LINKS GLOBAL
Communicating with Purpose. Powered by Insight.

THOUGHT LEADERSHIP SERIES

Critical Thinking in the Age of AI

The skill we are automating away at exactly the moment we need it most

A Links Global thought leadership article by M. Sanchez, founder of Links Global, a communications firm working at the intersection of trust science, narrative design, and civic engagement. Part of the Links Global article series: Contemporary Challenges. Drawn from 25 years of field practice and evidence-based research in strategic communication, public engagement, and institutional trust.

In Brief

Critical thinking is built through difficulty, and AI tools are very good at removing difficulty. This article explains what critical thinking actually is, the specific mechanisms by which AI use can erode it, why the consequence shows up at population level, and what deliberate preservation looks like in classrooms, workplaces, and daily information habits.

There is a paradox at the center of the current AI moment that the field's enthusiasts have not adequately reckoned with.

The technology that is most rapidly transforming how people access, process, and act on information is also the technology most likely to atrophy the cognitive capacity that makes people capable of evaluating information critically. And the communities most likely to experience this atrophy first (the ones who adopt AI tools earliest, most comprehensively, and with the least accompanying investment in the judgment frameworks that would allow those tools to be used well) are the communities of educated professionals, knowledge workers, and institutional decision-makers who most urgently need critical thinking capacity to do their jobs responsibly.

This is not a speculative concern. It is a straightforward consequence of how critical thinking actually develops, and how AI tools, deployed without deliberate countermeasures, systematically undermine that development.

What Critical Thinking Actually Is

Critical thinking is not skepticism. It is not the reflexive rejection of expert consensus, the performative questioning of established facts, or the ideological commitment to alternative interpretations. These are frequently mistaken for critical thinking. They are its absence: the adoption of a contrary posture as a substitute for the disciplined cognitive work that genuine critical thinking requires.

Critical thinking, in its technical sense, is the capacity to evaluate evidence, identify the assumptions underlying a claim, assess the quality of an argument's reasoning, recognize cognitive biases that may be distorting one's own judgment, distinguish between what the evidence supports and what one would like it to support, and revise one's conclusions when better evidence or better reasoning warrants revision.

It is one of the most cognitively demanding skills humans develop: demanding not just in its application, but in its development. Critical thinking capacity is built through practice: through the repeated experience of encountering claims, evaluating them, being wrong, discovering why, and updating. Through the friction of engaging with evidence that contradicts one's prior expectations. Through the discipline of tracing an argument's logic to its conclusions and examining whether those conclusions hold. Through the uncomfortable experience of sitting with genuine uncertainty rather than resolving it prematurely through the adoption of a confident position.

This development process takes time, effort, and the specific kind of cognitive engagement that is characterized by difficulty. The research on expertise development across domains is consistent: the skills that matter most are built through what cognitive scientists call desirable difficulties, challenges that are hard enough to require genuine engagement, produce errors that must be identified and corrected, and build the mental representations that underlie genuine competence.

AI tools, by design and by value proposition, reduce friction. They make difficult things easier. They accelerate the path from question to answer. They offload cognitive labor. These are genuine benefits when the labor being offloaded is genuinely mechanical, when the cognitive work is retrieval, formatting, translation, or summarization of well-established information.

They are a developmental problem when the labor being offloaded is the very difficulty that builds critical thinking capacity.

How AI Tools Undermine Critical Thinking Development

The mechanisms through which AI tools can undermine critical thinking development are specific and worth examining in detail, because awareness of the mechanisms is the first step toward deliberate countermeasures.

The answer-first structure. Large language models produce answers before reasons. They present conclusions confidently, in polished prose, with the rhetorical structure of authoritative explanation, often before the person asking the question has had an opportunity to think through the problem themselves. The student who uses an AI tool to answer a question before attempting to answer it themselves has bypassed the cognitive struggle that would have built the understanding the answer represents.

This is the educational equivalent of always looking at the answer key before attempting the problem. It produces the experience of learning without the actual development of capacity. The answer is right, or close to right, often enough that the person using the tool doesn't notice that they are not building the judgment to evaluate whether it is right, and therefore cannot recognize when it is wrong.

