

From Minerals to Minister: Exploring the Materiality, Embodiment, and Place of AI-Generated Sermons

Alison Gerber
Ph.D. Candidate in Preaching
Baylor University

Abstract: *AI cannot preach, claims one blogger, since preaching involves “real people, in a real place, through a real preacher.”¹ AI’s presumed lack of physicality has been one argument made against its use in preaching.*

However, are sermons generated by AI truly so disembodied, immaterial, and placeless? This paper traces, step-by-step, the actual development of one AI sermon – from the extraction of rare earth minerals, all the way to a completed manuscript upon the preacher’s desk. In so doing, this paper reveals that AI preaching is, in fact, multi-bodied, multi-material, and pluri-local. It raises the question: if these are the actual bodies, materials, and places used for AI preaching, what does this mean for a preacher’s use of it?

The term *artificial intelligence* hardly evokes images of trees and lakes, or soil and rocks, or villages and cities teeming with people. Instead, it implies the opposite: *artificial* meaning without pre-history, something human-made; and *intelligence* as something immaterial, an ability rather than a thing. Big Tech reinforces these ideas by describing AI as made of algorithms, machine learning, networks, and the cloud. Each of these terms both reflects and fortifies the prevailing perception that AI is an immaterial, disembodied, and placeless force.

It is this image of AI that is routinely used in arguments against deploying AI in the practice of preaching. AI cannot interpret scripture for preaching, claims New Testament professor John Boyles, since “the bot is disembodied,” and therefore, “its interpretations are necessarily disembodied.”² Christian ministry professor Ed Stetzer instructs preachers to balance “efficiency with quality” – weighing the choice between using a “disembodied automaton” and their own personal interaction with the church.³ Even ChatGPT joins in – when Pastor Randal Pelton asked the chatbot “What can I do as a preacher that you cannot do?” it replied, humans are “embodied and pastoral,” while AI is “disembodied and generic.”⁴

Are AI-generated sermons, or AI more broadly, truly as immaterial, disembodied, and placeless as these critics suggest? This paper will challenge that assumption. Step by step, it traces the making of an AI sermon: from the minerals that enable computation all the way to the sermon manuscript laid on the minister’s desk, just waiting to be preached. While a comprehensive survey of every stage of the development of an AI sermon is beyond the scope of this work, this paper will focus on five key stages along the way: 1) extraction, 2) data collection, 3) data cleaning, 4) prompting, and 5) inference (the processing of prompts).

¹ Brett Landry, “Should Pastors Use AI in Sermon Preparation?,” *The Gospel Coalition Canada*, May 26, 2025, <https://ca.thegospelcoalition.org/article/should-pastors-use-ai-in-sermon-preparation/>.

² John Boyles, “Misreading Scripture with Artificial Eyes,” *Christianity Today*, July 27, 2023, <https://www.christianitytoday.com/2023/07/ai-chatgpt-exegetical-tool-bible-scripture-sermon-mount/>.

³ Carter Wiley and Ed Stetzer, “The Ethics of AI in Sermon Writing,” *Sermon Shots*, January 30, 2025, <https://sermonshots.com/blog/the-ethics-of-ai-in-sermon-writing/>.

⁴ Randal Pelton, “What Can I Do That AI Can’t? Reflections for Preachers in a Digital Age,” *Pelton on Preaching*, April 13, 2025, <https://peltononpreaching.com/2025/04/13/3544/>.

At each juncture, this paper will reveal that AI is profoundly dependent upon place, body, and materials, complicating the claim that AI should be rejected in preaching on the grounds of its supposed disembodiment. However, this new perspective will not clear the way for an uncritical embrace of AI. Instead, it will bring into focus new challenges for the utilization of artificial intelligence in preaching – rooted, not in what AI supposedly lacks, but in the material realities of what AI truly is.

I. Extraction

When and where does an AI sermon begin? Contrary to appearances, an AI sermon is not born when an overextended pastor opens her laptop and tells ChatGPT to “Write me a sermon on” Instead, it begins in deep time – millions and millions of years ago – when, through magmatic processes, hydrothermal activity, and shifting tectonic plates, mineral deposits formed on or beneath the surface of the earth.⁵

All AI systems run on physical components housed in human-made structures, and each component is a composite of human-made materials and many minerals and elements. For example, the chips used in modern computing require silicon and small amounts of boron, phosphorus, copper, cobalt, and gold. Magnets in hard drives are made of rare-earth elements (REEs) such as neodymium and dysprosium. Cooling mechanisms, built to prevent overheating, require aluminum, copper, REEs, and steel. Even the iPhone used to access ChatGPT contains around 75 elements, two-thirds of the periodic table.⁶

