research note

The Quiet Cognitive Coup of Generative AI
Rewriting the Rules

Debora Pfaff, PhD

The author is an associate professor of research at the Ann Caracristi Institute, National Intelligence University.

Generative AI didn't ask permission. Instead, it began quietly rewriting the rules of how intelligence analysts think, decide, and make sense of the world.

When we think of analysis, we think of cognition—brain processes that humans use to interpret their surroundings, which include perception, learning, memory and decision making.[1] Dig a little deeper, and it's clear that cognition isn't just about an individual's thoughts and decisions; it's about a cognitive ecosystem—the interactive relationship between an organism and its environment and how that interaction impacts its thinking.[2] Put another way, humans don't think in a vacuum; we're heavily influenced by learned, symbolic, and social processes. Cognition evolves through constant adaptation, and it's irreducibly cultural.[3] As applied to the IC, one NIU study in 2024 found that intelligence work is a social process: "While intelligence professionals draw on their individual cognitive abilities to interpret information and make assessments, effective intelligence work also requires asking questions, sharing information, considering alternative points of view, coordinating assessments, and communicating those assessments to inform action."[4]

So what happens when a non-human mind is introduced to analytic workflows? Historically, technological advances—new sensors, more data, better computing—largely changed the inputs: bigger data streams, more efficient collection, faster processing. But the outputs largely stayed the same; humans deliberate over evidence, reason, debate interpretations, and apply tradecraft to produce a product that stems from their thinking. Introducing AI introduces a new synthetic cognitive agent; one capable of hypothesis generation, abstraction, summarization, pattern synthesis, and even narrative construction. But it cannot—critically—model intuitive, affective patterns of reasoning in the same manner that humans experience them. AI cannot think.[5] In fact, if you ask a large language model (LLM) whether it can reason, it will explain that it can generate patterns, but it cannot generate new knowledge.

Analysts draw on raw data collected from multiple intelligence sources, but the heart of their work lies in sensemaking: detecting patterns, forming hypotheses, adjudicating competing interpretations, assessing confidence, surfacing dissent, and building consensus. Sensemaking addresses the basic analytic questions—who, what, when, and where—as well as the more advanced why, how, and what next.[6] It ultimately leads to a "foraging for answers," and analysts must constantly

