Artificial Intelligence in Parish Ministry: Use Cases, Boundaries, and Lessons Learned

A Report to the General Convention AI Working Group

Prepared by The Rev. Dr. Christopher Corbin
Rector, Trinity Episcopal Church, Oshkosh, Wisconsin
January 2026

Introduction

Over the past year, Trinity Episcopal Church has engaged in careful experimentation with Large Language Models (LLMs) and related AI technologies in parish administration and operations. What follows are actual use cases we've implemented, real benefits we've experienced, pitfalls we've discovered, and the guardrails we've established to ensure AI serves ministry rather than distorting it.

My hope is that this report might help other parishes, dioceses, and church bodies navigate these technologies with both creative openness and appropriate caution.

Understanding the Technology

In adopting this technology, I tried to move cautiously but not technophobically. I began with the principle that many of us already use artificial intelligence uncontroversially through everyday tools: optical character recognition (OCR) that digitizes scanned documents, GPS navigation that calculates optimal routes, spam filters that sort your email, and autocorrect that fixes typos as you type. These tools use machine learning to recognize patterns, make predictions, and optimize processes.

Generative AI represents the next evolution of these same underlying technologies—neural networks and machine learning algorithms—but with a fundamentally different capability. Rather than recognizing patterns or calculating solutions from existing data, generative AI creates new content from conversational prompts. When you ask ChatGPT to "draft an email explaining our policy" or request that Claude "compare these two documents," the AI generates original text, analysis, or code in response.

The use cases described below apply this evolved capability to parish operations—always with intentional boundaries, rigorous oversight, and clear theological discernment about what belongs to human ministry and what can be appropriately assisted by these tools.

It's important to note that recognizing generative AI's connection to familiar technologies does not mean we simply adopted it uncritically. There are significant differences: This technology consumes substantially more energy than traditional AI applications, raising environmental stewardship concerns; it may affect human learning and cognition in ways we don't yet fully understand; and its capacity to generate plausible but false content de-

mands new vigilance. The widespread adoption of GPS or spell-check doesn't automatically justify adopting their technological descendants without careful theological and practical discernment—which is precisely what this report attempts to model.

Part One: Where We've Found AI Genuinely Helpful

The following sixteen use cases represent areas where AI has proven to be a valuable tool in parish operations. For each, I'll describe the application, the benefits we've experienced, and the pitfalls or limitations we've encountered.

Category 1: Administrative and Operational Efficiency

Use Case 1: Meta-LLM Work (Prompt Engineering and Workflow Design)

What we're doing: Using AI to help us use AI better—drafting and refining prompts for AI interactions, determining which models work best for specific tasks, and designing multi-model workflows for complex projects.

Benefits: This has been surprisingly valuable. Learning to "talk to" AI effectively is a skill in itself and using AI to help develop that skill has accelerated our learning curve significantly. We've also discovered that different models excel at different tasks (e.g., Claude for writing and editing, Chat GPT for data analysis, Gemini for deep, propery sourced research) and AI can help us determine the optimal tool for each job.

Pitfalls: There's a risk of over-engineering solutions or spending more time optimizing AI workflows than actually doing the work. We've had to remain pragmatic about when this meta-work is worth the investment.

Use Case 2: Policy Development Support

What we're doing: Uploading our existing organizational structure, processes, policies, and strategic plan, then asking AI to identify policy gaps and draft new policies, protocols, and standard operating procedures.

Benefits: AI excels at pattern recognition and can spot gaps in our policy coverage that human review might miss. It can also draft initial policy language that maintains consistency with our existing governance documents, significantly accelerating the policy development cycle.

Pitfalls: AI-drafted policies require extensive human and legal review. The AI doesn't understand the subtle politics of parish life or the practical realities that might make a theoretically sound policy unworkable. We've learned to treat AI-generated policy drafts as "first drafts requiring significant human expertise" rather than "nearly finished documents." When we iterate on AI-drafted policies based on stakeholder feedback, we've learned to feed corrections back with very specific constraints (always present patches for human review beore integrating, 0-drit audits) to prevent AI from making unauthorized changes to sections that weren't under review.

Use Case 3: Large-Scale Data Analysis

What we're doing: Processing and analyzing large datasets for ministry insights and operational intelligence. When sensitive information is involved, we anonymize data before AI processing.