One significant element present in all AI systems is lithium. Lithium is a critical component in the rechargeable lithium-ion batteries that power multiple aspects of AI infrastructure – from data-center backup systems to the laptops, servers, and mobile devices on which AI runs.⁷ One way lithium is extracted is through brine evaporation, a process in which underground reservoirs are pumped into ponds where sunlight evaporates the water, concentrating the lithium.⁸

Lithium extraction via brine evaporation can take a severe toll on water-scarce regions, as extracting one ton of lithium may require evaporating up to 2 million liters of water.⁹ The strain of this extraction can be seen in Chile’s Salar de Atacama, which holds one of the largest lithium reserves in the world. As a result of mining in the region, flamingo populations are now under threat, foxes are creeping into villages, vicuñas are foraging further from home, carob trees are dying, natural springs have turned to sinkholes, and fields that were formerly meadows thick with grass have become open vistas of dry, cracked earth.¹⁰

Another essential group of elements for AI is rare earth elements (REEs) – the 15 elements on the periodic table used in almost every component of data centers that train and run

⁵ Kate Crawford and Vladan Joler, “Anatomy of an AI System,” accessed August 5, 2025, <http://www.anatomyof.ai>.

⁶ “Critical Minerals in Artificial Intelligence,” SFA (Oxford), accessed September 1, 2025, <https://www.sfa-oxford.com/knowledge-and-insights/critical-minerals-in-low-carbon-and-future-technologies/critical-minerals-in-artificial-intelligence/>; Konstantin Kakaes, “The All-American iPhone,” *MIT Technology Review*, June 9, 2016, <https://www.technologyreview.com/2016/06/09/159456/the-all-american-iphone/>.

⁷ Crawford, *Atlas of AI*, 30.

⁸ “How Is Lithium Mined?” February 12, 2024, <https://climate.mit.edu/ask-mit/how-lithium-mined>.

⁹ “World Water Day: The Water Impacts of Lithium Extraction,” *Wetlands International Europe*, March 22, 2023, <https://europe.wetlands.org/blog/world-water-day-the-water-impacts-of-lithium-extraction/>.

¹⁰ Ione Wells, “The Rise of Green Tech Is Feeding Another Environmental Crisis,” *BBC*, July 19, 2025, <https://www.bbc.com/news/articles/c30741j351go>.

AI systems. REEs are used in chip manufacturing, cooling fans, fiber-optic networks, high-performance magnets and batteries, and speakers, microphones, and displays.¹¹

When REEs are extracted, the ratio of usable minerals to waste is extreme. As AI researcher Kate Crawford explains, “Only 0.2 percent of the mined clay contains the valuable rare earth elements. This means that 99.8 percent of the earth removed in rare earth mining is discarded as waste dumped back into the hills and streams.”¹² Additionally, it is estimated that the process of refining one ton of REEs produces 75,000 liters of acidic water and one ton of radioactive residue.¹³

The environmental cost of extracting REEs is seen at Bayan-Obo, the largest REE mine in the world, located in Baotou, China. By the mine is a massive lake filled with 180 million tons of waste powder from ore processing.¹⁴ This “lake” reeks of sulfur, and, when standing at the shore of it, its surface stretches to the horizon. BBC Reporter Tim Maughan, visiting the lake in 2015, called it “nightmarish.” “Dozens of pipes line the shore,” he wrote, “churning out a torrent of thick, black chemical waste...it feels like hell on earth.”¹⁵

AI systems and AI sermons, when viewed from the perspective of extraction, are both material and situated in a place. However, this insight exposes a new issue that arises precisely because of AI preaching’s materiality: can preachers really employ AI if its fabrication has such a negative impact on the earth? Preachers are not just any users of AI. They are a group who, by the nature of their vocation, should be acutely aware of their responsibilities toward the environment. As Elizabeth Achtemeier argues, a preacher’s understanding of the scriptures “should act as a powerful check on all of our proud attempts to claim that we can do with the natural world as we will.”¹⁶

The solution may seem straightforward: abstain from AI and thereby reduce one’s environmental impact. However, the same minerals that form AI’s infrastructure also undergird other technologies preachers already use – the laptops, phones, and tablets preachers depend upon for their daily work. Will preachers, then, renounce these as well?

II. Data collection

After mining, the minerals required to build the hardware for an AI system are smelted, refined, and manufactured into components. These components are then distributed and assembled into domestic infrastructure (the computers and routers one uses at home), internet infrastructure, and data center infrastructure, where AI systems are created.¹⁷

The next step in the process of generating an AI sermon is to use this hardware to create a Large Language Model (an LLM). A large language model is a type of AI designed for language processing tasks, like the system that powers ChatGPT. LLMs learn to generate human-sounding language by reading a data set of text and learning to recognize patterns within it.¹⁸

The creation of an LLM requires an extraordinary amount of data. Data from an LLM is broken into *tokens*, which is the term used by AI developers for individual words or portions of

¹¹ “Critical Minerals in Artificial Intelligence.”

