RADAR SPECIAL NO. 4A · 11 JULY 2026

Radar in-depth report

When artificial intelligence learned to collaborate

From ChatGPT to Codex, from intelligent agents to a new relationship between people and machines



Gruppo Radar · Comprendere oggi. Costruire domani
by effettoantropico.life

📌 **Did you know? — Why this special?** Every week Radar observes signals from technology, industry, finance, research and geopolitics. Some fade quickly; others change the conditions in which we work, study and choose. The evolution of conversational, multimodal and agentic systems belongs to the second category. This special does not tell the story of a single company or the latest version of a model: it observes the transition from a machine that answers to a machine that can collaborate, and the responsibilities this transition brings.

Introduction — When the relationship changes, not only the tool

The most visible novelty of artificial intelligence is not that a computer has learned to compose a sentence. It is that, to ask it to do something, today it is often enough to say what one wants to achieve. Writing a question, showing an image, dictating a note, uploading a document: these ordinary gestures replace some of the instructions, menus and technical languages that for decades have separated people from digital systems.

The change looks trivial if viewed as a new interface or with inexperience, but becomes more important when the system does not merely return "only" text: it can read a file, compare data, propose a plan, use an authorized tool, work on a file and return with a result to be checked. This is not autonomous operation in the human sense of the term, but an operating environment that translates a request expressed in a language that can be turned into a sequence of operations.

This proximity changes habits. A person who does not program can ask for a draft automation; someone studying can turn a question into a research path; a team can entrust a system with an initial survey of many documents, keeping the decision in the hands of those who know the case. The threshold for access falls, but access does not amount to understanding.

A smooth interface can create a misleading impression of competence. A language model does not automatically verify every fact, does not know the user's intention by itself and does not bear moral responsibility. It is important to make this clear. It produces plausible responses based on learned patterns; when it uses sources or tools, quality also depends on which tools it receives, how they are configured and who controls the process. Natural dialogue is therefore a useful achievement but also one more reason not to suspend judgment.

The decisive question is not only *"what will AI be able to do?"*. It is: ***how will the relationship between human beings and intelligent systems change?*** To answer, we need to look together at the model, the person using it, the organization introducing it and the infrastructure making it possible.

1. When everything seemed impossible

AI did not begin with chats. In 1950 Alan Turing proposed a conceptual experiment, the imitation game: instead of defining machine thought in the abstract, he asked whether an interlocutor could distinguish a machine from a person in a written conversation. *The so-called Turing Test has remained important because it shifts attention to observable behavior; it is not, however, a definitive measure of intelligence, consciousness or reliability*. [1]

A few years later, the 1955 proposal for the Dartmouth Summer Research Project gave the field the name *"artificial intelligence"*. The hypothesis was ambitious: aspects of learning and intelligence could be described precisely enough to be simulated by a machine. That promise guided a long period of research, but did not anticipate today's tools in a linear way. [2]

Early symbolic systems attempted to reason using explicit rules: if A happens, then consider B. Expert systems, widely discussed in the 1980s, encoded domain knowledge in large sets of rules. They were useful in narrow, well-defined settings; they struggled when the world became ambiguous, incomplete or different from the cases anticipated. The distance between a laboratory demonstration and reliable large-scale use led to funding cuts and reduced expectations: the so-called AI winters, periods of disinvestment and disillusionment.

The return was not a sudden twist, but the effect of three accumulations: *statistical learning methods, large quantities of digital data and growing computing capacity. Artificial neural networks—models made up of many parameters that change their behavior during training—became central again because they could extract patterns from examples instead of being given every rule one by one. The historical lesson is sober: progress in AI consists of insights, failures, financing, infrastructure and contexts of use. Not an inevitable march.*

📌 **Did you know?** Innovations that seem sudden are often the visible part of much longer research. The 2017 Transformer (a neural-network architecture based on an attention mechanism), for example, arises within a history of neural networks, language representations and parallel computing; it does not erase what came before, but makes a new scale possible. [3]

2. The breakthrough that changed everything

The recent breakthrough arises from the convergence of different elements. Deep learning uses neural networks with many layers to derive representations from large numbers of examples. Training them requires mathematical operations repeated billions of times. **GPUs** (graphics processing units), created to process many graphics operations in parallel, also proved suitable for this work: they are not "intelligent chips", but accelerators able to perform repeatable calculations quickly on large amounts of data.

