KPMG

The dawn of the reasoning machines: AI agents

KPMG.com.au

The dawn of the reasoning machines: AI agents

This symbiotic relationship between human ingenuity and artificial intelligence is not a distant dream or a figment of science fiction. Instead, it is the reality unfolding before our very eyes, driven by the dawn of AI agents that are set to transform our industries and revolutionise the very fabric of our work and society.

Picture, if you will, a world where the boundaries between the organic and the synthetic blur into insignificance, where machines not only perform tasks but think, learn, and evolve in concert with human beings.

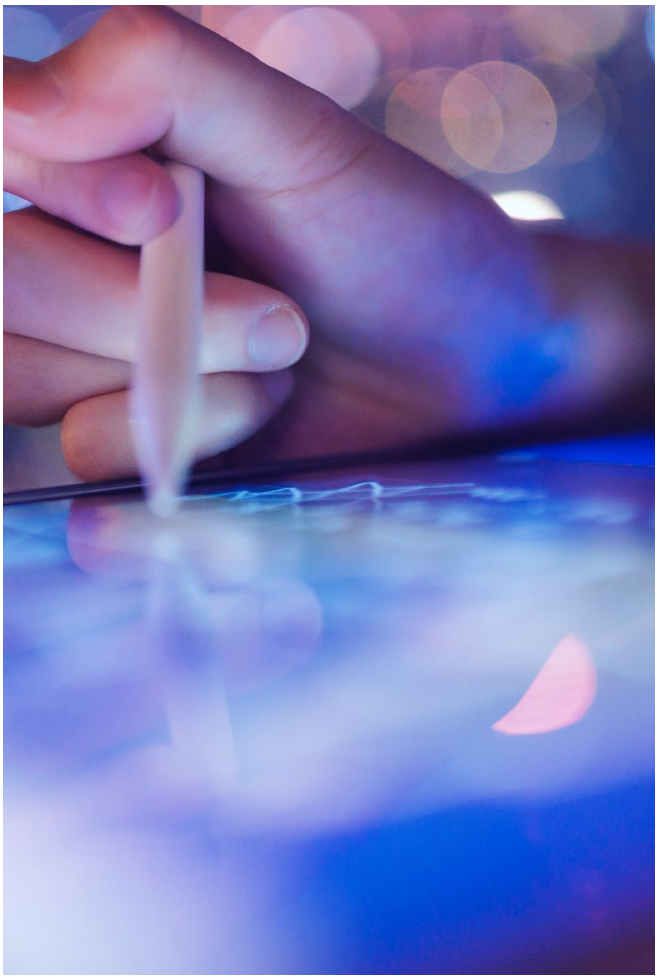

Today, we witness the emergence of AI evolve into sophisticated agents, engineered to orchestrate complex processes, learn from their environment, and amplify human potential to new heights relegated to the realm of science fiction. These machines have becoming indispensable to our technological ecosystem, signalling the dawn of an era brimming with unparalleled innovation and capability.

Within this article, we embark on an exploration of the evolution of AI agents, charting their development from basic computational models to advanced multi-agent systems that redefine the landscape of human endeavour. We will delve into the transformative impact of these technologies across various sectors, from IT and project management to HR and finance, illuminating the myriad ways in which AI agents are revolutionising operations and decision-making.

Furthermore, we will examine KPMG's Generative AI Maturity Pathway, a structured roadmap designed to guide organisations through the intricate process of integrating AI technologies effectively, responsibly and securely. Through this lens, we seek to uncover the vast potential and future opportunities presented by AI agents, ensuring that the synergy between machine intellect and human creativity can be harnessed to its fullest, shaping a future where the ability of innovation to push the frontiers of human empowerment knows no bounds.

The evolution of AI into thinking machines

The realm of artificial intelligence (AI) and machine learning (ML) has undergone a explosive evolution, progressing from its foundational roots to the cutting-edge innovations that define the current landscape. This journey can be traced through several distinct phases, each marked by remarkable advancements and transformative breakthroughs. In its earliest days, the journey of AI/ML began with rudimentary, model-based frameworks. These initial systems were built to perform narrowly defined tasks through predefined rules and logic. Early AI experiments included programs that could play chess or solve algebraic equations – tasks impressive for their time but limited in scope and adaptability. Expert systems represented a notable advancement, using rule-based algorithms to simulate the decision-making abilities of human experts within specific domains.

