

Agentic AI

What is it? Why do we need it?

Sept 2026 - Rick Rys P.E.



Where Agentic AI fits in the world of AI

The global overall AI market is estimated at \$539.5 billion for 2026, with Gartner forecasting full-stack AI spending to hit \$2.59 trillion

Agentic AI (~\$10–\$12 billion) represents just a tiny fraction of the broader AI economic footprint. It looks incredibly small because it only measures the software license and orchestration layer for autonomous workflows, completely excluding the staggering infrastructure powering it.

Four dominant players control 87% of the active enterprise agent market:

Enterprise Provider / Platform	Market share
Microsoft Copilot Studio	31%
Salesforce Agentforce	24%
Anthropic (Claude API)	18%
Google Vertex AI Builder	14%
ServiceNow AI Agents / Moveworks	7%

Agentic AI

- Agentic AI is about allowing AI models to act
- Going from a standard conversational query to Agentic AI is fundamentally a transition from an open-loop, single-pass text generation to a closed-loop feedback control system.



Examples of Agentic AI

Autonomous Robots, Cars, Planes, and Drones

Autonomous Software Engineering & Security

Autonomous Coding Engineers (Cognition Devin, Cursor, SWE-Agent)

Autonomous Cyber Defense & Penetration Testing: Platforms like DARPA's AI Cyber Challenge (AlxCC)

Biopharma & "Self-Driving" Laboratories

Closed-Loop Materials & Molecule Discovery (e.g., A-Lab, Emerald Cloud Lab)

Clinical Trial Operations & Prior Authorization - **Insilico Medicine**: They are the "pure play" leader: Drug discovered and designed entirely by AI (INS018_055 for Idiopathic Pulmonary Fibrosis)

Supply Chain & Operations Management


Autonomous Logistics Re-Routing

Predictive Maintenance Agents

Enterprise, Finance & Legal Operations

Transaction Compliance & AML (e.g., JPMorgan Chase)

Contract Analysis & Autonomous Redlining:



TNIK (Traf2- and Nck-interacting kinase).

AI terms

- **AI:** Computer systems or machines programmed to perform tasks that typically require human cognition.
- **AGI, SGI:** Artificial General Intelligence (AGI) describes hypothetical AI capable of matching human intellectual breadth across any arbitrary domain, whereas Specialized General Intelligence (SGI) refers to broad general-reasoning capabilities tailored to deep, complex vertical sectors like industrial automation or genomics.
- **Prompts:** The structured natural-language instructions, system context, and environmental feedback provided to an agent to trigger and guide its autonomous reasoning, tool selection, and multi-step execution loop.
- **Tokens, tokens/second:** Tokens are the atomic character- or sub-word fragments an AI processes, while tokens per second measures inference velocity—a key metric governing how quickly an agent can reason, query external tools, and react in real-time loops.
- **Agents:** Autonomous software systems that leverage AI reasoning models to decompose high-level objectives into subtasks, invoke external APIs or tools, evaluate environmental feedback, and iteratively correct course to achieve goals without constant human intervention.

AI Terms

- **Transformer model:** The foundational deep-learning architecture relying on self-attention mechanisms that acts as the core reasoning engine inside modern agents, allowing them to track dependencies across lengthy task histories, code bases, and multi-turn tool outputs.
- **Digital Twin (IN SILICO”:** A high-fidelity, synchronized virtual model of a physical asset, biological mechanism, or plant process that an agent can simulate, test hypotheses against, and directly manipulate via closed-loop control.
- **Recursive Training:** An iterative learning regime where a model or agent trains on self-generated synthetic reasoning traces, self-evaluates candidate outputs, and updates its weights or policies through continuous feedback loops.
- **J-space:** An internal, shared latent workspace (identified via Jacobian-based interpretability analysis) within a transformer where intermediate, unspoken multi-step reasoning converges before an agent produces visible outputs or executes actions.
- **Kill switch:** A deterministic software guardrail or physical interlock that instantly severs an autonomous agent’s API access, actuator commands, or execution authority if it violates predefined safety boundaries or operational envelopes.

