



Introducing the AI agents life cycle

A practical guide to understanding
a new era in AI



A resource for public and private organizations

This e-book is intended to help organizations prepare for what's next in AI.



Contents

- 01 **Foreword:** The agentic AI revolution
- 02 **Perception:** The eyes and ears of AI agents
- 03 **Cognition:** Making sense of the data
- 04 **Decisioning:** Choosing the best path forward
- 05 **Action:** Implementing the decision
- 06 **Learning:** Continuously learning over time
- 07 Industry use cases
- 08 Looking ahead

01

Foreword:

The agentic AI revolution

AI is undergoing a seismic transformation.

From the early days of rule-based automation to the rise of machine learning, each leap has redefined what's possible. The arrival of ChatGPT in 2023 marked a turning point, followed by the proliferation of multimodal AI in 2024, blending text, images and more into unified digital experiences.

Today, we stand at the threshold of the agentic era. Agentic AI represents a fundamental shift: It moves beyond content generation, considered the hallmark of GenAI, toward decision-oriented operations. Unlike traditional systems that require explicit prompts or operate within rigid boundaries, agentic AI acts as an intelligent teammate, capable of independent reasoning, memory and multistep action with minimal human guidance. The future is one where AI evolves from passive assistants to proactive, autonomous agents. It's an exciting time.

Yet, this promise is met with a challenge: the gap between AI hype and organizational readiness. Organizations worldwide, of all sizes, are experimenting with large language models (LLMs), seeking to harness organization wide intelligent decision making. **McKinsey estimates** a staggering US\$4.4 trillion in long-term productivity growth potential from AI across industries.

To unlock this value, leaders must rethink ROI. AI is not a bolt-on tool and employee reducer. It's a foundational capability shift, transforming how we think, decide and evolve. Its true impact lies in workplace innovation, faster responses, personalized interactions and agility in the face of rapid change.

Agentic AI is already reshaping industries: streamlining chronic disease management in health care, unlocking unstructured data in manufacturing, delivering actionable insights for insurance and empowering risk assessment in banking.

Looking ahead, we will see AI systems better balance strict instructions with real-world judgment, closing the gap between human common sense and machine logic. **Gartner predicts** pilot projects will expand into real deployments, ushering in a new era of practical, transformative AI.

In this e-book, we'll explore the components of an AI agent in a five-step life cycle and how agents combine to create powerful, autonomous systems.



“AI decisions impact people, industries and society. Decision intelligence ensures these decisions are robust, ethical and aligned with long-term goals.”

Marinela Profi,
Global Agentic AI Market Strategy Lead, SAS

02

Perception:

The eyes and ears of AI agents

An AI agent typically consists of five main components. Let's first break it down in perception.

At the heart of any AI agent is its ability to perceive the world around it. This is where the AI agent senses and collects information. Just like how we use our eyes to see and ears to hear, the AI agent gathers data from various sources (e.g., sensors, inputs, user interactions or databases) to understand the context in which it's operating.


In this phase, keep in mind that the quality and breadth of data gathered during perception play a huge role in the AI's overall effectiveness. If the AI lacks accurate or relevant data, its decisions will be based on incomplete information. Perception essentially sets the stage for everything that follows.

In the case of fraud detection, this is where the AI system collects data about every transaction that occurs in real time. This could include:

