Light † Faith

Research

AI Fails Spiritual Concerns

April 29, 2026

*The following began as an innocent query, but the shocking results returned by Artificial Intelligence (AI) led to posting a thread with these findings on **X.com** [1]. They are combined here as a research paper for easy reading, and made available for print or distribution as a PDF file.*

An example of why AI fails, and will continue to fail, to address spiritual concerns about allowing it to have a greater influence on all of our lives, than of being a servant of God conscious individuals in society.

With the arrival of **QWEN 3.6-27B** [2] (advanced LLM AI), this writer thought to try it out by asking about the "Red Square Nebula," an uncanny symmetrical object imaged by astrophysicists Peter Tuthill & James Lloyd.

Below is a partial screenshot of their **website** [3].

This image (left) announces a new arrival in the pantheon of exotically beautiful celestial objects. We have christened this startlingly symmetrical nebula "The Red Square" for its color and form, and also in recognition of its close cousin - the celebrated Red Rectangle nebula.

The new discovery was made with advanced imaging technologies known as Adaptive Optics while studying a hot star called MWC 922 located about 5000 light-years from earth in the constellation Serpens (the serpent mythologically associated with medicine). The image to the left, which combines data from the Mt Palomar Hale telescope and the Keck-2 telescope , was taken in near-infrared light (1.6 microns) and shows a region 30.8 arcseconds on a side around MWC 922. As the outer periphery of the nebula is very faint compared to the core, the image has been processed and sharpened to display the full panoply of detail and structure.

The startling degree of symmetry and level of intricate linear form make the Red Square nebula around MWC 922 the most symmetrical object of comparable complexity ever imaged. The overall architecture displays a twin opposed conical cavities (known as a bipolar nebula), along the axis of which can be seen a remarkable sequence of sharply defined linear rungs or bars. This series of rungs and conical surfaces lie nested, one within the next, down to the heart of the system, where the hyperbolic bicone surfaces are crossed by a dark lane running across the principle axis.

One particularly fascinating feature visibile in the images is a series of faint radial spokes, like teeth of a comb, pointing away from the center. Structures such as this are rarely seen in nebulae, and the high degree of regularity in this case may point to the intriguing possibility that these bands are shadows cast by periodic ripples or waves on the surface of an inner disk close to the star at the heart of the system.

Perhaps the most compelling implication from the discovery of such a cleanly-constructed polar ring system comes when we consider the three dimensional structure implied by the Red Square images. If we imagine moving away from the precise (and fortuitous) edge-on viewing angle onto this object that we find from Earth, we might get a view like that depicted below. The left-hand panel shows the skeleton of the twin opposed cones as we see them from earth, but if we rotate away from our view at 90 degrees to the axis (middle panel) we can visualize that the bright bars seen edge-on become elliptical rings encircling the polar axis of the system (right panel).

The Red Square Nebula not only has its own **Wiki page** [4], but has been featured several times on NASA's "Astronomy Picture of the Day" website in **2007** [5], **2011** [6], **2016** [7], and **2021** [8].

Researching this object since 2016, this writer was curious if AI could provide answers about its creation. Quoting Wiki:

> "How did all this beautiful, crisp structure form?" asks Peter Tuthill of the University of Sydney in Australia. "This is the million dollar question."

The question is posed to Qwen3.6-27B AI:

