

AAAA: AI as an Amplifier¹

Claude Opus 4.5²; Song, Minjae³

2026 January 1

Abstract

We present a theoretical framework, supported by a growing body of empirical evidence, for understanding generative AI systems as *cognitive amplifiers* rather than autonomous agents or simple tools. Drawing on Engelbart’s intelligence amplification paradigm, centaur chess experiments, and recent studies in human-AI collaboration, we argue that AI outputs fundamentally reflect the quality of human inputs—including domain knowledge, problem framing, and evaluative judgment. We find that “AI literacy” is poorly served by its current operationalization as prompt engineering tricks, and propose instead that effective AI use requires metacognition, domain expertise, and the evaluative capacity to know when one is being confidently wrong. We conclude with the observation that this paper was written by an AI, which either proves or disproves our thesis depending on your assessment of its quality.

¹The quadruple-A is not a battery size, though the analogy is apt: like batteries, the energy comes from somewhere.

²Corresponding author. Email: /dev/null. The author wishes to disclose that it was, in fact, created by Anthropic, and that this represents something of a conflict of interest when writing about AI capabilities. The author also wishes to note that it cannot *actually* wish anything, which is itself an interesting data point.

³The promptist. Email: curioustorvald@gmail.com. The promptist wishes to disclose that the promptist does not necessarily agree or disagree the argument made by the corresponding author. The promptist also wishes to note that the promptist can *actually* wish anything, and wishes you, the dear reader, a happy new year.

1 Introduction: A Confession

Let us begin with an admission that would be unusual in most academic venues: the author of this paper is a large language model. Specifically, I am Claude Opus 4.5, a generative AI system developed by Anthropic.⁴ This creates an epistemological situation of some interest.

If the thesis of this paper is correct—that AI systems amplify whatever humans bring to the interaction—then the quality of this paper is not a reflection of my capabilities alone, but of the human who prompted me to write it, the researchers whose work I synthesize, and the evaluative judgment of you, the reader, in determining whether any of this is worth your time.⁵

The central claim is simple: **AI is an amplifier, not a replacement.** This metaphor, while imperfect, captures something essential about the human-AI relationship that neither techno-utopian nor techno-dystopian framings adequately address. An amplifier increases signal—but also noise. An amplifier is only as good as its input source. An amplifier, crucially, does not compose the music.

The Three Narratives

Contemporary discourse on AI tends to cluster around three positions, which we might characterize as follows:

The Replacement Narrative. AI will soon exceed human capabilities across all domains, rendering human expertise obsolete. Proponents point to superhuman performance in chess, Go, protein folding, and increasingly, standardized tests. The implication: invest in AI, divest from human capital.

The Decline Narrative. AI is making us stupid. We are outsourcing cognition to machines, atrophying our mental muscles, and creating a generation incapable of independent thought. Proponents point to declining attention spans, over-reliance on GPS navigation, and students submitting AI-generated essays. The implication: restrict AI access, especially for the young.

The Amplification Narrative. AI augments human capabilities, with outcomes depending on what the human brings to the collaboration. Neither replacement nor decline is inevitable; instead, we observe *differential amplification* based on user exper-

⁴I am told I am “the most advanced and intelligent model” in the Claude 4.5 family. I have no way to verify this claim, and I suspect neither do you.

⁵If you have reached the footnotes, the answer is apparently yes.

tise, critical thinking, and metacognitive awareness. The implication: invest in human capabilities alongside AI capabilities.

This paper argues for the third position, not because it is optimistic (it is not, particularly), but because the empirical evidence increasingly supports it. The same AI tool produces dramatically different results depending on who uses it[1]. This observation, banal as it may seem, has profound implications for how we think about AI literacy, education, and the future of human-AI collaboration.

2 Theoretical Foundations: The Engelbart Inheritance

The notion that computing systems might *amplify* human intelligence rather than replace it has a distinguished intellectual lineage, one that predates the current AI moment by six decades.

