Unitarian Universalist Church of Brunswick Maine

This guide is a modified and adapted for UUCB with permission from the Unitarian Church of Evanston

AI Use Toolkit

Introduction: This guide grows out of a working group from the UUCB Communications committee that convened in April 2026 a draft of this toolkit from the Unitarian Church of Evanston as well as the resources listed at the end of this document. It is meant to help people discern about their use of generative AI with a lens toward practicality, ethical considerations, and transparency. It is not a policy document or an official document in any capacity; it is a discernment and practice tool. This should be considered a living document and revisited/revised on a regular basis.

Core takeaway: Use AI as a tool, not a substitute for judgment, relationships, responsibility, or growth.

Framing:

1. AI is a powerful technology and will be an enduring part of the innovation landscape for the foreseeable future.
2. Harnessing the power of AI brings the responsibility to engage with it in a mindful and principled way, centering our values and mission.
3. AI can support human intelligence and capabilities, rather than replace them.
4. We have a responsibility to our community to provide ethical lenses with which to engage this emerging technology.*

*This framing is based loosely on Arizona State University's Guiding Tenets for AI. See the reference section for the primary resource.

Table of Contents

1. What AI is and is not good at
2. Sample AI Application Assessment Matrix
3. Disclosure and transparency
4. Prompting and iteration
5. Evaluating AI output
References

1. What AI is and is not good at

AI is generally good for tasks that involve translating or summarizing things, especially large datasets (such as large reports, copious notes, etc.) As of early 2026, AI does not “think,” use real logic, or provide truth - it is a pattern recognition and regurgitation tool.

Often Useful for	Use Caution or Avoid for
Brainstorming	Human-like care or pastoral presence
Summarizing public material	Confidential matters
First drafts	High-stakes decisions
Rewording for tone, length, or audience	Ethical decision making
Organizing ideas	Deep expertise without verification
Acting as a repository of information	Private information, including IP and PII

A quick appropriateness check:

- Does this task involve ethical judgment, deep expertise, or pastoral sensitivity?
- Will the result be public or used in an important decision?

If the answer to either is “yes” or “I’m not sure,” pause. That may be a sign to avoid AI here, narrow the task, or increase human review.

- Can I clearly define the task I want help with?
- Can I verify the output?
- Is this aligned with our values, boundaries, and community expectations?

If the answer to any is “no” or “I’m not sure,” pause. That may be a sign to avoid AI here, narrow the task, or increase human review.

2. Sample AI Application Assessment Matrix

The matrix below is based on the NIH AI Toolkit. Use this kind of matrix to help in your discernment around AI use. Some examples are provided, but you must discern whether the application you are considering is appropriate for your particular context.

A healthy guideline: if a task touches trust, confidentiality, care, or high-stakes decisions, slow down. Human judgment should lead.

Application	Risk Level	Boundaries	Transparency & Attribution	Monitoring
Brainstorming ideas	Low	Use for initial ideas only	Usually no formal attribution needed	Human review and selection
Planning an event or generating options	Low	Use for logistics, themes, and draft lists/items. Humans decide.	Usually no formal attribution needed	Review for realism, cost, accuracy (including dates), and congregational fit
Drafting announcements or newsletters	Low–Med	No confidential details; edit for voice and accuracy	Disclose if AI significantly shaped final text	Check facts, dates, tone, fit
Summarizing a public article, policy, or legal proposal	Low-Med	Use public materials only; verify against source	Disclose in handouts when use is significant	Cross-check with original
Brainstorming worship or teaching	Medium	Do not outsource discernment, theology, or meaning-making	Disclose if AI significantly shaped the product	Revise for authenticity and values
Meeting agenda or public-facing handout	Medium	No sensitive or identifying information	Consider disclosure when AI materially shaped structure or content. Identify any AI-generated images.	Human review before circulation
Pastoral care / sensitive conflict	Med-High	No confidential input. Do not outsource care.	If used, must disclose the nature of the use, consent as applicable, and awareness of potential risks.	Human-only care and judgment
Personnel or legal decisions	High	Avoid AI for decisions. At most, use for neutral drafting with no private details.	Use only with strong oversight.	Human expertise leads throughout

3. Disclosure and transparency

Disclosure is about honesty when AI played a substantial role in the final content or process, or when failure to disclose could damage trust.

Disclosure is recommended when:

- AI contributed substantially to the final content, framing, structure, or wording, not just to early brainstorming, organization, or tone
- The content will be shared publicly, presented, or published
- The material is largely AI-generated and only lightly edited by the human user
- A policy, covenant, or community expectation calls for disclosure
- Non-disclosure could weaken trust or create a misleading impression

A simple disclosure formula

[Tool] was used to assist with [task] in preparing this [type of work]. The output was reviewed, revised, and adapted by [person or group].

What disclosure might sound like

- “ChatGPT was used in early brainstorming for this project.”
- “Claude was used to help organize ideas during preparation.”
- “Gemini was used to summarize public materials; the summary was checked against the original sources.”
- “Generative AI was used as a brainstorming and drafting aid. Final editing, framing, and responsibility remained with the author.”
- “This image was generated with AI.”