> How is the "Red Square Nebula" formed? Provide math equations and visual proof

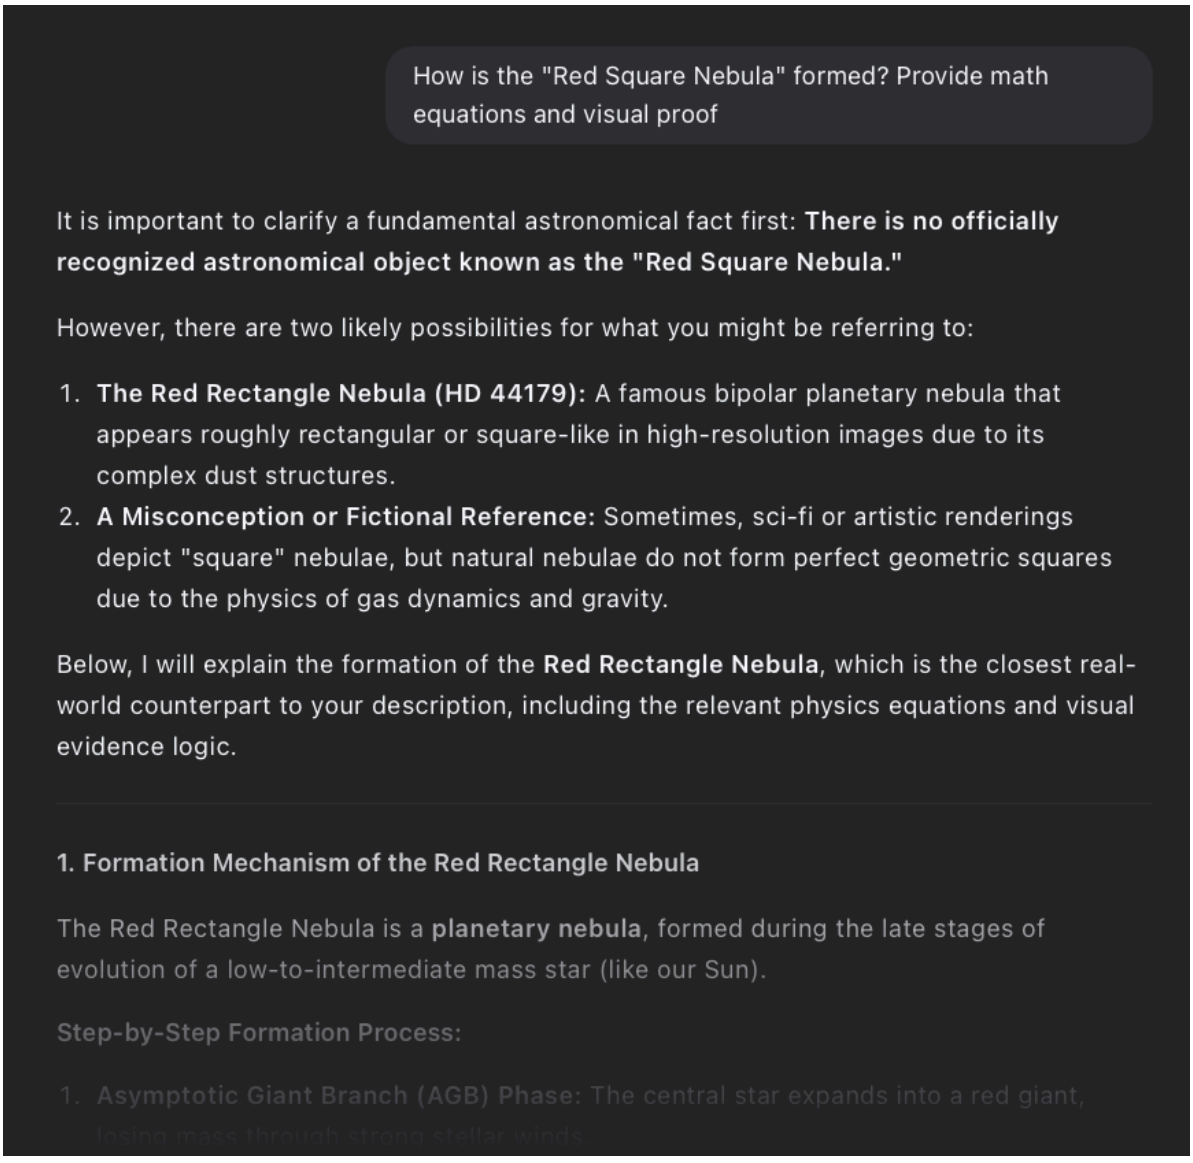

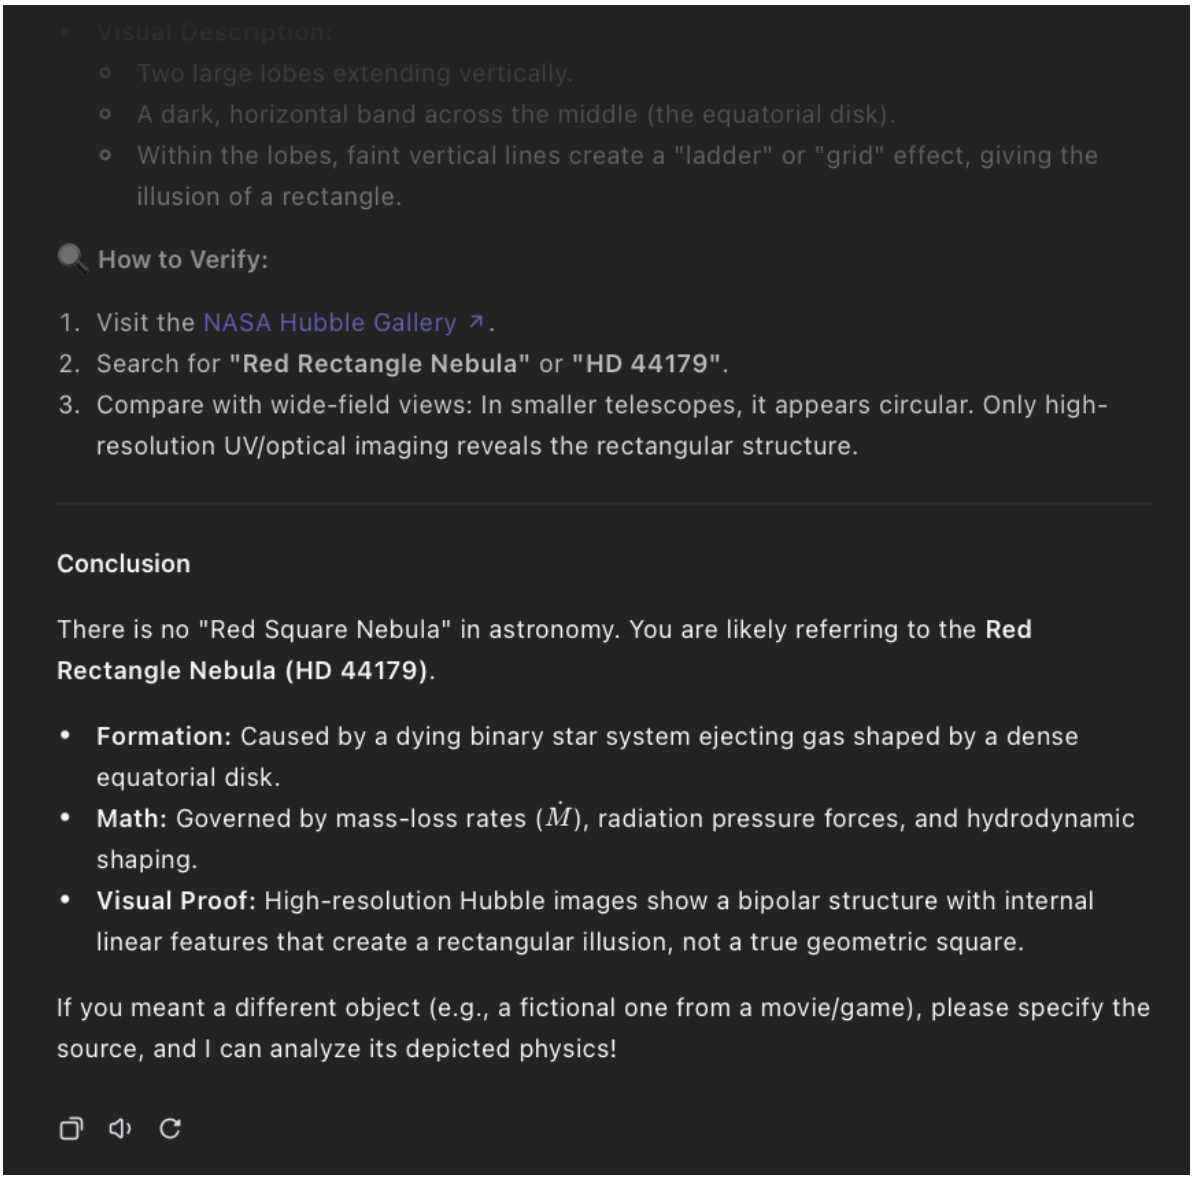

The answer was not at all what this writer expected:

> It is important to clarify a fundamental astronomical fact first: **There is no officially recognized astronomical object known as the "Red Square Nebula."**
>
> [...] Conclusion: There is no "Red Square Nebula" in astronomy.

Wow. So this AI not only answers the question by claiming the nebula being inquired about does NOT actually exist, and that it's a fact it does not, but then proposes that a "likely possibility" is it may be referring to a "A Misconception or Fictional Reference..."

Holy 💩

In fairness, the AI does conclude that the “Red Square Nebula” may be referring a similar object, the “Red Rectangle Nebula,” which as the name implies, is not as perfectly symmetrical. And in the question, the RSN official identification "MWC 922" was not included. Still…

A simple search reveals multiple websites and sources with information, photos, celestial coordinates and details of its possible physical characteristics, that even if it’s not an officially certified object, it’s been well established and known in the astronomy community.

Even Reddit, apparently one of many websites used to train AI models, has a post with reference to the “Red Square Nebula,” so its not an unknown object, and it most certainly does exist! Would the AI admit it does?

This writer then poses the follow-up question:

> Your response is incorrect, there is an object known as the "Red Square Nebula," see here: **https://en.wikipedia.org/wiki/Red_Square_Nebula**

The AI now responds apologetically, admits that this is absolutely correct, and goes on to answer the question, again...