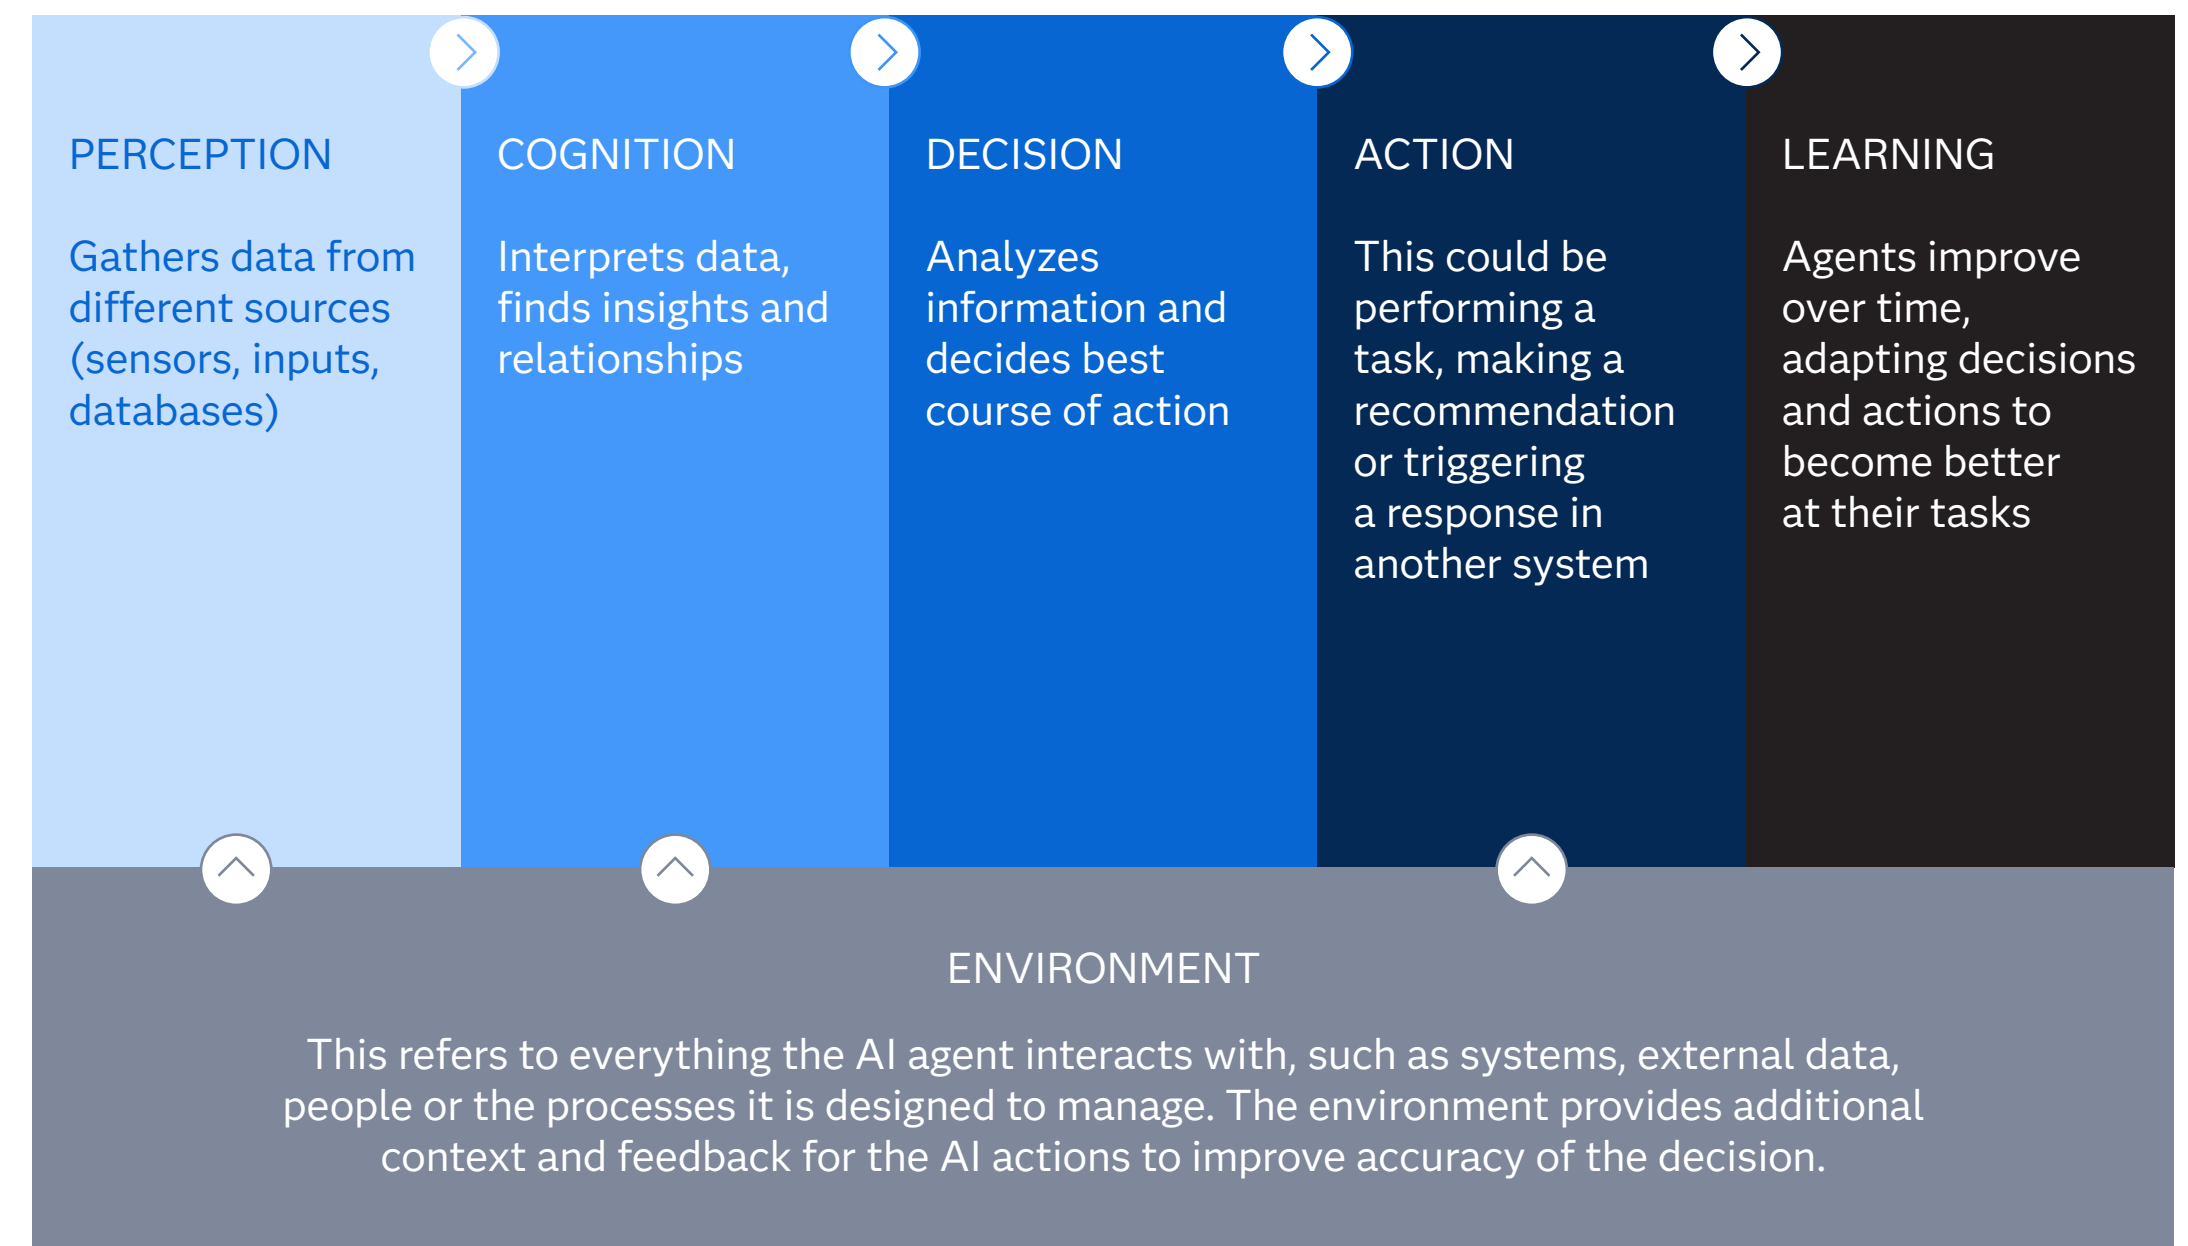
- The transaction amount.
- The location where the purchase is made.
- The time of the transaction.
- The cardholder's previous transaction history.
- The device or IP address used for the transaction.