However, these systems encountered significant limitations. Their lack of flexibility meant that any deviation from predefined rules required manual reprogramming. Moreover, they could not learn from new data or adapt to changes in their environment. These constraints highlighted the need for more advanced technologies that could bridge the gap between static rule-following and dynamic understanding.

Advancements and the rise of machine learning

The advent of machine learning marked a turning point in the evolution of AI.

Unlike earlier systems, machine learning models could learn from data, identifying patterns and making predictions without explicit programming for every scenario. This shift enabled AI to handle more complex and diverse tasks. The increased computational power and the proliferation of data facilitated the development of sophisticated machine learning algorithms, such as decision trees, support vector machines, and eventually, neural networks.

As we entered the era of deep learning, AI systems became even more capable. Deep learning algorithms, inspired by the human brain, allowed computers to process and analyse vast amounts of data through multi-layered neural networks. This breakthrough enabled significant advancements in image recognition, natural language processing, and even game-playing with superhuman proficiency.

The emergence of generative AI

The next phase in AI's evolution was the advent of generative AI. Generative models, such as generative adversarial networks (GANs) and transformers, brought a new dimension to AI capabilities. These models could generate new, coherent content such as text, images and music, by learning from existing data. The creation of sophisticated language models like GPT (generative pre-trained transformer) exemplifies this leap, demonstrating the ability to produce human-like text and engage in meaningful conversations. These advancements pushed the boundaries of creativity and opened up new possibilities in various fields, from art and entertainment to healthcare and education.

Transition to knowledge and information retrieval

As AI continued to develop, its applications expanded into knowledge and information retrieval. AI systems became adept at understanding, indexing, and retrieving vast amounts of information – from academic research and legal documents to customer queries and product recommendations. This capability to transform data into actionable knowledge empowered industries to make informed decisions, streamline operations, and enhance customer experiences.

From agentic workflows to reasoning and thinking machines

The evolution of AI/ML did not stop at retrieving and processing information. The development of agentic workflows marked another significant milestone. In this phase, AI agents evolved from single-purpose entities into multi-agent systems capable of orchestrating complex operations. These agents could communicate, collaborate, and adapt in real time, mirroring the dynamics of human teams.

The ultimate frontier in AI's evolution is the creation of machines that not only process information but also understand context; reason through problems; and make autonomous decisions. These advanced AI agents leverage deep learning, natural language understanding, and advanced reasoning algorithms to navigate complex scenarios with a level of cognitive capability that closely resembles human thought processes. The potential of these agents spans numerous domains, including healthcare, finance, logistics, and beyond, where they can augment human intelligence and drive unprecedented innovation.

KPMG's Generative AI Maturity Pathway

Imagine embarking on a journey to harness the transformative power of generative AI within your organisation. KPMG's Generative AI Maturity Pathway serves as your trusted guide, offering a structured roadmap to navigate this complex landscape.

This pathway is designed to help businesses unlock AI's potential while upholding better practice and strict standards of reliability, security, and efficiency.

The journey unfolds through five progressive stages: Exploration of Public LLMs, Experimentation with Prototypes, Establishing a Foundation Production solution, Scaling into Agentic Flows, and Optimise use of AI capabilities through a Workbench. KPMG's Generative AI Maturity Pathway presupposes the existence and embedding of a responsible AI governance approach, the sophistication of which matches the complexity of the AI being implemented and including provision for robust and safe experimentation.

