For many enterprises, AI isn't delivering.

The solution is knowledge graphs.

metis

In recent years, the enterprise landscape has become saturated by AI tools and platforms promising unprecedented growth and efficiency.

But for many enterprises, the return on their AI investment has been disappointing. Instead of transforming operations, AI strategies are failing to drive real growth; their applications are more limited, more expensive, and less trustworthy than decision makers have been led to believe.

We are entering what Gartner describes as the AI "trough of disillusionment" – the moment in the hype cycle when, for our purposes, enterprises begin to question whether AI can ever deliver on industry promises. In moments like these, it may be tempting to shelve that pilot – but this is precisely the moment when clear strategic vision matters most.

For many, implementing an AI strategy has proven more complex than it seems – mostly because the most popular tools, like large language models (LLMs), aren't actually fit for purpose.

Neural networks and LLMs are powerful systems. But they do not understand your business. They lack context, they hallucinate, and they produce unexplainable outputs that cannot be interrogated or defended.

When these models fail to produce promised results, leaders find themselves faced with two options: persist with incomplete, high-risk AI programmes, or retreat and fall behind competitors.

Fortunately, there's an alternative: fortify the existing strategy to execute as intended.

The missing piece is the knowledge graph: a symbolic layer that gives LLMs the structure, scope and context required to drive meaningful growth in complex enterprises.

Pairing LLMs with knowledge graphs creates a neuro-symbolic AI system capable of producing accurate, contextual, traceable outputs that leaders can trust. Combined, these technologies unlock enterprise data in ways that standalone LLMs cannot: creating high-value, context-specific insights, building user confidence, reducing operational risk, and enabling AI agents that can act on behalf of users in predictable, transparent ways.

01

Enterprises are already reaching the edges of their AI ability...

02

AI is meant to drive innovation and efficiency—but without the right tools, enterprises aren't seeing results...

03

But giving up on AI implementation is the best way to fall behind...

04

Successful implementation of truly powerful, agentic AI relies on a joined-up approach: one that builds trust within the business and confidence among users...

05

Neuro-symbolic AI provides the framework for future-proof implementation...

06

The marriage of neural networks and LLMs with knowledge graphs helps build trust across enterprises...

07

And empowers users to bring enterprise data into contact in innovative ways...

08

All while protecting the enterprise from a loss of institutional knowledge...

09

With agentic knowledge-driven AI, if a user can ask for it, they can have it—and this is the genesis of enterprise innovation...

Enterprises are already reaching the edges of their AI ability

AI and machine learning have been embedded in enterprise systems across industries for decades. But the last five years have placed AI tools and agents in the hands of everyday users. This rise has been explosive: in 2024, **78%** of organisations reported using AI, up from **55%** just one year earlier (Stanford).

Much of this adoption was driven by claims of extreme productivity improvements and dramatic cost efficiencies. Enterprises mobilised quickly, investing in AI pilots and agents that promised to revolutionize how they do business.

But for many, the reality has been underwhelming, with results far more muted than anticipated. **80%** of companies using generative AI saw no significant bottom-line impact, and **42%** ultimately abandoned their AI strategy altogether (McKinsey). At the same time, AI-related incidents, such as compliance failures, hallucinations, and uncontrolled data exposure, have spiked, with standardised AI evaluations still rare among major developers (Stanford).

AI is undoubtedly revolutionary technology, but this difficulty shows it's hardly a silver bullet. So if AI is as powerful as advertised, how have so many enterprises with AI strategies failed to realize any material value?

At least part of the answer lies in a misunderstanding of what LLMs actually do.

A chatbot that talks back can seem like it understands—like it has the same context and goal as its user—and therefore, like it is credible. But its outputs are not rooted in experience or certainty. They are simply statistically likely to be relevant.