The availability of digital data and cloud infrastructure (computing resources available over a network) completed the picture. **Training** a model means showing it many examples and iteratively adjusting its parameters so that it produces better predictions. This is not equivalent to transferring an entire ordered library to it, nor does it guarantee that every future answer will be correct. Quality depends on the data, training objectives, checks and the context in which the model is used.

In 2017 the research paper Attention Is All You Need presented the Transformer architecture. Its attention mechanism allows the model to weight relationships between different parts of a sequence: in a sentence, for example, it helps connect a word to terms that clarify its meaning. The architecture is particularly suited to parallel computing and has contributed to the growth of large-scale language models. [3]

An **LLM** (large language model) is a system trained to predict the next part of a text sequence and then fine-tuned to follow instructions. Translation, summarization, programming and dialogue emerge from this predictive ability, but so do convincing errors: the model can construct a grammatically sound statement without a verified basis to support it. This is why "it can write" does not mean "it can guarantee truth".

ChatGPT made this transformation cultural. From November 2022, a natural-language conversation became for millions of people the most immediate way to encounter a generative model. The impact does not lie only in technical performance: AI moved from being laboratory material or a hidden function in an application to becoming an accessible interlocutor. [4]

🧠 **Radar Insight** The simplest interface can conceal the most complex machine: data, models, chips, memory, networks, energy, cloud, capital and human labour. Democratizing use matters; it must not make us forget where control of infrastructure is concentrated.

3. When AI began to converse

A **chatbot** responds to a request. An **assistant** can keep track of several requests, read materials supplied by the user and adapt the form of its response. A **collaborator works** towards a shared objective: it proposes alternatives, uses authorized tools and makes steps and limits visible. Finally, **an agent** organizes a sequence of actions—searching, reading, calculating, filling in, passing the result to another application—within boundaries defined by people and control systems.

These are useful categories for understanding gradual change, not magical levels of autonomy.

Multimodality means the ability to process multiple kinds of input and output, such as text, images, audio and documents. GPT-4o made this transition more visible by integrating capabilities across text, vision and audio; its safety documentation also recalls that such capabilities require specific assessments, especially for voice. [5] *Contextual memory*, the ability to use information present in the conversation or available materials, makes dialogue more continuous but does not turn the system into an infallible archive of the user's life.

In these first days of July, OpenAI presented **GPT-Live**, *a voice system with continuous interaction: it can listen and speak at the same time and delegate tasks requiring research or reasoning to supporting models*. This is a significant step because it reduces the rigidity of turn-taking, not because it constitutes *"the first conversational AI"* in history. The continuity of voice makes the interface more natural; precisely for this reason, the sensation of presence must be distinguished from human understanding. [6]

GPT-5.6, made available on 9 July 2026, is also significant not as a definitive milestone, but as an example of models designed for longer tasks, computer use and coordination of multiple agents. Its availability in ChatGPT, Codex and APIs (interfaces that allow different programs to communicate and use functions) shows how the line between conversation and operational work is thinning. [7] Versions will change; the trend to observe is the combination of language, tools and context.

The landscape is not dominated by one actor. **Anthropic's Claude** has developed tool-use and code-execution functions for agentic workflows; **Google DeepMind's Gemini** works on multimodal reasoning; **Mistral** offers open models and services also distributed on controllable infrastructure; **DeepSeek** documents the use of external functions through APIs. [8][9][10][11] Differences between products, licenses, data location, capabilities and terms of use matter more than cheering for a brand.

Conversation, then, is an interface. It can facilitate access, support learning and make software less hostile. It does not prove that the system has intentions, empathy or authority. The responsibility not to anthropomorphize it is also a form of protection: it prevents us from asking a machine to assume what only a person, institution or verifiable source can take responsibility for.

From chatbot to agent

Form	What it mainly does	What must remain visible
Chatbot	Answers questions and generates content.	Sources, limits, possibility of error.
Assistant	Maintains context and helps organise a task.	What data it reads and retains.
Collaborator	Uses tools and works on an assignment with the user.	Objective, approvals, record of changes.
Agent	Executes sequences of actions within defined boundaries.	Permissions, stopping, verification and final responsibility.

4. When AI began to collaborate

Here the change of scale is more concrete. A **generative model** can produce a draft; an **agentic system** can try to carry a task forward. To do so, it combines a model with tools: web search, document access, code execution, spreadsheets, a development environment, application interfaces. **The agent does not *"enter the world"* on its own**: it acts only through permissions and tools someone has made available to it. **This is a decisive distinction because it shifts attention from stories about the machine to boundaries designed by human beings**.