Figure 1: KPMG's Generative AI Maturity Pathway

Late 2022	Early 2023	Late 2023	Late 2024	Current
1. Explore Public LLMs	**2. Experiment** Prototype	**3. Foundation** Production self-serve knowledge	**4. Scale** Agentic Flows	**5. Optimise** Workbench
Explore the use of public LLMs through use of their chat interfaces, leveraging a variety of capabilities to test the value and applications to a range of use cases.	Prototype gen AI application using LLM services (OpenAI, Bedrock, Gemini) which support experimentation into leveraging AI for targeted use cases.	Production gen AI application using LLM services (OpenAI, Bedrock, Gemini) to support the development of internal and external virtual assistants and application integrations.	Organisational gen AI services which enable business users to develop gen AI enabled use cases.	Organisational gen AI services which enable business users to develop gen AI enabled use cases.
Accuracy: Low–Medium	Accuracy: Low–Medium	Accuracy: High	Accuracy: High	Accuracy: Very High

1. Explore
Public LLMs

The adventure begins with the exploration of public large language models (LLMs). In this initial phase, organisations delve into the capabilities of generic LLM models through their chat interfaces. This stage is all about testing the waters, understanding the value, and identifying potential applications of generative AI across various use cases and in align with business needs and strategy. Although the accuracy at this stage is relatively low to medium, it lays the groundwork for deeper AI integration.

Core features:

- Generic LLM model
- Restricted to public information
- Limited configuration
- Limited data access restriction

Accuracy:

Low–Medium

KPMG's journey timeline: Late 2022

2. Experiment
Prototype

As the journey progresses, organisations move to the Experiment Prototype stage. Here, they prototype generative AI applications using services like OpenAI, Bedrock, and Gemini. This phase is a playground for experimentation, allowing businesses to refine their AI strategies and develop prototypes that showcase the potential benefits of generative AI in specific contexts.

Core features:

- Compliance library vector store
- Q&A user interface
- Prompt suggestions

Accuracy:

Low–Medium

KPMG's journey timeline: Early 2023

3. Foundation
Production self-serve knowledge

The third stage involves developing foundational production applications using LLM services. This includes creating internal and external virtual assistants and integrating AI into various applications. At this stage, the focus is on ensuring high accuracy and robust data security. Organisations can leverage a range of capabilities, including role-based APIs, vector databases, data dictionaries, and data quality assessors. The goal is to build a solid foundation for AI applications that can be scaled and optimised in subsequent stages.

Core features:

- Compliance library vector store
- Q&A user interface
- Prompt suggestions
- Role-based APIs, vector database, data dictionary, LLM and model settings
- Data quality assessor
- AI persona studio
- Data segregation and security
- Prompt confidence scoring

Accuracy:

Medium–High

KPMG's journey timeline: Late 2023

4. Scale
Agentic Flows

In the fourth stage, Scale Agentic Flows, organisations focus on scaling their generative AI services. This stage empowers business users to develop AI-enabled use cases by creating bespoke user interfaces and embedding AI workflows into business processes. The emphasis is on metering and telemetry to monitor AI usage and performance, driving innovation and efficiency across various functions.

Core features:

- Bespoke UI (beyond Q&A)
- Embedded workflow
- Metering and telemetry

Accuracy:

High

KPMG's journey timeline: Late 2024

User interface:

Business-driven and third party application

5. Optimise
Workbench

The final stage of the journey, Optimise Workbench, is where organisations achieve peak AI performance. This stage involves optimising the AI workbench to attain very high accuracy and efficiency. By implementing a multi-model architecture and centralised AI microservices, organisations ensure seamless integration and performance across different AI applications. The goal is to create a robust and scalable AI infrastructure that supports a wide range of use cases and delivers significant business value.

Core features:

- Multi-model/multi-LLM
- Centralised AI microservices architecture
- Agentic architecture
- Cross application metering and telemetry

Accuracy:

Very High

KPMG's journey timeline: Current

KPMG Australia proudly leads the charge in AI excellence, as the first firm globally to receive ISO 42001 certification by BSI for AI Management Systems.

User interface:

Business-driven

KPMG's Generative AI Maturity Pathway offers a clear and structured approach for organisations to adopt and integrate generative AI technologies. By following this pathway, businesses can explore AI's potential, experiment with prototypes, develop foundational applications, scale their AI capabilities, and optimise their AI infrastructure. Each stage builds upon the previous one, ensuring that organisations can leverage AI's capabilities while maintaining high standards of accuracy, security, and efficiency. This comprehensive approach enables businesses to drive innovation, enhance performance, and achieve significant business value through the adoption of generative AI.

