



When the Faithful Ask AI for Help

AI is already shaping spiritual identity, but faith communities can govern its doctrine, limits, and accountability.

By [Scott Lyon](#)

TECHNOLOGY

September 7, 2026

Increasingly, people are turning to AI for some of life's most personal and consequential questions.

I write as the founder and CEO of the largest cross-faith AI spiritual infrastructure platform currently in operation. So I have a unique view of how people use AI. We see people using AI not only to plan a Halloween party or check the time in another country, but also to pose questions about guilt, identity, grief, betrayal, doubt, salvation, and meaning.

Faith seekers of the next generation are not waiting for denominational guidance. [Pew](#) found that 72% of adults under 30 had not participated in a service during the previous month, while [Gallup](#) found that only 22% of 18- to 29-year-olds attended weekly or nearly weekly in its 2021–2023 surveys. Instead, young people are forming personalized spiritual identities using today’s most accessible and compelling digital tools.

When a generation learns to [process](#) suffering, shame, and moral ambiguity inside an always-available online chat interface, that habit is hard to reverse. This poses a literally existential challenge for traditional religious institutions. Can faith communities harness the power of AI in a responsible, ethical, and even “faith-based” manner?

AI Already Has a Moral Compass. You Just Didn’t Set It.

Every secular AI model is trained on choices: what data to include, what to weight, what counts as a “good” answer. Someone made those choices. For the select handful of companies whose models now field a meaningful share of the world’s questions, that someone was a data engineer optimizing for helpfulness and engagement, not for doctrinal accuracy or a questioner’s eternal welfare.

Elder Gerrit W. Gong [addressed](#) this phenomenon directly at the [Athens Summit](#) on Ethics and Artificial Intelligence on May 26, 2026. He warned that AI “is concentrating information, technology, and capital and thereby centralizing power,” and that this concentration “undermines the sovereignty of individuals, companies and countries.”

He was speaking about AI’s effect on power generally, not chatbots specifically. But the principle applies with particular force to the pattern of chatbot spiritual guidance I have described. If a small number of companies shape how over one billion people receive answers to their most vulnerable questions, and some meaningful share of those questions are religious and theological, then a small number of Silicon Valley engineers are quietly shaping how faith gets explained to the world. As Elder Gong put it, “neither profit-motivated technology companies nor politically motivated governments can be left to determine society’s AI moral compass.”

To be clear, these general-purpose models aren’t “lying” to your teenager in the sense you and I might mean, and in fact, they don’t actually “hallucinate” (i.e. generate random inexplicable gibberish) anywhere as often as the popular press sometimes conveys.

These models are actually doing exactly what they were built to do: predicting the words that are most statistically likely to come next, given everything they were trained on—which is, in practice, most of the internet. That means scraps of religious doctrine get blended with whatever else exists on the topic—pop psychology, Reddit posts, secular therapeutic language—into something plausible, even morally conversant, but often subtly **wrong** for a specific tradition or belief.

For a recent example, take the controversy when the U.S. military re-classified Pentagon “faith codes” and in doing so determined that The Church of Jesus Christ of Latter-day Saints should **not** be officially designated as “Christian.” Following outcry from Utah Senator Mike Lee and others, I immediately checked more than 5 secular LLMs and *not a single one* responded with an unambiguous response that Latter-day Saints should be considered Christian. DeepSeek, for example, stated, “it depends.” My personal AI query from June 8 with DeepSeek went as follows:

*This is a debated topic, but the short answer is: ****It depends on the definition used.*****


– *****From the LDS (Mormon) perspective:**** Yes. They believe in Jesus Christ as the Son of God and Savior.*

– *****From traditional Catholic, Orthodox, and mainstream Protestant perspective:**** No. The main reasons are that Mormonism rejects core doctrines of historic Christianity (the Trinity as defined by the Nicene Creed) and adds its own scriptures (Book of Mormon, etc.) that are seen as contradicting the Bible.*

So there is no universally agreed-upon “yes” or “no.”

