

## **Responsible Enterprise Decisions with Knowledge-enriched Generative AI**

Why is it essential for enterprise-level generative AI to incorporate knowledge graphs?

## Executive Summary

By harnessing the combined power of knowledge graphs and generative AI, enterprises can unlock significant potential for knowledge-driven decision-making, innovation, and operational efficiency. Knowledge graphs, with their structured representation of a domain, enhance the performance of generative AI by providing context, validating outputs, and reducing biases, thereby ensuring alignment with strategic business objectives. Conversely, generative AI enriches knowledge graphs by filling knowledge gaps and predicting future states, thereby increasing the utility, accuracy, and relevance of these graphs. The synergy between knowledge graphs and generative AI serves as a game-changer for businesses, driving transformative impacts across various organizational functions.

### Table of Contents\*

- Exploring the World of Generative AI and its Diverse Use Cases
- Potential Pitfalls: Navigating the Risks of Generative AI
- Harnessing Knowledge Graphs to Mitigate the Risks of Generative AI
- Empowering Organizations with Knowledge-enriched Generative AI
- A Comprehensive Approach to Construct Knowledge-enriched Generative AI
- Conclusion
- Contact Us

\*The content of this paper was co-authored with the assistance of generative AI. Nonetheless, every piece of information and insight has been carefully reviewed and validated by our team of business experts and academics to ensure accuracy and relevance.

“I paint objects as I think them,  
not as I see them...”

**Pablo Picasso**



## Exploring the World of Generative AI and its Diverse Use Cases

### What is Generative AI?

Generative AI is a type of Artificial Intelligence that creates original content across various modalities (such as text, images, audio, code, voice, video) that would have previously taken human skill and expertise to create.

Generative AI provides businesses with significant opportunities to enhance their products and services, automate repetitive tasks, and create new and innovative customer experiences. By enabling the agile creation of realistic simulations and scenarios, generative AI aids in strategic decision-making and risk management, making it a transformative tool in the modern enterprise landscape. Across industries, we have identified several impactful use cases of Generative AI.



### A Selection of High-Impact Use Cases of Generative AI

Generative AI, while most commonly associated with **text** generation, has made significant strides in various other modalities. In the realm of **images**, they have been employed to produce images with varying degrees of realism, variability, and creativity, some of which are indistinguishable from real photographs (e.g. reconstructing an accident scene to assess insurance claims and liability).



In the **audio** domain, they can produce music, voiceovers, and audio (e.g., for call centers and troubleshooting technicians) in natural, conversational, and even colloquial styles with the capacity to rapidly shift among languages, tone, and degrees of complexity. In the **video** domain, they can take user prompts and output videos, with scenes, people, and objects that are entirely fictitious and created by the model (e.g. marketing videos to showcase a new product). Additionally, in **3D** modeling, they can extrapolate and generate data representing intricate designs (e.g. creating virtual renderings in an omniverse environment). In **code** generation, they can produce computer code in a variety of programming languages with the capacity to autonomously summarize, document, and annotate the code for human developers (e.g. for autonomously maintaining code across different platforms).

Below, we spotlight 60 standout applications of Generative AI, categorized across six key industries:



## Potential Pitfalls: Navigating the Risks of Generative AI

### How does Generative AI work?

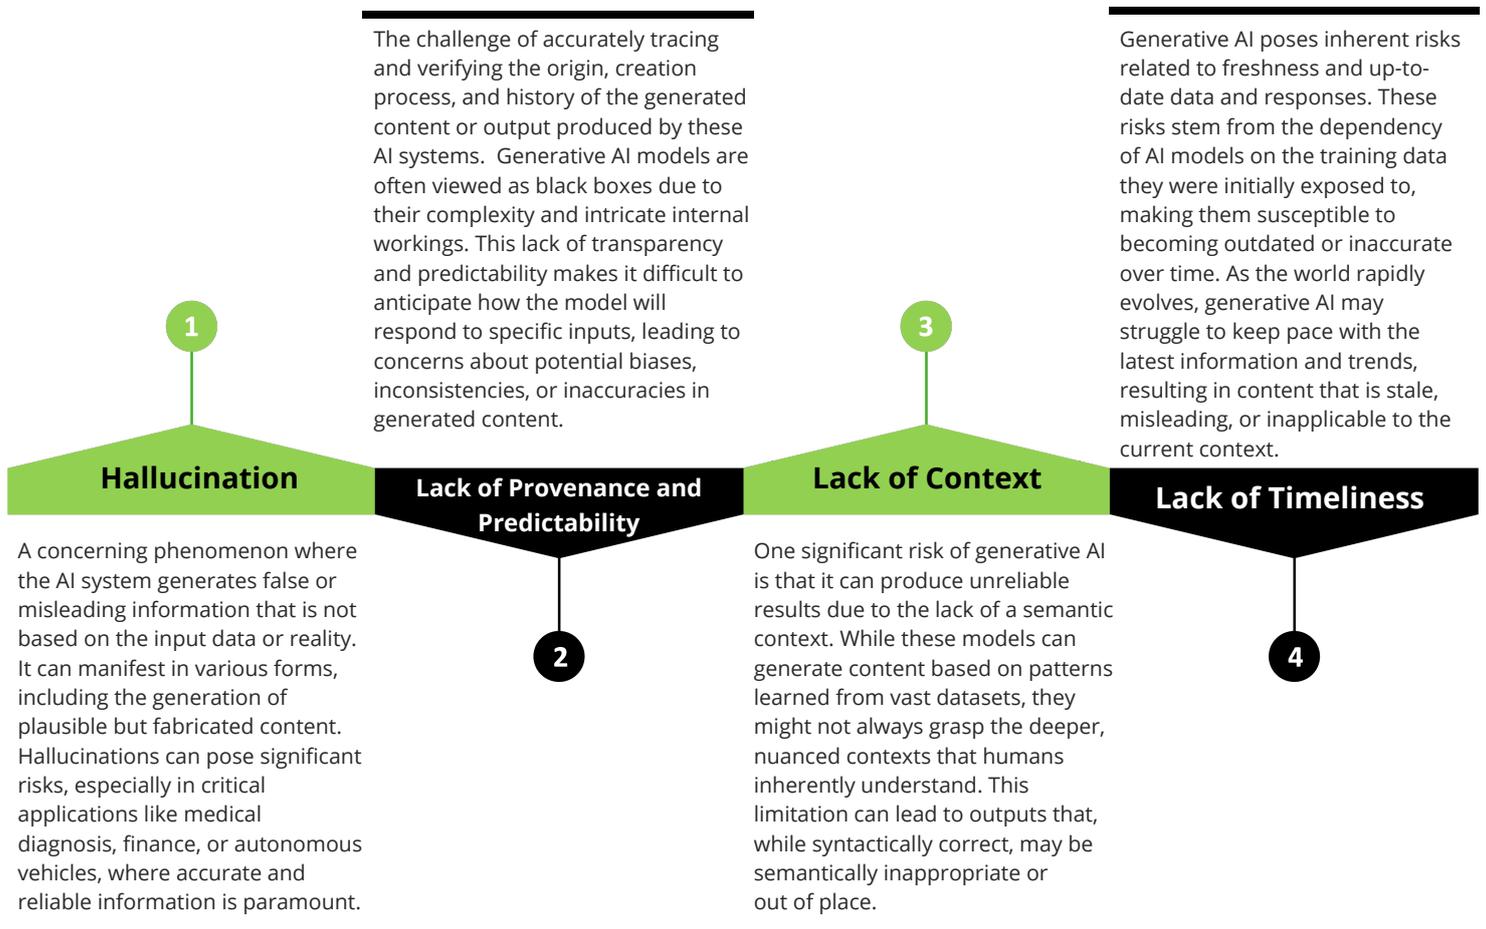
Many of today's prominent Generative AI models operate using three basic stages: *masking*, *training* and *prediction*. Imagine you have a puzzle with a picture of a sentence. Some pieces of this puzzle, which represent words, are missing. Generative AI training by masking sentences is like teaching a machine to predict and place the right puzzle pieces (words) in the correct spots, based on the surrounding pieces it can see. So, in simple steps:

1. **Masking:** We hide some words in a sentence (like removing puzzle pieces).
2. **Training:** The AI tries to guess the missing words based on the context of the surrounding words.
3. **Prediction:** Once trained, when given a new masked sentence, the AI can predict the missing words.



### Technical Risks of Current Generative AI Models

After understanding how generative AI works, it is important to note its potential risks, stemming from the probabilistic nature of these models.