A few people developing AI



Sir Demis Hassabis
Google Deep Mind



Sam Altman
OpenAI ChatGPT



Dario Amodei
Anthropic -Claude




Jensen Huang
Nvidia


Core Traits of Production-Grade Agentic Systems

Capability	Static GenAI / Chatbot	Agentic AI System
Execution Trigger	Single prompt → Single response	Objective-based → Multi-step execution loop
Tool Usage	Text generation only	Invokes APIs, SQL queries, terminals, and web scrapers
Error Handling	Hallucinates or stops	Inspects error output, refines prompt, and retries
Human Interface	Direct conversationalist	Operates asynchronously with "Human-in-the-Loop" escalation

Simple? level feedback control




Traditional PID Feedback Control




Manufacturing Digital twins

Pharma Manufacturing Reactor Digital Twin




Vision, Experience, Answers for Industry, Energy, & Infrastructure

L3 Harris Full Scale Nuclear training simulator



Going from a standard conversational query to Agentic AI is fundamentally a transition from an open-loop, single-pass generation to a closed-loop feedback control system. Agentic AI uses JSON as its core communication language to connect unstructured Large Language Model (LLM) thinking with deterministic software systems.



Benchmarks

- **General & Dynamic Intelligence**
 - Livebench
 - LMSYS –Chatbot Arena
- **Advanced Reasoning & Science**
 - GPQA (Graduate-Level Google-Proof Q&A)
 - ARC-AGI (Abstraction and Reasoning Corpus)
- **Real-World Coding & Software Engineering**
 - SWE-bench (Verified & Pro):
 - LiveCodeBench:
- **Mathematics & Agent Automation**
 - FrontierMath:
 - OSWorld / AutoBench:

Key Trends in Agent Automation

- **Anthropic Dominates Code Agents:** Anthropic's Claude 5 ecosystem (Opus, Mythos, and Fable) controls the top of the SWE-bench chart.
- **OpenAI Leads Computer Use:** OpenAI's newer GPT-6 Astra leads the difficult OSWorld 2.0 chart, showing a stronger capability for navigating desktop environments and long-horizon tasks.
- **Open-Weight Gains:** Open-weight models like Meta's Muse Spark 1.3 and Ornith-1.5 are now actively competing with or beating previous-generation proprietary models.

SWE-bench Verified Standings (Software Engineering Agents)

Rank	Model	Provider	Type	% Resolved
1	Claude Opus 5	Anthropic	Closed-source	96.0%
2	Claude Mythos 5	Anthropic	Closed-source	95.5%
3	Claude Fable 5	Anthropic	Closed-source	95.0%
4	Claude Opus 4.8	Anthropic	Closed-source	88.6%
5	Claude Opus 4.7 (Adaptive)	Anthropic	Closed-source	87.6%
6	Ornith-1.5-397B	Ornith AI	Open-weight	86.0%
7	Claude Sonnet 5	Anthropic	Closed-source	85.2%
8	GPT-5.3 Codex	OpenAI	Closed-source	85.0%

This chart tracks an AI's ability to autonomously operate across desktop OS ecosystems, files, and browsers to complete open-ended tasks.

Rank	Model	Provider	Success Rate
1	GPT-6 Astra	OpenAI	72.60%
2	Claude Opus 5	Anthropic	70.60%
3	Muse Spark 1.3	Meta	66.90%
4	GPT-5.6 Sol	OpenAI	62.60%
5	Gemini 3.8 Flash	Google	59.00%
6	GPT-5.6 Terra	OpenAI	50.20%
7	Gemini 3.7 Flash	Google	47.90%

Agentic - four software layers wrap the base model

1. Tool Declaration (Function Calling)

- The LLM is provided with a list of callable tools defined in structured schema formats (e.g., JSON Schema or OpenAPI).
- Instead of returning conversational prose, the model decides: *"To answer this, I need to execute query_database(site_id=4)."*
- It outputs a structured payload containing the exact function name and arguments.

