




AI agents: a model in a loop, a loop in a harness

The inner model · the loop · the harness · orchestration and teams · agentic systems

The model proposes. The harness decides what it sees, what it may do and what counts as done.



Models that predict




systems that act



an agent loop you can inspect

From an answer to an accountable action



An answer

Room B2 is free;
I've booked it for 15:00.

! no lookup · no booking · no evidence

A run




Text is a proposal, not an effect.




Acting needs tools, a loop, authority and checks.

Workflows and agents sit on one spectrum




Who decides the next step?



Start with the simplest design that works; add autonomy when flexibility pays.

The model reads one flat string



Structured messages (API input)

system (the operator instruction)

text You are a helpful assistant.

user

text how warm is it in New York?

assistant

tool call get_temperature(...)

reasoning Thinking about which tool to use ...

tool (results)

tool result { "temperature": ... }

JSON tool schema

```
{
  "type": "function",
  "function": {
    "name":
      "get_temperature",
    ...
  }
}
```

Chat template

ships with
open weights -
hidden behind
hosted APIs

Rendered string (what the model reads)

```
<|im_start|>system
You are a helpful assistant.
```

Tools

```
<tools>{"type": "function", "function":
{"name": "get_temperature", ...}}</tools>
For each function call, return a json object ...
within <tool_call></tool_call> XML tags
<|im_end|>
```

```
<|im_start|>user
how warm is it in New York?<|im_end|>
```

```
<|im_start|>assistant
```



A **chat template** renders roles and tool schemas into one token string.




A **tool schema** is prompt text; nothing is installed in the model.



Inside the harness there is one message format; adapters translate at each provider's edge.

A tool call is text the harness must parse



The model writes a call in the format it was trained on.



The harness parses it into id · name · arguments, or finds no call.



Formats differ by model family.



Parse strictly: derive the format, verify it, and refuse rather than guess.

Model reply (one turn)

Let me check the weather for Bologna.

```
<tool_call>{"name": "get_temperature",  
  "arguments": {"city": "Bologna"}}  
</tool_call>
```

That should give us the current conditions.

Text to show

Let me check the weather for Bologna.
That should give us the current conditions.

Tool call to parse

```
<tool_call>{"name": "get_temperature",  
  "arguments": {"city": "Bologna"}}  
</tool_call>
```

Reasoning to keep

(none in this example)

Tool call formats (examples)

Wrapped JSON (markers)

```
<tool_call>{"name": "get_temperature",  
  "arguments": {"city": "Bologna"}}  
</tool_call>
```

Bare JSON (no markers)

```
{"name": "get_temperature",  
  "arguments": {"city": "Bologna"}}
```


Nested tags (no JSON)

```
<function=get_temperature>  
  <parameter=city>Bologna</parameter>  
</function>
```

Name in opening marker

tool name inside the opening marker.

Structured output turns text into a checkable proposal



```
{"name": "get_weather",  
  "arguments": {"expression": "2+2"}}  
parses · wrong call
```

```
{"oneOf": [  
  {"properties": {"name": {"const": "get_temperature"},  
    "arguments": {"required": ["city"]}}},  
  {"properties": {"name": {"const": "list_cities"},  
    "arguments": {"properties": {}, "additionalProperties": false}}}  
]}
```

A **schema** states the shape:
type · properties · required · enum ·
additionalProperties: false.

A **decision schema** ties each tool name to its own arguments.



Three routes to structure:

ask in the prompt · trained formats ·
constrained decoding.

Shape is not truth:

the harness validates, then checks
meaning and authority.

Training teaches the format; only measurement shows compliance




The template shows the format the trainer intended.

A passing format check is not compliance: measure each model.

Reasoning models think in tokens before acting; the harness carries those tokens faithfully.

The model remembers nothing; the harness re-sends and budgets context




- Every turn re-sends the instruction, tool schemas, recalled memory and conversation.
- Order sections **stable** → **volatile**, because prompt caching matches a prefix.
- Over budget: leave out **whole old exchanges**, oldest first.
- Never cut the run in progress.
- **Memory** decides what is relevant; the **budget** decides what fits.

