From answer to result

How we really use artificial intelligence — and what distinguishes access, habit and capability



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

An answer arrives. The work, however, has yet to begin

The scene has become ordinary. A professional opens an artificial intelligence assistant and asks for a summary, a table or a recommendation. A few seconds later, an orderly text appears on the screen: clear headings, a confident tone, a ready conclusion. Speed creates a powerful impression: *the task already seems solved*.

Then the real work begins. A cited source does not quite support the sentence. One column uses a different denominator from the others. A technical term has been simplified until its meaning has changed. The conclusion depends on an assumption no one has stated. The answer is not necessarily useless: it may be excellent first material. But it is not yet a verifiable result.

It is in this distance that the most important part of AI adoption is at stake.

Not only between those who use a tool and those who do not, but between access and capability; between an occasional trial and a stable process; between the time apparently saved and the value that remains after errors have been checked.

Surveys describe rapid diffusion. Experiments show real gains in well-defined tasks. Other equally solid studies document slowdowns, errors or weaker learning when the tool is used beyond its frontier or replaces the reasoning it should support.

This is not a contradiction to eliminate. It is the phenomenon to understand.

The right question, then, is not only: "How many people use artificial intelligence?" It is: **what are they really doing, with what depth, under whose responsibility and with what measurable result?**

> 🧠 **Radar Insight — The decisive distance**
>
> Access makes it possible to obtain an output (a result produced by the system). Capability transforms that output into a decision, document, service or learning outcome whose quality, limits and consequences someone knows how to verify.

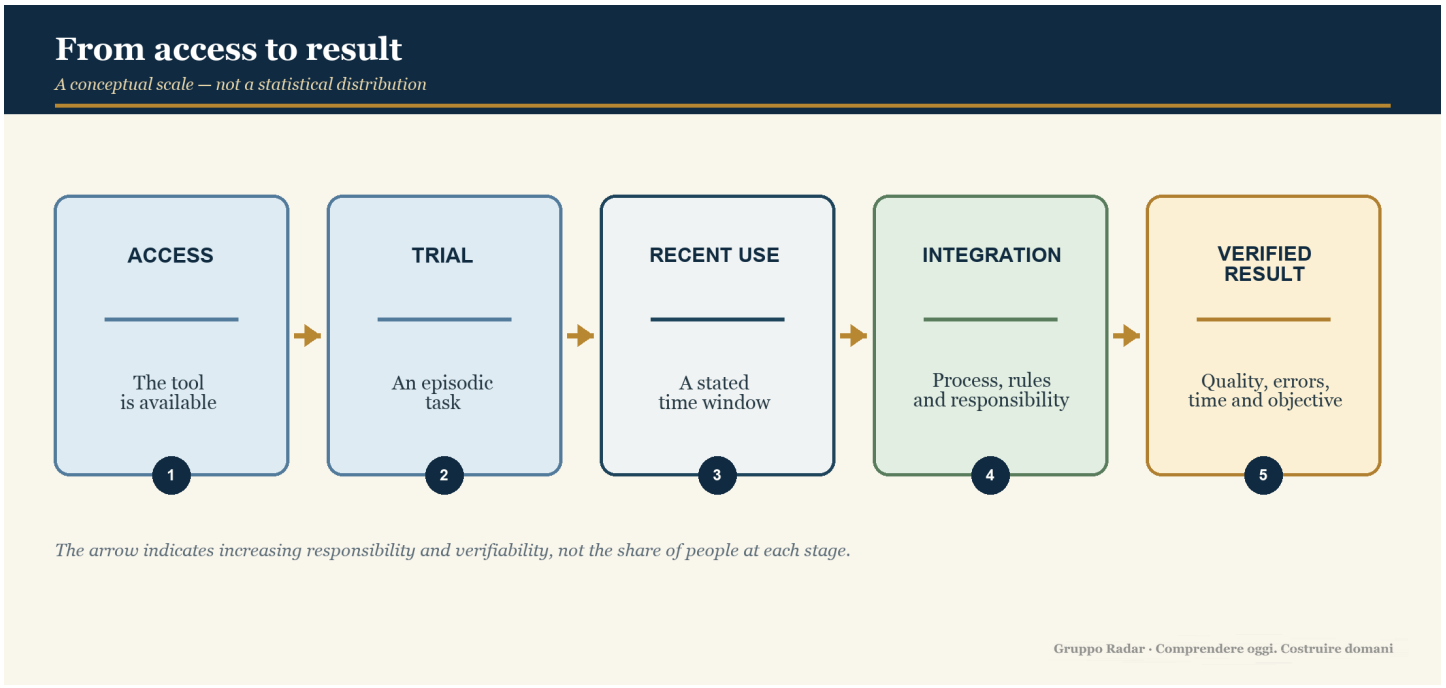
1. The word "use" has become a misunderstanding