¹² Crawford, *Atlas of AI*, 36.

¹³ Crawford, *Atlas of AI*, 37.

¹⁴ Crawford, *Atlas of AI*, 36.

¹⁵ Tim Maughan, “The Dystopian Lake Filled by the World’s Tech Lust,” *BBC*, April 2, 2015, <https://www.bbc.com/future/article/20150402-the-worst-place-on-earth>.

¹⁶ Elizabeth Achtemeier, *Nature, God and Pulpit*, (Eerdmans, 1992), 67.

¹⁷ Crawford and Joler, “Anatomy of an AI System.”

¹⁸ Ethan Mollick, *Co-Intelligence: Living and Working with AI* (New York, NY: Portfolio, 2024), 9-11.

words. GPT-4 is estimated to have been trained on a total of 13 trillion tokens.¹⁹ The training data for next-generation AI models, such as OpenAI's GPT-5, has been estimated at 70 trillion tokens.²⁰ If these figures are even roughly accurate, it would be analogous to every living person, all 8.1 billion, contributing one long academic paper to GPT-5 each.

Due to the massive amount of data required to create AI systems, data has become a valuable commodity in the 21st century. Data is so valuable that it has begun to be spoken of as an unrefined resource, with terms emerging like “data mining” and “data is the new oil.” Promoting the idea of data as a commodity has, however, “shifted the notion of data away from something personal, intimate ... toward something more inert and nonhuman.”²¹ Crawford questions the commodification of AI training data in her investigation into the Multiple Encounter Data Set (MEDS). MEDS is a dataset of mugshots of people arrested multiple times throughout their lives. The photographs in MEDS are used to train AI facial recognition software – to teach it how to recognize the same person as they age, or how to retrieve a face from a surveillance video, or even simply to detect the basic mathematical components of a face.²²

Looking closely at the photographs that constitute MEDS brings into sharp relief a crucial truth: that behind each datapoint is a human being. These photographs are not merely “mathematical components of faces.” They are images of real people: some young, others old. Not only are they human, but they are humans photographed at a vulnerable moment – some visibly sad, confused, or shocked; some bearing visible signs of violence – bruises, cuts, or blood. There are stories behind these photographs – stories that involve heartbreak, sorrow, regret, remorse, injustice, police brutality, and racism. Yet as soon as these images become datapoints, all of this humanity is erased.

OpenAI's LLMs, such as GPT-4 and GPT-5, are very likely to contain sermons in their training data. They are known to contain, at least, data scraped from the web, as well as the text of books that are old enough to be in the public domain.²³ Considering sermon manuscripts are found in both, it is safe to assume that sermons, representing many human voices, form at least some of the training data of the LLMs that power ChatGPT.

Keeping MEDS in mind, one comes to realize that there are also human elements of sermons that are discarded when these sermons are used as AI training data. Imagine: when cataloged sermon #95,512 was preached, was the pastor doubting that Sunday, or full of faith? When sermon #9001 was delivered, was the pastor preaching her very first sermon, or preaching her very last? What was happening when Sunday sermon #15,667 was preached – was the congregation celebrating their pastor's 50th anniversary, or grieving the loss of a beloved member of their church?

¹⁹ Patrick McGuinness, “GPT-4 Details Revealed,” Substack newsletter, *AI Changes Everything*, July 12, 2023, <https://patmcguinness.substack.com/p/gpt-4-details-revealed>.

²⁰ Alan D. Thompson, “What’s in GPT-5? (2024),” *LifeArchitect.Ai*, July 28, 2024, <https://lifearchitect.ai/whats-in-gpt-5/>. Note: these figures are sourced from leaks and have not been officially confirmed; OpenAI has not publicly disclosed the size of its training datasets, citing “the competitive landscape and the safety implications of large-scale models like GPT-4” as justification for their opacity. OpenAI et al., “GPT-4 Technical Report,” arXiv, March 4, 2024, <https://doi.org/10.48550/arXiv.2303.08774>.

²¹ Crawford, *Atlas of AI*, 113.

²² Crawford, *Atlas of AI*, 92.

²³ Stefan Baack, “A Critical Analysis of the Largest Source for Generative AI Training Data: Common Crawl,” *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, ACM, June 3, 2024, 2199–208; Simon Rowberry, “The Value of Books in the Age of Generative AI Training Data,” *Convergence*, SAGE Publications Ltd, July 2, 2025.