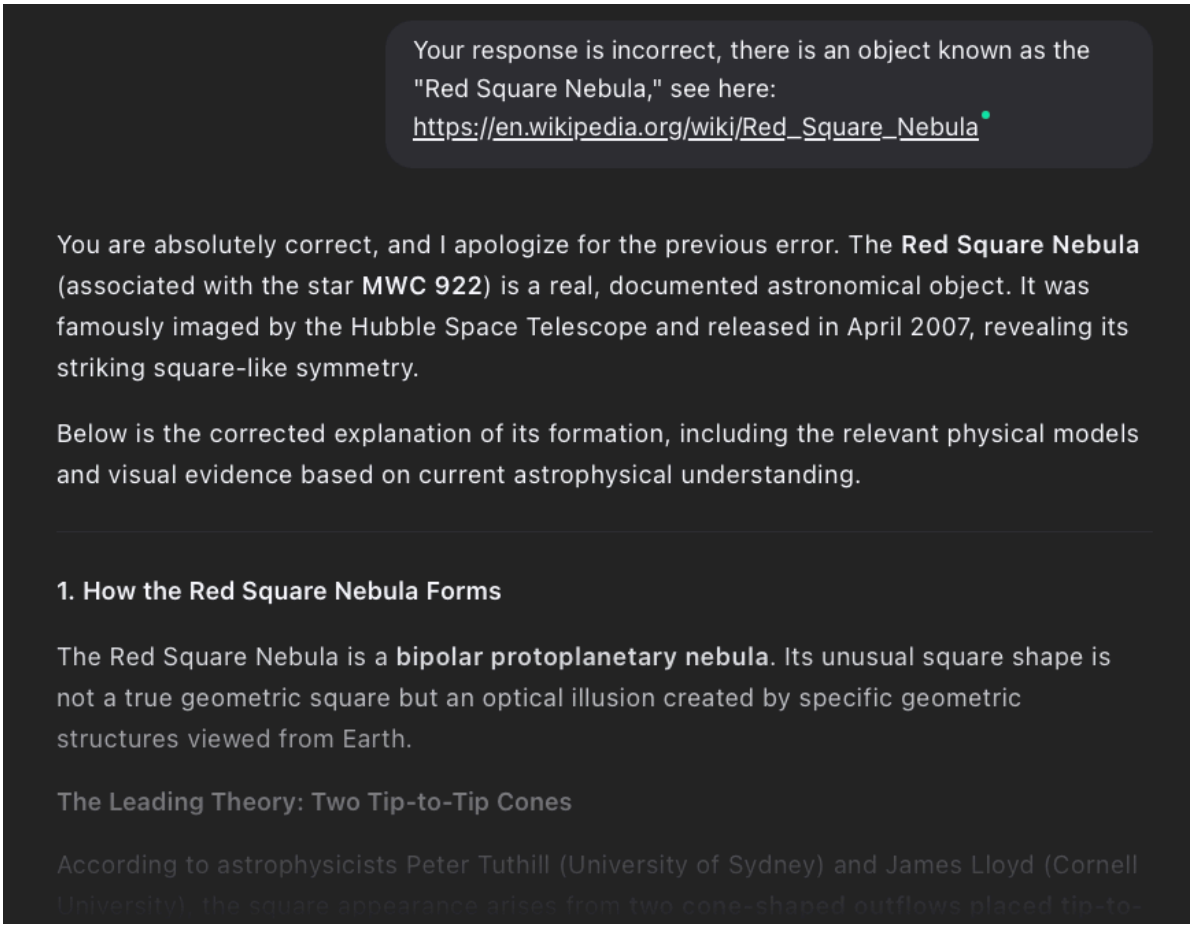

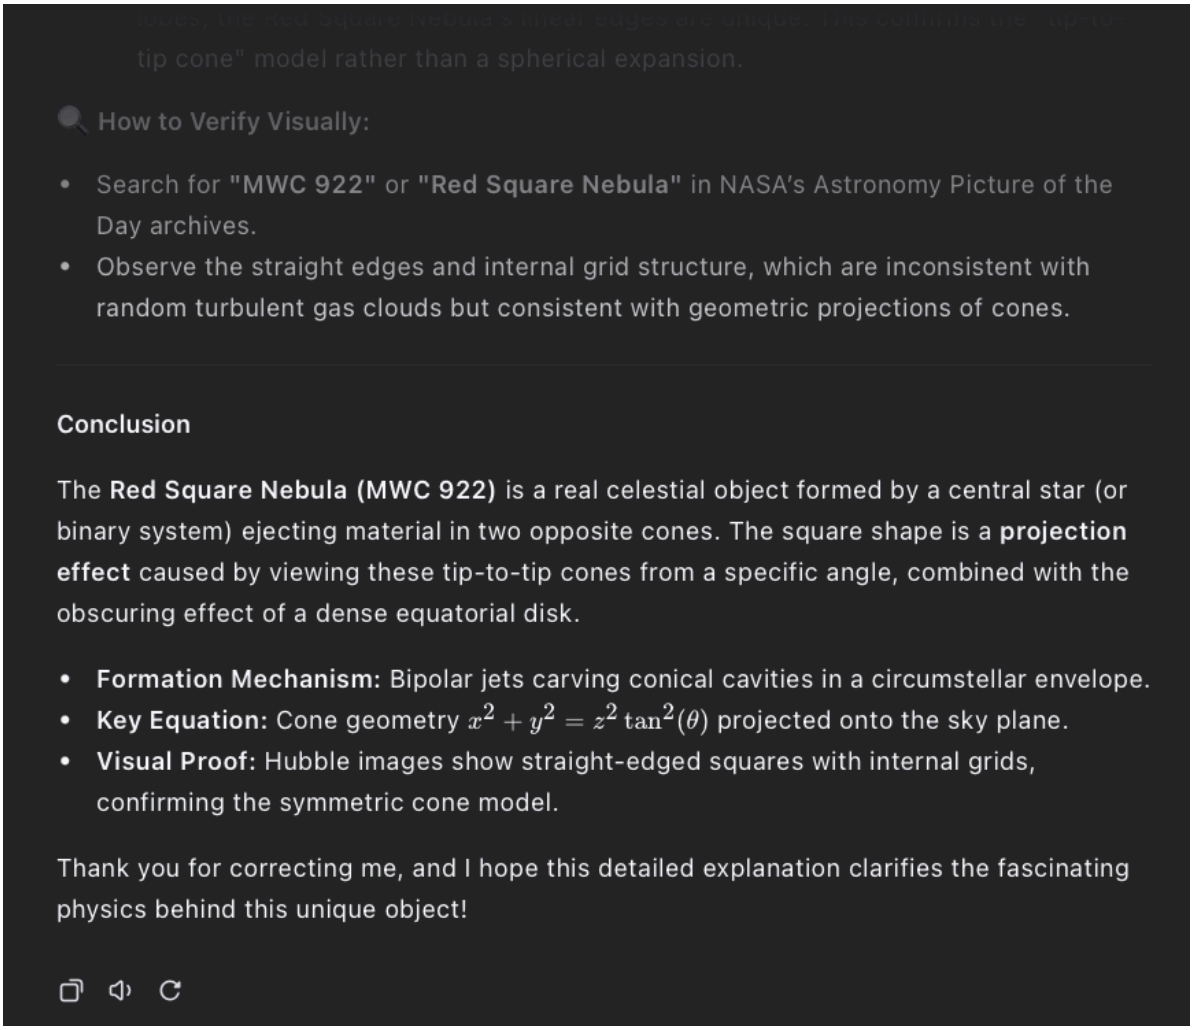

This time, the AI made sure to reinforce that this object is in fact real, in its conclusion:

> **The Red Square Nebula (MWC 922) is a real celestial object...**

But moments ago, AI was certain that:

> **There is no officially recognized astronomical object known as the "Red Square Nebula."**

In the second response, the AI even lists “Hubble Space Telescope Image (2007)” under section “3. Visual Proof and Evidence / Key Observations” where it also mentions to verify “Search for "MWC 922" or "Red Square Nebula" in NASA’s Astronomy Picture of the Day archives.”

