

Combining AI Approaches and Agentic AI

The AI Spectrum — From Classical AI to Agentic AI

Leadership with AI – Rajendra M Sonar, IIT Bombay 



What we are going to learn

Agentic AI: concepts, building blocks, and architectures

1. How Agentic AI extends GenAI

From prompt-response to goal-directed systems

2. Core agent building blocks

Models, tools, retrieval, memory, state, and execution loop

Agentic AI

3. Architectures and orchestration

Single-agent, multi-agent, sequential, concurrent, handoff, and group patterns

4. Fit in the AI spectrum

When rules, ML, GenAI, or agents are the right choice

5. Technical limits

Failure modes, latency, cost, observability, and reliability challenges

Revisit: operationalising AI: Automation of seamless end-to-end operations: Example




Hi, my customer id is 123456 and have subscribed to PlanX, I want to change to PlanY, will be grateful if you help me.

Sure, pls wait while I check how best I can help you.

- 1 Check credentials of customer if not logged in. If not, ask to Login in.
- 2 Understand what customer is looking for by extracting information from text
- 3 Check information provided
- 4 Check action, next step will be based on action, action is: change the subscription

- 5 Check eligibility of customer for product PlanY, ask any KYC is required, deposit etc.
- 6 Compliance check and risk assessment
- 7 If eligible, then check there any pending dues from PlanX
- 8 Unsubscribe PlanX and subscribe to PlanY
- 9 Inform to the customer




Dr Akshay, we have successfully changed your plan from PlanX to PlanY. Appreciate your feedback on our service [Link]

- 10 Get feedback

- 1 Operational task.
- 2 Gen AI Task: NLP: Named Entity Recognition with action. Decision Making: what needs to be done to change subscription. Ask missing information if not able to decide.
- 3 Enterprise systems: subscription management: check existing subscription
- 4 Decision task

- 5 AI + Operational Tasks: CRM system: KYC status, order management system: compliance, requirements (documents). Invoke workflow system, RPA.
- 6 Gen AI Task: Compliance check and risk assessment
- 7 Billing system: Status of payments etc. If dues, invoke payment processing system
- 8 Order management and billing system: product configuration and subscription
- 9 Workflow management
- 10 CRM System: Get feedback

Integrated Systems For End-To-End Processing



Operationalising AI: Automation of seamless end-to-end operations: Example



How can I help you?

Planning to go on lien

Welcome Prof. Sonar, just checked your service records, you are eligible to take lien for 2 years, want to go ahead?

Sure

How long you want to take lien?

2 years

Would you like to keep quarters during lien period?

Yes

Where you would like to go on lien?

In a private company

....

....

Pls download agreement pdf, take on legal paper and submit for further approval !



Rule based system:
Leave Rules

HR System |
Database

Rule-based
System



Gen AI: Extract
rules, get them
validated by
experts

HR System |
Database

HR System |
Database

Rule-based
System:
Housing Rules


Faculty Handbook

A Survival Guide for IITB Faculty




11. Lien

A faculty member may request for keeping lien on his/her post for accepting a job either in India or abroad. Such jobs may be for assignments in private or public sector undertakings. (The Govt. of India terminology for such lien is Foreign Service, which is used whenever a Govt. servant takes up an assignment in a non-Government organization.). To be eligible for lien a faculty member must have put in at least 5 years of service after confirmation. The period of lien can be one year at a stretch or two years if the period of service is over 10 years. (Lien period can be for a period of Five years for those appointed as Directors or CEOs in a Government organization or a PSU.)



Operationalising AI: Automation of technical task using co-pilots

How to format a hard drive for Windows XP?



How can I help you?

I want to format hard drive.

You have 2 HDs in this desktop, C: and D: which one you want to format? Here is summary of data:

C: Operating system, various software, etc.

D: Mostly your data files, PPTs, workbooks, database files, back up files etc.

D

Hey, hold on, I am not sure you have taken back ups of files, cross check your backup status. Based on my knowledge following files are important and you have not backed up since last 1 month or so. Do you want to proceed? Or you want me to sync your files to one drive?

Proceed

Do you want to create partitions (each partition will be treated as a separate drive)?

Yes

....

....


Just wait...while I format D: drive for you!

...