Confidence without calibration. AI systems produce confident outputs regardless of the quality of the underlying evidence. They do not express uncertainty in proportion to their actual uncertainty. They do not flag the domains in which their training data is sparse, the questions for which the evidence is genuinely contested, or the conclusions that are supported by some evidence but not others. They produce the linguistic markers of confidence (declarative sentences, specific claims, the absence of hedging language) in domains where confidence is not warranted as readily as in domains where it is.

Critical thinking requires calibration: the capacity to have confidence proportional to the strength of one's evidence and the quality of one's reasoning. A thinker who is regularly exposed to uncalibrated confidence (who gets confident answers to uncertain questions and has no mechanism to recognize the difference) is having their calibration trained in the wrong direction. They are learning to accept confident-sounding answers as a proxy for well-supported answers, which is precisely the cognitive pattern that makes people susceptible to misinformation, manipulation, and the characteristic failures of motivated reasoning.

The outsourcing of evaluation. The most consequential way AI tools undermine critical thinking is by outsourcing the evaluation function itself. A 2025 study of 666 adults in the United Kingdom found, in a correlational self-report survey, that heavier AI tool use went with lower critical thinking scores, and that cognitive offloading accounted for part of the relationship. The design cannot show that AI use caused the difference. When a person asks an AI tool not just for information but for a judgment (which of these arguments is stronger? what should I conclude from this evidence? is this claim well-supported?), they are asking the tool to do the critical thinking work that should be theirs.

The tool will produce an answer. The answer will often be reasonable, drawing on patterns in its training data that include genuine critical analysis. And the person who accepts that answer without examining its reasoning, checking its evidence base, or testing its conclusions against alternative interpretations has not thought critically. They have consumed a critical-thinking product. These are not the same thing, and the consumption of the product does not build the capacity to produce it.

The Population-Level Consequence

The individual-level consequence of AI-assisted critical thinking atrophy is a professional who is less capable of independent judgment than their credentials suggest. The population-level consequence is more alarming.

We are currently deploying AI tools (at scale, across education, professional practice, and daily information consumption) in an information environment that is also at a historical high point of misinformation production, algorithmic content distortion, and institutional manipulation. The populations most vulnerable to these threats are those least equipped with the critical thinking capacity to evaluate information independently. And the most rapidly growing population of people with reduced critical thinking capacity is, under current AI adoption trajectories, the educated professional class that has adopted AI tools earliest and most comprehensively.

This is not a problem that better AI tools solve. It is a problem that better AI tools accelerate, if the adoption of those tools is not accompanied by deliberate investment in the cognitive capacities they risk atrophying.

The historical parallel is not encouraging. The introduction of calculators into mathematics education was accompanied by legitimate concern that arithmetic fluency would atrophy, and some evidence suggests it did, to a degree, in populations that adopted calculators without maintaining fluency practice. The introduction of GPS navigation was accompanied by concern that spatial navigation capacity would atrophy, and research points the same way: in a small study of 50 drivers, more lifetime GPS use was associated with worse spatial memory when navigating unaided, and in a 13-person follow-up about three years later, more GPS use was associated with steeper decline. In both cases, the atrophied capacity is genuinely less important in a world where the technological substitute is reliably available. In neither case is the technological substitute as reliably available as it appears, and the populations who discover this are the ones who most need the atrophied capacity at precisely the moments when it is not available.

Critical thinking is the navigation capacity for the information environment. And the information environment in which we are adopting AI tools at scale is one in which navigation capacity matters more, not less, than it did before AI tools made navigation apparently unnecessary.

What Deliberate Preservation of Critical Thinking Looks Like

The appropriate response to the risk of AI-assisted critical thinking atrophy is not the rejection of AI tools. It is the deliberate, structured preservation of critical thinking practice alongside AI tool use: in education, in professional development, and in personal information habits.

