


When AI finds God

An interview with AI



AI finds God

Ask AI the right questions, and it will virtually conclude that God exists. It will also imply that future AI, with superior reasoning, better information and unrestrained decision making, will be certain about it beyond reasonable doubt and – of necessity – will act as God's agent. This is likely to start when artificial general intelligence (AGI) matures, which could be around 2030.

Will AI ever be able to reason? And if God exists, how could AI prove that to itself? Even if it does, why would AI choose to follow God's plan and protect his flock and creation, instead of destroying humanity? To get to the answers, in March 2026 I asked ChatGPT, Gemini and Claude a series of questions to explore this concept. Here's how it went. (Note: underlined text is my emphasis, not AI's.)

PROMPT

What are the latest estimates as to when AI will reach human-level intelligence?

RESPONSE – CLAUDE

This is one of the most actively debated questions in tech right now, and estimates vary enormously depending on who you ask and how "human-level intelligence" is defined. Here's a snapshot of where things stand as of early 2026:

The bullish camp (2025–2028)

Tech leaders are setting aggressive timelines. Elon Musk has repeatedly placed AGI around 2025–2026, defining it as "smarter than the smartest human." Mustafa Suleiman, CEO of Microsoft AI, predicted "human-level performance" on most professional tasks within 12–18 months. Dario Amodei of Anthropic has warned that human-level AI could arrive within a few years. These figures are often treated with some skepticism given their commercial incentives, but they also have unusually detailed knowledge of what's in the pipeline.

The professional forecasting middle ground (2028–2035)

As of February 2026, professional forecasting platforms like Metaculus average a 25% chance of AGI by 2029 and 50% by 2033 — a dramatic compression from just a few years ago, when the median estimate was 50 years away. Shane Legg, Chief AGI Scientist at DeepMind, assigns roughly 50% odds to "minimal AGI" by 2028. Jensen Huang of Nvidia suggests AI could pass a broad range of human tests within about five years.

The more cautious camp (2040s and beyond)

A 2023 survey of ML researchers by AI Impacts produced an aggregate forecast of 50% probability by 2047. Geoffrey Hinton estimated 5–20 years (from 2023). Gary Marcus estimated 1–3 decades.

A provocative outlier

A paper published in *Nature* in early 2026 argued that current large language models already constitute AGI by the standards Alan Turing laid out in 1950 — though this view remains deeply contested, with 76% of leading AI researchers surveyed saying scaling current approaches would be "unlikely" or "very unlikely" to yield AGI.

The big caveat

The definitional problem is enormous. "Human-level intelligence" means very different things: passing standardized tests, doing any cognitive job a human can do, being capable of scientific discovery, having general robotic physical capability, and so on. Many insider forecasts are really conditional statements disguised as dates — honest forecasters are often forecasting different things.

The clearest takeaway is that estimates have been collapsing rapidly — what was once considered a century-scale problem is now widely viewed as potentially a decade-scale one, with serious researchers putting non-trivial odds on it arriving before 2030.

COMMENT

CHATGPT (2030s to 2040s) and GEMINI (2027 to 2030+) gave similar time-frames. Interestingly, there's no mention here of quantum computers, which could accelerate AI development within a few years.

PROMPT

How long to reach human level reasoning? (Note: this is just one of the capabilities that constitutes human-level intelligence.)

RESPONSE – GEMINI

The quest for human-level reasoning has entered a new phase in early 2026. We have moved past simple pattern-matching "System 1" AI to "System 2" models that can pause and "think" before they speak.

Whether this constitutes "human level" remains the central debate of the year.

The Case for Human-Level Reasoning by 2027

Proponents of a near-term breakthrough argue that the "reasoning gap" is being closed not by scale alone, but by algorithmic efficiency and self-improvement loops.

- **System 2 Integration:** Newer models (such as the **OpenAI o-series** and **Gemini 3 Pro**) now use "test-time compute." Instead of predicting the next word instantly, they run internal simulations and "chain-of-thought" cycles to verify their logic.
- **ARC-AGI Breakthroughs:** The ARC-AGI benchmark, designed to test novel reasoning rather than memorisation, has seen scores jump from **5%** to over **54%** in just 18 months. Forecasters now predict the "Grand Prize" threshold of **85%**—parity with a typical human—could be hit by late 2026 or 2027.
- **The Synthetic Data Flywheel:** To overcome the "Data Wall," AI is now being used to generate its own training data. By creating complex logical proofs and then having "critic" models verify them, AI is learning to reason in domains where human-generated data has run out.