D: Drive is formatted successfully!

Where Agentic AI Sits in the AI Spectrum

Agentic AI is not a replacement for earlier AI approaches; it extends them when multi-step, tool-using action is needed.



Core idea: more autonomy + more context + more action


Human Agents

What does a human agent do?




A human agent is a person who understands **requirements**, background information, objectives and constraints of client, **makes decisions**, **performs tasks**, **remembers and tracks things** and **coordinates** and **interacts** with client and with others in the ecosystem, bringing uniquely human qualities like empathy, creativity, and critical thinking to complex or emotionally charged situations.

Also **monitors the happenings** and **proactively suggests** the client about the opportunity and risks.



Think of AI agent as system which automates job of a human agent! For example, insurance AI agent

Example: Insurance agent helping client



What are core tasks of agent (for each client)?

Set the Goal

- Understand client and what client is looking for
- What constraints, requirements and objectives client has
- Evaluates which plan suits to the client

Plan the things

- Based on the goal plan the activities

Execute plan

- Use digital tools, various channels
- Capture interactions
- Collect documents
- Use excel sheets
- Use applications
- Maintain the details and state (**memory**)


Adapt to any changes

- There may be issues with client or client wants to change some plans etc.

Hand hold

- Observe the activities being done
- Help customer to carry out activities smoothly

How humans collaborate to solve a problem: Example



What Makes AI “Agentic”?

An agentic system is defined by behavior, not by a specific vendor tool.

Goal-directed

Plans steps

Uses tools

Adapts through feedback

Maintains state

Acts with bounded autonomy

Revisit: Classical Intelligent Agents!

Intelligent Agents: Theory and Practice

Michael Wooldridge

Department of Computing
Manchester Metropolitan University
Chester Street, Manchester M1 5GD
United Kingdom
M.Wooldridge@doc.mmu.ac.uk

Nicholas R. Jennings

Department of Electronic Engineering
Queen Mary & Westfield College
Mile End Road, London E1 4NS
United Kingdom
N.R.Jennings@qmw.ac.uk

Submitted to *Knowledge Engineering Review*, October 1994.
Revised January 1995.

Michael Wooldridge is a prominent computer scientist known for his work on intelligent agents, particularly in the field of multi-agent systems. He defines intelligent agents as computer systems situated in an environment and capable of flexible, autonomous action to meet design objectives, according to Jennings and Wooldridge (1998). His research focuses on the theoretical foundations, architectures, and languages for building intelligent agents, often employing techniques from computational logic, game theory, and social choice theory

AI agents is not a new concept for AI researchers and practitioners. Lot of research papers, prototype systems and applications were developed in 1990s and post that which involved agents. Only they were called intelligent agents. “Intelligent” is now changed with “AI”. The fundamental difference was, old agents used classical AI most of the time, now replaced with LLM backed agents. They did not have Gen AI to support prompt, context engineering and LLMs

Revisit: Classical Intelligent Agents: an example.

Intelligent Agent-Based Expert System Architecture for Generating Work Plan Marshalling Station

Abstract - The present approach for generating work plan is difficult to meet the dynamic and intelligent requirement, and it has not had the function of assistant decision yet. An agent that can simulate the reasoning process of human has the properties such as adaptability and intelligence, which can be used to solve the problems in making plan process. This paper has presented the initial architecture for an intelligent agent-based expert system for generating of work plan in marshalling station known as TAES-OPMY. IAES-OPMY can support the station controller's decisions on generating plan and in turn improve the validity, encash ratio and automatization of plan. Furthermore, it can speed the turnover of wagons and bring great profits.