Codex is a useful example of this transition. In its contemporary form it is presented by OpenAI as an agent for writing, reviewing and delivering code; it can read a project, modify files and use a development environment, with levels of approval and access controls. [12] Its interest is not limited to programming. It shows how a system can work within an operational context rather than merely suggest lines to copy into a chat window.

The same logic extends to research, design and document management. A well-configured assistant can collect materials, compare versions, extract points to verify, prepare an initial structure and make tedious tasks repeatable. In a **company** it can help piece together scattered information, provided access to data is legitimate and the process traceable. In **medicine** it can support document synthesis and administrative workflows, but not replace clinical assessment, informed consent or professional responsibility. In **education** it can offer personalized explanations and exercises; it cannot decide by itself what is educational nor replace the teaching relationship. This is *augmented intelligence*: the human being retains direction, intention, judgment, critical thinking and responsibility; the system expands speed, analytical capacity, organization and operational possibilities. The formula is not a rhetorical compromise. It is a practical criterion for deciding where to delegate and where not to. Delegating a search for a citation may be reasonable; delegating without oversight the sending of a sensitive communication, a healthcare decision or an administrative act may not be.

A **reliable agent** must therefore make at least five things controllable: **objective**, meaning the requested result; **sources and data**, meaning what it may read; **tools and permissions**, meaning what it may do; **action trace**, meaning what it has done; stopping point, meaning when it must ask for confirmation. Agent platforms are building precisely around these elements: search tools, files, computer use (using a computer through its interfaces) and execution tracing. [13]

> 🧠 **Radar Insight — Collaboration does not mean surrendering control** Human oversight is not the residue of a still-imperfect system. It is a condition of quality. As the autonomy and impact of an action increase, so too must clarity of permissions, the ability to interrupt, independent checks and the responsibility of whoever approves the result.

In **creativity,** collaboration can speed up experimentation: variants, drafts, translations, classifications and prototypes. The risk is confusing quantity with direction.

In **research**, it can help explore literature and data; the risk is turning a plausible synthesis into proof.

In **programming**, it can speed up repetitive parts; the risk is introducing vulnerabilities or changes that are not understood.

In every case the right question is the same: who can verify the work and who answers if something goes wrong?

5. The cloud has weight

“Cloud” is a very light word for a physical system. Behind an answer generated in a few seconds are data centers, servers, networks, cooling systems, uninterruptible power supplies and global industrial supply chains. Contemporary models require computing capacity for training and, once deployed, for inference: the moment when they process a real request.

GPUs and other AI accelerators perform much of this computation. **HBM** (high-bandwidth memory) is designed to transfer data very quickly between memory and processor; without an adequate flow of data, even a powerful chip waits. **NVIDIA** and **AMD** design accelerators; **TSMC** and Samsung are significant nodes in advanced manufacturing; **Micron** and other memory makers compete in essential components. This is not a list of logos or big names: it is a chain in which design, manufacture, materials, production capacity and networks determine who can build and use systems at scale.

Then come the cloud and data centers. **Microsoft, Google, Amazon, Meta, OpenAI, Anthropic** and numerous specialized operators invest in computing capacity in different forms: direct ownership, cloud contracts, partnerships, shared infrastructure. The interface can be global and simple; major investments, by contrast, are concentrated. Stanford’s AI Index report notes that almost 90% of notable models in 2024 came from industry, a figure that invites us to distinguish the spread of tools from the distribution of the power to develop them. [14]

Energy is part of the same story. The **International Energy Agency** stresses that there is no AI at scale without electricity for data centers and that their growth raises questions of grid availability, location, efficiency and planning. [15] It is therefore not enough to ask how much energy “AI” consumes in the abstract: the model, number of requests, type of hardware (physical components of computing systems), data centre efficiency, local electricity mix, hours of use and how capacity is built all matter. Water and cooling also require attention to local conditions, not slogans.

Infrastructure is geopolitical because it is not distributed evenly. Controlling advanced chips, memory, manufacturing capacity, stable energy, networks and cloud services means influencing possibilities for research, defense, industry and public services. At the same time, reducing everything to a race between states would be misleading: supply-

chain interdependence is high and the decisions of companies, regulators and local communities matter as much as strategic declarations.

The material ecosystem of AI

Operational AI is not at the centre of an abstract cloud, but of a system of dependencies:

models ↔ chips/GPUs ↔ HBM memory ↔ data centers ↔ energy ↔ cloud and networks ↔ data ↔ users ↔ regulation

Each arrow carries a question: who owns the resource, under what rules, with what transparency, and who can access it? This is where technology, economics and freedom meet.