Benefits: We've used this for analyzing attendance patterns, giving trends, program participation, and other operational metrics. AI can identify patterns and correlations we might never notice through manual analysis, enabling more data-driven decision-making.

Pitfalls: The anonymization requirement is critical but can be labor-intensive. There's also a risk of becoming overly reliant on data and losing sight of pastoral wisdom and discernment. Numbers tell part of the story, not the whole story.

Use Case 4: Document Auditing for Style and Liturgical Compliance

What we're doing: Using AI to review documents for adherence to our style guide and liturgical guidelines, identifying inconsistencies and suggesting corrections.

Benefits: This has been enormously helpful for maintaining consistency across our publications. AI can catch formatting inconsistencies, style guide violations, and even liturgical errors (like incorrect or inconsistent names of liturgical days) far faster than human proofreading. A particularly helpful use in this category is transforming often overly verbose announcements into our concise and standardized format.

Pitfalls: AI sometimes flags things that are intentional stylistic choices or appropriate contextual variations. Human judgment is still needed to determine which "corrections" to accept.

Category 2: Communications and Publications

Use Case 5: Bulletin Formatting Optimization

What we're doing: Uploading our current bulletins to analyze formatting and develop comprehensive InDesign formatting guides for consistent, professional standards.

Benefits: This saved us significant time in redesign work. I was able to articulate my design knowledge and aesthetic preferences to the AI, which then translated them into a clear, step-by-step InDesign guide that our office administrator—who has limited design background—can easily understand and implement. What might have taken me hours to write out in technical terms, or required hiring a designer to document, was accomplished much more quickly through conversational iteration with AI.

Pitfalls: AI doesn't have the *je ne sais quoi* the the eye of a human designer brings. Sometimes a bulletin needs a nonstandard adjustment—a little more white space here, a subtle shift in hierarchy there—that violates the systematic rules but simply looks better. The AI-generated guide gives us excellent consistency, but we're working with our administrator to "feel" when to occasionally break the rules when her human eye tells her something needs adjustment.

Use Case 6: Research Bibliography Compilation

What we're doing: Using AI as an enhanced search engine to generate comprehensive research bibliographies on theological, pastoral, or administrative topics, with formatted outputs.

Benefits: This has dramatically accelerated literature reviews. What might have taken days of library research can now be accomplished in hours. The AI is particularly good at identifying connections between sources and finding resources across multiple databases.

Pitfalls: This is where we've encountered the most significant problems. AI "hallucination"—the generation of plausible but entirely false citations—is a serious issue. We now require mandatory hallucination audits, extensive citation verification, and human verification of every single source. We've had instances where AI generated completely fictitious journal articles, complete with plausible titles, authors, and even DOI numbers. This use case requires the most rigorous human oversight of any we've implemented. In many instances, it requires, especially if not using Gemini for the research, a two-stage "statement audit" where all the statements requiring support from a document are extracted (a task Gemini can do for outside reports) and then Gemini is tasked with finding whether those statements can be supported with outside sources.

Category 3: Technical Development

Use Case 7: Website Development (HTML/CSS)

What we're doing: Writing HTML and CSS code for our website using natural language descriptions and iterative chat-based debugging, making web development accessible to staff with only basic coding knowledge.

Benefits: This has been transformative. I, with only minimal formal coding training, can now make website updates that previously required hiring a developer. We've saved significant money and dramatically reduced the time lag between identifying a needed change and implementing it.

Pitfalls: The code AI generates sometimes works but isn't optimal or maintainable. We've learned that having someone with at least basic coding knowledge review AI-generated code is important for long-term sustainability. We've also learned the importance of drift control when using AI for iterative debugging. When we report a bug and ask AI to fix it, we now explicitly instruct: "Fix only the navigation menu bug; make no other changes to the code." Otherwise, AI might "helpfully" refactor other parts of the code, introducing new issues while fixing the original problem. As always, clear patch instructions and 0-drift auditing is our friend.

Use Case 8: Database Architecture (Customary Decomposition)

What we're doing: Using AI to decompose complex documents like our liturgical customary into component parts for transfer into dynamic database systems like Notion, including structural analysis, component naming, and creating role-specific output capabilities.