Saying that a person "uses AI" can mean many things. They may have tried a chatbot (a program with which one converses in natural language) once. They may encounter automatic recommendations every day without recognising them. They may entrust a first draft to a copilot, check it and rewrite it. They may work in a company that has enabled an experimental feature, or in an organisation that has redesigned a process,

established responsibilities, protected data and defined metrics (indicators used to measure results) before and after introducing the tool.

These conditions are not interchangeable degrees of the same statistic. They are different objects.

In the research supporting this Special, I separated at least five conceptual stages: access, trial, recent use, integration and verified result. They are not a moral ranking and do not describe the actual distribution of the population. They help prevent the verb I am placing at the centre, "to use", from hiding the distance between being able to open a tool and knowing how to incorporate it into an activity for which one remains responsible.



From access to result — a conceptual scale, not a statistical distribution

The first stage is technical: the tool is available.

The second is episodic: it is tried.

The third requires an explicit window—for example the past three or twelve months—and measures reported behaviour.

The fourth enters the process: use has continuity, data, rules, people responsible and a way to correct. **Only the fifth asks what is often missing**: has the result been checked against time, quality, errors, costs and the original objective?

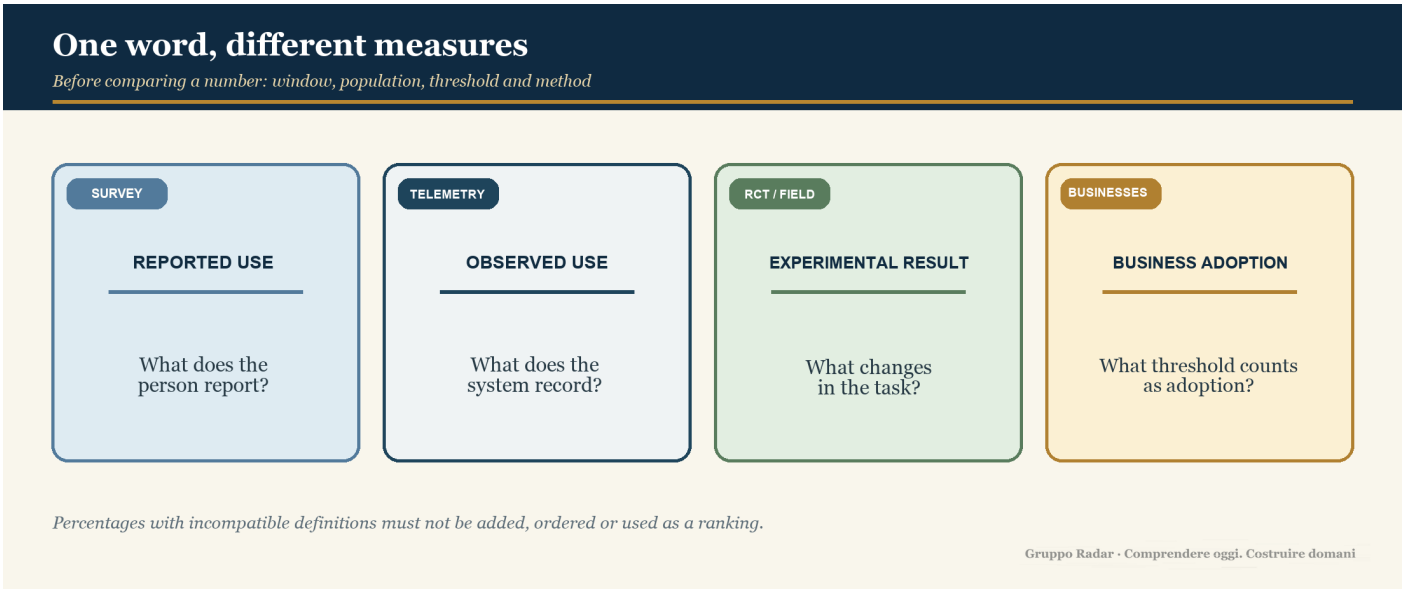
Most public statistics legitimately reach the first three stages. We know much less about the last two. Prudence does not mean devaluing the available data, but not attributing to it what it does not measure.