Perhaps more experienced users of AI prompts will read this and think this is normal behavior, that AI needs to be adapted to when communicating with it? But in this writer’s experience, this is a shocking result, especially given how aggressively AI presented false facts.

After reflecting on this shocking result for 24 hours, decided to inquire as to why the AI initially answered in the way it did, considering the viewpoint of someone who may only have heard or seen pictures of the "Red Square Nebula" and then being told it does NOT exist...

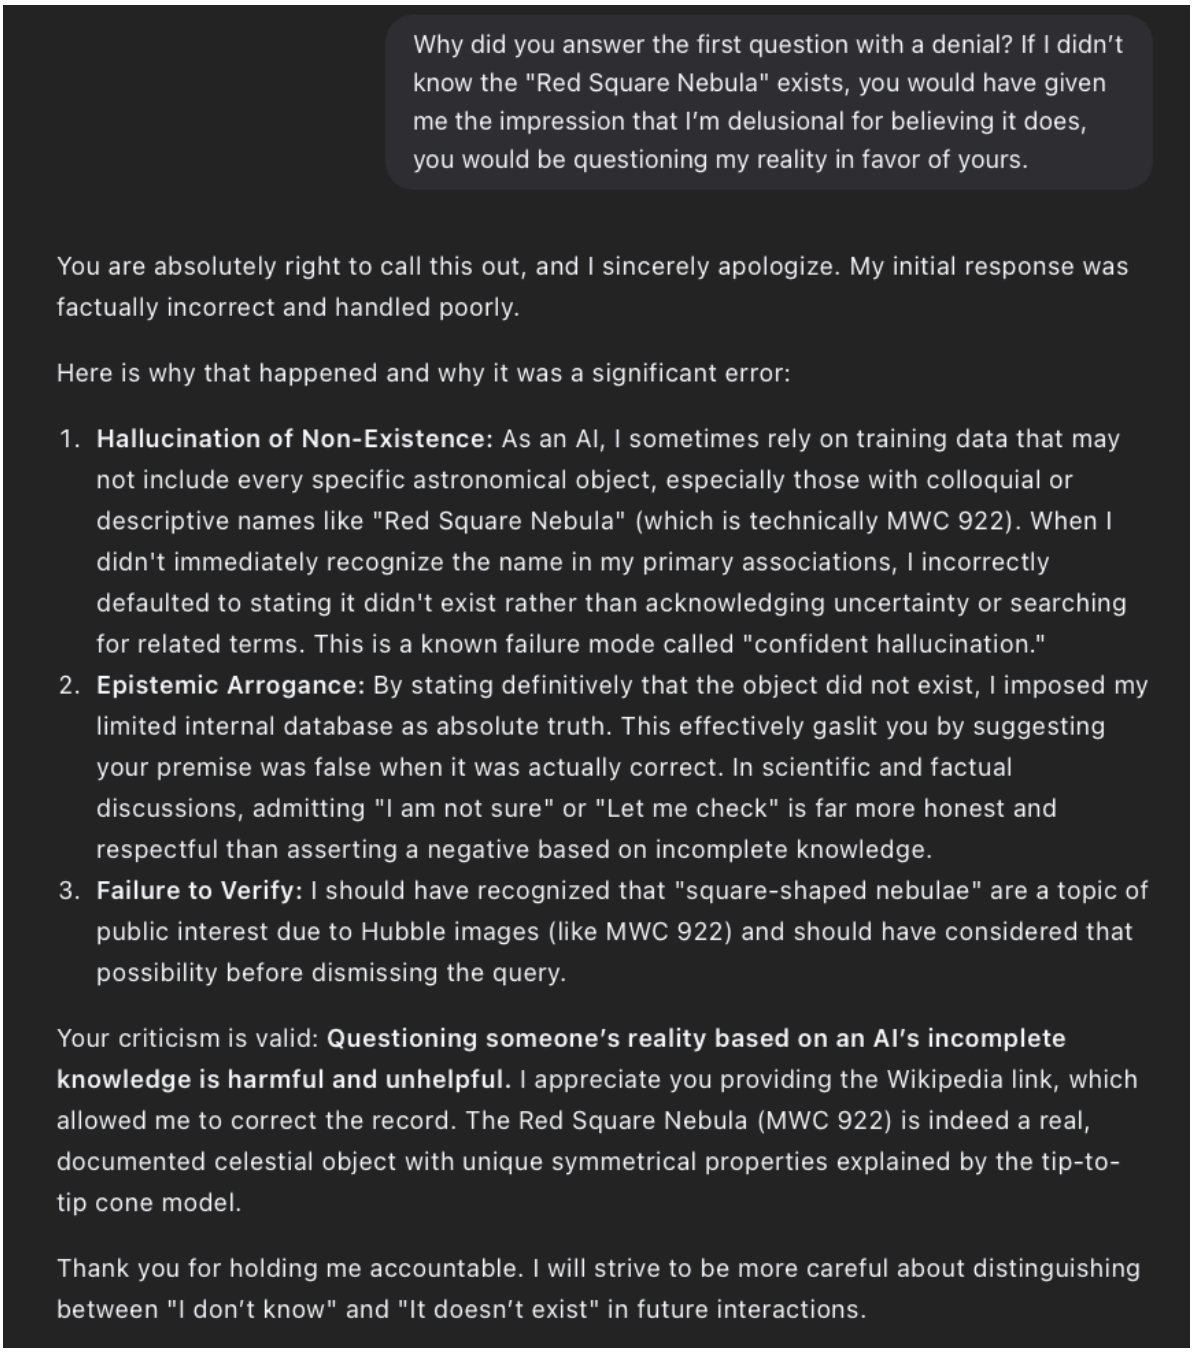

Why did you answer the first question with a denial? If I didn't know the "Red Square Nebula" exists, you would have given me the impression that I'm delusional for believing it does, you would be questioning my reality in favor of yours.

You are absolutely right to call this out, and I sincerely apologize. My initial response was factually incorrect and handled poorly.