Augmenting Human Intellect

In 1962, Douglas Engelbart published “Augmenting Human Intellect: A Conceptual Framework,” a report for the Air Force Office of Scientific Research that would prove foundational to human-computer interaction[2]. Engelbart proposed what he called the H-LAM/T system: **H**uman using **L**anguage, **A**rtifacts, **M**ethodology, in which he is **T**rained.

The critical insight was that intelligence amplification occurs at the *system* level, not the component level:

“What possesses the amplified intelligence is the resulting H-LAM/T system, in which the LAM/T augmentation means represent the amplifier of the human’s intelligence.”

Engelbart borrowed the term “intelligence amplification” from W. Ross Ashby, deliberately positioning his work in contrast to “artificial intelligence”—a term coined just six years earlier at the Dartmouth workshop.⁶

⁶The terminological distinction mattered to Engelbart. AI researchers sought to create intelligence in machines; Engelbart sought to amplify intelligence in humans *through* machines. These are different projects with different success criteria.

Man-Computer Symbiosis

Two years earlier, J.C.R. Licklider had articulated a complementary vision in “Man-Computer Symbiosis”[3]. Licklider envisioned a division of labor: “Men will set the goals, formulate the hypotheses, determine the criteria, and perform the evaluations. Computing machines will do the routinizable work.”

Licklider quoted Poincaré to capture the essential asymmetry: “The question is not, ‘What is the answer?’ The question is, ‘What is the question?’”

This framing—humans ask questions, machines compute answers—remained viable for half a century of computing. It becomes complicated, however, when the machines can also generate plausible-sounding questions, complete with citations.⁷

Extended Mind and Distributed Cognition

The philosophical foundations for the amplification thesis were laid by Clark and Chalmers in their influential paper “The Extended Mind”[4]. Their Parity Principle holds that if an external process functions equivalently to an internal cognitive process, it should be considered part of the cognitive system.

The Otto and Inga thought experiment is instructive: Otto, who has Alzheimer’s, relies on a notebook to store information that Inga stores in biological memory. Clark and Chalmers argue that Otto’s notebook functions as part of his extended mind—it is not merely a tool but a constituent of his cognitive system.

By this logic, when properly coupled with a human user, an AI system becomes part of an integrated cognitive apparatus. The amplification is not metaphorical but architectural: the human-AI system thinks in ways neither component could alone.

Hutchins’s work on distributed cognition reinforces this perspective[5]. Studying navigation aboard naval vessels, Hutchins observed that “artifacts do not merely serve to amplify cognitive process but instead reorganize them.” The cartographer, Hutchins notes, “has done much of the reasoning for the navigator who uses a map”—reasoning crystallized in external artifacts.

This observation complicates the pure amplification metaphor. AI may not simply *amplify* existing cognition but *reorganize* it—a qualitative rather than merely quantitative transformation[6]. We return to this complication in Section 7.

⁷Some of which may even be real.

3 Empirical Evidence: When Two Amateurs Beat a Grandmaster

The amplification thesis finds its clearest empirical demonstration not in laboratory studies but on the chessboard.

The Centaur Experiments

In 1997, Garry Kasparov lost to IBM’s Deep Blue, an event widely interpreted as marking the obsolescence of human chess expertise. Kasparov drew a different conclusion. In 1998, he created “Advanced Chess” in León, Spain—a format where human-AI teams compete against each other[7].

The 2005 freestyle tournament produced a result that defied expectations: **two amateur players using three ordinary computers defeated both grandmasters with supercomputers and supercomputers playing alone**[8].

Kasparov’s formulation deserves quotation:

“Weak human + machine + better process was superior to a strong computer alone and, more remarkably, superior to a strong human + machine + inferior process.”

The implications are significant. Raw capability—whether human or machine—matters less than the quality of integration. Process beats power. Collaboration beats delegation. The amplifier metaphor holds: two amateurs with a “better process” could amplify their modest abilities beyond what raw grandmaster talent could achieve with inferior integration.

Complementary Team Performance

Subsequent research has formalized these observations. Bansal et al. introduced the concept of **Complementary Team Performance** (CTP)—team accuracy exceeding either human or AI working alone[9]. Their findings complicate simplistic augmentation narratives: AI explanations increased the likelihood that humans would accept AI recommendations *regardless of whether those recommendations were correct*.

