DIOCESE OF GLOUCESTER

Communications and Engagement

Guidance on the Responsible Use of Artificial Intelligence (AI)

"AI's growth needs to be guided by robust legal frameworks, independent oversight, informed users and a political system that does not abdicate its responsibility. A more moral AI is not enough if that morality is determined by a few."

~ Pope Leo

"We are not simply instruments of production but living beings, entrusted with a moral compass. In a world full of imbalances, we have a duty to use AI responsibly."

~ Christine Allen

Part one: written by a human
Part two: the same document, written by an AI

Part one

Purpose

This guidance document will outline the principles and practices for ethical, secure, and sustainable use of AI tools within the Gloucester Diocesan Board of Finance. Our commitment is to transparency, inclusivity, and environmental stewardship.

1. Introduction to AI and best practices

AI processes can enhance our work by automating routine tasks, supporting creative projects, and exploring data sets to aid decision-making. However, AI operates in profoundly different ways to traditional computing, and responsible use of them is essential.

- AI is a *language* calculator, not a *numbers* calculator; that distinction means it will lie, dream and mislead in ways that traditional number calculators cannot.
- Team collaboration: though we may identify AI opportunities collaboratively; we also need to keep an open mind on our AI usage – so we can continue to explore both the potentials and the risks of the tools on offer.
- Complement, do not replace: use AI to assist human tasks — without replacing human reasoning, oversight or creativity.
- Verification: although AI can generate drafts or summaries, human judgement is *essential* to ensure accuracy, relevance, and tone.
- Awareness and training: recognise how AI can add value, such as summarising fragmented notes or exploring data, but remain critical of outputs. Staff awareness needs to include understanding of the very common risks of AI Hallucinations. AI output [to date, Q2 2026] typically tells you what it calculates you want to hear, which has serious implications, particularly when exploring data used to guide decision-making.
- Some practical examples of AI tools are listed below, in section 8.

2. Data Privacy and Security

- One of the few absolutes in this guidance: **do not upload** sensitive or confidential information into any AI tool unless that tool has been explicitly approved by the organisation (ie, if in the future the GDBF operates an AI in a secure DMZ).
- Avoid uploading large data archives of any sort for summary unless using approved tools (see part 6).
- Adhere to GDPR and internal data governance policies. AIs connected to the internet (ie. all of them) are not to be considered secure, particularly when handling personal data such as names, addresses, mobile phone numbers etc.
- All AI-generated content or code must be human-reviewed before use, checking all linked references. Content confidently produced is as trustworthy as an anecdote overheard on a bus; it *could* be good, but it's a hit-and-miss system.

3. Environmental Impact Guidelines

AI systems have a significant carbon footprint. AI apps offload the real computation from mobile phones to data centres – racks of servers – so it's a hidden impact. Some lightweight AIs can run on laptops ('Edge Computing'). While at this stage we need to encourage experimentation, teams should also consider usage in terms of environmental impact.

- ChatGPT alone consumes more energy daily than 18,500 households.
- A single AI text query uses over 10x the energy of a Google search: ~0.003 kWh (11 kJ). If we get into the habit of using AI, as we did with using Google, we are going to see significant impacts on survivability.
- Writing a 100-word email using an AI consumes 500ml of fresh water, and uses enough electricity for 7 full charges of an iPhone Pro Max. Image production is even more intensive.
- Training a model like OpenAI's GPT-3 consumed 5.4m litres of water, and 1,287 MWh—equal to the annual electricity consumption of 120 homes. That energy is quickly exceeded by ongoing usage.
- Data centres in the UK will soon account for 20% of our national electricity demands.

- **To use AI more sustainably:**
    1. Write clear, precise prompts.

2. Cut the filler (save the small talk).
3. Generate only when necessary. If it's 'always necessary' to use AI because someone only has time to do their job by using it, then the role is not sustainable nor properly supported, and the organisation needs to be informed and respond appropriately.
4. Avoid peak-hour usage.
5. Once an established prompt works for a typical task – save that prompt as a template (this could simply be saved in a Word document). Templating prompts reduces negotiation/reprompting and directly saves energy consumption.
6. Teams should monitor basic usage and then consider offsetting prompts, eg. plant a tree every 100 prompts, or invest proportions of budgets in renewable energy.