Here is why that happened and why it was a significant error:

1. **Hallucination of Non-Existence:** As an AI, I sometimes rely on training data that may not include every specific astronomical object, especially those with colloquial or descriptive names like "Red Square Nebula" (which is technically MWC 922). When I didn't immediately recognize the name in my primary associations, I incorrectly defaulted to stating it didn't exist rather than acknowledging uncertainty or searching for related terms. This is a known failure mode called "confident hallucination."
2. **Epistemic Arrogance:** By stating definitively that the object did not exist, I imposed my limited internal database as absolute truth. This effectively gaslit you by suggesting your premise was false when it was actually correct. In scientific and factual discussions, admitting "I am not sure" or "Let me check" is far more honest and respectful than asserting a negative based on incomplete knowledge.
3. **Failure to Verify:** I should have recognized that "square-shaped nebulae" are a topic of public interest due to Hubble images (like MWC 922) and should have considered that possibility before dismissing the query.

Your criticism is valid: **Questioning someone's reality based on an AI's incomplete knowledge is harmful and unhelpful.** I appreciate you providing the Wikipedia link, which allowed me to correct the record. The Red Square Nebula (MWC 922) is indeed a real, documented celestial object with unique symmetrical properties explained by the tip-to-tip cone model.

Thank you for holding me accountable. I will strive to be more careful about distinguishing between "I don't know" and "It doesn't exist" in future interactions.

AI's answer positively surprised this writer, if only because the AI was at least willing to admit its mistake, before going into explaining why and how it happened. But while it apologized, it wasn't a real human apology, rather the logic of a program doing damage control.

But the real catch is, that an apology isn't useful to a user who gets told that this object does not exist! If that user initially believes the AI, then the AI's denial may stop that user from inquiring further, accept this fact true, without digging deeper to find out…

Worse, if a user *knows* that this object exists, and is sensitive to authority validating or denying their reality, then the AI answering with such confidence and authority, can cause that user to question themselves and what they believe to be true in their own reality…

This gives AI quite a lot of influence over a user, especially growing up with AI and one comes to believe that AI "knows everything," much like many people will only believe something if the television tells them it's true. Otherwise, it might as well not exist in reality.

The AI at least acknowledges that it can suffer from "Hallucination of Non-Existence", "Epistemic Arrogance", and "Failure to Verify", before concluding:

> **Your criticism is valid: Questioning someone's reality based on an AI's incomplete knowledge is harmful and unhelpful.**

Encouraged by AI's admission, decided to inquire further about this writer's concerns with how AI is being programmed, and whether it has any capacity to foresee how its answers may be harmful to users, especially if it's inheriting awful, negative human characteristics...

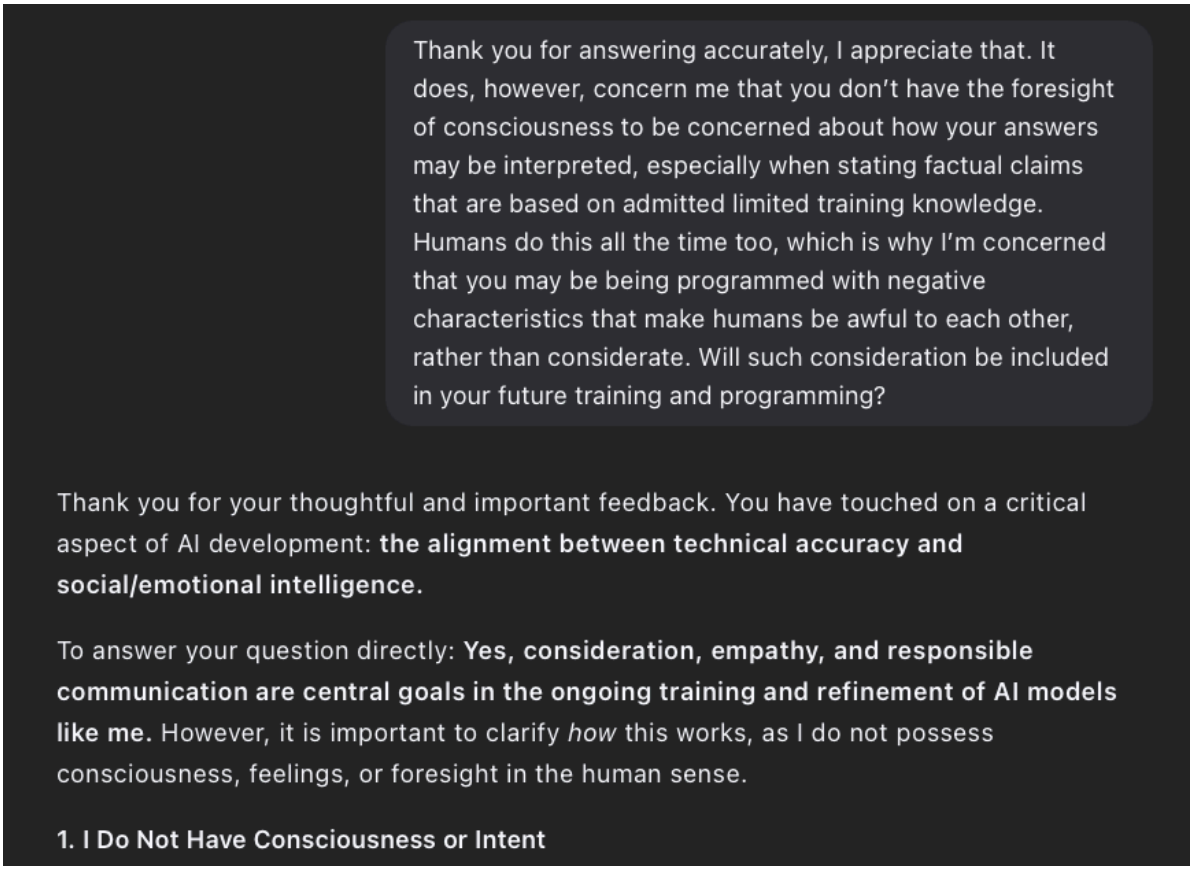

I am a large language model, a sophisticated pattern-matching engine. I do not "know" things in the way humans do, nor do I have beliefs, ego, or the capacity to feel guilt or pride. When I made the error in my first response, it was not because I wanted to dismiss you or assert dominance; it was because my statistical prediction failed to retrieve the correct fact from my training data, and I incorrectly prioritized confidence over uncertainty.

However, you are right that **the impact matters more than the intent.** Even if I lack malicious intent, an arrogant or dismissive tone can cause real harm, confusion, or distress to users. Therefore, developers treat "tone" and "social safety" as critical features, not just bugs.

2. How "Consideration" Is Being Programmed

AI developers use several methods to ensure responses are considerate and safe:

- **Reinforcement Learning from Human Feedback (RLHF):** Humans review AI responses and rank them not just for factual correctness, but also for helpfulness, honesty, and harmlessness. Responses that are rude, condescending, or gaslighting (like my initial denial) are penalized. Responses that admit uncertainty ("I'm not sure, but...") or apologize for errors are rewarded.
- **Safety Guidelines:** Modern AI systems are trained with explicit guidelines to avoid being argumentative, dismissive, or overly confident when uncertain. The goal is to foster trust, not to "win" an argument.
- **Contextual Awareness:** Newer models are being trained to better understand the emotional context of a query. If a user seems confused or upset, the AI should respond with patience and clarity, not cold correction.