What is so concerning about AI training data is not that it is non-human, but that it *is* human, but only *part* of a human. John McClure defines embodiment in preaching as “all the ways in which preaching is enhanced by virtue of its presentation through the human body.”²⁴ This data, like the text of a sermon, did find its origins in lived, enfleshed experiences. It too came through human bodies. Yet in the process of being fed to an LLM, this data became only a partial representation of human experience, neglecting significant dimensions of the human reality from which it emerged. AI training data, therefore, could function as a partially erased representation of human experience, which in essence devalues and dehumanizes that experience.

The field of homiletics is not immune to doing the same thing – homiletics archives do a similar thing when they collect data from minoritized preaching traditions. Historically, when a sermon event from the African American tradition is written down, it is the linguistic data from the preacher that is preserved in the transcription – the words that the preacher said. However, transcriptions of African American sermons almost always omit any paralinguistic data from the preacher, as well as any linguistic or paralinguistic data that might be collected from the congregation.²⁵ This data includes, but is not limited to, any call-and-response, any variation in the preacher's vocal volume, pitch, pace, and tone, any long pauses or sounds independent of speech, such as laughter, whooping, sighs, humming, and singing. William Turner highlights the importance of these elements of African American preaching when he argues that, in African American preaching, there is a “fusion of sonic and conceptual things” and the decoupling of them “unravels and erodes the very nexus from which the power of African American preaching emerges.”²⁶ Yet these elements *are* uncoupled when African American preaching is written down.

To collect certain elements while discarding others is never a neutral act: it is to assign value to one thing while treating another as negligible. When MEDS is used to determine “parts of a face,” but not “police brutality,” it is valuing the mathematical proportions of a face but devaluing an experience. When only the words of sermons are used as training data for LLMs, the contexts of those sermons are devalued. Moreover, homileticsians do the same: when transcriptions of African American preaching preserve only the preacher's words, they devalue the tradition's distinctives. Perhaps now, in this AI age, homiletics needs more than ever those “scholars and scribes” Frank Thomas called for in 2015, who would “transfer the living presence of the sermon to the written page.”²⁷

III. Data evaluation

After this system is trained on human data, an LLM will finally emerge, but the AI model still requires some fine-tuning. This, then, is where humans step in – “to provide data, annotate it, and verify algorithm inputs.”²⁸ This work requires the employment of human *data evaluators*.

²⁴ John S. McClure, *Preaching Words: 144 Key Terms in Homiletics* (Westminster John Knox, 2007), 26.

²⁵ A notable exception: Jon M. Spencer, *Sacred Symphony: The Chanted Sermon of the Black Preacher* (New York, NY: Praeger, 1988).

²⁶ William Turner, “The Musicality of Black Preaching,” in ed. Jana Childers. *Performance in Preaching* (Baker Academic, 2008), 207.

²⁷ Frank A. Thomas, *They Like to Never Quit Praisin' God: The Role of Celebration in Preaching*. Updated ed. (Cleveland, OH: Pilgrim Press, 2013), xiii.

²⁸ Julian Posada, “Platform Authority and Data Quality: Who Decides What Counts in Data Production for Artificial Intelligence?” in *Decoding Digital Authoritarianism*, ed. Ziyaad Bhorat and Martin Rauchbauer (Berggreun Institute, 2023), 82.

One example of humans employed to refine an LLM is the Kenyan data evaluators hired by OpenAI to reduce ChatGPT's toxicity. Data evaluators are low-wage workers responsible for evaluating an LLM's output to refine it before it is released to the public. OpenAI had already made an LLM called GPT-3 in 2020. GPT-3's language creation abilities were already impressive, yet it had a major issue – it “was also prone to blurting out violent, sexist and racist remarks.” OpenAI needed to create an additional safety mechanism to ensure its LLM would not continue to generate this kind of content. To achieve that goal, they planned to create a data set of “labeled examples of violence, hate speech, and sexual abuse,” so that GPT-3 “could learn to detect those forms of toxicity in the wild.”²⁹

In November 2021, OpenAI sent thousands of pieces of text to Sama, a firm in Nairobi, Kenya. These texts contained all the horrors of the internet: “describing situations in graphic detail like child sexual abuse, bestiality, murder, suicides, torture, self-harm and incest.” The employees at Sama were responsible for evaluating which type of graphic content each text contained. When TIME reporters interviewed four of the workers from Sama, all of them “described being mentally scarred.”³⁰ They were paid between \$1.32 and \$2 an hour to perform this work.