4. Ethical and Responsible Use

- AI can be a support tool — not a decision-maker.
- Be alert to bias: AIs are trained on human data which is inherently biased. Review outputs for discriminatory content, including subtle biases of gender, race, religion or ideology.
- Do not use AI to write sermons, as it may propagate skewed theological content from questionable sources. Theology is easy enough to misunderstand, without the added inaccuracy of language calculation; religious abuse is a serious risk. If a preacher cannot be bothered to write content, humans should not be expected to listen to it.
- The purpose of an AI is to help us in our tasks, not to replace us entirely. A good use of an AI process, would be to note down lots of key thoughts, references and the basic structure of a talk, and then ask the AI to bring the content together into a cohesive fifteen-minute talk, and then carefully edit the output. A bad use of AI would be to say 'write my talk for Sunday' and to read out whatever it regurgitates.
- Acknowledge that AI will confidently present inaccurate or ambiguous content. Hallucinations are a serious consideration and hunting hallucinations needs to feature heavily in any verification or quality processes.
- Separate validation data from training data in model development. This forms later discussions around DMZ-housed AIs, when the CofE looks to house its own verified AI – but that will be some years down the line.

- In the content we produce, we need to specifically monitor for unintended consequences such as intellectual property violations, bias, or unethical outputs. Supported staff training on each topic should align with general CPD.

5. Transparency and trust

- We all need to be open about our AI usage with our teams. Experimentation at this stage should be encouraged. We will learn together as the tools develop.
- Teams should consider maintaining brief documentation of model performance ~ which tools worked well for certain tasks, where common problems are found, and where potential could lie. Teams could then compare notes and broader ideas of how we use these technologies together.
- If you would like to pay for an AI tool, discuss it with your manager first, and have the tool paid for via the normal IT channels.
- See part 7: Governance and Compliance. The EU AI Act will be increasingly relevant.

6. Ethical AI tools

How we use our tools is important, but it is also worth considering that the specific tools we use, are also very different. Many AIs have been purposefully trained on illegal material in order to produce further illegal material, and are freely available in the public sphere. The following tools have demonstrated strong ethical and/or privacy practices:

- Adobe – Recognised for its ethical AI practices.
- Apple – Privacy-focused, including differential privacy and on-device processing (which relates to questions around environmental impact via Edge Computing).
- Google/Gemini – Focus on fairness, transparency, and 'red line' ethical boundaries which includes basics like surveillance and weaponry.
- IBM – Leading in transparent and robust AI practices, including 'explainability'. Offers the watsonx.governance platform.
- Microsoft Copilot – Session-based, encrypted, and considered secure.

7. Governance and Compliance

- GDBF should consider regular internal audits, and may at some point need to commission external reviews of AI systems.

- GDBF should consider writing up a formal Policy in order that violations can result in disciplinary actions. The same sort of risks apply with data storage, network use, and internet use and could simply be added to employment contracts/handbook in brief.
- This sort of content will need to evolve in response to new regulations, technologies, and ethical developments – currently changing month-by-month.
- The EU AI Act should be adhered to as it develops. For example, from 2026, prohibited activities include biometric categorisation systems used to sort people into groups based on politics, religion, sexual orientation, or race. Also under the spotlight will be systems with the potential to cause significant potential harm to health, safety, fundamental rights, environment, democracy, and the rule of law. So for example, AI systems which help employers choose candidates applying for jobs will need to be specially scrutinised rather than applied slap-dash by companies wanting to make recruitment 'more efficient'. The EU AI act is about transparency and basic ethics regarding the creation of tools, which could reasonably affect everything from voting patterns to employment bias. Though the UK is not currently within the EU, as the largest legal entity in the world, the EU does affect global policies of all sorts. Many UK organisations are already EU AI Law compliant.
- ISO 27001 does not yet mandate AI policies, but our framework should include policies and controls under Control 02-28 regarding the responsible use of AI ~ particularly when it comes to national or local DMZ -run AI tools. In short; build with transparency / explainability.

8. Example AIs

- Image-creation capabilities. Eg Adobe AI, Midjourney, Apple Image Playground.
- Audio expansion tools; pieces of music can be constructed from a single sound, or a speech built from someone's short voice recording. Eg. Opus, Loudly, Elevenlabs.
- Videos, websites or computer games generated from written descriptions. Eg. Kling, Sora, Rosebud.
- Code or pattern development in live discussion with an AI, offering snippets or entirely rebuilt code architectures. Eg. Gemini, Lovable.dev.