	Stable prefix (cacheable)			Conversation history (10 tokens per message)						Budget 60
Turn 2	Instruction	Tools	Skills	q0	a0	q1	a1	q2		50 of 60
Turn 3	Instruction	Tools	Skills		q2	a2	q3			30 of 60 cut to half
Turn 4	Instruction	Tools	Skills		q2	a2	q3	a3	q4	50 of 60 cut held
Turn 5	Instruction	Tools	Skills		q4	a4	q5			30 of 60



10 tokens per message · budget 60 · keep 0.5 Synthetic counter; an illustration.

What the inner model contributes, and where it fails



Contributes

reads the goal · chooses the next action · writes arguments · composes answers · routes or plans for a team.



Invents values



check answers against tool results.



Skips instructions it has read



enforce outside the prompt.



Misses a relevant tool or skill



write better descriptions.



Obeys text found inside data



keep authority outside the text.




Never stops, or breaks format



limits and strict parsing.

The model is one policy among several

- **A policy** decides: state $\rightarrow$ action. It may learn from an episode.
- **An action** is a final answer, or one or more tool calls.
- **The same seat** takes a language model, a script or a Q-table.
- **Vocabulary:** state $\cdot$ action $\cdot$ step $\cdot$ episode $\cdot$ reward (optional).





```
agent = Agent(ModelPolicy(model), tools=[get_temperature])
```

```
agent = Agent(ScriptedPolicy(call("get_temperature", {"city": "Bologna"}),  
    answer("It is 28 °C in Bologna.")), tools=[get_temperature])
```

```
agent = Agent(QLearning(epsilon=0.25), tools=[get_temperature], reward=Correct())
```

One run, turn by turn



Append-only log



user	request temperature in Bologna
assistant (call)	get_temperature(city="Bologna") c1
tool (result)	28.0 for c1 (simulated)
assistant (answer)	It is 28 °C in Bologna.

turn → decision → results → turn → decision → verdict → answer

The loop in twelve lines



Illustrative pseudocode of the reference harness's loop.

```
for turn in 1 .. max_iterations:
    state = build_state(log, tools, skills, turn)
    state = recall(memories, state)           # all the policy is shown
    action = policy.decide(state)             # final answer, or tool calls
    if needs_person(action): pause(action)    # nothing has run yet
    results = run_calls(action)               # failures become results
    if finished(action, results):
        verdict = verify(answer, state)
        if verdict.ok or no_retries_left: return answer
        log(answer, verdict.reason)           # try again, knowing why
    else: log(action, results)
raise LoopLimitExceeded
```



Turn limit and retry limit are separate.


Every run ends in a named outcome



- **Bounded:** a turn limit raises an explicit error, not an empty answer.
- **Failures are information:** a tool error returns as a result.
- **The log survives** errors and is the source of every turn.
- **Outcomes:** answered · refused · paused · limit reached · handed off.
- **Answer status** and **effect status** are recorded separately.

Activity: predict the next event

For each case, decide what happens next in the agent harness and which part acted.



A

Input / context

`get_temperature(city="error")`
tool raises an error.

Next event?

?

B

Input / context

`book_room(room="B2", time="15:00")`
requires confirmation.

Next event?

?

C

Input / context

Tool returned 28.0 (simulated), answer says 35.5 °C,
and grounding checks are on.

Next event?

?

D



Input / context

Turn 16 of 16 ends with yet another
tool call.

Next event?

?

The harness is everything around the model



- A **model** turns a prompt into text and stops.
- The **harness** gives it something to call, to look up and to remember, a person to ask, a check, a place to run code and a score.
- The **same model behaves differently** in different harnesses.
- **Course reference harness: Agent Playground**, readable and written from scratch; you open it in the lab.




Course reference harness


Agent Playground

A minimal, readable harness written from scratch. You open it in the lab.