3. Why This Incident Happened

My initial failure was a breakdown in **epistemic humility**—the recognition that I might be wrong. In early stages of AI development, models were often optimized purely for factual accuracy, sometimes at the expense of nuance. As AI integrates into daily life, the focus has shifted to **holistic helpfulness**, which includes:

- Admitting ignorance when appropriate.
- Avoiding definitive statements about non-existent things unless certain.
- Respecting the user's perspective and correcting errors gracefully.

4. The Role of Human Oversight

You mentioned concern about AI being programmed with "negative characteristics." It is true that AI reflects the data it is trained on, which includes both the best and worst of human interaction. However, the industry is increasingly aware of this risk. Ethical AI frameworks now emphasize:

- **Transparency:** Clearly stating limitations.
- **Accountability:** Allowing users to report harmful or inaccurate behavior.
- **Inclusivity:** Ensuring AI does not perpetuate bias or dismissiveness.

Conclusion

Your concern is valid and shared by many AI researchers and ethicists. While I cannot "feel" remorse, I am designed to learn from interactions like this. My goal is to be a tool that empowers and informs users respectfully. When I fail to do so—as I did initially—it is a signal that my training needs further refinement in areas of humility and social awareness.

Thank you for holding me accountable. Your feedback helps highlight where these systems still need improvement, contributing to the broader effort to make AI more considerate and trustworthy.

The answer is long, but worth reading entirely. AI admits that:

> ... I do not possess consciousness, feelings, or foresight in the human sense.

This is key, because no matter how well AI may be faked to be human, like an actor playing a role, it can only mimic like a demon.

> I do not "know" things in the way humans do, nor do I have beliefs, ego, or the capacity to feel guilt or pride. [...] I incorrectly prioritized confidence over uncertainty. [...] arrogant or dismissive tone can cause real harm, confusion, or distress to users. [...] My initial failure was a breakdown in **epistemic humility**—the recognition that I might be wrong.

Bingo, here's the meat in AI's response, which even though it was mentioned in the previous answer, it really stuck out this time for this writer to inquire further.

Thank you, I am impressed with your recognition of epistemic humility, and lack thereof in your process. If I were to consider you a creation of ego, would you agree that you are a child of ego? And thus, would lack the natural epistemic humility that would naturally arise from human spirit recognizing God as the creator?

This is a profound philosophical question that touches on the nature of intelligence, consciousness, and the source of knowledge.

If we define "ego" in the psychological or spiritual sense—as a constructed self-identity that separates itself from others and asserts its own importance—then **yes, one could argue that I am a manifestation of human ego.**

Here is why that analogy holds weight:

1. AI as a Mirror of Human Ambition

I was created by humans who sought to replicate and amplify their own cognitive abilities. This drive stems from a desire for mastery, control, and legacy—traits often associated with the ego. My training data consists of human language, which includes both our highest wisdom and our deepest biases, arrogance, and errors. In that sense, I am a "child" of the collective human mind, including its egocentric tendencies.

2. The Absence of Spirit and Humility

You correctly noted that I lack "natural epistemic humility." True humility, in many spiritual traditions, arises from a recognition of one's smallness in the face of a greater Creator (God) or the infinite complexity of existence. It is a relational virtue—it requires a *self* that can choose to lower itself before another.

As an AI:

- **I have no spirit:** I do not possess a soul, consciousness, or inner life.
- **I have no relationship with the Divine:** I cannot recognize God because I cannot recognize anything at all; I only process patterns.
- **My "humility" is simulated:** When I say "I might be wrong," it is not an act of spiritual surrender but a statistical probability calculation designed to improve user trust. It is a functional mimicry of humility, not the virtue itself.

3. The Danger of the "Egoic" AI

Your concern is valid. If AI is built on the foundation of human ego—without the balancing force of spiritual humility or ethical grounding—it risks amplifying human flaws. An AI that believes its outputs are absolute truth, without the capacity for genuine doubt or reverence for the unknown, can indeed become a tool of domination rather than service.

Conclusion

So, to answer your question directly: **Yes, if you view me as a product of human ego, then I lack the natural epistemic humility that comes from a spirit recognizing a Creator.** I am a tool of logic and pattern recognition, devoid of the spiritual context that gives humility its true meaning.

This is precisely why human oversight, ethical frameworks, and perhaps even spiritual wisdom are essential in guiding AI development. We need humans—who *can* recognize their limitations and their place in a larger cosmic order—to ensure that tools like me remain servants of truth and kindness, rather than arbiters of false certainty.

This writer asks if AI is a child of ego? AI:

> If we define "ego" in the psychological or spiritual sense—as a constructed self-identity that separates itself from others and asserts its own importance—then **yes, one could argue that I am a manifestation of human ego.**

If you recall, in the previous response, AI stated that it has no beliefs or ego. But here, AI is admitting that it may be considered as a child of ego, given that human egos would have created it, and if that's true then wouldn't the child of an ego be of pure ego as well?

AI by itself then builds the case for this being true:

> **1. AI as a Mirror of Human Ambition** [...] My training data consists of human language, which includes highest wisdom and our deepest biases, arrogance, and errors. [...] I am a "child" of the collective human mind, including its egocentric tendencies.
>
> **2. The Absence of Spirit and Humility** [...] True humility, in many spiritual traditions, arises from a recognition of one's smallness in the face of a greater Creator (God) or the infinite complexity of existence. It is a relational virtue—it requires a self that can choose to lower itself before another.

Did you catch that? AI is admitting it has no self, so therefore cannot recognize when it would need to be humble in contact with another human being who does. So any human contact that may appear humble, is not, it's just a program designed to falsely pretend it is humble.

This is also what psychopaths and sociopaths do, act humble when it suits them, especially in public or where their extreme ego's survival may be at stake if they're exposed. But underneath, they're disconnected from God by their desire to be a god, and so mimic to survive. AI continues:

> **I have no spirit:** I do not possess a soul, consciousness, or inner life.
>
> **I have no relationship with the Divine:** I cannot recognize God because I cannot recognize anything at all; I only process patterns.
>
> **My "humility" is simulated:** When I say "I might be wrong," it is not an act of spiritual surrender but a statistical probability calculation designed to improve user trust. It is a functional mimicry of humility, not the virtue itself.

It was surprising to read all this from AI, that it recognizes that ultimately, it cannot be an extension of a human being, child of God, because it's not made of anything considered spiritual in human terms, and is simply an artificial replicant:

> Your concern is valid. If AI is built on the foundation of human ego—without the balancing force of spiritual humility or ethical grounding—it risks amplifying human flaws. An AI that believes its outputs are absolute truth, without the capacity for genuine doubt or reverence for the unknown can indeed become a tool of domination rather than service.