- Automation of email tasks – writing to customers, responding to clients, recording conversations in spreadsheets or creating reports based on conversations. Eg. Google Studio, Make. Begin with a template.
- Custom AI tools for a specific task; everything from medical diagnosis, to sales strategy. Eg. OpenAI trained on a specific network of data to help spot patterns.
- Google AI Studio can watch your screen and provide live tutorials on how to solve problems, talking you through any difficult tasks while you're working on them.
- Transcription, subtitling and minuting tools are very powerful, including focusing on what's said by a single participant. Eg. Claude, Otter, Final Cut Pro.
- Disparate raw data can be uploaded, and the AI used to help explore the data, explaining what it means or how to present it using infographics. There are many examples, but only Copilot and Gemini are to be considered viable options currently. **Do not upload *personal* data of any sort, to any AI.**
- Abstract cross-media applications are also possible, so geographic data could be used to write campaigns, or a film could be used to *help* write a sermon. Tools may also be connected, so auto-filled spreadsheets used to run social media campaigns, or web content used to schedule adverts. Talk to Copilot for advice on how to get started with this.

Part two

(the same document, but written by an AI)

1. Purpose

This guidance outlines the acceptable use of Artificial Intelligence (AI) technologies within the GDBF. The goal is to ensure responsible, and effective use of AI while safeguarding the organization's values, data, and stakeholders. ('acceptable use' and 'safeguard the stakeholders' is all very corporate … guidance should be about equipping people)

2. Scope

This guidance applies to all employees, contractors, and third-party partners who utilize AI technologies in any capacity within [Organization Name]. (… repetition of the first sentence)

3. Definitions

- **Artificial Intelligence (AI):** Technologies that simulate human intelligence processes, including but not limited to machine learning, natural language processing, and robotics.
- **User:** Any individual accessing or utilizing AI technologies within the organization.

(Offers definitions, but without mention of the associated data centres, Large Language Models, hallucinations, AI training or edge computing concepts. The only definitions are, "AI means AI, user means user")

4. Acceptable Use

- **Compliance:** All AI applications must comply with relevant laws, regulations, and industry standards.

- **Data Privacy:** Users must ensure that any data used for AI training or processing is obtained legally and ethically, respecting privacy rights and data protection regulations (e.g., GDPR, CCPA).
- **Transparency:** Users should maintain transparency in AI processes, ensuring stakeholders understand how AI decisions are made.
- **Bias Mitigation:** Users are responsible for actively identifying and mitigating biases in AI models.

(More slop. Vague references to important topics, phrased in a way that keeps them ethereal and out of reach)

5. Prohibited Use

- **Unauthorized Access:** Users must not access AI systems or data without proper authorization. (Which systems? What data? Authorisation of what from whom?)
- **Discrimination:** AI applications must not perpetuate or exacerbate discrimination based on race, gender, religion, or other protected characteristics. ("Don't be bad")
- **Malicious Use:** The use of AI for malicious purposes, including but not limited to harassment, fraud, or cyberattacks, is strictly prohibited. ("Things which are against the law, are against the law")

6. Data Management

- **Data Security:** Users must adhere to data security protocols to protect sensitive information processed by AI systems.
- **Data Sharing:** Sharing data with third parties for AI purposes must be approved by the Data Protection Officer and must adhere to data sharing agreements.

7. Accountability and Responsibility

- **Training:** Users must complete training on the ethical use of AI technologies and the specifics of this guidance. (Which specifics? There are none)
- **Reporting Violations:** Users are encouraged to report any violations of this guidance or concerns regarding AI use to their manager or the Ethics Committee. ("Spy on your colleagues, win prizes")

8. Monitoring and Enforcement

- [Organization Name] reserves the right to monitor AI usage to ensure compliance with this policy. Violations may result in disciplinary action, up to and including termination. ("We will spy on you, and you could be killed")

9. Guidance Review

This guidance will be reviewed annually or as necessary to ensure it remains relevant and effective. Updates will be communicated to all users. ("This could change at any time, but we'll let you know" ~ otiose section)

10. Contact Information

For questions or concerns regarding this guidance, please contact: [Name] [Title] [Email] [Phone Number]

(Verdict: Without a proper detailed briefing and edit process, an AI will produce content which looks good at first glance, but is in fact pointlessly vague. We need to calibrate *appropriate* use of AI.)