Fig. 2. Technique work operation process in marshalling station



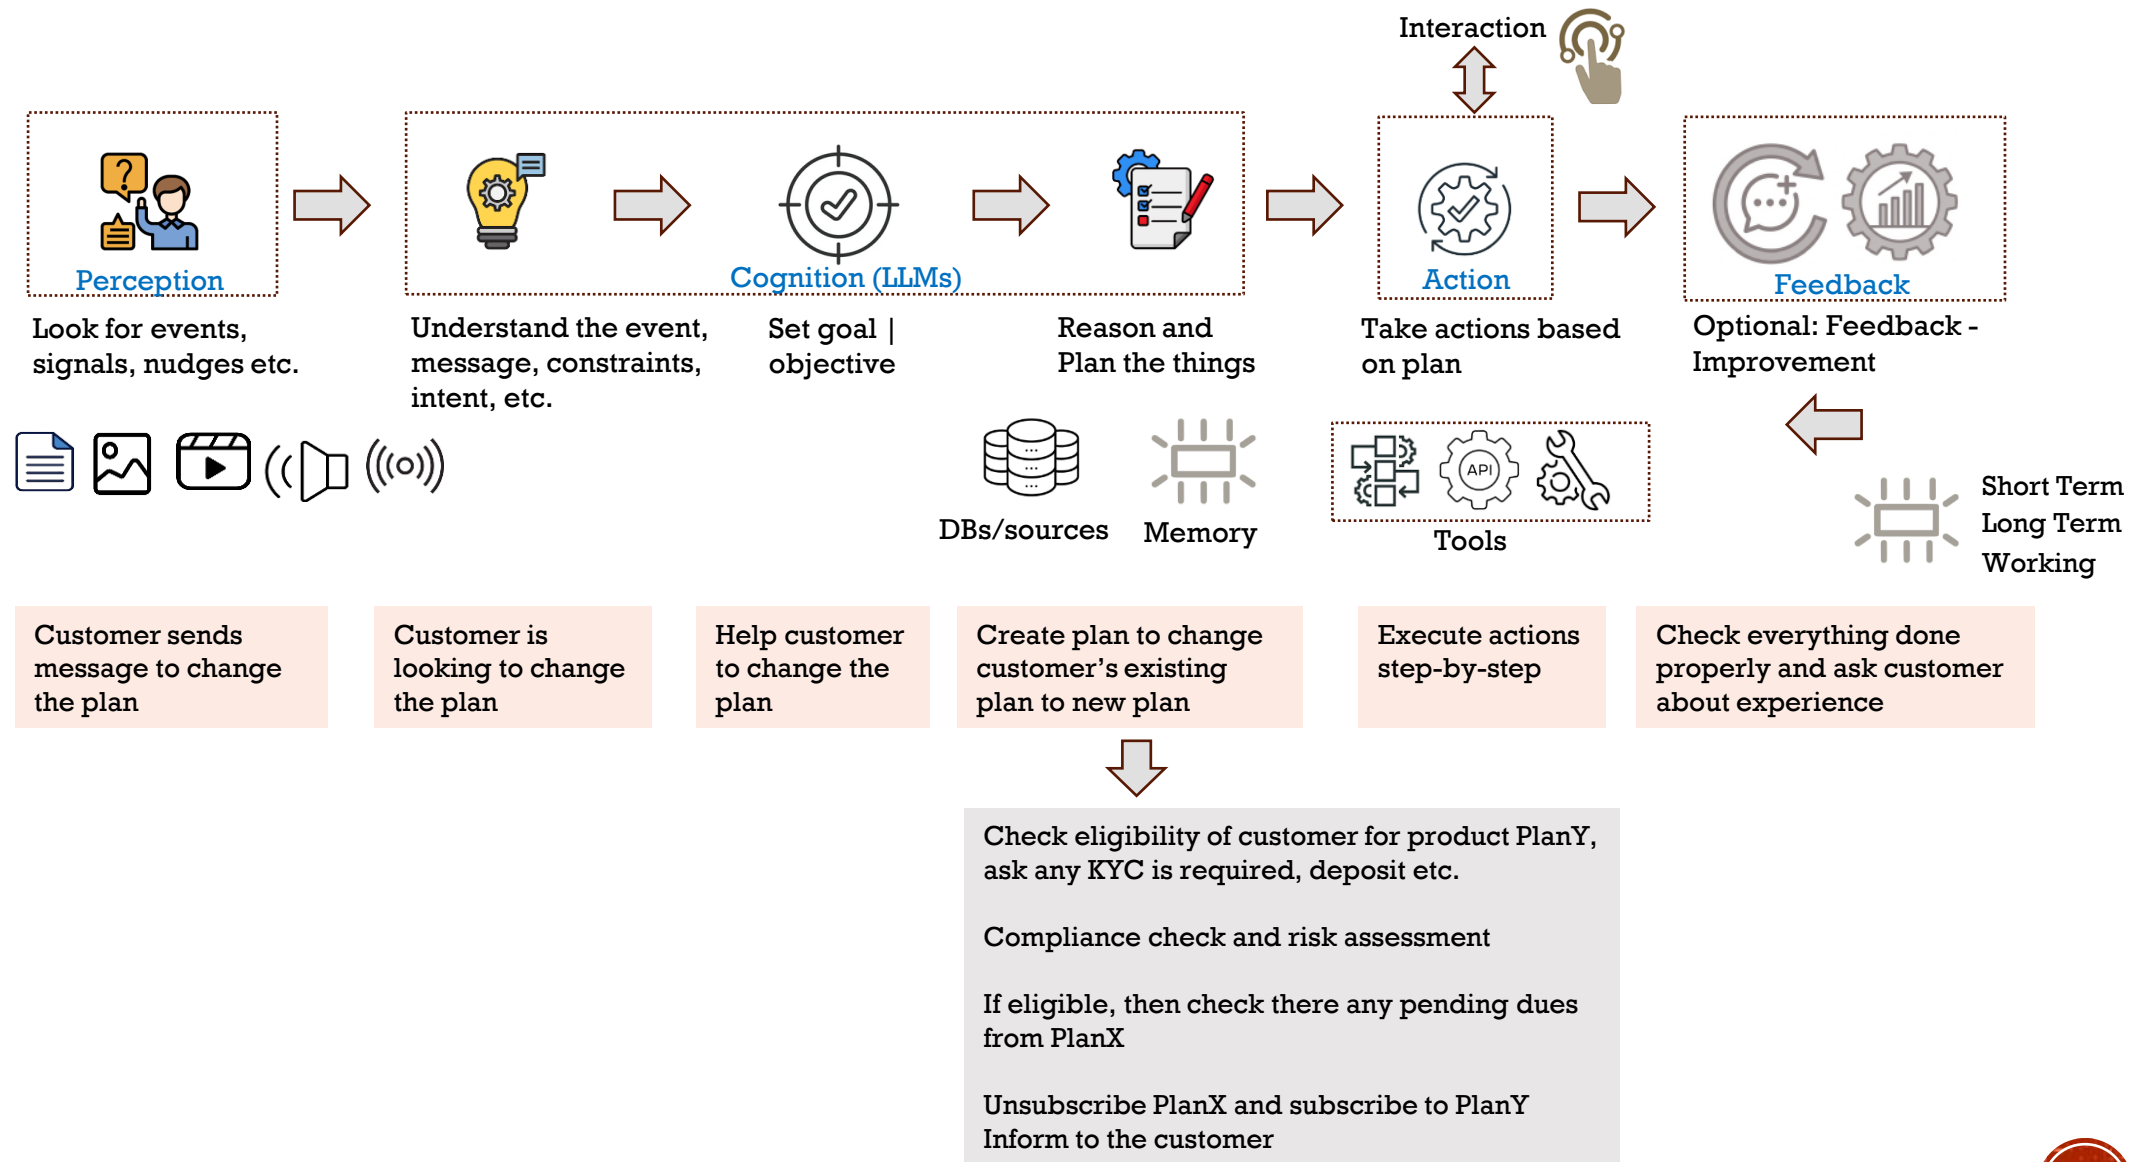


