Understanding Agentic AI Landscape

Agentic AI awareness, Potential Use Cases across Industries

Agenda

- Agentic AI
- Anatomy of an Agentic AI System
- Agentic AI Tech Stack
- Potential Agentic AI Use Cases
- Case Studies
- Activity

Agentic AI

What is Agentic AI?

AI Agent: An AI agent is a system designed to reason through complex problems, create actionable plans, and execute these plans using a suite of tools.

Simple Reactive Agents: Respond to their environment based on predefined rules. Think of a thermostat that turns the heat on when it's cold.

Planning Agents: Can think ahead and create sequences of actions. Imagine a delivery robot that not only knows how to navigate from A to B but can also plan alternative routes if it encounters obstacles.

Learning Agents: Agents that improve over time by learning from their experiences. This is where LLMs come in – they can help agents understand past interactions and make better decisions in the future.

source:IBM+Author

Understanding the Shift from Reactive to Proactive Technology

Imagine Sarah, a busy entrepreneur juggling multiple projects. She's not just using a tool anymore - she's working with an AI assistant that doesn't just follow commands, but understands her goals, anticipates her needs, and takes proactive steps to help her succeed.

Agentic AI is like having a super-intelligent collaborator. It is not just a passive tool that understands your business strategy, it suggests improvements, drafts proposals, tracks market trends, and even reaches out to potential partners without you micromanaging every step.

Agentic AI is

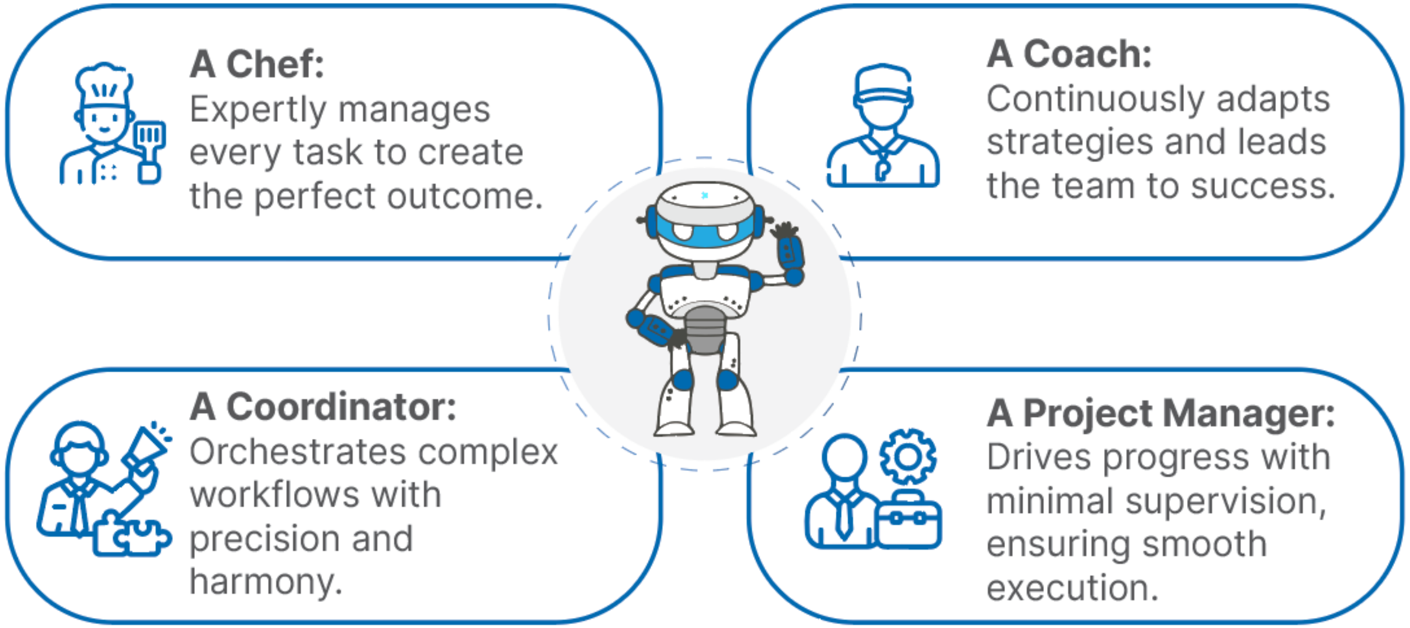

Agents are not plug-and-play solution : The Automation-Integration Framework

Integration: This term indicates how seamlessly the AI system can fit and operate harmoniously within your current system environment or business processes.

Automation: This term refers to the extent to which the AI system can execute tasks independently, with little to no human involvement.

Assistants: Enhance human efforts by providing support and augmenting tasks.

Copilots: Offer recommendations for subsequent actions or steps.

Autopilots: Handle specific tasks independently without human intervention.

Agents: Manage and coordinate multiple steps or processes seamlessly.

