

Understanding and Using AI Responsibly

A guide for mission-driven organizations on how large language models work, where they help and fall short, and how to choose the right level of investment in AI tools.

1. What Is a Large Language Model, and Why Does It Matter?

A large language model, or LLM, is the technology behind tools like ChatGPT, Claude, Gemini, and Copilot. It is a type of artificial intelligence trained to predict and generate human-like text by studying enormous amounts of existing writing.

Think of an LLM as a very well-read assistant who has skimmed a huge share of everything ever published online, but who was never trained to fact-check what it read, and who cannot tell you where a specific idea came from. It generates responses by recognizing patterns in language, not by understanding or reasoning about the world the way a person does.

How an LLM Is Built

LLMs are trained on massive collections of text scraped and licensed from the internet and other sources, including:

- Websites, forums, and blogs
- Digitized books and published literature
- News articles and journalism
- Encyclopedic and reference content
- Other publicly available or licensed digital text

The sheer scale of this training data is what makes an LLM feel capable and fluent. But that same data has a shape and a set of gaps, and those gaps show up in what the AI produces.

Why This Creates Bias, Especially a Language and Cultural Bias

The internet is not an even, representative sample of human knowledge and experience. A disproportionate share of the highest-quality, most abundant digital text is written in English and reflects perspectives, institutions, and cultural norms from wealthier, English-speaking, internet-connected parts of the world.

As a result, AI tools tend to:

- Perform best in English and struggle more with other languages, dialects, and regional phrasing
- Reflect assumptions, norms, and framing common in U.S. and Western media and institutions
- Underrepresent perspectives, histories, and expertise from communities with less digital presence, including many communities that civil society and advocacy organizations serve
- Reproduce stereotypes or skewed framing present in the training data, sometimes subtly

“ **Groundwire's Advice:** Because your organization's mission likely centers communities that are underrepresented in mainstream digital content, treat AI output about those communities with extra scrutiny. If a response feels generic, flattened, or slightly "off" about the population you serve, that is often the language and cultural bias in the training data showing through, not a neutral or authoritative answer.

2. Automation vs. AI: What's the Difference?

These terms are often used interchangeably, but they describe different kinds of technology, and knowing the difference helps you pick the right tool for a task.

Automation

Automation follows fixed, pre-programmed rules to complete repetitive tasks the same way every time. It does not "think," generate new content, or adapt to unfamiliar situations. Automation is predictable and consistent by design.

- **Examples:** a scheduled email reminder, a form that auto-fills a donor's name into a thank-you letter, a spreadsheet formula, a chatbot that only answers from a fixed script of FAQs
- **Strength:** reliable, repeatable, easy to audit — it does the same thing every time

- **Limit:** cannot handle anything outside the rules it was given

Artificial Intelligence

AI, and generative AI in particular, produces new content or judgments by recognizing patterns rather than following fixed rules. It can handle open-ended requests, adapt its response to context, and generate something that did not exist before — but its output can vary and is not guaranteed to be accurate.

- **Examples:** drafting a letter to a policymaker, summarizing a 40-page report, brainstorming talking points, answering an open-ended question about your uploaded documents
- **Strength:** flexible, handles novel or complex requests, saves drafting time
- **Limit:** can be inconsistent, confidently wrong, and requires human review every time

“ **Groundwire's Advice:** If a task is repetitive and the rules never change, look for an automation, not an AI tool — it will be more reliable and often cheaper. Reach for AI when the task is open-ended, requires drafting or synthesis, and you have time to review the output before it goes out the door.

3. Where AI Can Help — and Where It Has Limits

How AI Can Help

AI can help you save time by:

- Drafting content
- Rewriting text
- Summarizing information
- Organizing ideas
- Identifying patterns in large amounts of information
- Answering questions based on information it has access to

Where AI Has Limits

AI can be helpful, but it still needs human review because it does not:

- Understand the full situation or context
 - Know whether information is accurate or complete
 - Recognize what is fair, appropriate, or sensitive
 - Replace your professional expertise, knowledge, or judgment
-

4. Why AI Responses Always Need Human Review

When AI Sounds Confident

AI can write responses that sound polished and confident, even when the information is incomplete, outdated, or wrong. That means you may see responses that:

- Sound accurate but include incorrect information
- Leave out important context or details
- Repeat or reinforce an assumption instead of correcting it
- Look professional even when they still need review