2. The Runtime Execution Engine

- The model itself does not execute code or touch networks—the **host runtime (orchestrator)** does.
- The runtime intercepts the model's structured tool request.
- The runtime executes the API call, database query, Python script, or external device command.
- The raw response (data, error code, or status) is captured.

3. Feedback Loop & Observation (The ReAct Pattern)

- The runtime injects the tool output back into the model's context as a new message labeled Observation:.
- **Thought:** *"The query returned an error: 'Connection timed out on Port 502'."*
- **Action:** *"I will retry using the backup IP address via ping_node(ip='192.168.1.50')."*
- This cycle of **Reason** → **Act** → **Observe** → **Re-evaluate** repeats until the model determines the objective is satisfied.

4. State & Memory Management

- **Working Memory:** The active context window holding the current execution trace, intermediate outputs, and error logs.
- **Persistent / Long-Term Memory:** External vector databases or graph stores used to store past learnings, user preferences, or multi-day state machines.

Safety, Guardrails & Determinism

- Because agentic loops can drift or enter infinite execution cycles, production implementations enforce deterministic boundaries:
- **Termination Conditions:** Hard limits on max iterations, runtime timeouts, and budget/token caps.
- **Schema Validation:** Strict parsing (e.g., Pydantic) to ensure the agent cannot pass malformed parameters to downstream APIs.
- **Human-in-the-Loop (HITL) Gates:** Requiring explicit human confirmation before executing high-consequence actions (e.g., issuing buy orders, modifying database records, or sending external emails).
- **Environmental consequences of AI**
- **Existential threat of AI**



Can AI's climate benefits outweigh its costs?

Bottom line is developing AI models uses energy and has a negative impact on climate. If AI is successful at improvising the economy, that will have a much larger negative impact on the climate if it expands the economy in general.

Some papers and simulations for more details:

- <https://mitsloan.mit.edu/ideas-made-to-matter/can-ais-climate-benefits-outweigh-its-costs>
- <https://support.climateinteractive.org/support/solutions/articles/47001301880-data-centers-and-artificial-intelligence-in-en-roads>
- https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6689098

Future challenges - Controlling high consequence processes

- Driving, flying and operating equipment (robotics)
- Automotive Manufacturing
- Process manufacturing
- Electric grid management
- Cyber security
- Military operations
- Medical diagnostics



AI poses catastrophic and existential risks.

Threats include cyber warfare, pathogen synthesis, autonomous battlefields, and loss of control (kill switches that don't work)

- Current Safeguards



- Export restrictions on advanced semiconductor manufacturing equipment
- Pre-deployment red-teaming on frontier models
- common definition of Lethal Autonomous Weapons Systems (LAWS)
- Industrial physical kill switches are air-gapped from networks

- Gaps

- Universal DNA/RNA Synthesis Verification Treaty
- A Legally Binding Autonomous Weapons Treaty:
- International AI Governance Body (IAEA Model):
- Enforce strict civil and criminal liability on frontier developers for downstream catastrophic damages

Intelligence Evolution

Human Evolution



How does intelligence evolution evolve?

The value of industrial AI?



AI is x10 more significant than the industrial revolution

AI is happening x10 faster than the industrial revolution

The AI developers have a goal to get to AGI without major human disruptions



Manufacturing and warehouse logistics increasingly use AI for scheduling/robotics

Process industry uses digital twins for operations and predictive maintenance

GE foresees a day when AI will control our electric grids



In Pharma Generative Biology is creating new physical products (molecules)

Innovation: DeepMind, Isomorphic Labs, IsoDDE (Drug Design Engine), BioNeMo (NVIDIA), Lilly-NVIDIA Co-Innovation Lab

The Convergence: AI + CRISPR The synergy between structural biology (AI) and genomic engineering (CRISPR) has turned drug discovery into an engineering problem

Thank You

Rick Rys P.E.
www.R2Controls.com