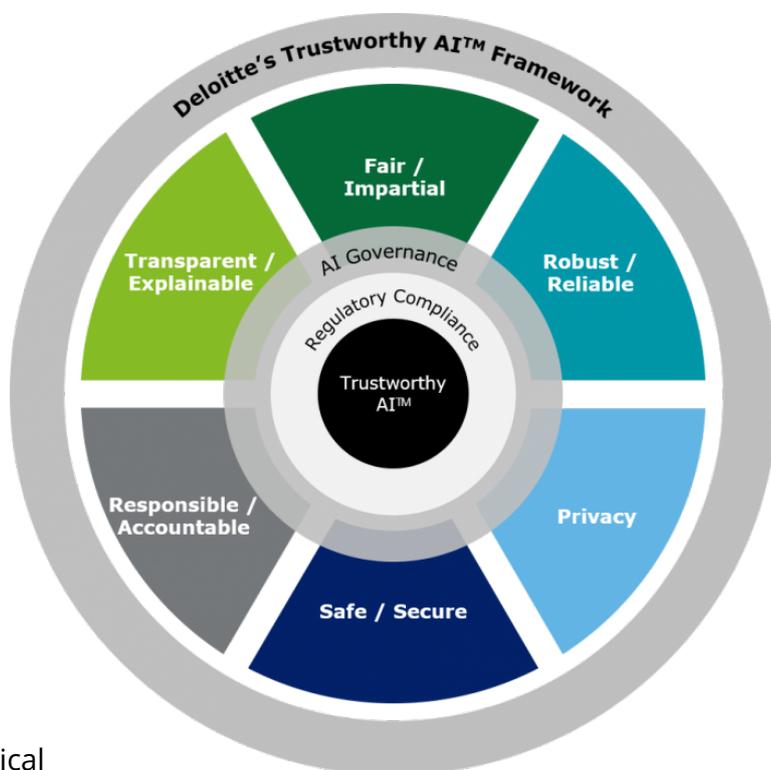


Furthermore, when it comes to Large Language Models (LLMs), the “**Language dependency**” also becomes a technical risk for text-based Generative AI. This relates to challenges in performance across multiple languages. These AI models often excel in generating coherent and accurate content in languages similar to those present in their training data. However, their performance diminishes significantly when faced with languages that differ substantially in syntax, grammar, or vocabulary from the training set. This limitation can lead to inaccurate translations, incoherent responses, or complete failure to understand and generate content in the non-native language.

## Regulatory Risks of Current Generative AI Models

As generative AI gains traction across various industries, especially within regulated sectors, top-level executives are finding themselves confronted with the challenges and potential regulatory risks associated with responsible implementation and integration of this novel technology in their organizations. To aid in this process, [Deloitte's Trustworthy AI Framework](#) offers a comprehensive roadmap on ensuring ethical and responsible usage of AI technologies, including generative AI in organizations.

In terms of regulatory compliance, [the AI Act](#) issued by the European Union has charted an innovative course by incorporating a risk-based strategy. This allows it to outline distinct yet critical provisions for the operation of generative AI systems which harmonize aptly with the dimensions encapsulated within Deloitte's Trustworthy AI framework.



### **Article 13**

#### **Transparency and provision of information to users**

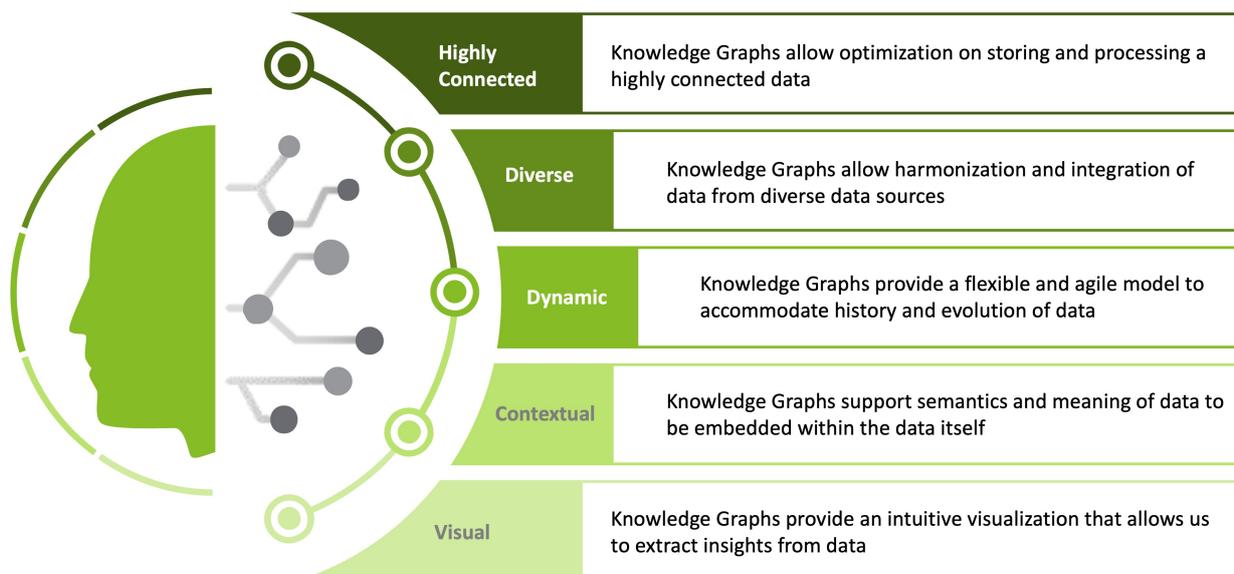
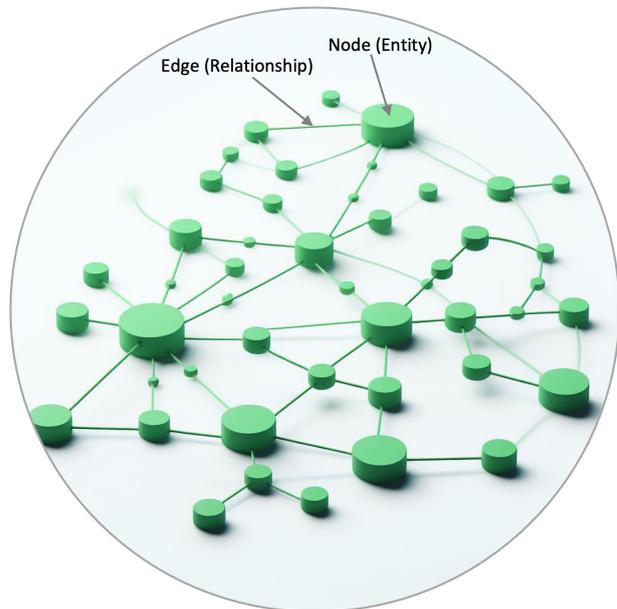
“High-risk AI systems shall be designed and developed in such a way to ensure that their operation is sufficiently transparent to enable users to interpret the system’s output and use it appropriately.”

## Harnessing Knowledge Graphs to Mitigate the Risks of Generative AI

### What is a Knowledge Graph?

A Knowledge Graph (KG) is a means to connect and represent **knowledge** in a domain of interest using a **graph** structure. It is typically built on top of existing data sources of an organization to link data together at web-scale, combining both structured and/or unstructured information. As opposed to the more commonly used relational data models, a graph model is built as a collection of concepts or entities and the relationships between them.

Knowledge graphs form the foundation of modern data analytics by connecting **complex** data and metadata in a meaningful way.



Data lineage, root-cause analysis, 360 view of data, compliance & data quality analysis, tracking changes overtime, impact analysis, data cleansing & migration, recommendation (link prediction), metadata management, semantics analysis & reasoning, entity resolution, semantic search and query answering are few use cases of knowledge graphs within companies.