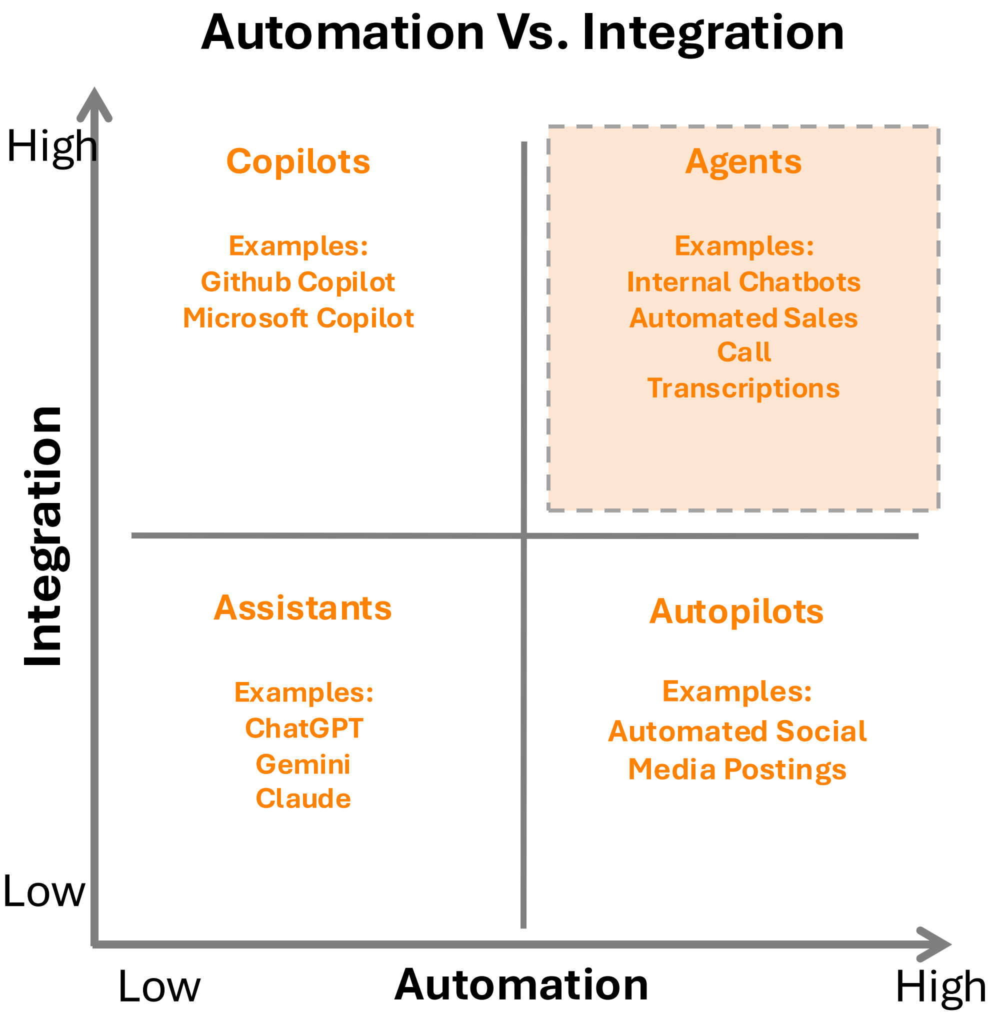

source:Author

Agent - Example: Highly personalized emails

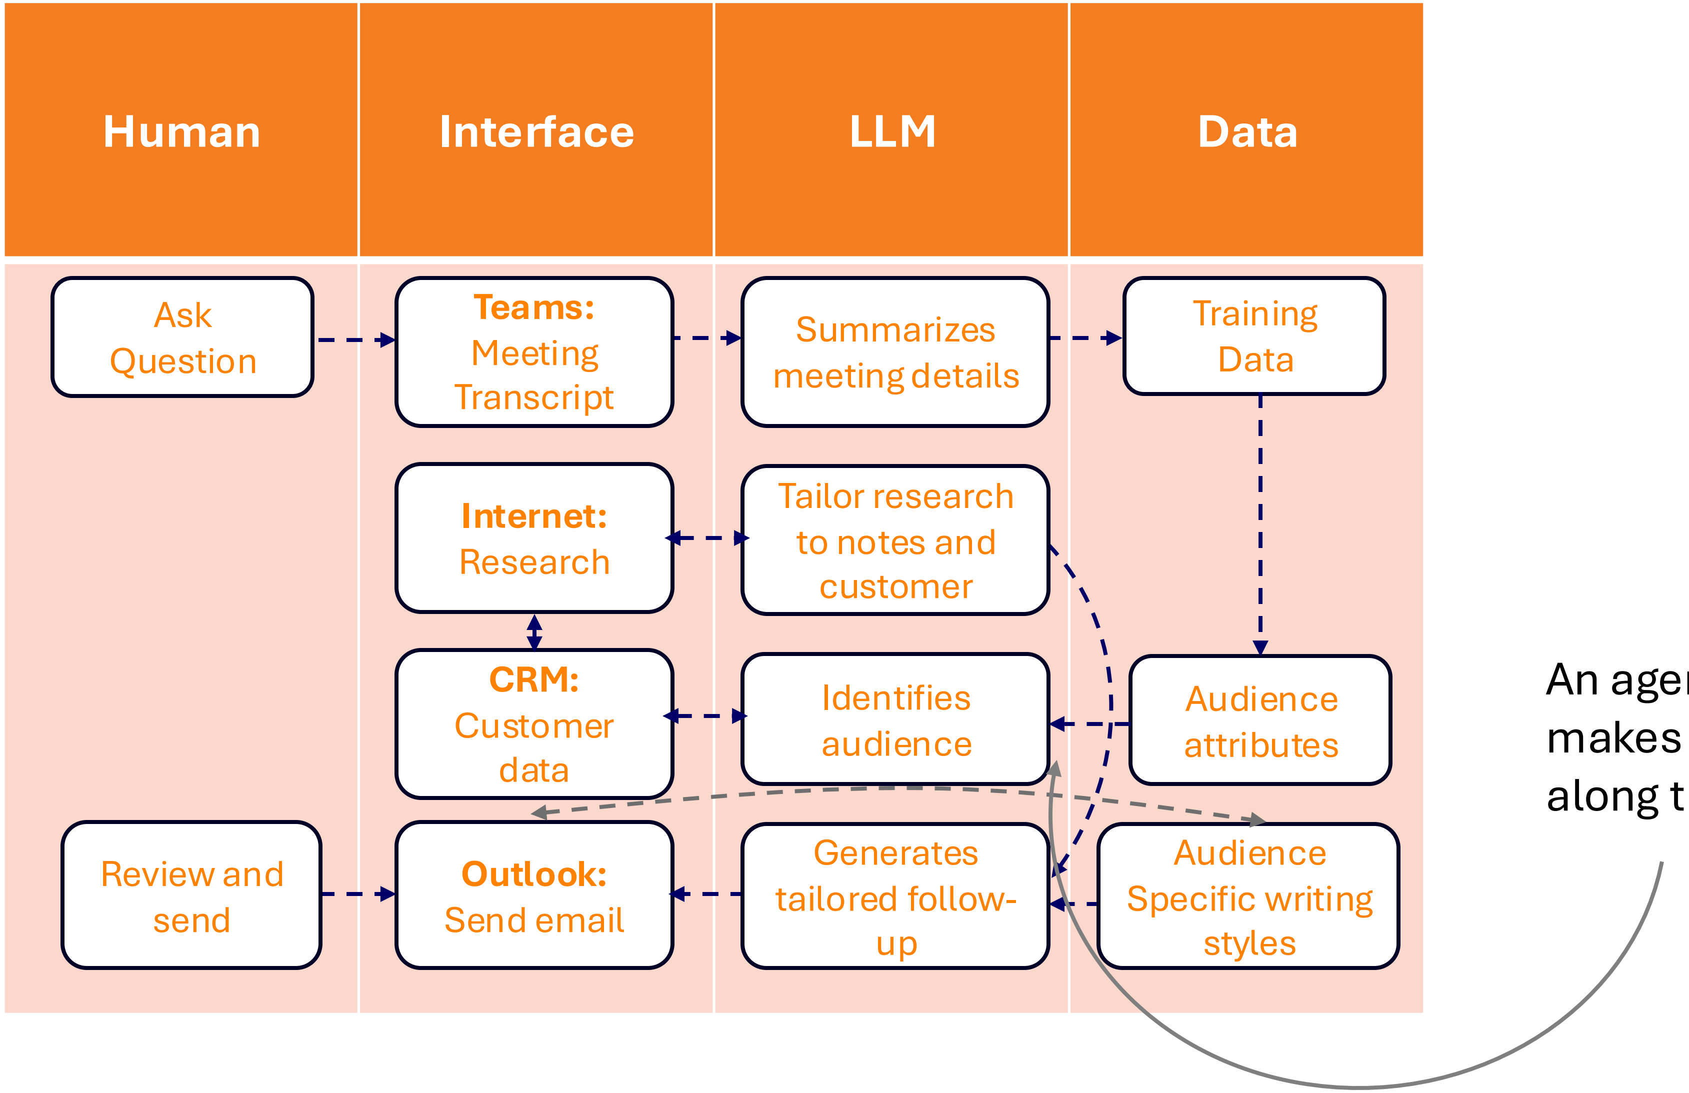

An agent makes choices along the way

Components of Agents

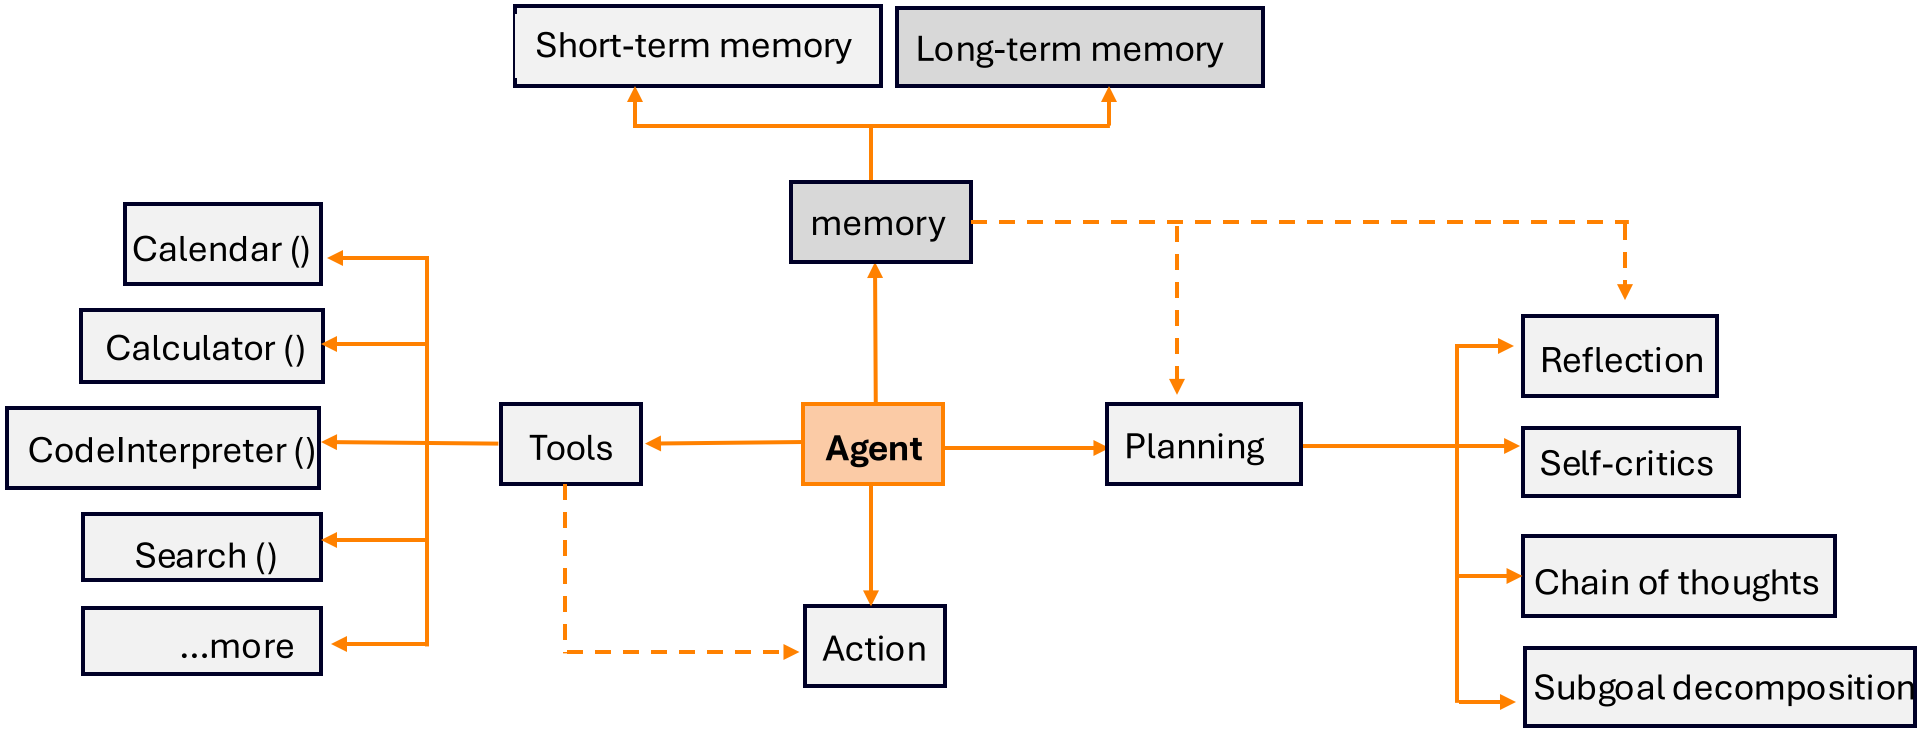

1. **Agent Core:** The Agent Core functions as the command center, orchestrating and harmonizing all system operations. It serves as the primary hub where different components intersect and interact.

2. **Memory Module:** Acting as the system's knowledge repository, the Memory Module captures, preserves, and accesses vital information. This ensures seamless operation across time periods while maintaining contextual awareness and operational continuity.

3. **Tools:** The Tools component encompasses a diverse array of external capabilities and interfaces. These resources empower the agent to execute specialized operations and interact with various platforms and services.

4. **Planning:** Within the Planning Module, complex challenges are dissected and evaluated. This component crafts strategic approaches by analyzing situations, identifying potential solutions, and developing systematic action plans to achieve desired outcomes.

Capabilities of Agents

Advanced Problem Solving

AI agents possess the ability to plan and execute intricate tasks.

Tool Utilization

AI agents can leverage external tools to assess the quality of their output.
This enables them to identify errors and suggest improvements.

Self Reflection and Improvements

AI agents can assess their own outputs, pinpoint issues, and offer constructive feedback. By integrating this feedback and repeating the critique/revision cycle, agents can consistently enhance their performance.