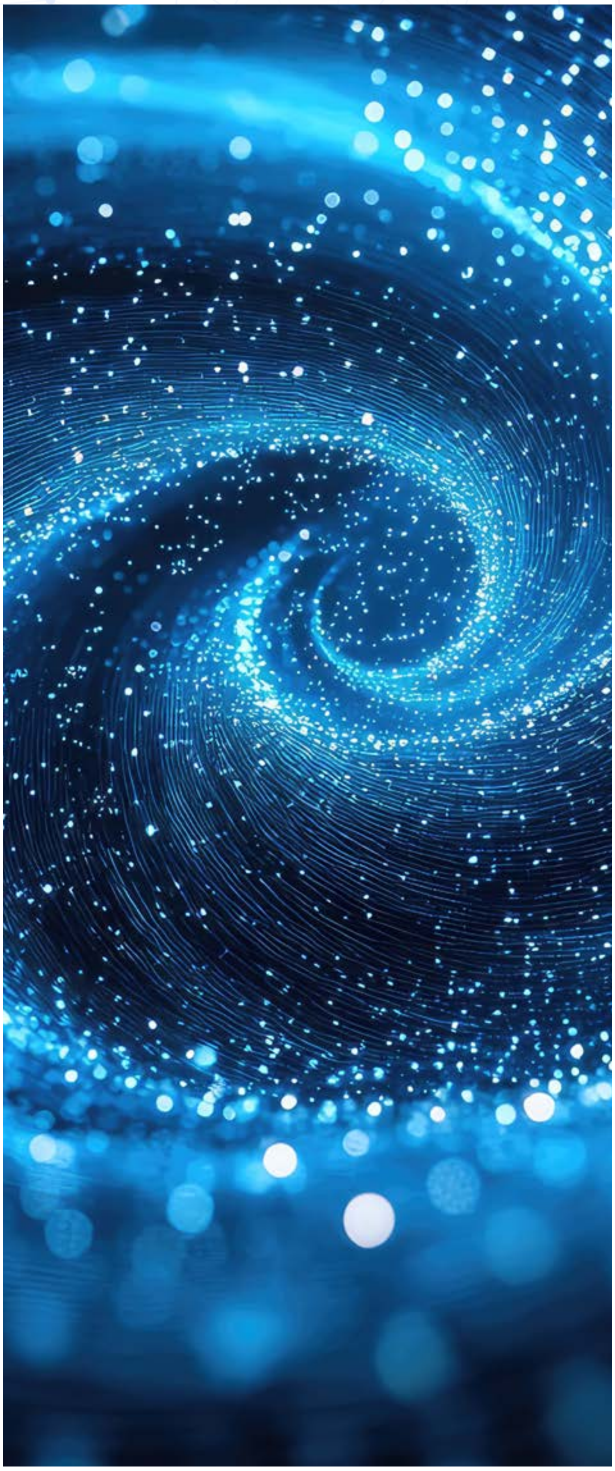

Without defining the terms and scope the AI is meant to draw from and act within—effectively grounding the model in enterprise knowledge—these systems and their outputs are unpredictable and difficult to control. This is how the illusion of AI understanding can lead organisations to deploy this technology in ways that introduce risk rather than remove it.

As a result, enterprises are encountering the limits of LLM-only approaches long before they have tapped into the real potential of a robust AI strategy.

AI is meant to drive innovation and efficiency—but without the right tools, enterprises aren't seeing results

You don't have to look far to find proof of this trend. An MIT study found that **95%** of AI pilot projects failed across major enterprises. At the same time, the Harvard Business Review recently documented the rise of "workslop": content that appears polished but contributes nothing to the actual task at hand. More than **40%** of US full-time employees reported receiving AI-generated content which they say did not help them in achieving their business objectives.

LLMs are meant to drive efficiency and innovation. But when outputs are erroneous, or require careful fact-checking and interrogation, they can actually decrease the speed and efficacy of existing workflows.

It's a slippery slope. AI outputs look correct but carry subtle errors. Users lose trust because results are inconsistent. Risk increases because hallucinations are difficult to spot. Compliance teams can't verify or audit decision pathways. The costs add up. And corresponding growth does not.