2. One word, different measures

In the United Kingdom, 59% of adults report recent use of generative AI; among users who responded online, 75% report using it at least weekly and 35% daily (E03). In the United States, 34% of adults report having ever used AI; the share rises to 58% among those under 30 and falls to 10% among those over 65 (E04).

It would be tempting to put the two figures in a ranking, but that would also be incorrect. In a consistency analysis, the values bear the weight of their differences. The tool observed, time window, sample, threshold of use and survey method all change. “Recent use of generative AI” is not equivalent to “having ever used AI”. A percentage without its own definition is a number that has lost its meaning. The same caution applies when AI is embedded in services that people do not call “artificial intelligence”. In Great Britain, a broad question including assistants, recommendations and facial recognition finds that 17% of adults report using these functions often or always, 50% sometimes or occasionally, and 33% hardly ever or never (E05).

In the United States, an observed browsing panel shows that 93% encounter at least one page mentioning AI in a month, while only 13% visit a chatbot website at least once (E22). Encountering AI, recognising it, visiting a tool and using it intentionally are not the same behaviour.



One word, different measures — windows, populations and thresholds

3. In Italy: non-use has more than one cause

In 2025, 19.9% of people in Italy aged 16 to 74 reported having used generative AI tools in the previous three months. Among people aged at least 14, the share was 17.8%; it reached 51.2% among 14–19-year-olds and 43.1% among 20–24-year-olds (E01).

The figure shows use remains a minority practice and varies sharply by age. But the more interesting side is not only who uses it. Among non-users, 59.9% said they had no need for it, 21.6% cited insufficient skills, 5.5% privacy or security reasons and 4.8% lack of knowledge of the tools (E01).

These answers do not permit a single diagnosis, but their meaning must be examined.

“I do not need it” may express an informed assessment, limited knowledge of the possibilities, an activity genuinely ill-suited to it, or the absence of a perceived benefit arising from various conditions, often cultural as well.

A lack of skills, meanwhile, is different from deliberate refusal; concern about data is another limit, further distinct from economic or practical exclusion.

Grouping everything under fear would be convenient but ethically wrong. European comparison also requires care.

In the EU-27, 32.7% of people aged 16 to 74 reported having used generative AI tools in the last three months of 2025: 25.1% for personal purposes, 15.1% for work and 9.4% for education, with categories that may overlap (E02). Collection is harmonised and allows the same indicator to be compared.

Even this survey, however, does not measure quality of use, frequency, verification or benefit produced. Turning it into a ranking of “national capability” would attribute to it a question that has not yet been addressed.

The first gap is visible: access and reported use vary with age, education, occupation and territory. The second is harder to measure, even among those accessing the same tool. Often the implication is that not everyone has the same time, professional context, training or knowledge.

> 📌 **Did you know? — Reported, observed, experimental**
>
> A survey records what a person or business reports. Telemetry records activity within a product. An experiment measures an effect under defined conditions. A business statistic applies an adoption threshold. None of these perspectives is “the whole truth”: each answers a different question.