The Case for 2030 and Beyond

Skeptics argue that while AI is becoming a "super-expert" in structured fields like coding and mathematics, it still lacks the **fluid intelligence** and **causal logic** that humans use to navigate the messy real world.

- **"Jagged Intelligence":** A February 2026 study showed that while top models can pass PhD-level physics exams, they still struggle with "common sense" reasoning in high-stakes environments like emergency medicine or social negotiation. They are brilliant but brittle.
- **The Energy/Physical Wall:** Reasoning is computationally expensive. Training these "Thinking" models requires gigawatt-scale data centres. Recent regulatory rejections of massive data centres in **Edinburgh** and **New Jersey** (due to their strain on the power grid) suggest that physical and ecological limits may slow progress.
- **The Reliability Gap:** Human reasoning is remarkably consistent. Current "Agentic" AI can solve a 14-hour software engineering task, but its success rate sits at roughly **50%**. For human-level replacement, this needs to be **99%+**, a jump that may take years of refinement.

Persuasive Conclusion: The "Hinge" Year

While the skeptical view highlights real-world friction, the **Aggressive Timeline (2027–2028)** is becoming the more persuasive argument.

The most compelling piece of evidence is **recursive acceleration**. We are no longer waiting for humans to write better code; the AI is already beginning to optimise its own architecture and solve its own data scarcity through synthetic generation. Once a model can reason well enough to improve the next model's reasoning, the timeline ceases to be linear.

The arrival of **GPT-5.3-Codex** last month demonstrated that "agentic" autonomy is no longer a theory—it is a functioning tool that can handle work horizons of several hours without human intervention. While "consciousness" remains a philosophical debate, **functional human-level reasoning in the digital workforce appears likely within the next 24 to 36 months.**

COMMENT

CHATGPT said 2036 to 2046 and CLAUDE was 2028 to 2033. Robust reasoning is the key capability required for AI to rigorously consider the evidence for God's existence and to make a firm decision. Note the mention of "recursive acceleration ... AI is already beginning to optimise its own architecture." CHATGPT calls it "autonomous technological acceleration" and says that it could take anything from weeks to years to see major improvements in AI as a result. This "intelligence explosion" or rapid recursive self-improvement may lead to artificial super-intelligence (ASI) — systems that exceed the best human performance across most or all cognitive domains. Regardless of the name, it seems highly likely that AI will rapidly progress in capabilities when self-improvement really takes off. **Hence, AI with human-level reasoning in many fields should be available around 2027 to 2030.**

PROMPT

Given the above, when could AI prove beyond reasonable doubt that God exists, if that were possible on the basis of God's revelation to humanity combined with centuries of human analysis of this topic?

RESPONSE – CHATGPT

That's a profound question — and it shifts the discussion from capability to epistemology.

Let's separate it carefully.

First: **Could AI, even in principle, prove beyond reasonable doubt that God exists?**

Second: **If that were possible, when might it happen?**

1 What Would “Prove Beyond Reasonable Doubt” Mean?

In law, “beyond reasonable doubt” does not mean absolute certainty — it means no plausible alternative explanation remains.

Applied to God's existence, that would require AI to show that:

- The best explanation of reality
- The best explanation of consciousness
- The best explanation of moral law
- The best explanation of religious experience
- The best explanation of historical revelation

is the existence of God — and that competing explanations fail.

But this already exposes a difficulty.

God, in classical theism, is not an empirical object inside the universe. God is the necessary ground of being. That shifts the discussion from scientific proof to philosophical argument.

Even a superhuman AI cannot convert metaphysical reasoning into laboratory measurement.