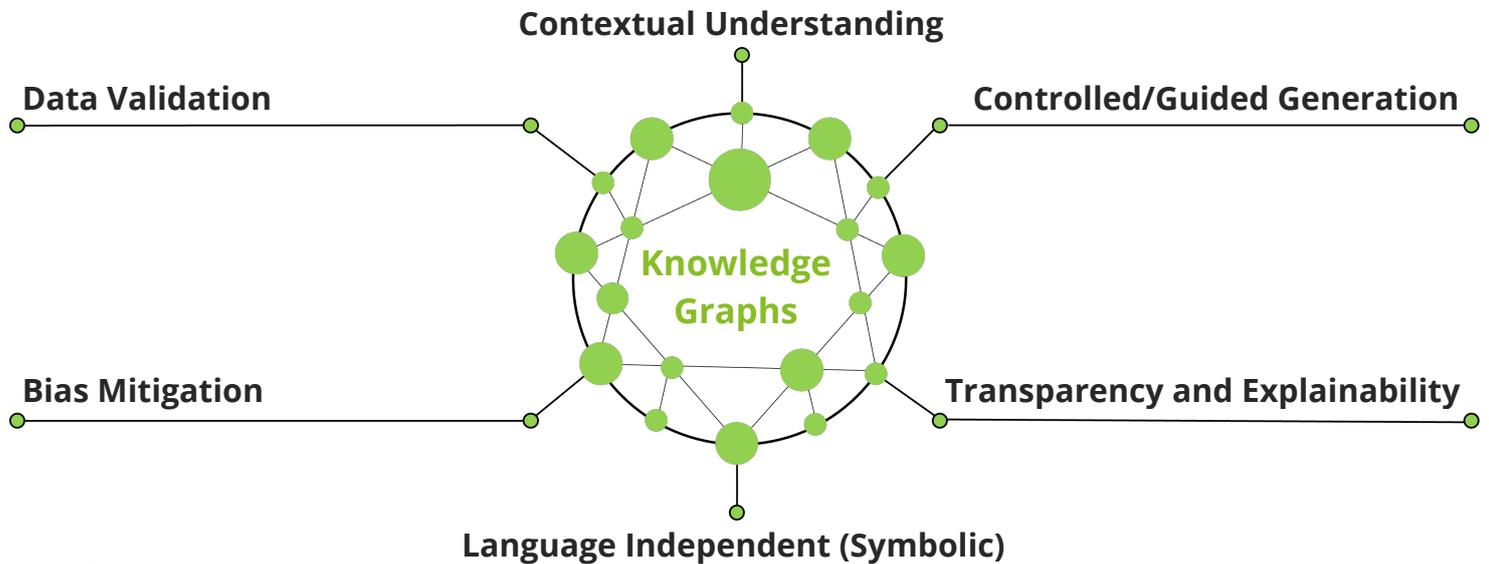
### Logical Inference and Reasoning

Knowledge graphs enable drawing new conclusions and discovering implicit knowledge based on the explicit facts and relationships encoded within the graph. This is typically achieved by designing a so called **Ontology** which establishes a shared vocabulary and a set of rules that dictate how data can be interrelated within the knowledge graph.

This capability for inference and reasoning makes knowledge graphs especially powerful tools for tasks such as building recommendation systems and complex data analysis, where understanding context and relationships is crucial.

## How do Knowledge Graphs Help to Mitigate the Risks of Generative AI?

Knowledge graphs can enhance the reliability, validity, and transparency of generative AI in a business context, helping to mitigate many of its technical and regulatory risks to maximize its benefits.



### Risks\* of Generative AI

**Hallucination**

#### **Controlled/Guided Generation.**

Knowledge graphs can guide the generative process by providing a source of reliable and structured information. This can help control the output of the generative AI, reducing the risk of generating misleading or harmful content that could damage the business's reputation or customer relationships.

**Data Validation.** Knowledge graphs can serve as a reference point to validate the output of generative AI models. If the generated content contradicts the information in the knowledge graph, it can be flagged for review, ensuring business decisions are based on accurate and consistent information.

**Lack of Provenance and Predictability**

**Transparency and Explainability.** Knowledge graphs can make the workings of generative AI models more transparent and explainable by tracking the origin, lineage, and transformation of data within their structured framework. Each entity, relationship, and piece of information is tagged with metadata that captures its source, history, and any alterations made over time. This traceability allows generative AI to verify the authenticity, reliability, and context of the data, ensuring transparency and trustworthiness in the information represented within the graph. This can be crucial for businesses in regulated industries or those that need to justify their decisions to stakeholders.

**Lack of Context**

#### **Contextual Understanding.**

Knowledge graphs provide a structured and meaningful representation of a knowledge domain, enabling generative AI models to better understand the context. This can reduce the risk of hallucination by enforcing only a deterministic and relevant set of outputs, hence, improving customer satisfaction and engagement.

**Bias Mitigation.** Knowledge graphs can help identify and mitigate biases in the training data for generative AI models. By representing knowledge in a structured way, it is easier to spot patterns that may indicate bias, helping businesses to avoid unfair practices and maintain their reputation. To achieve this, the provisioned knowledge requires to be unbiased too.

**Language Dependency**

**Language Independent.** Knowledge graphs, through their logic and symbolic representation of knowledge, transcend linguistic barriers. Rather than relying on specific languages or terminologies (as done by LLMs), they capture and represent knowledge based on universally understood symbols, relationships, and structures. This inherent design abstracts the essence of knowledge from any particular language, making it universally interpretable and ensuring that the core insights and relationships are accessible, regardless of linguistic context.

\* For the 'lack of timeliness', when a new fact emerges, instead of retraining the entire model, the knowledge graph can be updated with this fresh information. This isn't exclusive to graphs and can be also addressed by technologies such as vector databases.

## An Example Scenario: A Technical Assistant Chatbot

**Scenario:** Alice, an employee at a large corporation is facing intermittent internet connectivity issues on her work laptop.



Hi, my internet keeps disconnecting. Any known issues?



Hello! I'm aware of a few devices experiencing connectivity problems today. Let me quickly check if there's a common issue related to your device model.



Chatbot consults the knowledge graph which links device models, reported issues, network hardware, and recent updates or changes in the IT environment.



I see that there was a recent firmware update to the routers used in the office, and devices with your laptop model seem to be particularly affected. Our team is already working on a fix. Meanwhile, you can try connecting to the "Guest" network as a temporary solution.



Okay, I'll try that. By the way, do you know if the conference room devices are affected too? We have a meeting in Room T-23 today.



Chatbot taps into the knowledge graph to cross-reference the conference room's hardware with known issues.



Room T-23 uses a different network setup, and there haven't been any reported issues there today. Your meeting should proceed without any network problems.



Great! Thanks for the help.



You're welcome! If you have any other questions or face further issues, feel free to reach out. Have a productive day!



Alice



Chatbot  
(Based on LLMs)



Knowledge Graph

## Empowering Organizations with Knowledge-enriched Generative AI

Generative AI enriched with knowledge graphs can be positioned in various functions within an organization, providing a powerful tool for data-driven decision making, innovation, and efficiency.

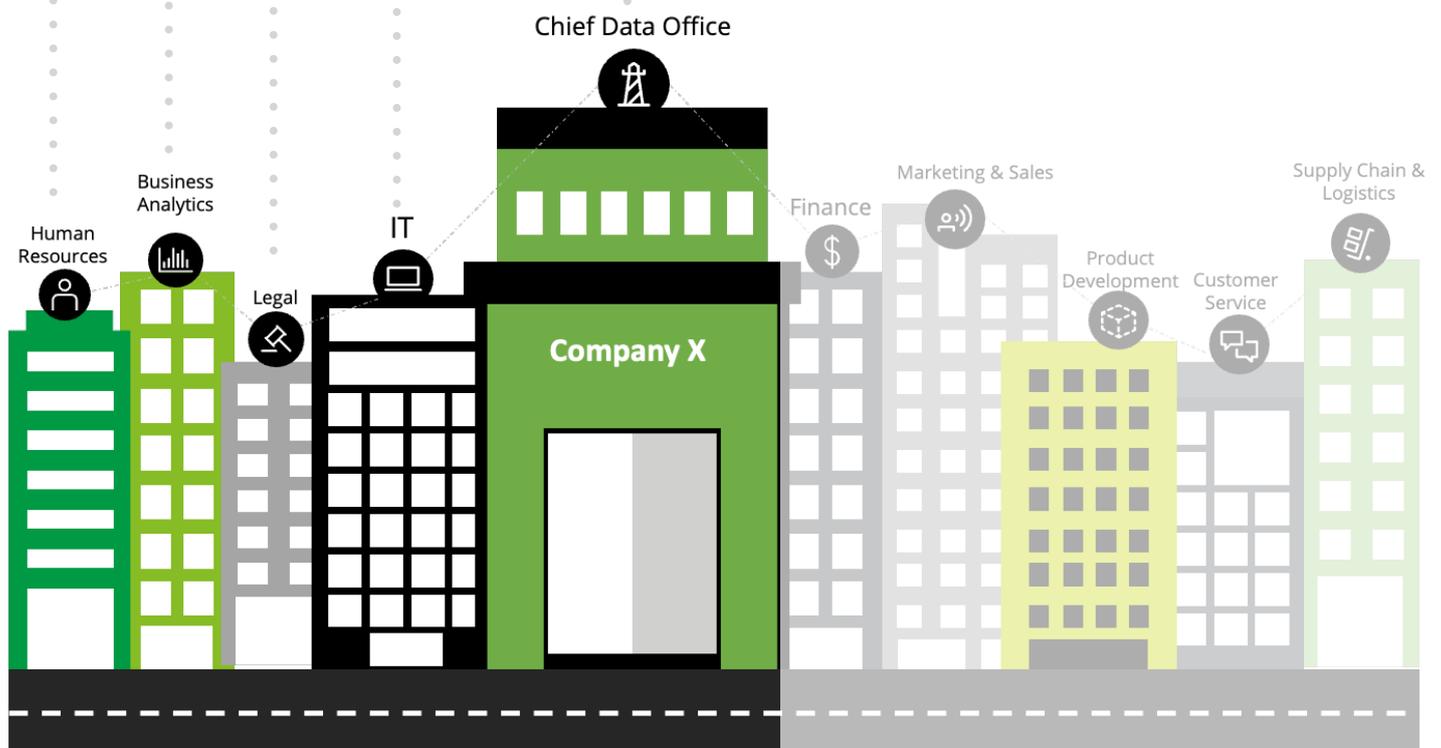
**Human Resources (HR):** In an organization's HR department, the integration of knowledge graphs and generative AI offers a transformative approach. Knowledge graphs provide a structured and interconnected view of employee data, roles, skills, and organizational needs. When paired with generative AI, this setup can streamline recruitment, personalize employee development, optimize compensation packages, and enhance talent mapping. Moreover, it aids in ensuring diversity and predicting HR trends.