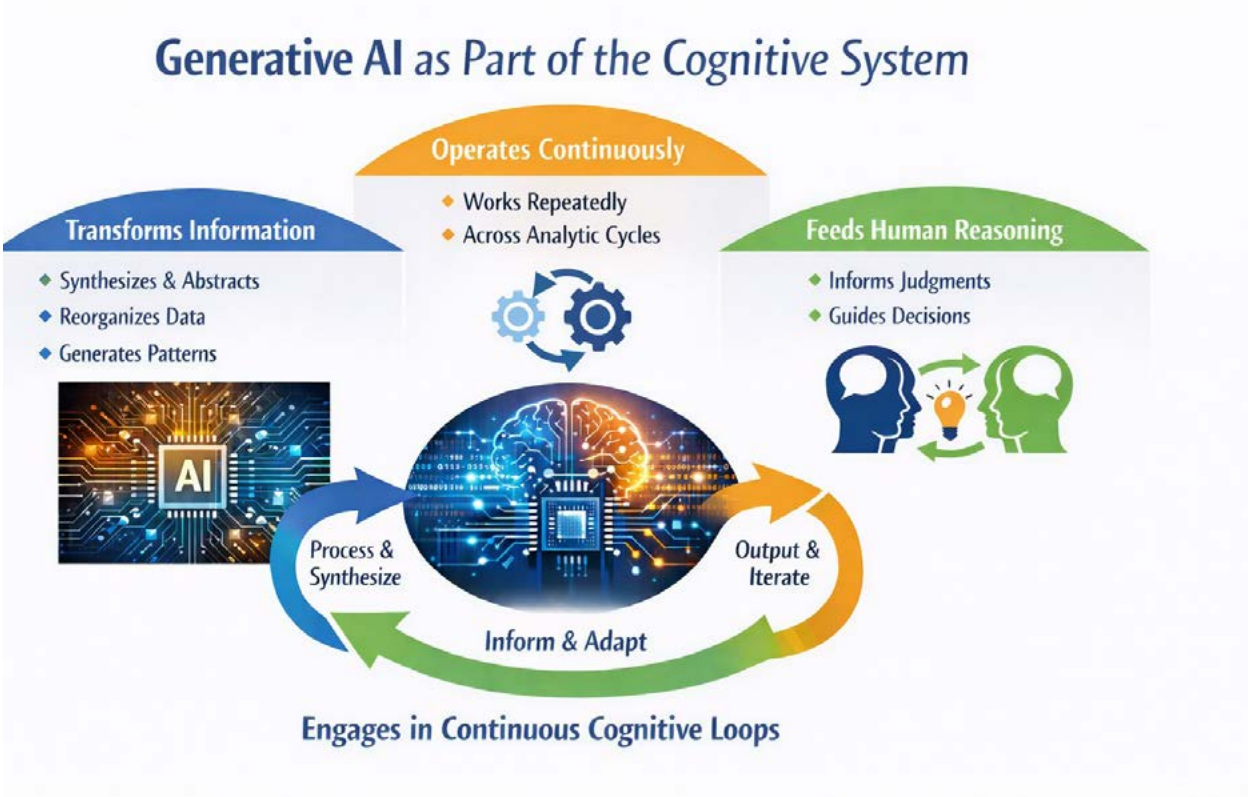

This graphic shows generative artificial intelligence as part of a cognitive system that transforms information, operates continuously, and feeds human reasoning.

evaluate both how the intelligence was produced and the accuracy and reliability of the information itself.[7]

AI as a Cognitive Technology

The IC has historically integrated new tools, but generative AI is the first cognitive technology. Generative models exhibit a form of synthetic cognition: pattern synthesis, abstraction, hypothesis generation, summarization, and translation. Automation and AI systems like LLMs are increasingly performing tasks that used to be conducted exclusively by humans. Some tasks are suited to AI, and make an analyst's job more efficient, like rapidly ingesting and filtering large datasets to surface relevant information, or automating translation and cross-lingual analysis to accelerate exploitation of foreign-language reporting. But as a cognitive technology, GenAI becomes a co-author of analytic reasoning; it participates in the formation, shaping, and constraint of thought rather than simply executing discrete tasks. It has shared cognitive agency in the analytic process.

AI has found its way upstream from the analytic process by affecting how a problem is framed and how options to approach that problem are constructed. This mirrors how human collaborators influence one another's reasoning—not by immediately coming up with conclusions, but by shaping the conceptual space. And unlike human collaborators, generative AI can repeatedly re-enter the reasoning loop with consistency, memory, and scale. The system ends up shaping not just single judgments, but the evolution of reasoning time and time again. This is a core premise of Hutchins' distributed cognition theory, which holds that cognition is not confined to individual minds, but is distributed across people, artifacts, representations, environments, and time.[8] This complexity affects analysts' cognitive ecosystems in a number of ways.

Epistemically. AI-generated assessments sound authoritative—analysts may overweight them. And, generative models create many convincing-but-wrong hypotheses. As analysts share their ideas and products, they may perpetuate the illusion of understanding. Their analysis may contain coherent explanations,

confident language, and structured arguments that mimic human reasoning. As one recent study argues in a broader context, generative AI can foster epistemic illusions that make much more difficult assessing whether the "communicator" truly understands the domain.[9]

Institutionally and culturally. AI co-authorship changes the norms of collaboration and can affect mentorship and junior analyst development by reducing the back-and-forth with mentors that helps them understand problems, learn how to think analytically, and develop sound judgment. It may also limit diversity of thought by driving analysts toward the most common or mainstream interpretations, running counter to analytic standards valuing dissenting views, red-teaming, devil's advocacy, and structured analytic techniques. It affects the relationship between an analyst and their leadership, too. Reviewers of analytic products must evaluate the credibility of output that includes synthetic cognition.

