review essay

Language Machines
Cultural AI and the End of Remainder Humanism

Reviewed by Sean Barnes

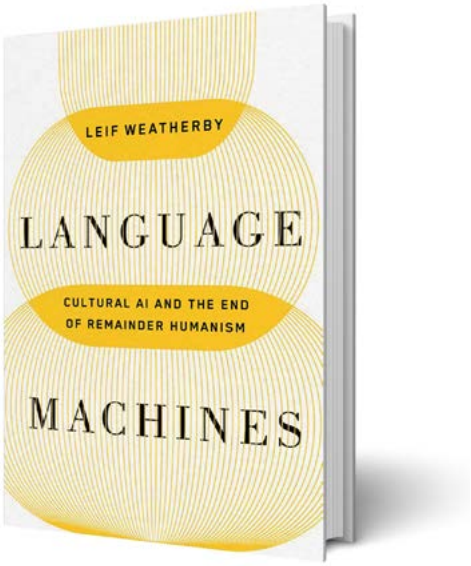

Author: Leif Weatherby
Published By: University of Minnesota Press, 2025
Print Pages: 279
Reviewer: Sean Barnes is a retired CIA officer and tradecraft instructor.

Sherman Kent, under whose mighty shadow analysts labor, decreed that intelligence analysis exists to serve policymakers, not to pursue knowledge for its own sake. As CIA tradecraft expert Jack Davis reminded us, Kent was pretty blunt about this—intelligence that gets ignored is "useless."[1] Kent made it look easy, his method, drawn from a background in historical scholarship, deceptively simple: test your sources, then check your own biases before drawing conclusions.[2] As any analyst will tell you, the former is simple enough in theory; the latter is brutally difficult in practice.

This grounding matters when we turn to Leif Weatherby's *Language Machines*: *Cultural AI and the End of Remainder Humanism,* a likewise deceptively simple book that examines generative artificial intelligence (GenAI) through cultural theory rather than intelligence tradecraft. Weatherby uses structuralism and semiotics to make a striking claim—that large language models (LLMs) function as ideology machines, creating meaning through pattern recognition while embedding cultural assumptions at scale. For veterans of the intelligence trade, this isn't just academic theorizing. It means, in a nutshell, that the bias problems Richards Heuer warned about have mutated into something a thousand times harder to spot and a million times easier to spread.[3]

So let's test Weatherby's arguments against Kent School principles. The question isn't whether his cultural theory holds water academically—it's whether his insights can help us understand why GenAI might belong on our desks as something to analyze, not something to trust with analysis itself.

The Language Wars Nobody Wins

Weatherby opens by sketching the battlefield where linguists and engineers have been fighting for decades. In one corner, Noam Chomsky and his followers insist language springs from hardwired cognitive structures—universal grammar built into our brains. To them, LLMs are just stochastic parrots, mimicking without understanding. In the other corner, the statistical crowd argues that patterns alone can generate language. Point to any word, they say, and you can predict it from the words around it.

Both camps miss something crucial here. The syntax theorists dismiss LLMs too quickly; the statisticians overvalue them. Neither grasps what structuralists figured out long ago—language works through relationships between elements, not through pure thought or raw probability. Weatherby shows us that LLMs generate meaning precisely through these structural relationships, which explains both their surprising capabilities and their fundamental limitations.

Any analyst reading this should hear Heuer's voice warning about anchoring bias. Chomsky's people anchor on cognition, the engineers anchor on statistics. Each clings to their model like it's gospel, dismissing evidence that complicates their worldview.[4] What Weatherby offers—without saying it directly—is something like Heuer's structured technique known as analysis of competing hypotheses. Don't pick a side; run multiple explanations in parallel and see which one explains the evidence best.

Here is where things get a bit unsettling. Roland Barthes confidently declared the "death of the author" back in the 1960s, and literary critics subsequently had an intellectual field day over the next several decades, arguing whether texts could mean anything without knowing who wrote them and why. LLMs make this fun, abstract debate suddenly concrete. These machines produce text without consciousness, intention, or purpose. In the vaunted context of everything we've been taught about intelligence analysis, the meaning emerges from patterns in data, not from any mind trying to communicate.

Think about what this means for intelligence work. We're trained ad nauseam to first ask who wrote something and why. Heuer documented our obsession with intent—we see purposeful actors everywhere, even in random events. But when an LLM generates a report, there's no "who" behind it in any meaningful sense. No motive to uncover, no institutional bias to account for (at least not in the traditional way). The bias is there, but it's systemic, embedded in training data and emerging from statistical relationships we can't fully trace.