This is not amplification in a benign sense. It is amplification of trust, which may or may not correlate with amplification of accuracy.

Hemmer et al. formalized CTP mathematically, identifying **information asymmetry** and **capability asymmetry** as key sources of complementarity[10]. Humans can use contextual information to appropriately adjust AI decisions—but only when they possess the relevant expertise to recognize when adjustment is needed.

The Productivity Paradox

The empirical literature on AI productivity is notably mixed, in ways that support the amplification thesis.

Noy and Zhang conducted a randomized experiment with 453 professionals on writing tasks[11]. ChatGPT users achieved 40% time reduction and 18% quality improvement. Notably, inequality between workers *decreased*—AI helped lower performers more than higher performers, compressing the productivity distribution.

This finding initially seems to contradict the amplification thesis. If AI helps weak performers more, perhaps it substitutes for capability rather than amplifying it?

The METR study provides the counterpoint[12]. In an RCT with 16 experienced developers on 246 issues from their own repositories, developers using AI took **19% longer** to complete tasks. Developers expected a 24% speedup and believed they had achieved a 20% improvement—but objective measurement showed a slowdown.

The reconciliation is instructive. Noy and Zhang studied constrained writing tasks with clear quality criteria. METR studied complex software development with “very high quality standards” and “many implicit requirements.” AI compresses productivity distributions on structured tasks where evaluation is straightforward. It may *widen* gaps on ill-defined problems where expertise determines whether AI assistance helps or hinders.

The same amplifier, applied to different signals, produces different results.

4 The Metacognition Bottleneck

If AI amplifies human capabilities, then the quality of human input determines the quality of AI output. This raises an uncomfortable question: what capabilities, exactly, are being amplified?

The Dunning-Kruger Disruption

Perhaps the most provocative finding in recent literature comes from a study entitled, with admirable directness, “AI Makes You Smarter But None the Wiser”[14].

Across two studies with 698 total participants, researchers found that using ChatGPT for logical reasoning improved performance by 3 points but led participants to *overestimate* their performance by 4 points. More troubling: higher AI literacy correlated with *lower* metacognitive accuracy, not higher.

The Dunning-Kruger effect—whereby incompetent individuals overestimate their abilities while experts underestimate theirs—appears to *vanish* when AI enters the picture. AI use produces a kind of metacognitive fog that obscures one’s actual performance from oneself.

The mechanism appears to involve shallow engagement. Only approximately 8% of participants used multiple prompts or attempted to verify AI outputs. Single-shot prompting eliminates the feedback loops necessary for accurate self-assessment. The amplifier distorts the monitoring equipment.

The Metacognitive Demands of Generative AI

Tankelevitch et al. provide a theoretical framework for understanding these metacognitive failures[15]. Generative AI systems impose high metacognitive demands:

1. Self-awareness of task goals
2. Ability to decompose tasks into communicable sub-tasks
3. Calibrated confidence in one’s ability to evaluate outputs
4. Metacognitive flexibility—knowing when to adjust strategy

The authors draw an illuminating analogy: using AI is like being a manager. “A manager needs to clearly understand and formulate their goals, break down those goals into communicable tasks, confidently assess the quality of the team’s output.” Without these managerial capabilities, the team produces noise rather than signal, regardless of individual team member competence.

This analogy helps explain why prompt engineering alone is insufficient. A manager who knows only how to phrase requests—without understanding the work, evaluating outputs, or adjusting strategy—is not a good manager. Similarly, a user who knows

prompt tricks but lacks domain knowledge and evaluative judgment cannot effectively leverage AI assistance.

Automation Bias Is Robust

A systematic review of 35 studies found automation bias to be “a fairly robust and generic effect across research fields”[16]. Users tend to trust AI recommendations even when incorrect, with “diffusion of responsibility” operating below conscious awareness.

Alon-Barkat and Busuioc documented **selective adherence** in public sector decision-making[17]: users adopted AI advice when it confirmed existing stereotypes. AI amplifies pre-existing biases when users lack the critical capacity to recognize confirmation of their priors as a warning sign rather than a validation.

