

BrokenSystems.io



Issue No. 004

NEWSLETTER

July 2026

What do we really know about AI? The Minab, Iran bombing

Artificial intelligence has moved from myth to a mythology of its own making. A short time ago, the term was generic, a placeholder for machines that thought. Today it has names: Claude, Gemini, ChatGPT, Grok, each with its own reputation, its own devotees, its own controversies. When precision matters less than category, the acronym LLM gathers them back into a single family. Yet naming a thing is not the same as knowing it. Each of us wants to know what AI can do and what it will eventually do, and each of us is, strictly speaking, unqualified to answer. What circulates as certainty is conjecture dressed in the confidence of expertise. Reports document its capabilities well enough, and its reshaping of the workplace is already visible, though where the reshaping ends remains anyone's guess.

The press reports what AI does, but the press is not a disinterested instrument. Its reliability is uneven, its errors frequent, and its typical account rarely rises to the evidentiary standard readers grant it out of habit. That gap, between what a story claims and what a story has actually proven, should send a careful reader toward primary sources before belief takes hold. It rarely does. Verification is the discipline most invoked and least practiced, precisely when the story concerns something as unproven and as easily embellished as artificial intelligence's future. An unknown future is fertile ground for inventing one.

Consider, then, AI's use by government, a subject that troubles observers well beyond this publication. Anthropic itself named specific concerns rather than vague ones. That concerns were raised is easy to establish. That the events those concerns anticipated actually occurred is a separate question, and a harder one. Take the bombing of the Shajareh Tayyebbeh elementary school in Minab, Iran. That the strike happened is not contested by anyone. What happened inside the decision that led to it is another matter, resting on claims of varying strength, some documented, some merely reported. Before accepting any account of why the school was struck, it is worth asking a question courts ask as a

matter of course: would this evidence be permitted to prove guilt? If the answer is no, the account deserves the same caution in a newsletter that it would receive in a courtroom.

The Minab strike's physical facts rest on strong footing: a Tomahawk strike hit the Shajareh Tayyebbeh school on February 28, killed somewhere between 150 and 175 people, and struck a building that satellite imagery confirms had been physically separated from the adjacent IRGC naval base years earlier. None of that depends on anyone's unverified account of what happened inside a targeting cell.

The explanation for the strike is a different matter entirely, and it is worth being exact about how different. Every version of the DIA-outdated-data narrative, from the original NYT reporting to CNN's account of overridden MIDB and MARS warnings, rests on unnamed people describing an internal investigation whose actual findings the Pentagon has never released. No official has put a name to any part of this explanation. If offered in court to prove DIA data caused the targeting error, it would be hearsay on hearsay: an anonymous source's account of an unreleased inquiry, repeated by a reporter who witnessed neither the targeting decision nor the inquiry itself. It fits no exception under Rule 802¹ for the same reason it can't be cross-examined: no one will attach their name to it.

This is not a mark against the reporting. It may be exactly what happened, and anonymous sourcing is frequently the only route by which an institution's own failures become visible before the institution is ready to admit them. But a publication built on diagnosing broken feedback channels cannot treat an unreleased, anonymously sourced explanation as though it carries the same evidentiary weight as Cameron Stanley's sworn declaration^{II} or Kirsten Davies's testimony under oath. The most consequential fact in this episode, why 150 to 175 people died, has not been tested by an oath, a signature, or a courtroom. It has only been described. That the death toll is certain and its cause is not itself worth sitting with.

Grok's Military Role Disclosed via Court Filing

The Same Standard, Applied Evenly

A claim does not earn trust by comparison alone. That Grok's court disclosure outranks the anonymously sourced Minab mechanism is not the same as saying the disclosure is beyond question, and a publication that applies scrutiny selectively has not really applied it at all.

Cameron Stanley's declaration differs from the Minab sourcing in a way that matters. It is not a reporter's account of what someone said. It is the thing itself: a document filed under penalty of perjury, bearing a named federal official's signature, entered into an active federal case as Exhibit A. Every outlet that reported it, WIRED, the Independent, Yahoo, Common Dreams, Benzinga, quoted the same filed language, which means a reader can verify the claim against the court record directly rather than trusting any single publication's paraphrase. That is what a primary source offers and an anonymous source cannot: a fixed text, open to inspection, that does not depend on anyone's credibility but its own signatory's.

Yet a sworn statement is not a disinterested one. Stanley did not volunteer this account to a journalist investigating military AI use. He filed it as part of the government's own defense in litigation the government itself has an interest in winning, arguing that Grok's operational value to national security is a reason the courts should let a Mississippi power plant keep running unpermitted. A declarant with a stake in the outcome has a reason, however small, to

describe that value generously. The statement may be entirely accurate. It may also be shaded by the very motive that produced it, filed to persuade a judge, not to inform the public.

Set beside it, Kirsten Davies's account of Claude'sⁱⁱⁱ role took a different shape. She did not offer it. Senator Reed extracted it under adversarial questioning, in a hearing where the more comfortable answer for the Pentagon would have been silence or evasion about a program that directly contradicted its own supply chain risk designation. An admission given under pressure, against the interest of the institution making it, carries a different kind of weight than one filed in support of that institution's own case. Both statements are sworn. Both are named. Only one was volunteered to help the speaker's side win.

None of this reopens whether Grok or Claude were used in Epic Fury; both facts stand on solid ground, better ground than the Minab mechanism claim by a wide margin. What it reopens is the comfortable habit of treating "sworn and on the record" as a single undifferentiated tier of trust. It is not. A sworn statement made to help its author's case and a sworn statement extracted against its author's interest belong to the same evidentiary family, but they do not belong to the same evidentiary weight class, and a careful reader owes both a slightly different level of scrutiny.