KPMG HAS BEEN ENABLING OUR CLIENTS TO JUMP FROM LEVEL 1 TO LEVEL 5 IN LESS THAN 8 WEEKS

In the dynamic landscape of modern business, AI agents are emerging as transformative tools that redefine how organisations operate.

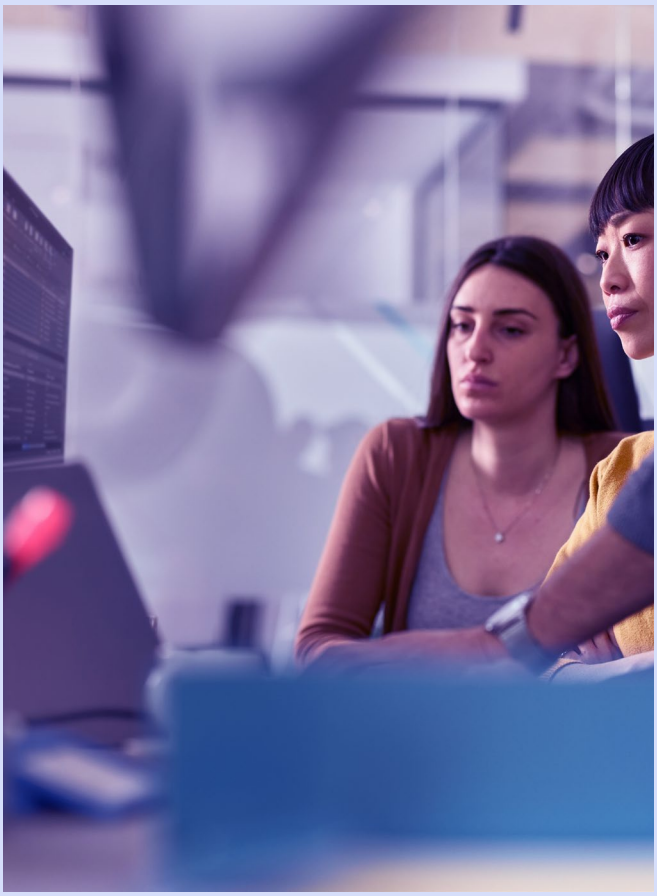

Embracing the future: the evolution and impact of AI agents

These sophisticated programs are designed to automate and execute business processes, working alongside or on behalf of individuals, teams, or entire organisations. The journey of AI agents from simple task-based systems to advanced, autonomous entities capable of handling complex tasks and interactions is nothing short of revolutionary. This article delves into the development, use cases, and significance of AI agents, offering insights into how organisations can harness their potential for enhanced efficiency and innovation.

The evolution of AI agents

The development of AI agents has been marked by significant advancements, each contributing to their growing capabilities and applications. Initially, AI agents were model-based, relying on predefined models to perform specific tasks. These early agents were limited in scope but laid the foundation for more sophisticated systems.

As technology progressed, monolithic agents emerged. These standalone agents could handle tasks independently, offering greater flexibility and functionality. However, the true potential of AI agents began to unfold with the advent of multi-agent systems.

These systems involve multiple agents working collaboratively to achieve a common goal, often incorporating human-in-the-loop (HIL) mechanisms to enhance performance. This collaborative approach allows for more complex and dynamic interactions, paving the way for AI agents to tackle a broader range of business challenges.

AI agents are being deployed across various domains, streamlining operations, improving efficiency, and enhancing decision-making. Here are some notable use cases that highlight their versatility and impact.

Sample transformative use cases of AI agents

1 IT helpdesk agent

These agents assist users with technical issues, such as connecting to corporate networks, providing quick and efficient support that reduces downtime and enhances productivity.

2 Project tracker agent

By monitoring project status, budgets, and timelines, these agents ensure that projects stay on track, enabling better resource management and timely completion.

