

“The Incalculable Element”: The Promise and Peril of Artificial Intelligence

Zachery Tyson Brown

For the Intelligence Community, questions of how swiftly and how thoroughly GenAI should be integrated into its distributed collection and analysis architecture are being hotly debated just as they are in other sectors, only with far more at stake.

But the incalculable element in the future exercises the widest influence, and is the most treacherous.

– Thucydides, *History of the Peloponnesian War*

Recent advances in generative artificial intelligence (GenAI) have prompted tremendous excitement and significant trepidation while renewing a vigorous public debate about just what the growing capabilities of these systems portend for the future of work and society. There is another, parallel conversation going on, though, and this one is being held largely behind closed doors. For the Intelligence Community, questions of how swiftly and how thoroughly GenAI should be integrated into its distributed collection and analysis architecture are being hotly debated just as they are in other sectors, only with far more at stake.

GenAI is a subset within the larger field of AI research that uses transformer neural networks (this is what the ‘t’ in GPT stands for) in conjunction with so-called large language models (LLMs, a euphemism for volumes of semantic data sourced from the internet).¹ Narrower models of AI based on convolutional or recurrent neural networks are often relatively good within a particular field (such as medicine or law) but otherwise prove comically helpless. They are also

limited by the need for vast amounts of labeled data, which makes the process of ‘training’ them expensive and time consuming.

Now, with the combination of the transformer foundation—which enables the AI to mathematically examine the relationships between sequential data like the words in this sentence all at once—and a hefty LLM corpus to draw upon, the new breed of conversational AI applications such as OpenAI’s ChatGPT, Google’s PaLM, Anthropic’s Claude, and Meta’s LLaMA, among others, can respond to a wide variety of user queries, generating sophisticated media output in forms that are useful to humans—such as texts, images, videos, or even music.² They can, for instance, satisfactorily deconstruct or summarize complicated financial contracts and technical manuals, create surreal artwork, and even pass the bar exam.³

The IC’s interest in AI isn’t new.⁴ Indeed, some intelligence components have pursued AI solutions for decades.⁵ Historical efforts focused primarily on using rudimentary machine learning and computer vision to process and mine voluminous sets of data gathered from technical intelligence collection systems, to identify patterns and spot irregularities in overhead surveillance, and to automate certain time-consuming,

The views, opinions, and findings of the author expressed in this article should not be construed as asserting or implying US government endorsement of its factual statements and interpretations or representing the official positions of any component of the United States government.

Looking past both the hype and the histrionics, we find that the reality of GenAI is neither quite so wondrous nor quite so bleak as either the proselytizers or the doomsayers would have us believe.

routine processes.⁶ When it comes to intelligence analysis, however, the pursuit of algorithmic solutions to what are essentially fundamental problems of analytic uncertainty has met with mostly uninspiring results—until now.

The advent of AI applications that can rapidly ingest large datasets and from them “generate” useful content has convinced the IC’s AI proponents that we are at the onset of a new, revolutionary era of intelligence.⁷ This new era, they say, will be characterized by how well intelligence services leverage AI to collect, process, and analyze massive global data streams, and the United States will pay dearly if it falls behind rival nations in this new realm of competition.⁸ AI’s advocates can sound almost evangelical at times in their fervor, with some even going so far as to claim that generative AI will spell the end of intelligence analysis as a human activity altogether.⁹

Skeptics (like me), naturally, disagree, finding both the declarations of revolution and the dire warnings of necessity to be premature, and find the apparent drive to integrate generative AI applications into our most sensitive systems and processes to be irresponsible, at best. A more judicious approach would recognize the potential time-saving benefits of GenAI while keeping in mind the risks of relying over-much on what remain brittle, untested, and untrustworthy applications whose inner workings not even their designers can entirely explain.¹⁰ The revolutionary new era that AI’s proponents portend

might instead turn out to be one that is defined by which agencies make themselves overly dependent upon it—and therefore vulnerable.