6. Real opportunities

AI creates value when it enters a process with a clear objective, carefully managed data, people able to interpret the output and an identifiable person responsible. This is not a theoretical condition: *a model can make a good process faster or speed up errors, prejudices and unnecessary steps. Automating a flawed process does not improve it; it merely makes it run faster*.

In scientific research, generative systems can help organize literature, formulate hypotheses, write exploratory code and analyze large volumes of data. Their value increases when results, data and methods remain reproducible. In **medicine** they can ease administrative tasks, support documentation and make some information more accessible; clinical use, by contrast, requires validation in context, privacy, risk assessment and professionals able to challenge the system. Stanford's AI Index pays increasing attention to research and medicine precisely because the point is not only technical capacity, but the quality of evidence and adoption. [16]

In **education** a conversational tutor can offer explanations, alternative examples and immediate feedback. It is particularly useful if it encourages the student to explain their reasoning, cite sources and compare answers. If it becomes a shortcut to avoid effort, it can weaken learning. Schools should not choose between prohibition and fear on one side and blind enthusiasm on the other: they must teach when to use a tool, how to verify it and how to disclose its use.

For **businesses and industries**, value often emerges in ordinary activities: extracting data from documents, summarizing steps, helping technical support, searching internal information, generating software tests, making complex procedures accessible. The hardest phase is not buying a model. It is redesigning workflows, deciding which data may circulate, training people and measuring whether the result is truly better. Responsible adoption is not a separate department: it is organizational design.

In **public administration** the potential concerns simpler language, accessibility, translation, document management and citizen-oriented services. Here the margin for error is often less tolerable, because a wrong answer can affect rights and access to benefits. Appeals procedures, traceability, human alternatives and oversight

📌 **Did you know? — Where value really arises** 1. Clear objective: the system addresses a defined problem, not a generic hope. 2. Competence: those using it recognize useful results, errors and exceptional cases. 3. Oversight: controls and responsibilities are part of the process, not an afterthought. 4. Organization: data, roles and procedures make it possible to turn a trial into stable improvement.

Accessibility deserves a central, not peripheral, place. Speech synthesis, image descriptions, language simplification, subtitles and translations can expand many people's autonomy. Here too, however, a useful function does not exempt us from verification: an incorrect description or imprecise translation can create new barriers. Technology expands possibilities when it is designed with the people involved, not when it is simply installed on them. This translates into "amplified thinking".

7. Real risks and constructed fears

The first concrete risk is plausible error. So-called hallucinations are invented or inaccurate answers produced in a confident tone. They are not quirks of the machine, but a predictable limit of systems that generate probable sequences without an intrinsic mechanism guaranteeing truth. For important information, the rule is simple: trace it back to the primary source, check dates and context, and do not treat a well-written answer as a definitive document.

Then there are deepfakes: *synthetic images, audio or video capable of imitating people and situations*. They can fuel fraud, defamation and manipulation. The remedy is not to

believe every item is false, but to build habits of verification, content provenance, official channels and training. Similarly, AI can increase the scale of certain harmful cyber activities, but can also assist defense, vulnerability analysis and incident response. The point is who can use the tools, within what limits and with what controls.

Privacy and surveillance are not synonyms. An organization can process data excessively, opaquely or discriminatorily even without a generative model; AI can make the analysis of that data cheaper and more pervasive. Biometrics, automated classification of people and high-impact decisions require particular attention to fundamental rights. The AI Act, which entered into force in 2024, introduces a European risk-based framework and bans on some practices considered unacceptable. [17]

Bias (systematic distortion) is not merely a technical defect: it can enter data, problem definitions, metrics (criteria and measures used to assess a result) and decisions by those using the output. Correcting a model without questioning the process may be insufficient. Likewise, opacity does not concern code alone: it concerns contracts, training sources, evaluation criteria, the possibility of challenging a decision and concentration of infrastructure.

Work will change, but not in a single direction. Some activities will be automated, others will change form, and still others will emerge around oversight, integration, quality and security. Saying that “everyone will be replaced” is too simple a forecast; saying nothing will change is equally so. Consequences will depend on contracts, training, work organization, distribution of productivity gains and institutions’ ability to accompany transitions, but above all on each person’s ability to understand the potential of this great opportunity.