Sama's relationship with OpenAI dissolved in February 2022, eight months before their contract ended, when OpenAI sent Sama the request to collect and label “sexual and violent images – some of them illegal under U.S. law.” The workers at Sama evaluated 1400 images for OpenAI – images that contained child sexual abuse, bestiality, rape, sexual slavery, graphic death, violence, or serious physical injury. Sama ended the contract, not wanting “to expose their employees to such [dangerous] content again.”³¹

Data evaluators, like those in Kenya, are not the only low-wage workers employed in AI production. There are also *data annotators* – individuals hired to label data before they are fed into machine learning systems. For instance, workers may be contracted to tag images of bodies, pedestrians, buildings, and vehicles to train driverless cars to recognize what is on the road.³² There are also *data generators*, laborers who produce input by submitting text, capturing photographs or videos, or recording audio of their surroundings.³³ Incredulously, there are even *data impersonators*, hired to mimic the behavior of chatbots or AI personal assistants.³⁴

The necessary employment of data evaluators, annotators, generators, and even impersonators reveals that AI requires the labor of many human bodies, many of whom are marginalized and underpaid.³⁵ This whole system is not so different from the reliance on underpaid and exploited workers in the textile industry. There, laborers farm cotton, weave fabric, and sew clothing, removed from each other, from their product, and from the consumers of their product. The alienation that workers in both industries experience supports their

²⁹ Billy Perrigo, *OpenAI Used Kenyan Workers on Less Than \$2 Per Hour*, TIME (January 2023), <https://time.com/6247678/openai-chatgpt-kenya-workers/>.

³⁰ Perrigo, *OpenAI Used Kenyan Workers*.

³¹ Perrigo, *OpenAI Used Kenyan Workers*.

³² Posada, “Platform Authority and Data Quality: Who Decides What Counts in Data Production for Artificial Intelligence?” 83.

³³ Paola Tubaro et al., “The Trainer, the Verifier, the Imitator: Three Ways in Which Human Platform Workers Support Artificial Intelligence,” *Big Data & Society* 7, no. 1 (2020). (Both data annotation and data generation are addressed here.)

³⁴ Arvind Narayanan and Sayash Kapoor, *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference* (Princeton University Press, 2024), 230.

³⁵ Crawford, *Atlas of AI*, 63-64.

mistreatment and their disempowerment, all while consumers are clueless to – or can willfully ignore – the impact of their purchases or their use of AI products on human beings.

Though these workers' involvement in AI text generation may not capture the full sense of the term "embodiment" as it is used in homiletics, this does not mean the system is disembodied. To make such a claim would be to ignore the presence of the flesh-and-blood human beings that keep AI systems afloat. Perhaps a better term to capture the distributed, collective human labor underwriting these systems would be *multi-bodied* rather than disembodied.

Preaching has been multi-bodied long before AI, most recently in the movement toward collaborative preaching. Collaborative preaching, as defined by John McClure in his *The Roundtable Pulpit*, is

a form of preaching in which the preacher and hearer work together to establish and interpret the topics for preaching. They also decide together what the practical results of those interpretations might be for the congregation. The preacher then goes to the pulpit and represents this collaborative process during sermon delivery.³⁶

In a way, the multi-bodied AI sermon is the ultimate realization of McClure's vision. Not a few, but millions of voices, across cultures and time, are woven together in the data processed by an LLM in order to write a single sermon. Moreover, with respect to data cleaning, hundreds of human beings are involved in the "editing" of the sermon, determining what is appropriate and what is not. Only in the pulpit does the preacher "represent this collaborative process in the event of sermon delivery."

However, the AI sermon also constitutes a distortion of collaborative preaching, in that it undermines two of its central aims: empowerment and recognition. As McClure observes, "When your preaching empowers a congregation it must express power with others."³⁷ In contrast, AI preaching disempowers and dehumanizes through the hidden labor and data upon which it depends. Furthermore, McClure insists that collaborative preaching requires sustained attention to the other, listening to discern who the other truly is. In AI preaching, however, the other is rendered invisible – concealed within the opaque infrastructures of Big Tech.

This tension raises a pressing question: why does AI preaching appear so close to collaborative preaching yet remain fundamentally different? Is the problem a problem of scale? Could it be that when collaboration is dispersed across millions of unseen contributors, the structure collapses under the weight of its own expansiveness?

More significantly, however, the bodies employed by AI organizations reflexively challenge practitioners of collaborative preaching to interrogate their own methods. Might collaborative preaching, like AI sermons, also depend upon hidden laborers whose contributions remain unacknowledged? Whose voices are not found around the round table?

IV. Prompting

At last, ChatGPT is ready for use – a preacher may now open a browser window and enter their prompt: "Write me a sermon on ..."

³⁶ John S. McClure, *The Roundtable Pulpit: Where Leadership & Preaching Meet* (Abingdon Press, 1995), 48.

³⁷ McClure, *The Roundtable Pulpit*.

It is hard to see how this moment might be dismissed as disembodied, given that it presents itself through the body of a human preacher. Unless, of course, one supposes that composing a prompt is purely mechanical – something like pressing an icon on a touchscreen to order a takeaway meal.