3 Employee self-service agent

These agents provide information on company policies and process leave of absence requests, simplifying HR processes and improving employee satisfaction.

4 Budget management agent

Managing financial planning and reviewing outstanding purchase orders, these agents help organisations maintain financial health and make informed budgeting decisions.

5 Sales order agent

Processing sales orders and identifying issues requiring attention, these agents streamline sales operations and enhance customer service.

6 Testing agent

Generating test data and conducting tests based on application acceptance criteria, these agents ensure software quality and reliability, reducing the risk of errors and enhancing user experience.

The importance of AI agents

AI agents represent a paradigm shift in how businesses operate, offering several critical benefits that underscore their importance:

1 **Increased efficiency**: By automating repetitive and time-consuming tasks, AI agents free up human resources for more strategic activities, boosting overall productivity.

2 **Enhanced decision-making**: AI agents provide real-time data and insights, enabling better-informed decisions that drive business success.

3 **Cost savings**: Automation of routine tasks reduces operational costs, improving return on investment and allowing organisations to allocate resources more effectively.

4 **Personalisation**: AI agents can deliver personalised experiences and interactions, improving customer satisfaction and engagement by tailoring services to individual needs.

5 **Accuracy**: AI agents can deliver significantly higher accuracy and consistency of results due to the focused knowledge and specialised skills.

Recommendations for organisations

To leverage the full potential of AI agents, organisations should consider the following guiding principles:

1

People-first mindset

Foster a culture of "People-first mindset" empowering teams to identify the most effective and ethical solutions for specific business needs, delivering groundbreaking advancements.

2

Establish an 'AI Agent library'

Each project is an opportunity to create and/or improve your organisation's re-usable digital asset (in the form of AI agents).

3

Implement robust data quality security measures

Addressing concerns related to inadequate data quality, data leakage and oversharing is essential. Implementing stringent security protocols can protect sensitive information and maintain trust.

4

Experimentation and innovation

Encouraging safe experimentation with different AI models can help identify the most effective solutions for specific business needs. This approach fosters innovation and allows organisations to stay ahead of the competition.

5

Stay informed on regulatory compliance

Keeping abreast of evolving AI regulations and ensuring compliance can help organisations avoid legal and ethical pitfalls. Staying informed also positions organisations to adapt to new requirements seamlessly.

The integration of AI agents into business operations is not just a technological advancement; it is a strategic imperative. By embracing AI agents, organisations can drive significant improvements in efficiency, decision-making, and overall business performance. As AI technology continues to evolve, the potential for AI agents to transform industries and create new opportunities is boundless. Organisations that proactively adopt and integrate AI agents will be well-positioned to thrive in the future, leveraging the power of automation and intelligence to achieve their goals.

KymX Platform

KPMG has developed a reference architecture which is designed to focus on organisational personas and enhance mission execution through a structured, multilayered approach. At its core, **Missions** represent intricate tasks that are broken down into manageable subtasks. **Agents** are autonomous entities that execute these tasks, utilising their specialised **Capabilities** such as research, planning, and tool usage. Intelligence encompasses the accumulated knowledge and insights that guide decision-making. Assets include foundational technologies like AI models and cloud solutions that empower agents. Finally, **Collaboration** integrates human and AI strengths to drive innovative solutions.

Figure 2: KymX Platform & Agent Library

Personas & Missions

Personas represent a digital worker who is tasked with performing an activity. An activity is defined as a complex task requiring multiple steps or multiple people or teams to complete.

Leadership	Change & Training	Operations	HR & Talent	Procurement	ITSM	Legal & Compliance	Finance	Customer Support	Sales
Review	Change Management	Incident report	Develop JD	Review Contract	Create Change Request	Review Contract	Create Finance Report	Helpdesk Support	Sales Assistant
Strategy	Change Strategy	Process & Procedures	Recruitment	Market Research	Request Environment	Risk Assessment	Spend Analysis	Call Analysis	Product & Services Agent
Reports	Communication Plan	Support	HR Helpdesk	Procurement Helpdesk	Raise Problem Ticket		Tax Assessment		Call Centre Review

Agents

Agents are autonomous entities that perceive, reason, and act within an environment. They collaborate, execute tasks, and adapt dynamically, leveraging memory, tools, and orchestration to fulfil missions.