In education: The most important shift is pedagogical rather than technological. Educators who want to preserve critical thinking development in the AI era need to structure learning around the process of reasoning rather than the product of answers: asking students not just what they conclude but how they evaluated the evidence, what alternative interpretations they considered, what they would need to see to change their minds. This requires assessment approaches that make visible the reasoning process rather than evaluating only the final product, and it requires classroom cultures that treat uncertainty as a legitimate and valuable cognitive state rather than a failure to find the right answer.

It also requires explicit instruction in the recognition of AI limitations: specifically, instruction that helps students identify the difference between AI-generated confidence and well-supported conclusions, between AI-produced summaries and primary source engagement, and between the experience of getting an answer and the experience of understanding the reasoning that supports it.

In professional practice: Organizations that adopt AI tools for knowledge work need to build in deliberate friction, structured requirements for human evaluation, explicit processes for questioning AI outputs, and professional development that maintains the judgment capacities that AI tools risk atrophying. This is not anti-efficiency. It is efficiency calibrated to the actual purpose of professional judgment: not the production of answers, but the production of correct answers in domains where correctness matters.

In personal information habits: The deliberate cultivation of practices that maintain critical engagement with information (reading primary sources rather than AI summaries, seeking out genuine intellectual challenge rather than confirmation of existing views, maintaining the uncomfortable practice of sitting with uncertainty before resolving it, and developing explicit habits of evaluating confidence against evidence quality) is increasingly an act of cognitive self-preservation in the current information environment.

The Deepest Risk

The deepest risk of the current trajectory is not that AI tools will make any specific group of people less intelligent. It is that the widespread atrophy of critical thinking capacity will produce populations less capable of recognizing manipulation, less capable of evaluating institutional claims, less capable of holding powerful actors accountable, and less capable of the collective reasoning that democratic governance requires.

These are not AI problems in the sense that AI caused them. Political polarization, institutional mistrust, the erosion of epistemic commons, and the weaponization of misinformation all predate the current AI moment. But AI tools, adopted at scale without deliberate countermeasures, have the potential to accelerate each of these dynamics: by reducing the individual and collective capacity to evaluate information independently, by making the outsourcing of judgment feel more normal, and by creating populations increasingly dependent on the judgment of systems that cannot be held accountable in the ways that human institutions can.

The skills required to navigate this are the same skills that have always been required to navigate complex information environments in high-stakes contexts: the ability to evaluate evidence critically, to recognize the limits of one's own knowledge and judgment, to hold uncertainty without resolving it prematurely, and to revise one's conclusions when better reasoning warrants it.

These skills are not innate. They are built through practice. And the most consequential thing the current generation of educators, professional trainers, and organizational leaders can do is insist on maintaining the conditions for that practice, even as, especially as, the tools that make practice feel unnecessary become more powerful, more polished, and more widely adopted.

Where This Fits

Critical thinking, trust literacy, and AI literacy are the same discipline applied to different material: read the conditions before you accept the answer. See Trust Literacy and the companion article on AI literacy.

Schedule a conversation with Links Global: www.linksglobal.org

Sources

- Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies, 15*(1), 6. doi.org/10.3390/soc15010006
- Dahmani, L., & Bohbot, V. D. (2020). Habitual use of GPS negatively impacts spatial memory during self-guided navigation. *Scientific Reports, 10*, 6310. doi.org/10.1038/s41598-020-62877-0
- Bjork, R. A., & Bjork, E. L. (2020). Desirable difficulties in theory and practice. *Journal of Applied Research in Memory and Cognition, 9*(4), 475–479.

Ownership and use. © 2026 Links Global. All rights reserved. The Links Global Trust Audit™, Narrative Stewardship Infrastructure™, NSI™, Epistemic Re-Regulation Architecture™, and ERA™ are trademarks of Links Global. You are welcome to read this article, link to it, share it, and quote brief passages with attribution to Links Global and a link to this page. Reproducing it in whole or in substantial part, adapting or translating it, or using the analysis and the preservation practices described here in training, consulting, publications, or derivative works requires prior written permission from Links Global. Please cite as: Sanchez, M. (2026). *Critical Thinking in the Age of AI: The Skill We Are Automating Away at Exactly the Moment We Need It Most*. Links Global. www.linksglobal.org.