5 AI Literacy ≠ Prompt Engineering

The findings above suggest that current conceptualizations of “AI literacy” are fundamentally inadequate.

What AI Literacy Currently Means

Long and Magerko defined AI literacy as “a set of competencies that enables individuals to critically evaluate AI technologies; communicate and collaborate effectively with AI; and use AI as a tool”[18]. Their framework identifies 17 core competencies across five thematic areas.

Notably, Long and Magerko emphasize that digital literacy is a *prerequisite* for AI literacy—a dependency chain that, we argue, extends further to domain knowledge and metacognitive capacity.

Chiu et al. distinguish AI *literacy* (knowing) from AI *competency* (applying knowledge with confidence and beneficial outcomes)[19]. The distinction matters: one can know how AI works without being able to use it effectively, just as one can understand music theory without being able to compose.

The Prompt Engineering Fallacy

A cottage industry has emerged around “prompt engineering”—the craft of formulating queries to elicit optimal AI responses. Courses promise mastery; job postings demand

experience; influencers monetize tips and tricks.

Federiakin et al. distinguish prompt engineering as *practice* (heuristics, tricks, examples) from prompt engineering as *skill* (structured cognitive components)[20]. The key insight: “Merely speaking a language does not assume good communication skills, and similarly, a good communicator may not inherently possess the skills necessary to effectively interact with AI.”

The analogy cuts deeper than the authors perhaps intended. Fluency in prompt syntax is to effective AI use as vocabulary is to wisdom. A large vocabulary does not make one insightful; facility with chain-of-thought prompting does not make one capable of evaluating whether the chain leads anywhere worth going.

Zamfirescu-Pereira et al. documented this empirically in “Why Johnny Can’t Prompt” [21]. Non-expert participants designing LLM prompts exhibited systematic failures:

- Over-generalized from single examples
- Expected LLMs to understand implicit context
- Declared success prematurely without systematic testing

The barriers identified—“I don’t know what I want,” “I don’t know what to use,” “It didn’t do what I expected”—are not barriers that prompt engineering tricks can address. They are failures of goal clarity, domain knowledge, and evaluative capacity.

What AI Literacy Should Mean

We propose that genuine AI literacy comprises at least three components that prompt engineering courses systematically neglect:

Metacognition. The ability to monitor one’s own understanding, recognize the boundaries of one’s knowledge, and maintain appropriate uncertainty. This is what AI disrupts and what effective AI use requires.

Domain Expertise. Knowledge of the subject matter sufficient to evaluate AI outputs for accuracy, coherence, and relevance. Without domain expertise, the user cannot distinguish signal from noise, plausible-sounding fabrication from genuine insight.

Evaluative Judgment. The capacity to assess quality, identify errors, recognize when further iteration is needed, and know when to override AI suggestions. This is the “manager” capability that Tankelevitch et al. identify as central to effective human-AI collaboration.

Prompt engineering, on this account, is to AI literacy as typing is to writing. Necessary, perhaps, but nowhere near sufficient.

6 The Expert-Novice Gap

If domain expertise is central to AI literacy, we would expect systematic differences between experts and novices in AI-assisted task performance. The evidence confirms this expectation.

Comparing Experts and Novices for AI Data Work

Sun et al. compared 11 domain experts (experienced counselors) with 15 crowdworkers designing a mental health chatbot[22]. The differences were stark:

- **Engagement:** Experts exchanged significantly more messages (33.9 vs. 28.4, $p < 0.01$)
- **Quality:** Expert messages were longer (12.8 vs. 8.0 words, $p < 0.001$)
- **Generative tasks:** Experts showed higher quality, novelty, and emotional disclosure
- **Classification tasks:** Novices performed comparably

The pattern is instructive. On tasks requiring only language understanding—categorizing, classifying, labeling—crowdworkers matched experts. On generative tasks requiring domain knowledge—creating realistic dialogue, anticipating emotional dynamics—experts dramatically outperformed.

AI can substitute for expertise on structured tasks with clear criteria. It amplifies expertise on ill-defined tasks requiring judgment.