Comparison: Was Claude Also Used in Epic Fury?

What Confirmation Does Not Cover

Claude's presence in Epic Fury reached the public record through two separate channels rather than one, which is itself a stronger pattern of corroboration than the Grok disclosure enjoyed. The Wall Street Journal reported CENTCOM's use of Claude in the Iran campaign first, through its own reporting process. Weeks later, under Senator Reed's direct questioning, Kirsten Davies confirmed the same fact from the opposite direction, an official statement that did not depend on the Journal's sourcing and would have stood on its own had the Journal never published a word. Two independent paths converging on the same fact is a stronger position than either path alone, and stronger still than Stanley's single filed declaration.

The context surrounding that confirmation deserves equal weight. The Pentagon had designated Anthropic a supply chain risk on March 6, after Dario Amodei declined to grant the military unrestricted use of Claude for domestic surveillance or autonomous weapons. A 180-day removal window followed. Whatever Davies confirmed, she confirmed about a system her own department had, weeks earlier, formally labeled unacceptable, an admission with the same adversarial texture as the one examined above, extracted

from an institution with every incentive to minimize it rather than volunteer it.

But confirmation has a boundary, and the boundary matters more here than anywhere else in this issue. What is established is that Claude was operationally embedded in Epic Fury as a whole. What is not established, by any source, named or anonymous, is whether Claude played any role in the Minab strike specifically. Amodei has said plainly that Anthropic itself lacks visibility into how Claude was used operationally once deployed, and does not know whether it touched the Minab targeting chain at all. That is not evasion. It is the most honest sentence available on the subject, an acknowledgment that the party best positioned to know cannot answer the question everyone actually wants answered. Some aggregator and opinion sources have closed that gap anyway, framing Claude's confirmed presence in the broader operation as though it settled the narrower question of Minab. It does not. Primary sources establish the first claim. No source, primary or otherwise, establishes the second, and a publication built on evidentiary discipline cannot borrow the confidence of one claim to cover a debt the evidence has not paid on the other.

The Same Word, Three Different Weights

Three claims, three different evidentiary architectures, one word applied loosely to all of them: confirmed. The Minab mechanism rests on people who cannot be named, testifying to an inquiry that has never been released, a claim that would not survive a hearsay objection in any courtroom in the country. Grok's role rests on a named signature, filed under oath, entered as an exhibit any reader can inspect, yet offered by a declarant arguing his own side's case, testimony strengthened by its documentary form and weakened by its motive. Claude's role rests on the same sworn, named, on-record footing, reached independently through two separate channels, extracted from an official with every institutional reason to withhold it rather than volunteer it, and bounded honestly by an admission of what even its own maker does not know. None of these claims are equally trustworthy and treating them as equally trustworthy because each arrived wearing the same word is precisely the failure this issue set out to examine.

What separates them is not whether someone said something happened. It is who said it, under what obligation, against what interest, and subject to what test. A statement volunteered to a friendly audience is not the same as a statement extracted by an adversary. A statement a reporter attributes to unnamed others is not the same as a statement its own author signed under penalty of perjury. And a confirmed fact about an operation as a whole is not the same as a confirmed fact about one incident within it, a distinction the Minab-Claude question makes unavoidable, since the honest answer to "was Claude involved in the strike on the school" is that no one, including Anthropic, currently knows. A publication that diagnoses broken feedback channels for a living owes its own readers the same channel discipline it asks institutions to practice, report what the evidence supports, mark plainly where it runs out, and resist the pull to let one well-sourced claim lend its confidence to a neighboring claim that has not earned it on its own.

Abbreviations & Definitions

Cameron Stanley. Currently, Chief Digital and Artificial Intelligence Officer at Chief Digital and Artificial Intelligence Office. <https://www.ai.mil/>

CENTCOM - U.S. Central Command. <https://www.centcom.mil/>

Dario Amodei - Current Anthropic CEO

DIA - Defense Intelligence Agency. <https://www.dia.mil/>

Epic Fury - On Feb. 28, 2026, the U.S. military commenced Operation Epic Fury under the direction and direct orders of the president of the United States. U.S. and partner forces are striking targets to dismantle the Iranian regime's security apparatus, prioritizing locations that pose an imminent threat.

IRGC - The Islamic Revolutionary Guard Corps. <https://www.cfr.org/backgrounders/irans-revolutionary-guards>

Kirsten Davies – Current Department of War Chief Information Officer. ([Web](#))

Jack Reed – Current U.S. Senator for Rhode Island. <https://www.reed.senate.gov/>

MARS - Machine-assisted Analytic Rapid-repository System. <https://www.gao.gov/assets/gao-21-57.pdf>

MIDB - Modernized Integrated Database. <https://www.gao.gov/assets/gao-21-57.pdf>

NYT - New York Times. <https://www.nytimes.com/>

Tomahawk - a precision weapon that launches from ships, submarines, and ground launchers and can strike targets precisely from 1,000 miles away, even in heavily defended airspace. <https://www.rtx.com/raytheon/what-we-do/sea/tomahawk-cruise-missile>

End Notes

ⁱ Rule 802. The Rule Against Hearsay. https://www.law.cornell.edu/rules/fre/rule_802

ⁱⁱ <https://www.congress.gov/119/meeting/house/119184/witnesses/HHRG-119-AS35-Wstate-StanleyC.pdf>

ⁱⁱⁱ <https://www.c-span.org/program/senate-committee/defense-department-officials-testify-on-it-security-and-operations/676113>