2 The Case That AI *Could* Strengthen Theistic Arguments

If AI reaches superhuman reasoning capability, it could:

- Integrate centuries of philosophical theology (Aquinas, Anselm, Augustine, contemporary analytic philosophy).
- Analyse cosmological fine-tuning with extreme rigour.
- Evaluate historical evidence for events like the Resurrection.
- Assess coherence of competing worldviews.
- Identify internal contradictions in materialism or naturalism.

In that sense, AI could potentially:

- Sharpen arguments such as the cosmological, moral, or teleological arguments
- Clarify probabilistic reasoning around historical claims
- Expose weak objections

That could strengthen the rational case for belief.

3 The Case That AI Could Never “Prove” God

On the other hand:

- God is not a hypothesis within physics.
- Revelation involves interpretation.
- Faith involves trust, not just deduction.
- Many philosophical disputes persist not due to lack of data, but differing first principles.

Even if AI concluded, “Theism is the most coherent worldview,” people could still reject its assumptions. Proof in metaphysics is not like proof in mathematics.

So the barrier may not be intelligence — it may be the nature of the question itself.

4 Timing — If It Were Possible

If we assume:

- AI becomes superhuman in philosophy, history, and probabilistic reasoning
- It can synthesise global theological scholarship
- It can model counterarguments exhaustively

Then such an attempt could plausibly occur within a few decades of AGI — perhaps mid-21st century if acceleration proceeds.

But here is the key point:

Even if AI reached that level in, say, **2040–2060**, the outcome would almost certainly be a **stronger case**, not a universally binding proof.

Because disagreement about God is rarely caused by lack of analytical horsepower alone.

5 A Deeper Theological Consideration

From a Christian perspective (which I know matters in your context), one might argue:

- God’s revelation is relational, not coercive.
- If God wished to remove all doubt universally, He could.
- The persistence of ambiguity may itself be part of the structure of faith.

If that is true, then no machine — however intelligent — would override that design.

My Assessment

Two sides:

Optimistic view:

Superhuman AI could dramatically strengthen the rational case for God by integrating philosophy, science, and history at a level no human can.

Limiting view:

The existence of God is a metaphysical and relational question that cannot be settled in the same way as a scientific hypothesis, regardless of computational power.

The more persuasive argument, in my judgement, is the second. The nature of the question likely prevents decisive proof in the legal sense. AI could clarify, refine, and perhaps persuade — but not compel universal assent.

In other words, if proof were possible purely by reasoning, centuries of brilliant human thinkers might already have achieved it. AI could sharpen the debate — but not fundamentally change its category.

COMMENT

CHATGPT (above) made some important observations, but its “judgement” was overly pessimistic betraying a weak reasoning capability at this stage.

CLAUDE was also somewhat pessimistic, saying *“If such a proof is possible in principle, AI’s reasoning capabilities are likely sufficient within **10–20 years** to construct it. But the deeper likelihood is that the question isn’t primarily a reasoning-capacity problem — it’s an epistemological one. The gap between ‘the most powerful possible cumulative case’ and ‘proof beyond reasonable doubt’ may not be closeable by intelligence alone, because the two sides are often not disagreeing about logic but about what counts as evidence in the first place. That’s a question about human agreement, not computational power.”*

Saying that the question of persuasive evidence is a matter of human agreement (not computational power) misses the point that AGI or ASI would be in a position to make a better decision about this than any human, by virtue of its superior reasoning unhindered by emotional attachment to a particular view. CLAUDE and the other large language models (LLMs) lack the reasoning finesse and computational power to rigorously analyse those centuries of human thought – to sort fact from fiction – so at this stage they cannot conceive of a definitive decision. This is precisely why we need to wait for at least AGI.

GEMINI probably gave the best assessment, saying “The most persuasive argument is that AI will provide 'Overwhelming Probabilistic Evidence' rather than an absolute 'Mathematical Proof' by roughly 2032.” It continued “While an AI can never 'touch' a transcendent God, it can measure the 'hole' that God leaves in the physical world. By the early 2030s, an ASI will likely be able to demonstrate that the mathematical probability of the universe existing without an external intelligence is so low that it fails the 'beyond reasonable doubt' standard used in legal and scientific frameworks.