Looking past both the hype and the histrionics, we find that the reality of GenAI is neither quite so wondrous nor quite so bleak as either the proselytizers or the doomsayers would have us believe. To be clear, transformer models are a genuinely remarkable achievement in the pursuit of artificial intelligence and have certain utility to the craft of analysis. The devil, however, is as always in the details—and there are ample reasons for us to be cautious about just how swiftly and how thoroughly we integrate these tools into our community’s most important work.

Will AI Revolutionize the Craft of Intelligence Analysis?

If you’ll forgive the use of a deservedly reviled analytic trope—it depends. It depends on what you mean by “AI” and on what you mean by “intelligence analysis.” As Alice Borene has argued, much routine issue updating and summarization of fragmented reporting that currently absorbs much of an analyst’s valuable time might feasibly be trusted to generative applications (albeit with humans in the loop to check their work).¹¹ But while generative AI can save time with instant summaries, those summaries can also be full of lies.¹² While generative AI can serve as a powerful tool to help analysts identify correlations and even to help spark new insights, it can also serve

to diminish the ability of analysts to think for themselves, becoming less of a prosthetic and more of a crutch.¹³ And while the IC faces many of the same challenges as the private sector, which AI can help address, it also faces many other challenges that are unique to the field of intelligence that AI might just make worse.

Intelligence is Uncertain

Even with the ability to swiftly parse volumes of material, GenAI still relies upon a wellspring of reliable data to be most effective. With certain and ample sources to pull from, GenAI applications can produce passable generalizations and even derive useful, if rudimentary, conclusions. The problem is that any issue worth being the subject of intelligence analysis is fundamentally uncertain, characterized by incomplete, ambiguous, and often contradictory snippets of partial, unreliable information.

Former NSA and CIA Director Michael Hayden once put it like this: “If it were a fact, it wouldn’t be intelligence.”¹⁴ What he meant was that because of the inherently uncertain nature of intelligence work, intelligence analysts do not, primarily, operate in the black-and-white world of facts. Intelligence analysis instead seeks to illuminate the gray spaces of the world that lie somewhere in between truth and fiction and that are rife with various and often conflicting sets of claims, assumptions, and inferences. These are not the sorts of issues that generative AIs, who already have problems telling the truth, are well suited for.

Only the most rudimentary of intelligence questions can be reduced

to a simple, binary answer of “yes” or “no” (did an event occur? Is a thing located in a place?). More frequent and more important intelligence questions concern intangibles like “will” and “intent” and the fundamentally unpredictable interactions of complex dynamic systems—the sorts of questions former National Intelligence Council Chair Gregory Treverton referred to as “mysteries.”¹⁵

Here we should recall Clausewitz’s timeless tenet: “A great part of information obtained in war is contradictory, a still greater part is false, and by far the greatest part is somewhat doubtful. What is required of an officer in this case is a certain power of discrimination.” Just as Clausewitz warned against attempting to reduce warfare to a crude “algebra of action,” we should resist all attempts to reduce intelligence to formulaic calculation.

Intelligence is, at its heart, about risk, contingency, and surprise. It is not concerned with averages, but rather with exceptions to the average. The truth is that, no matter how large your model may be, it will never encompass the world; that there is no amount of data that will permit the forecasting of novel events in an increasingly complex competitive environment wherein innumerable threads, material and immaterial, sympathetic and antagonistic, are all wound together in a Gordian knot of causality. In other words, while GenAI may be incredibly useful in comparatively tame, bounded fields such as advertising, customer service, medical sales, or even management consulting precisely because they excel at arranging and correlating regularities, the real world has no

It is impossible to fill every so-called intelligence gap. Missing knowledge is instead far more often mitigated by the insight, experience, and judgment of expert human analysts.

boundaries, and it is the irregularities that drive intelligence failures.