For instance, when a customer uses their credit card to make a purchase, the system instantly captures this data and prepares it for analysis. In the context of fraud detection, it's crucial to gather as much relevant data as possible to assess the legitimacy of each transaction.

WHAT SOME THINK AI AGENTS LOOK LIKE



WHAT THEY ACTUALLY LOOK LIKE



03

Cognition:

Making sense of the data

Once the AI has gathered data through perception, it needs to process and interpret that information.


This is where the cognition phase comes into play. Here, the AI agent looks for patterns, identifies trends and draws conclusions from the data it has collected. It can use a combination of analytics, machine learning, linguistic rules, inference and LLMs.

Keep in mind that in this phase, the AI agent essentially “thinks” about the data, weighing different outcomes based on rules, probabilities or learned behavior. This gives the agent the foundation (cognitive understanding) to proceed to the next phase, where it will decide what to do next. The more effectively it can process and understand the data, the better its decisions will be.

In a fraud detection system, the AI analyzes the transaction and compares it against historical data and known fraud patterns. Specifically:

Pattern recognition: The AI looks for discrepancies or unusual patterns. For example, if a credit card has never been used in another country before and suddenly a purchase is made from a foreign location, the system flags this as potentially suspicious.

Risk assessment: The system evaluates the risk of the transaction by considering various factors. For example, it might analyze whether the transaction amount is unusually high for that particular user, whether it's a typical purchase for the cardholder or if there are any signs of unusual behavior.



04

Decisioning: Choosing the best path forward

The decisioning phase is where the agent determines the best course of action based on the insights gained during cognition. It's like when we make a decision based on the information we have available, whether it's choosing a business strategy, making a hiring decision or reacting to a customer's needs.

Decisioning is a pivotal moment in an AI agent's life. The decisions it makes drive its actions and ultimately determine its effectiveness. In business, a poor decision made by an AI agent could have financial, operational or reputational consequences. Having a well-defined decision intelligence framework ensures that **AI agents** can make the right choices, even in complex environments.


In the fraud detection system, based on its analysis of the transaction and the perceived risks, the agent will make a decision. For example, after processing the data, it might have several options:

- **Approve the transaction:** If the transaction is deemed legitimate based on historical patterns, it will go through.

- **Flag the transaction:** If the AI is unsure, it might flag the transaction for further review by a human fraud analyst.
- **Decline the transaction:** If the AI identifies clear signs of fraud, such as a sudden purchase from an unknown location or a suspiciously high transaction amount, it may immediately block the transaction.

The decision-making process in agentic AI is not left to casualty. It should be based on a set of human predetermined rules that help the system assess the best decision. That is why LLMs alone are not sufficient to build AI agents. You need a decisioning framework that combines LLMs, business rules, analytics, machine learning and AI governance.

[Read more](#) about the importance of decision intelligence in the age of agentic AI.



05

Action: Implementing the decision

Once the agents make a decision, they take action. This is the phase where the system executes the chosen course of action. This could be performing a task, making a recommendation or triggering a response in another system or agent.

In the fraud detection agent, the actions could be:

- **Immediate response:** If the transaction is flagged as fraudulent or suspicious, the AI agent could block the transaction, alert the cardholder or notify the fraud detection team for manual review.
- **User notification:** If a transaction is declined or flagged for review, the cardholder might receive a real-time notification asking them to verify the transaction.

For example, if an international purchase was flagged, the system might decline the transaction, send an alert to the cardholder and ask for confirmation through their phone or email.

The action phase is where the AI agent delivers its value. It's no longer just thinking or analyzing – it's doing. For AI agents to be useful in business, their actions must align with strategic goals.