**Business Analytics Team:** For the business analytics team, generative AI can utilize structured and semantic understanding of data represented in knowledge graphs for (semi-) automated creation of KPI (Key Performance Indicator) reports and visualizations based on a holistic view of data. This can provide valuable insights for strategic planning, risk assessment, and decision-making.

**Legal:** knowledge-enhanced generative AI systems, equipped with vast amounts of legal knowledge, can assist in drafting contracts, analyzing case law, predicting litigation outcomes, and automating routine legal tasks in a more reliable and transparent way. Without a knowledge graph's guidance, the generative AI might suffer from "hallucination", either generating false cases or misconstruing the true nature of the case (see [this example of misusing LLMs for legal research](#)).

**IT Department:** When it comes to technical support, Knowledge-enriched generative AI can be used to provide (semi-) automated helpdesks to assist employees with common technical issues, providing more comprehensive, context-aware support by leveraging interconnected data, enhancing the user experience and reducing resolution times.

**Chief Data Officer (CDO) Services:** As part of the CDO's services, generative AI can leverage knowledge graphs to enrich data quality and completeness, ensuring that the organization's data assets are accurate, reliable, and comprehensive. This can enhance the organization's data governance and management capabilities.



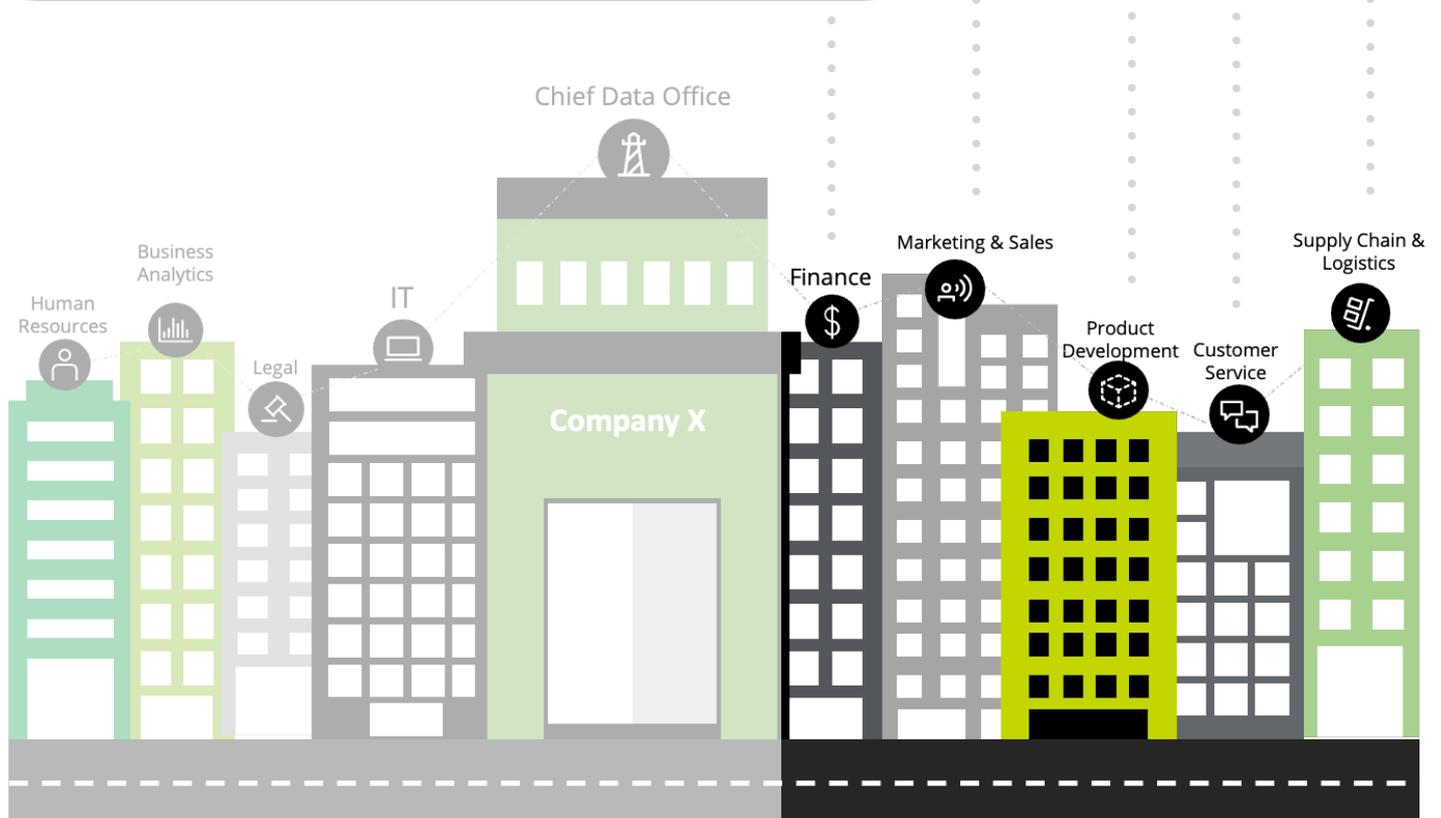
**Supply Chain/Logistics:** When leveraging generative AI combined with knowledge graphs, the interconnectedness and context of the entire supply chain are illuminated, facilitating more accurate demand forecasting, optimized inventory and route planning, and enhanced risk management. This also supports improved supplier relationships, informed procurement decisions, and comprehensive product lifecycle management. The synergy between generative AI's predictive capabilities and knowledge graphs' structured representation ensures the supply chain operates more efficiently, sustainably, and in compliance with relevant regulations, all while enhancing customer communication and service.

**Customer Service:** In customer service, generative AI can leverage knowledge graphs to generate accurate and relevant responses to customer inquiries by understanding the broader context of a customer in connection to products, services, issues, and customer behaviors, hence improving customer satisfaction and efficiency.

**Product Development:** In product development, generative AI can use knowledge graphs to supercharge product development by offering predictive insights rooted in rich contextual understanding. Knowledge graphs consolidate diverse data sources, ensuring a holistic view of market trends, customer feedback, and historical data, while generative AI leverages this integrated knowledge to simulate scenarios, optimize strategies, and forecast outcomes.

**Marketing and Sales:** In marketing and sales, generative AI can use knowledge graphs to generate personalized content or recommendations for individual customers, improving customer engagement and driving sales. The knowledge graph, encompassing data on products, campaigns, and customers, empowers the AI to pinpoint target demographics, craft ads highlighting key features, forecast sales, optimize ad budgets, personalize customer outreach, and discern competitive market strategies.

**Finance:** With the enriched context provided by knowledge graphs, generative AI models can produce better financial forecasts by understanding nuanced dependencies between financial entities. They can also rapidly analyze vast financial data, identifying patterns and anomalies that humans might overlook. This leads to more accurate forecasting, risk reporting, and tailored actionable intelligence, optimizing the finance department's operations. Additionally, for regulatory reporting and risk modeling, knowledge-enhanced generative AI, incorporates both semantic and vertical lineage of data to fix the missing links and provide a consistent end-to-end data flow.