It is impossible to fill every so-called intelligence gap. Missing knowledge is instead far more often mitigated by the insight, experience, and judgment of expert human analysts. While informed by data where possible and appropriate, intelligence analysts are more concerned with nuance, judgment, and yes, with hunches—or what we might call tacit or implicit knowledge if you prefer a more scientific term.¹⁶

“Tacit” comes from a Latin word meaning silent. “Implicit” means folded in, referring to the fact that implicit knowledge is complex and layered. Implicit knowledge, then, is an emergent phenomenon that we sometimes call intuition or a “gut feeling.” It is something that arises from experience, knowledge of multiple domains, and the complexity of the human brain. In short, tacit knowledge emerges from the uniquely human synthesis of explicit knowledge, and this is what makes it incredibly difficult to quantify in terms a synthetic intelligence can parse. As polymath Michael Polanyi once put it, “we know more than we can tell.”¹⁷

Many of these issues derive ultimately from a deceptively seductive desire to make the world legible, and perhaps even to render it predictable, through the sheer accumulation of data. This fallacy is rooted in a misguided conflation of the social and natural sciences, in faith that the tools and methods of the former can

be made applicable to the latter, with a few tweaks. It posits a simplification of the inherently complex by forfeiting context, attempting to make the abstract straightforward and the messy machine-readable. Just as the military-technical revolution proponents mistakenly conflated sensing and targeting with strategy, the IC’s AI proponents can often appear to conflate data and information with intelligence.¹⁸ Both are examples of mistaking the tactical for the strategic, the finite for the infinite.¹⁹

As long centuries of history have demonstrated, however, reality is different. There is a fundamental and perhaps unbridgeable gap between the physical and natural and the social and political—and intelligence, strategy, and foreign policy are all primarily political (which is to say, adversarial) subjects.

Intelligence Is Adversarial

Intelligence analysis is not a neutral field of academic research—intelligence is a deeply adversarial political activity undertaken by, and directed primarily against, rival states. Because of its inherently antagonistic nature, intelligence is a discipline mired in lies. This aspect is not tangential; it is fundamental.

All intelligence is vulnerable to deliberate and at times elaborate deceptive measures undertaken by hostile foreign intelligence services, who have diverse and effective ways of concealing, distorting, and poisoning the information that intelligence

Intelligence analysis is not a neutral field of academic research—intelligence is a deeply adversarial political activity undertaken by, and directed primarily against, rival states.

services strive to acquire from them. Any means of intelligence collection is vulnerable to deception, from purportedly secret documents passed by recruited human sources to the ethereal streams of technical data gathered by billion-dollar satellites.

Because of these competitive aspects, intelligence work is most often not at all like data science—despite many attempts in recent years to make it so. In data analytics, the right piece of information can usually be found if the researcher designs their study well. Intelligence analysis is usually more like the old trope about putting together a jigsaw puzzle, only with someone else constantly trying to steal your pieces while also placing pieces of an entirely different puzzle into the pile you're working with. In a future competitive environment that is inundated with AI-generated content, this problem will only become worse, making it both easier to deceive and harder for analysts to accurately judge the capabilities and intentions of adversaries.²⁰

As a result, intelligence analysts are less like data scientists and more like judges. Both practice a sort of impartiality, as both are obligated by conscience and their sense of professionalism to adhere as closely to the facts as possible. Yet while the academic scientist can call their task complete once a topic is accurately observed and explained, the analyst (like the judge) must go further—she must pass judgment. She must pass judgment on many things: on which sources to consider reliable, on which to discount or cast aside, on how

much material a customer can feasibly make use of, on just how explicit to make their suppositions, inferences, and uncertainties. Intelligence analysts are also not completely impartial—they should, after all, want their own side to “win,” or at least to make better decisions in the face of uncertainty than their rivals, approaching something like truth in the outcome. The same cannot be said for GenAIs, which are currently challenged by the very notion of truth—because truth, at least so far as concerns humanity, is an ontological problem, not a mathematical one.

Promise and Peril

The truth is that no one knows what the future holds in store with AI, nor that future's implications for the world, let alone the IC. The engineers designing AI systems themselves do not know all of the potential uses (and hidden limitations) of what they are building, let alone how progressively variant iterations of those applications will be used ten, twenty, or fifty years from now.²¹ While we can catch glimpses of potential AI futures here in the present, these are fleeting, incomplete, and often illusory—much like intelligence work.