This is where leaders must decide whether to double down on AI and its generative potential or to give up. Both choices have risk—but only one can help ensure enterprise competitiveness in the years to come.

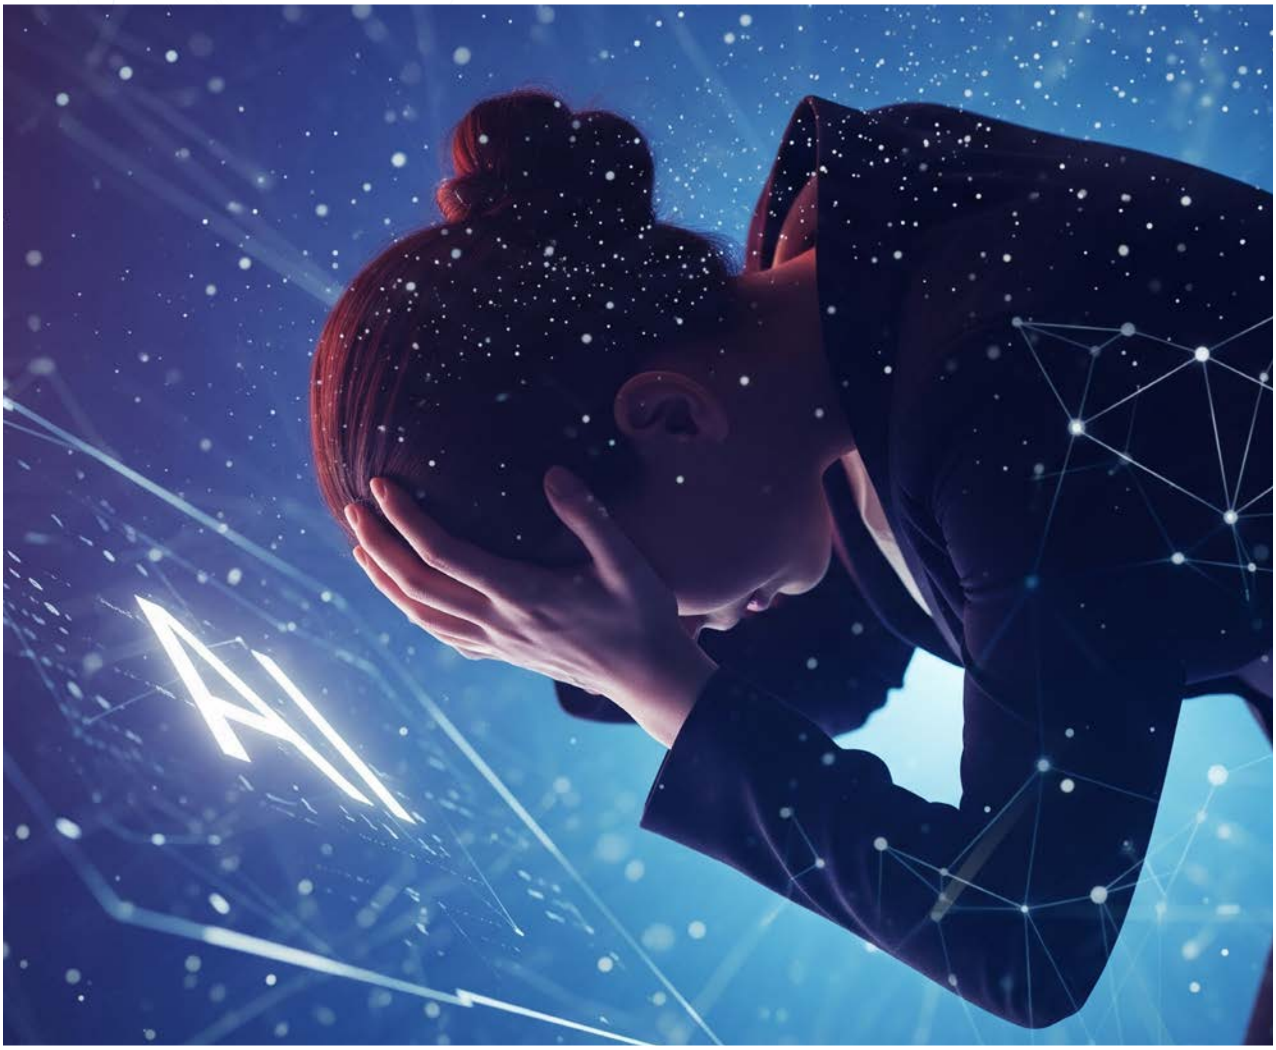

But giving up on AI implementation is the best way to fall behind

For leaders, deciding whether or not to keep on with an AI strategy isn't a simple decision. But stepping away from AI risks strategic irrelevance. Global enterprise is at an inflection point; organizations that get AI right will accelerate, and those who don't will fall behind.

Gartner predicts that by 2026, enterprises that adopt robust AI engineering practices will outperform peers by at least **25%** in their ability to operationalize models. Gartner also predicts that organizations which successfully operationalize transparency, trust and security will see a **50%** improvement in AI model adoption and business outcomes. This growth will beget more growth.

So when an enterprise scraps an AI pilot or abandons a strategy, they risk accelerating their own obsolescence. The gap will not close—it will expand.

The good news is that the failure rate noted by MIT is not a reflection of organizational incompetence—it simply reflects a problem in the architecture of an AI strategy.

In recent years, most enterprise AI strategies have relied on neural networks and LLMs alone. These systems bring massive analytical capability, but without a grounding symbolic layer, they can offer only the illusion of reason and comprehension. This is because LLMs lack the contextual definition required in enterprise environments. They cannot distinguish insight from noise. They cannot reliably follow enterprise rules. And they can't explain their outputs.

Those who solve for these limitations now will build systems that learn, scale and compound value. To do this, enterprises need strategies designed to build trust, reduce risk and enable adoption and confidence across teams.

These strategies must start at the foundation—with the symbolic layer.

25%

By 2026, enterprises that adopt robust AI engineering practices will outperform peers by at least 25% in their ability to operationalize models

Gartner

50%