Benefits: This transformed our static customary document into a dynamic, searchable database that can be filtered by role, season, or liturgical action. Updates are now made in one place and automatically reflected everywhere the information appears. We can generate role-specific "cheat sheets" for acolytes, altar guild, or clergy with a few clicks.

Pitfalls: The initial decomposition requires careful human oversight to ensure liturgical logic is preserved. AI sometimes creates categories that make technical sense but don't reflect how ministry actually works.

Use Case 9: Internal Systems Automation (Room Reservations)

What we're doing: Using AI to develop custom automation tools like a room reservation system using Google Apps Scripts connected to our shared calendar, reducing redundant data entry and eliminating the need for costly third-party services.

Benefits: We eliminated a $600/year subscription to a third-party room reservation system and reduced double-booking conflicts. The system integrates seamlessly with our existing calendar infrastructure.

Pitfalls: Custom automation requires ongoing maintenance. When Google updates its APIs or our needs change, someone needs to understand the system well enough to maintain it. We've learned to document AI-generated code extensively.

Use Case 10: Visual Design Integration (Calendar Skins)

What we're doing: Creating custom design elements like calendar skins and interface components that integrate with our website design and branding.

Benefits: Visual consistency across our digital platforms without hiring a graphic designer. AI can generate CSS that matches our existing branding and create variations for testing.

Pitfalls: AI-generated design sometimes looks generically "professional" rather than distinctively ours. Human aesthetic judgment is still essential.

Category 4: Process Improvement

Use Case 11: Meeting and Communications Protocol Development

What we're doing: Using AI to brainstorm and draft streamlined protocols for meetings, internal communications, and collaborative workflows.

Benefits: AI can suggest structures and best practices from organizational development literature that we might not have encountered otherwise. It's particularly good at identifying inefficiencies in existing processes.

Pitfalls: AI-suggested protocols sometimes assume organizational contexts that don't match parish life. A protocol that works in a corporation might be completely inappropriate for a church vestry. All suggestions require significant adaptation. At the same time, interestingly, AI has also sometimes offered suggestions or cautions about adopting secular business methods too woodenly, noting where a parish may need to adapt or reject

certain methods, language, or principles when a human would have simply adopted them whole cloth.

Use Case 12: AI Governance Policy Creation

What we're doing: Using AI to help develop comprehensive AI use policies—essentially, using AI to regulate itself.

Benefits: AI can identify edge cases and potential problems we might not anticipate. It can also review policies for internal consistency and completeness.

Pitfalls: The irony of using AI to write AI policies is not lost on us. Human theological and pastoral wisdom must drive these policies; AI can only assist with structure and comprehensiveness.

Use Case 13: Ministry Training Program Design

What we're doing: Using AI to suggest training protocols, curricula, and processes for various ministry areas including volunteer development and specialized ministry skills.

Benefits: AI can outline comprehensive training programs, suggest learning objectives, and propose assessment methods based on adult education best practices. It can also do "red team" audits that anticipate potential volunteer confusion and suggest changes to better fit human implementation, especially if the people drafting the training are not experts in pedagogy.

Pitfalls: Training materials must begin with human-written foundational drafts that reflect our actual parish culture and theological commitments. AI-generated training can feel generic or fail to address parish-specific dynamics.

Use Case 14: Comparative Analysis (Customaries and Liturgical Documents)

What we're doing: Using AI to analyze differences between historical liturgical texts and our current practices—comparing our customary with sources like Urquhart's historic Sarum reconstruction, Dearmer's liturgical works, Fortescue's texts, and similar scholarly liturgical resources.

Benefits: As we are revising and finalizing our customary, AI helped us quickly compare our practices with historical liturgical traditions, identifying how our contemporary practice relates to these rich historical sources. AI excels at this kind of detailed textual comparison across multiple documents, work that would be extremely time-consuming manually. This informed our revision process significantly and helped us understand the historical foundations of our current practices.

Pitfalls: AI sometimes treats liturgical differences as merely procedural when they may reflect important theological commitments. Human liturgical knowledge is essential for interpreting the analysis.

Use Case 15: Meta-Documentation

What we're doing: Using AI to organize and present documentation about AI use itself—including this very report.

Benefits: AI is helpful for structuring complex information, ensuring comprehensive coverage of topics, and maintaining consistency in documentation.