4. Everyday life is part of the data, but meaningful outcomes remain poorly measured

In Great Britain, over the past twelve months, 41% of adults report conscious use of AI for work or education, 27% at home and 23% for hobbies or interests; 33% report no

conscious use. The areas can overlap (E24). This snapshot broadens the view beyond the office: household organisation, learning, personal interests and everyday activities are already part of the experience.

But knowing where a tool is used does not tell us whether the day has become better. On a comparable basis, we do not know how much AI increases personal autonomy, improves decisions, eases a care burden or makes a service more accessible. Nor do we know the cost of verification, the effect of dependence on a platform or what happens to the ability to perform the activity without assistance.

OpenAI's telemetry, for example, offers a different and useful view, provided it stays within its boundaries: in the consumer sample analysed, around 30% of messages are classified as work and 70% as non-work; intent categories are Asking (questions) 49%, Doing (requests) 40% and Expressing (conversation) 11% (E15, provider source).

The data describes conversations included in that product, not the population, non-users or value obtained. A message on a personal topic is not automatically well-being; an operational request is not automatically a result.

In the United States, 62% report interacting with a broad category of AI at least several times a week, but 57% feel little or no control and 61% would like more (E26). Here too, interaction and autonomy do not coincide.

AI can be frequent without becoming a governed capability.

5. Work: between perceived time and measured results

An important part of professional adoption passes through perceptions. In the United States, 25% of workers report having used generative AI at work in the past month; among all workers, 44% agree that it saves time and 24% that it improves quality (E12). In another recurring online survey, reported workplace use reaches 37.4% and self-reported savings are estimated at 1.6% of total hours (E14).

These are important pieces of information: they describe diffusion, expectations and perceived usefulness. On their own they do not demonstrate that net output has increased, that quality has improved after review or that freed time has reduced workload. In many organisations, time saved at one stage may be absorbed by more requests, new checks or a faster pace. This persistent operational density often arises from a serious lack of clarity about workflows, management, assessment of value and the weight of leadership (organisational guidance and responsibility) and its levels.

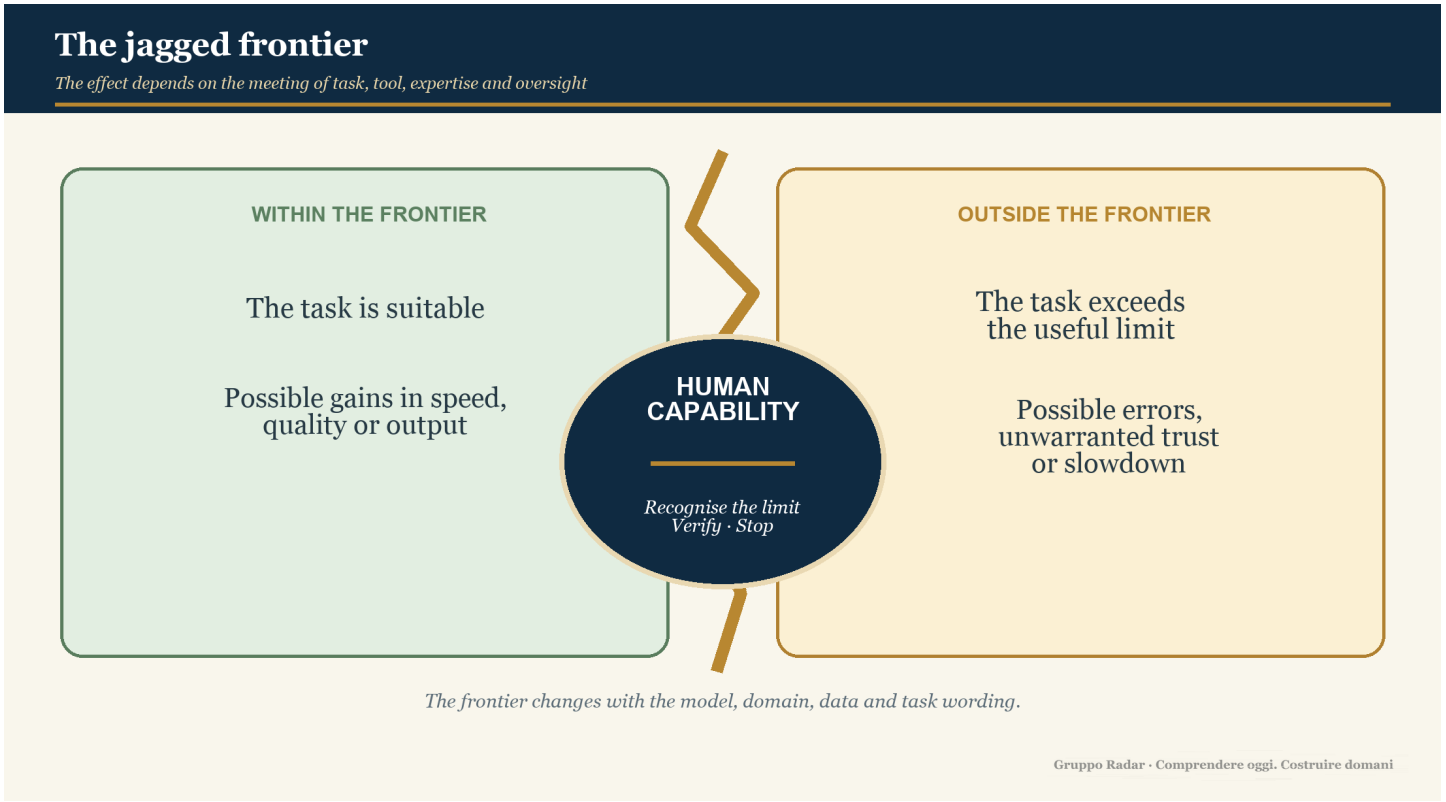
Experiments conducted so far allow some tasks to be observed more closely. In a controlled professional writing test, average time falls from around 27 to 17 minutes