06

Learning:

Continuously improving over time

The final component of an AI agent is learning. Unlike traditional systems that require manual updates or adjustments, AI agents can learn and improve over time by analyzing the outcomes of their actions.

After the agent takes action, it assesses the results. Did the action lead to the expected outcome?

If the action is successful, the AI agent strengthens the model and continues making similar decisions in the future.

If it fails, the agent adjusts its models to improve. For instance, if a fraud detection system incorrectly flags a legitimate transaction, it will learn from that mistake to avoid similar errors in the future.

Learning is what sets AI agents apart from traditional software. It enables AI to adapt, evolve and get better with each interaction. Over time, they become more accurate, more efficient and more aligned with business goals.

The role of the environment


While we've covered the five key components of an AI agent, it's also important to recognize the importance of the environment in which the AI operates.

This refers to everything the AI agent interacts with, such as systems, people or the processes it is designed to manage. The environment provides the context and feedback for the agent's perception, cognition and actions, directly affecting the quality of its decisions and its ability to learn and improve in the learning phase.

Key environmental factors include:

- External data sources (third-party fraud databases, news about data breaches, etc.).
- Regulatory frameworks (legal requirements for handling sensitive customer data).
- Customer behavior and feedback can inform future decisions and adjustments.
- Market conditions may impact fraud risk at different times (e.g., a rise in online shopping during the holiday season).

The AI system relies on data from its environment to stay informed about potential fraud risks, and this external context helps refine its decision making.



60%


of Fortune 100 companies will appoint a Head of AI Governance in 2026 to cope with emerging AI regulations and risk management. – IT Wire

How AI adapts to detect fraud

Imagine a scenario where a fraud detection AI agent detects a large transaction being made in a location far from where the cardholder usually shops. The environment impacts the process through:

- **External data source:** The agent might check with the bank's transaction history to see if this behavior fits the cardholder's usual spending patterns. It may also consider publicly available data, such as news reports or records of recent data breaches in the region or industry, which could indicate increased fraud risk in certain locations or times.
- **Contextual factors:** The agent might notice that there has been a recent data breach in the region and that many cards have been compromised, increasing the likelihood that this transaction is fraudulent.
- **Human interaction:** The customer receives an alert, and they confirm that it was not them making the purchase. The environment in this case (the human feedback) leads the system to flag the transaction as fraudulent.

- **Regulatory compliance:** Compliance with financial regulations (such as GDPR, PCI-DSS, etc.) is part of the environment. The fraud detection system must be aware of legal constraints when it collects, processes and stores sensitive data. For example, it may have to handle user data in specific ways depending on the country of origin or the financial regulations in place. The system might also ensure that its actions comply with data privacy laws, notifying the customer and handling their data securely.



AI governance and trustworthy AI work together. AI governance is what organizations do to accelerate innovation and ensure AI is trustworthy.

“Fully autonomous and LLM-only based AI agents are not enough in an enterprise context. To succeed, AI agents need determinism, guardrails and governance by design.”

Marinela Profi, Global Agentic AI Market Strategy Lead, SAS

07

Industry use cases

Insurance

CHALLENGE – Insurers possess mountains of information that employees must sift through to understand even the simplest of answers, like status of a policy application or claim. This disparate data exists in multiple places in the four walls of insurance companies, and getting those answers fast can mean the difference between a happy policyholder or the next policy cancellation.

SOLUTION – Agentic AI allows insurers to retrieve information across various data sources and make decisions and recommendations to create meaningful, lasting experiences for policyholders. A simple question, such as, “Can you tell me the status of my claim?” answered in real time with real information is the expectation.

However, agentic goes beyond a passive response to trusted advice. Agent response: “Your claim has been approved by our adjusting team. I see you’ve not used your rental coverage to secure a vehicle while your vehicle is being repaired in our Blue Ribbon repair network. Here are three options in your area. May I schedule you an appointment to pick up a rental vehicle?”