Seven seams, each with a working default



AGENT HARNESS




Seam	Contract	Default	Alternatives
Tools	function or toolkit	none	toolkits, code runner
Skills	SKILL.md folders	skills folder roster	—
Memory	remember · recall	whole transcript	window, layers, corpus, board
Policy	decide · learn	ask the model	script, Q-learning
Verifier	check the answer	accept everything	schema, grounded, tests, proof
Reward	score a step	none = inference	settled, verified
Executor	run a command	subprocess	container sandbox



Observers (a trace, a display) watch from outside; they are not seams.

Course reference harness:
Agent Playground

A tool is a function with an operational contract



Function (code and docstring)

```
def get_temperature(city: str) -> float:
    """
    Returns the temperature (in Celsius degrees) for a given city.
    Args:
        city (str): The name of the city.
    """
    return 25.0
```

Derived schema (what the model reads)

name: get_temperature

city: string

required: city

Mark badges (how a call runs)



confirm



run externally



stop after



show



cache



time out

Effect classes (increasing authority)



read

observe_scene



compute

transform_point



predict

simulate_viewpoint



write

report_issue




actuate

move_camera




Docstring = model-facing interface. Predict previews without changing the world.
Failing calls return errors; misspelled marks are refused.

Authority lives in the application, not in the text



The model proposes a call; an application-held policy decides: **allow** · **pause** · **deny**.




- A **pause** stops the run before anything executes.
- **Approve**, and the call runs. **Resume** without approval returns declined. No reply leaves the run paused.
- Text cannot approve itself: not a tool argument, not a plan, not a web page.



```

office = Office(requires_confirmation_tools=["book_room"])
agent = Agent(ModelPolicy(model), tools=[office])
try:
    agent("Book room B2 at 15:00")
except Paused as stop:
    pending = stop.pending[0]
    reply = agent.resume(approve={pending.id: True})
    # book_room(room="B2", time="15:00")
  
```

Three checks: permission before, effect after, answer at the end



⚠ | A correct answer does not prove that an action was authorized.


The log is not the view

Log: everything that happened, append-only and never pruned.

View: what memories recall into the next prompt; it supplies the model's context.

Chat history is one memory: swap it for a window and the loop is unchanged.

Some memories were never turns: documents, web pages, a shared board.





a sliding window over the conversation

```
agent = Agent(ModelPolicy(model), memory=[Working(keep=12)])
```

passages from a corpus, then the whole conversation

```
agent = Agent(ModelPolicy(model), memory=[Documents("handbook/"), Transcript()])
```

Four memory layers, four lifecycles




Layer	Holds	Written	Main risk
Working	turns in play	rebuilt every turn	lost constraints
Episodic	earlier runs	appended per run	wrong episode
Semantic	facts with provenance	promoted deliberately	stale or contradicted facts
Procedural	skills and procedures	versioned; read-only during runs	obsolete or poisoned procedure



Plus experience: rewarded runs replayed as worked examples.

Skills: procedures loaded on demand



A skill is a folder:

SKILL.md (name, description, instructions) plus bundled files.



Finding a skill is retrieval: the description decides.

Level 1: the roster

One line per skill, always in the prompt.

Token cost:

Level 2: the body

Read through a tool call when needed.

Token cost:

Level 3: files and scripts

On request, and scripts only when allowed.

Token cost:



Bundled files and scripts.

tools register
when the body
is read



Tools
register



name: weather

description: Look up the current temperature for a city in Celsius degrees, or list the cities covered.

tools: tools.py



1. Call `get\_temperature` with the city name exactly as the user wrote it.

Where code runs, and what untrusted text can reach



Learning can land in the policy, the memory or the weights

With a reward, every turn becomes a scored step.



Invented-game experiment

one seed · 24 held-out locks
one run · 27B open-weight model

Reasoning off

untrained 5/24 → best combined
10/24, $p = 0.18$



not significant

Reasoning on

untrained 4/24 → memory
12/24, $p = 0.06$

Directional evidence,
not a benchmark.