Engineer	Environments	Customer	Procurement	Finance	Incident	IT Service
Contracts	Teacher	Compliance	Policy	Modelling	Problem	Business
Trainer	Safety	Paralegal	Enterprise	Lawyer	Change	

Common Agents

Analyst	Manager
Planner	Data Collector
Overseer	

Knowledge

Knowledge represents the accumulated information used to reason, learn, and adapt. It includes structured data, contextual memory, and dynamic insights that enhance decision-making and mission execution.

Assets	Policies	Education	Risk & Compliance	Legislation	Standards	Finance	Customer & Products
Manuals	Human Resources	Curriculum	Procedures	Telco	Energy	Financial History	Products
Capacity	Procurement	Assessments	Controls	Energy Market	Safety	Contracts	Support
Performance	Customer		Risk	Health & Safety	Finance	Tax	Services
	IT				Accounting	Business Cases	
	Finance				ITIL		

Technology

Assets serve as foundational technologies including models, frameworks, APIs, and computing infrastructure. These assets empower agents with intelligence, enabling efficient task execution, adaptability, and scalability in mission-driven environments.

Large Language Models		Traditional AI		Business Intelligence	Data Management
Open AI 4o	DeepSeek R1	Optimisation	Supervised Learning	Reporting	Data Lineage
OpenAI o1	Llama	Classification	Anomaly Detection	Analytics	Metadata
Claude Sonnet	Sora	Clustering	Decision Trees	Trends	Structured Data Query
Gemini Pro	Dalle	Forecasting		Data Modelling	Data Dictionary
Eleven Labs		Linear Regression			

Data Sources		Interfaces
Customers	Service Management	API
Workforce	Finance	Middleware
Payroll	HR	FTP
Assets	GIS	ODBC
Sales	Payroll	File

KPMG has consolidated all its learnings and experience from generative AI application development into a library of AI agents ready for client use cases.

1. Personas & Missions:

Missions are complex tasks requiring multi-step reasoning, planning, and execution. Agents decompose these missions into subtasks, leveraging orchestration, memory, and adaptive decision-making to achieve objectives efficiently. Examples include developing comprehensive reports and articles (Editorial & Content Creation), implementing new software training programs (Change & Training), conducting safety audits and compliance checks (Operations & Safety), managing recruitment processes and talent development (HR & Talent), and handling customer inquiries and support tickets (Customer Support).

2. Agents:

Agents are autonomous entities that perceive, reason, and act within an environment. They collaborate, execute tasks, and adapt dynamically, leveraging memory, tools, and orchestration to fulfil missions. Examples include engineers designing and maintaining technical systems, trainers conducting employee training sessions, customer support agents assisting customers with queries, finance specialists managing financial transactions and reports, and lawyers providing legal advice and drafting contracts.

3. Knowledge:

Intelligence represents the accumulated knowledge agents use to reason, learn, and adapt. It includes structured data, contextual memory, and dynamic insights that enhance decision-making and mission execution. Examples include developing training curricula (Curriculum), creating detailed lesson plans (Lesson Plans), maintaining a database of past financial contracts (Financial Contracts History), conducting risk assessments (Risk Assessments), and ensuring compliance with health and safety regulations (Health & Safety).

4. Technology:

Assets serve as foundational technologies, including models, frameworks, APIs, and computing infrastructure. These assets empower agents with intelligence, enabling efficient task execution, adaptability, and scalability in mission-driven environments. Examples are utilising advanced AI models (OpenAI Models), leveraging cloud-based solutions for data storage and processing (Claude), implementing machine learning frameworks (Sonnet), using advanced analytics tools for data insights (Gemini Pro), and employing specialised tools for data analysis and reporting (Eleven Labs).