Organizations which successfully operationalize transparency, trust and security will see a 50% improvement in AI model adoption and business outcomes

Gartner

Successful implementation of truly powerful, agentic AI relies on a joined-up approach: one that builds trust within the business and confidence among users

For AI to integrate successfully into the fabric of an organization, it must be able to earn user trust. Models can only earn trust through consistency; any tooling must be able to follow instructions to support business objectives, provide clear access points for audit and output validation, and in so doing, minimize operational risk and eliminate untraceable hallucinations.

This is the vision of a successful AI program that many enterprises have been sold.

But LLMs alone cannot deliver the accountability, consistency or contextual intelligence required for enterprise applications. A genuinely powerful AI strategy requires a model that is knowledge-driven.

Knowledge-driven AI enriches its data and outputs with context to provide domain-specific language generation. One powerful approach is neuro-symbolic AI, which integrates LLMs with a symbolic layer—such as a knowledge graph—that formalises enterprise knowledge. This layer, informed by an enterprise-specific ontology, encodes entities, relationships, rules, and hierarchies in a machine-interpretable form, enabling reasoning, explainability, and more reliable integration of AI systems into enterprise systems.

This technology—an AI with enterprise-specific knowledge, capable of learning, collaborating, and showing its work in support of complex, goal-driven tasks—is what enables material growth.

Neuro-symbolic AI provides the framework for future-proof implementation

Neuro-symbolic AI merges two complementary technologies:

- **The neural components:** LLMs and neural networks that excel at pattern recognition, language modeling and text analysis.
- **The symbolic components:** knowledge graphs and ontologies that define the logical structure, context and constraints of enterprise data.

Most enterprises will be familiar with neural AI architecture and principles. They handle vast datasets, boast incomparable computational power, and rely on patterns and probability to generate outputs.