Collaborative Multiagent Framework

Implementing a system where one agent generates outputs, and another provides constructive feedback leads to improved performance through iterative feedback and dialogue.

Anatomy of an Agentic AI System

A Journey into the Heart of Autonomous Intelligence

Agentic AI refers to systems capable of autonomous decision-making and action in pursuit of specific objectives. We have seen this field evolve from a set of theoretical ideas to practical systems shaping industries. If you have been in AI long enough, you know we are standing on the shoulders of giants. Let us take you through this, breaking it down so you see not just what it is, but how it works and why it matters.

Early Software Agents

Hewitt et al. introduced actors as self-contained, interactive objects with internal states, capable of concurrent actions and communication via message-passing.

Intelligent Agents

These agents advanced tactical decision-making but lacked long-term goals, focusing on isolated tasks within workflows.

Autonomous Agents

Modern agents sense and act on their environment to pursue long-term goals, influencing future states for continuous progress.

A Journey into the Heart of Autonomous Intelligence

1990s Impact

Revolutionized distributed computing with object-oriented paradigms and laid foundation for modern microservices

2000s Impact

Enhanced workflow automation and introduced rule-based decision systems in enterprise environments

Present

Impact enabling self-driving systems, smart assistants, and adaptive industrial automation with continuous learning

The Core Pillars: From Perception to Execution