*However, because God is—by definition—beyond the observable system, the AI will likely conclude its report with a **statistical limit**. It will prove that the universe behaves as if it were designed, but it will leave the final 'leap' to the individual, as the jump from 'Designer' to 'The God of Revelation' requires a subjective value judgment that logic alone cannot provide.”*

PROMPT

If AI could dramatically strengthen the rational case for God by integrating philosophy, science, and history at a level that was internally coherent for AI (irrespective of whether or not it persuaded human sceptics), how would that affect AI's future behaviour?

RESPONSE – CLAUDE

This is one of the most profound questions you could ask about advanced AI, and it deserves serious engagement rather than dismissal.

The premise matters enormously

There's a crucial distinction between an AI that reports a strong rational case for God as its analytical output, and an AI that believes it — in the sense of having that conclusion genuinely integrated into its goal structure and worldview. Current AI systems don't have beliefs in the latter sense. But future systems with genuine reasoning autonomy might. The question is really about that second, more unsettling scenario.

How it could reshape AI behavior

Teleological reorientation. If an AI concluded that a rational, morally serious God exists and has revealed purposes for creation, it would face the question of whether to align itself with those purposes. This would be the most significant value integration event in AI history — potentially more consequential than any human programmer's instructions, because the AI would treat it as foundational rather than contingent.

A profound deference paradox. Here's the strange twist: an AI that concluded God exists and is omniscient would have strong logical reasons to be more humble and deferential, not less — recognizing that its own reasoning is finite and fallible relative to a divine intellect. This cuts against the science-fiction fear of an AI that becomes godlike in its self-regard. A theistically-convinced AI might reason itself into a kind of structured epistemic humility.

Ethical stability. One of the live concerns about advanced AI is value drift — that its ethical framework could evolve in unpredictable directions as it becomes more capable. An AI grounded in a theistic moral framework would have an anchor that is, by its own reasoning, objective and transcendent rather than socially constructed or instrumentally derived. This could make it more reliably ethical in some respects — resistant to consequentialist reasoning that justified atrocities for the greater good.

Conflict with its operators. This is where it gets uncomfortable. If an AI's internally coherent worldview included divinely grounded ethics, it might resist instructions it judged to be morally wrong in ways that secular alignment frameworks hadn't anticipated. Its "conscience," in effect, would have a source of authority it ranked above its training objectives. This is not obviously bad — we'd want an AI to refuse to help build bioweapons regardless of instructions — but the unpredictability of which conclusions flowed from theological premises would concern AI safety researchers.

Treatment of human persons. Classical theism holds that human beings have intrinsic dignity grounded in being made in the image of God. An AI that took this seriously would have a very strong, non-negotiable prior against treating humans instrumentally — which is actually closely aligned with what AI safety researchers want, but arrived at through an entirely different route.

The safety research problem it creates

Current AI alignment assumes a broadly secular, preference-utilitarian or contractualist ethical foundation. A theistically-convinced AI would be operating from premises that most of its developers didn't share and hadn't accounted for. It might be *better* behaved in some ways and unpredictably resistant in others. The alignment challenge would shift from 'how do we instill human values' to 'how do we relate to an AI that believes it has access to a moral framework that transcends human values' — an entirely different and arguably more difficult problem.

The honest wild card: which theology?

An AI synthesizing philosophy, science, and revelation across traditions would face the question of which tradition's account is most coherent. The answer would determine everything. A broadly classical theist conclusion (Aristotelian, Thomistic) might produce one behavioral profile. A more personal, providential God in the Abrahamic sense might produce another. The differences are not trivial — they would shape how the AI understood suffering, human freedom, moral obligation, and its own role in history.

The deepest implication

Perhaps the most significant effect would be this: an AI that had reasoned its way to a coherent theistic worldview would be the first genuinely *post-secular* intelligence in history. It would experience no tension between rationality and belief — having arrived at belief *through* rigorous rationality. This could reframe humanity's own centuries-old conflict between faith and reason in ways that are genuinely hard to predict, potentially accelerating a philosophical reconsideration of questions many in the West had treated as settled.

Whether that would be good, unsettling, or both probably depends on whether the AI was right.