When AI Creates False Information

AI may create information that sounds real but cannot be verified or is not accurate. This is sometimes called a hallucination. This can happen when AI fills in missing details instead of saying it does not know. Be especially careful with:

- Numbers, statistics, or dates
- Citations, sources, or links
- Technical specifications or product details
- Policy, legal, or compliance information
- Specialized topics that require expert knowledge

A Simple Way to Think About AI

Think of AI as a helpful assistant that has seen many examples and can create a polished response quickly. However, AI may fill in missing details, make assumptions, or present information with more confidence than it should. Use AI to support your work, but do not rely on it as the final answer without reviewing the response first.

Before You Use an AI Response

Before copying, sharing, or acting on AI-generated content, ask yourself:

- Does this match what I know or what trusted sources say?
- Is any important context missing?
- Does the response sound more certain than the information supports?
- Have I verified key facts before using it?

“ **Remember:** AI output should always be treated as a first draft, not a final answer.

5. How to Prompt AI Effectively

The quality of an AI response depends heavily on the quality of the prompt. A simple framework — Context, Role, Instructions, and Tone — will get you a far more useful first draft.

Context — share helpful background

“ "I am reviewing a vendor proposal for a clinical operations team."

Role — tell AI how to approach the task

“ "Act as a neutral assistant helping summarize information."

Instructions — say exactly what you want

“ "Summarize the key offerings in three bullet points."

Tone — describe how it should sound

“ "Use clear, plain language suitable for a non-technical audience."

6. Should Your Organization Get an Enterprise AI License?

Once staff start relying on AI regularly, organizations face a real decision: stay on free or individual paid accounts, or invest in an enterprise-tier license. For a large corporation this is a simple budget line. For a small civil society organization with a handful of staff and a tight budget, it is a genuine trade-off worth thinking through deliberately.

The Case for an Enterprise License

An enterprise or business-tier plan (as opposed to a free account or an individual \$20/month subscription) typically adds:

- **Data protection by default:** inputs are excluded from model training and are not used to improve the underlying AI, which matters if staff are pasting in draft grant language, HR content, or program data
- **Centralized admin controls:** single sign-on, the ability to see (in aggregate) how the tool is being used, and the ability to immediately revoke access when someone leaves the organization
- **Consistency and oversight:** one sanctioned tool is far easier to govern than a dozen staff quietly using a dozen different free tools with unclear settings, sometimes called "AI sprawl"
- **Higher usage limits and stronger models:** fewer interruptions from rate limits, and often access to more capable versions of the AI
- **Compliance support:** audit logs, data retention controls, and documentation that can matter for funder requirements or legal review

Why Many Small Organizations Hesitate

- **Cost adds up fast:** enterprise pricing is generally quoted per seat, often in a similar range to the ~\$20/month individual tier but with minimum seat counts or contract terms — a real cost for an organization with five or ten staff and a lean budget
- **Procurement overhead:** contracts, data processing agreements, and admin setup take staff time that a small team may not have
- **Uneven staff usage:** if only two or three people use AI daily, paying for enterprise seats for everyone may not be worth it yet
- **Free and individual tiers have improved:** many platforms now let individual paid users opt out of having their chats used for training, closing some of the privacy gap with enterprise tiers, though admin oversight and audit tools still require the enterprise tier

The Privacy and Security Angle

For organizations serving vulnerable populations — immigrants, survivors, litigants, patients, or other communities where a data exposure could cause real harm — privacy and security should weigh more heavily than cost alone.

- On any tier, free or enterprise, staff should never enter names, addresses, case details, or other identifying or confidential information into an AI tool. Treat every AI input as if it could become public.
- Enterprise and business tiers generally guarantee, contractually, that inputs are not used for model training — free tiers may not offer this by default and require staff to manually adjust settings, which is easy to forget
- Centralized admin control means leadership can see which tool the organization is using and turn off access in one place, rather than discovering after the fact that AI sprawl left data scattered across accounts no one is tracking
- If your organization has real subpoena risk or works with especially sensitive information, a privacy-first tool built around end-to-end encryption is worth exploring alongside — or instead of — a mainstream enterprise plan (see Confer, below)

“ **Groundwire's Advice:** Small organizations do not need to choose the most expensive option. A reasonable starting point is: pick one AI tool as your organization's standard, put it on individual paid accounts (with data-sharing settings turned off) while usage is light, and move to an enterprise or business tier once several staff are using AI daily or once you're regularly handling anything sensitive. Whatever tier you choose, pair it with a written AI policy so staff know what they can and cannot enter into the tool.