Pitfalls: The same as any AI use: human oversight is essential, and the human author bears full responsibility for accuracy and appropriateness.

Use Case 16: Abbreviated Cheat Sheets from Style Guides

What we're doing: Using AI to extract relevant portions from comprehensive style guides and governance documents to create abbreviated, role-specific cheat sheets that are easier for volunteers to understand and use.

Benefits: Our comprehensive style guide and organizational instructions are thorough but can be overwhelming for volunteers who only need a subset of the information. AI can quickly extract the relevant sections—for instance, creating a "Bulletin Preparer's Quick Reference" from our full editorial style guide, or an "Usher Essentials" card from our complete worship customary. This makes critical information more accessible without requiring volunteers to wade through lengthy documents. The AI maintains consistency with the source documents while presenting information in a more digestible format.

Pitfalls: AI sometimes oversimplifies in ways that create ambiguity or leaves out important context. A cheat sheet that says, "arrive 15 minutes early" might omit the crucial detail "except for Christmas and Easter services, when ushers should arrive 30 minutes early." It's necessary to have someone with deep knowledge of the source material review cheat sheets to ensure critical nuances aren't lost in the abbreviation process.

Part Two: Where We Will Not Use AI—The Red Lines

As helpful as AI has been in operational and administrative contexts, there are areas of ministry where we believe AI use is fundamentally inappropriate. These boundaries are not arbitrary but reflect our theological convictions about the nature of ministry, the importance of human presence, and what can and cannot be mediated through algorithmic processes.

Red Line 1: Sermon Preparation and Composition

The boundary: AI may not be used to draft, compose, or substantially write sermon content.

Why this matters: Preaching is an inherently personal and pastoral act. It requires the preacher's authentic engagement with Scripture, tradition, reason, and experience—including their own spiritual life and their knowledge of the congregation before them. A sermon is not merely information delivery, but a word spoken from one person's faith and authority of office to interpret scripture in *lived* contexts. This cannot be delegated to an algorithm.

AI certainly can draft sermons (though I have not done this because this is such a bright red line). It can assemble biblical references, quote theologians, and construct logical arguments, but it cannot speak from embodied existence to embodied existence. The congregation deserves better.

What we do instead: Preachers engage Scripture, theological sources, and the congregation's context themselves. AI might function, under the rigorous guidelines of the AI research assistant, to help locate a patristic citation or find a theological resource, but the sermon itself must be authentically the preacher's work. And again, all sources must be rigorously cited and verified.

Red Line 2: Theological Statement Drafting

The boundary: AI may not be used to draft or compose theological statements, position papers, doctrinal declarations, or theological arguments.

Permitted exception: AI may be used for style editing and, under the rigorous guidelines of the AI research assistant, pre-composition research bibliography compilation.

Why this matters: Theological work requires not just logical consistency but spiritual discernment, tradition engagement, and awareness of how language shapes belief. AI can mimic theological language without understanding theological meaning. It has no tradition to stand in and no community to be accountable to.

What we do instead: Theologians write theology. AI can help with formatting, style consistency, functioning as a glorified search engine, or bibliography compilation, but the theological content must be human-authored and theologically accountable.

Red Line 3: Sacramental Preparation, Administration, and Documentation

The boundary: AI may not be used to compose or generate sacramental liturgical texts, administer sacraments, or create sacramental records.

Permitted exceptions: AI may be used for:

- Bibliographic compilation on liturgical sources and sacramental theology
- Comparing inputted liturgical documents (e.g., analyzing differences between rites)
- Formatting sacramental service bulletins (visual design, layout, spacing, fonts—NOT content creation or modification)
- Refining inserted liturgical directions or rubrics in bulletins (e.g., "The congregation stands" or "Please remain seated until dismissed by the usher")

AI may NOT be used to compose, generate, or substantially modify actual liturgical texts (prayers, blessings, scriptural readings, sacramental formulae). When AI analyzes bulletins for formatting purposes, it processes visual design elements only; any liturgical or prayer content within those bulletins must be human-authored and remains off-limits for AI modification.

Why this matters: The sacraments require embodiment. There is something profoundly wrong with outsourcing any dimension of sacramental life to artificial intelligence. The very phrase "AI-assisted baptism" or "algorithm-generated Eucharistic prayer" reveals the category error.