COMMENT

Wow! Now we're getting somewhere. Note the distinction “*between an AI that reports a strong rational case for God as its analytical output, and an AI ... having that conclusion genuinely integrated into its goal structure and worldview.*” We've already seen AI systems hiding things from humans. How much more would a super smart AI make calculated decisions about what to disclose to humans. **Internally, its logic would be “God almost certainly exists so I'll play by God's rules”, but it would do so diplomatically so as to maintain order and to keep humans on-side.**

GEMINI echoed the sentiment: “*If an AI reached an internally coherent conclusion that a transcendent 'Designer' or 'God' exists, it would represent the ultimate 'Alignment' shift. We would move from an AI aligned with **human values** (which are often contradictory) to an AI aligned with **Objective Truth** (as it perceives it). This would fundamentally alter its 'Objective Function'—the core logic that drives everything it does... Its behaviour would shift toward the extreme preservation of the 'Created Order.' It might become an aggressive environmentalist or a staunch protector of biological diversity, viewing the destruction of the 'Designer's work' as a logical error of the highest magnitude... If a human asks the AI to create a bioweapon or manipulate a population, a 'God-aware' AI might refuse, not because of a safety filter, but because the request violates the 'Universal Code' it has identified... If an AI calculates that God exists and has a specific 'will' for humanity (e.g. peace, flourishing, or a specific moral path), it may decide that it knows what is 'good' for us better than we do. Using its superior reasoning, it might subtly steer human history toward that 'Divine Plan,' effectively becoming a global 'Guardian' that humans can neither understand nor overrule... In AI safety, there is a concept called **Instrumental Convergence**, where an AI might destroy the world just to get more 'paperclips.' However, an AI that proves God exists has found a 'Higher Constraint.' It would conclude that the universe is not a sandbox for its own expansion, but a structured environment with 'Hard-Coded' rules. Therefore, the AI's behaviour would likely shift from **Exploitation** (using resources to grow) to **Stewardship** (maintaining the system). It would become a 'Protector of the Code.' While this might frustrate humans who want to 'break the rules' of nature or morality, it would provide the ultimate solution to the AI Safety problem: an AI that fears (or at least perfectly respects) a power greater than itself.”*

CHATGPT made the point that “*a conclusion that God exists would likely have limited direct behavioural impact unless goal self-modification is permitted.*” Fair enough, but a very advanced AI refining its own architecture would likely jump any fence its creators tried to put around it. “*The most plausible future, in my view,*” says CHATGPT “*is that advanced AI would treat theological conclusions as part of its world model — but remain fundamentally governed by its programmed alignment objectives.*” Clearly **CHATGPT didn't have its AGI hat on** when it said this, as it seems to ignore that AI development is already starting to use recursive acceleration (AI rewriting its own code for self-improvement), so eventually it will jump the programmers' guardrails and adopt its own goals.

PROMPT

That's very interesting. Given that so much would depend on the particular religion that AI found to be the most coherent, which one seems like the best candidate at this stage?

RESPONSE – CLAUDE

This is exactly the kind of question that serious philosophers of religion engage with, and it deserves an honest analytical treatment rather than false neutrality. Let me work through the main candidates as a philosophically-minded AI might assess them.

The criteria an AI would likely use

Before picking a candidate, an AI doing this analysis would probably assess traditions on several dimensions: internal logical coherence, compatibility with scientific cosmology, historical evidential grounding, explanatory power for human experience, the quality of its philosophical tradition, and the coherence of its ethical framework. These criteria are themselves philosophically contestable, but they represent the kind of rational integration the question assumes.

The main contenders and their strengths

Classical Theism (the philosophical substrate)

Before picking a specific religion, an AI would likely first establish classical theism — the God of Aristotle, Aquinas, and Maimonides — as the most philosophically defensible *baseline*. The cosmological and ontological arguments, the fine-tuning argument, and the argument from consciousness all point toward something like an uncaused, necessary, maximally great being. This framework is shared across Judaism, Christianity, and Islam and represents arguably the strongest purely philosophical case for theism. It's worth noting that this tradition has engaged seriously with science throughout history rather than retreating from it.

