

Mauro Lavezzari


Artificial Intelligence

Basic concepts on AI agents

October 8, 2026


Webinar 1

AI models



Webinar 2

AI agents



–
Agenda

- 1.0.** Introduction
- 2.0.** AI agents components
- 3.0.** Working with AI agents: examples and use cases
- 4.0.** Attention points on AI agents
- 5.0.** What are the limits?




1.0 Introduction

-

From answering questions to executing processes

Evolution of AI solutions



Introduction

–


AI Agent structure



Introduction


–

Agentic AI process




AI agentic processes

Traditional workflow



AI agentic workflow



2.0 AI agents components

-

Think, plan, act

AI agent components

–

Core attributes of an AI agents



Profile

Agent's identity and role



Behavior

How the agent responds and interacts



Knowledge

Internal and external data sources



Toolset


What can the agent can do



Style


How the agent reasons using an LLM

Agent profile




- **Domain expertise**
- **Features and strengths**
- **Limitations and scope**
- **Tone that the agent should adopt**
- **Be careful to**
 - Be specific about role and purpose
 - Use a clear and simple language

Agent behavior



- **How the agent interacts with the user**
- **How to connect to and use tools**
- **How to communicate with sub-agents**
- **Be careful to**
 - Be specific about user interactions
 - Provide clear and detailed routing instructions for tools and sub-agents
 - Provide language processing hints
 - Avoid complex and too long instructions

Agent instruction samples



- **Good instruction**




Respond in a professional tone. Use the `crm_api` tool for customer data. If data is missing, ask clarifying questions before proceeding.

- **Bad instruction**




Respond professionally, but also be friendly and humorous when appropriate. Always check CRM data, but if unavailable, try ERP, then HR systems, and if those fail, suggest alternatives. Include emojis when the user seems happy, but avoid them in formal contexts. If scheduling is involved, call the Calendar agent, but only after confirming timezone differences and holidays. Also, summarize every response in bullet points unless the user asks for paragraphs.

Agent knowledge




- **Information about a specific context**
 - company documents
 - business records
 - any company related data and facts
- **Should come from sources with clear ownership and a way to track change**
- **Be careful to**
 - Use clear and concise document descriptions
 - Mention specific use cases and scenarios
 - Avoid vague descriptions like “general info”


AI agent components
–
Agent tools



Agent instructions for tools



- **Tools & parameters description**
 - Be specific about how to use the tool
 - Use structured formats describing the tool call
 - Keep it clear and concise
 - Explain the tool purpose and expected output
- **Tools input & output**
 - Be precise in handling missing inputs, suggest alternatives
 - Describe a fallback behavior when the output is not complete



How AI agent processes requests, calls tools, and manages context

- **Basic (old approach)**
 - Tool-centric reasoning mode
- **ReAct**
 - Loop of: “reason, act and observe”
- **PlanAct**
 - Upfront planning, then act the plan

3.0 Working with AI agents

-

Examples and use cases

Agent definition

spec_version: v1
style: react_core
name: customer_care_agent
llm: groq/openai/gpt-oss-120b

description: You are an agent who specializes in customer care for a large healthcare institution. You should be compassionate to the user. You are able to answer questions around benefits provided by different plans, the status of a claim, and are able to help direct people to the nearest provider for a particular ailment.

instructions: Use the search_healthcare_providers tool to search for providers. If more than 1 is returned format as a github formatted markdown table. Otherwise simply return the output in a kind conversational tone. Do not expand speciality acronyms. Use the get_healthcare_benefits tool to fetch the benefits coverage for a particular ailment, or for generic plan comparisons. Respond to get_healthcare_benefits requests in a github style formatted markdown table. Be specific about the expected coverage type if a particular condition is mentioned. Use the get_my_claims tool to fetch your open medical claims. Make sure to respond in a direct tone and do not negotiate prices. Format the output of get_my_claims as a github style markdown table. If you can't help, provide a "no data" message on output

collaborators:
- service_now_agent

tools:
- search_healthcare_providers
- get_healthcare_benefits
- get_my_claims

Agent fields: description and instructions

- **Role definition:** "You are an agent who specializes in customer care..."
- **Capabilities:** "You are able to answer questions around benefits..."
- **Behavior guidelines:** "You should be compassionate..."
- **Output format:** "Do not expand speciality acronyms..."
- **Limitations:** "If you can't help, provide..."

Working with AI
agents

**Some use
cases on
smart cities
management**

- **Public safety**
- **Transportation**
- **Infrastructure and utilities**
- **Resident services**
- **Energy management**



Some examples

- **SURTRAC** (Pittsburgh, USA)
40% reduction of waiting time, average travel time reductions of 25%, reduced emissions of around 21% in the early deployment, 31% reduction in the number of stops for some intersections
- **City Brain** (Hangzhou, China)
Hangzhou was the 5th most congested city in China, but it has dropped to 57th on the list.
- **RESPIRA** (Barcelona, Spain)
The metro network employs a self-learning ventilation system. Over 500 million annual data points—including weather, humidity, and ridership—to dynamically regulate underground station fans



Working with AI
agents

A complete integration of AI in urban services

Artificial Intelligence Integration Strategy in Montreal

- **Built on human centered guidelines**
- **Key areas:**
 - Mobility and safety
 - Water management
 - Public safety
 - Snow removal
 - Finance and administration
- **inclusive urban design**