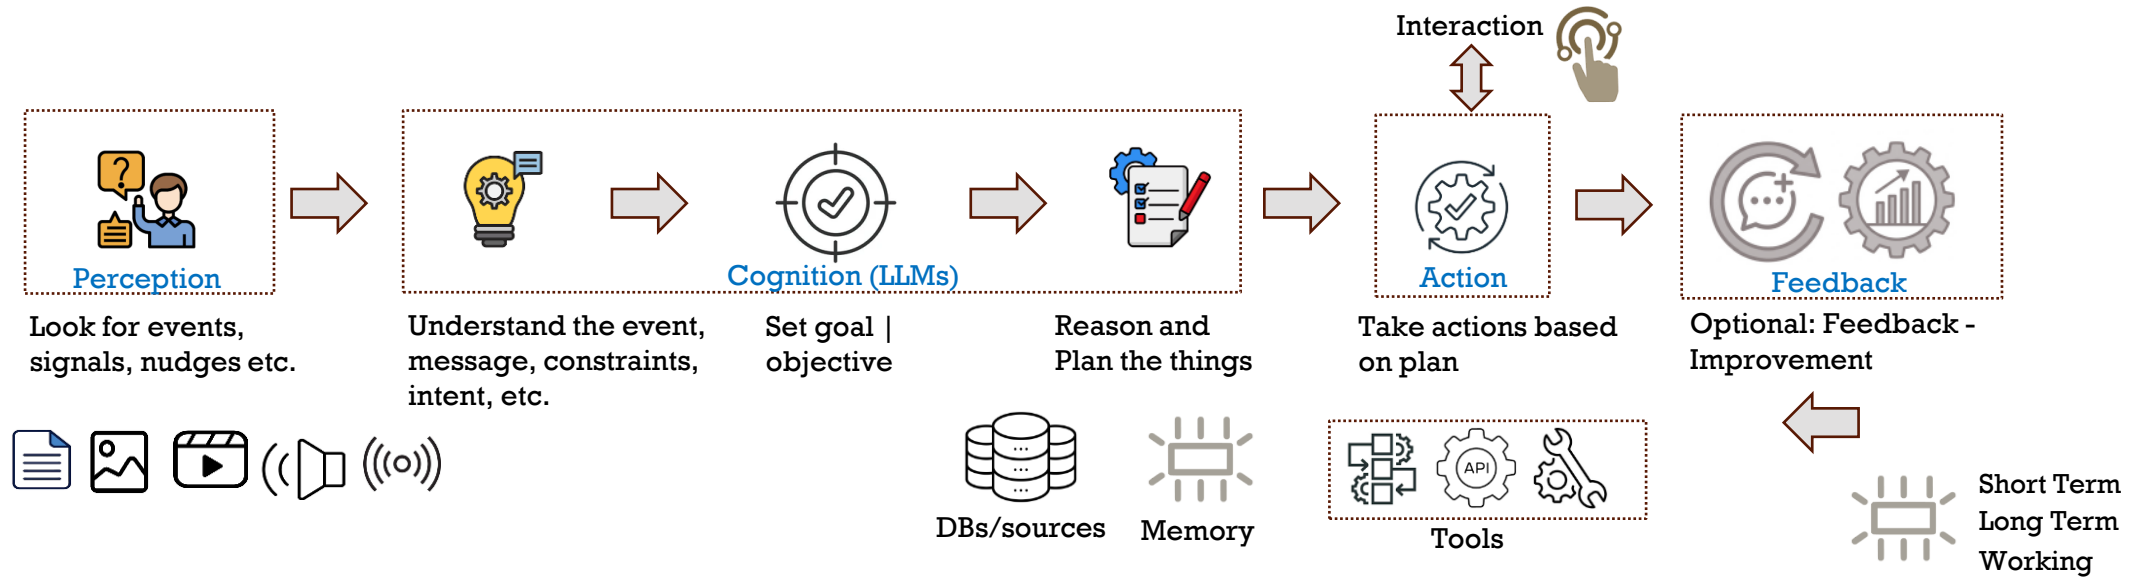
Fig. 1. IAES architecture

This shows agentic AI kind of frameworks existed earlier too. As mentioned earlier, agents were backed by conventional AI.