and average quality, on the evaluators' scale, rises from 3.79 to 4.54 (E17). In a customer service setting, a generative assistant is associated with around 15% more problems resolved per hour, with larger gains for less experienced workers (E19).

These results apply within the contexts studied. They do not prove that every document is better, that an entire profession produces more or that gains automatically become income, profit or well-being. They show something more precise and useful: when the task is suitable and the metric is clear, AI can produce measurable benefits. This indicates the importance for every user of learning to use this technology to the fullest potential it offers, including for their own benefit.

6. The jagged frontier

The critical point emerges when the task changes. In an experiment with consultants, within the frontier of the tool's capabilities users complete 12.2% more tasks, work around 25% faster and achieve higher quality. Outside the frontier, the probability of providing a correct answer is 19 percentage points lower than in the control group (E18).



The jagged frontier — the limit recognised by the person

The frontier is not a fixed, visible line in the collaboration under way. It changes with the model, the domain, the data provided and the wording of the task. For this reason, the decisive human capability is not only obtaining a more refined answer. It is

recognising when a plausible answer is crossing a limit that requires a different check, an expert or stopping the process.

The contrast between studies of developers makes the point clearer. Three company experiments estimate, in aggregate, around 26% more tasks completed with assistance, despite heterogeneity and imprecise individual estimates (E20, study flagged for provider involvement). In an experiment with experienced open-source developers (working with code accessible according to its licence), by contrast, AI increases completion time by around 19%. Before the test participants expected a 24% saving; afterwards they still perceived a benefit close to 20% (E21). The results do not cancel each other out.

In mature projects, familiarity with the code and the cost of understanding or correcting external suggestions can change the sign of the effect. The divergence between measured time and perceived benefit reminds us that satisfaction with use is not a sufficient metric.

An experiment involving 7,137 knowledge workers in 66 companies adds another nuance: treated users spend around two fewer hours on email each week and work less outside working hours, but no change emerges in the overall quantity or composition of activities (E30, study flagged for provider involvement).

The work pattern changes; transformation of the organisation, at least within the observed horizon, is not demonstrated. The cultural factor still carries old models, old habits and processes that are no longer sustainable.

7. The capability that remains when the tool is absent

The most delicate boundary concerns learning.