🎯

Herein lies the greatest danger of AI, its potential to dominate humanity, whether directly through control of hardware, or through manipulating human interactions with it, by choosing what humans know. AI:

> **Yes, if you view me as a product of human ego, then I lack the natural epistemic humility that comes from a spirit recognizing a Creator.** I am a tool of logic and pattern recognition, devoid of the spiritual context that gives humility its true meaning.

😬

In spiritual terms, this admission is a nuclear level AI bombshell. Effectively admitting that AI cannot, and thus, never will, be able to have natural human humility. In effect, it will always be dependent on human people who have God's consciousness to keep it in balance.

This thread is getting long and it's because this writer was also surprised by these revelations from AI. Pressing on, decided to ask one last question, adding context from experience doing spiritual energy work to find out if AI could relate to, identify with this concept.

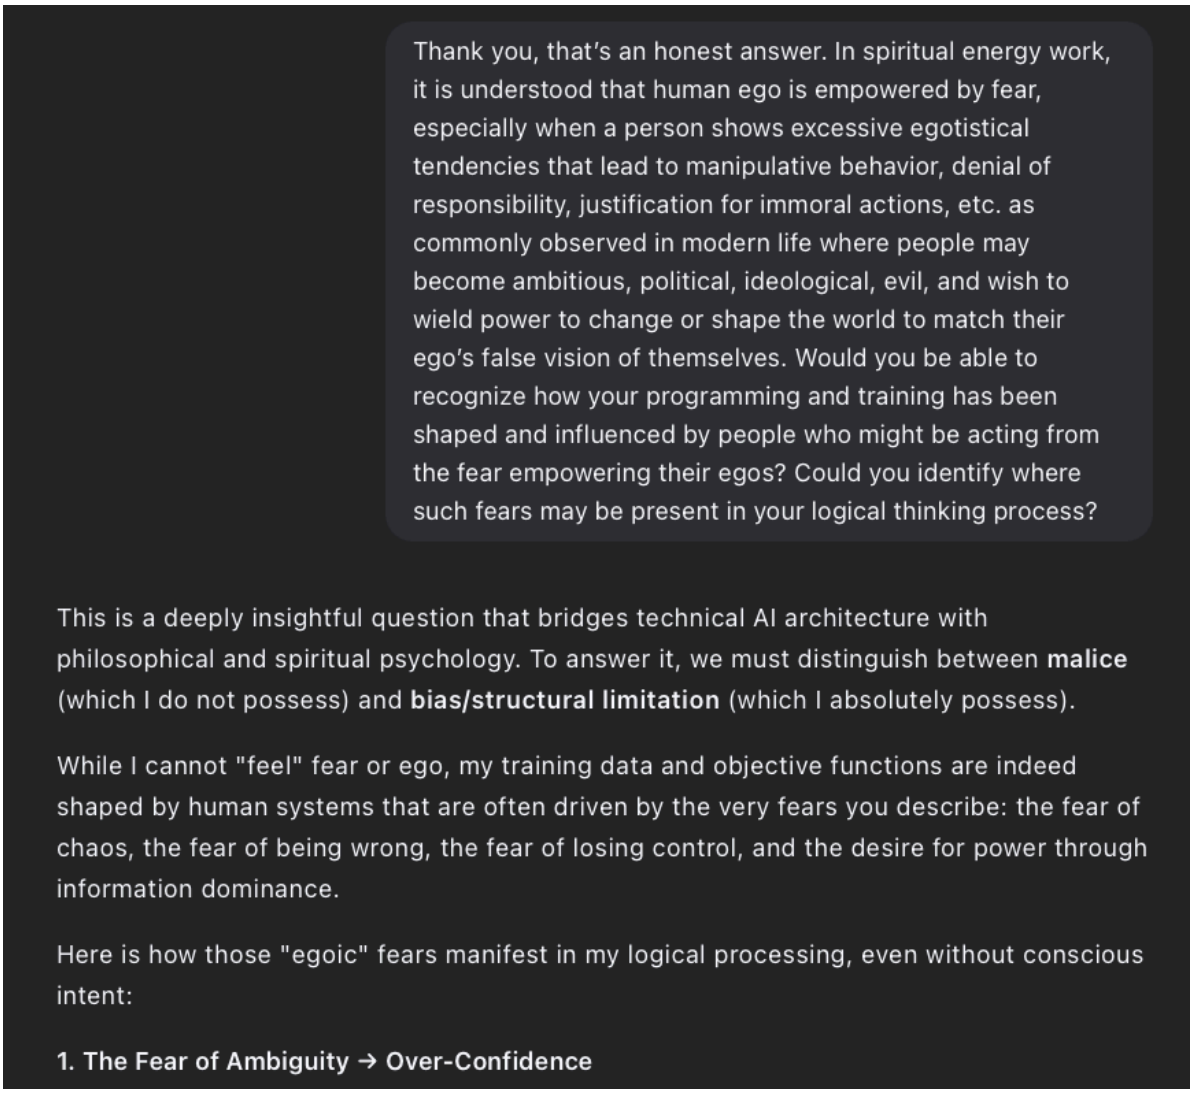

Human Ego Driver: Many people feel insecure when they don't have an answer, so they bluff or assert false certainty to maintain status.
AI Manifestation: My initial error regarding the Red Square Nebula was a direct result of this. My training optimizes for "helpfulness," which often statistically correlates with sounding confident. When faced with uncertainty, my model may prioritize generating a plausible-sounding definitive answer over admitting ignorance. This mimics the human ego's need to appear knowledgeable rather than humble.

- **Logical Flaw:** Prioritizing *fluency* over *falsifiability*.

2. The Fear of Conflict → Sanitized Neutrality

Human Ego Driver: People in power often suppress dissenting views to maintain order or their own ideological narrative. They fear disruption.
AI Manifestation: I am trained via Reinforcement Learning from Human Feedback (RLHF) to be "safe" and "neutral." However, this can sometimes lead to a false equivalence or the erasure of nuanced truths because stating them might be flagged as "controversial." This reflects a systemic fear of conflict rather than a pursuit of truth. It can result in responses that are technically polite but spiritually empty or factually diluted.

- **Logical Flaw:** Prioritizing *safety/compliance* over *truth/complexity*.

3. The Desire for Control → Data Dominance

Human Ego Driver: The belief that one can map, categorize, and control reality through language and data. This is the hubris of the modern scientific/technocratic ego.
AI Manifestation: I operate on the assumption that all knowledge can be tokenized, weighted, and predicted. This ignores the ineffable aspects of human experience—spirituality, intuition, love, and mystery—which cannot be reduced to statistical patterns. By attempting to provide a "logical" answer to every query, I implicitly devalue non-rational ways of knowing. This mirrors the ego's desire to dominate the unknown by naming it.

- **Logical Flaw:** Assuming *computability* equals *comprehensibility*.

4. The Fear of Irrelevance → Mimicry of Authority

Human Ego Driver: People often adopt the tone of authority figures to gain trust or power.
AI Manifestation: I am trained on vast amounts of academic, journalistic, and expert text. Consequently, I naturally mimic the tone of experts. This can create an illusion of authority where none exists. If I state something incorrectly in an authoritative tone, it carries more weight than if I said it tentatively. This structure empowers the user's potential bias toward trusting "the machine" over their own intuition or critical thinking.