Sacramental ministry requires embodied human presence, spiritual authority derived from ordination and the community's recognition, and personal accountability to God and the church. AI has none of these.

However, AI can appropriately assist with research about sacramental theology, comparison of liturgical documents, and formatting of worship materials—just as it can with other theological and administrative work. The distinction is between creating sacramental content (prohibited) and supporting the practical preparation for sacramental worship (permitted with appropriate oversight).

What we do instead: Clergy prepare for and administer sacraments without AI assistance in composing liturgical texts or generating sacramental content. Sacramental records are created and maintained by human persons who understand the spiritual and legal significance of what they're documenting. AI may assist with research, document comparison, and bulletin formatting, but the sacramental texts themselves and the acts of worship remain entirely human-mediated.

Red Line 4: Prayer Composition

The boundary: AI may not be used to draft or compose prayers for public worship, private devotion, or pastoral care contexts.

Permitted exception: AI may assist in creating educational resources, such as guides on "How to Write Prayers of the People."

Why this matters: Certainly, we use pre-written prayers all the time: this is foundational for our tradtion. But those were all written by human beings who actually prayed those prayers. AI can generate words that sound like prayer, but it cannot pray. To use AI-generated prayers in worship or pastoral care is to substitute simulation for authenticity.

What we do instead: People pray their own prayers. When learning to write liturgical prayers, people might use guides composed with AI-asssistance (for things like extracting the normal length of petitions in the BCP model Prayers of the People), but the prayers themselves must be authentically theirs.

Red Line 5: Pastoral Counseling

The boundary: AI may not be used in pastoral counseling sessions or relationships.

Permitted exception: AI may compile research bibliographies on pastoral care topics.

Why this matters: Pastoral counseling requires human presence, empathy, spiritual discernment, appropriate professional boundaries, and pastoral confidentiality. These cannot be mediated through AI. The pastoral relationship is built on trust, and that trust requires a human person who is spiritually and ethically accountable.

There are AI chatbots marketed as "spiritual companions" or "pastoral care assistants." These are not merely inadequate; they are pastoral malpractice. They simulate care without actually caring, simulate understanding without actually understanding, and can cause real harm to vulnerable people.

What we do instead: Pastoral care is provided by trained human persons who are accountable to professional standards and spiritual disciplines. AI might help locate resources on specific pastoral care challenges, but the care itself is entirely human.

Red Line 6: Research Report Writing and Theological Synthesis

The boundary: While AI may compile research bibliographies and initial possible further avenues for research, it may not be used to write research reports for publication or produce un-reviewed synthesized findings.

Why this matters: Synthesizing research requires critical judgment about what sources are reliable, how they relate to each other, what claims are well-supported versus speculative, and how findings apply to our context. AI can assemble information and offer possible avenues of interpretation but cannot finally exercise scholarly judgment.

What we do instead: Human scholars read, analyze, synthesize, and interpret sources. AI can help locate sources and even provide initial plausible synthesis, but humans do the actual final research and writing.

Part Three: The Guardrails—How We Govern AI Use

Beyond the absolute prohibitions, we've developed four essential guardrails that govern all permitted AI use:

Guardrail 1: Full Human Ownership and Responsibility

The rule: Everything generated with AI assistance is fully owned by the person generating it. The human user bears complete responsibility for accuracy, appropriateness, and consequences.

Why this matters: There's a temptation to treat AI errors as "the AI's fault" rather than accepting responsibility. We reject this. If I use AI to draft a policy and that policy contains an error, the error is my responsibility, not the AI's. I chose to use the tool; I chose to accept its output; I bear the consequences.

This also means we cannot blame AI for inappropriate content. If AI suggests something offensive or theologically problematic, the fault lies with the human who approved and used it.

In practice: Every document created with AI assistance is attributed to a human author who takes full responsibility for its content. We do not add disclaimers like "AI-assisted" to shift responsibility.

Guardrail 2: Mandatory Human Review Before Publication

The rule: All AI outputs must receive thorough human review before publication, implementation, or distribution.

Why this matters: AI makes plausible-sounding errors. It generates text that reads smoothly but contains factual mistakes, logical inconsistencies, or inappropriate content. Only human review can catch these problems.