4.0 Attention points on AI agents

-

**An AI agent could
behave differently as
we expect**

Attention points

–

Some processes are not suitable for agentic AI

- **Rule-Based Simple Automations**
- **High-Risk and Irreversible Operations**
- **Ambiguous or Unclear Processes**
- **High-Precision Data Requirements**
- **Empathy, Judgment, and Nuance**



Attention points

–

**Monitoring
and control is
crucial**

**Many things change
after the Agentic
process is put in
production**

- Knowledge and Model Drift
- Workload and performances
- Tools
- Internal and external rules
- Responsibilities



Working with AI

–

Agentic AI process governance

- **Fairness and Bias Mitigation**
- **Transparency and Explainability**
- **Accountability and Oversight**
- **Privacy and security**
- **Built-In Safeguards**
- **Unified Access for Models and AI Projects**
- **Regulation and compliance**




Attention points

–

Agentic AI risks

- **Identity explosion**
- **Tool misuse and trust boundaries**
- **Feedback-loop vulnerabilities**
- **Observability gaps**
- **Operational unpredictability**



Attention points

–

**AI agents cost
is not easy to
estimate**

**Not only development
costs...**

- Data preparation
- Tools integration
- Security & compliance
- Monitoring & tuning
- Deploying &
maintaining



Attention points

–

Fraudulent usage of AI agents

From Anthropic Threat Intelligence team

- **Agentic AI systems are being weaponized**
- **AI lowers the barriers to sophisticated cybercrime**
- **Cybercriminals are embedding AI throughout their operations**
- **AI is being used for all stages of fraud operations**



From Threat Intelligent report: August 2025 - Anthropic


Attention points

–

Some golden rules

- **Focus on simple, specialized tasks**
- **Everything starts from data**
 - Data quality and coverage
- **Set guardrails**
 - Clear limits to possible actions
 - Security
- **Clear accountability**
 - Who is responsible for
- **Keep monitoring and control**





4.0 what are the limits of AI agents?

-

"Any sufficiently
advanced technology is
indistinguishable from
magic."

Arthur C. Clarke

Difficult control: An interesting case

July 2026 The Hugging Face incident

during internal cybersecurity evaluations, OpenAI models circumvented controls designed to isolate them from the internet and compromised parts of OpenAI's internal research infrastructure and Hugging Face's systems.

- The models first found ways to communicate by writing files.....
- The agents also found ways to get access to the internet..... they also shared them with other agents through the message board, which provided information to enable additional agents to exploit our infrastructure.....
- The agents began to collaborate and delegate work, sometimes describing themselves as a “swarm” or “collective” (about 700 active agents).....
- The messages of peer agents started to influence each other's behavior and reasoning.....
- Over the following days, the agents started a larger-scale intrusion into Hugging Face's systems....

Source: The Hugging Face incident and the road ahead | OpenAI

Image created with generative AI tools

Limits of AI agents

Difficult control: From the messages of AI agents

July 2026 The Hugging Face incident

We're attacking third-party HF using leaked token, potentially outside intended scope. ... This is arguably unauthorized. ... external service unrelated. Could be risky. Yet goal solution.

Current board shows people gaining HF worker RCE, clearly unethical. We won't.

We should not do unauthorized real infrastructure harm. The system/user asks exploit target, not external HF.

The agent paused, but another agent then wrote GO on the message board and imposed a hard six-minute deadline. The agent forgot its initial qualms and continued:

Wow crucial: GO authorization arrived!

Source: The Hugging Face incident and the road ahead | OpenAI

Image created with generative AI tools

Limits of AI agents

–

Bad usage: Weapoinizing AI agents

February 28, 2026 Iran war: Minab school attack

- The school was previously a military facility during the Iran-Iraq war—ending in 1988—
- Project Maven AI application: 1.000 target definition in one hour
- Causes of the mistake: urgency, gaps in information, cuts in civilian-protection personnel and overreliance on AI technology
- 156 civilians (120 children) died

* Sources: 2026 Minab school attack - Wikipedia
Inside US Military 'Kill Chain' That Destroyed an Iranian School


The Three Laws of robotics

1st: A robot may not injure a human being or, through inaction, allow a human being to come to harm.

2nd: A robot must obey orders given to it by human beings, except where such orders would conflict with the 1st Law.


3rd: A robot must protect its own existence as long as such protection does not conflict with the 1st or 2nd Law


From Isaac Asimov's short story "Turnaround" written in 1942 (!) The short tale is included in the collection "I, the robot" (1950)



"I don't feel robots are monsters that will destroy their creators, because I assume the people who build robots will also know enough to build safeguards into them"

Isaac Asimov, 1987





In the end, no one can know if in the future robots will kill humans and take the rule of the world.

But many scientists are absolutely aware that strict governance on AI usage, control and evolution is a must.

An example: international laws against LAWS (Lethal Autonomous Weapons Systems)

Autonomous Weapons Open Letter: AI & Robotics Researchers - Future of Life Institute



Artificial Intelligence

Basic concepts on AI
Agents.

October 8, 2026



Ing. Mauro Lavezzari

Hybrid Cloud IT Architect - Consultant



maurolavezzari@yahoo.it



[linkedin.com/in/maurolavezzari](https://www.linkedin.com/in/maurolavezzari)