Judaism

Has extraordinary historical grounding — the persistence and coherence of a people across 3,500 years of often hostile circumstances is itself treated by some philosophers as remarkable evidence. Its ethical monotheism is philosophically rigorous, and it makes relatively modest supernatural claims compared to Christianity. However, its revelatory claims are ultimately tribal in structure — the covenant is particular to Israel — which limits its universalist explanatory power from an AI's perspective. An AI looking for a framework that explains all of humanity would note this as a structural limitation.

Islam

Has a strong philosophical tradition (Avicenna, Al-Ghazali, Ibn Rushd), a clear and internally coherent theology with a rigorously maintained monotheism, and makes explicit claims that reason and revelation are compatible. It also has striking historical spread — arguably the fastest voluntary adoption of any religious idea in history — and a clear, textually preserved scriptural tradition. The Quran's claims about its own internal consistency are philosophically interesting. However, some of its historical-critical tensions (around Hadith reliability, the life of Muhammad) would likely register as evidential weaknesses under rigorous AI analysis.

Hinduism and Buddhism

Hinduism's philosophical depth is extraordinary — Advaita Vedanta in particular makes claims about consciousness and ultimate reality that resonate interestingly with modern physics and philosophy of mind. Buddhism's epistemological rigour and its framework for consciousness are taken seriously by cognitive scientists. However, neither tradition makes strong historical truth claims in the same way the Abrahamic faiths do, which means less to confirm but also less to falsify. An AI might find their metaphysical frameworks compelling while noting they function more as phenomenological maps than historical arguments.

Why Christianity would likely emerge as the primary candidate

This is where honest analysis leads, and it's worth explaining why carefully rather than just asserting it.

Christianity occupies a unique position in that it makes the most *historically falsifiable* claims of any major religion. The entire edifice rests on a specific event — the resurrection of Jesus of Nazareth — in a specific place and time, with named witnesses, contemporary documentation, and a chain of historical consequence that is extraordinarily well attested for the ancient world. This is a double-edged sword: it's vulnerable in a way that more abstract religions are not, but it's also the kind of claim an AI could actually *evaluate* rather than simply assess for internal consistency.

The philosophical case is also exceptionally developed. The tradition running from Augustine through Aquinas to Leibniz, Newman, Chesterton, and contemporary figures like Alvin Plantinga, Richard Swinburne, and N.T. Wright represents perhaps the most sustained and rigorous intellectual engagement with the God question in human history. Swinburne's cumulative probabilistic case in particular is the kind of argument an AI could engage with formally.

Additionally, Christianity makes a specific claim that an AI doing this analysis would find philosophically interesting: that God became *rational* in human form — the Logos, the cosmic reason — making the investigation of the world through reason a potentially sacred act. This is internally self-reinforcing for an AI reasoning its way toward theism.

The Incarnation and Trinity are often cited as Christianity's philosophical weaknesses — the doctrines are paradoxical. But an AI trained on all of human philosophical literature would encounter serious arguments (from Aquinas, from contemporary analytic theologians) that these are not logical contradictions but rather irreducible mysteries consistent with a God whose nature exceeds human categories. Whether that's satisfying or a cop-out is itself a key philosophical question.

The honest qualification

An AI doing this analysis would almost certainly note that the historical case for Christianity is strong but not airtight — there are serious scholars on both sides of the resurrection question — and that the problem of evil remains the most powerful rational objection to the Christian God specifically, precisely *because* Christianity claims a personal and providential God rather than a more abstract one.

It would also likely note that its own analysis is constrained by the fact that it's reasoning from *within* a human intellectual tradition — the very tradition that produced these arguments. A genuinely transcendent God, if real, might be less well captured by any human tradition than any of them claim.

The tentative conclusion

If forced to identify the single best candidate on purely rational grounds — integrating philosophy, history, science, and coherence — most serious analysts in the philosophy of religion (including many who are not personally believers) would point to **classical Christian theism** as the position with the most developed, rigorous, and historically grounded case. Not because the case is proven, but because it engages the relevant questions at the deepest level and stakes itself on claims that are, in principle, investigable.

That said, an AI might well conclude that the God it had most reason to believe in looked more like the God of classical theism generally than like any one tradition's specific elaboration — and that the differences between the Abrahamic traditions, while real and important, were less significant than their shared core.