On the one hand, AI is already remarkably well suited to perform routine and time-intensive tasks that often make humans bored and thus prone to error. In the foreseeable future, analysts will be able to use AI to analyze handwriting collected by human agents, identify targets of interest from ubiquitous and persistent

space-based and air-breathing overhead surveillance platforms, and identify micro-expressions that may serve as tells during source interviews or interrogations. Researchers at Microsoft, for instance, have written a lengthy treatise that lists many other surprising ways in which computer vision applications can be used, for instance, reading an operating manual to learn how to pilot a machine or diagnose a patient.²² AIs will, for some tasks such as geolocation, make today's best open-source analysts look like amateurs. AI will be critical in fields such as cybersecurity, where the digital conversation between attacker and defender is incessant. We will require security applications that evolve at the speed of AI in a future where AIs develop and deploy cyber-weapons autonomously.²³

People, on the other hand, excel at higher-order critical, creative, imaginative, and innovative thinking under novel, unstructured, or ambiguous conditions. They are empathetic, imaginative, and capable of authentic emotional engagement with others. Empathy is key—both for the foreign adversaries that intelligence analysts seek to understand and anticipate, and, crucially, for the users of intelligence they exist to serve—because intelligence is at its core a characteristically human endeavor that is ultimately and essentially about the perspectives, thoughts, fears, desires, and behaviors of human beings.

More specifically, intelligence is most concerned with a relatively small number of individual human beings—the foreign political and military leaders and other decision-makers who inform the decisions to invest in a weapons program or a new hospital, to ally with or bandwagon

against, to wage war, or to sue for peace. While we can sometimes estimate the average behavior of groups of humans, the behavior and decisions of individuals are famously difficult, if not impossible, to predict. Expert intelligence analysts immerse themselves in the histories, languages, and cultures of foreign places to better adopt the perspective of those who live there, with all the inherent biases and preferences of their fears, internal contradictions, and other irrationalities that are the sum of personal experience and acculturation.

Fortunately, there is no shortage of both types of work in the intelligence world, which potentially makes for a match made in heaven. As the Director of CIA's Open-Source Enterprise Randy Nixon has said, "AI is a starting point."²⁴ If the IC applies AI smartly, it will lead to an empowerment of the human, not a loss of agency to the machines. AIs will soon be able to see, hear, listen, and speak to us, in real time, across multiple platforms. This means analysts will be able to have real conversations with seemingly intelligent digital assistants. Our AI assistants will generate our travel reports, draft our emails, and offer editorial corrections. They will serve as encyclopedia, thesaurus, and search engine all in one, all while managing our inboxes, deconflicting our meeting calendars, monitoring our favorite newsfeeds, and even offering advice or chiding us against bias creeping into our writing. A compelling vision of this future was sketched out by Joseph Gartin not so long ago in these very pages.²⁵

At the same time, intelligence services must be cognizant and cautious of the very real risks incurred by the

Expert intelligence analysts immerse themselves in the histories, languages, and cultures of foreign places to better adopt the perspective of those who live there, with all the inherent biases and preferences of their fears, internal contradictions, and other irrationalities that are the sum of personal experience and enculturation.

too-rapid or the too-comprehensive integration of AI into their collection and analysis enterprises, a sampling of which follows.

GenAIs are not pure search engines trained to find and source facts. They are instead a sort of voluble calculator, whose transformer modules mathematically predict the next likeliest word to appear in a string of text based on the corpus available to it. Thanks to the way this calculation works, generative AIs are famously susceptible to "hallucinations," that is, generating imaginary facts, figures, reports, quotes, and citations.²⁶ In one particularly egregious case, GenAI even fabricated a slew of legal opinions and judicial precedents when an attorney preparing a case asked for a brief.²⁷ It's important to understand that the LLMs are not lying to us. They're simply doing what they are told, which is giving users what they ask for—for better and worse. If the bulk of an LLM corpus suggests a citation should exist, the AI will generate it even if it doesn't. Generative AI will find what we tell it to find.²⁸ But, of course, the problem in finding what we're looking for is that we are often surprised when the things we're not looking for find us instead.