Cognitive dependence is a less visible risk: if a person always delegates writing, research and judgment, they may lose practice and autonomy. But not every use of a tool reduces skills. A calculator can free time for more complex problems; it can also prevent learning the basics if used too early. The difference lies in the design of the experience and

⚠️ **To watch — The same technology, different outcomes** A system that summarises documents can help an official find their bearings or conceal a decisive error. A voice system can make a service accessible or collect unnecessary data. Outcomes depend on purpose, governance (rules, powers and decision-making processes), access, transparency, competence and democratic control.

Constructed fears deserve consideration. The machine described as a subject with a will of its own shifts responsibility away from designers, companies, governments and users. The idea of total and immediate replacement erases differences between jobs, rules and contexts. The opposite idea (that the future is already decided by technology) is a convenient abandonment of politics. Maturity lies in rejecting both naïveté and fatalism.

8. The next evolution concerns us

The skill of the future will not be merely knowing how to use a model. It will be knowing how to formulate a problem, distinguish a source from an opinion, check a result, protect data and people, and understand when not to automate. It will also mean recognizing that a quick answer is not always a good decision and that an efficient system is not necessarily fair. To name just a few.

For organizations the test is not adopting the latest novelty, but building contexts in which people working there can flag errors, understand limits, refuse improper uses and contribute to the design. For institutions it means investing in literacy, infrastructure, open research, enforceable rules and oversight capacity. For people it means not giving up curiosity, verification and responsibility merely because the interface has become pleasant.

AI will continue to evolve. There is no need to pretend to know every form it will take. We need to decide which principles must remain: *plurality of sources, dignity of work, protection of rights, access to knowledge, and the ability to understand and challenge decisions that affect people's lives*.

For decades we asked whether machines would become intelligent like us. Today a more useful question may be this: **will we become mature enough to collaborate with increasingly capable systems without giving up the judgment, responsibility and freedom that make us human?**

Verified sources and further reading

1. Alan M. Turing, *Computing Machinery and Intelligence*, **Mind**, 1950. https://doi.org/10.1093/mind/LIX.236.433
2. John McCarthy, Marvin Minsky, Nathaniel Rochester, Claude Shannon, *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*, 1955. http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf

3. Ashish Vaswani et al., *Attention Is All You Need*, NeurIPS, 2017. https://arxiv.org/abs/1706.03762
4. OpenAI, *Introducing ChatGPT*, 30 November 2022. https://openai.com/index/chatgpt/
5. OpenAI, *GPT-4o System Card*, 2024. https://openai.com/index/gpt-4o-system-card/
6. OpenAI, *Introducing GPT-Live*, 8 July 2026. https://openai.com/index/introducing-gpt-live/
7. OpenAI, *GPT-5.6: Frontier intelligence that scales with your ambition*, 9 July 2026. https://openai.com/index/gpt-5-6/
8. Anthropic, *Introducing Claude 4*, 2025. https://www.anthropic.com/news/claude-4
9. Google DeepMind, *Gemini Robotics brings AI into the physical world*, 2025. https://deepmind.google/blog/gemini-robotics-brings-ai-into-the-physical-world/
10. Mistral AI, *Models documentation*. https://docs.mistral.ai/models/
11. DeepSeek, *Function Calling — API Docs*. https://api-docs.deepseek.com/guides/function_calling/
12. OpenAI, *Using Codex with your ChatGPT plan and Introducing upgrades to Codex*, 2025–2026. https://help.openai.com/en/articles/11369540-using-codex-with-chatgpt ; https://openai.com/index/introducing-upgrades-to-codex/
13. OpenAI, *New tools for building agents*, 2025. https://openai.com/index/new-tools-for-building-agents/
14. Stanford HAI, *The 2025 AI Index Report*, Research and Development chapter. https://hai.stanford.edu/ai-index/2025-ai-index-report/research-and-development
15. International Energy Agency, *Energy and AI*, 2025. https://www.iea.org/reports/energy-and-ai
16. Stanford HAI, *The 2026 AI Index Report*. https://hai.stanford.edu/ai-index/2026-ai-index-report
17. European Commission, *AI Act enters into force*, 1 August 2024. https://commission.europa.eu/news-and-media/news/ai-act-enters-force-2024-08-01_en

Methodological note: information on products and models was verified against official primary sources as of 11 July 2026. Product features may change; the text focuses mainly on structural transformations, not performance rankings.

© 2026 Barbara Gattobigio · Gruppo Radar

Comprendere oggi. Costruire domani · by effettoantropico.life

All rights reserved on original content. Free access does not authorize reproduction or republication without consent, except for uses permitted by law. Third-party materials remain the property of their respective owners. To share this issue, use the link to the article.