

ACCA7026

AI Agents and Agentic AI

Dr. Seena Joseph

Last time: what makes AI “agentic”



PILLAR 1

Multi-step reasoning

Break a problem into sub-problems, explore paths, revisit alternatives: Chain and Tree of Thought.



PILLAR 2

Persistent memory

Remembers beyond the context window, so earlier sessions and experience still count.



PILLAR 3

Tool use

Acts through APIs, code execution, databases and knowledge bases, not just text.




PILLAR 4

Active interaction

Engages with users, with environments such as browsers or robots, and with other agents.

Today we open the box: what is each pillar made of?

The architecture at a glance




IN OUR EXAMPLE...

New articles and lab charts come in (1). Collector, analyser and citation roles split the work (2). It remembers paragraphs under revision, prior articles and background science (3), plans the review (4), and uses a database, statistics software and an editor (5). The researcher's feedback is stored for the next version (6).

Our running example: the scientific report assistant




Meet the scientific report assistant




It reads and summarises new articles, extracts and organises data and figures, manages references and generates the report.

Perception: interpreting the environment

INPUTS AND WHO HANDLES THEM



THREE RESEARCH DIRECTIONS



IN OUR EXAMPLE...

Reads new articles with language models, and also interprets charts and laboratory images, so it grasps the structure of the research and not only the text.

Perception: the report assistant at work

WHAT THE PAPER SAYS

Section 3.1

Perception

Language models read and analyse newly published articles, and multimodal models interpret charts or laboratory images. Combining visual and textual data lets the assistant grasp the structure of the research, not only summarise the text, and use that understanding in later report writing or analysis.

INPUTS → WHAT THE ASSISTANT UNDERSTANDS

ILLUSTRATIVE DEMO



Article text

Chen 2024 (PDF)

Reports a rise in microplastics downstream



Chart

Figure 2, bar chart

Bars are higher at sites 4 to 6



Lab image

Microscope photo of a filter

Fibre-like particles are visible

Structured understanding of the study

QUESTION

Where do microplastics build up?

METHOD

Water sampled at several sites

FINDING

Higher downstream of the town

FIGURE EVIDENCE

Figure 2: bars higher at sites 4 to 6

Handed to memory and planning

WHAT IT ADDS

Perception turns articles, charts and images into one structured understanding.

Role: who is this agent supposed to be?

A role bundles goals, skills and constraints, so agents can specialise and be audited.

WHO DECIDES THE ROLES? MORE HUMAN CONTROL → MORE FLEXIBILITY

Set in advance

A person writes the roles before the system runs. Some designs add roles for oversight.

FOR EXAMPLE

CrewAI a software framework

HADA a research architecture with an ethics agent

Discovered by the system

Nobody writes the roles. The system finds useful ones through trial and feedback.

FOR EXAMPLE

R3DM a published research method

Changes as it runs

The agents rewrite their own profiles and roles as feedback arrives.

FOR EXAMPLE

MorphAgent a published research method

In our example Three roles, each with its own goals and constraints: a collector of new articles, an analyser of laboratory data, and a manager of citations and references.

Role: the report assistant at work

Role

The assistant defines distinct roles: a collector of new articles, an analyser of laboratory data and a manager of citations and references. Each role has its own objectives and constraints, and together they coordinate to produce a precise, well-structured report. Roles help avoid confusion across tasks while keeping specialisation.

THREE ROLES, EACH WITH A GOAL AND A CONSTRAINT

ILLUSTRATIVE DEMO



Collector

GOAL

Find new relevant articles each week

CONSTRAINT

Peer-reviewed sources only



Analyser

GOAL

Process the lab CSV of particle counts

CONSTRAINT

Never edit the raw data



Citation manager

GOAL

Keep one consistent reference list

CONSTRAINT

Every claim has a source

A well-structured report

Without roles: tasks get muddled

WHAT IT ADDS

Roles give each part of the report its own goal and rules, so the tasks stay separate and specialised.

Memory: three kinds, used together



Parametric memory

INSIDE THE MODEL

Training knowledge compressed into the model's weights. Hard to update online; high risk of catastrophic forgetting.
What the model learned in training (inside the model, hard to change).



