




Chatbot, workflow, or agent?

A one-page cheat sheet, in plain English

ONE COMMON WAY TO EXPLAIN IT

Our plain-English framing, with wording from Anthropic's "Building effective AI agents" (a December 2024 article) and OpenAI's developer docs. Companies word these ideas differently. Running example: booking a table for Friday.




1



Chatbot: you ask, it answers

In plain terms: it writes a reply, and the next move is yours.

Try: Ask for restaurant ideas. You still do the booking.



2



Workflow: a path set in advance

Anthropic: LLMs and tools orchestrated through predefined code paths. Predictable for well-defined tasks.

Try: Every morning: read emails, summarize, save to a folder.



3



Agent: it directs its own steps

Anthropic: LLMs dynamically direct their own processes and tool usage. OpenAI's docs: an AI system with instructions, guardrails and tools to take action on the user's behalf.

Try: Goal: a table for four on Friday. It picks the steps and tools.




4



Is a simpler solution enough?

Anthropic recommends the simplest solution possible, and notes that agentic systems often trade latency (waiting) and cost for better task performance.

General good practice: ask whether one good prompt or a fixed workflow could do this.




5



Can you check the result, and what if it is wrong?

Anthropic warns of the potential for compounding errors. In plain terms: one early mistake can carry into later steps.

General good practice: start with low-stakes tasks.



6



Set limits and test safely first

Anthropic recommends extensive testing in sandboxed environments, with appropriate guardrails. It also describes agents pausing for human feedback at checkpoints.

General good practice: decide in advance what it may not do.



Quick rule (general good practice): a chatbot answers, a workflow follows a path, an agent chooses its path. The more steps it can take alone, the more you should test and check.