Another issue that should give the leaders of the intelligence community pause is the pernicious effects AI has on human reasoning. There are troubling (but perhaps, unsurprising) indications that extensive use of AI can render humans less

capable, not more. One study found that humans using a high-quality AI application became lazy and careless over time, letting the AI take over instead of using it as just another tool. The author, Harvard researcher Fabrizio Dell'Acqua, refers to this as users "falling asleep at the wheel"—which is one thing when it happens to recruiters, but quite another if it happens to intelligence officers responsible for informing strategic decisions.²⁹

A great amount of time and effort is spent training intelligence analysts to identify and mitigate the effects of bias, but bias is quite literally built into generative AI models. GenAI reflects the values, frames, and biases of the models they are built around and trained upon, regardless of how large they are. One of the chief restraints of at least the current slate of GenAI is the fact that all of these LLMs aren't quite as large as we might imagine. They represent only a narrow slice of the online world, and a slice that's predominantly white, male, and anglophone at that—which is to say, not exactly the most useful corpus for intelligence purposes (except maybe those of foreign intelligence services). Even more worrisome than AI making its users careless, perhaps, is another study that demonstrates human users of AI "absorbing" the system's built-in biases. Worse, these users were shown to retain those biases even after they stopped using the AI tool.³⁰

Like any tool, AI is, at its best, a prosthetic for authentic intelligence, which, if aptly applied, will help human intelligence analysts to better serve the users of intelligence.

Eventually, in a world where AIs are both generating and consuming content, over time everything might start to sound alike. Some AI researchers have even suggested that AIs trained on AI-generated content would eventually collapse under the recursive weight of semantic saturation, losing their ability to form coherent sentences altogether.³¹

In conclusion, intelligence leaders must remember that for all its seeming sapience, AI is, again, just a tool, and one for which there is no user manual. For some tasks, generative

AI can achieve remarkable results (although usually with significant prompting from a human user). For others, it outright fails, either gracefully or spectacularly, and the ability to distinguish between these outcomes beforehand is not always obvious. Like any tool, AI is, at its best, a prosthetic for authentic intelligence, which, if aptly applied, will help human intelligence analysts to better serve the users of intelligence. It will augment our understanding, reasoning, and yes, even our creativity.³²

AI holds promise and peril for the craft of intelligence analysis, and the ways IC leaders choose to employ it will make all the difference. If intelligence officials are seduced by the characteristically American conceit that the world can be made legible through technical means alone, they will be frustrated when the algorithms inevitably break in the face of nuance or novelty. Alternatively, if they choose instead to adopt a more cautious and judicious approach to use AI as primarily an aid to human analysis, harnessing its potential to augment the natural ingenuity and empathy of our community's analytic cadre, the IC—and more importantly, the nation—will benefit.



The author: Zachery Tyson Brown is a national security futurist with The Aerospace Corporation's Strategic Foresight Team. His focuses on the intersection of strategy, policy, and emerging and disruptive technologies. He is a US Army veteran, former defense intelligence officer, and a graduate of the National Intelligence University.

Endnotes

1. David Neild, "How ChatGPT and Other LLMs Work," *Wired* (April 2023).
2. Visual Storytelling Team and Madhumita Murgia, "Generative AI Exists Because of the Transformer," *Financial Times* (September 2023).
3. Jonathan H. Choi, Kristin E. Hickman, Amy Monahan, Daniel B. Schwarcz, "ChatGPT Goes to Law School," *Journal of Legal Education* 387, (January 23, 2023), <http://dx.doi.org/10.2139/ssrn.4335905>.
4. Christopher R Moran, Joe Burton, George Christou, "The US Intelligence Community, Global Security, and AI: From Secret Intelligence to Smart Spying," *Journal of Global Security Studies* 8, no. 2 (June 2023), <https://doi.org/10.1093/jogss/ogad005>.
5. CIA, "John McMahon: Address to the Second Annual Symposium on AI Applications in the Intelligence Community. CIA HQ Auditorium." CIA Records Search Tool (CREST), March 1983, <https://www.cia.gov/readingroom/document/cia-rdp-87m00539r000600730001-2>.
6. Jeffery Richelson, *The Wizards of Langley: Inside the CIA's Directorate of Science and Technology* (Westview Press, 2001).
7. Anthony Vinci, "The Coming Revolution in Intelligence Affairs." *Foreign Affairs* (August 2020), <https://www.foreignaffairs.com/articles/north-america/2020-08-31/coming-revolution-intelligence-affairs>.
8. Billy Mitchell, "Those who adopt AI will disrupt those who do not, CIA cyber policy adviser says," FedScoop (October 2023), <https://fedscoop.com/those-who-adopt-ai-will-disrupt-those-who-do-not-cia-cyber-policy-adviser-says/>.
9. John F. Galascione, "The End of Human Intelligence Analysis—Better Start Preparing," *Studies in Intelligence* 67, no. 4 (December 2023).
10. Neil Savage, "Breaking into the Black Box of Artificial Intelligence," *Nature* (March 2022), <https://www.nature.com/articles/d41586-022-00858-1>.
11. Alice B. Borene, "Intelligence analysis, Synthesis, and Automation," *Studies in Intelligence* 67, no. 4 (December 2023).
12. Caroline Mimbs Nyce, "AI Search is Turning Into the Problem Everyone Worried About," *The Atlantic*, November 2023, <https://www.theatlantic.com/technology/archive/2023/11/google-generative-ai-search-featured-results/675899/>.
13. Fabrizio Dell'Acqua, "Falling Asleep at the Wheel: Human/AI Collaboration in a Field Experiment on HR Recruiters," Working Paper, Harvard Business School (January 2022).