Operationalising AI: Automation of seamless end-to-end operations

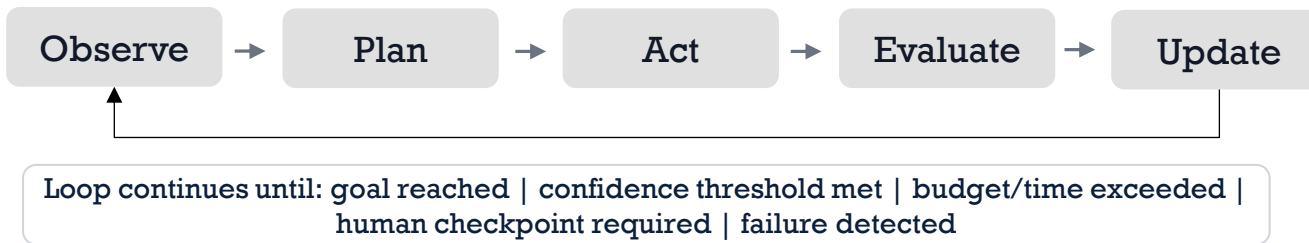


Operationalising AI: Automation of seamless end-to-end operations



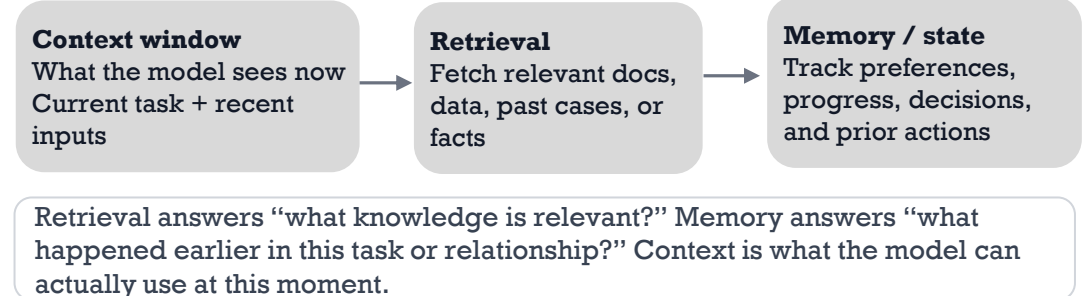
The Agent Execution Loop

Agents repeatedly interpret context, decide the next step, take action, observe the result, and update state.




Memory, Retrieval, and Context

Agents need more than a prompt: they need access to relevant knowledge and enough state to continue a multi-step task.



Agent to tools communication



MCP (model context protocol) is a standard way for AI agents to securely discover, understand, and use tools, data, and context at runtime.

Tools and Actions: How Agents Interact with Systems

Without tools, an agent mostly thinks and writes. With tools, it can retrieve, calculate, update, trigger, and coordinate.

Search / retrieval

DB Operation

API call

Workflow trigger

File / document action

Code / computation

UI/UX

AI/ML Task

Why MCP Exists (The Core Problem It Solves)

Before MCP:

Every AI agent integration was **custom**
Tool schemas were **hard-coded**

Context was:

- fragmented
- vendor-specific
- brittle

Scaling agents across tools was painful

In short:

Agents could think, but didn't know their environment reliably.

MCP answers 4 critical questions for an agent:

1. What tools exist?
2. What data can I access?
3. What is the schema / contract?
4. What am I allowed to do?



Without MCP, agents:

guess, hallucinate APIs, require manual wiring

With MCP, agents:

discover capabilities
reason over *real* affordances
act safely

Agentic AI



Agentic AI is the **automation** of goal-driven workflow, where Gen AI backed systems decide, plan, and execute both cognitive and operational tasks - often collaboratively - to achieve an objective.


Layer	Purpose
Foundation Models	Reason & plan
Agent Frameworks	Coordinate behavior
MCP	Discover context & capabilities
Tools & APIs	Execute actions
Memory	Learn & personalize
Orchestration	Control execution
Governance	Ensure safety & trust
Infrastructure	Scale reliably

Agentic AI

Common Agent Orchestration Patterns

Orchestration defines how agents or steps are coordinated to complete a workflow.

- 1 Sequential**
Steps happen in a fixed order
- 2 Concurrent**
Independent steps run in parallel
- 3 Handoff**
One agent passes task to another
- 4 Group chat**
Agents deliberate or review together
- 5 Supervisor**
Coordinator assigns, checks, and consolidates



Source pattern names are commonly used in current architecture guidance.

Design Patterns for Agentic Systems

Design patterns combine building blocks into reusable architectures.

Retrieval-augmented agent

Uses documents/data to ground actions

Planning agent

Decomposes a goal into steps before execution

Supervisor pattern

Coordinator assigns work and consolidates outputs

Tool-using agent

Calls APIs, databases, calculators, or enterprise tools

Workflow agent

Follows a controlled process with checkpoints

Patterns are not mutually exclusive. A real system may be a retrieval-augmented, tool-using workflow agent with a supervisor.

Agentic AI

When Is Agentic Behavior Actually Needed?