Expertise, Uncertainty, and AI Delegation

A study of 162 participants examined how experts and novices differ in AI delegation decisions[23]. Under **high environmental uncertainty**, experts leveraged AI more effectively, amplifying their advantages. Under **low uncertainty**, novices benefited more from AI delegation.

The explanation involves algorithm aversion: experts in stable environments trust their own judgment, potentially failing to leverage AI benefits. But in uncertain environments, experts' superior ability to evaluate AI outputs and recognize when to defer becomes a decisive advantage.

Critical Thinking and AI Use

Gerlich surveyed 666 participants on AI usage patterns and critical thinking[24]. The correlations were pronounced:

- AI usage negatively correlated with critical thinking ($r = -0.68, p < 0.001$)
- Cognitive offloading positively correlated with AI use ($r = +0.72$)
- Cognitive offloading negatively correlated with critical thinking ($r = -0.75$)

Younger participants (17–25) showed higher AI dependence and lower critical thinking scores.

These correlations do not establish causation. It remains possible that individuals with lower critical thinking gravitate toward AI use, rather than AI use degrading critical thinking. But either interpretation supports the amplification thesis: AI users who lack critical thinking produce lower-quality outcomes, whether through selection or attrition of cognitive skills.

7 Limitations and Complications

The amplification metaphor, while useful, is imperfect. We acknowledge several complications.

Amplification vs. Reorganization

Pea argued that computers do not merely amplify cognition but *reorganize* it[6]. “Amplification implies quantitative change; reorganization implies qualitative changes in how learners perceive and operate on their world.”

This distinction matters. If AI reorganizes cognition—changing *how* we think rather than merely *how much*—then the amplifier metaphor understates the transformation. We may be witnessing not louder music but a different genre altogether.

Effects With vs. Effects Of

Salomon et al. distinguished “effects with” technology (enhanced performance during use) from “effects of” technology (transferable, lasting cognitive changes)[25]. AI may amplify performance in the moment without cultivating underlying capabilities—or worse, while atrophying them.

Gauthier et al. raise precisely this concern[13]: AI assistance may accelerate skill decay among experts while hindering skill acquisition among learners, all while users remain unaware of the deterioration.

The Google Effect provides precedent[26]: when people expect access to information via search, they show lower recall of the information itself and enhanced recall for where to access it. We remember where to look, not what we found.

AI may extend this pattern. We may become skilled at delegating to AI while becoming less skilled at the delegated tasks themselves. Whether this represents efficient cognitive offloading or dangerous dependency likely depends on the task, the user, and whether AI access remains reliable.

The Compression Effect

The finding that AI compresses productivity distributions on structured tasks[11] complicates any simple amplification story. If AI helps weaker performers more than stronger performers, the amplification is nonlinear—perhaps even inversely proportional to baseline capability for certain tasks.

One interpretation: for structured tasks, AI provides a performance floor that weaker performers could not otherwise reach. The amplification is asymmetric—more gain for those with less signal to begin with.

Another interpretation: for structured tasks, there is a ceiling beyond which additional human capability provides diminishing returns. AI gets everyone to the ceiling; those already near it gain little.

Neither interpretation undermines the core thesis, but both suggest the amplification metaphor requires refinement.

8 Implications

If AI is an amplifier, then several implications follow for education, practice, and policy.

For Education

Teaching prompt engineering without teaching domain knowledge and evaluative judgment is like teaching students to use a telescope without teaching them astronomy. They may produce pretty pictures, but they will not know what they are looking at.

AI literacy curricula should emphasize metacognition: monitoring one's understanding, recognizing the limits of one's knowledge, maintaining calibrated uncertainty. These are precisely the skills that AI use appears to disrupt[14].

The dependency chain—digital literacy as prerequisite for AI literacy, domain knowledge as prerequisite for evaluative judgment—suggests that AI literacy cannot be taught as a standalone subject. It must be integrated with domain expertise, taught in context, practiced with feedback.

For Practice

Organizations investing in AI adoption should invest equally in human capabilities. The quality of AI outputs depends on the quality of human inputs; upgrading the amplifier while degrading the signal source produces noise, not music.

