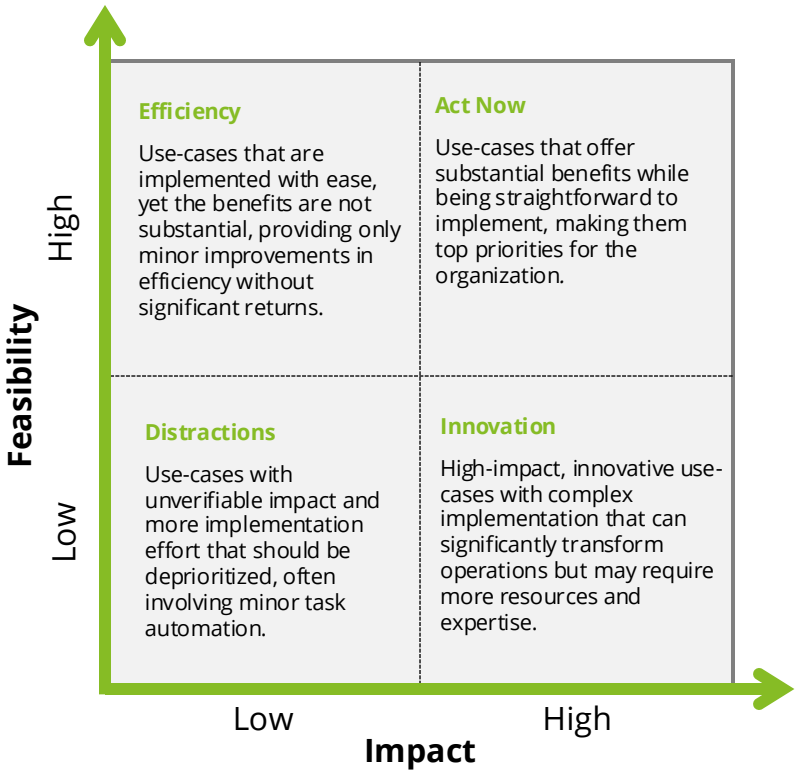

Following the identification of ideal candidates for automation, addressing the follow-up questions below can help organizations ensure they are well-positioned for success in their journey, by gaining clarity on the requirements and challenges they may face.

- *What **types of data** are necessary to support the workflows, and how will the **quality**, **availability**, and **granularity** of this data be ensured?*
- *How will the Knowledge Graph be **designed** and **maintained** to effectively capture context and semantics?*
- *What **resources** and **skills** are required within the team to successfully develop, implement, and maintain the agentic workflows and Knowledge Graphs?*
- *How will the effectiveness of the workflows be measured, and what key performance indicators (**KPIs**) will be used to track progress and success?*

From Initial Assessment to Scaling Solutions

This implementation roadmap identifies the key phases and activities essential for the effective integration and optimization of knowledge-enriched agentic AI workflows within an organization. By following this structured approach, all stakeholders can ensure alignment with business objectives, while establishing a clear pathway from initial assessment to scaling and expansion.

The roadmap distinguishes five essential phases: Assessment & planning, Design & development, Deployment & integration, Monitoring & optimization, and Scaling & expansion. Each phase encompasses specific actions and milestones that contribute to the overall success of the project, ensuring that the solution not only meets the immediate needs of the organization but also supports long-term growth and adaptability in a rapidly changing business landscape. By following this roadmap, organizations can effectively leverage advanced technologies, enhance operational efficiency, and drive sustainable value creation.

1. **Assessment & planning**
   - **Identify Business Objectives:** Define goals such as efficiency, customer satisfaction, and cost reduction.
   - **Stakeholder Engagement:** Collaborate with stakeholders for input and alignment on scope and requirements.
   - **Current State Analysis:** Assess existing workflows and technology to find gaps and opportunities.

2. **Design & development**
   - **Architecture Design:** Create an architecture that integrates LLMs with a Knowledge Graph, ensuring scalability.
   - **Data Strategy Development:** Develop a strategy for data sourcing, cleaning, and governance.
   - **Prototype Development:** Build a prototype to test functionalities and gather feedback.

3. **Deployment & integration**
   - **Pilot Testing:** Validate workflows through pilot tests and refine based on feedback.
   - **Integration with Business Processes:** Ensure compatibility with existing systems for smooth transitions.
   - **Training and Support:** Provide training and establish support systems for users.