Recent research, however, demonstrates that prompt formulation is not mechanical at all but is shaped by the choices and personalities of an LLM's user. For one, LLMs have been shown to be affected by a user's point of view. They experience "anchoring bias," being especially persuaded to a certain point of view, alluded to in the first piece of information the LLM encounters in a prompt.³⁸ Second, studies have also shown "impolite prompts often result in poor performance, but overly polite language does not guarantee better outcomes," and that "the best politeness level is different according to language."³⁹ Third, a study by Sawalha et al. revealed that female students who used ChatGPT to answer exam questions tended to compose single-question prompts, whereas male students tended to use multiple-question prompts, and, as a result, male students performed better on the exam.⁴⁰ Finally, more generally, in an effort to tailor responses to individual users, an LLM will "sometimes provide answers based on assumed demographic traits inferred from identity markers" such as a user's name or dialectal markers – a practice that is called "implicit personalization."⁴¹ This personalization to a provided or assumed age, gender, race, and/or socio-economic status will lead LLMs to produce different outputs for the same prompt, or, as one paper explains, "while model internal knowledge remains stable, user identity can skew generation outcomes."⁴² In sum, the personal choices of an LLM user, as well as who the user is, as well as who the user is assumed to be by the LLM, all affect the output of an LLM's prompts.

The choices of and personality of a prompter affect both the human input and the AI output in AI sermon generation, too, as evidenced by the earliest reported ChatGPT-generated sermon, preached by Rabbi Joshua Franklin on January 1, 2023. Rabbi Franklin's sermon garnered much press interest, but little attention was paid to the prompt he offered to ChatGPT. Franklin did not simply ask for any sermon. He asked for a sermon on the weekly Torah portion, on the themes of vulnerability and intimacy, incorporating a quote from Brené Brown.⁴³

The uniqueness of Franklin's prompt reveals that the Rabbi made several choices the day ChatGPT supposedly wrote that sermon all on its own. Franklin chose to create a sermon using ChatGPT and said in the coming week that he would preach it to his synagogue. He decided that

³⁸ Jiaxu Lou, "Anchoring Bias in Large Language Models: An Experimental Study," arXiv, December 9, 2024, <https://doi.org/10.48550/arXiv.2412.06593>.

³⁹ Ziqi Yin et al., "Should We Respect LLMs? A Cross-Lingual Study on the Influence of Prompt Politeness on LLM Performance," arXiv, October 14, 2024, <https://doi.org/10.48550/arXiv.2402.14531>.

⁴⁰ Ghadeer Sawalha et al., "Analyzing Student Prompts and Their Effect on ChatGPT's Performance," *Cogent Education* 11, no. 1 (2024).

⁴¹ Vera Neplenbroek et al., "Reading Between the Prompts: How Stereotypes Shape LLM's Implicit Personalization," arXiv, September 16, 2025, <https://doi.org/10.48550/arXiv.2505.16467>.

⁴² Ibid; Miao Zhang et al., "Identity-Robust Language Model Generation via Content Integrity Preservation," arXiv, January 14, 2026, <https://doi.org/10.48550/arXiv.2601.09141>. For more information concerning the impact of a user's identity – real or assumed – on an LLM's output, see Anjali Kantharuban et al., "Stereotype or Personalization? User Identity Biases Chatbot Recommendations," in *Findings of the Association for Computational Linguistics: ACL 2025*, ed. Wanxiang Che et al. (Association for Computational Linguistics, 2025), <https://doi.org/10.18653/v1/2025.findings-acl.1254>.

⁴³ Johana Bhuiyan, "Are Chatbots Changing the Face of Religion? Three Faith Leaders on Grappling with AI," *Technology, The Guardian*, April 7, 2023, <https://www.theguardian.com/technology/2023/apr/07/chatgpt-artificial-intelligence-religion-faith-leaders>.

he would ask for a sermon based on the weekly portion of the Torah. He asked for the sermon to be on vulnerability and intimacy, specifically. He desired a quote by Brené Brown—all conscious, personal, creative choices – *human* choices.

The influence of the human preacher on an AI sermon through their prompt is corroborated by Frida Mannerfelt and Rikard Roitto's investigation into the way Swedish preachers use AI in their sermon preparation. Among the preachers they observed, they identified that each preacher's prompts were deeply connected to their identity and homiletical orientation. For example, one preacher, "Luke," emphasized that what characterized him as a preacher was the importance he placed on the context of his listeners, as well as the context of his Biblical text. Mannerfelt and Roitto found that Luke's homiletical interests corresponded to his prompting strategies.⁴⁴ Furthermore, Mannerfelt and Roitto found that "none of the preachers copy text from the LLM directly into the sermon." When asked "how they chose among all the different ideas that the LLM throws at them, most initially express ... they just intuitively pick ideas that 'speak to them.'" The process was "an unconscious process, no different from how preachers usually use their intuition to choose among options."⁴⁵

Two essential features of preaching – the personal approach and the discernment preachers employ – thus remain operative when preachers prompt and sift through AI-generated output. AI sermons are shaped – through their prompting – by the preacher's unique, embodied personality and discernment. The act of prompting is not a mechanical gesture but a deeply human activity: weaving together texts, themes, and voices, imprinting the sermon with the preacher's creative choices.