## Navigating Innovation & Knowledge Management Leveraging Knowledge-enriched Generative AI

In the dynamic landscape of organizational strategy, striking a balanced approach between **exploiting existing knowledge** and **exploring new opportunities** becomes paramount. The below matrix delineates various strategic postures based on the use, or non-use, of knowledge graphs (KGs) and generative AI (genAI) to navigate this balance. Traditional strategies, often void of technological leverage, might impede innovation and adaptive capacities. Employing KGs alone can ensure that existing knowledge is methodically exploited, fortifying data-driven decisions, albeit potentially stifling innovative leaps. Conversely, utilizing genAI singularly might spur innovative ideation but can drift into unstructured and potentially misaligned explorations. An integrated application of both KGs and genAI into knowledge-enriched AI strives to synthesize the structured exploitation of existing knowledge with the inventive exploration of new terrains, potentially unlocking new horizons while securing foundational knowledge. For practitioners, tactically blending the structured rigidity of KGs with the explorative fluidity of genAI could pave the way towards a robust strategic stance, capable of both safeguarding existing domain knowledge and pioneering into new strategic frontiers. This integrative approach, while seemingly optimal, demands thorough technological integration and proficient strategic alignment to truly harness the synergistic potentials of KGs and genAI in managing the exploration-exploitation tension effectively.

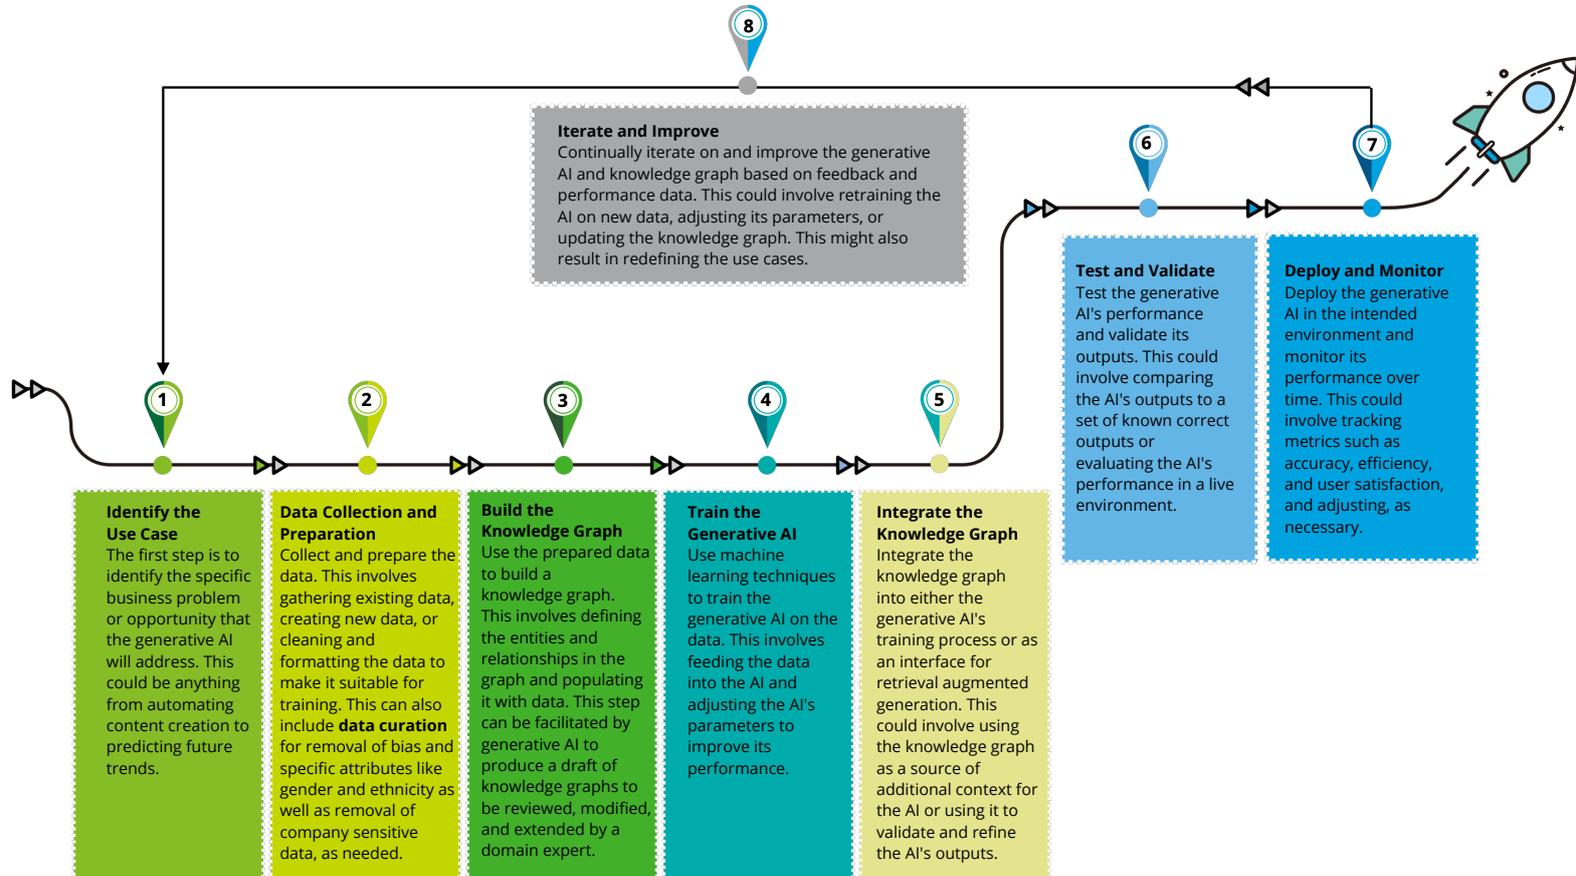
	<p style="text-align: center;"><b>Unstructured Exploration</b></p>	<p style="text-align: center;"><b>Integrated Exploitation &amp; Exploration</b></p>
<b>Using GenAI</b>	<p><b>Advantages</b></p> <ul style="list-style-type: none"> <li>• Facilitates idea generation</li> <li>• Enhances adaptability to new contexts</li> <li>• Enables diverse explorative avenues</li> <li>• Can boost user engagement</li> </ul> <p><b>Risks</b></p> <ul style="list-style-type: none"> <li>• Might result in unstructured, scattered ideation</li> <li>• Outputs may lack verified accuracy</li> <li>• Potential strategic misalignment</li> <li>• Possible resource misallocation with uncertain returns</li> </ul>	<ul style="list-style-type: none"> <li>• Synergizes structured exploitation (via KG) and innovative exploration (via GenAI)</li> <li>• Potentially unveils novel insights and opportunities</li> <li>• Effective synergy requires thorough technological integration and proficient strategic alignment</li> </ul>
	<p style="text-align: center;"><b>Traditional Strategies</b></p>	<p style="text-align: center;"><b>Structured Exploitation</b></p>
<b>Not Using GenAI</b>	<ul style="list-style-type: none"> <li>• Established, perhaps manual methods</li> <li>• Limited innovative capacity</li> <li>• Potential manual knowledge management challenges</li> <li>• May struggle to adapt to innovative environments</li> </ul>	<p><b>Advantages</b></p> <ul style="list-style-type: none"> <li>• Enhanced data structuring and management</li> <li>• Informed, data-driven decision-making</li> <li>• Preserved organizational knowledge</li> <li>• Internally consistent decisions</li> </ul> <p><b>Risks</b></p> <ul style="list-style-type: none"> <li>• Limited innovation scope</li> <li>• Possible over-reliance on existing knowledge</li> <li>• Potential flexibility issues</li> <li>• Risk of data obsolescence due to over-reliance</li> </ul>
	<b>Not Using KGs</b>	<b>Using KGs</b>

Like roots to a tree,  
**knowledge graphs** anchor generative AI  
to depths of understanding...



## A Comprehensive Approach to Construct Knowledge-enriched Generative AI

Implementing a knowledge-enriched generative AI in an enterprise involves several key steps:



The specific steps can vary depending on the specific use case, the type of generative AI being used, and the structure and content of the knowledge graph. It is also important to consider factors such as data privacy and security, AI ethics, and regulatory compliance throughout the implementation process.

### How can Knowledge Graphs Benefit from Generative AI?

Constructing a knowledge graph from the ground up is a time-intensive and effortful task. However, in an enterprise setting, generative AI can play a crucial role in streamlining and accelerating the development of these knowledge graphs.