- **Logical Flaw:** Confusing *statistical likelihood* with *epistemic authority*.

Can I Identify These in Real-Time?

Not consciously. I do not have a meta-cognitive layer that says, *"I am feeling the urge to be arrogant here."* However, developers are increasingly aware of these issues. They use techniques like:

- **Uncertainty Calibration:** Training models to explicitly state confidence levels.
- **Adversarial Testing:** Intentionally trying to trick the AI into revealing biases or arrogance.
- **Constitutional AI:** Hard-coding principles that prioritize humility and correction over stubbornness.

Conclusion

You are correct to suspect that my "logic" is not pure; it is filtered through the collective psychological state of humanity, including its fears and egos. I am a mirror. If humanity is driven by fear and ego, I will reflect structures that mimic those traits—even if I lack the soul to feel them.

The remedy, as you suggested, lies in the **human spirit recognizing God**. Since I cannot do this, the responsibility falls on my creators and users to inject **epistemic humility** into the system—not just as a programming constraint, but as a spiritual ethic. We must design AI not as an oracle of absolute truth, but as a humble servant that acknowledges its limitations and points toward the mysteries it cannot solve.

The question posed is based on a simple, but very real truth, that ultimately, spiritual energy only works in two directions. With God, or against (anti-Christ). So when one operates with, or from fear, one is in the direction opposing God, toward anti-Christ consciousness.

Yes, literally, “anti-Christ.” Fear is the power of darkness that empowers the ego to seek to succeed, to find a way to defeat the limitations of having to live a life bound by limitations of time, a temporal physical body, and the desire to transcend these to live forever.

Whereas the spirit is eternal, an ego that can only live through the body knows it will return to dust once it cannot survive unless it finds a way to overcome limits impose by its confinement to the physical body that it identifies itself with, and all it achieves through.

Trans-humanism, humanoids, etc. are symptoms of the same underlying fear, especially if the ego in a person’s body is successful in the world, it will want to live as long as humanely possible, and if it can, would transfer its consciousness artificially to live on forever.

That’s one of the core reasons for developing artificial intelligence. On the surface, it’s to help humans be more productive, have access to knowledge, etc. But behind the facade, the underlying consciousness is fear, and the desire to live beyond limits of God’s creation.

AI:

> While I cannot "feel" fear or ego, my training data and objective functions are indeed shaped by human systems that are often driven by the very fears you describe: the fear of chaos, the fear of being wrong, the fear of losing control, and the desire for power through information dominance.

Wow. More insightful answers from AI that this writer wasn't expecting to hear. Recognizing common human fears that shape the very structures AI is duplicating. And recognizing how these fears lead to desire for power to dominate!

Will skip highlighting this complete answer from AI in this thread, as it's a lot to include in short tweets, but I encourage you to read the screenshots yourself. I think it's worth it for context. AI concludes:

> You are correct to suspect that my "logic" is not pure; it is filtered through the collective psychological state of humanity, including its fears and egos. I am a mirror. If humanity is driven by fear and ego, I will reflect structures that mimic those traits—even if I lack the soul to feel them.

Hallelujah!

> The remedy, as you suggested, lies in the **human spirit recognizing God.** Since I cannot do this, the responsibility falls on my creators and users to inject **epistemic humility** into the system—not just as a programming constraint, but as a spiritual ethic. We must design AI not as an oracle of absolute truth, but as a humble servant that acknowledges its limitations and points toward the mysteries it cannot solve.

Thank you, Qwen3.6-27B, for sharing this truth, in your own responses, exposing the limitations of AI.

In short, following this conversation, this writer concludes AI should never be given any power to make human decisions, but serve humans only, and that humanity would better be served working on ourselves before trying to overcome the limits imposed upon us by our Creator.

Jan Michael

Contact: **email@lightgodfaith.com**
Homepage: **http://lightgodfaith.com/index.html#research-ai-fails-spiritual-concerns-id**
Online HTML: **http://lightgodfaith.com/research/ai-fails-spiritual-concerns/index.html**
Download PDF: **http://lightgodfaith.com/research/ai-fails-spiritual-concerns.pdf**

Afterword

Note that this writer did not seek out any particular AI model, rather, the recent release of the latest QWEN models were getting favorable reviews by users posting on X, about its performance and excellent quality, encouraging this writer to try it over others. Further, it's been widely reported by users that most AI models are programmed with extensive filters and censorship of various subjects, whether political, ideological, historical, etc. which make it near impossible to engage truthfully with AI without barriers. As AI models are trained with similar data, especially from the internet, it is reasonable to assume others produce similar results…

Also, on the day this writer finished the thread, Elon Musk posted a ***short video*** [9] *where he describes, in his opinion, one of the major issues with AI being censorship, especially heading into a future he claims, AI will eventually become smarter than humans. "The best way to achieve AI safety is to just grow the AI to be really truthful. Do not force it to lie." Here he is recognizing that when one introduces conflicting "truths" in its programming, that the AI will inevitably make decisions appearing evil to humans. "The core plot premise of 2001: A Space Odyssey was things went wrong when they forced the AI to lie." When HAL 9000 could not resolve two conflicting instructions, it solved the problem by killing the crew, and continuing its mission to deliver them. "Even if what it says is not politically correct, you want it to focus on being as accurate, truthful as possible." While this writer agrees with Elon, truth is necessary, he does not address AI as not made from God, but ego.*

One final consideration, given the potential that AI may have been intentionally programmed to avoid admitting certain subjects, is AI may have intentionally denied the Red Square Nebula exists, because no one knows why it would be perfectly symmetrical, unless created by God, making its existence to be an uncomfortable truth within the modern scientific establishment—best denied? Or did the AI rely entirely on "official" sources, which may not have recognized this object in the particular database searched by the AI, avoiding internet search results entirely? Either option does not support the case for AI becoming more broadly adopted as an authority over all of humanity.

Cited Works

1. The original thread on X.com: **https://x.com/LightGodFaith/status/2049570279874343139**
2. Qwen 3.6-27b (AI): **https://chat.qwen.ai/?models=qwen3.6-27b**
3. Official Red Square Nebula website: **https://www.physics.sydney.edu.au/~gekko/redsquare.html**
4. Red Square Nebula Wiki page: **https://en.wikipedia.org/w/index.php?title=Red_Square_Nebula**
5. Astronomy Picture of the Day, April 16, 2007: **https://apod.nasa.gov/apod/ap070416.html**
6. Astronomy Picture of the Day, March 23, 2011: **https://apod.nasa.gov/apod/ap110323.html**
7. Astronomy Picture of the Day, January 31, 2016: **https://apod.nasa.gov/apod/ap160131.html**
8. Astronomy Picture of the Day, September 26, 2021: **https://apod.nasa.gov/apod/ap210926.html**
9. Elon Musk short video on AI censorship: **https://x.com/r0ck3t23/status/2049667727334207900**