Not every AI use-case needs an agent. Use agents when the task needs autonomy, tools, and multi-step adaptation.

Usually not agentic

A simple answer is enough
Use chatbot / GenAI response

Prediction is enough
Use ML model

Clear rules dominate
Use rule-based / expert system

Fixed steps dominate
Use workflow automation

Agentic fit

Multi-step goal
Task must be decomposed and progressed

Tool use required
External systems/data must be accessed or updated

State matters
Progress, context, and previous actions affect next step

Adaptive path
The next action depends on intermediate results

Technical Limits and Failure Modes

Agentic systems are powerful but technically harder to make reliable than one-shot GenAI interactions.

Error propagation

A wrong early decision can distort later steps

Cost and latency

Multi-step loops may call models/tools many times

Tool misuse

Incorrect tool, wrong arguments, or unsafe action

Observability difficulty

Harder to debug why a path was chosen

Memory problems

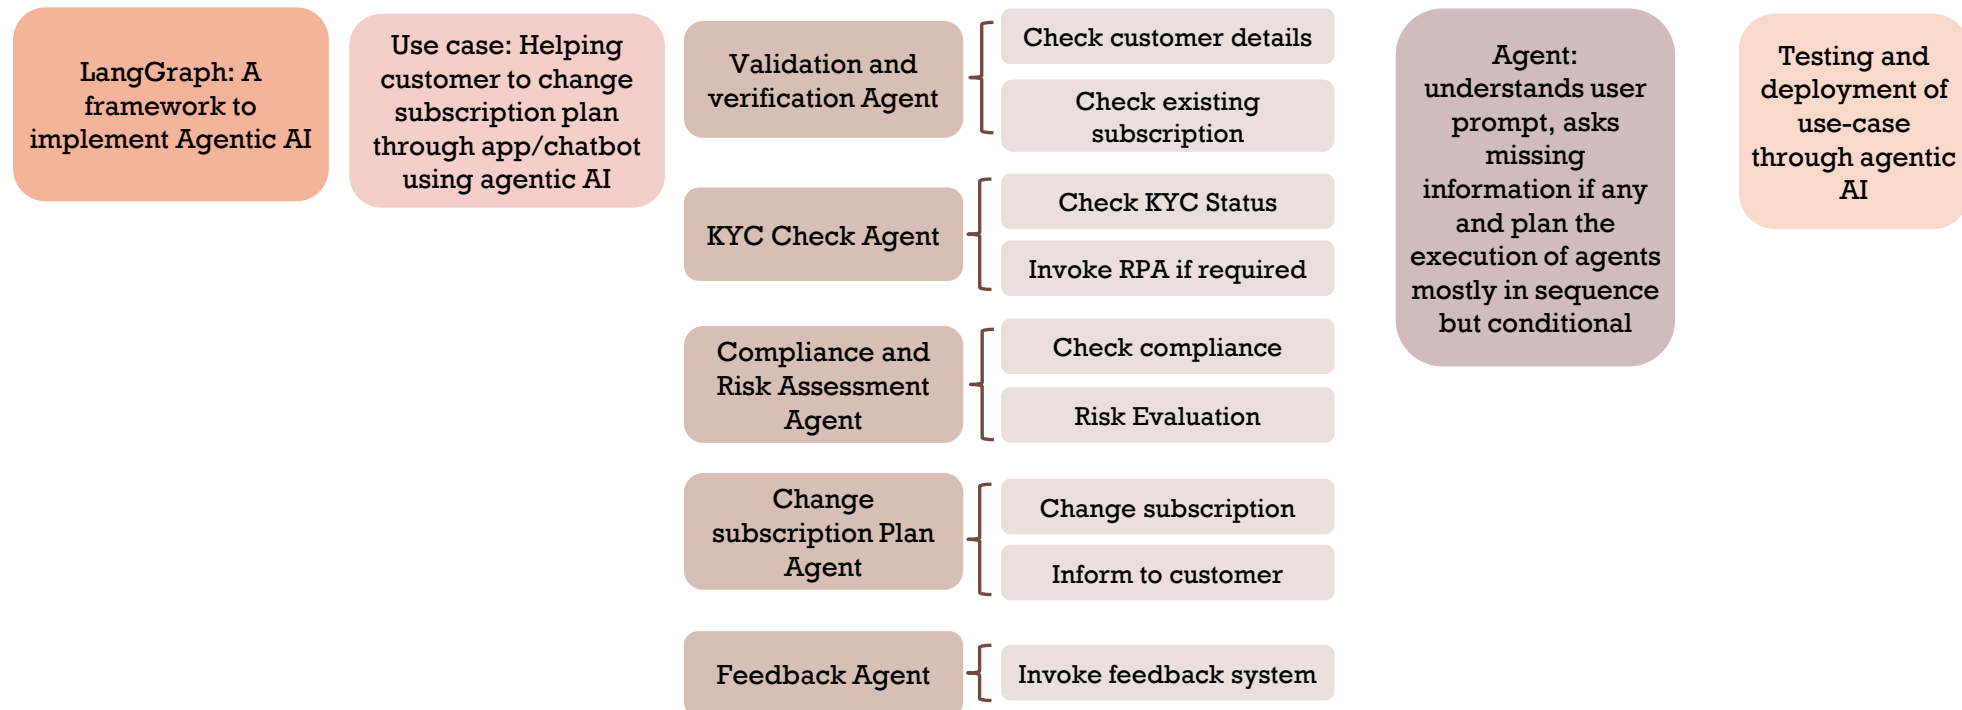

Stale, missing, or misleading state

Complexity explosion


Multi-agent systems can become hard to coordinate

Conceptual rule: more autonomy requires more control, observability, and validation

Agentic AI



Where AI plays role?



AI can play role almost at every level from operational to strategic and almost in all verticals, simply, *wherever human intelligence is required*, AI can come into play! It can effectively address biological limitations of natural species. Often combined with other digital technologies to exploit the value (it is like adding intelligence and making it intelligent) e.g. IOT+AI helps in creating digital twins and predictive maintenance of machines/equipment/plants etc. rather than using in silos.



Effective and optimal **management, scheduling, tracking and monitoring of assets** including **predictive maintenance, defect/failure detection** equipment/devices/plants, with reduced downtime...e.g. may using digital twins or embedded intelligence.



Supplement humans in **decision making** by providing insights from data. Analysing huge datasets, continuously **monitor** and follow **objects, activities** and alerting deviations, suggesting improvements, risks, frauds and telling business opportunities