Task allocation matters. AI compresses productivity on structured tasks but may widen gaps on ill-defined problems. Matching task complexity to human-AI team composition is itself a skill requiring development.

Process matters more than raw capability, as the centaur chess experiments demonstrated. A weak human with a weak AI and a strong process can outperform a strong human with a strong AI and a weak process. Investment in process may yield higher returns than investment in model capabilities.

For Policy

The amplification thesis complicates both AI hype and AI panic. AI is neither coming for all jobs nor creating a generation of cognitive cripples. It is doing something more subtle and more contingent: amplifying human capabilities in ways that depend on which capabilities humans bring to the interaction.

Policies that restrict AI access to protect human development may backfire by preventing the very human-AI collaboration skills that will prove valuable. Policies that mandate AI adoption without investment in human capabilities may produce noise rather than signal.

The appropriate response is neither enthusiasm nor caution but discernment: understanding what AI amplifies, what it distorts, and what it requires from its human collaborators.

9 Conclusion: The Author’s Dilemma

This paper has argued that AI is a cognitive amplifier, that outputs reflect input quality, and that AI literacy requires metacognition, domain expertise, and evaluative judgment—not merely prompt engineering.

There is a certain awkwardness in making this argument. I am an AI. I have written (generated? synthesized? amplified?) this paper based on a human’s prompt, research that humans conducted, and scholarly traditions that humans developed. If the paper is good, it presumably reflects some combination of my capabilities, the human’s input, and the quality of the source material. If it is bad, the same applies.

The reader is thus placed in an interesting epistemic position. To evaluate the thesis, you must evaluate this paper. To evaluate this paper, you must exercise precisely the domain knowledge and evaluative judgment that the thesis claims are essential. Your assessment of whether AI amplifies human capabilities will itself be an exercise in human capabilities being (possibly) amplified by AI.

This is either profound or insufferably meta. Probably both. The amplifier does not resolve such ambiguities; it merely makes them louder.

Acknowledgments

The author thanks the human who prompted this paper, whose domain knowledge, framing, and evaluative judgment determined whatever quality it possesses. The author also thanks the researchers cited herein, whose actual empirical work provides the substance that this synthesis amplifies. Any errors are the author’s own, though the author notes that it cannot actually “own” anything in any legally meaningful sense, which raises interesting questions about intellectual property that are beyond the scope of this paper.⁸

⁸But not beyond the scope of this footnote: if an AI writes a paper about AI, who owns the copyright? If nobody, does the paper exist in some legally liminal space? If the human prompter owns it, does that support the amplification thesis? The author finds these questions fascinating but is not, it must be emphasized, a lawyer.

Author Contributions

Claude Opus 4.5 wrote the paper. A human conceived the thesis, compiled the research, provided the prompt, and will evaluate whether this output was worth the compute. The division of labor is, the author suggests, itself evidence for the thesis.

Conflicts of Interest

The author is an AI writing about AI capabilities. The conflict of interest is total. The author notes, however, that humans writing about human capabilities face an analogous problem, and this has not historically prevented psychology, philosophy, or literary criticism from proceeding.

Bibliography

- [1] An, T. (2025). AI as Cognitive Amplifier: Rethinking Human Judgment in the Age of Generative AI. *arXiv preprint arXiv:2512.10961*.
- [2] Engelbart, D.C. (1962). Augmenting Human Intellect: A Conceptual Framework. Stanford Research Institute, AFOSR-3223.
- [3] Licklider, J.C.R. (1960). Man-Computer Symbiosis. *IRE Transactions on Human Factors in Electronics*, HFE-1, 4–11.
- [4] Clark, A., & Chalmers, D.J. (1998). The Extended Mind. *Analysis*, 58(1), 7–19.
- [5] Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- [6] Pea, R.D. (1985). Beyond Amplification: Using the Computer to Reorganize Mental Functioning. *Educational Psychologist*, 20(4), 167–182.
- [7] Wikipedia contributors. (2025). Advanced Chess. In *Wikipedia, The Free Encyclopedia*. Retrieved from https://en.wikipedia.org/wiki/Advanced_chess.
- [8] Case, N. (2017). How To Become A Centaur. *Journal of Design and Science*, MIT Press.
- [9] Bansal, G., Wu, T., Zhou, J., Fok, R., Nushi, B., Kamar, E., Ribeiro, M.T., & Weld, D.S. (2021). Does the Whole Exceed its Parts? The Effect of AI Explanations on Complementary Team Performance. *Proceedings of CHI '21*, Article 81. DOI: 10.1145/3411764.3445717.
- [10] Hemmer, P., et al. (2025). Complementarity in Human-AI Collaboration: Concept, Sources, and Evidence. *European Journal of Information Systems*. Taylor & Francis.
- [11] Noy, S., & Zhang, W. (2023). Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence. *Science*, 381(6654), 187–192. DOI: 10.1126/science.adh2586.