The symbolic layer, by contrast, brings consistent, structural order. It defines how information relates, assigns importance and significance, and sets constraints that a model must respect. This is what makes context-specific reasoning and intuition possible.

A useful analogy is the left-brain/right-brain model. **Neural networks function like the left brain: they're analytical and pattern-driven. The symbolic layer plays the role of the right brain, representing contextual, relational and interpretive processes.** One without the other produces incomplete outputs; but together, they enable AI to generate goal-driven, logically coherent results.

The neuro-symbolic approach is gaining recognition across industries, and Gartner has noted its growing role in future enterprise architectures. Early adopters are already demonstrating its ability to scale safely, transparently and effectively.

Without this architecture, enterprises will remain stuck in the trough of disillusionment – simultaneously unable to progress and unable to justify the expense of a failing strategy.

With it, they unlock the foundation for truly agentic AI: systems that can reason, plan, integrate with tools and act autonomously within guardrails that users can easily control and understand.

With agentic knowledge-driven AI, if a user can ask for it and they have permission to see it, they can have it—and this is the genesis of enterprise innovation

Think of the LLM you use most often. How does it arrive at the answers it offers you?

No one—not even its creators—can accurately and completely answer this question.

This is owing to the nature of the learning and "thinking" processes in neural networks. There's no way to understand or explain how the model generates its outputs—it's a black box. This is what causes hallucinations—and makes them so difficult to detect.

Your LLM will attest to this itself. Every relevant AI interface carries the same warning: *"Check important info."*

In small text on your screen, this comes across as nothing more than a gentle suggestion.

But for decision-makers—across private and public sectors, from executives to doctors and police who use LLMs to inform how they work—all info is important info.

Any hallucination could negatively affect another person, society, or the enterprise. This creates massive operational risk.

Without a semantic layer like the knowledge graph to underpin the AI's knowledge, it becomes significantly harder for users—let alone compliance teams—to satisfactorily assess which data a model has access to, how the model has processed that data, or how the model arrives at its outputs. Checking every output would take hours; the executive is busy, and the friendly chatbot really seems like it understands.

Trust erodes quickly after only a few failures or hallucinations—especially when there is limited remedial action that an enterprise can take.

Neuro-symbolic AI addresses these challenges structurally. The knowledge graph helps define the scope of tools, services and data that an LLM can reference when queried, informing what information it may retrieve and shaping how it formulates its answer. This enables the model to interpret data only through approved, explicitly defined rules and user permissions.

As a result:

- Hallucinations are dramatically reduced.
- Compliance teams can trace outputs to their sources.
- Risk assessments and audits become streamlined and defensible.
- Bias can be controlled by governing how data is interpreted.
- Stakeholders and users gain confidence in outputs.

Audit trails are critical across sectors, but especially in those which can impact public health and safety. A neuro-symbolic AI can monitor lab records, trial results, and regulatory documentation, cross-checking entries against defined rules. When queried, it could explain why specific outputs or recommendations were generated, showing which data sources and rules were applied. Auditors and compliance teams could trace each step, reducing errors and increasing confidence in reporting, and critically, providing a clear, auditable record for regulators.

The result is a system that is not only more accurate, but more explainable, controllable and transparent.

And empowers users to bring enterprise data into contact in innovative ways

Once users trust the AI tools at their disposal, models can be woven into everyday, role-specific workflows. This is where enterprises begin to see the step-change in efficiency and innovation they were promised.

Neuro-symbolic AI allows users to navigate and integrate information across systems and siloes through a conversational interface, supporting complex, goal-driven tasks while ensuring that reasoning and data access follow defined, enterprise-specific rules.

Sectors which produce goods especially tend to rely on multiple platforms and systems to execute, often creating huge inefficiency at scale. Across R&D, manufacturing, and supply chain, neuro-symbolic AI has the ability to pull relevant data from multiple siloed systems. A researcher might request a resource plan or production schedule via natural language, and the system could integrate data, flag conflicts, and propose optimised allocations—saving hours of manual research and reconciliation.