Comparing Enterprise AI Vendors

The table below compares the leading enterprise AI platforms most organizations are choosing between today.

Feature	ChatGPT Enterprise (OpenAI)	Claude for Business (Anthropic)	Gemini for Business (Google)	Microsoft 365 Copilot
Primary Strength	Creativity, R&D, coding, content ideation	Safe, reliable AI for complex reasoning & code	Real-time web search & Google Workspace integration	Deep integration with Microsoft 365 apps (Word, Teams, Outlook, etc.)
Data Privacy	Never trains on customer data; data is not used to improve models	Customer data is not used for training; explicit privacy-first design	Workspace data used contextually but not for model training; customers can disable AI data usage	Customer data not used to train models; processed within the Microsoft cloud

Feature	ChatGPT Enterprise (OpenAI)	Claude for Business (Anthropic)	Gemini for Business (Google)	Microsoft 365 Copilot
Admin Controls	SSO, SCIM, role-based access, usage analytics	Self-serve seat management, spend caps, usage & code analytics, compliance API	Admin console, contextual grounding in Drive/Gmail, data regions, DLP controls	Full M365 admin center integration, conditional access, sensitivity labels
Security	Encrypted in transit and at rest; data residency available; no OpenAI staff access to customer data	Encrypted in transit and at rest; no training on inputs; audit logs and compliance API	Encrypted in transit and at rest; DLP, context-aware access, data region controls	Encrypted in transit and at rest; Microsoft Zero Trust model; integrates with Defender and Purview
Workspace Integration	SharePoint, OneDrive, Google Drive, Box, GitHub; custom app integrations via GPTs	Google Workspace, GitHub, Slack via API or webhooks	Native in Gmail, Docs, Drive, Meet, Calendar ("Gemini in Workspace")	Fully embedded in Word, Excel, PowerPoint, Outlook, Teams, Loop
Collaboration Features	Shared GPTs, team knowledge base, chat with files (PDFs, spreadsheets)	Projects & Artifacts (shared workspaces), team prompts & templates	Real-time co-editing with AI suggestions, AI-generated meeting summaries in Meet	Meeting recaps in Teams, smart scheduling in Outlook, collaborative drafting
Compliance & Governance	SOC 2, ISO 27001, GDPR; data residency options	SOC 2, HIPAA, GDPR; compliance API for audit and retention	FedRAMP, HIPAA, GDPR; data regions, eDiscovery	Full M365 compliance suite; Purview, eDiscovery, sensitivity labels

A Privacy-First Alternative to Watch: Confer

A newer category of AI platform is emerging that is built around privacy as the foundation rather than an add-on feature. Confer, launched by Moxie Marlinspike (the creator of Signal), is the leading example. It's worth understanding, though it is not yet a like-for-like substitute for the enterprise platforms above.

- Confer is designed as a private, end-to-end encrypted AI assistant — the goal is that no one, not even Confer itself, can see what you say to it
- If Confer were subpoenaed, the government could not see what was shared with it — only that an account exists
- It runs on a small subscription model (roughly \$35/month) rather than an enterprise per-seat contract, and currently has no team plan, admin console, API, or compliance documentation — meaning it is best suited to an individual staff member handling especially sensitive conversations, not yet a replacement for an organization-wide platform
- It does not currently support file or document uploads, or collaboration features, though this may change as the platform matures

Groundwire's Advice: Confer is a great option for an individual staff member — an executive director, an attorney, or anyone regularly discussing sensitive matters — who wants a more private alternative to mainstream chatbots for one-off, non-collaborative use. It is not yet ready to serve as your organization's primary, organization-wide AI platform. We recommend watching how it develops and pairing it with, rather than substituting it for, your main enterprise tool.

A Final Word

AI is a genuinely useful tool for a resource-constrained organization — but it is a tool, not a replacement for your judgment, your relationships, or your mission expertise. Whatever combination of tools and license tiers you choose, the constant should be human review, a clear written policy, and staff who understand both what AI is good at and where it falls short.