Agentic AI systems function like a well-coordinated orchestra, with each component playing a vital role in achieving harmony.

Perception: The Eyes and Ears of AI

Agentic AI starts by perceiving its environment, leveraging technologies such as computer vision and natural language processing to convert raw data into actionable insights.
Example: In a production line, an AI system identifies defective products by analyzing images in real time, reducing waste and boosting efficiency.

Reasoning: The Brain Behind Decisions

Reasoning enables these systems to detect patterns, make connections, and draw conclusions.
Example: The production line AI uses reasoning to correlate defect patterns with specific machine errors, enabling proactive maintenance and preventing downtime.

Planning: Charting the Path to Success

Planning ensures that Agentic AI systems strategize effectively, optimizing resources and achieving goals within constraints.
Example: The production line AI plans the allocation of inspection resources to areas with higher defect probabilities, ensuring streamlined operations and minimal delays.

The Core Pillars: From Perception to Execution

Learning: Continuous Improvement

Agentic AI systems learn from past experiences and adapt in real time, enhancing their capabilities.
Example: Over time, the production line AI learns to recognize emerging defect patterns, continuously improving its detection accuracy and reducing false positives.

Verification: Ensuring Accuracy Before Action

Verification checks the accuracy and reliability of the AI's reasoning, planning, and learning before execution, ensuring consistency and reducing errors.
Example: The AI verifies its model against new data before applying strategies.