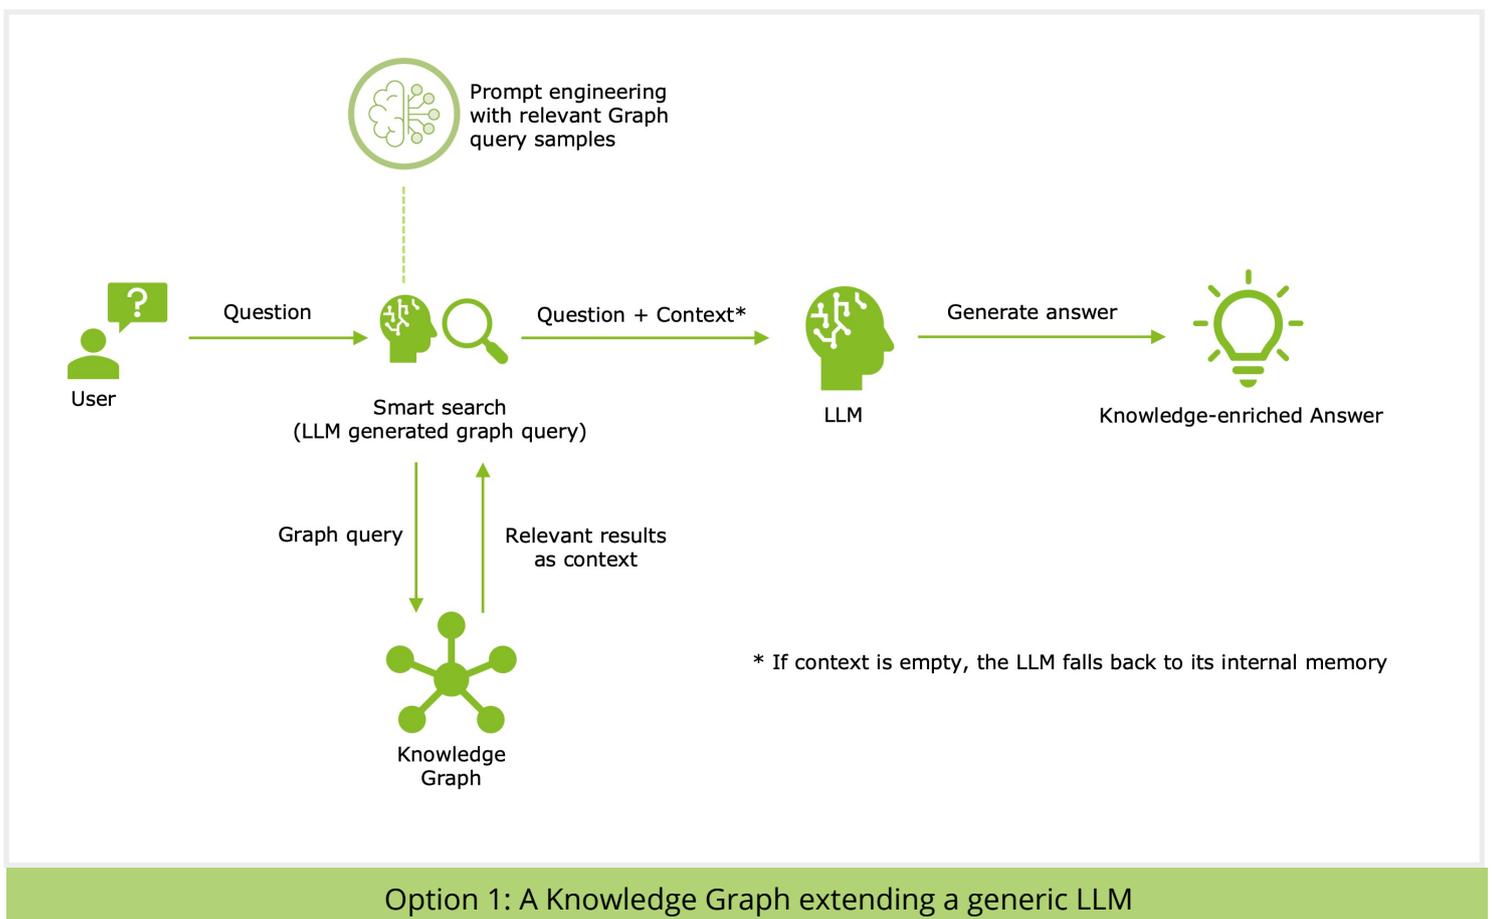
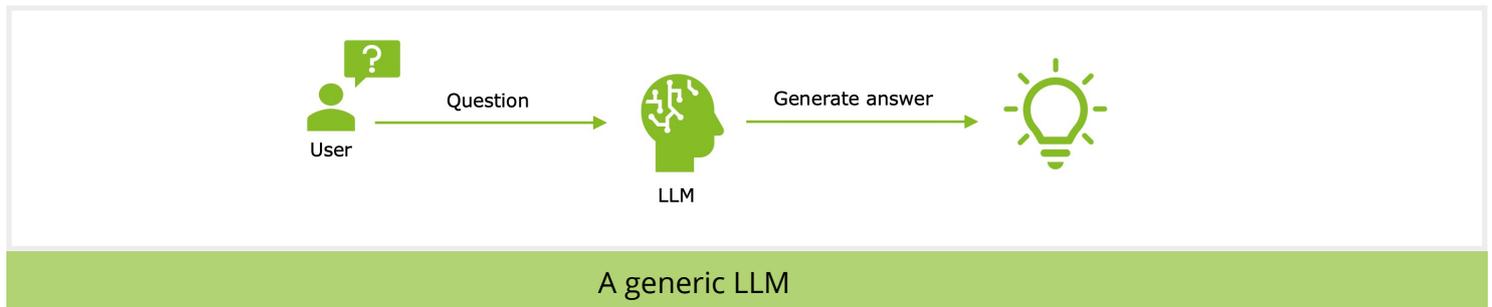
When it comes to **Populating a Knowledge Graph**, generative AI can be used to extract instances of entities and relationships from text, given an ontology or schema for the graph.

With regards to **Data Enrichment for Graph Completion**, generative AI can be used to fill gaps in a knowledge graph by generating new nodes and relationships based on the existing structure and content of the graph. This can help to enrich the knowledge graph and make it more comprehensive. In practice, this task could involve vetting of the added content, to ensure its accuracy.

For **Predictive Analysis**, generative AI can be used to predict future states of the knowledge graph based on historical data. This can be particularly useful for strategic planning and decision-making in a business context.

## How to Realize Knowledge-enriched LLMs: Architectural Designs

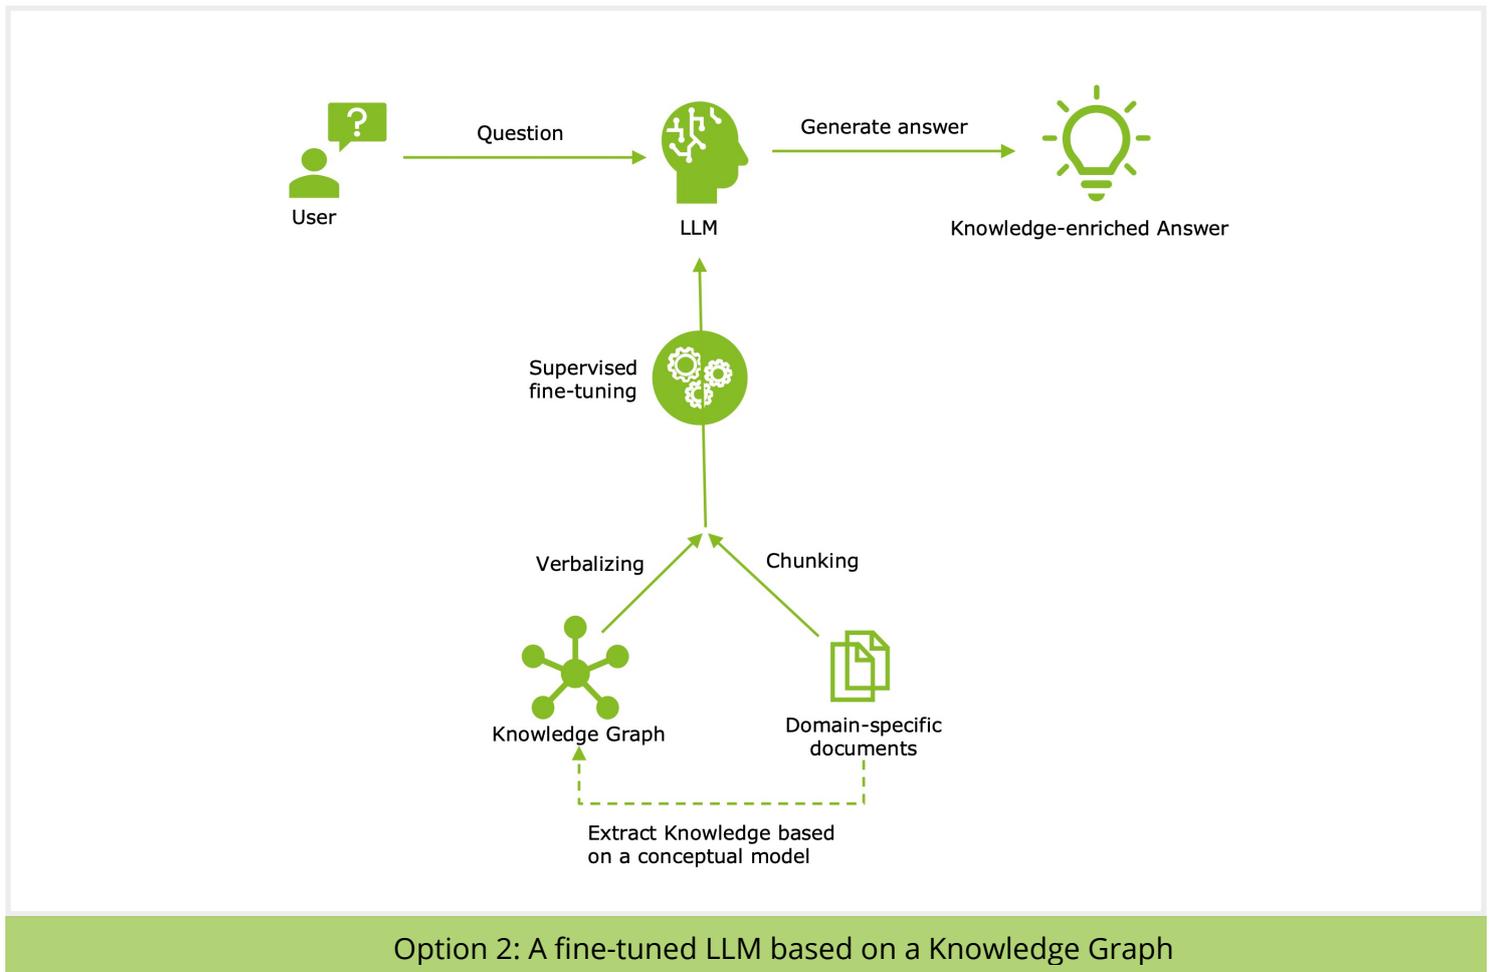
In the following section, we introduce a variety of architectural designs that merge knowledge graphs with LLMs. While these serve as a source of inspiration, tailoring them to your organization's unique needs and specifications is essential.