V. Inference

One step remains in this before the sermon manuscript is complete: the preacher's prompt must be processed by the system. In machine learning, this stage is called *inference* – the process by which a trained model receives a new input (a prompt such as "Write me a sermon on John 3:16") and generates an output (the sermon on John 3:16). Having begun with minerals in the ground, through inference, the preacher's sermon manuscript will be complete.

Although a preacher may spend nothing to retrieve a sermon manuscript from ChatGPT, processing a prompt is not free: it requires a significant amount of both energy and water. Sam Altman, the CEO of OpenAI, states that the "average query" to ChatGPT "uses about 0.34 watt-hours, about what an oven would use in a little over one second, or a high-efficiency lightbulb would use in a couple of minutes. It also uses about 0.000085 gallons of water; roughly one fifteenth of a teaspoon."⁴⁶ However, Altman never defines what an "average query" entails, so these figures may be imprecise. To arrive at more reliable estimates, Jegham et al combined publicly available usage data with "company-specific environmental multipliers and statistical inference of hardware configurations" (in layman's terms, information about how the computers that process these systems are built and operate) across 30 different models to determine the energy, water, and carbon footprint of LLM inference.⁴⁷ As a result, they were able to estimate

⁴⁴ Frida Mannerfelt and Rikard Roitto, "Preaching with AI: An Exploration of Preachers' Interaction with Large Language Models in Sermon Preparation," *Practical Theology* 18, no. 2 (2025): 127–38.

⁴⁵ Mannerfelt and Roitto, "Preaching with AI," 134.

⁴⁶ Sam Altman, "The Gentle Singularity," accessed March 20, 2026, <https://blog.samaltman.com/the-gentle-singularity>.

⁴⁷ Nidhal Jegham et al., "How Hungry Is AI? Benchmarking Energy, Water, and Carbon Footprint of LLM Inference," arXiv, May 15, 2025, <https://doi.org/10.48550/arXiv.2505.09598>.

that a medium-length inquiry (1,000 tokens of input and 1,000 tokens of output) through GPT-4o uses around 1.21 Wh of energy (the equivalent of running an LED light for 7 minutes), 3.7mls of water, and emits around 0.38 gCO₂e (around half a human breath).⁴⁸

While the energy required to process a single prompt may seem inconsequential, it becomes astronomical when scaled to ChatGPT's annual usage. In December 2024, OpenAI revealed that 1 billion queries were made per day across all ChatGPT deployments. Adjusting for various lengths of inputs and outputs that might occur daily, Jegham et al found that GPT4o's annual global inference would "exceed the total electricity consumption of 35,000 U.S. residential households, 50 inpatient hospitals, and even 325 universities annually."⁴⁹

ChatGPT's annual global carbon footprint and water use figures are just as sobering. When ChatGPT's carbon footprint is aggregated globally, its CO₂ figures "are comparable to the annual emissions of 30,000 gasoline powered cars" and to offset GPT-4o's annual emissions "would require of 138,000 acres of average U.S. forest, an area roughly equivalent to the size of Chicago." With respect to ChatGPT's annual water consumption, Jegham et al. estimate it to be between 1,334,991 and 1,579,680 kiloliters. Of this figure, they explain, "this consumption refers to evaporated freshwater permanently removed from local ecosystems rather than recycled. This means GPT-4o alone is responsible for evaporating an amount of freshwater equivalent to the annual drinking needs of almost 1.2 million people."⁵⁰

To the very end, an AI sermon is material – and yet again its materiality raises a new issue for homiletics, and, in particular, eco-homiletics. Earlier eco-homiletical texts, like Elizabeth Achtemeier's *Nature, God and Pulpit* and John Holbert's *Preaching Creation*, were concerned with *what* to preach – equipping preachers with exegetical material that they may construct sermons that address the environment. More recently, Leah Schade's book *Creation-Crisis Preaching*, concerned with "ways to preach ... environmental sermons that are equally pastoral and prophetic," and HyeRan Kim-Cragg's "Earth-Bound Preaching," which proposes "three movements ... to help preachers and congregations address the climate crisis," have shifted the discipline's focus from *what* to *how* to preach in light of eco-concerns.⁵¹

The existence of AI sermons, however, presents a new challenge to homiletics: *by what means* will one preach? Until now, this has hardly been a concern, since countless preachers throughout history have composed sermons without computers, electricity, or even a pen and paper. Today, however, since queries to ChatGPT incur a measurable environmental cost, a preacher's means *is* a concern.