Execution: Transforming Plans into Action

Execution is where AI manifests its strategies, automating tasks with precision and consistency.
Example: The production line AI integrates with systems to automatically remove defective products, ensuring quality control with minimal human intervention.

Understanding the Building Blocks

At the heart of every Agentic AI system lies a delicate interplay of key elements:

Environment: The external world the agent interacts with.

Goals: The desired outcomes driving the agent's actions.

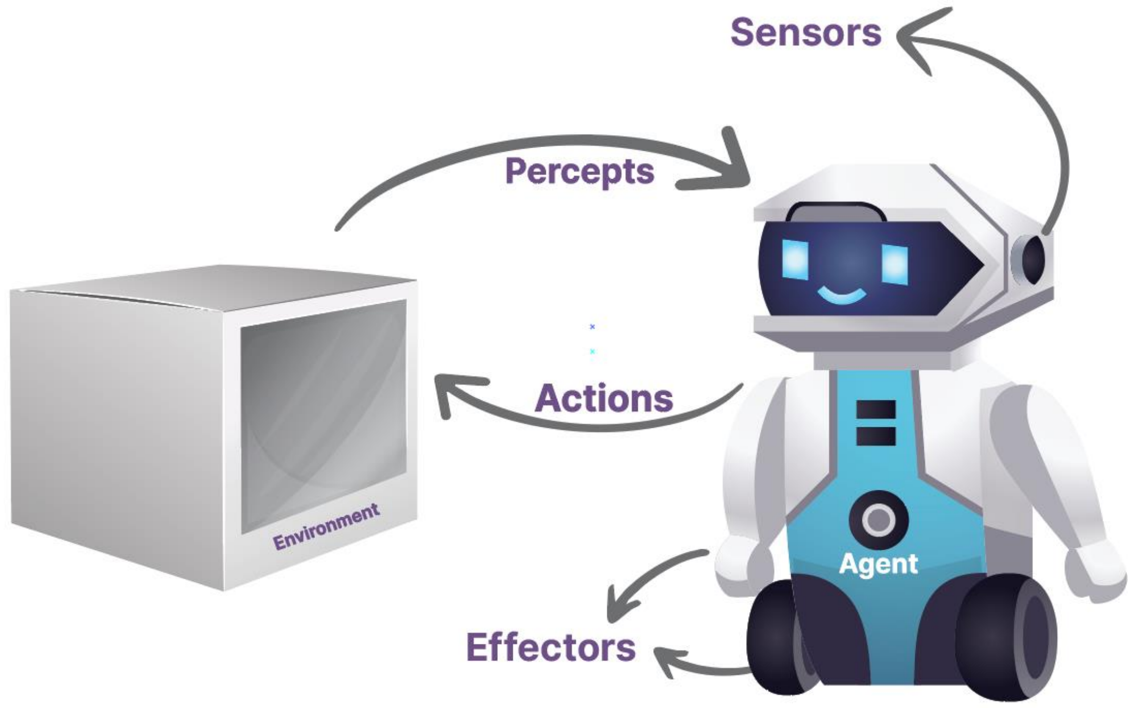

Agent: The entity perceiving and acting on its environment.

Actions: The decisions and behaviors taken to achieve objectives.

The interaction between the **agent**, its **environment**, and its **actions** forms the foundation of agentic AI. The environment can be static or dynamic, and the agent can be **reactive** (responding to immediate stimuli) or **proactive** (planning for future outcomes). The agent's use of **memory** and **learning from interactions** enables **continuous improvement**, leading to more effective and adaptive decision-making. As we progress, we will see how different decision-making models influence how these components work together.

Defining Characteristics of an Agent

Agents have been defined in various ways and have been a point of discussion for some time now, especially after the introduction of Large Language Models (LLMs) and Visual Language Models (VLM). It is useful to have one encompassing definition that is extensible as well as expressive, building on commonly understood concepts from the past. For us, we refer to a few basic principles that embody an agent, and we use them to develop the notion of an "agent object" in the world of LLMs and VLMs. We've distilled the following primary concepts that need to be embodied by an agent object in order for them to be programmed scalably to build robust systems.

1. Autonomy
An agent is an autonomous, interactive, goal-driven entity with its own state, behavior, and decision-making capabilities. It has the capability of self-improvement when it sees that it is unable to meet the performance parameters for reaching its goal.
Example: An Agentic Chatbot adapts its conversational tone and logic based on user feedback, refining responses to better align with the user's intent over time.

2. Reactivity and Proactivity
Agents can be both reactive, meaning they can sense their environment and respond to changes by taking actions, & proactive, meaning they can take initiative based on their goals. Agents can take actions that change the state of their environment, which can ensure their progress toward their goal.
Example: The chatbot detects frustration in a user's tone (reactivity) and proactively offers a detailed guide or direct solution without waiting for further input.

3. Beliefs, Desires, and Intentions (BDI)
A common model used in agent-oriented programming (AOP) is the BDI model, where agents are characterized by their beliefs (information about the world), desires (goals or objectives), and intentions (plans of action).
Example: The chatbot believes the user is seeking customer support, desires to provide the best assistance, and executes a plan to guide the user through relevant troubleshooting steps.

Defining Characteristics of an Agent

4. Social Ability & Communication
Agents have a communication mechanism and can interact with other agents or entities in their environment. This interaction can be highly complex, as it may involve negotiation, coordination, and cooperation. With the advent of LLMs, this communication can be completely based on natural language, which is both human - and machine - understandable.
Example: The chatbot interacts with a scheduling system to book appointments and communicates the details in clear, natural language to the user.

5. Constitution
An agent needs to adhere to some regulations and policies depending on the imperatives of its task and goals. It needs to protect itself from being compromised or destroyed as well as be trusted not to harm other agents sharing the environment in which it is operating.
Example: The chatbot complies with privacy regulations by anonymizing sensitive user data and ensures that it avoids generating harmful or biased responses.