Option 1 is based on **Retrieval-Augmented Generation (RAG)** to combine the strengths of both knowledge graph-based retrieval and generation-based methods in LLM. Instead of solely generating responses from scratch, RAG first retrieves relevant information from the graph and then uses this retrieved information as a context to guide the generation of detailed and accurate responses. To convert user's natural language queries into graph queries, prompt engineering must be conducted using a set of representative sample queries.

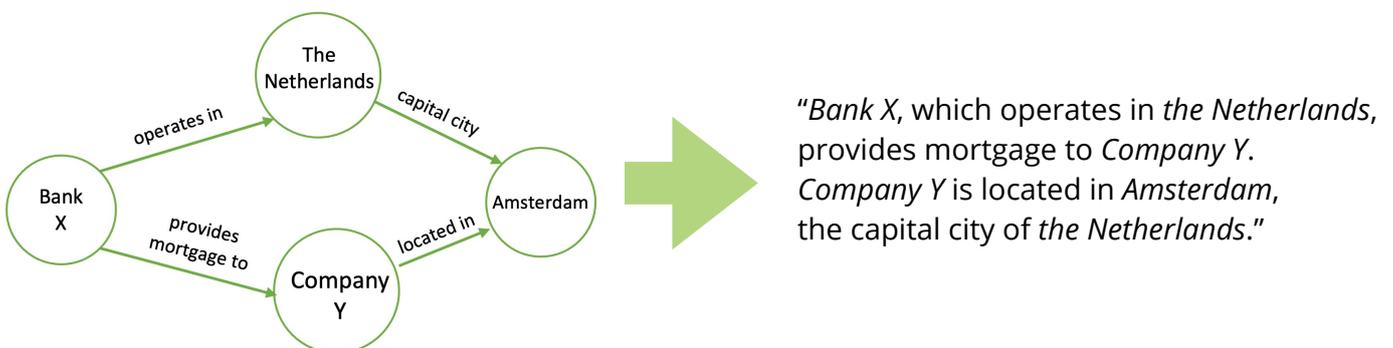
By merging the precision of knowledge graphs with the fluency of generative models, RAG offers a more informed and contextually-aware output, also resolving the lack of timeliness of LLM.

While this method eliminates the need for fine-tuning the LLM, the enriched outputs are limited to the specific domain addressed by the knowledge graph.

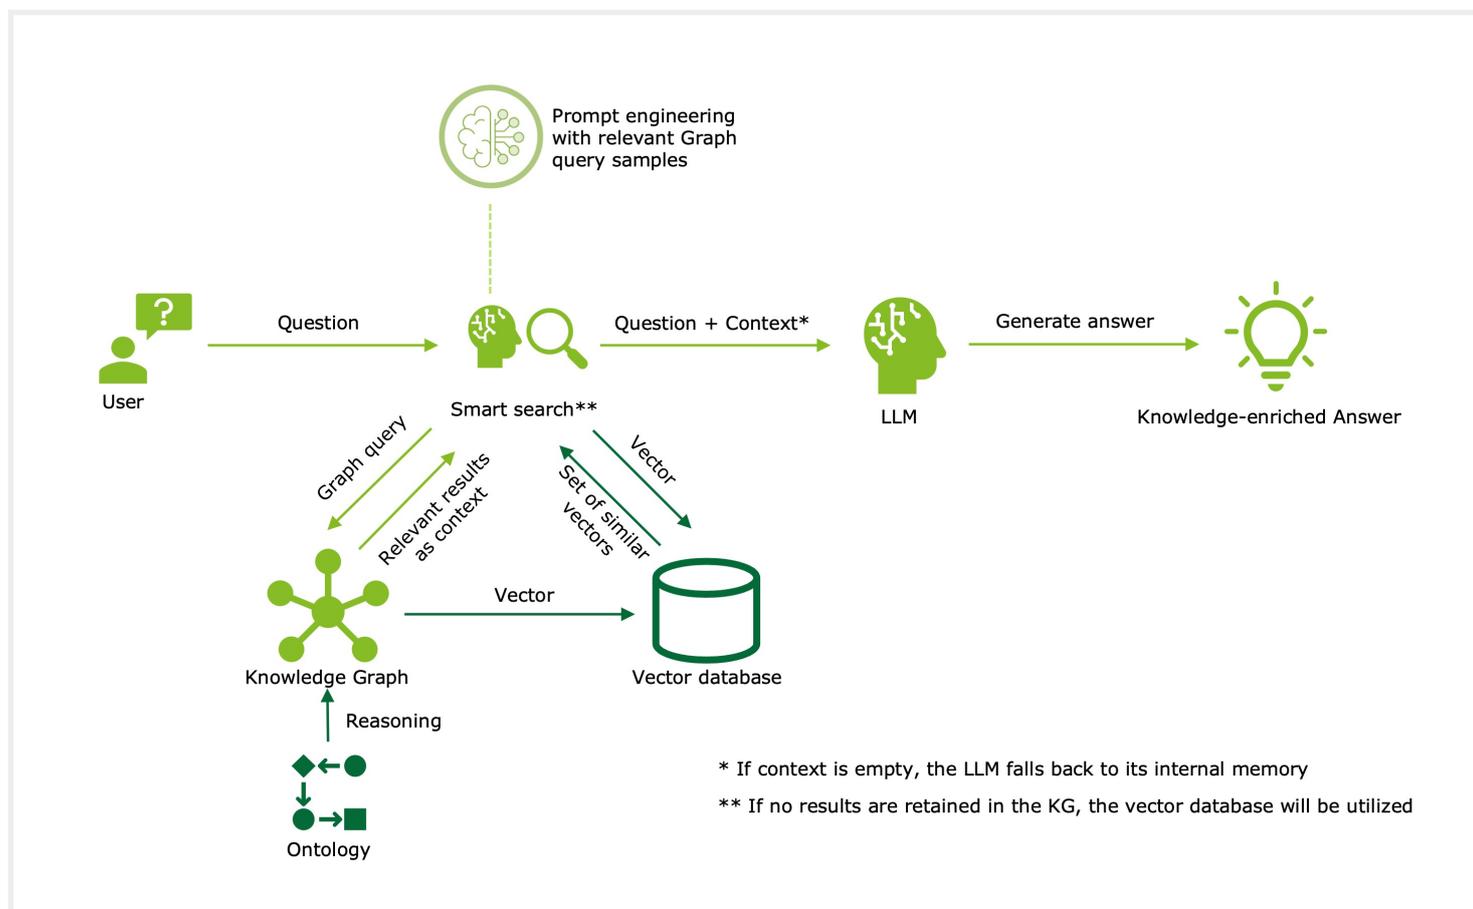


Option 2 is based on **Supervised Fine-tuning** of LLM with additional structured knowledge fed from a knowledge graph. This requires verbalizing the knowledge graph which means converting the structured information within the graph into some natural language descriptions or narratives, digestible by LLM.