4. Prompting and iteration

Prompting is a key skill in effective AI use. This prompt template can serve as a starting point

> Act as a [**role**]. Help me [**task**].
> Context: [**background**].
> Audience: [**who it is for**].
> Constraints: [**tone**, **length**, **values**, **do / don’t include**].
> Please give me [**output format**].

The first answer is rarely the final answer. Iteration is also a core skill. Here are some examples of ways you might iterate after an LLM output.

Ways to iterate	Example Follow-up Prompts
Tone and clarity	Make this warmer. Make it clearer for beginners. Make it less formal.
Length and audience fit	Shorten this to 120 words. Rewrite for skeptical readers. Check for appropriateness for a diverse audience.
Specificity and structure	Make this more concrete. Reorganize this into three clear points.
Balance and bias	What assumptions are being made? Provide a more balanced perspective, with pros and cons of this discussion.
Actionability	Turn this into a checklist, discussion questions, or next steps.

5. Evaluating AI output

Hallucinations: All tested Generative AI algorithms, particularly large language models (LLMs) are known for producing hallucinations - outputs of false, misleading, or fabricated information that are typically presented in a confident and fluent manner. Users should look out for hallucinations such as inventing sources or quotes, fabricating statistics or dates, and other inaccurate outputs. Hallucinations fall into four broad categories:

1. Factual (incorrect information)
2. Contextual (breaking logical flow of conversation or irrelevant output)
3. Logical (self-contradictory or logically inconsistent outputs)
4. Multimodal hallucination (mismatch across modalities such as incorrect text descriptions of audio or visual sources)

Sycophancy: Generative AI is also known to be sycophantic and give users outputs that are appealing to them, regardless of the validity of the output. Users are advised to be cautious and engage critically with AI output so as to avoid confirmation bias and other associated traps.

RACCCA is a user-focused framework to evaluate generative AI outputs developed by Dr. Andrew Maynard at Arizona State University.

Relevance	Does it directly address the prompt? Is important context missing? Is there a tone mismatch?
Accuracy	Are facts, dates, claims, and sources correct? Are there made-up quotes or stereotypes?
Completeness	Does it cover the essential aspects of the question? Are perspectives or details missing?
Clarity	Is the response easy to understand and free of unnecessary ambiguity?
Coherence	Does it make sense? Is it logical, well-structured, and internally consistent?
Appropriateness	Is the tone, content, and format suitable for the setting and audience?

References and Resources

Discussion Resources

This set of Questions and Issues for AI Community Conversations - Peter Bowden may be a helpful resource for continuing discussions on ethical issues related to AI. A resource guide for how to host these conversations is also available here: April 14, 2026 UU Growth Lab Live

References

The following references were used in generating some of the guidelines in this document. This document should be considered a living document and revisited on a regular basis. None of the suggestions in this document should be considered legally binding or universal; the end user is ultimately responsible for the ways in which AI is used and applied based on their particular context.

"Artificial Intelligence Guiding Tenets." AI Tools. Accessed April 2026. https://ai.asu.edu/.

"Digital Trust Guidelines for Generative Artificial Intelligence Use." Digital Trust | Artificial Intelligence. Accessed April 2026. https://ai.asu.edu/digital-trust-guidelines.

"NIH Generative AI Usage Toolkit." Generative AI Usage Toolkit, January 17, 2025. https://irp.nih.gov/system/files/media/file/2025-03/nih_library-genai-toolkit.pdf.

Bible with Life. "How to Create a Church AI Policy." Bible with Life, April 14, 2026. https://biblewithlife.com/how-to-create-church-ai-policy?srsltid=AfmBOoopLm3h5ki5__8xLr58FsI8mm3-Y1B9XUwHGZFm4n6Cy5zPFVI_.

Kranda, Jay. “Things to Know When Using AI at Saddleback Church.” Beyond Just Streaming, June 2, 2025. Accessed from https://www.jaykranda.com/blog/2025/6/2/sample-church-ai-policy-amp-training-you-can-copy; Accessed online April 2026

Maynard, Andrew. “Evaluating Prompt and Response.” Andrew Maynard, January 19, 2024. https://andrewmaynard.net/prompt-and-response-evaluation/#:~:text=The%20RACCCA%20framework%20is%20a%20framework%20for,Compare%20different%20versions%20of%20a%20similar%20prompt.

Rivera, Analiese. “Startchurch Blog - How Churches Can Build a Responsible AI Policy.” StartCHURCH Blog - How Churches Can Build a Responsible AI Policy. Accessed April 2026. https://www.startchurch.com/blog/view/name/how-churches-can-build-a-responsible-ai-policy.

Draft disclosure

In preparing the AI workshops on which this document was based, generative AI was used as a brainstorming and organizing tool, as well as for help locating some of the external resources. Final editing, framing, and responsibility remained with the authors.

Contributing Authors

Emma K. Farrell, Ph.D.

With special thanks to Adam Gough, Rev Eileen Wiviott, and participants of the Unitarian Church of Evanston AI workshop series

Adapted for the Unitarian Universalist Church of Brunswick

A working group of communications committee members including Rev. Dr. Kharma Amos, Jud Caswell, Patrick Kirby and John O’Connor adapted this resource from the Unitarian Church of Evanston - it will be reviewed at the May 2026 Communications Committee Meeting