Anatomy of an employee

By using an employee as the reference model, we delve into the aspects that empower them to excel in their role and perform their tasks. By understanding these elements, we map them to the KymX Architecture's layers, facilitating a harmonious collaboration between AI agents and human employees, and thereby cultivating a cohesive and efficient working environment.

Enable organisations to forge a digital workforce.

Purpose

An employee's purpose acts as a guiding compass, directing their actions towards meaningful and impactful outcomes. This sense of direction is akin to the missions assigned to AI agents within the KymX Architecture.

Learning

For employees, learning is a dynamic process of exploring, absorbing, and integrating new experiences and knowledge. This continuous growth and adaptability are essential for personal and professional development, similar to how AI agents refine their models through iterative learning.

Experience

An employee's experience is the cumulative tapestry of interactions and emotions that shape their understanding of the world. This experience informs their choices and drives personal growth, much like how AI agents accumulate and leverage past interactions to enhance decision-making.

Processes and procedures

Employees rely on structured frameworks that transform abstract goals into concrete actions. These processes ensure clarity, consistency, and continuous improvement in their pursuits, mirroring the systematic operations of AI agents.

Capabilities

Employees possess a blend of innate talents, acquired skills, and learned experiences that empower them to innovate, adapt, and overcome challenges. These capabilities are analogous to the specialised skills and techniques used by AI agents to execute tasks efficiently.

Tools

Employees utilise tangible and digital tools that extend their natural abilities, transforming abstract ideas into concrete actions. These tools amplify their creativity and productivity, similar to the suite of integrated systems and technologies that enhance the functional reach of AI agents.

Collaboration

Collaboration for employees involves uniting diverse talents and perspectives, fostering open dialogue, and mutual support to achieve common goals. This collaborative effort drives innovation, akin to the synergistic exchange of knowledge between human partners and AI agents.

In conclusion, *The dawn of the reasoning machines: AI agents* **underscores the transformative potential of AI agents in reshaping the modern business landscape. As we navigate the 21st century digital revolution, AI agents have evolved from rudimentary tools to sophisticated systems capable of executing complex tasks and enhancing decision-making processes. Their integration into various sectors will not only streamline operations, but also drive significant cost savings and operational efficiencies.**

The journey of AI agents, from basic model-based frameworks to advanced multi-agent systems, highlights the relentless pursuit of innovation and efficiency within organisations. These agents are becoming indispensable across diverse domains, from IT helpdesks and project management to HR and finance, demonstrating their versatility and impact.

The collaborative nature of multi-agent systems, incorporating human-in-the-loop mechanisms, further amplifies their potential, enabling them to tackle a broader range of business challenges with higher efficacy.

KPMG's Generative AI Maturity Pathway provides a structured roadmap for organisations to harness the power of AI, ensuring a comprehensive approach to adoption and integration in a way complementary to and building upon the foundation of a responsible AI governance strategy. By following this pathway, businesses can explore, experiment, and optimise AI capabilities, maintaining high standards of accuracy, security, and efficiency.

As AI agents become more pervasive, their role in driving operational excellence and strategic advancements will continue to grow. Organisations poised to embrace this technological frontier will not only enhance their competitive edge but also unlock new opportunities for innovation and growth. The dawn of the thinking machines heralds a new era where the synergy between human ingenuity and artificial intelligence promises to reshape industries, economies, and societies at large.

Contact us

Dhawal Jaggi
Partner, AI Foundations & Transformation Offering Lead
KPMG Australia
E: dhawaljaggi@kpmg.com.au

John Draper
Director, National AI Solutions Architect
KPMG Australia
E: jdraper2@kpmg.com.au

Daniel Ferguson
Partner, Data and AI Offering Portfolio Lead
KPMG Australia
E: dferguson4@kpmg.com.au

John Munnelly
Chief Digital Officer
KPMG Australia
E: jmunnelly@kpmg.com.au

Jessica Wyndham
Trusted AI Lead
KPMG Australia
E: jwyndham@kpmg.com.au

Richard Boele
Chief Purpose Officer
KPMG Australia
E: rboele@kpmg.com.au

KPMG.com.au

June 2025. 1638102630ALL