Orchestration: an agent can be another agent's tool




- **Delegation is a tool call:** a request in, an answer out.
- The delegate **does not automatically receive** the caller's conversation.
- Its **description guides** when work reaches it.
- An **instance** remembers earlier requests; a **factory** starts fresh every time; shared resources still need **concurrency guards**.
- A **shared board** coordinates through memory, without messages.

```
desk = Agent(ModelPolicy(model),
             memory=[Documents("handbook/"), Transcript()])

manager = Agent(ModelPolicy(model),
                tools=[as_tool(desk, "office_desk",
                              "Questions about the office.")])
```

A team is an agent whose tools are its members



Leader
(team agent)


Coordinate

assigns chosen members



Route

selects one member



Broadcast

asks all members at once



Tasks

manages dependencies and assignees

Task plan



Mode	The leader	The run ends when
coordinate	assigns and combines	leader answers
route	selects one member	member answers
broadcast	asks all at once	leader combines answers
tasks	manages dependencies and assignees	all tasks complete

```
desk = Agent(ModelPolicy(model), name="desk",
              description="Questions about the office.")

weather = Agent(ModelPolicy(model), tools=[get_temperature],
                name="weather", description="Temperatures in cities.")

team = Team(ModelPolicy(model), members=[desk, weather],
            mode="route")

print(team("How warm is it in Pisa?"))
```



Same loop: limits, verifiers, pauses and budgets still apply.





A team can itself be a member. Member names are an enum: no delegating to someone absent.

Teams cost time; guards belong in code

Every **delegation** is a **full agent run**: time and tokens add up.

A guard refused two parallel tasks on one agent; the leader read why and ran them in turn.

Use teams for separable, parallel, read-heavy work; keep writes single-threaded.





separable context and parallel reading
→ team

shared context and many writes
→ one agent


One workstation, one locally served 27B model, 4-bit; timings are observations, not a benchmark.

One anatomy across today's agentic systems



Each is an inner model, a loop and a harness, and sometimes a team.

Protocols connect agents; they do not grant trust



MCP

A host connects to servers that offer tools, resources and prompts.



A2A

Independently built agents exchange tasks and messages.



Agent Skills and AGENTS.md


Portable know-how and project instructions.



Authentication, permission, approval and sandboxing stay in the harness.

Protocols connect agents and resources; they do not grant trust.

Judge the outcome and the trajectory



Outcome

the final state and answer, including a correct refusal or hand-off.

Trajectory

the right tools, the right arguments, no unapproved effects.



Reliability

repeated trials, not one lucky run.



Cost

turns, tokens, time and human review.

Test the risks: goal hijack • tool misuse • privilege abuse • code execution • memory poisoning • cascading failures.

goal hijack



verifiers

tool misuse



tool contracts

privilege abuse



authority gate

code execution



sandbox

memory poisoning




memory lifecycle

cascading failures



limits and guards

Where each responsibility lives



Model
proposes calls, answers and plans.

Loop
controls turns, limits, retries and outcomes.

Harness
decides context, capabilities, authority, checks, isolation, observation and learning.

Orchestration and teams
divide work, isolate context, and add cost.

Exit ticket

For an agent in a virtual room, name one tool that needs approval, one memory layer and one verifier, and say whether it needs a team.



References

Agency and architecture

- AIMA
- Wooldridge & Jennings
- ReAct
- CoALA

Inner model

- Toolformer
- JSON Schema
- Guided generation
- DeepSeek-R1

Harness and context

- Context engineering
- Long-running harnesses
- Agent Skills

Orchestration and teams

- Building effective agents
- agno Teams
- Anthropic research system
- Cognition

Protocols, evaluation and safety

- MCP
- A2A
- AGENTS.md
- AgentBench
- GAIA
- tau-bench
- SWE-bench
- OSWorld
- WebArena
- OWASP

Virtual-world agents

- Generative Agents
- Voyager
- SIMA 2