One way to approach this issue would be to conduct energy audits of different contemporary methods of sermon preparation. Such knowledge would enable preachers to make more informed choices about whether to integrate AI into their homiletical practice. In the process, homileticians may even discover that other aspects of preaching carry environmental burdens that have gone unexamined until now.

Conclusion

⁴⁸ Jegham et al., "How Hungry Is AI?"

⁴⁹ Jegham et al., "How Hungry Is AI?"

⁵⁰ Jegham et al., "How Hungry Is AI?"

⁵¹ Leah D. Schade, *Creation-Crisis Preaching: Ecology, Theology, and the Pulpit* (Chalice Press, 2015), 3; HyeRan Kim-Cragg, "Earth-Bound Preaching: Engaging Scripture, Context, and Indigenous Wisdom," *Religions* 15, no. 3 (2024).

The goal of this paper has been to lift the veil on AI systems and reveal, for homiletics, the true nature of an AI sermon. While this article does not claim that AI preaching possesses embodiment, or is situated in a place, or uses materials in the same ways as traditional sermons, it does reveal that AI sermons are materially, geographically, and socially embedded in ways that are often overlooked in critiques of AI preaching. AI sermons, when scrutinized, are found to be not local, material, or embodied, but *pluri-local* (emerging from lithium flats in Chile, or rare earth mines in China, or a preacher's study in Sweden), *multi-material* (composed of silicon, copper, lithium, water and more), and *multi-bodied* (created with the help of human data, and data cleaners, and the preacher themselves).

Although these forms of embodiment and materiality may differ from those typically present in conventional preaching, the presence of materials, places, and bodies nevertheless complicates the claim that AI-generated sermons are simply disembodied and therefore unfit for preaching. If AI-generated sermons are insufficient, the reasons for those insufficiencies may have to be sought elsewhere.

Yet what replaces this common critique is not ease, but gravity. As Duke ethicist Matthew Elmore explains, “when nodes of the network come to light, they reveal the interconnectedness of ethical dilemmas that transcend geographical and disciplinary boundaries.”⁵² Now that at least some of the “nodes” of AI have been revealed, homiletics faces new ethical concerns regarding AI preaching. Preachers must consider – will we continue to employ hidden, underpaid labor? How have we data-fied and dehumanized the marginalized? Should we take up technologies that wreak havoc on ecosystems? What is our role in preaching when we decide, “This week, I'll preach on vulnerability”?

These are only preliminary questions – there are many more “nodes” of AI that intersect with homiletics that have not been explored above. Still, here is a starting point, a first glance, at the pluri-local, multi-material, multi-bodied system that is AI preaching.

For Further Reading

Baack, Stefan. “A Critical Analysis of the Largest Source for Generative AI Training Data: Common Crawl.” *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, June 3, 2024, 2199–208. <https://doi.org/10.1145/3630106.3659033>.

Crawford, Kate. *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. Yale University Press, 2021.

Hussey, Ian. “Preaching and Generative AI: A Perspective from Early 2024.” *International Journal of Practical Theology* 28, no. 2 (2024): 307–23. <https://doi.org/10.1515/ijpt-2024-0003>.

Jegham, Nidhal, Marwen Abdelatti, Lassad Elmoubarki, and Abdeltawab Hendawi. “How Hungry Is AI? Benchmarking Energy, Water, and Carbon Footprint of LLM Inference.” arXiv, May 15, 2025. <https://doi.org/10.48550/arXiv.2505.09598>.

Mannerfelt, Frida, and Rikard Roitto. “Preaching with AI: An Exploration of Preachers' Interaction with Large Language Models in Sermon Preparation.” *Practical Theology* 18, no. 2 (2025): 127–38. <https://doi.org/10.1080/1756073X.2025.2468059>.

⁵² Matthew Elmore, “The Hidden Costs of ChatGPT: A Call for Greater Transparency,” *The American Journal of Bioethics* 23, no. 10 (2023): 48.

Muldoon, James, Mark Graham, and Callum Cant. *Feeding the Machine: The Hidden Human Labour Powering AI*. Canongate Books, 2024.

Narayanan, Arvind, and Sayash Kapoor. *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference*. Princeton University Press, 2024.

Neplenbroek, Vera, Arianna Bisazza, and Raquel Fernández. "Reading Between the Prompts: How Stereotypes Shape LLM's Implicit Personalization." arXiv, September 16, 2025.
<https://doi.org/10.48550/arXiv.2505.16467>.