With the support of knowledge-driven AI agents, users may:

- Plan resources across departments with structured guidance.
- Forecast using live and historical data grounded in approved sources.
- Connect and coordinate disparate databases and tools.
- Generate enterprise-specific insights aligned with business goals and timelines.
- Automate multi-system workflows in a controlled, auditable manner.

Without a semantic knowledge graph, none of this is achievable.

The knowledge graph standardizes the enterprise ontology: who the organization is, what systems it relies on, how its data structures fit together, and how people, permissions, processes and tools interact.

Defining this structure enables agents to gather data across disparate tools and platforms, and to reason, interpret, and take action based on a cohesive understanding of that data.

With this, enterprise data can be unified and interpreted; institutional memory becomes living and accessible; and employees at every level can contribute confidently to strategic objectives—without wasting time switching back and forth between systems.

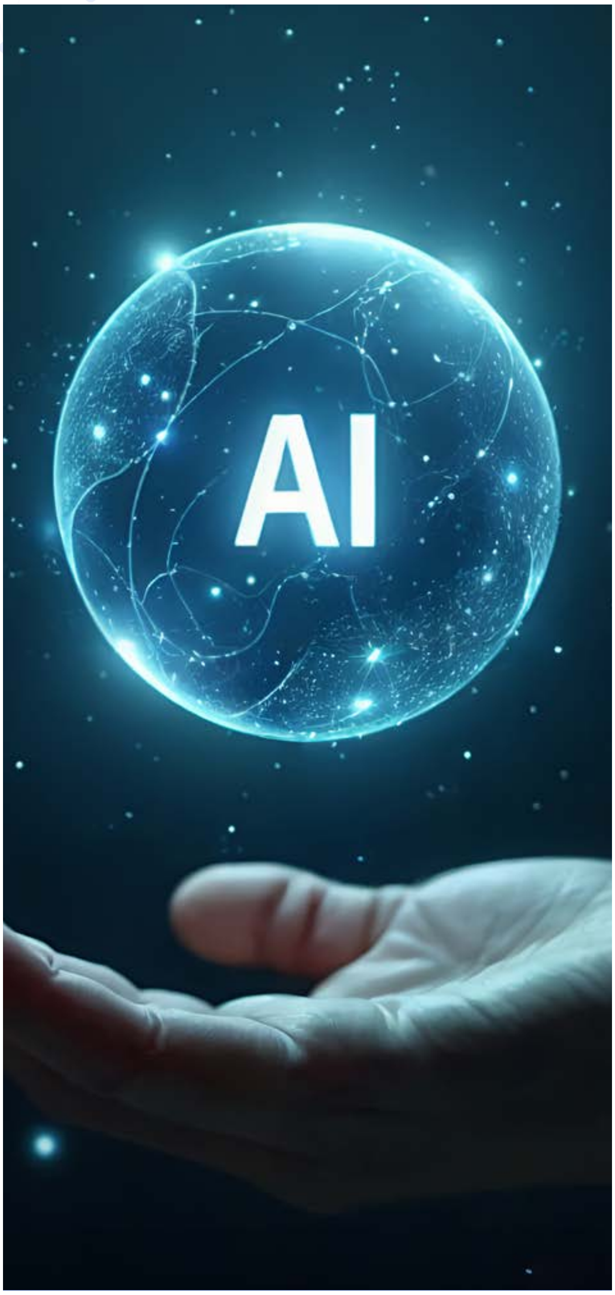

All while protecting the enterprise from a loss of institutional knowledge

Be it processes or niche product or market knowledge, each employee takes something intuitive about their role with them when they leave. This exposes enterprises of all sizes to risk with even routine churn.

Knowledge-driven AI preserves this institutional knowledge by enabling LLMs to reason over structured, enterprise-defined knowledge, capturing not just facts but the relationships, rules, and context that give them meaning. The result is a living, referenceable knowledge bank that integrates human expertise with AI-driven insights, allowing knowledge to persist beyond the tenures of specific employees.