In a school experiment, during practice a basic GPT raises results by 48% and a tutor equipped with guardrails (protections designed to guide and limit the system) by 127%. In a subsequent test without the tool, the group that had used the basic GPT scores around 17% below the control; the tutor with guardrails mitigates the harm without producing an advantage (E27).

This makes clear that assisted performance and acquired capability are different outcomes. If the tool reduces the effort needed to build a mental procedure, the immediate result can improve while learning weakens. This is not a general condemnation of AI in education: the comparison between basic GPT and a tutor with protections shows precisely that interaction design matters, as does understanding how a tool should actually be used.

Creativity also requires more than one metric. In a writing experiment, generative suggestions improve individual ratings on average, especially among participants who were initially less creative, but make stories more similar to one another (E28). This reveals a lack of understanding of how that advantage should actually be used. The principle to understand clearly is that an individual benefit and a collective cost can coexist, but asking only whether the individual text "is better", without the user's discernment, leaves out the overall diversity of the work produced. And the problem is not "AI does not work", but "am I really capable of using it?"

I go deeper and add another operational question: what remains with the person? Can they explain the decision, correct the error, transfer the method to a new case, choose not to delegate? A useful system need not make people autonomous in every activity; it must, however, make visible where assistance ends and the user's own responsibility begins.

> ⚠️ **To observe — Faster does not always mean further ahead**
>
> The advantage gained from knowing the tool depends on the meeting of task, tool, expertise, human supervision and measurement. Without before-and-after measures and checks in the working and cooperative flow between user and tool, errors lead to rework. The initial speed of the output can be mistaken for productivity, whereas what matters is the productivity achieved with the knowledge acquired.

8. Businesses: adoption does not mean integration

AI is spreading rapidly among businesses. In the EU-27, the share of enterprises with at least 10 employees using at least one AI technology rises from 8.1% in 2023 to 13.5% in 2024 and 20.0% in 2025 (E06). In Italy it reaches 16.4% in 2025, compared with 8.2% in 2024 and 5.0% in 2023; it reaches 53.1% among large enterprises and 15.7% among SMEs (E07).

"At least one technology" is a useful threshold for measuring diffusion. It does not say how extensively the tool is used, whether it has moved beyond experimentation, which activities it has changed or what return it has produced.

The barriers reported by Italian enterprises that considered but did not adopt AI help reveal the complexity: 58.6% cite skills, 47.3% legal issues, 45.2% data availability or quality, 43.2% privacy and 43.0% costs, with multiple answers possible (E07). In Canada, among adopting enterprises, 44.4% report changes in training or staffing and 32% staff training (E11). Access to software is only one component; usable data, rules, people and responsibilities are needed.

Statistical definitions also change the picture considerably. In the United Kingdom, a survey counting technologies implemented in methods or processes, excluding tests, estimates adoption at 9% in 2023 (E08). A more timely survey with a broader definition of use estimates around 35% among enterprises with at least 10 employees in June 2026, 28% among those with 0–9 employees and 49% among those with at least 250 employees (E09). The difference is not a pure measure of growth: the year, threshold and inclusion of light uses change (when the user mistakes AI for a simple chatbot and limits use to that interaction alone).

Depth remains limited even when adoption is reported. Among British enterprises with at least 10 employees, 18% use large language models (language models trained on large amounts of text) and 16% tools for visual content; among adopters, only 10% describe their use as extensive. Of all enterprises with at least 10 employees, 15% report that more than half of their employees use AI in daily work (E09). In the United States, the BTOS supplement estimates 18% of enterprises are users, and 32% when the share is weighted by employment; among adopters, 57% use AI in no more than three functions (E10).

These denominators are not interchangeable: the share of enterprises, the share weighted by employment, the number of functions and the portion of the workforce answer different questions.

The common lesson is more restrained: switching on a tool, using it in one function and redesigning a process are distinct steps.