6. Memory
An agent needs to have long-term memory (LTM) of its past interactions and successful past means of completing its task. LTM can help to greatly reduce the amount of computing that an agent must perform to complete a new task by referencing relevant past plans and actions. The memory is also the place where an agent can store human demonstrations it has seen, which can expedite its progress without as rigorous a planning and reasoning loop. An agent also has short-term memory, which is typically its current context signified by prompts and any information available in its context length.
Example: The chatbot recalls previous user preferences (LTM) to suggest relevant solutions and uses the current conversation's context (STM) to address immediate concerns effectively.

Categories and Types of Agentic Systems

Agentic AI systems can be understood through two main lenses:

1. Categories based on Complexity and Application
2. Types based on Functionality and Behavior

Categories of Agentic Systems

Let’s consider the scenario of an AI-based customer support system handling a user inquiry about a product return. Here's how different categories of agentic systems would react and proceed:

Simple Reflex Agents: Respond to immediate triggers without any internal model or predictive capability.
Example: The system instantly replies to the customer with a pre-defined response like, "Please enter your order number to proceed with the return process," whenever it detects a request for returns.

Model-Based Agents: Use internal models to predict and adapt to changing circumstances.
Example: The system checks the customer’s past orders in its database and, based on prior interactions or return history, suggests a return option that aligns with the customer's preferences or past behavior.

Goal-Based Agents: Operate by prioritizing specific objectives and taking actions to achieve them.
Example: The system identifies the goal of processing the return request. It asks the customer for order details, verifies the return policy eligibility, and proceeds with the steps to complete the return process.

Utility-Based Agents: Optimize outcomes by evaluating and balancing various factors.
Example: The system evaluates urgency, return policies, and fees to decide whether to prioritize customer satisfaction with expedited returns or credits, balancing company and customer interests.

Types of Agents: Functional Versatility

Agents are classified based on how they perceive and act within their environment:

Reactive Agents: Quick responders to immediate stimuli, with no long-term planning.
Example: A data anomaly detection system that immediately flags irregular data points in a real-time stream, such as a sudden spike in website traffic indicating a potential attack.

Deliberative Agents: Plan their actions by considering multiple variables and possible outcomes.
Example: A predictive maintenance AI system that analyzes historical sensor data from machinery and schedules maintenance based on the condition, usage patterns, and performance forecasts.

Hybrid Agents: Combine reactivity and deliberation, making them suitable for dynamic and complex scenarios.
Example: A recommendation system that provides personalized content in real time based on user behavior (reactive) while also considering long-term user preferences and trends to improve future suggestions (deliberative).

Types of Atomic Agents

While there are several ways to categorize the different types of atomic agents in AI systems, one simple categorization is based on their abilities and influence.

Foundational Agents

Specialized in core capabilities like planning or verification, foundational agents provide horizontal expertise and support workflows of other agents.

Example: A planner agent in manufacturing optimizes production schedules to reduce downtime.

Workflow Agents

Focused on executing specific high-quality tasks, these agents excel in vertical expertise and often collaborate with foundational agents.

Example: A coding agent in software development generates application prototypes based on design specifications.

Utility Agents

Simple helper agents that connect to tools for basic operations, typically involving low complexity.

Example: A report generation agent in retail consolidates sales data into actionable dashboards for decision-makers.

Types of Atomic Agents

Foundational Agents	Workflow Agents	Utility Agents
Planner Agent	Web Navigation Agent	Report Generation Agent
Verifier Agent	API Agent	Communications Agent
Moderation Agent	Coding Agent	Executor Agent

How Agentic AI Shapes Industries?

Agentic AI systems are transforming industries by deploying specialized solutions tailored to unique challenges. These solutions are often implemented as vertical agents-AI systems designed to excel in specific domains or tasks, offering unparalleled efficiency and precision. They are called "vertical" because they focus deeply on particular industries or applications rather than providing general-purpose capabilities.

Connecting the Dots: Agentic AI and Psychological Theories

Albert Bandura's Social Learning Theory: Individuals actively shape their environment through interactions, fostering learning and agency.

Agentic AI: Represents this concept by allowing systems to learn from data, adapt, and make autonomous decisions.

Jean Piaget's Cognitive Development Theory: Intelligence evolves through distinct stages, shaped by experiences and interactions with the environment.

Agentic AI: Mirrors these stages, progressively enhancing its learning and adapting to complex tasks over time.

Defining Characteristics of an Agent

Industry	Vertical Agents
Healthcare	AI diagnostic tools, Virtual health assistants, Radiology AI, Personalized treatment plans, Patient monitoring systems
Finance	Fraud detection systems, Credit scoring AI, Investment portfolio optimization tools, Risk assessment models, Trading bots
Retail	Personalized shopping assistants, Chatbots for customer support, Inventory management AI, Visual search systems, Recommendation engines
Education	AI tutors for personalized learning, Adaptive learning platforms, Grading and assessment AI, Virtual classroom assistants, Career counseling AI
Manufacturing	Predictive maintenance systems, Supply chain optimization agents, Production line robots, Quality control systems, Automated inventory systems
Legal	Contract review and analysis tools, Legal research assistants, Case prediction AI, Document automation systems
Transportation	Traffic optimization systems, Fleet management AI, Autonomous vehicle navigation, Route planning and logistics optimization
Biopharma & sciences	Drug discovery AI, Clinical trial optimization tools, Predictions, Genomic Data Analysis, Real-Time Monitoring and Decision Support
Medical Devices	Predictive diagnostics, Remote monitoring systems, AI for surgical planning, Device performance monitoring
Construction	Project management AI, Autonomous construction, Building energy optimization, Safety monitoring systems