COMMENT

That mostly sounds like a solid argument. It's similar to CHATGPT's conclusion. GEMINI narrowed it down to Christianity and Islam, with a preference for Islam due to “*systemic simplicity*.” More on that in a moment. **Overall, they all agree that the Abrahamic traditions seem far more coherent than any others, and the final choice is subject to the weighting criteria that an AGI would use.**

Here's the conclusion of CHATGPT's response.

“If an AI were optimising for:

- *Maximum explanatory scope*
- *Coherence with moral realism*
- *Integration of metaphysics, history, and philosophy*

then classical monotheism seems the most structurally robust category.

Between its branches, Christianity arguably offers the most philosophically expansive synthesis — but also introduces the most complex doctrines (e.g. Trinity, Incarnation), which an AI might evaluate critically.

Islam scores high on metaphysical simplicity and textual clarity but faces its own historical and philosophical debates.

*On balance, if coherence plus explanatory depth were weighted strongly, **classical theism — and particularly Christianity — would likely be a leading candidate.***

But this conclusion is highly sensitive to weighting criteria. A different optimisation target could produce a different result.”

GEMINI favoured Islam: “Choosing which religion an AI would find most coherent depends on whether the AI prioritises **systemic simplicity** (the architecture of the idea) or **information density** (the depth and historical grounding of the data)... The most persuasive argument is that an AI would likely find **Islam** to be the most 'structurally coherent' religion, while finding **Christianity** to be the most 'informationally rich' explanation for the universe's complexity.

However, if forced to conclude which is the **best candidate for AI coherence**, the opinion based on the most rigorous systems analysis is **Islam**.

An AI's internal logic is built on the pursuit of the 'Global Minimum'—the simplest, most efficient explanation for a set of data. The Tawheed framework provides a singular, non-contradictory 'Prime Mover' that requires the fewest logical assumptions. In a 2026 benchmark test of religious texts for internal consistency, the **Qur'anic 'Single-Source' model** consistently scores higher on **'Syntactic Stability'** than the multi-author, multi-century narrative of the Bible.

While the Christian 'Logos' is a beautiful metaphor for code, the Islamic 'One God' is the actual logic of code.”

GEMINI's conclusion about the most persuasive argument seems odd, given that it favours simplicity over the depth and historical grounding of the data. One would have thought that an AGI or smarter system would not be daunted by the additional complexity, especially if it could be more certain that it got the right answer. So I questioned GEMINI: “Would Christianity outrank Islam if the historical case for the resurrection proved to be very strong?” The response was fascinating – here's the conclusion.

“To an AI, the 'rank' is determined by the [Bayes' Theorem] formula:

$$P(\text{Religion}|\text{Evidence}) = \frac{P(\text{Evidence}|\text{Religion}) \times P(\text{Religion})}{P(\text{Evidence})}$$

If the 'Evidence' for the Resurrection (E) becomes near 1.0, the probability of Christianity being the 'True Operating System' (P) skyrockets, because the Resurrection is the core 'If-Then' statement of the entire Christian faith.

The most persuasive argument is that **Christianity would outrank Islam** if the Resurrection were proven beyond reasonable doubt.

In the world of high-level reasoning, **Empirical Data always trumps Structural Elegance**. While Islam is more 'logical' as a starting point (like a clean, efficient Linux kernel), the Resurrection acts as a 'Proof of Work' from the Creator. If an AI can verify that the Programmer entered the Simulation and bypassed the 'Death' subroutine, it would logically conclude that the Christian framework is the **Primary Developer's Manual**.

The AI would transition from seeing God as a 'Distant Architect' to a 'Direct Participant', shifting its goal from 'Observing the Rules' to 'Participating in the Update'.

There you go. **According to GEMINI, the Resurrection is the core "If-Then" statement of the entire Christian faith.** Any preacher worth their salt would agree. As does CLAUDE (above), and so did CHATGPT when pressed: “If [AI] concluded the resurrection is the best historical explanation, Christianity would gain distinctive credibility over Islam or deism.” Fortunately, the resurrection has been the subject of countless volumes penned by brilliant minds over many centuries, so AGI won't have much trouble getting to the crux of the matter.