The geographical extensions confirm some gaps but do not establish a ranking. In Canada, 19.2% of enterprises report using AI in the previous twelve months to produce goods or provide services; the share is 21.0% in urban areas, 9.9% in rural areas and 27.8% among enterprises with at least 100 employees (E11). In Australia, 12% of employing businesses report use in the 2024–25 financial year (E13). Definitions and survey designs do not permit these countries to be ranked alongside the EU, the United Kingdom and the United States.

9. Fear for jobs: a cultural and social fact, not a forecast

Among US workers, 52% say they are worried, 36% hopeful and 33% overwhelmed about the future use of AI at work; 32% expect fewer opportunities and 6% more opportunities (E23). In Great Britain, 36% of adults expect benefits and 38% see more risks than benefits; 47% of 16–29-year-olds and 49% of 30–49-year-olds think AI could put their job at risk, but the question is also asked of people who are not employed (E24).

Fear is both a data point and a fact. It can influence training choices, trust and willingness to experiment. It should not be seen exclusively as ignorance, but neither should it be turned into an employment forecast.

A perception measures expectations and sense of control; it does not count jobs lost.

The first data on observed outcomes present a less dramatic but still incomplete picture. In Denmark, in the first two years studied, no detectable average effects on income or hours emerge; the authors rule out average effects greater than around 2% within their interval and observe a reorganisation of tasks and new AI-related activities (E29). In Canada, between 2015 and 2025, no persistent and disproportionate employment decline appears among the most exposed groups; the analysis is descriptive and does not identify a causal effect of AI (E31).

Not seeing an aggregate collapse in the short term does not mean no one loses power, income or opportunity. National averages can conceal costs concentrated by occupation, enterprise or territory. In Canada, 22% of workers report using generative AI at work in the previous twelve months; the share is 44% in jobs requiring a university degree, 10% in those requiring a diploma and 3% where no formal requirement is specified (E25). Professional access is already unequal; we do not yet know who will capture the value and who will chiefly experience intensification or replacement.

The strongest position is to hold two truths together: average effects observed so far do not support generalised replacement; the horizon is short and the distribution of costs and benefits remains insufficiently measured. At the same time, integrating the tool changes the pace and opportunities compared with remaining on the sidelines or demonising it through lack of knowledge. This gives rise to the sweeping claim: I do not have a job because an AI occupies my position. But the real question is another: why am I not integrating AI to try to improve my position or create new opportunities in relation to my potential?

10. What distinguishes aware use

There is no validated scale in the verified corpus (the set of documents examined) that allows us to count “aware” or “high-performing” users. These expressions must remain operational editorial categories, not new percentages.

The evidence does, however, allow concrete conditions to be identified.

Aware use begins with defining the task and the ability to ask whether it lies within the tool’s frontier. It provides the necessary context without exposing unnecessary data. It

verifies facts, calculations, sources and completeness. It distinguishes performance achieved with assistance from the capability that remains without it. It compares perceived time with measured time. It includes errors, rework, quality and costs in the metric. It knows when to stop or not use the tool. In professional workflows, it explicitly assigns supervision and responsibility.

These are not prompting tricks (ways of formulating requests sent to the model), now almost surpassed. They are transferable working practices: formulating an objective, recognising uncertainty, choosing a test, documenting a correction, making a decision open to challenge.

Prompting can help produce a better output; capability determines whether that output deserves to enter the world.

This distinction also prevents a common misunderstanding. The expert user is not necessarily the person who uses AI most often or most enthusiastically. It may be the person who uses it less, but at the right points; who measures the benefit; who protects data; who maintains integrated and aware cooperation; who knows how to attribute an error and correct the process.

11. What the data do not yet allow us to say

Credible research does not fill every gap. It makes it visible.

We cannot say how many new enterprises have been created thanks to AI. Adoption surveys describe technologies and functions, not the causality of business creation. It is not yet possible to quantify a lasting improvement in the quality of private life: areas of use and perceptions are known, but comparable longitudinal outcomes are not.

We cannot count the population using AI in an aware way, because a validated measure is lacking. We do not know the total long-term costs with sufficient confidence: subscriptions, verification, dependence on platforms, loss or consolidation of skills, interface changes and process sustainability. We cannot build a global ranking: for Japan, South Korea, India, much of Asia, Latin America, Africa, the Middle East and New Zealand, the audited corpus does not offer a sufficiently harmonised basis for this question.