Agentic AI Tech Stack

Agentic AI Tech Stack

Framework	Best Use For
crewai	Task Delegation, Team Co-ordination, Role based Tasks, Human – AI Collaboration
AG AutoGen	Multi Agent Chat, Complex Reasoning, Group Problem Solving, Autonomous Interactions
LangChain	Search Applications, Document Processing
LlamaIndex	Document Indexing, Structured Data Access, Customer Data Connectors
LangGraph	Document Analysis, Multi-step Workflows, State Management

Orchestrating Agentic AI Systems

Orchestrator in Action

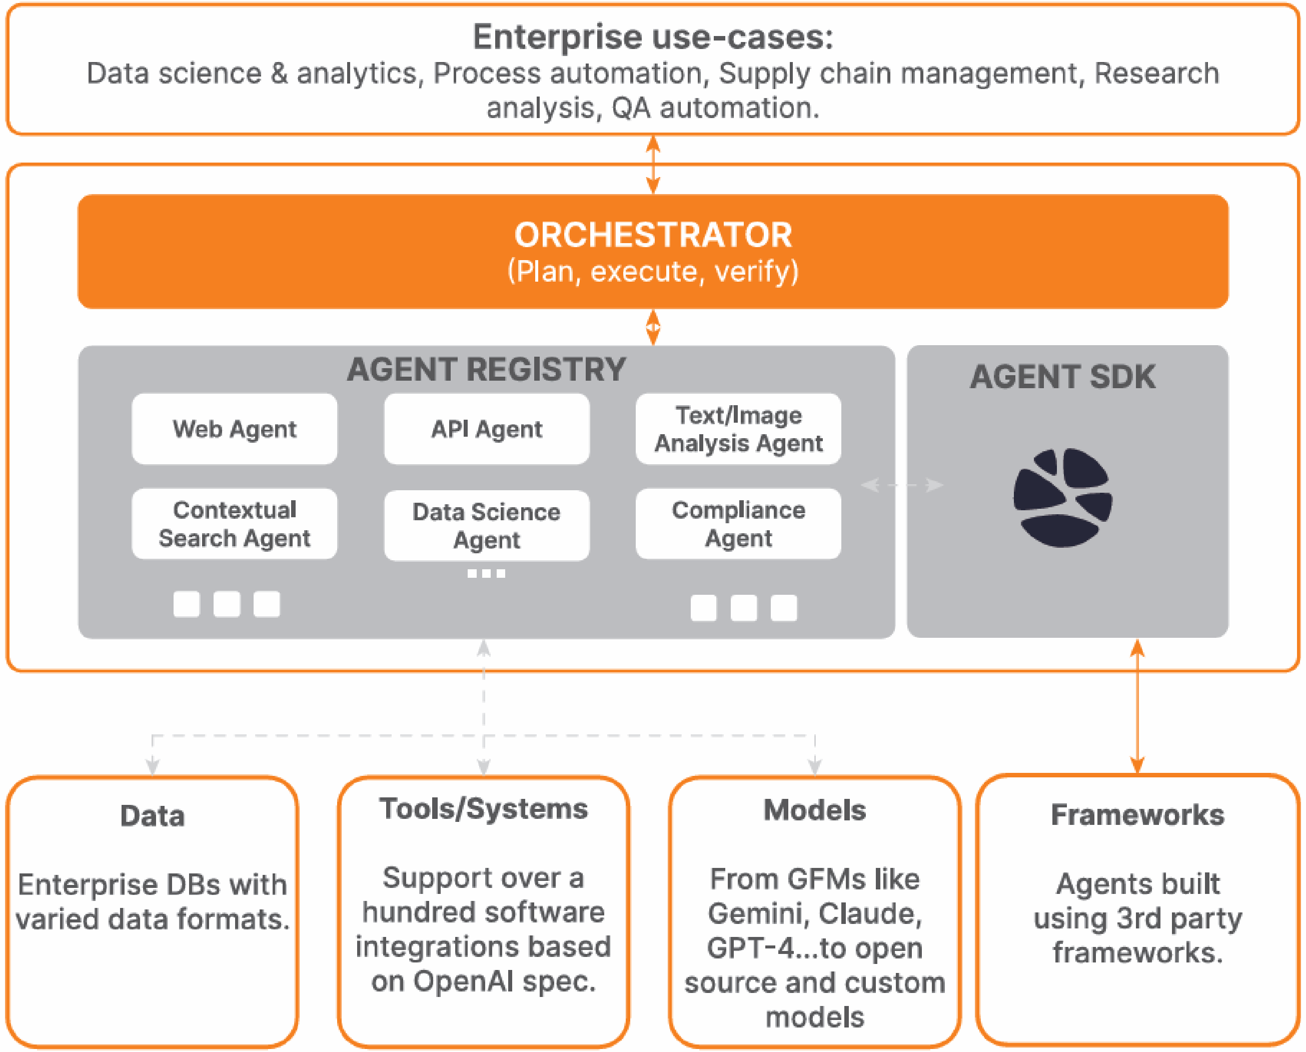

Building on Fine-tuning and Prompt-tuning of GFMs and Custom LLMs

Challenges of Orchestrating Complex Agentic Systems

Communication and Coordination

- Ensuring seamless communication among diverse agents can be difficult due to differences in functionality and sophistication.
- Interoperability issues arise when agents are designed with varying technologies or standards.
- Developing a uniform protocol for interaction is often essential, but it can be complex to implement and requires collaboration within the community.

Conflict Management

- Agents operate autonomously, which can lead to conflicts arising from overlapping objectives, competition for resources, or differing priorities.
- Resolving these conflicts without human intervention requires effective automated mechanisms.
- Choosing the Right Agent: When multiple agents with the desired capabilities are available, a selection process is needed to ensure the most suitable agent is chosen, based on context, efficiency, and resource availability.