Operationally. Analytic products become a synthesis of what AI considers the most crucial arguments. Under time pressures or a heavy workload, an analyst may accept AI-generated content uncritically. In this way, the role of a human analyst shifts from being the sole sensemaker to an arbiter of synthetic cognition.

Policy and Tradecraft Recommendations

The integration of generative AI demands new tradecraft. Current policy and tradecraft frameworks are not equipped to answer questions like: How do we validate machine-sourced judgments? How do we audit AI's influence? How do we prevent skill deterioration and an over-reliance on AI?[10]

This is not a subtle shift. And given that the IC is generally slow to change, intelligence leaders should quickly begin to consider the following suggestions for adapting its operations to a technology that evolves faster each day.

1. Adopt "epistemic transparency" protocols: Require delineating AI contributions from human input, and describe the prompts used to generate the output.

2. Reinforce human-centered tradecraft: Emphasize or enforce mandates requiring mechanisms for dissent, devil's advocacy, alternative analysis, and structured analytic techniques that keep human thinking at the center of analysis.

3. Support multiple ways of thinking: AI tools shouldn't all push analysts toward the same answers. Use multiple AI models and human judgment together to avoid converging on the same conclusions.

4. Invest in AI-literacy for analysts: Offer training to understand AI limitations, like hallucinations, bias, and overconfidence, and teach analysts to critically evaluate AI outputs rather than treat them as authoritative.

5. Introduce audit and accountability frameworks: Create institutional oversight (audit logs, validation steps, after-action reviews) to examine when AI-derived judgments led to mistakes or inaccurate threat assessments.

Conclusion

Adoption of generative AI in the Intelligence Community isn't just a matter of automation or efficiency. It's a profound shift, one that places non-human cognitive actors alongside human analysts to make sense of the uncertain. If we treat AI as just another tool, it will certainly offer productivity enhancements. But that will patently ignore the infiltration into the analytic cognitive ecosystem which buttresses the Intelligence Community's objectivity and rigor.

The IC has focused on how to adopt and incorporate AI, but that's not the real challenge. The real challenge is learning how to govern an analytic environment where AI has changed how analysts think, reason, and exercise judgment. ■

Endnotes

1. Claudia Mettke-Hofmann, "Cognitive Ecology: Ecological Factors, Life-Styles, and Cognition," *WIREs Cognitive Science* 5, no. 3 (2014): 345–60.
2. Edwin Hutchins, "Cognitive Ecology," *Topics in Cognitive Science* 2, no. 4 (2010): 705–15, https://doi.org/10.1111/j.1756-8765.2010.01089.x.
3. Harry Heft, "An Ecological Approach to Psychology," *Review of General Psychology* 17, no. 2 (June 2013): 162–67, https://doi.org/10.1037/a0032928.
4. Julie Ferringer, "The Science of Teamwork in the Intelligence Community," NI Press Research Short, National Intelligence University, August 22, 2024, https://www.ni-u.edu/wpcontent/uploads/2024/08/NIUShort_20240822_DNI_2024_02540.pdf.
5. Teppo Felin and Matthias Holweg, "Theory Is All You Need: AI, Human Cognition, and Causal Reasoning," *Strategy Science* 9, no. 4 (2024): 346–71, https://doi.org/10.1287/stsc.2024.0189.
6. Kristin Warren et al., *Improving Sense-Making with Artificial Intelligence* (RAND Corporation, 2025), https://www.rand.org/pubs/research_reports/RRA3152-1.html.
7. David T. Moore, *Sensemaking: A Structure for an Intelligence Revolution*, Clift Series on the Intelligence Profession (Center for Strategic Intelligence Research, National Defense Intelligence College, 2011).
8. Edwin Hutchins, *Cognition in the Wild*, (MIT Press, 1995).
9. Siavosh Sahebi and Paul Formosa, "The AI-Mediated Communication Dilemma: Epistemic Trust, Social Media, and the Challenge of Generative Artificial Intelligence," *Synthese* 205, no. 3 (2025): 1–24, https://doi.org/10.1007/s11229-025-04963-2.
10. Raja Parasuraman, Dietrich Manzey, "Complacency and bias in human use of automation: An attentional integration," *Human Factors* 52, (2010): 381–410. ■