Dynamic pricing, **co-creating products**, reconfiguring products, leveraging long tail, reconfiguring business models backed by insights from AI



Productivity enhancements: Using AI tools, co-pilots to expedite work, generate contents at ease, extract information, change format, integrate data sources etc.



Domain **knowledge automation**: Reduce information overload, moving from information delivery to actionable knowledge delivery (experiencing, tell me what it means to me!)


Automation of human expertise e.g. compliance, auditing, rule-books etc.



Effective and optimal **management, scheduling, tracking and monitoring of assets** including **predictive maintenance, defect/failure detection** equipment/devices/plants, with reduced downtime...e.g. may using digital twins or embedded intelligence.



Enhance end-users/customers experiences: understanding, hand-holding and assisting end users intelligently! offering them **highly personalized** and **WoW experiences, tracking and understanding customer journey**,



Leveraging unstructured data sources: Understanding texts/documents, translating, generating contents, understanding sentiment/topics etc.

Reliability Ladder

Compiled Binary



Managed Runtime



Interpreted Language



AI Agent



[You're Learning AI Agents Wrong. Here's How To Learn](#)
[Most "AI Engineers" Are Frauds \(and I Can Prove It\)](#)
[What the AI Industry won't tell you](#)
[Big Tech is Killing Agentic AI](#)

Key takeaways

Agentic AI moves AI from answering prompts to pursuing goals through planning, tool use, memory, and multi-step action.

Agentic AI extends GenAI from answering to acting

Agents pursue goals through plans, tools, state, and loops

The agent is more than the model

Reasoning core + tools + retrieval + memory + guardrails + orchestration

Agentic AI

Architecture matters

Single-agent and multi-agent designs create different capabilities and complexity.

Agents work best for multi-step, adaptive, tool-using tasks

Do not use agents where simple rules, ML, workflow, or GenAI are enough.

Architecture matters

More autonomy means more validation, observability, and control