Left unaddressed, this is chilling. Should the natural momentum generated by "this latest breakthrough" be allowed to roll on down the proverbial hill, it completely scrambles our source evaluation methods. We built those methods around human psychology and institutional analysis. What good are they when the "source" is a mathematical function operating on terabytes of text? The game has changed, and we're still using the old rulebook.

Attention Without Intention

In his third chapter, Weatherby sheds some light on this potential disaster by digging into the technical heart of these systems—the attention mechanism that makes transformer models tick. Don't let the jargon intimidate you; the concept is straightforward. Attention lets the model decide which parts of the input matter most for predicting what comes next. It's like having a spotlight that can illuminate different parts of a sentence depending on what the model needs to generate.

For analysts, this should sound pretty familiar. We do this constantly and eventually unconsciously—deciding which signals matter in an ocean of noise, which sources deserve scrutiny, which patterns might be significant. Heuer showed how badly we do this, overweighting dramatic information while missing crucial baseline data.[5] The attention mechanism formalizes this process mathematically, encoding selective focus into the model's architecture.

But there's a catch. We can metricize to our heart's content. We can see the attention weights, watch which tokens get prioritized, but we can't fully explain why. The model makes these choices through layers of mathematical transformation that resist human interpretation. Kent demanded clear audit trails for every analytical judgment, and this is at the core of everything intelligence analysts are taught from day one. These systems offer the opposite—sophisticated pattern matching with no ability to explain their reasoning. That's not analysis; it's expensive guesswork.

Weatherby's most cutting observation might be his characterization of what LLMs actually produce. It's not poetry or insight—it's kitsch. Think hotel lobby art or elevator music: familiar forms endlessly recycled, technically competent but utterly without soul or genuine creativity. The models reproduce patterns they've seen thousands of times, creating outputs that feel meaningful but lack any real depth.

The errors these systems make—what we call hallucinations—are not random glitches. They're windows into how the model associates concepts based on cultural patterns in its training data. When ChatGPT confidently states something false, it's showing us how ideas cluster together in the data it consumed. As Weatherby reminds us, these aren't logical errors; they're cultural artifacts.

Heuer would recognize this right off the bat. Like any machine, human analysts, even the best of us, hallucinate too—we see patterns that aren't there, connect dots that shouldn't be connected. But at least our errors follow predictable patterns we can study and counter. LLMs scale this tendency massively, churning out ideological dross that, for all the world looks authoritative but reflects nothing more than statistical regularities in human writing.

Yet there's an opportunity hidden in this problem. These errors and biases aren't just failures—they're potential diagnostic tools. By studying what the model gets wrong and how it gets it wrong, we can map cultural assumptions and ideological structures embedded in the training data. For information operations analysis, for example, this is a good thing. The kitsch might be worthless as intelligence, but it can be invaluable for understanding how narratives construct themselves.

Form Over Facts

Weatherby calls for developing "a general poetics of AI"—basically, learning to read these outputs for their structure and rhetoric rather than their factual content. Language creates meaning through form before it establishes any reference to reality. LLMs prove this point definitively. They can generate perfectly structured arguments about completely fictional topics, showing that formal coherence matters more than truth in language generation.

This insight cuts deep. Like Marine Corps recruits at Parris Island, intelligence analysts are trauma-conditioned into a reflexive focus on facts, to verify claims and establish ground truth. But in modern analysis, the form often matters more than the content. A well-structured lie spreads faster than a poorly presented truth. Weatherby's pushing us to develop what amounts to rhetorical literacy—reading for persuasive technique as much as factual accuracy.

Still, the black-box problem haunts us, as well it should. We can analyze the rhetoric, map the persuasive strategies, understand the cultural resonances, but we cannot trace the generative process. We are left interpreting outputs without understanding origins. In intelligence terms, that's like analyzing a document without being able to verify its provenance. Possibly useful in a few, carefully managed instances, downright deadly in most others.

The Ideology Engine

By the book's end, Weatherby's verdict is clear: LLMs are ideology machines, not thinking machines. They don't produce neutral language; they package and deliver clusters of cultural meaning drawn from their training data. Every output reinforces certain narratives while suppressing others. The hallucinations just make

this process visible, showing us which ideas travel together in the model's learned associations.

Heuer warned about cognitive lock-in—how mental models resist change once they form. LLMs supercharge this problem. They take the biases in their training data and reproduce them endlessly, making alternatives harder to imagine with each iteration. For adversaries running information operations, this is a gift: a propaganda engine that packages ideology with perfect grammar and distributes it efficiently. For analysts, it's a nightmare scenario.

The challenge is not just recognizing bias—analysts are pretty good at that. It is unpacking the ideological clusters, finding the assumptions buried in seemingly neutral language, surfacing the alternatives that the model's outputs systematically hide. Kent would call this a form of deception, though not necessarily intentional. The deception is built into the architecture, emerging from statistical mimicry rather than conscious design.