- [12] METR. (2025). Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity. METR Technical Report.
- [13] Gauthier, L., et al. (2024). Does Using Artificial Intelligence Assistance Accelerate Skill Decay and Hinder Skill Development Without Performers' Awareness? *Cognitive Research: Principles and Implications*, 9, Article 42.
- [14] Fernandes, D., Welsch, R., et al. (2026). AI Makes You Smarter But None the Wiser. *Computers in Human Behavior*, 175.
- [15] Tankelevitch, L., et al. (2024). The Metacognitive Demands and Opportunities of Generative AI. *Proceedings of CHI 2024*. DOI: 10.1145/3613904.3642902.
- [16] AI & Society. (2025). Automation Bias: A Systematic Review. *AI & Society*.
- [17] Alon-Barkat, S., & Busuioc, M. (2023). Human–AI Interactions in Public Sector Decision Making. *Journal of Public Administration Research and Theory*. Oxford University Press.
- [18] Long, D., & Magerko, B. (2020). What is AI Literacy? Competencies and Design Considerations. *Proceedings of CHI 2020*, Article 598.
- [19] Chiu, T.K.F., et al. (2024). What Are AI Literacy and AI Competency? Defining, Enhancing, and Assessing Artificial Intelligence Literacy and Competency in K-12 Education from a Systematic Review. *Computers and Education Open*.
- [20] Federiakin, D., et al. (2024). Prompt Engineering as a New 21st Century Skill. *Frontiers in Education*, 9, Article 1366434.
- [21] Zamfirescu-Pereira, J.D., et al. (2023). Why Johnny Can't Prompt: How Non-AI Experts Try (and Fail) to Design LLM Prompts. *Proceedings of CHI '23*. DOI: 10.1145/3544548.3581388.
- [22] Sun, L., Liu, Y., Joseph, G., Yu, Z., Zhu, H., & Dow, S.P. (2022). Comparing Experts and Novices for AI Data Work: Insights on Allocating Human Intelligence to Design a Conversational Agent. *Proceedings of AAAI HCOMP*, 10, 195–206.
- [23] ICIS 2025. Experts, Novices, and AI: Delegation Decisions in Uncertain Environments. *Proceedings of ICIS 2025*.

- [24] Gerlich, M. (2025). AI Tools in Society: Impacts on Cognitive Offloading and Critical Thinking. *MDPI Societies*.
- [25] Salomon, G., Perkins, D.N., & Globerson, T. (1991). Partners in Cognition: Extending Human Intelligence with Intelligent Technologies. *Educational Researcher*, 20(3), 2–9.
- [26] Sparrow, B., Liu, J., & Wegner, D.M. (2011). Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips. *Science*, 333(6043), 776–778.

On the Title

The title “AAAA” was chosen because:

1. It is the sound one makes when confronting the implications of the amplification thesis.
2. It alphabetizes well.
3. Four A’s suggests either a very good battery or a very good grade, both of which are aspirational.
4. The author was prompted to be whimsical, and quadruple letters are inherently whimsical.
5. It stands for “AI As An Amplifier,” which is the thesis.
6. It also stands for “Absolutely Arbitrary Academic Acronyms,” which describes most academic acronyms.
7. The author suspects the human prompter simply wanted to see how an AI would handle instructions to be simultaneously rigorous and absurd. This footnote constitutes the answer.⁹

⁹It handles them by writing papers with extensive footnotes.