Episodic (transient) memory

THE CONTEXT WINDOW

Holds only a limited number of tokens, like human working memory. Retains a small part of recent conversation.
What is in the conversation right now (small, temporary).



Retrieval-based memory

VECTOR STORES

Stores large volumes and retrieves what is relevant. Enables RAG, extends knowledge beyond training data, and can reduce hallucinations.
What the agent looks up from an external store (large, updatable).

HYBRID

Hybrid memory combines internal and external. It preserves past dialogue and feedback across turns, so replies stay coherent.

WHY EXTERNAL MEMORY EXISTS

LLMs can only attend to a limited number of preceding tokens, so long documents and conversations are hard to handle.




IN OUR EXAMPLE...

Short-term: the paragraphs under revision.
Retrieval-based: prior articles and research data. Parametric: foundational scientific knowledge.

Memory: organising and managing it

MEMGPT-STYLE HIERARCHY (OPERATING-SYSTEM INSPIRED)



IN OUR EXAMPLE...

Latest draft in active memory, partly organised articles in mid-term, key sources in long-term. Summarise old sessions and drop low-importance data to keep the report concise.

FIVE MANAGEMENT POLICIES



Storage

Record events that matter for long-term goals.



Summarisation

Keep compressed summaries, not full transcripts.



Deletion

Clear redundant items from short-term memory.



Re-indexing

Re-index periodically for optimal, unbiased retrieval.



Sensitive-data purging

Securely remove personal or confidential data after use.

Memory: the report assistant at work

WHAT THE PAPER SAYS

Section 3.3

Memory

The assistant combines memory types: short-term memory for the most recent paragraphs under revision, retrieval-based memory for prior articles and research data, and parametric memory for foundational scientific knowledge. Policies such as summarisation and removing low-importance data keep reports concise, accurate and up to date.

THREE LANES: WHICH MEMORY ANSWERS WHICH QUESTION

ILLUSTRATIVE DEMO



Short-term

Recent paragraphs

ASSISTANT ASKS

Which paragraph am I editing?



LIVES IN

Context window



Retrieval-based

Prior articles and data

ASSISTANT ASKS

What did Okafor 2025 report?



LIVES IN

Article and data store



Parametric

Foundational science

ASSISTANT ASKS

What is a confidence interval?



LIVES IN

Model weights

MEMORY POLICIES



Summarise last week's session



Drop low-importance notes

WHAT IT ADDS

Memory lets the assistant keep the current draft in view, recall prior work and stay concise.

What should the report assistant remember?

The scientific report assistant reads new articles, manages references and keeps a researcher's report up to date.

Store

Summarise

Delete

Re-index

Purge

i

The researcher's preferred citation style

☐☐☐☐☐

ii

The full transcript of last week's working session

☐☐☐☐☐

iii

Study participants' identifying details in a lab dataset

☐☐☐☐☐

iv

Last month's shortlist for a section since dropped

☐☐☐☐☐


Pick a policy for each and defend it. Privacy is a design constraint, not an afterthought.

Planning: decide, search, reflect




Chain of Thought

Linear intermediate steps, one after another.



Tree of Thought

Branch, evaluate options and revisit alternatives.



Self-consistency

Several reasoning chains, then a majority vote.

REFLECTION HAS TWO PARTS

Evaluation: review intermediate outputs and correct the reasoning if errors appear.

Improvement: use feedback or external guidance to optimise the result.

Longer is not always better. Accuracy rises, peaks, then falls as errors accumulate; a “length-filtered vote” is more stable.

Planning = finding an ordered sequence of actions to reach a goal. **Search** = the computational process of carrying it out.

IN OUR EXAMPLE...

Categorises articles by domain or year and plans a logical path through them. If it spots an overlooked key article, it revises its path.

Planning and reflection: the report assistant at work



WHAT IT ADDS

Planning orders the work and reflection catches the mistakes, so the report improves before it is delivered.

Action: tools turn text into effects



ReAct

Interleaves reasoning traces with action calls.



Toolformer

Learns by self-supervision when and how to call APIs: calculator, search, translation.



WebGPT

Browses the web and gathers citations; improved by human feedback.



GUI agents

Screen Agent, Cradle: act through screenshots, mouse and keyboard (plan, act, reflect).