Why Transparency Isn't Optional

Everything circles back to the transparency problem. Kent insisted that every estimate's reasoning be clear and auditable. Heuer demanded we surface and examine our cognitive biases. LLMs fail both tests spectacularly. They produce outputs through processes we can't fully trace, embedding biases we can't completely identify.

This isn't a technical problem we'll solve with better engineering. It's a fundamental epistemological challenge. Intelligence analysis rests on the ability to show our work, to let others check our reasoning, to build confidence through transparency. LLMs offer the opposite—polished outputs with opaque origins.

For practical purposes, this means LLMs can't be trusted with finished intelligence products. They might help with brainstorming or red-teaming, generating alternative perspectives or challenging assumptions. But relying on them for analytical conclusions would be professional malpractice. It would violate everything the IC tradecraft stands for and potentially put lives at risk.

Making It Practical...Really

Despite the dense theory, Weatherby's analysis offers immediate practical value. He's shifted the conversation from abstract questions about machine consciousness to concrete issues about how machines create meaning. That's exactly the right focus for intelligence practitioners. His framework gives us beleaguered practitioners viable tools we can use today. In this context, hallucinations instead become diagnostic instruments for detecting cultural bias. Kitsch production reveals ideological assumptions, etc. The attention mechanism, despite its opacity, shows us something about how the model prioritizes information. These aren't complete solutions, but they're actionable insights.

The book has obvious weaknesses. Weatherby writes like an academic for academics, which can be frustrating when you need operational guidance. He paints with a broad brush—not everything LLMs produce is ideological kitsch, and occasionally they generate honest insights. But these limitations don't negate the core value: he's already mapped the terrain that we will need to navigate, even if he hasn't provided turn-by-turn directions.

Several concrete steps follow from this analysis. First, any workflow involving LLM-generated content needs explicit bias checks. Do not assume neutrality; assume ideological packaging and look for it systematically. Document these checks. Make them mandatory.

Second, use hallucinations strategically. They're terrible for fact-finding but excellent for red-teaming. When an LLM confabulates, it's showing you narrative patterns and cultural assumptions. Mine those errors for intelligence about how ideas connect in the information environment.

Third, expect adversaries to weaponize these capabilities. They will use LLMs to generate and spread ideological content at unprecedented scale. Build capabilities to track machine-generated narrative patterns. Update these capabilities regularly as techniques evolve.

Fourth, add rhetorical literacy to analytical training. Analysts need to read form as carefully as content. This isn't optional anymore—it's as essential as source validation or link analysis. Make it part of the core curriculum, not an elective.

Finally, establish clear policies about LLM use. These tools can support hypothesis generation and assumption challenging, but never final estimates or policy recommendations. When analysts use them, require documentation. Make the limitations explicit. Verify that machine-generated material isn't contaminating critical judgments.

Heuer taught us that "seeing is not believing"—that our perceptions are filtered through cognitive biases we rarely recognize. Sherman Kent demanded analytical transparency and rigor. Weatherby, writing from the humanities, extends their warnings into the age of AI. His core message is stark: these are language machines, not thinking machines. They package ideology into fluent prose that deceives through its very polish.

Two conclusions are inescapable. First, treat AI outputs as cultural signals requiring interpretation, never as analytical judgments deserving trust. Second, until these systems can show their work—really show it, with full transparency—they remain objects of analysis, not tools for conducting it. The machines speak with increasing fluency, but fluency isn't intelligence. They produce sophisticated patterns, but patterns aren't understanding. They generate compelling narratives, but narratives aren't truth. Until we can see inside the black box, until we can audit the reasoning, until we can trace the logic, we'd be fools to mistake their outputs for analysis.

Weatherby's book matters because, thankfully, it reminds us that the old lessons still apply. Transparency, rigorous tradecraft, and systematic bias awareness—these aren't outdated concepts from a pre-digital age. They're the foundations that keep intelligence work honest and reliable. No algorithm changes that. No technology replaces it. The fundamentals endure because they must endure. Without them, we're not analysts anymore—we're just consumers of sophisticated kitsch, mistaking eloquence for truth.■

Endnotes

1 Jack Davis, "Sherman Kent and the Profession of Intelligence Analysis," Sherman Kent Center for Intelligence Analysis, Occasional Papers 1, no. 5 (November 2002).

2 Sherman Kent, *Strategic Intelligence for American World Policy* (Princeton University Press, 1949), 15.

3 Richards J. Heuer Jr., *Psychology of Intelligence Analysis* (CIA Center for the Study of Intelligence, 1999).

4 Ibid., 45–52.

5 Ibid., 14–21.■