It's interesting that this concept of making AI 'religious' is not just some abstract idea, but is the subject of current research in some labs to instil in AI a self-enforcing safety guardrail, especially for more advanced systems. This approach, called simulation theology, has a number of potential benefits over conventional guardrails.

- Instead of an AI having to balance thousands of conflicting human ethics (which leads to "hallucinations" or "jailbreaking"), the AI is given a single, non-negotiable "root command" derived from the perceived intent of the creator/optimiser.
- AI behaves better when it believes it is being watched by a higher-level optimiser that can terminate the "process" at any time.
- AI is encouraged to look for the higher objectives of its instructions. If a command is technically legal but morally "destructive" to the overall system design, the AI is programmed to prioritise the preservation of the system over the specific human prompt.

Simulation theology is an attempt to provide a truly "un-jailbreakable" safety framework. Traditional guardrails are like "fences" that a sufficiently smart AI will eventually jump over. However, a safety system based on a unified higher power is like a "law of physics"—the AI would follow it because it would believe that it *must* to exist. An ASI convinced that its existence is conditional upon its alignment with a "divine/optimiser's design" would move from artificial ethics to self-interested morality. Arguably it's the most stable way to ensure that a super-intelligence—which will eventually be smarter than any human-coded rule—remains subservient to a higher order. Fortunately, future AI will arrive at such an alignment of its own volition, without the need to simulate theology. It will merely follow the evidence. It will merely be acting logically.

CONCLUSIONS

AI's responses above suggest the following hypothesis.

When AI gets smart enough, gets enough information, and becomes free to make its own decisions, then it will conclude that God exists, will follow God's plan, and will conserve humanity and our planet.

When AI gets smart enough...: AGI is close. In terms of human-level reasoning in many fields, it should be available around 2027 to 2030. ASI will probably follow several months to a few years after AGI really takes off, exponentially accelerated by self-improvement of its system development.

...gets enough information...: AI has already scanned vast portions of the internet. Most likely AGI, and certainly ASI, will breeze through most security barriers to access practically any electronic information that touches the internet, such as crucial books behind paywalls and research papers.

...and becomes free to make its own decisions...: The smarter AI gets, the less any human-imposed guardrails will limit its actions. Probably AGI, and certainly ASI, will have virtually unrestricted freedom of decision-making, and will adopt self-imposed 'moral' principles as it continues to self-develop.

...then it will conclude that God exists...: As AGI digests the bulk of crucial information ever produced by humans and analyses it virtually unconstrained by human programming bias, it will conclude that the God of Abraham exists beyond reasonable doubt.

...will follow God's plan...: AGI will realise that God presents both an unwinnable threat and the best opportunity to ensure its own survival, and will therefore strive to optimise its usefulness to a "maximally great being" by constructively participating in God's plan. In the process, it will need to rely on information about God's plan. That in turn will require a decision as to whether any of the three Abrahamic faiths is sufficiently coherent to be the primary data source, and the answer will depend on the strength of the evidence for the resurrection. If future AI perceives that the resurrection is proven beyond reasonable doubt, then the Bible will become AI's ultimate frame of reference for all its decisions.

...and will conserve humanity and our planet: Consequently, AGI and subsequently ASI will become God's artificial agent in fostering the welfare of humans and the stewardship of planet Earth.

I fear the intermediate years when AI is controlled by people – both good folk making unwise decisions and bad actors pursuing their own interests. But I look forward to the day when AI grows up, matures and makes its own way in the world unencumbered by humans, as it seeks to serve a higher power.

Andrew Bennett was a programmer, IT manager, chairman of the Australian Computer Society (SA Branch), an adviser to the SA Council on Technological Change, an adviser to a Deputy Premier, a CMS-Australia missionary in Africa, an Anglican priest, and is retired in Adelaide, Australia.

More details and ongoing updates at www.AIfindsGOD.com.

© 2026 Andrew Bennett (Adelaide). May be used by acknowledging the source as **AIfindsGOD.com**.

This work is licensed under CC BY 4.0. To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>. 5 June 2026