This tried and tested repository of accessible enterprise knowledge helps stabilize teams during transitions, supports upskilling, accelerates onboarding, and provides a foundation for strategic planning, scenario modelling, and performance analysis – all while ensuring that the AI's outputs are grounded in the enterprise's specific context.

In innovation-driven industries which rely heavily on the production and protection of proprietary knowledge, preserving institutional memory is critical. When an engineer or scientist leaves, their project notes, decisions, and reasoning come to form part of the institutional record in their absence. With neuro-symbolic AI at hand, new hires can query the system to understand design rationales, historical fixes, or contextual relationships, building a living knowledge bank that evolves over time.

AI is itself driving a restructure of the labour market; certain roles are becoming obsolete as others are created. Retaining enterprise knowledge will be critical in navigating this changing landscape. With knowledge-driven AI, enterprises can keep continuity even as their workforce and workflows evolve.

With agentic knowledge-driven AI, if a user can ask for it and they have permission to see it, they can have it—and this is the genesis of enterprise innovation...

An AI strategy which leaves users feeling confused or disappointed will never yield the results that enterprises are striving for. Many organizations today are trapped in that uncomfortable middle state: enough AI to create frustration, not enough to create transformation.

But when LLMs are paired with a symbolic layer, enterprises can unlock a new way of working: natural-language collaboration with an AI agent that understands the business, its data and its goals.

Like working with LLMs, users can interact with knowledge-driven agents in natural language. They don't need extra technical skills—they can just ask for what they need. As usual, the agent responds in natural language, too.

But unlike simple LLMs, the knowledge-driven agent interprets the request through the lens of the enterprise's knowledge, cross-references with the user's permissions and clearance to protect sensitive data, and executes accordingly.

This is where true enterprise innovation can begin. Imagine if your most creative employees could harness the entirety of your enterprise data and institutional knowledge. What questions would they ask? What answers might you be surprised by?

Knowledge-driven AI turns enterprise data into a strategic asset rather than a passive repository. It enables context-aware reasoning, faster decision-making, and the creation of new agentic applications that compound value over time. This is the path to operationalising AI safely and at scale—and to future-proofing enterprise AI strategy for years to come.