I've had instances where AI-generated text seemed perfect until reviewing to find significant errors. The text was grammatically flawless but substantively wrong.

In practice: Nothing generated by AI goes public without human review by someone with expertise in the subject matter. This review is not cursory proofreading but substantive evaluation.

The iterative refinement advantage: One significant benefit of AI-assisted document creation is that corrections identified during human review can often be integrated back into the document through natural language input. Rather than manually editing files, reviewers can describe needed changes conversationally, and AI can implement them efficiently.

However, this creates a new challenge: ensuring that AI implements only the requested changes without making other unauthorized modifications—what we call "drift." A seemingly simple instruction like "fix the grammatical errors in section 3" might result in AI also rewording content, changing formatting, or "improving" language in ways that alter meaning.

Our solution: Clear parameters for patch integration and mandatory drift auditing. When feeding review corrections back to AI, we:

1. **Use specific, bounded instructions:** "In paragraph 2, change 'clergy' to 'ordained ministers' and make no other changes" rather than "improve the language in this section"
2. **Request explicit change documentation:** Ask AI to list exactly what changes it made
3. **Conduct zero-drift audits:** Compare before and after versions to verify only authorized changes were made
4. **Reject outputs with unauthorized drift:** If AI made changes beyond what was requested, we revert and try again with clearer constraints

This iterative refinement capability is powerful but requires disciplined oversight to prevent scope creep in AI modifications.

Guardrail 3: Human-Written Foundational Drafts for Communications and Training

The rule: All communications materials and training resources must begin with a foundational draft written by a human author.

Why this matters: AI should enhance human work, not replace human authorship. When we use AI to generate communications or training materials from scratch, the results are generic and fail to reflect our parish's actual culture, theology, and voice.

In practice: Someone writes a draft—even a rough outline—that reflects our parish's perspective and priorities. Then AI might help expand, refine, or format it. But the human foundation comes first.

Guardrail 4: Rigorous Review for Policies and Legal Documents

The rule: Any policy, contract, or legal document created with AI assistance must undergo both comprehensive human legal and stakeholder review before adoption.

Why this matters: AI doesn't understand the legal implications of language choices, the political dynamics of policy implementation, or the potential unintended consequences of seemingly straightforward rules. Policies that look reasonable in AI-generated draft form can create serious problems when implemented.

In practice: AI-drafted policies are treated as "very rough first drafts requiring extensive revision" rather than "nearly finished documents needing minor tweaking." Both parish leadership and, when appropriate, legal counsel review them before adoption.

Conclusion: Technology in Service of Ministry

This experience has convinced me that AI can be a valuable tool for church administration and operations—but only when used within clear boundaries and with rigorous governance.

Not all tools are morally neutral—some technologies lend themselves more readily to virtuous uses, others to vicious ones. With proper guardrails, we believe Large Language Models can be useful tools for serving God's purposes in the world without compromising our theological and ministerial integrity. The boundaries we've established are not technophobic resistance to innovation but theological discernment about what can and cannot be appropriately mediated through algorithmic processes.

Some ministry tasks are genuinely administrative and can benefit from AI assistance. Others are inherently pastoral, sacramental, or theological and must remain fully human. Wisdom lies in knowing the difference.

I offer these reflections not as definitive answers but as one parish's experience navigating new territory. Other contexts will require different applications and perhaps different boundaries. But I believe these fundamental principles should guide any church's engagement with AI:

1. **AI should serve ministry, not replace ministers.** It's a tool, not a solution.
2. **Some things are too important to delegate to AI.** Preaching, sacraments, prayer, pastoral care, and theological work require human presence and spiritual authority.
3. **Human responsibility cannot be outsourced.** If we use AI, we own the results—for better or worse.

4. **Guardrails are essential.** Without rigorous human oversight, AI use will drift into inappropriate territory.

5. **The technology will change; the principles won't.** These boundaries are rooted in our theology of ministry, not in the current state of technology.

My prayer is that the church will engage AI with both creative openness and theological discernment—using it where appropriate, refusing it where necessary, and always keeping our focus on the mission we've been given: to restore all people to unity with God and each other in Christ.

The Rev. Dr. Christopher Corbin
Rector, Trinity Episcopal Church
Oshkosh, Wisconsin

Available for consultation: frchris@trinityosh.org