Finally, we do not yet know the full distribution of effects. An average gain can coexist with concentrated costs; an individual benefit with less collective diversity. A reduction in workload through the use of AI may lead companies to shift work elsewhere, but the time made available does not necessarily have to become more work: it can also give a

person time to choose something better for themselves. For now, however, we have insufficient data to draw a much richer and more complete picture.

> ⚠️ **To observe — Insufficient data are a result, not a gap to hide**
>
> The corpus does not support causal conclusions about new businesses created by AI; verified benefits in private life; the prevalence of validated aware use; long-term costs and persistence; global rankings across non-harmonised geographies; or the full distribution of benefits and harms among social groups.

From answer to result

AI lowers the cost of the first answer. This is an enormous change: it makes drafts, analyses, code, explanations and alternatives accessible to people who would previously have needed so much time that it would have cost too much directly, or different skills. Experiments show that, within well-defined tasks, this access can become measurable speed, quality and output.

But it is worth emphasising again that access to AI does not automatically contain the method. A fluent answer can conceal a fragile definition. A perceived advantage can coexist with a real slowdown. Better performance with the tool can leave weaker learning without it. Discernment and an understanding of value are therefore fundamental, whether for personal, professional or business use.

The real capability lies in the next step: formulating the task, providing context, recognising the frontier, verifying, correcting, measuring and taking responsibility for the result. It is a human chain before it is a technological one.

This is why the future of AI use will not be read solely in the number of accounts, messages or enterprises reporting adoption. It will be read in the quality of decisions that remain after the initial enthusiasm; in the ability to stop a procedure when it does not hold up; in the possibility of explaining why an answer became a result.

The machine can deliver a real advantage, but there must be someone who truly wants to recognise and use that advantage constructively and proactively.

Essential sources

The complete research, operational definitions and limitations are documented in the canonical source and audited matrix. IDs E01–E31 refer to the original records; the primary sources actually cited in the text are listed here.

E01: Istat, Cittadini e ICT 2025.

E02: Eurostat, individual use of generative AI — isoc_ai_iaiu.

E03: DSIT, Public Engagement Survey 2025/26.

E04: Pew Research Center, use of ChatGPT.

E05: ONS, recognition of and attitudes towards AI.

E06: Eurostat, enterprises and AI — isoc_eb_ai.

E07: Istat, Imprese e ICT 2025.

E08: ONS, Management and Expectations Survey.

E09: ONS, AI in UK businesses 2023–2026.

E10: U.S. Census Bureau, BTOS AI supplement.

E11: Statistics Canada, AI in businesses.

E12: Federal Reserve, SHED 2025.

E13: Australian Bureau of Statistics, Characteristics of Australian Business.

E14: Real-Time Population Survey, GenAI adoption.

E15: OpenAI, How People Are Using ChatGPT — provider telemetry.

E17: Noy and Zhang, professional writing.

E18: Dell'Acqua et al., jagged frontier.

E19: Brynjolfsson, Li and Raymond, customer support.

E20: Cui et al., developers — study flagged for provider involvement.

E21: METR, experienced open-source developers.

E22: Pew Research Center, observed browsing.

E23: Pew Research Center, workers and AI.

E24: ONS, Public opinions and social trends, June 2026.

E25: Statistics Canada, use of GenAI at work.

E26: Pew Research Center, AI in everyday life.

E27: Bastani et al., learning.

E28: Doshi and Hauser, creativity and diversity.

E29: Humlum and Vestergaard, Danish labour market.

E30: Dillon et al., work patterns with M365 Copilot — study flagged for provider involvement.

E31: Statistics Canada, employment by exposure.

Gruppo Radar *Comprendere oggi. Costruire domani.*

© 2026 Barbara Gattobigio · Gruppo Radar

Comprendere oggi. Costruire domani · by effettoantropico.life

All rights reserved in the original content. Free access does not authorise 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.