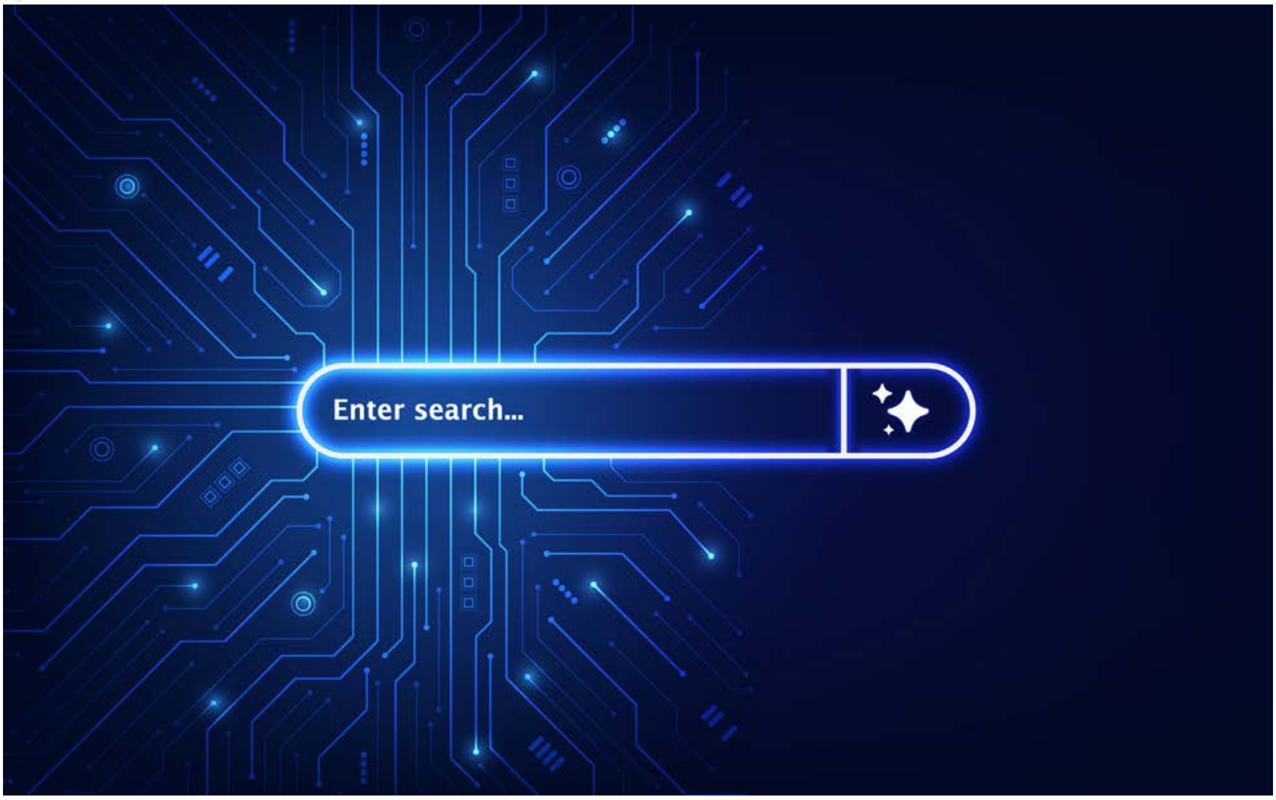

Book a demo to learn about metis—your knowledge-driven AI platform

Enterprises can't afford to settle for AI strategies that deliver only half the vision. A workforce empowered with conversational access to trusted, contextualized enterprise knowledge is within reach.

metis is a knowledge-driven AI platform purpose-built for enterprise. The metis platform brings together the strengths of LLMs and knowledge graphs to deliver enterprise-ready AI agents that combine generative power with semantic precision and return explainable, trustworthy and contextual insights.

About metaphacts

metaphacts is a semantics & AI company delivering innovative solutions that help global enterprises transform data into consumable, contextual and actionable knowledge.

metaphacts supports customers across a range of industries and use cases, such as building a semantic layer for enterprise information architecture, creating digital twins, or building trustworthy AI apps for knowledge discovery.

Sources:

Stanford HAI. (2025). The 2025 AI Index Report. Stanford, California: Stanford University. https://hai.stanford.edu/ai-index/2025-ai-index-report hai.stanford.edu

McKinsey & Company. (June 13 2025). Seizing the agentic AI advantage: A CEO playbook to solve the gen AI paradox and unlock scalable impact with AI agents. https://www.mckinsey.com/capabilities/quantumblack/our-insights/seizing-the-agentic-ai-advantage#/ McKinsey & Company

Niederhoffer, K., Rosen Kellerman, G., Lee, A., Liebscher, A., Rapuano, K., & Hancock, J. T. (September 22 2025). AI-generated "workslop" is destroying productivity. Brighton, Massachusetts: Harvard Business Review. https://hbr.org/2025/09/ai-generated-workslop-is-destroying-productivity

Gartner. (2025). What Is Artificial Intelligence? Sydney, Australia: Gartner. https://www.gartner.com/en/topics/artificial-intelligence

Gartner. (2025). Gartner Hype Cycle Identifies Top AI Innovations in 2025. Sydney, Australia: Gartner. https://www.gartner.com/en/newsroom/press-releases/2025-08-05-gartner-hype-cycle-identifies-top-ai-innovations-in-2025

Challapally, Aditya; Pease, Chris; Raskar, Ramesh; Chari, Pradyumna. (2025). The GenAI Divide. State of AI in Business 2025. Cambridge, Massachusetts: Massachusetts Institute for Technology. https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf

Contact us

metaphacts GmbH
contact@metaphacts.com
www.metaphacts.com

metaphacts-gmbh

metaphacts

metaphacts