SWE-agent

Interface adapted to IDEs and repositories for large software tasks.



Voyager

A retrievable library of executable skills in Minecraft, with self-review.



Caution: tool-enabled agents are vulnerable to prompt injection. Defences: restricted tool permissions, user confirmation for sensitive actions, auditing of tool calls and continuous evaluation, all inside a clear and auditable control loop.

IN OUR EXAMPLE...

Queries a scientific database for new articles, runs statistical software on laboratory data, and formats the report with a text editor.

Action and tool use: the report assistant at work

WHAT THE PAPER SAYS

Section 3.5

Action and tool use

The assistant does not merely generate text for a report; it uses tools. It may query a scientific database for new articles, use statistical software to process laboratory data, or use a text editor to format the report automatically. This makes the report more efficient, more accurate and aligned with the researcher's needs.

A THOUGHT → ACTION → OBSERVATION TRACE

ILLUSTRATIVE DEMO

1

Thought Find new articles on the topic.**Action** `database.search("microplastics river 2026")`**Observation** 12 new articles**Database**

finds new articles

2

Thought Process the lab particle counts.**Action** `stats.run("particle_counts.csv")`**Observation** summary table saved**Statistics software**

processes lab data

3


Thought Format the report for the journal.**Action** `editor.apply_template("journal")`**Observation** report formatted**Text editor**

formats the report

WHAT IT ADDS

Tools turn text generation into doing work: finding, processing and formatting.

Learning: getting better while running



Retrieval-augmented memory

Store experiences, user feedback and results, then retrieve and reuse them.



Reinforcement learning

Choose reasoning strategies and when to call tools; learn from human or environmental feedback.



Hierarchical memory (MemGPT)

Drop what is unnecessary from short-term memory; keep what is essential.

Self-evolution (Zhao et al., 2026): agents refine planning and decision policies from outcomes, trajectories and evaluation loops.

Caveat: updating model weights online is hard and risks catastrophic forgetting, so much learning happens in memory.

IN OUR EXAMPLE...

Stores the researcher's feedback in long-term memory and uses it to improve the next draft's writing style and source selection.

Online learning: the report assistant at work

WHAT THE PAPER SAYS

Section 3.6

Online learning

After each report the assistant receives the researcher's feedback, for example on whether the sources are sufficiently covered or the figures correctly interpreted. By storing it in long-term memory and learning from past mistakes, it improves its writing style and source selection in later versions, becoming more accurate and efficient over time.

ONE FEEDBACK CYCLE



Version 1

First draft of the report, sent to the researcher

Dr Lee's feedback

"Two recent sources missing"

"Figure 3 misread"

ILLUSTRATIVE DEMO



Long-term memory

Feedback is stored for later versions

Version 2: what changed

- ✓ Added 2 sources
- ✓ Re-read Figure 3 correctly
- ✓ Style: shorter sentences

Each new version goes back to the researcher

WHAT IT ADDS

Feedback stored in memory means each version of the report is better than the last.

QUIZ

Which component?

a

It splits the task into sub-goals and tries three alternative routes.

b

It stores a summary of yesterday's session and re-indexes weekly.

c

It runs a statistics script and queries a database.

d

It has been told it is the "ethics reviewer" and may not approve spending.

e

It reads a photo of a lab chart.

f

The user's corrections change what it does next time.

THE SIX COMPONENTS

Perception

Role

Memory

Planning & reasoning

Action & tool use

Online learning

References

- [1] *W. X. Zhao et al., "A Survey of Large Language Models," Frontiers of Computer Science, vol. 20, no. 12, 2026, doi: 10.1007/s11704-026-60308-3*
- [2] *M. A. K. Raiaan et al., "A Review on Large Language Models: Architectures, Applications, Taxonomies, Open Issues and Challenges," IEEE Access, vol. 12, pp. 26839–26874, 2024, doi: 10.1109/access.2024.3365742*
- [3] *P. Peykani, S. Ghanidel, I. Javadi-Sisi, V. Snasel, and S. Mirjalili, "A Holistic Review of Agentic AI Frameworks, Applications, and Research Trajectories," Archives of Computational Methods in Engineering, 2026, doi: 10.1007/s11831-026-10675-8*