4. **Monitoring & optimization**
   - **Performance Monitoring:** Track AI performance using metrics for evaluation.
   - **Continuous Improvement:** Regularly review and optimize the system with new data.
   - **Feedback Loops:** Gather continuous feedback for iterative improvement.

5. **Scaling & expansion**
   - **Scale Operations:** Plan for scaling successful implementations across the organization.
   - **Explore New Applications:** Identify new use cases for agentic AI workflows in various departments.
   - **Maintain Alignment with Goals:** Ensure the AI strategy aligns with evolving business goals.

Conclusion: Integrating Knowledge Graphs boosts process automation

In today's fast-paced business environment, operational efficiency and innovation are essential. This whitepaper introduces a transformative approach to task automation through the integration of **Knowledge Graphs** into **agentic AI workflows**. This integration empowers AI agents to autonomously manage complex tasks while ensuring a robust **contextual understanding** of the knowledge at hand.

Knowledge-enriched agentic AI workflows represent a significant advancement in **business process automation**. By effectively tackling common challenges such as **(meta)data fragmentation**, **lack of interoperability**, the difficulty of maintaining consistent data models, and **limited contextual understanding**, Knowledge Graphs enhance the efficiency of AI agents. Empirical research demonstrates substantial improvements in response accuracy, cost efficiency, and operational effectiveness, highlighting the transformative potential of these technologies.

Forward-thinking organizations that aspire to remain at the forefront of AI innovations can greatly benefit from integrating knowledge-enriched agentic AI workflows into their operations. This technology not only improves **operational efficiency** and **decision-making** capabilities but also fosters a culture of continuous **innovation**. By exploring and adopting these advanced solutions, organizations position themselves to cultivate a more intelligent and agile future, making them better equipped to meet the challenges of an evolving business landscape.

Initiate your journey towards establishing a more capable and innovative organization today…

Get in touch

Dr. Ali Khalili
Knowledge-enriched AI lead, Responsible Data & Analytics
Deloitte Strategy Risk and Transactions Advisory, the Netherlands
alkhalili@deloitte.nl

Aïcha Mohazzab
Knowledge-enriched AI manager, Responsible Data & Analytics
Deloitte Strategy Risk and Transactions Advisory, the Netherlands
amohazzab@deloitte.nl

Marko van Zwam
Partner, Responsible Data & Analytics
Deloitte Strategy Risk and Transactions Advisory, the Netherlands
mvanzwam@deloitte.nl

Marc Verdonk
Partner, Digital Risk Solutions
Deloitte Strategy Risk and Transactions Advisory, the Netherlands
mverdonk@deloitte.nl

Prof. Frank van Harmelen
Head of Learning & Reasoning Research Group
Vrije Universiteit Amsterdam, the Netherlands
frank.van.harmelen@vu.nl

Contributors to this whitepaper:
Axel Blom, Liselotte Mol, Stan Schonk, Martijn Verspeek, Juliette van Ooijen, Dr. Daniel Daza

Suggestions for further reading:

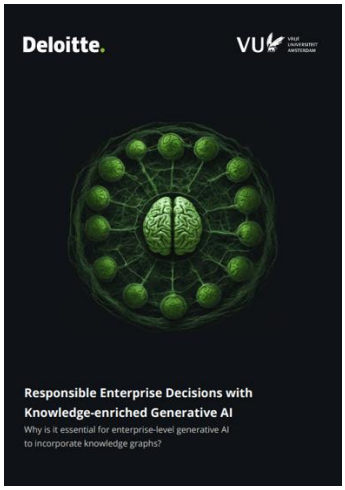

Responsible Enterprise Decisions with Knowledge-enriched Generative AI

Why is it essential for enterprise-level generative AI to incorporate Knowledge Graphs?

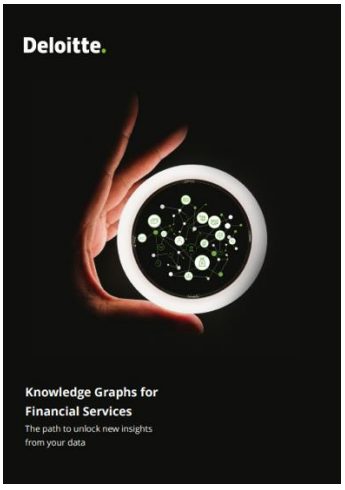

Knowledge Graphs for Financial Services

The path to unlock new insights from your data

Deloitte.