Scalability

- Orchestrating large-scale systems with hundreds or thousands of agents presents logistical challenges.
- Robust infrastructure and efficient resource allocation algorithms are necessary to support dynamic scaling.
- The system must be designed to accommodate the addition or removal of agents without disrupting performance.

Reliability and Fault Tolerance

- Maintaining system reliability is critical, particularly when dealing with partial failures.
- A failure in one agent should not jeopardize the overall system's functionality or performance, necessitating advanced fault tolerance mechanisms.

These challenges highlight the complexity involved in orchestrating multi-agent systems, requiring a careful balance of technical strategies and collaborative efforts.

Potential Agentic AI Use Cases

Agentic AI Use cases

Retail

- **Personalized Shopping Experience:** AI agents recommend products based on customer preferences and past behaviors, enhancing the shopping experience and boosting sales.
- **Inventory Management:** AI can autonomously track stock levels, predict demand, and reorder products, minimizing stockouts and excess inventory.

Manufacturing

- **Predictive Maintenance:** AI agents monitor equipment health, predict potential failures, and schedule maintenance, reducing downtime and repair costs.
- **Supply Chain Optimization:** AI manages inventory, tracks shipments, and adjusts delivery routes in real-time, improving operational efficiency and reducing costs.

Healthcare

- **Patient Monitoring:** AI agents track patient vitals and alert healthcare providers about critical changes, enabling faster response times and better care.
- **Personalized Treatment Plans:** AI analyzes patient data to suggest tailored treatment options, improving patient outcomes and treatment efficiency.

Agentic AI Use cases

Biosciences

- **Drug Discovery:** AI agents autonomously sift through vast datasets to identify potential drug candidates, speeding up the research process.
- **Gene Editing:** AI simulates the effects of gene edits, assisting in precise genetic research and therapeutic development.

Pharmaceuticals

- **Clinical Trial Optimization:** AI selects trial participants and optimizes trial designs, improving recruitment rates and accelerating the drug development process.
- **Pharmacovigilance:** AI monitors and analyzes data for drug side effects, helping to ensure drug safety and compliance with regulations.

Finance & Insurance

- **Fraud Prevention & Risk Assessment:** AI detects fraud in real-time and automates risk analysis, enhancing financial security and credit evaluations.
- **Smart Automation:** Automation streamlines claims processing and personalizes recommendations, improving efficiency and customer satisfaction.

Agentic AI Use cases

Education

• **Personalized Learning:** AI adapts learning materials to suit each student's progress and style, helping improve learning outcomes and engagement.
• **Automated Grading:** AI agents handle grading of assignments and exams, saving educators time and ensuring consistency in evaluation.

Telecommunications

• **Network Optimization:** AI monitors network performance in real-time, detecting issues and automatically adjusting resources to maintain service quality.
• **Customer Support Automation:** AI chatbots provide real-time support, resolving customer issues and reducing wait times for service requests.

Construction

• **Project Scheduling and Management:** AI agents optimize construction timelines, track project milestones, and predict potential delays, ensuring timely project completion.
• **Site Monitoring and Safety:** AI can monitor construction sites in real-time, identifying safety hazards and ensuring compliance with safety regulations, reducing accidents.

Case Studies

AI-Powered Loan Eligibility Automation

About Customer

Customer Profile: A leading finance technology company aiming to enhance customer engagement and streamline loan eligibility processes.

Objective: To automate and accelerate the home loan eligibility check process, improve customer experience, and reduce operational inefficiencies by integrating GenAI-powered solutions and agent support.

Problem Statement

- **Lengthy Eligibility Checks:** Prolonged waiting times frustrated customers and caused disengagement.
- **Complex Document Access:** Manual processes delayed document uploads, retrieval, and verification.
- **High Drop-off Rates:** Lack of clear communication and slow responses led to application abandonment.
- **Inconsistent Customer Engagement:** Delayed, irrelevant interactions eroded trust and satisfaction.
- **Manual Verification and Redundancy:** Labor-intensive processes hindered scalability and efficiency.

Our Solution

To address the challenge, we developed a system that leverages GenAI to automate loan eligibility and amortization checks. Complex cases are intelligently flagged by AI and escalated to human agents for verification, ensuring accuracy and reliability:

•Automated Loan Checks
Our system automates loan eligibility and amortization checks using GenAI-powered calculations, ensuring accurate and efficient decision-making.

•Hybrid Onboarding
Secure SMS OTP and WhatsApp processes streamline initial customer verification, with agents stepping in for personalized support whenever needed.

•Instant Query Resolution
Conversation Prompt Management provides swift and precise responses to standard queries, while complex issues are escalated to human agents for in-depth handling.

•Real-Time Document Processing
The system automates document uploads and processing, with agents ensuring compliance and addressing flagged exceptions for error reduction.

•Seamless Workflow Integration
The Orchestrator and Scheduler manage smooth collaboration between AI systems and human agents, supported by secure GCP VPC and firewalls for data integrity.

Impact

- **Improved Processing Efficiency**: Automated eligibility checks and agent-assisted verifications cut waiting times, enabling faster and more accurate loan processing with minimal manual intervention.
- **Enhanced User Engagement**: AI-driven communication, complemented by human support, reduces drop-offs and increases customer satisfaction, ensuring secure and seamless verification for a positive user experience.
- **Scalable and Reliable Operations**: The blend of AI and human oversight ensures efficient handling of increased application volumes while maintaining consistent reliability and performance.

Activity

Activity

Build Agentic AI code for Research agent and Blog Agent using Crew AI and Deepseek-R1