OUTCOME – AI agents create lifetime value for insurers – this, in turn, boosts retention, making premium growth goals more achievable and reduces the burden of new policy acquisition to replace policy cancellations.

Health Care

CHALLENGE – Chronic diseases are the leading driver of health care costs, accounting for a vast majority of expenditures globally. Managing chronic conditions requires ongoing monitoring, early detection of deterioration and timely interventions – tasks that can overwhelm clinical resources and care teams.

SOLUTION – By deploying AI agents, health care teams can continuously evaluate patients at the highest risk for chronic disease progression, integrating data from electronic health records, remote monitoring devices, payer databases and patient-reported outcomes. Using natural language processing (NLP), predictive modeling and machine learning pattern recognition, agents can interpret clinical notes, detect emerging health risks and track changes in biometric and utilization data in real time.

OUTCOME – AI agents identify individuals needing early interventions and alert care teams so resources can be directed where they have the greatest impact. The system supports remote care management, schedules follow-ups and provides concise, actionable insights that enable clinicians to respond proactively. This approach improves adherence to care plans, reduces preventable complications and reduces long-term health care costs through timely, data-driven decision support.

Manufacturing, Energy & Public Sector

CHALLENGE – More than 80% of enterprise data is trapped in unstructured formats (maintenance logs, product catalogues, service requests, etc.), causing longer research cycles and escalating content inaccuracy. Traditional Retrieval-Augmented Generation (RAG) solutions are helpful, but they are code-heavy, inflexible and hard to integrate.

SOLUTION – Organizations in manufacturing, energy and the public sector can use a retrieval agent manager to overcome these obstacles. They can securely launch and manage retrieval agents for smarter decision making.

OUTCOME – [According to Cervicorn Consulting](#), there is immediate value from using this approach (quantitative and/or qualitative):

- 50% reduction in research time.
- 40% increase in content accuracy.

[Learn more about SAS Retrieval Agent Manager.](#)

Banking

CHALLENGE – In commercial lending, a big pain point is the amount of manual documentation during underwriting – a slow and error-prone process.


SOLUTION – By combining OCR with AI, banking organizations can automatically extract, structure and integrate data from unstructured documents directly into risk processes, such as credit origination and risk models.

OUTCOME – AI agents speed up the underwriting process by reducing manual effort, improving customer experience and enabling a full workflow orchestration and data automation.

“Modern agentic AI systems don’t just answer questions; they plan, reason and act across complex, multistep tasks.”

Peter Williams

Head of Partner Technology,
Global Financial Services, Amazon Web Services



08

Looking ahead

Shaping the future of agentic AI

As we look ahead, agentic AI is poised to redefine the private and public sectors. Emerging trends such as agent marketplaces and composite agents will accelerate innovation, enabling organizations to deploy tailored AI solutions at scale. The focus will shift from simply gaining a competitive advantage to building true organizational resilience, like adapting rapidly to change, mitigating risks and thriving in dynamic environments.

Trust will overtake performance as the No. 1 AI KPI. We have entered an era where users trust GenAI models more than explainable ones – not because they feel more human. Moving forward, trust will become a measurable business metric. Organizations will compete not just on what their AI can do, but on whether employees, customers and regulators believe in it.

SAS is leading the way, preparing enterprises for what's next. By offering governance frameworks, a scalable agentic platform and industry-specific solutions, SAS empowers organizations to use the full potential of AI while ensuring trust, transparency and ethical decision making. **According to IDC**, by 2030, nearly half of organizations will run AI agents at scale, but those who neglect governance may face significant legal and regulatory setbacks.

The agentic era demands vision, agility and a commitment to responsible innovation. With the right strategy and partners, organizations can transform disruption into opportunities, unlocking new levels of resilience and growth. We're here to guide you along your journey.



Make a confident decision with agentic agents and systems. Learn more.

[Read the AI Blueprint: A Leader's Guide for Organizational Trust and ROI During Rapid Change](#)

[Read the Data and AI Report: The Trust Imperative](#)

[Discover ready-made AI models from SAS](#)



To contact your local SAS office, please visit: sas.com/offices