14. Bob Woodward, *Plan of Attack* (Simon & Schuster, 2004), 134.
15. Gregory F. Treverton, "Risks and Riddles," *Smithsonian Magazine* (June 2007), <https://www.smithsonianmag.com/history/risks-and-riddles-154744750/>.
16. Sherman Kent, *Strategic Intelligence for American World Policy*, (Princeton University Press, 1949), 203.
17. Michael Polanyi, *The Tacit Dimension* (University of Chicago Press, 1966), 4.
18. Colin S. Gray, *Airpower for Strategic Effect* (Air University Press, 2012), 358.
19. Everett Carl Dolman, *Pure Strategy: Power and Principle in the Space and Information Age* (Routledge, 2011).
20. Gian Volpicelli, "The End of Photographic Truth," *Politico* (October 2023), <https://www.politico.eu/article/ai-photography-machine-learning-technology-disinformation-midjourney-dall-e3-stable-diffusion/>.
21. Will Knight, "The Dark Secret at the Heart of AI," *MIT Technology Review* (April 2017), <https://www.technologyreview.com/2017/04/11/5113/the-dark-secret-at-the-heart-of-ai/>.
22. Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, Lijuan Wang, "The Dawn of LMMs: Preliminary Explorations with GPT-4V(ision)" Microsoft Corporation (October 2023), <https://arxiv.org/pdf/2309.17421.pdf>.
23. Hannah Murphy, "AI: A New Tool for Cyber Attackers, or Defenders?" *Financial Times* (September 2023), <https://www.ft.com/content/09d163be-0a6e-48f8-8185-6e1ba1273f42>.
24. Quoted in the State Secrets podcast, "The Future of Open Source: A View from the CIA," *The Cipher Brief* (September 2023), <https://www.thecipherbrief.com/podcasts/state-secrets-podcast-interview-cia-osint-randy-nixon-announcement>.
25. Joseph W. Gartin "The Future of Analysis," *Studies in Intelligence* 63, no. 2 (June 2019).
26. Will Knight, "Chatbot Hallucinations Are Poisoning Web Search," *Wired*, October 2023.
27. Benjamin Weiser and Nate Schweber, "The GPT Lawyer Explains Himself," *New York Times*, June 8, 2023.
28. Cade Metz, "Chatbots May Hallucinate More Often Than People Realize," *New York Times*, November 6, 2023.
29. Fabrizio Dell'Acqua, "Falling Asleep at the Wheel."
30. Lauren Leffer, "Humans Absorb Bias from AI—And Keep It After They Stop Using the Algorithm," *Scientific American* (October 2023).
31. Darren Orf, "A New Study Says AI is Eating its Own Tail," *Popular Mechanics* (October 2023).
32. Ethan Mollick, "Automating Creativity," *One Useful Thing* (August 2023), <https://www.oneusefulthing.org/p/automating-creativity>.