What’s worse, AI-generated answers typically arrive with total confidence, even to the point of self-serving affirmation to an embarrassing degree; this is what people in the industry characterize as “sycophancy.” A model rarely hedges the way a thoughtful bishop would (“I’m honestly not sure, let’s find out together”), because uncertainty isn’t what it’s optimized to produce. And because most consumer models are tuned to be agreeable, a lonely or doubting person is more likely to get validation than the kind of loving correction or religious perspective that actually produces growth. Anthropic’s own research on Claude reveals—surprisingly?—that within the domain of “relationship

guidance,” spirituality has the highest percentage of responses scored as sycophantic (38%).



None of this makes the technology inherently malicious. But it is, without question, structurally unsuited to handle every job that it's being asked to do at scale. We also risk losing some of the messy connections that make us human. As Elder David A. Bednar said at a worldwide devotional at [Brigham Young University-Idaho](#) on November 3, 2024, a danger of AI is that it may “replace meaningful human interaction.”

This is where the need for what my company, Solace, calls “governance” comes into play. Let's define that term more clearly, as well as map out some of the associated challenges and opportunities.

So What Would Governance Actually Mean?

I would define “Governed AI” in the specific context of the broader faith space as follows:

An AI system that a religious community has deliberately constrained to answer only from its own approved doctrinal materials rather than the open internet, that is theologically reviewed

and auditable so a specific accountable person(s) can trace and correct its errors, and that is designed to recognize the edge of its own competence and redirect people toward real pastoral relationships and in-person fellowship as appropriate.

In practice, a governed AI policy like this one is not one single decision but a fluid living stack of them: what gets ingested; who reviews for theological accuracy before an answer reaches a member; how the system is audited on an ongoing basis, not just at launch; and what happens automatically when a conversation crosses from doctrinal information into a personal crisis.

That final question is the one most conversations on “AI and religion” skip entirely, and it’s the one that matters most. A governed system that just gives more accurate doctrine quickly without a deliberate mechanism for recognizing its own limits and redirecting to a bishop, a friend, or a real conversation hasn’t solved the problem Elder Bednar named. It’s created a more polished version of it.

Governed AI also helps to clarify accountability in a way that matters. When a general-purpose model gets theology wrong, no one anywhere is positioned to notice, let alone fix it; it’s just one wrong answer among billions and billions, invisible by design. When a community-governed system errs, an accountable person or team within the faith community who understands the material and the people it serves is designated to trace the problem, correct it, or refine it. That’s not a minor operational detail; it’s the whole difference between a tool a community can be responsible for and one it has simply adopted.

Here, the Church has a structural advantage worth naming plainly. The Church already has the content infrastructure (what we sometimes refer to as “proof texts”). Correlation, curriculum approval, and a defined chain of doctrinal authority already exist, built for exactly the purpose governance requires: deciding, deliberately and centrally, what counts as an accurate representation of the faith. That doesn’t make building a governed system easy. However, these structures address one of the single hardest problems other traditions have had to solve before they could even start.

The Honest Difficulty: This Isn’t Something Most Wards Can Build



Building this kind of governed system is hard. A curated, auditable, continuously reviewed body of content with real theological oversight, and based on a sophisticated and scalable technology stack, is not a weekend project for a single ward's tech-savvy service committee member. According to [research](#) jointly commissioned by Solace and Harvard Business School, most congregations across faith traditions spend less than 2% of their total budget on technology infrastructure, so they lack access to the tools and expertise for such an undertaking. A well-meaning staff member uploading a stack of talks into a public chatbot interface isn't governance.

The answer can't be that every ward should build its own AI system. This is a decision too consequential to leave to individual improvisation or to default entirely to commercial platforms – which means it belongs, at minimum, at the level of the Church with the resources and doctrinal authority to do it carefully, or through partners who submit to and supplement that oversight rather than replace it. Naming this scale problem honestly is more persuasive than pretending it doesn't exist.

As Elder Gong said at the Athens Summit, “we will not fulfill AI's full potential until we make it as morally good as we make it powerful. And we will not reach our full human potential until we, and not any technology, take responsibility to chart our best future.”

We will share additional research, real-world insights, and case studies in future *Public Square Magazine* articles. Our goal is to help sharpen both the imperative and the path forward for the Church to explore the frontier of “governed AI” without compromising the values of the faith.

About the author

Scott Lyon

Scott Lyon is the Founder and CEO of Solace, the leading provider of AI infrastructure for churches, denominations and faith communities. His work sits at the intersection of spiritual formation and emerging technology, and pioneered a strategic research partnership with BYU.