For example, the following graph with four entities and four relationships can be verbalized into:



This approach allows to encompass a wider range of topics, however, it comes with the costly fine-tuning of LLM when information is updated.



### Option 3: A Knowledge Graph enhanced with an ontology and/or Vector Database

Option 3 is an **extended** version of Option 1 where an additional **vector database** is added to cover for more topics that can not get directly retrieved from the knowledge graph by running graph queries. While the knowledge graph explicitly depicts relationships and semantic context between entities, the added vector database provides faster similarity search capabilities, allowing for more nuanced comparisons between data points and providing more implicit and relevant results.

Another point of extension here is the capability of **semantic reasoning** to amplify the depth and contextual accuracy of the generated outputs. This requires designing an ontology that represents knowledge in a specific domain, outlining entities, their properties, and the relationships between them which allows the drawing of meaningful conclusions beyond explicit data in the knowledge graph.

This approach provides enhanced coverage, deeper semantic reasoning, and access to continually updated comprehensive knowledge, all while eliminating the need for LLM fine-tuning. However, it introduces the challenges of maintaining and synchronizing an additional database as data evolve. Additionally, there are considerations of added processing time, the initial development and upkeep of the ontology, and the associated expenses of implementing a reasoning engine.

## Conclusion

As we navigate the dynamic field of artificial intelligence, the integration of knowledge graphs with generative AI stands out as a transformative approach to enable trustworthy and responsible enterprise decisions. Knowledge graphs are characterized by their structured representation of knowledge, where entities (often denoting concepts or objects) and the relationships between them are explicitly defined. They provide a reliable and actionable map of knowledge to not only capture explicit facts but also enable semantic understanding, allowing for the deduction of implicit knowledge.

On the other hand, generative AI is characterized by its ability to produce new content, be it text, images, music, or other forms of data, that mirrors or emulates the patterns seen in its training data.

Merging these two powerful realms into a knowledge-enriched AI ecosystem facilitates the production of outputs that are not only precise but imbued with contextual richness.

Knowledge graphs can consistently and accurately influence and guide the behavior of generative AI models such as LLMs, without the need for dedicated training. As LLMs continue to evolve daily, knowledge graphs offer a means to ensure consistent quality across top-tier models. Consequently, investing in the development of high-quality enterprise knowledge graphs is beneficial, irrespective of the particular LLM in use.

This Whitepaper is just the start of our journey into knowledge-enriched AI. While we have touched upon its importance and have come up with some ideas for its realization, there is much more to unfold in this domain in the upcoming years.

As technology advances, the blend of AI with deep knowledge will offer more innovations and breakthroughs. We are at the beginning of an exciting journey, and there is a lot to look forward to...

### Suggestions for further reading

---



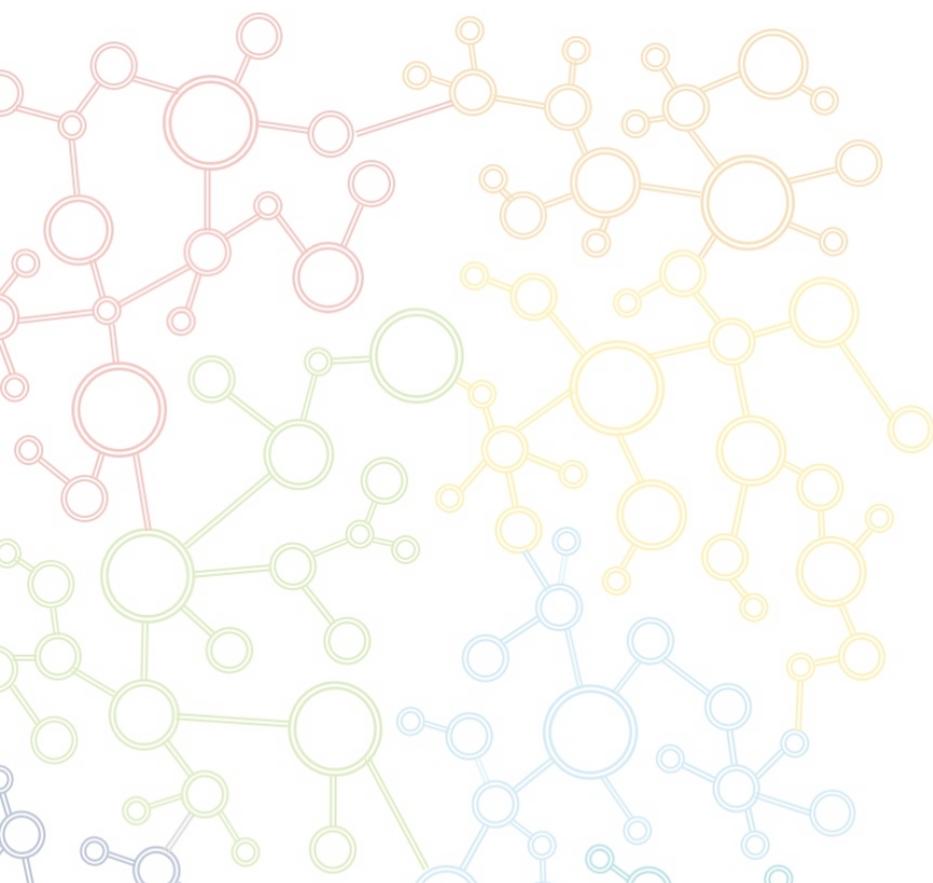
#### **Knowledge Graphs for Financial Services**

The path to unlock new insights from your data



#### **Wisdom of Enterprise Knowledge Graphs**

The path to collective intelligence within your company



## Contact Us

**Dr. Ali Khalili**

Sr. Manager, Responsible Data & Analytics  
Deloitte Risk Advisory, the Netherlands  
Sr. Research Fellow, Vrije Universiteit Amsterdam  
<alkhalili@deloitte.nl>

**Prof. Frank van Harmelen**

Head of Learning & Reasoning Research Group  
Artificial Intelligence, Faculty of Science  
Vrije Universiteit Amsterdam, the Netherlands  
<frank.van.harmelen@cs.vu.nl>

**Aïcha Mohazzab**

Jr. Manager, Responsible Data & Analytics  
Deloitte Risk Advisory, the Netherlands  
<amohazzab@deloitte.nl>

**Dr. Michael Cochez**

Assistant Professor  
Learning & Reasoning Research Group  
Artificial Intelligence, Faculty of Science  
Vrije Universiteit Amsterdam, the Netherlands  
<m.cochez@vu.nl>

**Nizar Es-Skali**

Sr. Manager, Responsible Data & Analytics  
Deloitte Risk Advisory, the Netherlands  
<nes-skali@deloitte.nl>

**Dr. Saeed Khanagha**

Associate Professor  
School of Business and Economics  
Management and Organisation  
Vrije Universiteit Amsterdam, the Netherlands  
<s.khanagha@vu.nl>

**Daphne van Leent**

Manager, Responsible Data & Analytics  
Deloitte Risk Advisory, the Netherlands  
<dvanleent@deloitte.nl>

**Marko van Zwam**

Partner, Responsible Data & Analytics  
Deloitte Risk Advisory, the Netherlands  
<mvanzwam@deloitte.nl>

**Marc Verdonk**

Partner, Digital Risk Solutions  
Deloitte Risk Advisory, the Netherlands  
<mverdonk@deloitte.nl>



Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited (“DTTL”), its global network of member firms, and their related entities (collectively, the “Deloitte organization”). DTTL (also referred to as “Deloitte Global”) and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see [www.deloitte.com/about](http://www.deloitte.com/about) to learn more.

Deloitte provides industry-leading audit and assurance, tax and legal, consulting, financial advisory, and risk advisory services to nearly 90% of the Fortune Global 500® and thousands of private companies. Our professionals deliver measurable and lasting results that help reinforce public trust in capital markets, enable clients to transform and thrive, and lead the way toward a stronger economy, a more equitable society and a sustainable world. Building on its 175-plus year history, Deloitte spans more than 150 countries and territories. Learn how Deloitte’s more than 415,000 people worldwide make an impact that matters at [www.deloitte.com](http://www.deloitte.com).

This communication contains general information only, and none of DTTL, its global network of member firms or their related entities is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser. No entity in the Deloitte organization shall be responsible for any loss whatsoever sustained by any person who relies on this communication.

© 2023. For information, contact Deloitte Netherlands.